

Optimizing the hybrid parallelization of BHAC

Salvatore Cielo^a, Oliver Porth^b, Luigi Iapichino^a, Anupam Karmakar^c, Hector Olivares^d, Chun Xia^e

^aLeibniz Supercomputing Centre of the Bavarian Academy of Sciences and Humanities (LRZ), Garching b. München, Germany

^bAnton Pannekoek Institute for Astronomy, University of Amsterdam, Science Park 904, 1098 XH, Amsterdam, The Netherlands

^cLenovo Infrastructure Solutions Group, Lenovo Global Technology, Meitnerstr. 9, D-70563 Stuttgart, Germany

^dDepartment of Astrophysics/IMAPP, Radboud University Nijmegen, P.O. Box 9010, 6500 GL Nijmegen, The Netherlands

^eSchool of Physics and Astronomy, Yunnan University, Kunming 650050, China

Abstract

We present our experience with the modernization on the GR-MHD code BHAC, aimed at improving its novel hybrid (MPI+OpenMP) parallelization scheme. In doing so, we showcase the use of performance profiling tools usable on x86 (Intel-based) architectures.

Our performance characterization and threading analysis provided guidance in improving the concurrency and thus the efficiency of the OpenMP parallel regions. We assess scaling and communication patterns in order to identify and alleviate MPI bottlenecks, with both runtime switches and precise code interventions. The performance of optimized version of BHAC improved by $\sim 28\%$, making it viable for scaling on hundreds of thousands of supercomputer nodes.

We finally test whether porting such optimizations to different hardware is likewise beneficial on the new architecture by running on ARM A64FX vector nodes.

Keywords: Science, Publication, Complicated

1. Overview

Supercomputers have become essential tools in most fields of modern research. Thanks to ever-growing computational resources, the boundary of knowledge and the possibilities of the scientific method have been pushed through otherwise impossible challenges. A very representative example is given by the COVID-19 pandemic: it has been shown that supercomputers are crucial tools for drug discovery and development¹. Such novel projects have had the additional merit of better introducing supercomputers to the general public, and to funnel new classes of users from industry and science. Both novel applications and the ones that are well-established since decades (among the latter, numerical simulations are the main representative ones) have however been facing several challenges in order to utilize at best this powerful resource.

On the verge of the Exascale Computing Era, we can confidently state that scientific applications, and astrophysical ones in particular, have –not without significant efforts– successfully managed several challenges in the field of High-Performance Computing (HPC): to keep up with the performance improvement dictated by Moore’s law (see e.g. [1]), new layers of parallelisation have been introduced in the architectures, and programming paradigms have been designed accordingly. Examples include the emerging of hybrid (distributed plus shared-memory) parallel codes to make efficient use of multicore nodes

while allowing good scaling, even up to full machine size; the introduction of vector registers and the simultaneous development of SIMD (Single Instruction Multiple Data) programming strategies; and structured memory hierarchy of modern processors, to ensure data streaming from the main memory into the computing units.

Yet the field is undergoing several major breakthroughs, well exemplified by the top places of the TOP500² ranking. The usage of GPUs as compute accelerators and the possibilities offered by the latest vector architectures are the most promising technologies in an Exascale perspective. Both of them however present challenges, respectively the necessity of programming the extra devices, and a thorough optimization work. Supercomputing centers are in charge of presenting these technologies to users from the different science fields, foster dialog with vendors and explore viable programming models.

In this work we present our experiences from one such project, the modernization of the BHAC code³, hosted by the Astro- and Plasma- physics Application Lab (AstroLab) of the Leibniz Supercomputing Center in the framework of high-level user support, and operated jointly with the code developers.

We begin by outlining the code baseline performance and defining the optimization goals (Section 3). We then characterize the initial code performance (SIMD, MPI layer, roofline analysis) at small and intermediate scales through a first series of scalings and tool-assisted runs (Section 4). In Section 5 we instead describe the main bottlenecks we identify, highlight the

URL: cielo@lrz.de (Salvatore Cielo), o.porth@uva.nl (Oliver Porth)

¹See for instance <https://www.combiomed.eu/> for a description of the efforts in the framework of a current EU-funded consortium in this field.

²<https://www.top500.org/>

³<https://bhac.science>

implemented code interventions and show the improved performance at scale (Section 5.2). We finally attempt a performance portability test on a A64FX architecture (Section 6), before summarizing the optimization roadmap and drawing some general conclusions (Section 7).

2. Code overview

BHAC is an open-source (GPL v3.0) general relativistic MHD code [2, 3], written in Fortran90, parallelized using MPI and is based on the MPI-AMRVAC Toolkit⁴ v1.0 [4, 5, 6]. BHAC offers a variety of numerical methods and fully adaptive block based (bi-, quad- or oct-) tree adaptive mesh refinement (AMR) using staggered mesh constrained transport in any covariant coordinate system [3]. The high-resolution shock-capturing finite volume scheme uses second order Implicit/Explicit timesteps combined with second or third order (PPM) spatial reconstruction. BHAC uses robust multi-dimensional non-linear solvers (Newton-Raphson and Newton-Krylov) for the conversion from the integrated conserved variables to the primitive variables with well tested backup strategies [e.g. 9]. This allows to perform simulations of challenging magnetically-dominated regimes as present in pulsar magnetospheres, black hole accretion and relativistic reconnection.

BHAC has been applied (among other codes like HARM, KORAL and H-AMR) to build GRMHD models for the simulation library used for the interpretation of the first “picture of a black hole” by the Event Horizon Telescope Consortium [10]. The veracity of GRMHD codes like BHAC has been established in a community spanning code comparison effort [11]. In [12], BHAC simulations were employed to show how the accretion process onto horizon- and surfaceless “black hole mimickers” called Boson stars differs from the black hole case. The study showed how the EHT can be used to distinguish between different classes of compact objects and thus affirm the black hole hypothesis.

Numerical simulations with BHAC are also used to unravel the origin of “flares” which are observed from the galactic center black hole SgrA* in the X-ray and infrared bands [e.g. 13, 14]. Two distinct physical scenarios are currently being investigated [15, 16, 17].

BHAC is also used to model jet formation in the context of the binary neutron star merger event GW170817 [18]. In [19, 20], BHAC simulations showed, for the first time, how self-consistently launched magnetized outflows provide accurate fits to the multi-frequency afterglow lightcurve of GW170817, allowing to deduce key parameters of the source.

Yet another research line uses BHAC to study magnetic reconnection in the relativistic regime. Capitalizing on the AMR capabilities, *effective* resolutions of 8000^3 have already been achieved which allows to resolve both the large scale plasma instabilities as well as the small scale current sheets forming in the non-linear evolution (e.g. [21] and Ripperda et al. in prep).

3. Initial performance and optimization goals

In the past few years, BHAC has undergone previous modernization projects within LRZ. These mainly targeted a refactoring of the solver scheme towards a task-based algorithm, which also allowed the addressing of bottlenecks in the memory access pattern. Hence the refactoring of important loops of the source code, finally complemented by further efforts towards parallel I/O challenges and SIMD usage.




Figure 1: Initial scaling on SuperMUC-NG. At 9600 cores, the pure MPI case requires in reality 400 nodes in order not run out of memory, due to the MPI buffer overhead. Hybrid MPI/OpenMP alleviates this problem, running successfully up to 76800 cores. However, there is a severe performance penalty when increasing the number of OpenMP threads per task.

Since then, the code has seen the introduction of a basic hybrid MPI/OpenMP parallelization, which has not yet been used in production, but has proven very promising in extending the scalability range of the code. Typical use cases so far have utilized up to 8000 cores, e.g. on SuperMUC, Cartesius/Netherlands or Breniac at VSC/Belgium. However, we have observed in previous scaling runs that the MPI overhead of pure MPI runs becomes prohibitive at $\sim 10\,000$ processes. This is demonstrated in Figure 1 which shows the BHAC strong-scaling on SuperMUC-NG (48 cores per node). The pure MPI run with 9600 processes in fact allocated 400 nodes, loaded with only half the processes to fit memory requirements.

Our main target machine is LRZ’s SuperMUC-NG, which consists of 6,336 Intel Xeon Skylake Platinum 8174 processor compute nodes each with 48 cores and 96 GB memory (thin nodes), as well as 144 nodes of the same processors each 48 cores and 768 GB memory per node (fat nodes). An Intel Omni-Path (OPA) high-performance fast interconnect provides up to 100G network bandwidth on these nodes. These compute nodes are bundled into 8 domains (islands). The OmniPath network topology within the islands is a ‘fat tree’ for highly efficient communication and the interconnection between the islands is pruned with a pruning factor 1:4.

The preliminary exploration shown in Figure 1 shows the potential of the hybrid parallelization: already with four or eight threads per task, the code is able to run beyond 400 nodes (tested up to 1600 nodes/76800 cores) and scaling is efficient until 800 nodes. However, in its initial state, the hybrid im-

⁴<http://amrvac.org>

plementation comes with a severe performance penalty, e.g. at eight threads per task, the runtime is twice the pure MPI case at the same node count.

To efficiently run large problems beyond 2000^3 cells as needed for example in the study of relativistic turbulence, it is necessary to go beyond 10 000 cores accessible with pure MPI parallelization. Goal of this AstroLab project is hence to understand and improve the performance in the hybrid implementation. Since this is the first investigation of its kind with the MPI-AMRVAC toolkit, we use a simple uniform grid setup and aim for a fairly complete characterization of the code infrastructure.

In order to explore all the code bottlenecks, at node level and beyond, we need to run BHAC at different scales, and thus test problems of different sizes.

We target single node, 16 nodes and 800 nodes, scaling up the same physical problem while keeping a constant workload of 65536 cells per core. Thus our problems range from $\sim 3 \times 10^6$ to over 2×10^9 computational cells.

4. Profiling of the initial version

4.1. single node characterization

For the first general profiling runs we used SuperMUC-NG at LRZ. As the OpenMP implementation was experimental, we ran a correctness analysis with the Intel Inspector tool; this check yielded several race conditions which we needed to address before attempting any optimization, but were very easily corrected. At this point we performed a few single node scaling tests and tool-assisted profiling runs to determine the optimal node configuration for the (mainly compiler-assisted vectorization and MPI/OpenMP ratio). The tools are more properly discussed in Section 5, so for now we just summarize these early results of ours.

Thanks to the previous modernization interventions, BHAC showed good SIMD efficiency, to the point that it was already convenient for us to start with the Xeon’s full AVX512 instruction set and 512-bit ZMM register, by automatic compiler vectorization (i.e. compiling with the flags `-xAVX512 -qopt-zmm-usage=high`). A 32% speedup from SIMD alone was still theoretically possible; however not very profitable, as the setup showed a prevalence of memory-bound kernels.

The hotspot analysis yielded no pressing need for action, suggesting that we could proceed with the optimization of the OpenMP threading efficiency. On this respect, node configurations with a high numbers of thread (8 to 32 per node) had actually higher parallel efficiency than with fewer threads (at least for the single node case); thus it was more convenient to run any blind test with such configurations.

This fact also strongly suggests that the problems are likely due the MPI/OpenMP interaction (as using less tasks per node is always more efficient when OpenMP is used). We later (see Section ??) found this issue to be key, the exact cause residing in the details of MPI’s shared memory intra-node communications. A better characterization of the MPI structure of the code was thus necessary.

4.2. MPI profiling of the 16 nodes test case

Using `mpi_trace` (Lenovo’s lightweight and extended MPI tracing tool), we evaluated the pure MPI scaling of the 16 node test case in the range from 80 to 3200 MPI tasks, in a strong-scaling fashion. These runs were performed at the Lenovo Benchmark Cluster in Stuttgart (“Lenox”) on Lenovo SD530 platform with dual-socket Intel Xeon Gold 6148 based nodes. These nodes have roughly 80 GB available physical memory for the running applications. An Intel Omni-Path (OPA) high-performance fast interconnect provides up to 100G network bandwidth on these nodes. For these baseline measurements Intel fortran compiler version 20.2 and Intel MPI 2019.8 were used.




Figure 2: Pure MPI scaling (solid lines) of the 16 node use-case on the Lenox cluster. The corresponding scaling efficiency declines in both direction from the designed optimal load of 16 nodes.

The results obtained are summarized in Figure 2 which shows an excellent strong scaling of up to a factor of 19 from 160 to 3200 processes. The run at 80 MPI tasks (2 nodes) significantly under-performed, likely due to caching effects. This experiment shows that the pure MPI implementation using non-blocking `ISEND/IRECV` is quite efficient, requiring no urgent improvements.

An analysis of the `mpi_trace` findings, Figure 3, shows how the relative time spent on MPI communications alters very slowly with the number of MPI tasks, whereas the memory usage rapidly drops due to greater parallelism. This, in a way, explains the good scalability behaviour achieved with node scaling of this use-case.

As an aside, the same setup was also used to compute the frequency scaling and energy efficiency of the code. Using 16 nodes, the frequency was scaled from 1 GHz to 2.4GHz and the time-to-solution and average node power (Watt) was recorded. The results are illustrated in Figure 4.

At an average of 277W per node, the minimum energy-to-solution (~ 309 kJ) was found for a clock speed of 2.2GHz, this is an optimum trade-off between performance and energy consumption. At a maximum clock speed, the energy-to-solution increases around 9%, however, it slightly lowers ($\sim 6\%$) the time-to-solution performance. Switching on the turbo frequency does not positively impact the overall performance, but




Figure 3: Computation/communication fraction (histogram, stacked bars) and memory footprint (solid line) for the 16 nodes, pure MPI setup (from trace analysis on the Lenox cluster). Once more the optimal computation, communication balance and memory requirement per task are achieved for the 16 nodes run.

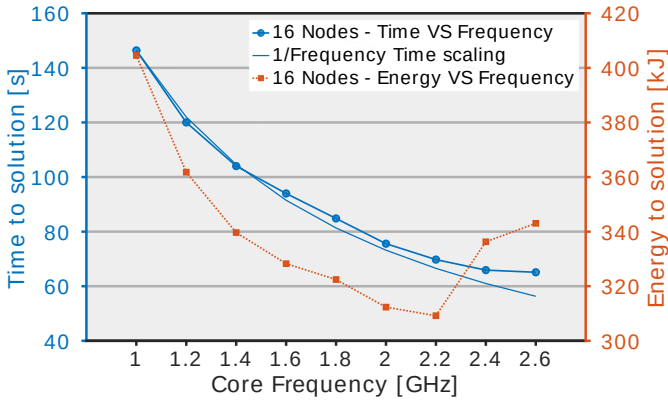


Figure 4: Frequency scaling (time to solution, light blue line) of the 16 node pure MPI test setup. The corresponding global energy to solution (orange line) has its clear minimum at a clock frequency of 2.2GHz.

only surges the energy to solution requirement.

With these numbers in hand, we can compute useful conversion factors indicating the energy and carbon footprint of full numerical simulations with BHAC. Since the majority of the electricity consumption in astrophysical institutes originates from the use of high-performance computing [e.g. 22, for the recent analysis of a representative case], it is important to be sensitive to the power- and climate- impact of numerical simulations. As argued by [23], computing related greenhouse gas emissions also surpass telescope operations in astronomy. We estimate the carbon footprint of simulations performed on hardware with similar characteristics to Lenox as follows: A valid metric for the work of grid-based codes such as BHAC is the number of the performed cellupdates, whereby we count each sub-step of the temporal integration multiplied by the total number of cells in the simulation. Our test-setup performed $\approx 5.3 \times 10^9$ individual cellupdates which yields an energy-per-million-cellupdate conversion of $E_{pc6} = 58 \frac{\text{J}}{10^6 \text{cellupdates}} = 1.6 \times 10^{-5} \frac{\text{kWh}}{10^6 \text{cellupdates}}$. Adopting the average CO₂-equivalent emis-

sion per kWh electrical energy in Europe (2019) ⁵ of $275 \frac{\text{gCO}_2\text{e}}{\text{kWh}}$, we obtain the emission coefficient $CO_{26} = 4.4 \frac{\text{mg}}{10^6 \text{cellupdates}}$. To put this in perspective: for a typical production run within the simulation library in [24], we perform 2×10^6 iterations, and simulate a grid of 6×10^6 cells. Using a two-step timeintegrator, the total simulation hence performed 24×10^{12} cellupdates. Adopting the characteristics found for the test run on Lenox, this corresponds to an energy to solution of 384kWh and a CO₂-equivalent of 106Kg. The same amount of greenhouse gasses will be emitted by a 820 km drive in an average new car in Europe (2019) ⁶.

4.3. Roofline analysis

Using Intel’s Application Performance Snapshot™ (APS) and Advisor™, the performance bottlenecks in the hybrid implementation were identified. Advisor is especially useful in the characterization of the node-level performance (safe for threading utilization, see below), SIMD leverage and memory access efficiency. This information is best summarized in the *roofline analysis*, such as the one we present in Figure 5 for the baseline version of BHAC.

In this plot, the performance of each individual relevant kernel is measured at runtime and plotted against the nominal hardware capabilities, in terms of total FLOPS executed (y-axis) versus *arithmetic intensity* (FLOPS executed per byte, i.e. a measure of how efficiently the memory hierarchy is used; x-axis). Kernel bottlenecks here appear as *roofs*, either horizontal (compute bottlenecks) or slanted (memory bottlenecks). Larger circles represent longer runtimes, SIMD kernels are in orange, scalar ones in light blue.

In the figure we also display the information Advisor provides on click (position in the source code, most relevant roof, arithmetic intensity extrema for all memories in the hierarchy, performance figures), for one of the most relevant kernels in the optimization, `fill_gc_sr1` (in red, with bold border).

The analysis shows a mixture of memory-bound and compute-bound kernels, the former being however prevalent. Since memory utilization also limits SIMD efficiency, even that one has room for improvement, confirming the APS results, and in concordance with the fact that vectorized kernels do not appear on average higher than serial ones.

5. Optimization achievements

5.1. Threading overhaul

For the proper OpenMP analysis, the most useful tool provided by Intel is instead VTune Profiler™. Its summary view reported an OpenMP imbalance which could be solved by adding dynamic scheduling to the loops with significant workload. This yielded an average performance improvement of $\sim 5\%$.

⁵https://www.eea.europa.eu/ds_resolveuid/7e1d0edbc6b4488a89fb1686d8239858

⁶<https://www.eea.europa.eu/data-and-maps/data/co2-cars-emission-18>

Besides this clear summary of the threading utilization, VTune is capable of showing detailed core utilization histograms over time. Idle times can be further distinguished in MPI wait times, OpenMP spin times, or serial OpenMP regions. The graphs can be arranged in several useful ways, and all kernel present in a given time interval are shown with relevant metrics.

An example of this analysis is shown in Figure 6: the considered run featured 2 MPI task per node and 24 OpenMP threads per task; the utilization histograms refer to a single task, with the master thread on top, followed by the other threads that depend on it (for brevity, only 6 working threads are shown).

Let us at first examine the top panel, which refers to the initial version. The diagrams shows clearly several idle regions, where all worker threads wait for the master to finish long serial calculations and MPI calls (see color legend in the figure). The parallel regions are always well aligned as a consequence of the dynamic scheduling, which improved work-sharing (although with the small impact we mentioned).

VTune’s threading analysis identified these large serial code blocks within kernels managing the ghost-cell exchange. A look at the source and assembly code, also through VTune’s proposed interface, identified a major OpenMP serial region in the ghost-cell exchange. The code could then be optimized by the developers via separating out the MPI-dependent and MPI-independent ghostcell operations. This optimization positively impacts performance in two ways: 1. the MPI-independent parts have been fully OpenMP parallelized and 2. the newly OpenMP parallelized operations are executed before the MPI-WAITALL calls of the MPI-dependent operations, allowing more overlap of computation and communication.

After the refactoring, we repeated VTune’s threading analysis (shown in the bottom panel of Figure 6). The idle and wait times are very much reduced, resulting in much more compact parallel regions and thus highly increased core utilization.

Figure 7 shows the single node scaling before and after the optimization. While pure MPI scaling (dark blue lines) is close to ideal, the initial pure OpenMP scaling (orange lines) was poor due to the increasing contribution of serial regions (Amdahl’s law). The updated OpenMP scaling (light blue) performs as well as pure MPI. The hybrid performance was investigated by varying the OpenMP/MPI contributions. In the figure we exemplify this by either choosing 8 OpenMP threads and increasing MPI tasks, or by fixing two MPI tasks and increasing the OpenMP thread count (see legend). This shows a generally acceptable scaling but also points out room for improvement when mixing MPI and OpenMP. The optimized OpenMP runs show a clearly improved scaling efficiency (in a few cases also above unity) which testifies the effectiveness of our interventions.

5.2. Final performance at scale

Using the updated code version, the initial large scaling run was repeated using 8 threads per task. The result is shown in Figure 8. Even on this scale, the data shows that improving the OpenMP parallelism led to an average performance increase of

27% compared to the initial code. This shows a significant reduction of the hybrid penalty. Trace analysis at 16 nodes has indicated that further modifications of the ghost-cell exchange for mixed MPI/OpenMP could improve this even more. Unfortunately a known bug of the default MPI version of SuperMUC-NG at the time of writing made pursuing a scaling to 1600 nodes and beyond not worth the effort. Yet the data collected so far give already enough directions for future optimizations.

Further exploration of the options offered by Intel MPI was very profitable: switching off *shared-memory intranode communication* by setting `export I_MPI_SHM=off` at runtime, we were able to efficiently port the single node speedups to a higher node counts (namely, a speedup of 28%). Finally we signal that switches to reduce the MPI memory footprint are also available, should the initial overhead problem be presented again (e.g. the family of `I_MPI_SHM_CELL_` and `I_MPI_SHM_HEAP_` environment variables, and the control on `I_MPI_MALLOCC`), which may also have positive impact on the performance.

6. Test run on A64FX architecture

We finally decided to explore whether our optimization work, so well-performing to Intel x86 architecture, offers some reward when the code is ported to different hardware.

Our target architecture for this experiment were the HPE CS500 nodes of LRZ’s BEAST (for *Bavarian Energy, Architecture, and Software Testbed*) cluster⁷ mounting A64FX processors. The A64FX is the 64-bit ARM architecture microprocessor designed by Fujitsu, in force of Fugaku⁸ the 537-TFlops supercomputer number one in the Top500 list at the time of writing (since June 2020). The processor is the first to feature the Armv8.2-A instruction set Scalable Vector Extension (SVE), capable of major performance enhancements on vectorizable operation. The (RISC) SVE – vectorizing operations of arbitrary length with a single instruction – operates on different principles than the (CISC/mixed) SIMD extension discussed so far, although the two have the same basic requirements (mainly, avoidance of vector dependencies within loops), and the fact that SIMD versions of some instructions are nevertheless present on the A64FX. It is thus legitimate to wonder how code optimized for one extension performs on the other. Proven benefits of SVE, besides high performance, include improved energy efficiency and ease of manufacturing; all valid motivational reasons which address the core bottlenecks of HPC and scientific computing for the Exascale Era.

During the setup of BHAC a few build issues occurred, for which fixes had to be found. We used Cray `ftn` compiler v.10.0.1, in an instance containing support for SVE but not for accelerators (not present anyways on the nodes). The `explain` command was rather useful in clarifying compilation errors, and made the porting process smoother. Nonetheless, compilation macros were adjusted to provide `.f90` extension to all source files, rather than `.f` as in the public version. Also, a

⁷https://www.lrz.de/presse/ereignisse/2020-11-06_BEAST/

⁸<https://www.r-ccs.riken.jp/en/fugaku/project>




Figure 5: VTune roofline analysis. The profiling shows a mixture of compute-bound and memory-bound kernels, with a prevalence of the latter. Such memory issues also hinder the SIMD efficiency of vectorized kernels (in orange), which do not particularly stand out from serial ones (in light blue). Point information for a relevant kernel are shown in the inset text panel (on-click feature of Advisor).




Figure 6: Timeline histograms of initial (top) and optimized (bottom) version, from VTune's Bottom-up view. 2 MPI tasks and 24 OpenMP per run. Only a few worker threads and one task/master thread (in blue) are shown. Wait and overhead time (see color legends) are enormously reduced as a consequence of parallelization of the ghost cell exchange region.

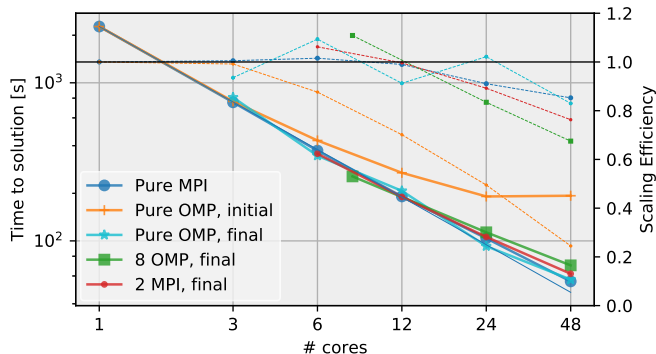


Figure 7: Single node scaling with varying contributions of MPI and OpenMP. Code restructuring has yielded a significant improvement in the pure OpenMP scaling, on par with the pure MPI case. The curve labeled as “8 OMP, final” uses the optimized code with fixed eight threads per task, varying the number of tasks. The curve “2 MPI, final” fixes 2 MPI tasks and investigates scaling by increasing the number of threads.

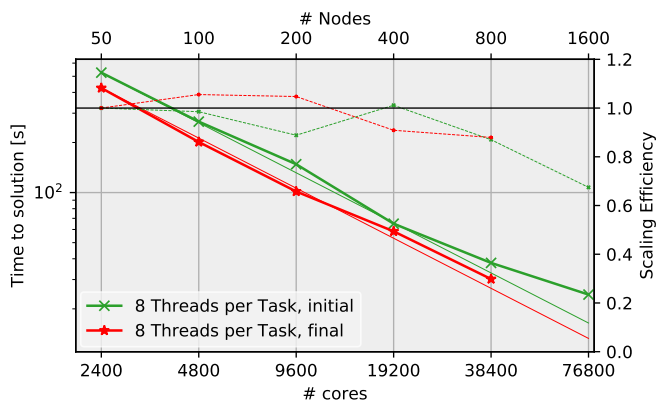


Figure 8: Comparison of the initial and final speedup for the large scale setup. The green curve is identical to the one shown in Figure 1.




Figure 9: Out-of-the-box scaling of the single node test on the A64FX node part of LRZ’s BEAST cluster. The nodes shows excellent scaling, especially for both pure paradigms, although overall inferior performance (about 3x slower than a single SuperMUC-NG node). Dedicated tuning of the application is likely to bring major improvement, but is outside our current scope.

compiler bug causing segmentation faults forced us to compile a small portion of BHAC with the option `-hipa0`, switching off procedural optimization for a few kernels. This issue has been notified, and will be fixed in future versions of the compiler (from v.12.0 on). All the other kernels have simply been compiled with `-O3` optimization.

We then ran the code interactively on a single node, performing a strong-scaling test similar to what is shown in Figure 7 (for the same code version). We tested pure MPI scaling from 1 to 48 cores (the maximum of the node; in blue), then a full-node pure OpenMP run (purple dot), and a hybrid configuration with 2 MPI tasks and increasing OpenMP threads (in yellow). Whenever running with OpenMP, we followed the advice given by the runtime to set `MV2_ENABLE_AFFINITY=0` to avoid potential thread overlap on individual physical cores; the results are displayed in Figure 9 (median of 15 measures, statistical errors not shown as always less than 1%).

A comparison with Figure 7 shows that the absolute performance is still about a factor 3 lower than on SuperMUC-NG. The pure paradigms perform best, with parallel efficiency around 90% (and note how in the pure MPI scaling the efficiency is higher in the full node than for 12 or 24 cores). These efficiency values appear well inline with those of SuperMUC-NG, though slightly higher for the pure paradigms, slightly lower for the hybrid run with 2 MPI tasks.

Inspection of the assembly code with the `perf record` and `perf report` tools showed very poor usage of the SVE set. The SIMD version of some instructions are used, although in limited amount when compared to SuperMUC-NG. Considering the high degree of SIMD vectorization of BHAC, even in its baseline version, the missing performance is not surprising. Additional reasons may include the fact that the node and compiler were never before tested on FORTRAN applications (otherwise administrators and expert users may have provided guidelines for compiler optimization), and the aforementioned procedural optimization switch off: even though limited to a

few kernels, those concern a rather important part of the grid tree management.

Finally, the measured performance of BHAC seems in line with other applications (e.g. benchmarks or other code optimization projects). The emerging pattern is that high speedups typically require much more involved optimization work, such as explicit loop unrolling and development of detailed performance models [25], which call for dedicated projects, when not dedicated staff; but can result in general optimization hints all users will benefit from.

We thus consider our test run concluded, having answered our most pressing question (i.e., astrophysical real-world applications optimized for x86 SIMD are unlikely to benefit from SVE out-of-the-box), and having laid the basis for porting BHAC to A64FX; future studies may start from here. The result is undesirable, but hopefully software/hardware co-design efforts will alleviate this discrepancy in the near future.

7. Conclusions and Outlook

In presenting our work on the BHAC code, we showcase some HPC tools and techniques of key importance in the optimization of codes' threading efficiency, especially in the case of newly introduced OpenMP parallelization layers aimed at production on modern supercomputers.

A characterization of the BHAC code infrastructure has been obtained using single node, 16 nodes and 800 node scenarios on the Intel Xeon based SuperMUC-NG supercomputer. The setup used for this initial investigation was deliberately chosen as simple as possible: we have resorted to non-AMR grids and IO has also been excluded. In this simple scenario, OpenMP correctness was addressed and imbalance was solved through dynamic scheduling (remaining OpenMP potential gain: 0.5% as reported by VTune).

Guided by VTune's hotspots and threading analyses and APS' communication pattern report, the OpenMP serial fraction was significantly reduced, leading to a performance increase of 27% at scale. We observed that the MPI node memory management was one of the main reason hindering the scale-up of the OpenMP optimizations to larger node numbers; we got the best results by switching Intel MPI's shared memory intra-node communication off. While there remains room for improvement, the project has shown that the hybrid implementation is capable to efficiently utilize over 30 000 cores allowing to study large scale problems in scientific production. The optimizations performed here have been added to the v1.1 release of BHAC⁹. We finally ported the code to a different architecture, an ARM A64FX vector processor part of the LRZ BEAST cluster. We observed with disappointment that the threading and SIMD optimizations (achieved in the course of several projects of which this work is only the last one) have little impact on the usage of vector instructions and thus on the final node-level performance, which remains about factor of 3 lower than on the x86 machine.

⁹<https://gitlab.itp.uni-frankfurt.de/BHAC-release/bhac/-/releases/v1.1>

The improvements described in this work are already merged into the staging branch of BHAC and will become part of the next public release.

The modernization work on BHAC is however not concluded. Among the main issues that remain open for further projects we can list: 1. addressing the remaining OpenMP serial fraction in case of mixed OpenMP/MPI parallelization. 2. Extending the optimizations to the AMR case. 3. Investigating the performance for a production level setup using AMR. Here, next to the issues due to OpenMP/MPI hybridization, potential load-imbalance on the MPI level will need to be addressed.

Acknowledgments

S.C. thanks Dr. Josef Weidendorfer, leader of LRZ Future Computing Group, for the support on the test runs with the A64FX nodes described in Section 6.

References

- [1] H. Sutter, The free lunch is over - a fundamental turn toward concurrency in software, *Dr. Dobbs' Journal* 30 (2005).
- [2] O. Porth, H. Olivares, Y. Mizuno, Z. Younsi, L. Rezzolla, M. Moscibrodzka, H. Falcke, M. Kramer, The black hole accretion code, *Computational Astrophysics and Cosmology* 4 (2017) 1. URL: <http://adsabs.harvard.edu/abs/2017ComAC...4....1P>. doi:10.1186/s40668-017-0020-2. arXiv:1611.09720.
- [3] H. Olivares, O. Porth, J. Davelaar, E. R. Most, C. M. Fromm, Y. Mizuno, Z. Younsi, L. Rezzolla, Constrained transport and adaptive mesh refinement in the Black Hole Accretion Code, *A&A* 629 (2019) A61. doi:10.1051/0004-6361/201935559. arXiv:1906.10795.
- [4] R. Keppens, Z. Meliani, A. van Marle, P. Delmont, A. Vlasiv, B. van der Holst, Parallel, grid-adaptive approaches for relativistic hydro and magnetohydrodynamics, *Journal of Computational Physics* 231 (2012) 718–744. URL: <http://www.sciencedirect.com/science/article/pii/S0021999111000386>. doi:10.1016/j.jcp.2011.01.020.
- [5] O. Porth, C. Xia, T. Hendrix, S. P. Moschou, R. Keppens, MPI-AMRVAC for Solar and Astrophysics, *ApJS* 214 (2014) 4. URL: <http://adsabs.harvard.edu/abs/2014ApJS...214....4P>. doi:10.1088/0067-0049/214/1/4. arXiv:1407.2052.
- [6] C. Xia, J. Teunissen, I. El Mellah, E. Chané, R. Keppens, Mpi-amrvac 2.0 for solar and astrophysical applications, *ApJS* 234 (2018) 30. URL: <http://adsabs.harvard.edu/abs/2018ApJS...234...30X>. doi:10.3847/1538-4365/aaa6c8. arXiv:1710.06140.
- [7] P. Colella, P. R. Woodward, The Piecewise Parabolic Method (PPM) for Gas-Dynamical Simulations, *Journal of Computational Physics* 54 (1984) 174–201. URL: <http://adsabs.harvard.edu/abs/1984JCoPh...54...174C>. doi:10.1016/0021-9991(84)90143-8.
- [8] F. Acker, R. B. de R. Borges, B. Costa, An improved WENO-Z scheme, *Journal of Computational Physics* 313 (2016) 726–753. doi:10.1016/j.jcp.2016.01.038.
- [9] B. Ripperda, F. Bacchini, O. Porth, E. R. Most, H. Olivares, A. Nathanail, L. Rezzolla, J. Teunissen, R. Keppens, General-relativistic resistive magnetohydrodynamics with robust primitive-variable recovery for accretion disk simulations, *ApJS* 244 (2019) 10. URL: <https://ui.adsabs.harvard.edu/abs/2019ApJS...244...10R>. doi:10.3847/1538-4365/ab3922. arXiv:1907.07197.
- [10] Event Horizon Telescope Collaboration, K. Akiyama, A. Alberdi, W. Alef, K. Asada, R. Azulay, A.-K. Baczko, D. Ball, M. Baloković, J. Barrett, D. Bintley, L. Blackburn, W. Boland, K. L. Bouman, G. C. Bower, M. Bremer, C. D. Brinkerink, R. Brissenden, S. Britzen, A. E. Broderick, D. Brogiere, T. Bronzwaer, D.-Y. Byun, J. E. Carlstrom, A. Chael, C.-k. Chan, S. Chatterjee, K. Chatterjee, M.-T. Chen, Y. Chen, I. Cho, P. Christian, J. E. Conway, J. M. Cordes, G. B. Crew, Y. Cui, J. Davelaar, M. De Laurentis, R. Deane, J. Dempsey, G. Desvignes,

- J. Dexter, S. S. Doeleman, R. P. Eatough, H. Falcke, V. L. Fish, E. Fomalont, R. Fraga-Encinas, W. T. Freeman, P. Friberg, C. M. Fromm, J. L. Gómez, P. Galison, C. F. Gammie, R. García, O. Gentaz, B. Georgiev, C. Goddi, R. Gold, M. Gu, M. Gurwell, K. Hada, M. H. Hecht, R. Hesper, L. C. Ho, P. Ho, M. Honma, C.-W. L. Huang, L. Huang, D. H. Hughes, S. Ikeda, M. Inoue, S. Issaoun, D. J. James, B. T. Jannuzi, M. Janssen, B. Jeter, W. Jiang, M. D. Johnson, S. Jorstad, T. Jung, M. Karami, R. Karuppusamy, T. Kawashima, G. K. Keating, M. Kettenis, J.-Y. Kim, J. Kim, M. Kino, J. Y. Koay, P. M. Koch, S. Koyama, M. Kramer, C. Kramer, T. P. Krichbaum, C.-Y. Kuo, T. R. Lauer, S.-S. Lee, Y.-R. Li, Z. Li, M. Lindqvist, K. Liu, E. Liuzzo, W.-P. Lo, A. P. Lobanov, L. Loinard, C. Lonsdale, R.-S. Lu, N. R. MacDonald, J. Mao, S. Markoff, D. P. Marrone, A. P. Marscher, I. Martí-Vidal, S. Matsushita, L. D. Matthews, L. Medeiros, K. M. Menten, Y. Mizuno, I. Mizuno, J. M. Moran, K. Moriyama, M. Moscibrodzka, C. Müller, H. Nagai, N. M. Nagar, M. Nakamura, R. Narayan, G. Narayanan, I. Natarajan, R. Neri, C. Ni, A. Noutsos, H. Okino, H. Olivares, G. N. Ortiz-León, T. Oyama, F. Özel, D. C. M. Palumbo, N. Patel, U.-L. Pen, D. W. Pesce, V. Piétu, R. Plambeck, A. PopStefanija, O. Porth, B. Prather, J. A. Preciado-López, D. Psaltis, H.-Y. Pu, V. Ramakrishnan, R. Rao, M. G. Rawlings, A. W. Raymond, L. Rezzolla, B. Ripperda, F. Roelofs, A. Rogers, E. Ros, M. Rose, A. Roshanineshat, H. Rottmann, A. L. Roy, C. Ruszczyk, B. R. Ryan, K. L. J. Rygl, S. Sánchez, D. Sánchez-Arguelles, M. Sasada, T. Savolainen, F. P. Schloerb, K.-F. Schuster, L. Shao, Z. Shen, D. Small, B. W. Sohn, J. SooHoo, F. Tazaki, P. Tiede, R. P. J. Tilanus, M. Titus, K. Toma, P. Torne, T. Trent, S. Trippe, S. Tsuda, I. van Bemmél, H. J. van Langevelde, D. R. van Rossum, J. Wagner, J. Wardle, J. Weintroub, N. Wex, R. Wharton, M. Wielgus, G. N. Wong, Q. Wu, K. Young, A. Young, Z. Younsi, F. Yuan, Y.-F. Yuan, J. A. Zensus, G. Zhao, S.-S. Zhao, Z. Zhu, J.-C. Algaba, A. Allardi, R. Amestica, J. Anczarski, U. Bach, F. K. Baganoff, C. Beaudoin, B. A. Benson, R. Berthold, J. M. Blanchard, R. Blundell, S. Bustamente, R. Cappallo, E. Castillo-Domínguez, C.-C. Chang, S.-H. Chang, S.-C. Chang, C.-C. Chen, R. Chilson, T. C. Chuter, R. Córdoba Rosado, I. M. Coulson, T. M. Crawford, J. Crowley, J. David, M. Derome, M. Dexter, S. Dornbusch, K. A. Dudevoir, S. A. Dzib, A. Eckart, C. Eckert, N. R. Erickson, W. B. Everett, A. Faber, J. R. Farah, V. Fath, T. W. Folkers, D. C. Forbes, R. Freund, A. I. Gómez-Ruiz, D. M. Gale, F. Gao, G. Geertsema, D. A. Graham, C. H. Greer, R. Grosslein, F. Gueth, D. Haggard, N. W. Halverson, C.-C. Han, K.-C. Han, J. Hao, Y. Hasegawa, J. W. Henning, A. Hernández-Gómez, R. Herrero-Illana, S. Heyminck, A. Hirota, J. Hoge, Y.-D. Huang, C. M. V. Impellizzeri, H. Jiang, A. Kamble, R. Keisler, K. Kimura, Y. Kono, D. Kubo, J. Kuroda, R. Lacasse, R. A. Laing, E. M. Leitch, C.-T. Li, L. C.-C. Lin, C.-T. Liu, K.-Y. Liu, L.-M. Lu, R. G. Marson, P. L. Martin-Cocher, K. D. Massingill, C. Matulonis, M. P. McColl, S. R. McWhirter, H. Messias, Z. Meyer-Zhao, D. Michalik, A. Montaña, W. Montgomerie, M. Mora-Klein, D. Muders, A. Nadolski, S. Navarro, J. Neilsen, C. H. Nguyen, H. Nishioka, T. Norton, M. A. Nowak, G. Nystrom, H. Ogawa, P. Oshiro, T. Oyama, H. Parsons, S. N. Paine, J. Peñalver, N. M. Phillips, M. Poirier, N. Pradel, R. A. Primiani, P. A. Raffin, A. S. Rahlin, G. Reiland, C. Risacher, I. Ruiz, A. F. Sáez-Madaín, R. Sassella, P. Schellart, P. Shaw, K. M. Silva, H. Shiohara, D. R. Smith, W. Snow, K. Souccar, D. Sousa, T. K. Sridharan, R. Srinivasan, W. Stahm, A. A. Stark, K. Story, S. T. Timmer, L. Vertatschitsch, C. Walther, T.-S. Wei, N. Whitehorn, A. R. Whitney, D. P. Woody, J. G. A. Wouterloot, M. Wright, P. Yamaguchi, C.-Y. Yu, M. Zaballos, S. Zhang, L. Ziurys, First m87 event horizon telescope results. i. the shadow of the supermassive black hole, *ApJL* 875 (2019) L1. URL: <http://adsabs.harvard.edu/abs/2019ApJ...875L...1E>. doi:10.3847/2041-8213/ab0ec7.
- [11] O. Porth, K. Chatterjee, R. Narayan, C. F. Gammie, Y. Mizuno, P. Aninos, J. G. Baker, M. Bugli, C.-k. Chan, J. Davelaar, L. Del Zanna, Z. B. Etienne, P. C. Fragile, B. J. Kelly, M. Liska, S. Markoff, J. C. McKinney, B. Mishra, S. C. Noble, H. Olivares, B. Prather, L. Rezzolla, B. R. Ryan, J. M. Stone, N. Tomei, C. J. White, Z. Younsi, K. Akiyama, A. Alberdi, W. Alef, K. Asada, R. Azulay, A.-K. Baczko, D. Ball, M. Baloković, J. Barrett, D. Bintley, L. Blackburn, W. Bolland, K. L. Bouman, G. C. Bower, M. Bremer, C. D. Brinkerink, R. Brisenden, S. Britzen, A. E. Broderick, D. Brogiere, T. Bronzwaer, D.-Y. Byun, J. E. Carlstrom, A. Chael, S. Chatterjee, M.-T. Chen, Y. Chen, I. Cho, P. Christian, J. E. Conway, J. M. Cordes, Geoffrey, B. Crew, Y. Cui, M. De Laurentis, R. Deane, J. Dempsey, G. Desvignes, S. S. Doeleman, R. P. Eatough, H. Falcke, V. L. Fish, E. Fomalont, R. Fraga-Encinas, B. Freeman, P. Friberg, C. M. Fromm, J. L. Gómez, P. Galison, R. García, O. Gentaz, B. Georgiev, C. Goddi, R. Gold, M. Gu, M. Gurwell, K. Hada, M. H. Hecht, R. Hesper, L. C. Ho, P. Ho, M. Honma, C.-W. L. Huang, L. Huang, D. H. Hughes, S. Ikeda, M. Inoue, S. Issaoun, D. J. James, B. T. Jannuzi, M. Janssen, B. Jeter, W. Jiang, M. D. Johnson, S. Jorstad, T. Jung, M. Karami, R. Karuppusamy, T. Kawashima, G. K. Keating, M. Kettenis, J.-Y. Kim, J. Kim, M. Kino, J. Y. Koay, Patrick, M. Koch, S. Koyama, M. Kramer, C. Kramer, T. P. Krichbaum, C.-Y. Kuo, T. R. Lauer, S.-S. Lee, Y.-R. Li, Z. Li, M. Lindqvist, K. Liu, E. Liuzzo, W.-P. Lo, A. P. Lobanov, L. Loinard, C. Lonsdale, R.-S. Lu, N. R. MacDonald, J. Mao, D. P. Marrone, A. P. Marscher, I. Martí-Vidal, S. Matsushita, L. D. Matthews, L. Medeiros, K. M. Menten, I. Mizuno, J. M. Moran, K. Moriyama, M. Moscibrodzka, C. Müller, H. Nagai, N. M. Nagar, M. Nakamura, G. Narayanan, I. Natarajan, R. Neri, C. Ni, A. Noutsos, H. Okino, T. Oyama, F. Özel, D. C. M. Palumbo, N. Patel, U.-L. Pen, D. W. Pesce, V. Piétu, R. Plambeck, A. PopStefanija, J. A. Preciado-López, D. Psaltis, H.-Y. Pu, V. Ramakrishnan, R. Rao, M. G. Rawlings, A. W. Raymond, B. Ripperda, F. Roelofs, A. Rogers, E. Ros, M. Rose, A. Roshanineshat, H. Rottmann, A. L. Roy, C. Ruszczyk, K. L. J. Rygl, S. Sánchez, D. Sánchez-Arguelles, M. Sasada, T. Savolainen, F. P. Schloerb, K.-F. Schuster, L. Shao, Z. Shen, D. Small, B. W. Sohn, J. SooHoo, F. Tazaki, P. Tiede, R. P. J. Tilanus, M. Titus, K. Toma, P. Torne, T. Trent, S. Trippe, S. Tsuda, I. van Bemmél, H. J. van Langevelde, D. R. van Rossum, J. Wagner, J. Wardle, J. Weintroub, N. Wex, R. Wharton, M. Wielgus, G. N. Wong, Q. Wu, K. Young, A. Young, F. Yuan, Y.-F. Yuan, J. A. Zensus, G. Zhao, S.-S. Zhao, Z. Zhu, T. E. H. T. Collaboration, The event horizon general relativistic magnetohydrodynamic code comparison project, *ApJS* 243 (2019) 26. URL: <https://ui.adsabs.harvard.edu/abs/2019ApJS...243...26P>. doi:10.3847/1538-4365/ab29fd. arXiv:1904.04923.
- [12] H. Olivares, Z. Younsi, C. M. Fromm, M. De Laurentis, O. Porth, Y. Mizuno, H. Falcke, M. Kramer, L. Rezzolla, How to tell an accreting boson star from a black hole, *MNRAS* 497 (2020) 521–535. URL: <https://ui.adsabs.harvard.edu/abs/2020MNRAS.497...521O>. doi:10.1093/mnr/51/1/521. arXiv:1809.08682.
- [13] K. Dodds-Eden, D. Porquet, G. Trap, E. Quataert, X. Haubois, S. Gillissen, N. Grosso, E. Pantin, H. Falcke, D. Rouan, R. Genzel, G. Hasinger, A. Goldwurm, F. Yusef-Zadeh, Y. Clenet, S. Trippe, P.-O. Lagage, H. Bartko, F. Eisenhauer, T. Ott, T. Paumard, G. Perrin, F. Yuan, T. K. Fritz, L. Mascetti, Evidence for X-Ray Synchrotron Emission from Simultaneous Mid-Infrared to X-Ray Observations of a Strong Sgr A* Flare, *ApJ* 698 (2009) 676–692. URL: <http://adsabs.harvard.edu/abs/2009ApJ...698...676D>. doi:10.1088/0004-637X/698/1/676. arXiv:0903.3416.
- [14] Gravity Collaboration, R. Abuter, A. Amorim, M. Bauböck, J. P. Berger, H. Bonnet, W. Brandner, Y. Clénet, V. Coudé Du Foresto, P. T. de Zeeuw, C. Deen, J. Dexter, G. Duvert, A. Eckart, F. Eisenhauer, N. M. Förster Schreiber, P. García, F. Gao, E. Gendron, R. Genzel, S. Gillissen, P. Guajardo, M. Habibi, X. Haubois, T. Henning, S. Hippler, M. Horrobin, A. Huber, A. Jiménez-Rosales, L. Jocou, P. Kervella, S. Lacour, V. Lapeyrère, B. Lazareff, J.-B. Le Bouquin, P. Léna, M. Lippa, T. Ott, J. Panduro, T. Paumard, K. Perraut, G. Perrin, O. Pfuhl, P. M. Plewa, S. Rabien, G. Rodríguez-Coira, G. Rousset, A. Sternberg, O. Straub, C. Straubmeier, E. Sturm, L. J. Tacconi, F. Vincent, S. von Fellenberg, I. Waisberg, F. Widmann, E. Wieprecht, E. Wiezorrek, J. Woillez, S. Yazici, Detection of orbital motions near the last stable circular orbit of the massive black hole SgrA*, *A&A* 618 (2018) L10. doi:10.1051/0004-6361/201834294. arXiv:1810.12641.
- [15] A. Nathanail, C. M. Fromm, O. Porth, H. Olivares, Z. Younsi, Y. Mizuno, L. Rezzolla, Plasmoid formation in global grmhd simulations and agn flares, arXiv e-prints (2020) arXiv:2002.01777. URL: <https://ui.adsabs.harvard.edu/abs/2020arXiv200201777N>. arXiv:2002.01777.
- [16] B. Ripperda, F. Bacchini, A. A. Philippov, Magnetic reconnection and hot spot formation in black hole accretion disks, *ApJ* 900 (2020) 100. URL: <https://ui.adsabs.harvard.edu/abs/2020ApJ...900...100R>. doi:10.3847/1538-4357/ababab. arXiv:2003.04330.
- [17] O. Porth, Y. Mizuno, Z. Younsi, C. M. Fromm, Flares in the galactic centre - i. orbiting flux tubes in magnetically arrested black hole ac-

- cretion discs, *MNRAS* 502 (2021) 2023–2032. URL: <https://ui.adsabs.harvard.edu/abs/2021MNRAS.502.2023P>. doi:10.1093/mnras/stab163. arXiv:2006.03658.
- [18] B. P. Abbott, R. Abbott, T. D. Abbott, F. Acernese, K. Ackley, C. Adams, T. Adams, P. Addesso, R. X. Adhikari, V. B. Adya, et al., GW170817: Observation of Gravitational Waves from a Binary Neutron Star Inspiral, *Physical Review Letters* 119 (2017) 161101. URL: <http://adsabs.harvard.edu/abs/2017PhRvL.119p1101A>. doi:10.1103/PhysRevLett.119.161101. arXiv:1710.05832.
- [19] A. Nathanail, R. Gill, O. Porth, C. M. Fromm, L. Rezzolla, On the opening angle of magnetized jets from neutron-star mergers: the case of grb170817a, *MNRAS* 495 (2020) 3780–3787. URL: <https://ui.adsabs.harvard.edu/abs/2020MNRAS.495.3780N>. doi:10.1093/mnras/staa1454. arXiv:2003.03597.
- [20] A. Nathanail, R. Gill, O. Porth, C. M. Fromm, L. Rezzolla, 3d magnetized jet break-out from neutron-star binary merger ejecta: afterglow emission from the jet and the ejecta, *MNRAS* 502 (2021) 1843–1855. URL: <https://ui.adsabs.harvard.edu/abs/2021MNRAS.502.1843N>. doi:10.1093/mnras/stab115. arXiv:2009.09714.
- [21] B. Ripperda, O. Porth, C. Xia, R. Keppens, Reconnection and particle acceleration in interacting flux ropes - II. 3D effects on test particles in magnetically dominated plasmas, *MNRAS* 471 (2017) 3465–3482. URL: <http://adsabs.harvard.edu/abs/2017MNRAS.471.3465R>. doi:10.1093/mnras/stx1875. arXiv:1707.08920.
- [22] K. Jahnke, C. Fendt, M. Founesneau, I. Georgiev, T. Herbst, M. Kaasinen, D. Kossakowski, J. Rybizki, M. Schlecker, G. Seidel, T. Henning, L. Kreidberg, H.-W. Rix, An astronomical institute’s perspective on meeting the challenges of the climate crisis, *Nature Astronomy* 4 (2020) 812–815. URL: <https://www.nature.com/articles/s41550-020-1202-4>. doi:10.1038/s41550-020-1202-4.
- [23] S. Portegies Zwart, The ecological impact of high-performance computing in astrophysics, *Nature Astronomy* 4 (2020) 819–822. URL: <https://ui.adsabs.harvard.edu/abs/2020NatAs...4..819P>. doi:10.1038/s41550-020-1208-y. arXiv:2009.11295.
- [24] Event Horizon Telescope Collaboration, K. Akiyama, A. Alberdi, W. Alef, K. Asada, R. Azulay, A.-K. Baczkó, D. Ball, M. Baloković, J. Barrett, D. Bintley, L. Blackburn, W. Boland, K. L. Bouman, G. C. Bower, M. Bremer, C. D. Brinkerink, R. Brissenden, S. Britzen, A. E. Broderick, D. Brogiere, T. Bronzwaer, D.-Y. Byun, J. E. Carlstrom, A. Chael, C.-k. Chan, S. Chatterjee, K. Chatterjee, M.-T. Chen, Y. Chen, I. Cho, P. Christian, J. E. Conway, J. M. Cordes, G. B. Crew, Y. Cui, J. Davelaar, M. De Laurentis, R. Deane, J. Dempsey, G. Desvignes, J. Dexter, S. S. Doeleman, R. P. Eatough, H. Falcke, V. L. Fish, E. Formalont, R. Fraga-Encinas, P. Friberg, C. M. Fromm, J. L. Gómez, P. Galison, C. F. Gammie, R. García, O. Gentaz, B. Georgiev, C. Goddi, R. Gold, M. Gu, M. Gurwell, K. Hada, M. H. Hecht, R. Hesper, L. C. Ho, P. Ho, M. Honma, C.-W. L. Huang, L. Huang, D. H. Hughes, S. Ikeda, M. Inoue, S. Issaoun, D. J. James, B. T. Jannuzi, M. Janssen, B. Jeter, W. Jiang, M. D. Johnson, S. Jorstad, T. Jung, M. Karami, R. Karuppusamy, T. Kawashima, G. K. Keating, M. Kettenis, J.-Y. Kim, J. Kim, J. Kim, M. Kino, J. Y. Koay, P. M. Koch, S. Koyama, M. Kramer, C. Kramer, T. P. Krichbaum, C.-Y. Kuo, T. R. Lauer, S.-S. Lee, Y.-R. Li, Z. Li, M. Lindqvist, K. Liu, E. Liuzzo, W.-P. Lo, A. P. Lobanov, L. Loinard, C. Lonsdale, R.-S. Lu, N. R. MacDonald, J. Mao, S. Markoff, D. P. Marrone, A. P. Marscher, I. Martí-Vidal, S. Matsushita, L. D. Matthews, L. Medeiros, K. M. Menten, Y. Mizuno, I. Mizuno, J. M. Moran, K. Moriyama, M. Moscibrodzka, C. Müller, H. Nagai, N. M. Nagar, M. Nakamura, R. Narayan, G. Narayanan, I. Natarajan, R. Neri, C. Ni, A. Noutsos, H. Okino, H. Olivares, T. Oyama, F. Özel, D. C. M. Palumbo, N. Patel, U.-L. Pen, D. W. Pesce, V. Piétu, R. Plambeck, A. Pop-Stefanija, O. Porth, B. Prather, J. A. Preciado-López, D. Psaltis, H.-Y. Pu, V. Ramakrishnan, R. Rao, M. G. Rawlings, A. W. Raymond, L. Rezzolla, B. Ripperda, F. Roelofs, A. Rogers, E. Ros, M. Rose, A. Roshaninshat, H. Rottmann, A. L. Roy, C. Ruszczyk, B. R. Ryan, K. L. J. Rygl, S. Sánchez, D. Sánchez-Argüelles, M. Sasada, T. Savolainen, F. P. Schloerb, K.-F. Schuster, L. Shao, Z. Shen, D. Small, B. W. Sohn, J. SooHoo, F. Tazaki, P. Tiede, R. P. J. Tilanus, M. Titus, K. Toma, P. Torne, T. Trent, S. Trippe, S. Tsuda, I. van Bemmelen, H. J. van Langevelde, D. R. van Rossum, J. Wagner, J. Wardle, J. Weintraub, N. Wex, R. Wharton, M. Wielgus, G. N. Wong, Q. Wu, A. Young, K. Young, Z. Younsi, F. Yuan, Y.-F. Yuan, J. A. Zensus, G. Zhao, S.-S. Zhao, Z. Zhu, J. Anczarski, F. K. Baganoff, A. Eckart, J. R. Farah, D. Haggard, Z. Meyer-Zhao, D. Michalik, A. Nadolski, J. Neilsen, H. Nishioka, M. A. Nowak, N. Pradel, R. A. Primiani, K. Souccar, L. Vertatschitsch, P. Yamaguchi, S. Zhang, First m87 event horizon telescope results. v. physical origin of the asymmetric ring, *ApJL* 875 (2019) L5. URL: <http://adsabs.harvard.edu/abs/2019ApJ...875L...5E>. doi:10.3847/2041-8213/ab0f43.
- [25] C. Alappat, J. Laukemann, T. Gruber, G. Hager, G. Wellein, N. Meyer, T. Wettig, Performance modeling of streaming kernels and sparse matrix-vector multiplication on a64fx, in: 2020 IEEE/ACM Performance Modeling, Benchmarking and Simulation of High Performance Computer Systems (PMBS), 2020, pp. 1–7. doi:10.1109/PMBS51919.2020.00006.