PUNCH4NFDI Report / WP 3.2
2023/001

Survey of Astrophysics Simulation Codes in Germany: An Initiative of the PUNCH4NFDI Consortium

Annika Hagemeier[1]*, Frank W. Wagner[1], Susanne Pfalzner[1], and Ramin Marx[2]

Abstract

PUNCH4NFDI is the consortium of particle and astroparticle physics, astrophysics, and hadronic and nuclear physics in the NFDI (national research data infrastructure), Germany. It aims at developing concepts and tools for efficient management of digital research products in fundamental physics and promotes the idea of FAIR data – which is to make scientific data sets findable, accessible, interoperable, and reusable. Here we concentrate on the aims and measures of PUNCH4NFDI in the context of astrophysics simulations. As a first step towards better understanding the software usage of the astrophysics simulation community in Germany, PUNCH4NFDI developed a 14-question survey. We distributed the survey through various channels (mailing list, conference flyer, personal communication) in 2022. In total, 130 computational astrophysicists responded to our survey. We found that predominantly codes able to simulate gravitational *N*-body problems and magnetohydrodynamics are used by the German astrophysics community. Computer programs typically associated with research in solar physics and numerical relativity turned out to be applied to a lesser degree. The degree to which the FAIR principles are already applied varies greatly. In many cases a basic software version is open access, however, the newest work is often based on modified and unreleased versions. The degree of practising the FAIR principles is often a question of available manpower. Some codes are developed and used by single local research groups, others by large research consortia spread around the globe. While smaller research groups are in principle willing to publish and openly share their simulation data, they often simply lack the manpower to do so. We also found that most astrophysicists in Germany view re-using other researchers' data sets as highly desirable.

Keywords
Astrophysics — Simulations — Survey

[1] *Jülich Supercomputing Centre, Forschungszentrum Jülich, Jülich, Germany*
[2] *Heidelberg-Königstuhl State Observatory, Heidelberg University, Heidelberg, Germany*
***Corresponding author**: a.hagemeier@fz-juelich.de

Contents

1. Introduction

Computer models have substantially changed scientific working methods and have long since become integral to modern research. The NFDI consortium of particle, astro-, astroparticle, hadron, and nuclear physics (PUNCH4NFDI; www.punch4nfdi.de) performed a survey to figure out which astrophysical simulation software is mainly used by the scientific community in Germany. The aim of this project is to make heavily used codes available in an easy-to-use form at several German HPC centres. The ultimate goal is to encourage scientists to share their codes, simulation results and diagnostics tools. For these codes, we will provide tools that support publishing all the information necessary to make simulation-based research compliant with the FAIR principles [1]. The FAIR principles stand for **f**indable, **a**ccessible, **i**nteroperable, and **r**eusable research data. A separate report on how we want to achieve this goal is in preparation.

Before improving on a situation, it is always advisable to determine the current status. Therefore, we designed an online survey and circulated it in the German astrophysical community to obtain an overview of the current situation and the needs and wishes with respect to simulation codes. In order to evaluate how far the code usage in Germany differs form the world-wide usage, we searched the astro-ph[1] at arXiv for these codes during the years 2020–2022. This was done by using a publication harvester for arXiv developed in task area 6 of PUNCH4NFDI. The first objective is to make available performance-optimized prototypes of these software packages on high-performance computing systems. First, we will implement these codes at two tier-1 HPC centres in Germany, i.e. the Jülich Supercomputing Centre (JSC) at Forschungszentrum Jülich (FZJ) and the Leibniz Supercomputing Centre (LRZ). In a second step, we plan that other HPC centres in Germany follow with similar services.

In summary, AREPO, GADGET, RAMSES, FLASH, and SWIFT were the most named astrophysics codes in our survey. However, the usage and development of simulation codes in Germany is diverse. On the one hand, there are codes developed and limited to a single research group; on the other hand, there are codes used by different research communities spread around the globe. Often a basic version of the codes is open access, but the researchers use more advanced versions, which still need to be made available. While, in principle, smaller research communities are often willing to share their codes and data, they often lack the human power to do so. Generally, many scientists would find re-using other researchers' results highly desirable.

2. Methods

The survey consisted of 14 questions grouped into three categories: (a) information on the simulation codes, (b) information on the simulation data, and (c) user information. In autumn 2022, we performed a test run at a meeting of the Virgo Consortium[2] to learn whether the survey works or requires modifications on specific points. After the test run, we circulated the survey via the mailing list of the German Astronomical Society[3] and distributed flyers at the 2022 Annual Meeting of the German Astronomical Society[4] in Bremen, Germany. However, personal emails to institute heads and collaborators of the survey team proved to be the most prolific method in convincing people to participate. In total, 130 scientists took part in the survey, 28 of them during the test run. 31 people responded to the call via mailing list. Flyers and personal correspondence added another 71 people. The majority of people (99) work in Germany. 13 participants answered only the first part of the survey, where we asked for information on the codes that are currently in use or will be in the future. They left the questions concerning data provision and personal information unanswered. For an overview of the complete survey see https://go.fzj.de/survey.

3. Results and Discussion

3.1 Simulation codes

In the first section of the survey we gathered information on computer simulations within the German astronomical community. The central question was which simulation software is used by astrophysicists in Germany. In addition, we inquired if these codes are open source, on which HPC systems they are executed, how they are parallelized, and which programs scientists anticipate to use in the future.

3.1.1 Q1: Which codes are you currently using?

We provided a list of 25 astrophysics codes that are currently used in HPC projects at FZJ and LRZ. We asked the survey participants to indicate whether they use any of those codes and add also software packages they use but where not part of this list. The participants added many programs that were not part of the original list, leading to a final list containing more than 70 different codes. On closer inspection it turned out that also some astrophysics codes and software packages that are primarily used for data processing in observational astronomy and/or not optimized for HPC systems were mentioned. These software tools we removed from the study. An example would be CASA – the Common Astronomy Software Applications package which is the primary data processing software for the Atacama Large Millimeter/submillimeter Array (ALMA) and NSF's Karl G. Jansky Very Large Array (VLA), and is frequently used also for other radio telescopes. Out of the remaining codes, about a dozen were mentioned more than twice, whereas the vast majority was only mentioned once. The original 70 codes also contained different developmental versions of certain codes, we grouped these together. Here, all available versions of GADGET and Athena are considered as a single code, respectively. Figure 1(a) shows the astrophysics codes that were mentioned more than 2 times. The top-5 codes used are AREPO, GADGET, RAMSES, FLASH, and SWIFT. Thus, these are the codes we will concentrate on first for providing additional services.

An interesting question, but difficult to obtain information about, is whether the codes used in Germany are also those

[1] https://arxiv.org/archive/astro-ph
[2] https://virgo.dur.ac.uk
[3] https://www.astronomische-gesellschaft.de/en
[4] https://ag2022.astronomische-gesellschaft.de

Table 1. Astrophysics simulation codes.

Name	Code site	ASCL record	References
AREPO	arepo-code.org	1909.010	[2, 3, 4]
Athena	github.com/princetonuniversity/athena-cversion	1010.014	[5]
Athena++	www.athena-astro.app	1912.005	[6, 7]
FLASH	flash.rochester.edu/site/flashcode	1010.082	[8]
GADGET-2	wwwmpa.mpa-garching.mpg.de/gadget	0003.001	[9, 10]
GADGET-4	wwwmpa.mpa-garching.mpg.de/gadget4	2204.014	[11]
NBODY6/6++	–	1102.006	[12, 13]
NBODY6++GPU	–	–	[14]
PLUTO	plutocode.ph.unito.it	1010.045	[15, 16]
POLARIS	portia.astrophysik.uni-kiel.de/polaris	1807.001	[17]
RAMSES	bitbucket.org/rteyssie/ramses	1011.007	[18]
SWIFT	swift.strw.leidenuniv.nl	1805.020	[19, 20]

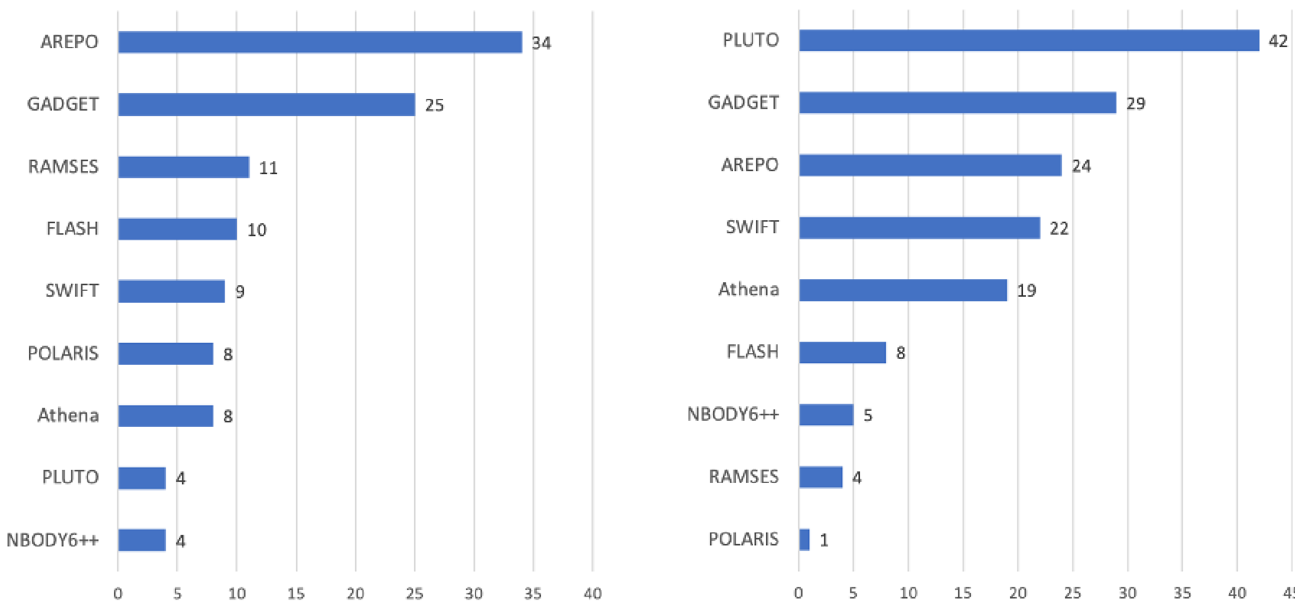

Figure 1. (a) Astrophysics code usage in Germany. (b) Number of arXiv citations.

most used world-wide. Therefore, we resorted to the number of citation referring to these codes as an indicator. Figure 1(b) shows the ranking of these codes according to their number of citations in the open-access repository arXiv. The numbers were obtained with the arXiv search tool[5] allowing a full text search of arXiv publications. For each astrophysics code it provides the number of publications that contain the code's repository or download link. Our search period covers the years 2020–2022. Comparing the rank order in Figure 1(a) and 1(b) reveals that the three codes PLUTO, RAMSES, and POLARIS reside on very dissimilar ranks. While PLUTO is heavily cited on arXiv, RAMSES and POLARIS are much less. PLUTO is either heavily used outside of Germany or has an extremely productive German user community. RAMSES and POLARIS are popular codes in the German astrophysics community, but much less visible abroad. At least the popularity of POLARIS in Germany is plausible given the fact that the code is developed and maintained at the University of Kiel.

For data protection reasons we did not ask the survey participants about their affiliation. Therefore, it is not possible to determine how many people from the same research group or institution completed the survey. This can of course lead to some biasing and might be an explanation why some codes are mentioned more often than others. That this issue can have an impact becomes obvious, when we compare the code usage with and without the test run data to each other (see Figure 2). 9 out of 130 survey participants stated that they use the SWIFT code. However, only 2 scientists use SWIFT without adding the 28 "test run" people to the survey results, clearly indicating that the majority of SWIFT users is from the test run.

[5] https://hess.science/arxiv

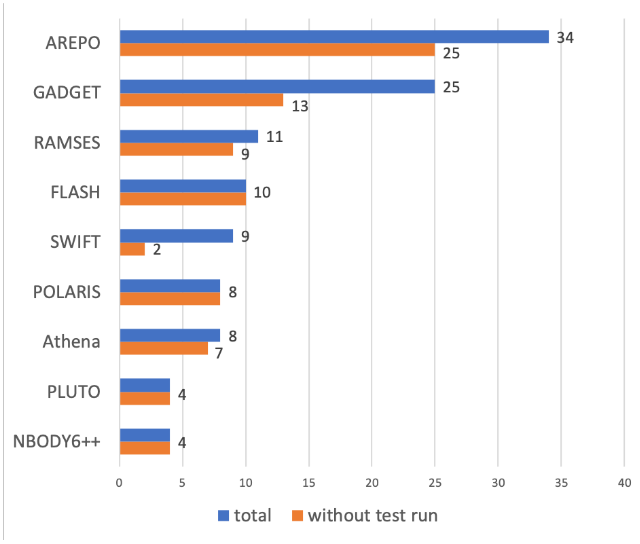

Figure 2. Astrophysics code usage with and without test run data.

3.1.2 Q2: Are these codes open source?

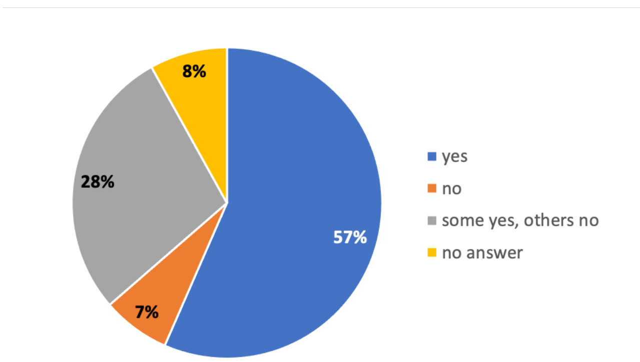

Figure 3. Percentage of scientists that use open source astrophysics codes.

Figure 3 shows the results regarding the question whether the codes that scientists are using are open source. 57% of the survey participants answered with „Yes“ and 28% chose at least the option „some yes, others no“. Only 7% of the participants answered with „No“. This leads us to the conclusion that most codes are indeed open source.

However, looking at it face-value might hide a serious problem that becomes obvious from the comments added by some of the survey participants. 10% of the people commented that a basic version of the simulation software is publicly available for download, but they use versions with proprietary modules or special features which are not yet part of a public repository or develop their own modules. Thus, the full reproducibility of scientific simulations is often hindered by not tracked modifications to open source programs. It is likely that the real percentage of not tracked modifications to open source software is higher because some of the participants answering with „yes“ might actually use such a modified version of the open source code.

3.1.3 Q3: On which HPC systems have you been running your simulation codes?

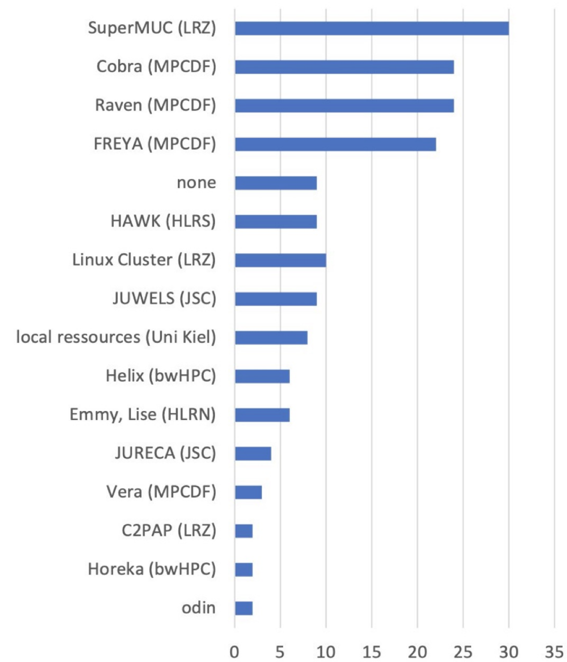

Figure 4. HPC clusters on which astrophysical codes are executed.

We provided a list of 9 German HPC clusters known to accept astrophysical projects to the survey participants. We did ask them to indicate whether they use any of those resources. In addition, the participants could name HPC systems not mentioned on this list. Figure 4 shows the HPC clusters named in the survey. SuperMUC at LRZ is the most commonly used system for astrophysical simulations. The cluster in Munich is closely followed by three different HPC systems at the Max Planck Computing and Data Facility (MPCDF) in Garching. JUWELS at JSC in Jülich was named less than a third as often as SuperMUC despite the two machines are comparable in computing power. It seems that the vicinity of the system in the Munich area seems to play an important role when deciding which system to use. A few university research centres focusing on computational astrophysics (e.g., Ludwig Maximilian University of Munich (LMU)) and a major non-university research organization (the Max Planck Institute for Astrophysics (MPA)) are located in southern Bavaria, Germany. It seems generally that often the local compute cluster is used rather than looking for the most efficient computer for the considered problem. Many participants listed their local compute resources as location where they run their codes.

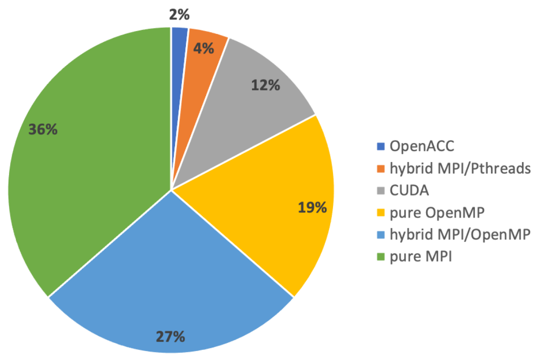

Figure 5. Astrophysics codes by parallelization strategy.

3.1.4 Q4: Which parallelization strategies are used in the codes that you are using?

Figure 5 shows the results regarding the question of the parallelization strategy that is used in codes used by the participants. The parallelization with pure MPI is the most common strategy directly followed by the hybrid strategy using MPI between the nodes of an HPC cluster and OpenMP on the shared memory nodes. Both parallelization paradigms are well suited to execute scalable simulation codes with a large number of processors. Thus, the users are aware of the advantages using these parallelization strategies. However, codes parallelized with pure OpenMP can only be run with a limited number of processors which is suitable for significant smaller simulations on clusters, workstations, or even laptops. Nevertheless, OpenMP is the third most frequently used parallelization option. Possible reasons could be that the simulations themselves are relatively small or that the limited local resources make OpenMP the only choice. CUDA for accessing GPUs is also used in some codes. Parallelization strategies such as Pthreads and OpenACC were rarely mentioned. It is also noticeable that 38% of the survey participants gave no answer to the question. The reason for this might be that these people simply use codes for performing simulation runs without knowing the underlying implementation details.

3.1.5 Q5: Which codes are you planning to use in the future?

In the last question concerning the simulation codes we wanted to know which computer programs are planned to be used in future projects of the participating scientists. Figure 6 shows the codes that were mentioned by more than one person. With 8 nominations AREPO is on the top followed by RAMSES with 4 mentions. SWIFT and GADGET are mentioned three times each and PLUTO, Nbody6++, FLASH, and Athena are mentioned two times each. Around 15 codes where mentioned only once, e.g. Rebound and Enzo. Some people gave a more general answer without naming a specific code. They are planning to develop their own code for example for using GPU acceleration or to model detectors in high energy astrophysics.

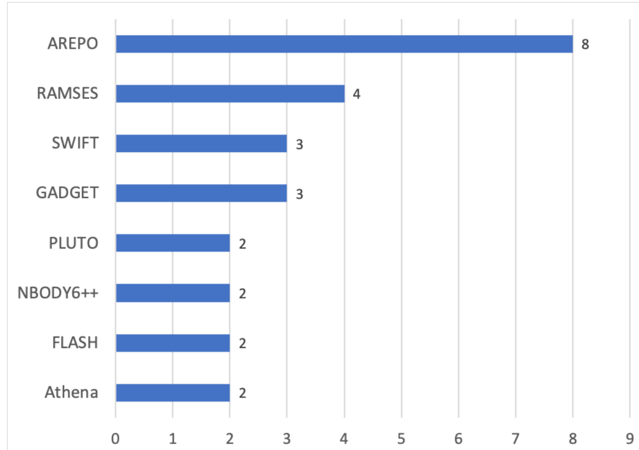

Figure 6. Astrophysics codes that are planned to be used in future projects.

3.2 Data management

Reproducibility of results is a fundamental aspect of scientific research. Therefore, we inquired in the second part of the survey about the data management of the simulation outputs. We wanted to know if simulation outputs are made publicly available and, if not, what are the main reasons for not publishing the output data. 10% did stop the survey at this point and did not finish the survey. So that the statistical significance is slightly reduced as only 117 people participated in this part.

3.2.1 Q6: Do you regularly make your simulation outputs (raw or post-processed) publicly available?

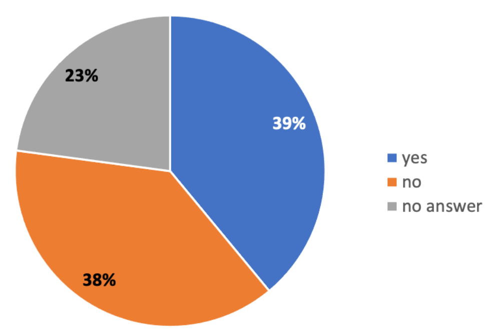

Figure 7. Percentage of scientists that share their simulation outputs publicly on a regular basis.

With the first question of the second part we wanted to find out how many scientists already publish their simulation outputs. Figure 7 shows the answers to question 6. It is illustrated that the ratio between publishing and not publishing is nearly equal, with 39% already publishing their simulation outputs on a regular basis and 38% currently not publishing their output data openly. Furthermore, 23% of the 117 participants did select the third option “no answer”.

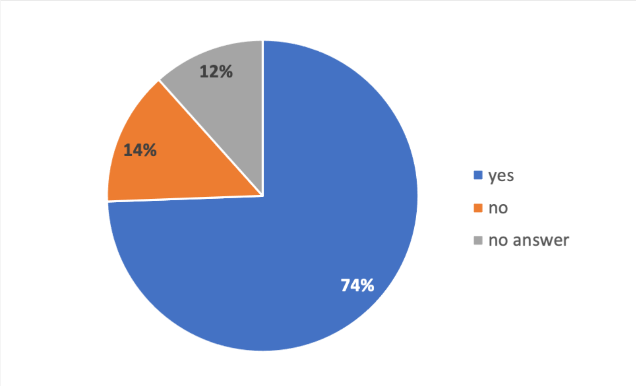

Figure 8. Percentage of scientists that are in principle willing to share their simulations outputs if an easy option is available.

3.2.2 Q7: Would you do that if there was an easy option?

This question was only shown to the fraction of survey participants who said that they do not make their simulation outputs publicly available. It should shed light on the reasons for not publishing these data. Figure 8 shows the answers to question 7. The majority or 74% of the 43 participants is willing to share their simulation outputs openly if publishing data becomes easier. Only 14% are not interested in publishing their simulation outputs even if an easy option would exist.

Furthermore, optional comments to this question gave some idea about how participants interpreted this question. People who are willing to publish their simulation outputs pointed out that they already publish them partly in the form of articles in scientific journals or post processed results in the form of plots or videos. However publishing raw data which are in most cases very large is difficult since no suitable infrastructure is currently available.

People who are not willing to publish their simulation results argued that most of the time only test runs are performed or only small or medium size simulations where there is generally no community interest in re-analyzing the simulations. If there would be an interest in the community they would think about making them public. This statement shows that even people who indicated that they are not willing to publish their simulation outputs are not strictly against publishing them. Instead if they see the need, i.e. if there is interest in the community, they would consider making them public. If there was an easy way to do it this may influence their motivation in a positive way.

One person who selected the "no answer" option remarked that they would in principle be willing to make raw simulation outputs public, but providing the meta data is an important aspect. However, they found that providing proper documentation of the initial settings and configurations is time consuming and can only be justified if there is an interest from the community.

3.2.3 Q8: Do you attach proper tags and metadata (e.g. code version, execution system, ...) to your data?

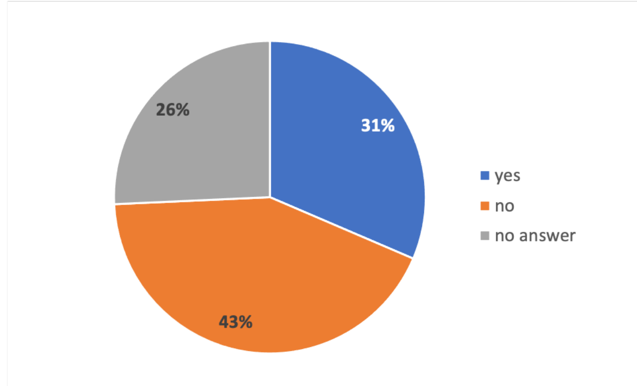

Figure 9. Percentage of scientists that attach proper tags and metadata when simulation outputs are published.

This question concerns the quality of the documentation provided for the published data. The statistical sample became again smaller because only people who publish their simulation results were included. Figure 9 shows the answers to question 8. 31% of 35 participants add proper tags and metadata to their simulation outputs, whereas the majority (43%) publishes their results without metadata or any additional descriptive information. Furthermore, 26% did prefer the third option "no answer".

In the comments section some scientists stated that they at least partially attach tags and metadata like e.g. the code version to their simulation outputs. Unfortunately, no one mentioned information such as the systems the codes were executed on, indicating that such information seem to be rather uncommon. It was said that some codes such as SWIFT collect metadata automatically. These metadata, including the simulation setup and output files, are dumped and documented in separate FITS files. Some scientists commented that they describe software and hardware details in their peer-reviewed publications.

One comment nicely summarizes the current situation: "We are working towards this and we acknowledge the need to make science data more usable to the community at large, but creating appropriate metadata is lots of work and often is neglected. Getting support here would be very helpful."

3.2.4 Q9: Would you be interested in using publicly available simulation outputs?

Figure 10 shows the answers to question 9. The majority or 79% of 99 participants is interested in using simulation outputs from other scientists that are publicly available. In the optional comments section several scientists indicated that they do already use simulation results provided by other researchers such as stellar evolution tracks and spectral libraries. A specific example that was mentioned are the data from the project IllustrisTNG[6]. Only 8% are not interested in publicly

[6] https://www.tng-project.org

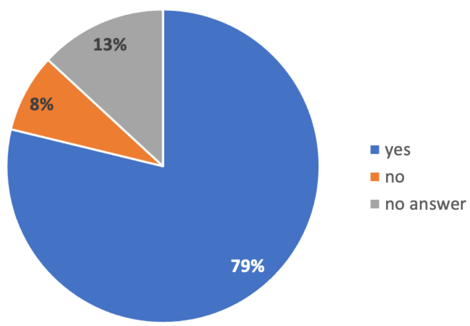

Figure 10. Percentage of scientists that are in principle interested in using publicly available simulation outputs.

available simulation outputs. One scientist commented that already too much data exist to be analyzed in a lifetime. Furthermore, 13% of the 99 participants did select the third option "no answer".

3.3 User information

In the last part of the survey we collected some statistics about the people who answered the survey. It emerged that 76% of the survey participants are currently working in Germany, 12% answered that they are working outside Germany and 12% selected the "no answer" option. 54% of the people who answered the survey are currently working at a research facility. 34% choose the answer "other" and again 12% selected the "no answer" option. The majority (70%) of the participants are male, 8% female, whereas 21% again selected the "no answer" option and 1% chose the "others" as answer.

We wanted to also obtain an impression about the HPC knowledge of the participants. Nearly half of the people (45%) stated that they are using HPC systems for more than 5 years, 15% for 3–5 years, 13% for 1–3 years and only 8% are beginners working with HPC systems for less than a year. 19% selected the "no answer" option. In the last question we asked the people to estimate their level of knowledge concerning HPC systems. 32% of the people judge their HPC knowledge as advanced. 25% have intermediate knowledge and 26% estimate their HPC expertise as basic. 2% chose the answer "I don't know" and 15% selected the "no answer" option.

4. Discussion

The survey results have to be taken with care because of (a) selection biases and (b) small number statistics.

4.1 Effects due to selection biases

Selection biases concerning the participating scientists stem from two sources. First, a strong selection effect occurred probably during the active mailing phase. While trying to cover all fields of computational astrophysics equally, some groups might have been missed out. However, the response rate of people knowing the sender of the survey invitation was higher than of those to whom the sender was unknown. The survey conductors are active in the fields of star and planet formation, galaxy dynamics and exoplanet research. Therefore, codes in these fields could be over-represented in question 1. Using the mailing list of the German Astronomical Society and the distribution of flyers at its 2022 annual meeting likely also introduced a bias. Experienced researchers are more likely to be members of the German Astronomical Society or attend the annual meeting than young researchers and students. Therefore, there might exist a general bias toward more experienced researchers. Second, the people responding to our survey represent only a fraction of the entire computational astrophysics community. For example, researchers directly contacted and/or people who personally knew the surveyers were more likely to take part in the survey.

5. Summary

In summary, AREPO, GADGET, RAMSES, FLASH, and SWIFT were the most named codes in our survey. However, the usage and development of simulation codes in Germany is diverse. On the one hand, there are codes developed and limited to a single research group; on the other hand, there are codes used by large communities spread around the globe. Often a basic version of the codes is open access, but many researchers use more advanced versions, which still need to be made available. While, in principle, the small communities are often willing to share their codes and data, they often lack the human power to do so. Generally, many scientists would find re-using other researchers' results highly desirable.

6. Outlook

We obtained much positive feedback for performing such a survey. It seems that many scientists are concerned about the current situation. They encourage us taking steps toward making simulation codes and data more open accessible. They particular look forward to PUNCH4NFDI providing tools that enable scientist to proceed themselves to making data and codes FAIRer.

Currently, we are working on a unified solution that makes the top-5 codes of the survey available on our tier-1 HPC systems at FZJ and LRZ. Regarding astrophysics codes in general, it is usually not very practicable to preinstall such programs in the software stack on an HPC system. The reason is twofold: (a) astrophysics codes are typically maintained by their developers on a part-time basis, therefore lack fixed release cycles known from commercial applications and (b) in most cases the software has to be compiled with a specific problem in mind; often no generic installation which would offer the entire performance range is possible. To mitigate these issues, our solution is to provide cluster-specific makefiles and installation scripts to assist scientists to perform automatic software installations on our compute facilities. Furthermore, we provide guides, working examples, and benchmark results to scientists to get started easily with the particular HPC cluster and to run first astrophysical simulations. Going forward, we will further offer tools to annotate simulation data and metadata assisting scientists to publish their research results according to the FAIR standard.

Acknowledgments

This work was performed in the framework of the PUNCH4NFDI consortium which will be supported by DFG fund "NFDI 39/1", Germany.

References

[1] Mark D. Wilkinson, Michel Dumontier, IJsbrand Jan Aalbersberg, Gabrielle Appleton, Myles Axton, Arie Baak, Niklas Blomberg, Jan-Willem Boiten, Luiz Bonino da Silva Santos, Philip E. Bourne, Jildau Bouwman, Anthony J. Brookes, Tim Clark, Mercè Crosas, Ingrid Dillo, Olivier Dumon, Scott Edmunds, Chris T. Evelo, Richard Finkers, Alejandra Gonzalez-Beltran, Alasdair J. G. Gray, Paul Groth, Carole Goble, Jeffrey S. Grethe, Jaap Heringa, Peter A. C. 'T Hoen, Rob Hooft, Tobias Kuhn, Ruben Kok, Joost Kok, Scott J. Lusher, Maryann E. Martone, Albert Mons, Abel L. Packer, Bengt Persson, Philippe Rocca-Serra, Marco Roos, Rene van Schaik, Susanna-Assunta Sansone, Erik Schultes, Thierry Sengstag, Ted Slater, George Strawn, Morris A. Swertz, Mark Thompson, Johan van der Lei, Erik van Mulligen, Jan Velterop, Andra Waagmeester, Peter Wittenburg, Katherine Wolstencroft, Jun Zhao, and Barend Mons. The FAIR guiding principles for scientific data management and stewardship. *Scientific Data*, 3:160018, 2016.

[2] Rainer Weinberger, Volker Springel, and Rüdiger Pakmor. The AREPO public code release. *The Astrophysical Journal Supplement Series*, 248(2):32, 2020.

[3] Volker Springel. *E pur si muove:* Galilean-invariant cosmological hydrodynamical simulations on a moving mesh. *Monthly Notices of the Royal Astronomical Society*, 401(2):791–851, 2010.

[4] Rüdiger Pakmor, Andreas Bauer, and Volker Springel. Magnetohydrodynamics on an unstructured moving grid. *Monthly Notices of the Royal Astronomical Society*, 418(2):1392–1401, 2011.

[5] James M. Stone, Thomas A. Gardiner, Peter Teuben, John F. Hawley, and Jacob B. Simon. Athena: A new code for astrophysical MHD. *The Astrophysical Journal Supplement Series*, 178(1):137, 2008.

[6] Christopher J. White, James M. Stone, and Charles F. Gammie. An extension of the Athena++ code framework for GRMHD based on advanced Riemann solvers and staggered-mesh constrained transport. *The Astrophysical Journal Supplement Series*, 225(2):22, 2016.

[7] Kyle Gerard Felker and James M. Stone. A fourth-order accurate finite volume method for ideal MHD via upwind constrained transport. *Journal of Computational Physics*, 375:1365–1400, 2018.

[8] Greg Weirs, Vikram Dwarkadas, Tomek Plewa, Chris Tomkins, and Mark Marr-Lyon. Validating the Flash code: Vortex-dominated flows. *Astrophysics and Space Science*, 298:341–346, 2005.

[9] Volker Springel. The cosmological simulation code GADGET-2. *Monthly Notices of the Royal Astronomical Society*, 364(4):1105–1134, 2005.

[10] Volker Springel, Naoki Yoshida, and Simon D. M. White. GADGET: A code for collisionless and gasdynamical cosmological simulations. *New Astronomy*, 6(2):79–117, 2001.

[11] Volker Springel, Rüdiger Pakmor, Oliver Zier, and Martin Reinecke. Simulating cosmic structure formation with the GADGET-4 code. *Monthly Notices of the Royal Astronomical Society*, 506(2):2871–2949, 2021.

[12] Sverre J. Aarseth. From NBODY1 to NBODY6: The growth of an industry. *Publications of the Astronomical Society of the Pacific*, 111(765):1333, 1999.

[13] Rainer Spurzem. Direct N-body simulations. *Journal of Computational and Applied Mathematics*, 109(1):407–432, 1999.

[14] Long Wang, Rainer Spurzem, Sverre Aarseth, Keigo Nitadori, Peter Berczik, M. B. N. Kouwenhoven, and Thorsten Naab. NBODY6++GPU: Ready for the gravitational million-body problem. *Monthly Notices of the Royal Astronomical Society*, 450(4):4070–4080, 2015.

[15] A. Mignone, G. Bodo, S. Massaglia, T. Matsakos, O. Tesileanu, C. Zanni, and A. Ferrari. PLUTO: A numerical code for computational astrophysics. *The Astrophysical Journal Supplement Series*, 170(1):228, 2007.

[16] A. Mignone, C. Zanni, P. Tzeferacos, B. van Straalen, P. Colella, and G. Bodo. The PLUTO code for adaptive mesh computations in astrophysical fluid dynamics. *The Astrophysical Journal Supplement Series*, 198(1):7, 2011.

[17] S. Reissl, S. Wolf, and R. Brauer. Radiative transfer with POLARIS - I. Analysis of magnetic fields through synthetic dust continuum polarization measurements. *Astronomy & Astrophysics*, 593:A87, 2016.

[18] R. Teyssier. Cosmological hydrodynamics with adaptive mesh refinement - A new high resolution code called RAMSES. *Astronomy & Astrophysics*, 385:337–364, 2002.

[19] Matthieu Schaller, Pedro Gonnet, Aidan B. G. Chalk, and Peter W. Draper. SWIFT: Using task-based parallelism, fully asynchronous communication, and graph partition-based domain decomposition for strong scaling on more than 100,000 cores. In *Proceedings of the Platform for Advanced Scientific Computing Conference*, pages 1–10, 2016.

[20] Matthieu Schaller, Josh Borrow, Peter W Draper, Mladen Ivkovic, Stuart McAlpine, Bert Vandenbroucke, Yannick Bahé, Evgenii Chaikin, Aidan B G Chalk, Tsang Keung Chan, Camila Correa, Marcel van Daalen, Willem Elbers, Pedro Gonnet, Loïc Hausammann, John Helly, Filip Huško, Jacob A Kegerreis, Folkert S J Nobels, Sylvia Ploeckinger, Yves Revaz, William J Roper, Sergio Ruiz-Bonilla, Thomas D Sandnes, Yolan Uyttenhove, James S Willis, and Zhen Xiang. Swift: a modern highly parallel gravity and smoothed particle hydrodynamics solver for astrophysical and cosmological applications. *Monthly Notices of the Royal Astronomical Society*, 530(2):2378–2419, 03 2024.