



Publication Year	2022
Acceptance in OA	2022-11-11T16:07:31Z
Title	All the Shades of the Cloud
Authors	TAFFONI, Giuliano, BERTOCCO, SARA, LANDONI, Marco, MOLINARO, Marco, RUSSO, Stefano Alberto
Handle	http://hdl.handle.net/20.500.12386/32718
Serie	ASTRONOMICAL SOCIETY OF THE PACIFIC CONFERENCE SERIES
Volume	532

All the Shades of the Cloud

Giuliano Taffoni,^{1,4} Sara Bertocco,¹ Marco Landoni,^{2,3} Marco Molinaro,¹ and Stefano Alberto Russo¹

¹*INAF - OATs, Trieste, Italy ; giuliano.taffoni@inaf.it*

²*INAF - O.A. Cagliari, Italy*

³*INAF - O.A. Brera, Milano, Italy*

⁴*IFPU Institute for Fundamental Physics of the Universe, Trieste, Italy*

Abstract. Cloud computing is a powerful technology that in the last decade revolutionised computing and storage in particular for Industry and Private Sectors. Today, large investments are ongoing to build Cloud Infrastructures at National or International level (e.g. the European Open Science Cloud initiative). Also, scientists are approaching commercial and private Clouds at different scales: single researchers test the Clouds for small research projects, at the same time large international collaborations are evaluating Cloud technology to collect, process, analyse, archive and curate their data.

In this paper, we discuss the use of Cloud in Astrophysics at different scales using some examples and we present future trends and possibilities that the use of Cloud computing and its convergence with high performance computing will open: from high-end data analysis to high performance data analytics, from scientific computing to data analytics.

1. Introduction

The experience of large scientific instruments in the last years demonstrates how experiments and observations are critically dependent on computing, data processing and storage infrastructures and our ability to utilise them through codes and algorithms.

The explosive growth in data generation today happens (and will happen in the future) in “edge environments”. These include major scientific instruments, experimental facilities, and satellites, but also the incredible welter of digital data generators from Internet of Things, social networking, smart cities and so on. The size and complexity of data is so large to open new technical challenges in data acquisition from “the edge” to data centers, data storage and post-processing. According to CISCO, by 2021, around 1000 Exabytes of data would be produced and stored in Cloud data centers and half of them will be processed in Cloud in using high-end data analysis and analytics tools and methodologies (CISCO 2020).

Also scientific experiments on large-scale instruments require significant data reduction, and updates of these instruments will produce orders of magnitudes more data in the near future. For example, the National Institutes for Health’s Brain Initiative focuses on high-throughput x-ray tomography of whole mouse brains producing hundreds of about 160TB size images to transport, reduce, analyse and store. Due to

the size and complexity of those data, High Performance Computing (HPC) facilities, methodology and algorithms have been already adopted for processing and analysis.

Also, in Astronomy and Astrophysics (A&A), new instruments (e.g. the Square Kilometer Array – SKA –, the Cherenkov Telescope Array – CTA –, the Extremely Large Telescope – ELT –, the James Webb Space Telescope – JWST –, the Euclid satellite, the Low Frequency Array – LOFAR –) and the advent of multi-messenger astronomy require significant high-end data reduction with exascale computing resources and methodologies typical of HPC environments (Taffoni et al. 2020b): distributed memory systems, hard platform optimization, use of accelerators, HPC system SW and libraries and more. At the same time accelerated Cosmological N-Body hydrodynamic codes (e.g. OpenGADGET, GADGET4, RAMSES) can generate 20 petabytes of data out of a single 10000^3 particles simulation to further post-process and compare with observed data.

This will lead to the use of new infrastructures where exascale HPC and Clouds will converge to answer new challenges of (Big-) data analysis and (Big-) data analytics (HPDA). New technologies (e.g. containerization) are driving the convergence of these “worlds” and the advent of Science Platforms (SPs) as a means to access data, storage and computing (but also SW and algorithms) is facilitating the use of novel Cloud infrastructures at different scales.

In this paper, we review the use of the Cloud in A&A with some examples and we introduce new challenges in Cloud computing and HPC arising from the novel requirements for HDA and HPDA. We discuss the value of Cloud and HPC convergence and we present the efforts in this direction based on novel Cloud models as Analytics-as-a-Service or HPC-as-a-Service.

The paper is organized as follows. In Sect. 2 we summarize Cloud computing concepts and deploying models used in science. We mainly focus our discussion on the role of Commercial Clouds for science and engineering. Sect. 2.1 focuses on the use of Public Cloud infrastructures by Astronomers with some practical example. In Sect. 3 we discuss the role of HPC and Clouds for scientific computing and data analysis and the differences and analogies between the two platforms. The importance and role of the Cloud and HPC convergence is discussed in Sect. 4, where we introduce a proof of concept of HPC-as-a-Service Cloud. Finally, we discuss the role of Clouds and SPs to facilitate the use of computing resources and storage for A&A.

2. Cloud Computing for science and technology

Cloud computing offers scalable, reliable, elastic computing and storage services. The resources used for these services can be requested on-demand and metered so users can be charged only for the resources they use. Data is stored closer to the computing and analysis site and in such a way that it is device and location independent, thus simplifying software (SW) development and data processing.

The Cloud services operated by Amazon, Google, Microsoft and other vendors are commonly referred to as public Clouds (by analogy with the public utilities) or even better as commercial Clouds. In contrast, a private Cloud is operated by a private institution(s) to provide computing, storage, and other services to a more limited audience, for example a company or a specific community (community Cloud) as the Astrophysics one. Of particular interest for Astronomers is the case of Open Science Clouds: Academic Clouds that embrace the Open Science concepts.

More recently, containerization technologies enable developers to shape their application's computing environment and encapsulate SW that can be deployed on Cloud infrastructures. Containerization also enhance scientific application portability, allowing workflows and SW to be packaged, shared, and redeployed without complex and often arduous configuration. Containers have emerged as a viable delivery platform for SW also on HPC environments (Ruiz et al. 2015), demonstrating to lower the usage barrier of complex HPC computing platforms and contributing to SW preservation and experiment reproducibility (coupling a specific SW environment with data and computing infrastructures). As we will discuss in Sect. 4, containerization technology is playing a central role in the HPC and Cloud convergence.

2.1. Examples of Cloud Computing in Astronomy

Scientists are using commercial and private Clouds to explore scientific data and run computer-based simulation and modeling (O'Driscoll et al. 2013; Hoffa et al. 2008; Wiley et al. 2011) stimulating also the deployment of academic and Open Science Clouds (European-Commission 2016).

The growing interest of the A&A community in Cloud computing is witnessed by papers that discuss tests (Bertocco et al. 2018; Berriman et al. 2010; Landoni et al. 2019a; Sabater et al. 2017; Timmes et al. 2020) or production environments (Hayden 2021; Sciacca 2021) developed in the Clouds. What connects many of those experiments is a similar use of the service model and approach: the development of Infrastructure-as-a-Service (IaaS) Clouds (most often based on containers) for computing, usually based on the use of spot instances. The spot instances represent Cloud excess capacity which a Cloud provider need to have available for any surges in customer demand. Public Cloud providers offer this excess capacity at a massive discount to drive usage. The drawback in using of spot resources is that providers can terminate spot instances with just a short term warning to reassign them to other customers.

The National Institute for Astrophysics in Italy promoted a Cloud test campaign focused on two commercial Cloud Platforms, namely Google Cloud Platform and Amazon Web Services (AWS), by offering to the Italian A&A community the possibility to exploit the computational power of the two services. For the former, we have proposed scenarios ranging from implementation of Workflow-as-a-Service aiming to offer reduction pipelines (in the context of Exoplanetary spectroscopy) in an on-demand fashion to small HTC clusters with GPU capabilities (mainly devoted to Adaptive Optics simulations). The platform has demonstrated to be resilient and scalable for HTC-based projects offering a viable solution for real-life astronomical application while the performance of HPC tasks were rather poor (Landoni et al. 2019b).

The experiments on AWS are still ongoing, although we already successfully tested a number of HTC applications, frequently coupled with Containers and related SW for their orchestration, like Amazon Batch. Among them, it is worth mentioning the implementation of Windows-based high end-nodes for ray-tracing in the context of the key project MAORY (Diolaiti et al. 2016) for the European Southern Observatory ELT telescope. This scenario foresees the use of high-end Virtual Machines (up to 64 vCPU and 0.5TB of RAM) to perform optical and mechanical analyses for the MAORY AO system, essentially triggered by project milestones (e.g. Design Reviews Phases). This prototypical scenario is a common practice in the context of the design of astronomical instrumentation and demonstrated that the use of Commercial Clouds (which offers HTC capabilities in non-Unix environment) could greatly reduce the cost of the

ownership (including licensing of non-Unix OS and SW) when the use of the resources is sporadic, and strictly coupled with the phases of a project.

3. The divergence of HPC and Data Analysis Infrastructures

In the last decades HPC and high-end data analysis infrastructures evolved independently driven by distinctly different optimization criteria to answer to different scientific and technical requirements. While the high-end data analysis requirements and technologies evolved in the Cloud environment, the HPC community focused on theory and numerical experiments (Asch et al. 2018).




Figure 1. Time evolution of the performance of the 500 most powerful HPC platforms expressed in Floating Point Operations per Seconds and measured with HPL benchmark. Blue line and squares are the average resources' performance, yellow line and triangles are the first more powerful resource performance, green line and circles are the overall performance of the whole 500 machines.

Major technical differences between the HPC and the data analysis ecosystems include SW development paradigms and tools, virtualization and scheduling strategies, storage models, resource allocation policies, and strategies for redundancy and fault tolerance. User requirements also differ. The HPC communities emphasize the efficiency of their infrastructures: for example application execution efficiency, system utilization

rate, and/or energy efficiency (Taffoni et al. 2019). On the other hand, Cloud and Big data communities tend to emphasize user experience and scalability.

The data center is the core facility for both HPC and Cloud infrastructures. The top500 project (<https://top500.org>) provides the performance of the 500 most powerful HPC platforms as shown in Figure 1. It is not possible to find similar information for public Cloud infrastructures, however we made an indirect estimate based on power consumption of Cloud data centers.

Cloud infrastructures are based on coordinated distributed facilities called “regions”. Each region is composed by multiple data centers in a redundancy group. To optimize efficiency and management, redundancy groups are based on 3 data centers minimum. For example, when Amazon AWS deploys a new region, it is usual that three new data centers be opened. The size of each data center depends on a cost-risk balance: the cost savings from scaling a single facility are logarithmic, whereas the negative impact of blast radius is linear. For this reason, AWS currently elects to build right around 32MW size facilities. Considering that the power consumption of the Riken Fujitsu Supercomputer Fugaku is 30MW and its peak performance is 442 PFlop/s (Kodama et al. 2020), we can extrapolate that **each AWS data center has more than 100 PFlop/s computing capacity** and more than 300 PFlop/s per region.

As the region scales, the number of data centers can easily escalate to far beyond ten. AWS already has regions scaled far beyond 10 data centers of 32MW, consequently we estimate that at least some AWS regions are able to offer more than Exascale computing capacity.




Figure 2. OSU MPI benchmark (5.3.2) for interconnect latency using OpenMPI 4.0. We compare the results of a real cluster with Infiniband 54 Gbps HDR and the 10 Gbps Ethernet, with the one of Azure HPC VM. Azure tests are executed on two VM nodes equipped with 200 Gbps HDR Infiniband. The results on the Cloud deployment are one order of magnitude higher (worse) than cluster Infiniband ones.

Even if from the hardware point of view the two environments are strikingly similar, in HPC the focus is to offer a unique homogeneous platform within a shared memory model, high performance interconnect and high performance parallel storage and accelerators (e.g. graphic processing units, or vector processing units). Clouds resources are designed for increasing proliferation of hardware, be it in processors, memory size, disk size and capabilities, accelerators, field-programmable gate arrays

(FPGAs), and more, to target divergent applications and computing scenarios. For example, we are witnessing the emergence of numerous hardware and system architectures for supporting deep learning based on FPGAs (Shawahna et al. 2018) used for advanced data analysis and analytics. FPGAs are also used to accelerate computations in numerical Astrophysics (Goz et al. 2020). Cloud infrastructures are designed to partition the facilities to offer resources to a large number of heterogeneous customers, including small and large enterprises more than offering a single supercomputer for exascale applications.

4. Big data: the Convergence of HPC and Data Analysis Infrastructures

The value of HPC and high-end data analysis infrastructures (in particular Cloud) convergence has been already recognized by different authors and authorities. In the European Data Infrastructure and the European Open Science Cloud (EOSC) position papers (Eudat et al. 2015) HPC has been completely absorbed into a broader digital single market and only appears as a component in this global system. Similarly, USA, Cina and Japan emphasized the importance of big data, data analysis, data analytics and HPC ecosystem convergence (AAVV 2015).

To produce new scientific results, HPC facilities must become nodes of a more complex network of Cloud resources able to support workflows consisting of multiple and divergent applications (e.g. classical HPC, machine learning, data analysis and visualization) to combine in the same environment both simulations and data reduction. In this scenario, moving data from an infrastructure to another is unfeasible as well as downloading data from the Cloud data lake to process with a custom SW: Astronomers must move computation close to data.

Commercial Clouds already recognize the need to deploy HPC Cloud resources. Microsoft Azure and AWS are offering Virtual Machines for High-Performance Computing equipped with a low latency/high throughput interconnect and GPU accelerators, on demand. The resources are accessible in a IaaS mode, where users can deploy an MPI based SLURM cluster with dynamic node allocation, to configure with their own SW and libraries. The performance of such infrastructures are very close to the one of a standard HPC Linux cluster (Figure 2).

4.1. A proof of concept of HPCaaS

An HPC-as-a-service (HPCaaS) approach enables an entire HPC facility up and running without the complexity of setting up a local HPC infrastructure, and to have it running close to the data. However, such approach requires to move the computation to the HPC platform itself, and there are a number of factors to consider in this respect.

First and foremost, the way in which to run and interact with the SW must not be forced to a specific environment (i.e. a Jupyter Notebook), since it may greatly vary in terms of usage patterns. Command line and GUI-based desktop applications are very common in the HPC domain, and in particular for data processing pipelines and interactive data analysis. In second place, there must be a reliable and secure way to bring and execute custom SW on the HPC platform, as the paradigm to move the computation close to the data does not (and should never) mean to restrict the computation to the SW available on the platform. The third aspect to consider is how the growing SPs concept relates a HPCaaS, as they usually target simpler use cases (i.e. use a data analysis library inside a Jupyter Notebook environment).

Software containerisation and microservices can provide a solution for these aspects. An architecture based on this approach has been proposed for the SP use-case in this volume in (Russo 2021), and is reported here in figure 3.




Figure 3. A SP architecture for running containerised user task microservices. This architecture is based on four main blocks: 1) a back-end with a web-based user interface; 2) a scheduler; 3) an agent; and 4) a registry for the microservice containers from where to pull them.

On top of this architecture, and in the framework of EU funded project ESCAPE, we are developing a tool that is a proof of concept of a HPCaaS platform and that respond to the requirements outlined in the previous paragraphs. The tool has been designed to execute containerized workloads (in Docker or Singularity) on a number of different computing elements that an organization can offer to the users, which can select the most appropriate one based on their needs: single computing nodes, HTC and HPC clusters, even public Clouds. Full remote desktop access, as well as web-based applications and custom interaction modes (i.e. SSH) are supported. Users can set up and run their own containers with their SW autonomously, as long as the container expose a viable interface on a TCP/IP protocol, thus bringing great flexibility. See Figure 4 for an example of a real user task.




Figure 4. Our prototype in action using Astrocook, a GUI-based quasar spectra analysis SW, running on a HPC cluster.

Emphasis should be given to the value of this approach to couple user’s own SW with computing resources into a task and mask the complexity of the use of the resources. Moreover, being a component of the ESCAPE SP (Bertocco 2021), that will also offer a data lake, in the future we will be able also to couple SW, computing and data. Our prototype is now used on our HPC clusters to make interactive data reduction for the LOFAR Italian community or as shown here to make spectral analysis using the Astrocook SW.

5. Conclusion and discussion

Cloud computing provides the infrastructure that is powering key digital trends such as mobile computing, the Internet of Things, big data, and artificial intelligence, thereby accelerating industry dynamics, disrupting existing business models, and fueling the digital transformation.

Recently, we have seen a rapid growth of the use of Cloud computing also for scientific research in particular in the field of high-end data analysis and analytics. Also, the A&A community is approaching with success Cloud technologies, in particular in the framework of IaaS or Container-as-a-service. On the other side, many applications and tools in Astronomy could be offered as SaaS, for example tools designed to visualize astronomical images or to manipulate and visualize data (e.g. Aladin, Topcat, Vizier, CASA). This could be particularly effective when SaaS applications access data stored on Cloud based data lakes (e.g. EOSC data lake) or thought IVOA standards¹. A research activity focused on SaaS in A&A deserves further and significant investments.

The size and complexity of data produced on the “edge” by new instruments is guiding scientific communities towards combining HPC and Cloud infrastructures, applications and methods into a common Cloud environment able to process large-scale analysis pipelines for data generated by simulations, experiments or observations. Exascale supercomputers must be viewed as the most important nodes in the very large network of heterogeneous computing and storage resources offered as Cloud services for example using a HPC-as-a-Service model. This approach requires coupling computing resources and SW into running tasks hiding all the complexity of setting up a super computing cluster as we implement in our proof of concept in Sect. 4.1.

Cloud platforms (Commercial, but also Private and Open Science) require skills and expertise far beyond the ones of Astronomers in particular for more complex set up like a HPC cluster, however they can simplify the development of SP designed to offer users a smoother experience when interacting with data and computing resources. The SP provides the framework for combining computing resources with SW, tools and services “marketplaces” or users custom SW, facilitating the achievement of “computing close to data”. To obtain such infrastructure, it is necessary to implement data lakes at the “bottom” of Cloud infrastructures.

Public Clouds are already offering some of the Data Lake capabilities based on object storage or other kind of distributed file system like HDFS. However, a Data Lake for science should integrate community (IVOA) standards for data annotation, access and preservation and combine the capabilities of commercial data Clouds with the requirements of the scientific communities. For example, a Data Lake should be

¹<https://www.ivoa.net>

able to facilitate data acquisition from the “edges” and their analysis. At the same time they should manage the transition between *cold* data sets into *hot* data sets: data set that suddenly become particularly interesting for the community and that therefore require fast reprocessing and continuous access. This may happen in case of some extraordinary Astrophysical event (e.g. a supernova explosion). A Data Lake should be able to seamlessly manage this transition for example replicating the data to improve availability or access from a large number of computing resources.

In conclusion, Commercial Clouds may become an important asset for scientific research. Researchers have usually access to a set of resources: national computing centres, institutional computing resources or private Clouds, local workstation and so on. Commercial Clouds may become an addition to this portfolio, particularly useful if we are able to identify what are the best use cases for their usage. This can be achieved using specific Cloud Economy metrics, however as general rule, Cloud is cost-effective when the use of the resources is sporadic and a limited amount of data is produced and stored in the Cloud.

This is an approach that we are pursuing in INAF, where Astronomers may institutionally request access to national super computing center – CINECA –, HPC and HTC local clusters – CHIPP – (Taffoni et al. 2020a) and AWS. INAF is offering these resources to the Astronomers on a proposed based evaluation since 2018.

Acknowledgments. This work was supported by the European Science Cluster of Astronomy and Particle Physics ESFRI Research Infrastructures project, funded by the European Union’s Horizon 2020 research and innovation programme under Grant Agreement no. 824064. We acknowledge the computing centre of INAF-Osservatorio Astronomico di Trieste, (Bertocco et al. 2019) and the CHIPP project, for the availability of computing resources and support.

References

- AAVV 2015, National strategic computing initiative executive order, Tech. rep., Executive Office of the U.S. President. URL <https://www.whitehouse.gov/the-press-office/2015/07/29/executive-order-creating-national-strategic-computing-initiative>
- Asch, M., Moore, T., Badia, R., et al. 2018, *The International Journal of High Performance Computing Applications*, 32, 435
- Berriman, G. B., Juve, G., Deelman, E., Regelson, M., & Plavchan, P. 2010, in *2010 Sixth IEEE International Conference on e-Science Workshops (IEEE)*, 1
- Bertocco, S. 2021, in *ADASS XXX*, edited by J. E. Ruiz, F. Pierfederici, & P. Teuben (San Francisco: ASP), ASP Conf. Ser.
- Bertocco, S., Dowler, P., Gaudet, S., Major, B., Pasian, F., & Taffoni, G. 2018, *Astronomy and Computing*, 24, 36
- Bertocco, S., Goz, D., Tornatore, L., Ragagnin, A., Maggio, G., Gasparo, F., Vuerli, C., Taffoni, G., & Molinaro, M. 2019, in *Astronomical Data Analysis Software and Systems XXIX*. arXiv:1912.05340
- CISCO 2020, *Global Network Trends Report*, Tech. rep., CISCO. URL https://www.cisco.com/c/dam/m/en_us/solutions/enterprise-networks/networking-report/files/GLBL-ENG_NB-06_0_NA_RPT_PDF_MOFU-no-NetworkingTrendsReport-NB_rpten018612_5.pdf
- Diolaiti, E., Paolo, C., Abicca, R., et al. 2016, in *Adaptive Optics Systems V*, edited by E. Marchetti, L. M. Close, & J.-P. Văcăŕan, International Society for Optics and Photonics (SPIE), vol. 9909, 768

- Eudat, Liber, OpenAIRE, Egi, & Geant 2015, European Open Science Cloud for research, Tech. rep., European Commission. URL http://libereurope.eu/wp-content/uploads/2015/11/OSC_Position_Paper-final-30.10.15.pdf
- European-Commission 2016, European Open Science Cloud, Open Science – Research & Innovation, Tech. rep., European Commission. URL <http://ec.europa.eu/research/openscience/index.cfm?pg=open-science-clou1>
- Goz, D., Ieronymakis, G., Papaefstathiou, V., Dimou, N., Bertocco, S., Simula, F., Ragagnin, A., Tornatore, L., Coretti, I., & Taffoni, G. 2020, *Computation*, 8, 34
- Hayden, B. 2021, in *ADASS XXX*, edited by J. E. Ruiz, F. Pierfederici, & P. Teuben (San Francisco: ASP), ASP Conf. Ser.
- Hoffa, C., Mehta, G., Freeman, T., Deelman, E., Keahey, K., Berriman, B., & Good, J. 2008, in *2008 IEEE Fourth International Conference on eScience*, 640
- Kodama, Y., Odajima, T., Arima, E., & Sato, M. 2020, in *2020 IEEE International Conference on Cluster Computing (CLUSTER)*, 484
- Landoni, M., Romano, P., Vercellone, S., Knödlseider, J., Bianco, A., Tavecchio, F., & Corina, A. 2019a, *The Astrophysical Journal Supplement Series*, 240, 32
- Landoni, M., Taffoni, G., Bignamini, A., & Smareglia, R. 2019b, in *Astronomical Data Analysis Software and Systems XXIX*. arXiv:1903.03337
- O’Driscoll, A., Daugelaite, J., & Sleator, R. D. 2013, *Journal of Biomedical Informatics*, 46, 774
- Ruiz, C., Jeanvoine, E., & Nussbaum, L. 2015, in *Euro-Par 2015: Parallel Processing Workshops*, edited by S. Hunold, A. Costan, D. Giménez, A. Iosup, L. Ricci, M. E. Gómez Requena, V. Scarano, A. L. Varbanescu, S. L. Scott, S. Lankes, J. Weidendorfer, & M. Alexander (Cham: Springer International Publishing), 813
- Russo, S. 2021, in *ADASS XXX*, edited by J. E. Ruiz, F. Pierfederici, & P. Teuben (San Francisco: ASP), ASP Conf. Ser.
- Sabater, J., Sánchez-Expósito, S., Best, P., Garrido, J., Verdes-Montenegro, L., & Lezzi, D. 2017, *Astronomy and computing*, 19, 75
- Sciacca, E. 2021, in *ADASS XXX*, edited by J. E. Ruiz, F. Pierfederici, & P. Teuben (San Francisco: ASP), ASP Conf. Ser.
- Shawahna, A., Sait, S. M., & El-Maleh, A. 2018, *IEEE Access*, 7, 7823
- Taffoni, G., Becciani, U., Garilli, B., Maggio, G., Pasian, F., Umana, G., Smareglia, R., & Vitello, F. 2020a, in *Astronomical Data Analysis Software and Systems XXIX*. arXiv:2002.01283
- Taffoni, G., Goz, D., Tornatore, L., Frailis, M., Maggio, G., & Pasian, F. 2020b, in *Astronomical Data Analysis Software and Systems XXVII*, edited by P. Ballester, J. Ibsen, M. Solar, & K. Shortridge, vol. 522 of *Astronomical Society of the Pacific Conference Series*, 49
- Taffoni, G., Tornatore, L., Goz, D., Ragagnin, A., Bertocco, S., Coretti, I., Marazakis, M., Chaix, F., Plumidis, M., Katevenis, M., Panchieri, R., & Perna, G. 2019, in *2019 15th International Conference on eScience (eScience)*, 403
- Timmes, F., Townsend, R., & Bildsten, L. 2020, arXiv preprints. arXiv:2001.02559
- Wiley, K., Connolly, A., Gardner, J., Krughoff, S., Balazinska, M., Howe, B., Kwon, Y., & Bu, Y. 2011, *Publications of the Astronomical Society of the Pacific*, 123, 366