INAF USCVIII - Calcolo Critico

Thursday, 15 June 2023 - Friday, 16 June 2023

Dipartimento di Fisica e Astronomia "Ettore Majorana"Universitá degli Studi di Catania Via S. Sofia, 64,

Book of Abstracts

ii

Contents

The MPI+CUDA Gaia AVU–GSR Parallel Solver towards next-generation Exascale Infras- tructures	1
PLUTO	1
"A new(ish) tool on the workbench": new version of the population synthesis code SEVN.	1
SuperBAM: characterization of Supernova transients through Bayesian Analytic Modeling	2
Unveiling the Complexity through High Performance Computing: the Link between Mas- sive Stars, Core-Collapse Supernovae, and their Remnants	2
Numerical Astrophysics at INAF-OAPa: an overview and future perspectives	3
Computational infrastructure for the Monitoring System of the Cherenkov Telescope Array	3
Computational infrastructure for the Monitoring System of the Cherenkov Telescope Array	4
Modeling young pulsar wind nebulae	4
Understanding the magnetic field evolution in supernova remnants: a crucial role of high- performace computing	5
MHD numerical simulations of nanojets and nanoflares in the solar corona for MUSE $\ . \ .$	5
HPC MHD modelling of unstable reconnecting plasma in the solar corona and EUV diagnostics with the MUSE mission	6
Bow-Shock Pulsar Wind Nebulae: a tale of trails	6
Relativistic Thermodynamics in CFD codes.	7
Challenges in forthcoming CMB data sets	7
The Arxes code suite: high-performance planet formation and astrochemistry \ldots .	7
Development of radio source detection and classification tools for SKA precursors \ldots	8
(INVITED) The Universe in a Box	9
High performance visualization for Astronomy & Cosmology	9

"From Vlasov-Poisson to Schr ⁻ odinger-Poisson: dark matter simulation with a quantum variational time evolution algorithm"	10
(INVITED) The Gaia legacy. From data reduction and analysis to data management and exploitation: HPC, HTC and Big Data issues.	10
GalaPy, the highly optimised C++/Python spectral modelling tool for galaxies $\ . \ . \ .$	11
(INVITED) Present and Future of Data E-Infrastructure to support computing in INAF $$.	11
Broad-band radio observations in the SKA era: the case of the galaxy group Nest200047	12
The cosmic dance	12
Energy efficiency in Data Reduction for Imaging in a Radio Astronomy pipeline	13
"The Dark Energy and Massive Neutrino Universe" cosmological simulations	13
The OPS4: towards a legacy Big Data system - A detailed view	14
Reducing MeerKAT data of the Galactic Plane	14
(INVITED) Computation in space missions: the Euclid case	14
Remote Visualization of Big Data: VisIVO as a Visualization Prototype for SKA Regional Centres	15
(INVITED) Addressing Computing Challenges for Imaging Air Cherenkov Telescopes .	15
What HPC really means in practice and how to love it	16
(INVITED) Esame Schede Inaf	16
(INVITED) SPACE	16
(INVITED) Radioastronomy: the path toward next generation computing $\ldots \ldots \ldots$	17
(INVITED) Documento Commissione Calcolo	17
(INVITED) Cineca	17
(INVITED) Centro Nazionale HPC, BigData & Quantum Computing	17
(INVITED) USCVIII	17

Population Synthesis and Infrared-Ultraviolet / 1

The MPI+CUDA Gaia AVU–GSR Parallel Solver towards next-generation Exascale Infrastructures

Author: Valentina Cesare¹

¹ Istituto Nazionale di Astrofisica (INAF)

Corresponding Author: valentina.cesare@inaf.it

We ported to the GPU with CUDA the Astrometric Verification Unit-Global Sphere Reconstruction (AVU-GSR) Parallel Solver developed for the ESA Gaia mission, by optimizing a previous OpenACC porting of this application. The code aims to find, with a [10,100] µarcsec precision, the astrometric parameters of ~10^8 stars, the attitude and instrumental settings of the Gaia satellite, and the global parameter y of the parametrized Post-Newtonian formalism, by solving a system of linear equations, $A \times x = b$, with the LSQR iterative algorithm. The coefficient matrix A of the final Gaia dataset is large, with ~10^11 × 10^8 elements, and sparse (Becciani et al., 2014), reaching a size of ~10-100 TB, typical for the Big Data analysis, which requires an efficient parallelization to obtain scientific results in reasonable timescales. The speedup of the CUDA code over the original AVU–GSR solver, parallelized on the CPU with MPI+OpenMP, increases with the system size and the number of resources, reaching a maximum of ~14x, >9x over the OpenACC application (Cesare et al., 2021, 2022c,b; Cesare, et al., submitted). This result is obtained by comparing the two codes on the CINECA cluster Marconi100, with 4 V100 GPUs per node. After verifying the agreement between the solutions of a set of systems with different sizes computed with the CUDA and the OpenMP codes and that the solutions showed the required precision, the CUDA code was put in production on Marconi100, essential for an optimal AVU-GSR pipeline and the successive Gaia Data Releases. This analysis represents a first step to understand the (pre-)exascale behaviour of a class of applications that follow the same structure of this code. In the next months, we plan to run this code on the pre-exascale platform Leonardo of CINECA, with 4 next-generation A200 GPUs per node, towards of a porting on this infrastructure, where we expect to obtain even higher performances.

Cosmological Simulation, Jet Pulsar Wind Nebulae (PWN), Mergers and Explosive Events, Other Simulations / 2

PLUTO

Author: Marco Rossazza¹

¹ Università degli Studi di Torino

Corresponding Author: marco.rossazza@unito.it

The PLUTO code, developed at the University of Torino in collaboration with the Osservatorio Astrofisico di Torino, is one of the most widely used public codes for

astrophysical fluid-dynamics and magnetohydrodynamics, both in the classical and relativistic regimes. The code is designed with a modular and flexible structure whereby different numerical algorithms can be separately combined to solve systems of conservation laws by using a finite volume or a finite difference approach, based on Godunov-type schemes.

We present the work done on the GPU porting of the code by using OpenACC. OpenACC is a programming model that uses high-level compiler directives and parallelizing compilers to exploit GPU technology. We will highlight the many code structure changes required by the new parallel programming paradigm. We finally show the results obtained in terms of the acceleration and efficiency on various systems such as Marconi100 and Leonardo at Cineca.

"A new(ish) tool on the workbench": new version of the population synthesis code SEVN.

Author: Giuliano Iorio¹

¹ Istituto Nazionale di Astrofisica (INAF)

Corresponding Author: giuliano.iorio@unipd.it

Binary population-synthesis codes are crucial tools to study the evolution of massive binaries and their impact for the demography of Black Holes binaries.

I present a new fully revised version of the population synthesis code SEVN (Stellar EVolution for N-body codes, Spera et al. 2019).

The stellar evolution is computed interpolating from look-up tables containing a grid of stellar evolution models for both Hydrogen and naked Helium stars that can be easily updated. The code includes binary evolution processes (wind and Roche Lobe overflow mass transfer, Common Envelope, stellar mergers, tidal evolution, gravitation wave decay), PPISNe and PISNe prescriptions and different models for SNe and remnants creation. The addition of new binary processes or new SNe models is made simple by the modularity of the code.

With respect to the last SEVN version, we introduce a new adaptive time-step method plus a checkand-repeat schema that guarantees the automatic reduction of the time-step when needed. The new implemented time step has a double advantage: it is capable to follow with extraordinary details the crucial phases of the binary evolution (e.g. RLO) and to drastically reduce the overall computational time (about ten times less with respect to the precedent version). The code is fully parallelised and it adopts a smart dynamic memory allocation. All these properties make SEVN perfectly scalable from personal laptops to large clusters.

In conclusion, SEVN represent a perfect balance between computational resources (both time and memory) and versatility due to the easy update of stellar evolution models. It is a precious tool for the investigation of black holes properties and the interpretation of current and forthcoming detections (e.g. Gravitational Waves or astrometric signals).

4

SuperBAM: characterization of Supernova transients through Bayesian Analytic Modeling

Authors: Maria Letizia Pumo¹; Stefano Pio Cosentino¹

¹ Istituto Nazionale di Astrofisica (INAF)

Corresponding Author: stefano.cosentino@dfa.unict.it

The launch of the Legacy Survey of Space and Time (LSST) is going to open the possibility to perform statistical studies on the physical properties of Supernova transients. These studies need sufficiently fast and accurate procedures for the characterization of these events. In this framework, we have developed a Bayesian Analytic Modeling procedure, which is able to simulate Supernovae Light Curves and infer the main physical parameters describing their progenitors at the explosion. This new procedure as well as its validation will be presented.

High Energy / 5

Unveiling the Complexity through High Performance Computing: the Link between Massive Stars, Core-Collapse Supernovae, and their Remnants

Author: Salvatore Orlando¹

Co-authors: Fabrizio Bocchino¹; Marco Miceli¹

¹ Istituto Nazionale di Astrofisica (INAF)

Corresponding Author: orlando@astropa.inaf.it

Core-collapse supernova remnants (SNRs) exhibit intricate morphologies and a highly non-uniform distribution of stellar debris. In the case of young remnants (less than 5000 years old), their characteristics offer insights into the inner processes of the supernova (SN) engine, including nucleosynthetic yields and large-scale asymmetries originating from the early stages of the explosion. Additional features stem from the progenitor star's internal structure at collapse and the interactions between the remnant and the circumstellar medium (CSM), which is shaped by the mass-loss history of the progenitor.

Hence, investigating the connection between young SNRs, parent SNe, and progenitor massive stars is of paramount importance. Firstly, it allows us to delve into the physics of SN engines by shedding light on the asymmetries that occurred during the explosion. Secondly, it provides an avenue to examine the final stages of massive star evolution and the elusive mechanisms that govern their mass loss.

Presently, our ability to study the progenitor-SN-SNR connection has greatly improved due to advanced 3D MHD models and the availability of adequate high performance computing resources. Now we have the ability to elucidate the long-term evolution from the progenitor star to the SN and subsequently to the SNR. Coupled with high-quality observational data spanning the electromagnetic spectrum, we can effectively constrain these models.

In this talk, I will offer a brief overview of recent advancements in modeling the progenitor-SN-SNR connection. The primary focus will be on investigations aimed at establishing connections between the observed physical and chemical properties of SNRs and their progenitor stars and SN explosions. By doing so, we gain valuable insights into the life and death of massive stars.

Cosmological Simulation, Jet Pulsar Wind Nebulae (PWN), Mergers and Explosive Events, Other Simulations / 6

Numerical Astrophysics at INAF-OAPa: an overview and future perspectives

Author: Fabrizio Bocchino¹

¹ Istituto Nazionale di Astrofisica (INAF)

Corresponding Author: fabrizio.bocchino@inaf.it

In this talk, I will outline the activities in the field of numerical astrophysics at INAF-OAPa and the role of the local Sistema Computazionale per l'Astrofisica Numerica (SCAN). The availability of local computational resources fostered the realization of serveral astrophysics projects, representing a springboard for large computational programs which used massive international HPC systems. I will present a few remarks on the future of such clusters in the PNRR and Exa-scale computing era, from a INAF research structure perspective.

High Energy / 7

Computational infrastructure for the Monitoring System of the Cherenkov Telescope Array

Authors: Federico Incardona¹; Alessandro Costa¹; Kevin Munari¹

Co-authors: Pietro Giuseppe Bruno¹; Stefano Germani¹

¹ Istituto Nazionale di Astrofisica (INAF)

Corresponding Author: federico.incardona@inaf.it

We hereby present the computational requirements of the infrastructure needed by the Monitoring System (MON) of the Cherenkov Telescope Array (CTA) in two different scenarios: the performance tests and the on-site deployment. The CTA will be composed of hundreds of telescopes working together to attempt to unveil some fundamental physics of the high-energy Universe. Along with the scientific data, a large volume of housekeeping and auxiliary data coming from weather stations, instrumental sensors, logging files, etc., will be collected as well. MON is the subsystem of ACADA that is responsible for monitoring and logging the overall CTA array. It acquires and stores monitoring points and logging information from the array elements, at each of the CTA sites. MON is designed and built in order to deal with big data time series and exploits some of the currently most advanced technologies in the field of the Internet of Things (IoT). The complex software architecture of MON would require large resources in terms of I/O throughput and the number of CPU cores, which will probably necessitate the full-exclusive use of several bare metal nodes.

8

Computational infrastructure for the Monitoring System of the Cherenkov Telescope Array

Authors: Kevin Munari¹; Alessandro Costa¹; Federico Incardona¹

Co-authors: Pietro Giuseppe Bruno¹; Stefano Germani¹

¹ Istituto Nazionale di Astrofisica (INAF)

Corresponding Author: kevin.munari@inaf.it

We hereby present the computational requirements of the infrastructure needed by the Monitoring System (MON) of the Cherenkov Telescope Array (CTA) in two different scenarios: the performance tests and the on-site deployment. The CTA will be composed of hundreds of telescopes working together to attempt to unveil some fundamental physics of the high-energy Universe. Along with the scientific data, a large volume of housekeeping and auxiliary data coming from weather stations, instrumental sensors, logging files, etc., will be collected as well. MON is the subsystem of ACADA that is responsible for monitoring and logging the overall CTA array. It acquires and stores monitoring points and logging information from the array elements, at each of the CTA sites. MON is designed and built in order to deal with big data time series and exploits some of the currently most advanced technologies in the field of the Internet of Things (IoT). The complex software architecture of MON would require large resources in terms of I/O throughput and the number of CPU cores, which will probably necessitate the full-exclusive use of several bare metal nodes.

High Energy / 9

Modeling young pulsar wind nebulae

Author: Barbara Olmi¹

¹ Istituto Nazionale di Astrofisica (INAF)

Corresponding Author: barbara.olmi@inaf.it

Pulsar Wind Nebulae are powered by the relativistic, magnetized and cold wind emanating from a rapidly rotating neutron star (the pulsar) that interacts with the ambient medium. They are visible as bright non-thermal sources at a very broad range of energies, from radio to gamma-rays, with a variety of different morphologies.

Pulsar Wind Nebulae are perfect places where to look at for extreme processes, and their relevance

goes beyond high-energy astrophysics: relativistic plasmas, magnetic dissipation, particle acceleration in unfavorable environments, massive injection of particles in the ambient and contribution to the cosmic rays spectrum.

Over the past two decades, the main tools for studying these sources have been relativistic magnetohydrodynamic numerical simulations, which can reproduce many of the properties observed in these fascinating systems, down to very fine details, shedding light on the nature of the pulsar outflow.

In this talk I will discuss the actual state of the art numerical models for pulsar wind nebulae, with a focus on newborn and young systems.

High Energy / 10

Understanding the magnetic field evolution in supernova remnants: a crucial role of high-performace computing

Author: Oleh Petruk¹

¹ Istituto Nazionale di Astrofisica (INAF)

Corresponding Author: oleh.petruk@inaf.it

Magnetic fields manifest themselves almost everywhere in the Universe. Their effects are visible through different kinds of electromagnetic radiation and in the spectra of cosmic rays. In our talk, we focus on the magnetic field in supernova remnants. We show how high-performance computing allows us to investigate the evolution of magnetic field in the remnants of different types of supernova and to uncover development of its three-dimensional spatial structures. We demonstrate on example of the remnant of supernova SN1987A that massive three-dimensional MHD simulations coupled with the radio polarization observations result in strong limitations on models of the presupernova circumstellar magnetic field in SN1987A and thus of the progenitor star itself.

High Energy / 11

MHD numerical simulations of nanojets and nanoflares in the solar corona for MUSE

Author: Paolo Pagano¹

Co-authors: Fabio Reale ¹; Gabriele Cozzo ¹; Antonino Petralia ¹; Costanza Argiroffi ²; Patrick Antolin ³

¹ Istituto Nazionale di Astrofisica (INAF)

² Università degli Studi di Palermo

³ Northumbria University

Corresponding Author: paolo.pagano@unipa.it

The solar corona shows inexplicably high temperatures, up to million degrees, when compared with the cold lower photosphere. The conversion of magnetic energy into thermal one through the magnetic reconnection has been chased for the last decades as the mechanism to explain this phenomenon. The reason why such mechanism remains elusive is because reconnection events are singularly too small and fast to be detected (nanoflares), whereas their collective action is sufficient to sustain the million degrees corona against thermal conduction and radiative losses.

The forthcoming MUSE mission of NASA, a new high cadence high resolution EUV spectrometer will be launched in 2027 with the aim of unveiling such elusive phenomena.

The strategy of the mission is to couple high cadence high resolution observations of coronal loops (arch like magnetic structures where the plasma is confined) with state-of-the-art numerical simulations which can synthesise MUSE observables and disentangle observations otherwise too dynamic and complicated to be understood with traditional inversion methods.

We perform magnetohydrodynamics (MHD) simulations of the dynamic counter part of nanoflares, i.e. the nanojets a byproduct of the magnetic reconnection.

We analyse the relationship between the nanoflare and the nanojet, explaining how the latter, when observed, could give away the occurrence of the former.

Our 3D MHD simulations are key to bridge the gap between idealised magnetic reconnection models and future MUSE observations.

High Energy / 12

HPC MHD modelling of unstable reconnecting plasma in the solar corona and EUV diagnostics with the MUSE mission

Author: Gabriele Cozzo¹

Co-authors: Alan W. Hood²; Fabio Reale¹; Jack Reid²; Paolo Pagano¹

¹ Università degli Studi di Palermo

² St. Andrews University

Corresponding Author: gabriele.cozzo@unipa.it

The solar corona consists of plasma confined by, and interacting with, the coronal magnetic field. The magnetic processes are highly dynamic and non linear, and their description requires time-dependent magnetohydrodynamic modelling on high performance computing systems.

Large-scale energy release in the corona may involve MHD instabilities such as the kink instability in a single twisted magnetic flux tube. The turbulent dissipation of the magnetic structure into small-scale current sheets converts into a sequence of a-periodic, impulsive, heating events. Since single twisted magnetic filaments are generally embedded in a multi-threaded structure, an unstable strand could trigger a global MHD instability, and a nanoflare cascade.

Using full 3D MHD simulations with the PLUTO code we show that avalanches are a viable mechanism for the storing and release of magnetic energy in the solar corona, as a result of photospheric motions. We provide a synthetic perspective for the NASA MUSE forthcoming mission. Many fine scale features as well as rapid changes in emissivity and Doppler shifts might be accessible to the MUSE spectrometer and shed new light on coronal heating mechanisms.

High Energy / 13

Bow-Shock Pulsar Wind Nebulae: a tale of trails

Author: Niccolo' Bucciantini¹

Co-author: Barbara Olmi¹

¹ Istituto Nazionale di Astrofisica (INAF)

Corresponding Author: niccolo.bucciantini@inaf.it

Pulsar Wind Nebulae (PWNe) constitute a magnificent lab to investigate high-energy astrophysics in its many facets, from non-thermal emission, to particle acceleration, from relativistic fluid dynamics to anti-matter creation. The group in Arcetri has always been one of the leading team in the study of PWNe, and through the years has developed a vast and advanced suits of numerical tools for their study.

I will present the more recent results, based on state of the art 3D relativistic MHD simulations of the so called Bow Shock Phase of PWNe, aimed at providing a full sampling of the vast parameter space characterizing these objects, in terms of spin-axis inclination, ISM magnetization, pulsar wind

energy distribution. I will describe how this has helped us not just to improve our knowledge of the dynamics of these objects, but also to understand the formation of misaligned X-ray trails, and possibly the even more mysterious TeV Haloes, that have catch much attention in the high energy astrophysical community.

High Energy / 14

Relativistic Thermodynamics in CFD codes.

Author: Vincenzo Antonuccio¹

¹ Istituto Nazionale di Astrofisica (INAF)

Corresponding Author: vincenzo.antonuccio@inaf.it

In numerical experiments of the propagation of relativistic jets produced e.g. by Supermassive Black Holes (SMBH) the covariant equation of state determines a relationship between density, pressure and temperature. While the former is proportional to the Lorentz factor and the second is invariant, one is left with different possibilities concerning temperature. The usual choice adopted in Computational Fluid Dynamics (hereafter \emph{CFD}) is: $T \equiv T (U^{\mu}p_{\mu})$, i.e. that also temperature is Lorentz invariant. Current covariant formulation of thermodynamics do not give unique predictions about the transformation of temperature among comoving reference systems. We have investigated the implications of these possibilities on the observed properties of highly relativistic components of AGNS, i.e. the infall region near the SMBH horizon and the relativistic jets. We suggest that these different environments could be testbeds of different RT models.

Cosmological Simulation, Jet Pulsar Wind Nebulae (PWN), Mergers and Explosive Events, Other Simulations / 15

Challenges in forthcoming CMB data sets

Author: Giuseppe Puglisi¹

¹ Universita' di Roma Tor Vergata

Corresponding Author: giuseppe.puglisi@uniroma2.it

One of the major challenges in the context of the Cosmic Microwave Background (CMB) radiation is to detect a polarization pattern, the so called B-modes of CMB polarization, that are thought to be directly linked to the quantum tensor fluctuations produced in the Universe during the inflationary phase. To date, several challenges have prevented to detect the B-modes partly because of the lower sensitivity of the detectors and because of the contamination of our own Galaxy acting as a foreground contamination in polarization at large scales. At smaller angular scales the data are instead contaminated by the emission of extra-galactic sources like radio quasars and dusty star forming galaxies. In this talk, I will list the computational challenges when it comes to analyze, simulate and reduce realistic amount of data from the forthcoming CMB polarization observations where a huge amount of receivers are going to be employed. I will, particularly, show how improvements thanks to novel machine learning techniques help in correctly accounting for the foreground contamination; particularly, in the context of the future CMB experiments (e.g. SO, LiteBIRD, CMB-S4), where high sensitivities will be achieved at both high (~1 arcmin) and low (~1 deg) resolutions.

The Arxes code suite: high-performance planet formation and astrochemistry

Author: Diego Turrini¹

Co-authors: Elenia Pacetti ¹; Eugenio Schisano ¹; Sergio Molinari ¹; Sergio Fonte ¹; Danae Polychroni ¹; Romolo Politi ¹

¹ Istituto Nazionale di Astrofisica (INAF)

Corresponding Author: diego.turrini@inaf.it

Thanks to the continuous expansion and improvement of the observational datasets on circumstellar disks, exoplanets and our own Solar System, planet formation and astrochemistry are becoming increasingly interconnected fields of study. The growth and migration processes that shape planet formation in the native circumstellar disks create a bidirectional link between planetary and disk composition. To explore the nature of this link, in the framework of the INAF program Arxes we developed a new suite of simulation and processing codes. The *Arxes* code suite includes:

- *Mercury-Arxes*, a parallel n-body code that simulates the growth and migration of forming planets within circumstellar disks and the interactions between planetary bodies and the disk gas;
- *JADE*, a paralled multi-language code that jointly reproduces the physical and chemical evolution of circumstellar disks setting the initial conditions of planet formation;
- *Hephaestus*, a compositional code that combines the outputs of Mercury-Arxes and JADE to quantify the composition of newly-formed planets from their accretion of gas and solids in the disk;
- *Debris*, a parallel collisional code that estimates from the output of Mercury-Arxes the collisional production of dust from the impacts of planetary bodies and its distribution within the disk.

The simulation and processing capabilities of the Arxes suite of codes are currently supporting the ESA mission Ariel, the ERC Synergy project ECOGAL and the NASA mission Juno, and are being expanded in the framework of the PNRR activities for the National Center for High-Performance Computing, with plans for future applications also in the framework of SKA.

Radio Astronomy / 17

Development of radio source detection and classification tools for SKA precursors

Author: Simone Riggi¹

¹ Istituto Nazionale di Astrofisica (INAF)

Corresponding Author: simone.riggi@inaf.it

The SKA precursor communities are currently developing new software to automate the processing of radio images for various tasks, including source extraction, object or morphology classification, and anomaly detection. These developments heavily rely on HPC processing paradigms and machine learning (ML) methodologies.

In this context, we are developing several tools to support the scientific analysis conducted within SKA precursor surveys. One of them, dubbed *caesar*, is a source finder for large radio-continuum maps, capable of distributing processing across multiple computing nodes in an HPC infrastructure. Another tool, *caesar-mrcnn*, employs trained deep neural networks to detect imaging artefacts and radio sources of different morphologies. *sclassifier* offers various ML methods for different applications, such as source classification from multi-wavelength images, or unsupervised/self-supervised learning of radio data representations. Furthermore, we are developing tools that use generative models to produce synthetic radio image data for data challenge or model performance

boosting scopes. These tools have been trained and tested on ASKAP EMU and MeerKAT GPS survey data.

An overview of the achieved results will be presented at the workshop, along with details on the employed computing resources, ongoing activities, and future prospects.

Cosmological Simulation, Jet Pulsar Wind Nebulae (PWN), Mergers and Explosive Events, Other Simulations / 18

(INVITED) The Universe in a Box

Author: Stefano Borgani¹

¹ Istituto Nazionale di Astrofisica (INAF)

Corresponding Author: borgani@oats.inaf.it

In my talk I will overview the current landscape of large simulations for the formation of cosmic structures. After overviewing the role that modern simulations have in the field of HPC, I will highlight their growing relevance as a theoretical framework for the interpretation of observational data, and as a support to ongoing and future observational facilities in astrophysics and cosmology. In this context I will discuss the points of strengths and the criticalities of the Italian community working in this field, and the perspectives offered by the new generation of computing facilities.

Cosmological Simulation, Jet Pulsar Wind Nebulae (PWN), Mergers and Explosive Events, Other Simulations / 19

High performance visualization for Astronomy & Cosmology

Author: Eva Sciacca¹

Co-authors: Fabio Roberto Vitello ; Nicola Tuccari¹

¹ Istituto Nazionale di Astrofisica (INAF)

Corresponding Author: eva.sciacca@inaf.it

Modern astronomy and astrophysics produce massively large data volumes (in the order of petabytes) coming from observations or simulation codes executed on high performance supercomputers. Such data volumes pose significant challenges for storage, access and data analysis. Visual exploration of big datasets poses some critical challenges that must drive the development of a new generation of graphical software tools, specifically: (i) *Interactivity* to deal with datasets exceeding the local machine's memory capacity, so for complex visualizations the relevant computations should be performed close to the data to avoid time consuming streaming of large data volumes; ii) *Integration* to be ideally fully integrated within the scientists'toolkit for seamless usage, abstracting from technical details related to the underlying high performance computing (HPC) resources, freeing scientists to concentrate in doing science; iii) *Collaboration* to facilitate visualization, processing and analysis of big data in a collaborative manner within e.g. science gateway technologies to allow collaborative activity between users and provide customization and scalability of data analysis/processing workflows, hiding underlying technicalities.

INAF-OACT has been developing and maintaining the Visualization Interface for the Virtual Observatory (VisIVO) and recently extended it with the ViaLactea Visual Analytic modules. VisIVO is developed adopting the Virtual Observatory standards and its main objective is to perform 3D and multi-dimensional data analysis and knowledge discovery of a-priori unknown relationships between multi-variate and complex astrophysical datasets. VisIVO has been already deployed using Science Gateways to access DCIs (including clusters, grids and clouds) using containerization and virtualization technologies, it has also been selected as one of the pilot applications deployed on EOSCpilot infrastructure demonstrating that the tools can be accessed using gateways and cloud platforms and it has been deployed on the European Open Science Cloud (EOSC), efficiently exploiting Cloud infrastructures and interactive notebooks applications.

During this period, and thanks to the collaboration within the SPACE EU Centre of Excellence, the H2020 EUPEX Project and the ICSC National Research Centre for High Performance Computing, Big Data and Quantum Computing, we are planning to adapt VisIVO solutions for high performance visualisation, remote, in-situ, in-transit visualisation of data generated on the (pre-)Exascale systems by HPC applications in Astrophysics and Cosmology (A&C) including GADGET (GAlaxies with Dark matter and Gas intEracT) simulated data and astrophysical fluid dynamics (PLUTO) simulations.

In this talk I will present the evolution direction and the related implementation activities tailored to pursue the following objectives: 1) Enhance the *portability* of the VisIVO modular applications and their resource requirements. 2) Foster *reproducibility* and maintainability; 3) Take advantage of a more flexible resource exploitation over *heterogeneous HPC facilities* (including also mixed HPC-Cloud resources); 4) Minimize data-movement overheads and *improve I/O performances*.

Acknowledgements

This work is funded by the European High Performance Computing Joint Undertaking (JU) and Belgium, Czech Republic, France, Germany, Greece, Italy, Norway, and Spain under grant agreement No 101093441 and it is supported by the spoke "FutureHPC & BigData" of the ICSC –Centro Nazionale di Ricerca in High Performance Computing, Big Data and Quantum Computing –and hosting entity, funded by European Union –NextGenerationEU"

Cosmological Simulation, Jet Pulsar Wind Nebulae (PWN), Mergers and Explosive Events, Other Simulations / 20

"From Vlasov-Poisson to Schr ödinger-Poisson: dark matter simulation with a quantum variational time evolution algorithm"

Author: Luca Cappelli¹

Co-authors: Francesco Tacchino²; Giuseppe Murante³; Ivano Tavernelli²; Stefano Borgani³

¹ INAF

² IBM Reaserch Zurich

³ Istituto Nazionale di Astrofisica (INAF)

Corresponding Author: luca.cappelli@inaf.it

"Recent studies showed an interesting mapping of the 6-dimensional+1 (6D + 1) collisionless fluid (Vlasov-Poisson)

problem into a more amenable 3D + 1 non-linear Schr odinger-Poisson (SP) problem for simulating the evolution of DM perturbations. This opens up the possibility of improving the scaling of time propagation simulations using quantum computing. We propose a rigorous formulation of a variational-time evolution quantum algorithm for the simulation of the SP equations to follow DM perturbations and investigate the transition of the SP dynamics towards the classical (h/m \rightarrow 0) limit."

Population Synthesis and Infrared-Ultraviolet / 21

(INVITED) The Gaia legacy. From data reduction and analysis to data management and exploitation: HPC, HTC and Big Data issues.

Author: Mario G. Lattanzi¹

¹ INAF-OATo

Corresponding Author: mario.lattanzi@inaf.it

As the Gaia mission approaches the end of its scientific life (June 2025) after 11 years of continuous in-orbit operations, this presentation addresses the work that Team INAF has started at our Data Processing facility (DPCT@ALTEC) to face the challenges of the mission legacy, as advocated and prototyped in the framework of The Living Sky Project. We focus on the computational aspects implied by the deep exploitation of the first Big Data system of the faint astronomical sky built from space-borne data.

Population Synthesis and Infrared-Ultraviolet / 22

GalaPy, the highly optimised C++/Python spectral modelling tool for galaxies

Author: Tommaso Ronconi¹

Co-authors: Alessandro Bressan²; Andrea Lapi²; Martina Torsello²

¹ Scuola Internazionale Superiore di Studi Avanzati (SISSA), Istituto Nazionale di Astrofisica (INAF)

² SISSA

Corresponding Author: tronconi@sissa.it

Fostered by upcoming data from new generation observational campaigns, we are about to enter a new era for the study of how galaxies form and evolve.

The unprecedented quantity of data that will be collected, from distances only marginally grasped up to now, will require analysis tools designed to target the specific physical peculiarities of the observed sources and handle extremely large datasets.

One powerful method to investigate the complex astrophysical processes that govern the properties of galaxies is to model their observed spectral energy distribution (SED) at different stages of evolution and times throughout the history of the Universe.

To address these challenges, we have developed GalaPy, a new library for modelling and fitting galactic SEDs from the X-ray to the radio band, as well as the evolution of their components and dust attenuation/reradiation. On the physical side, GalaPy incorporates both empirical and physicallymotivated star formation histories, state-of-the-art single stellar population synthesis libraries, a twocomponent dust model for extinction, an age-dependent energy conservation algorithm to compute dust reradiation, and additional sources of stellar continuum such as synchrotron, nebular/free-free emission and X-ray radiation from low and high mass binary stars.

On the computational side, GalaPy implements a hybrid approach that combines the high performance of compiled C++ with the user-friendly flexibility of Python, and exploits an object-oriented design via advanced programming techniques.

GalaPy is the fastest SED generation tool of its kind, with a peak performance of almost 1000 SEDs per second.

The models are generated on the fly without relying on templates, thus minimising memory consumption.

It exploits fully Bayesian parameter space sampling, which allows for the inference of parameter posteriors and thus facilitates the study of the correlations between the free parameters and the other physical quantities that can be derived from modelling.

The API and functions of GalaPy are under continuous development, with planned extensions in the near future.

In this talk, I will introduce the project and showcase the photometric SED fitting tools already available to users.

(INVITED) Present and Future of Data E-Infrastructure to support computing in INAF

Author: Andrea Bignamini¹

Co-author: Cristina Knapic¹

¹ Istituto Nazionale di Astrofisica (INAF)

Corresponding Author: andrea.bignamini@inaf.it

Scientific computing in astrophysics must certainly face challenges related to the implementation and optimization of algorithms, data analysis pipelines, and the management of computing resources. On the other hand using huge computational resources may imply enormous needs of storage space, temporarily saved on scratch storage. To reach optimal performance, scratch storage is not intended to be highly available and persistent, so a slower but much more robust storage area is needed to preserve data in long term repositories or even in structured archives. This is no less important issue we have to deal with when talking about critical computing, mainly in the context of FAIR Principles and Open Data. As a matter of fact, appropriate data management plans must be taken into account as well. Here, we want to show the current hardware resources, the architecture and all the services related to Big Data challenges available to the community and the future perspectives for supporting critical (and less critical) computational topics in the INAF.

Radio Astronomy / 24

Broad-band radio observations in the SKA era: the case of the galaxy group Nest200047

Author: Marisa Brienza¹

¹ Istituto Nazionale di Astrofisica (INAF)

Corresponding Author: marisa.brienza@inaf.it

In recent years, the remarkable capabilities offered by precursors and pathfinders of the Square Kilometre Array (SKA) have started revolutionizing our view even of previously well-known objects such as jetted Active Galactic Nuclei (AGN). Particularly in the MHz-frequency regime, observations are now able to uncover the oldest plasma injected by AGN jets into their surrounding environment, shedding light on the jet duty-cycle and their interaction with the external medium over very long timescales. The crucial advancement towards comprehending the physics of these sources lies now in the exploitation of broad-band spectro-polarimetric radio observations. Although the combination of multiple telescopes at different frequencies currently enables this approach, it remains computationally intensive and time-consuming.

In this context, I will present our recent study on the galaxy group Nest200047, which serves as an exemplary instance of recurring AGN jet activity in a low-mass system. Our investigation reveals unprecedented evidence of the evolution of the old AGN plasma into intricate filamentary structures over hundreds of Myr. By conducting a dedicated multi-frequency campaign spanning from 53 to 1500 MHz, using LOFAR, uGMRT, MeerKAT, and VLA telescopes, we have performed a unique, resolved spectro-polarimetric analysis of the system. I will discuss the computational and storage resources required to achieve these results, as well as the used algorithms and the overall processing time. This type of analysis is crucial in preparing for the full utilization of the SKA's potential once it becomes operational.

Cosmological Simulation, Jet Pulsar Wind Nebulae (PWN), Mergers and Explosive Events, Other Simulations / 25

The cosmic dance

Author: Alice Damiano¹

¹ Istituto Nazionale di Astrofisica (INAF)

Corresponding Author: alice.damiano@inaf.it

Cosmological hydrodynamic simulations are unique and successful tools for investigating the evolution of galaxies in a cosmological context. The introduction of BHs and their feedback into simulations - mandatory to model the intertwined growth of BHs and host galaxies - runs up against numerical limitations that are circumvented with ad-hoc sub-resolution techniques.

However, the accurate reconstruction of BH dynamics proves to be not only a necessary ingredient to recover the AGN feedback that influences structure growth, but also a powerful tool to fully exploit the new window offered by gravitational wave astrophysics for the study of the formation and evolution of cosmic structures.

To this aim, we have developed a new physically-based method to reproduce the dynamical friction force that binds BHs at the centre of galaxies and drives the early stages of mergers, and we want to present it and discuss the preliminary results.

Radio Astronomy / 26

Energy efficiency in Data Reduction for Imaging in a Radio Astronomy pipeline

Authors: Giovanni Lacopo¹; Giuliano Taffoni²; Luca Tornatore²

¹ Istituto Nazionale di Astrofisica (INAF)

² OATS-INAF

Corresponding Author: giovanni995@icloud.com

The effective exploitation of modern architecture is a key factor to achieve best performances in terms of both energy efficiency and run-time reduction.

We bring a specific example of this, by discussing the W-stacking gridder, an algorithm that tackles Radio imaging in massively parallel systems; its performance is limited by an all-to-all data reduction needed to pass from time-domain decomposition to space-domain decomposition.

To overcome this limitation, we have implemented a customized reduce operation built on explicitly numa-awareness.

We have found inside each computing node an increase in both performance and energy efficiency by a factor of 4 to 7 on different architectures.

High Energy / 27

"The Dark Energy and Massive Neutrino Universe" cosmological simulations

Author: Carmelita Carbone¹

¹ INAF IASF-MI

Corresponding Author: carmelita.carbone@inaf.it

I will introduce the DEMNUni simulation set, which accounts for different cosmologies with massive neutrinos and dynamical dark-energy. I will briefly present the scientific results obtained up to now. Then, I will focus on the computational resources (both ISCRA and MoU CINECA-INAF) needed to produce the DEMNUni suite, as well as the long-term storage required to support and mantain the

scientific exploitation of such a big amount of data, kindly provided by ISCRA, MoU CINECA-INAF, IA2 and CNAF.

Population Synthesis and Infrared-Ultraviolet / 28

The OPS4: towards a legacy Big Data system - A detailed view

Authors: Enrico Licata¹; Deborah Busonero¹; Mario G. Lattanzi²; Roberto Morbidelli¹

Co-authors: Rosario Messineo 3; Leonardo Tolomei 3; Lorenzo Bramante 3

¹ Istituto Nazionale di Astrofisica (INAF)

² INAF-OATo

³ ALTEC S.P.A.

Corresponding Author: enrico.licata@inaf.it

The OPS4 has been designed by a joint INAF-ALTEC team, as a legacy Big Data facility coming from the experience of the TLS prototype and the Gaia data reduction.

This is accomplished by taking into account the scalability and performance requirements necessary for the analysis and exploitation of Big Data dedicated to the investigation of the nearby Universe in the context of multimessenger astronomy.

The presentation concentrates on the definition of a coherent DM suitable to the expected use cases, and illustrates the reasons for choosing hybrid data management, with its metadata structured under Oracle DBMS (storage intensive), with the potential to exploit technologies such as in-memory and relational duality, while unstructured data (cold data) are residing on filesystem, being required for large-scale/deep data analysis.

The system implements a highly reliable and resilient file system, which allows to manage and organize data efficiently, while ensuring data security and integrity.

Radio Astronomy / 29

Reducing MeerKAT data of the Galactic Plane

Author: Adriano Ingallinera¹

¹ Istituto Nazionale di Astrofisica (INAF)

Corresponding Author: adriano.ingallinera@inaf.it

SKA precursors are giving us a first glimpse of the future capabilities of SKA. Designed to be the most sensitive radio telescopes ever, the precursors are planned to release large area surveys with arcsec resolution. However, the final image product is heavily influenced by the data reduction. Not only the huge data quantity makes a careful visual inspection and manual reduction of the data impractical, but using new data reduction techniques is mandatory to correct imaging errors that prevent reaching the theoretical sensitivity. The cost is the extreme computational demand. New data reduction pipelines have been developed to overcome this problem, running on HPC facilities. In this talk we present the science case of the data reduction of the SARAO MeerKAT Galactic plane survey as a particular case of Galactic data reduction. We will focus on all the peculiarities of Galactic data processing and on the challenges they pose. We will finally show some solutions we are adopting and what we expect in future surveys.

Population Synthesis and Infrared-Ultraviolet / 30

(INVITED) Computation in space missions: the Euclid case

Authors: Marco Frailis¹; Daniele Tavagnacco¹; Erik Romelli¹; Samuele Galeotta^{None}

Co-authors: Thomas Gasparetto ¹; Gianmarco Maggio ¹; Davide Maino ²; Federico Rizzo ¹; Thomas Vassallo ³; Andrea Zacchei ¹

¹ Istituto Nazionale di Astrofisica (INAF)

² UniMi

³ INAF

Corresponding Author: marco.frailis@inaf.it

The EUCLID project, which is part of the M-class ESA Cosmic Vision program, aims to map the geometry of the dark universe through two independent probes: weak gravitational lensing and galaxy clustering. It will conduct a survey of the extragalactic sky in the visible and near-infrared bands, resulting in the generation of up to 30 petabytes of data. The architecture of the Euclid Science Ground Segment (SGS) has been designed to distribute computation and data among 9 Science Data Centres (SDCs) and minimize data transfers between them. Key architectural components of the SGS include the Euclid Archive System for metadata inventory, a Distributed Storage System providing a unified view of the data stored in the SDCs, an Infrastructure Abstraction Layer for workflow management, and a Common Orchestration System for a balanced distribution of data and processing. Containerization, in conjunction with the CernVM-FS software distribution service, facilitates software deployment across all SDCs. The processing pipelines within the SGS primarily follow the HTC paradigm, with GPU acceleration being essential for a critical analysis step in production. Resource estimation is conducted based on simulations and periodic end-to-end tests, guiding the procurement strategies of institutes and agencies.

This presentation will provide an overview of the Euclid SGS computing and software infrastructure, highlighting the rationale behind certain design choices and the lessons learned during the software development and integration phases.

Radio Astronomy / 31

Remote Visualization of Big Data: VisIVO as a Visualization Prototype for SKA Regional Centres

Authors: Giuseppe Tudisco¹; Fabio Roberto Vitello^{None}

¹ Istituto Nazionale di Astrofisica (INAF)

Corresponding Author: giuseppe.tudisco@inaf.it

Radio astronomy is evolving toward ever larger and more accurate datasets.

As soon as the SKA telescopes are fully operational, hundreds of petabytes of data will be produced each year with unprecedented resolution and detail.

This rapid evolution drives toward the development of Big Data analysis and visualization tools and services, which will necessarily need to be supported by suitable infrastructure and computational capabilities to sustain this immense flow of data. This is being implemented in the context of SKA through the creation of a distributed global network of so-called SKA Regional Centres (SRC). In this talk, we will discuss the most recent developments of an interactive visualization tool that is a part of the VisIVO suite and is proposed as one of the visualization prototypes for SRCs.

In particular, we will discuss the transition from a local visualizer to a remote visualizer based on client-server architecture that was required as a basis for engaging with such large data. We are going to also discuss the challenges and advantages of remote visualization and/or parallel visualization, as well as the viability of running interactive visualization pipelines on HPC clusters.

(INVITED) Addressing Computing Challenges for Imaging Air Cherenkov Telescopes

Author: Ciro Bigongiari¹

¹ Istituto Nazionale di Astrofisica (INAF)

Corresponding Author: ciro.bigongiari@inaf.it

Imaging Air Cherenkov Telescopes (IACTs) play a crucial role in the field of ultra-high energy astrophysics (E > 10GeV), bringing a fundamental contribution to the study of cosmic gamma-ray sources. As the sensitivity and complexity of IACT systems increase, so does the demand for efficient data reduction techniques and computational resources to handle the ever-growing data volumes. This presentation focuses on the challenges posed by the data processing requirements of two major nextgeneration IACT projects: the Cherenkov Telescope Array (CTA) and the ASTRI miniarray.

Cosmological Simulation, Jet Pulsar Wind Nebulae (PWN), Mergers and Explosive Events, Other Simulations / 33

What HPC really means in practice and how to love it

Author: Luca Tornatore¹

¹ Istituto Nazionale di Astrofisica (INAF)

Corresponding Author: luca.tornatore@inaf.it

Many scientific codes or data-focused algorithms are being quickly rendered obsolete.

The gargantuan size of the problems that we face is the main responsible for that.

On one hand, the data sets coming from both ground- and space-based observations are already much larger than in the recent past and will still grow by at least an order of magnitude. On the other hand, the target cutting-edge numerical simulations of various kinds require a colossal computational effort to cope with the scientific challenges and will, in turn, produce an amount of data as much colossal.

That obsolescence happens for a number of reasons: it may descend from the lack of distributedmemory capacity (there always be a data set that could not fit in your ram ..), so that they are memory-bound, or it may be due to fundamentally now-inadequate basis (in terms of threading management or the algorithmic implementation) so that their run-time is sky-rocketing.

It is then becoming essential to acquire a solid understanding of how to achieve "high performance" - in its multiple meanings - on modern architectures, how to design and develop a natively parallel code, and how to assess its profile.

In this talk I try to convince of that every sceptic, also discussing two case studies.

Population Synthesis and Infrared-Ultraviolet / 34

(INVITED) Esame Schede Inaf

Corresponding Author: busonero@oato.inaf.it

Cosmological Simulation, Jet Pulsar Wind Nebulae (PWN), Mergers and Explosive Events, Other Simulations / 35

(INVITED) SPACE

Corresponding Author: giuliano.taffoni@inaf.it

Radio Astronomy / 36

(INVITED) Radioastronomy: the path toward next generation computing

Corresponding Authors: gheller@ira.inaf.it, botteon@ira.inaf.it

E-Infra / 37

(INVITED) Documento Commissione Calcolo

Corresponding Author: gianfranco.brunetti@inaf.it

E-Infra / 38

(INVITED) Cineca

Corresponding Author: m.guarrasi@cineca.it

E-Infra / 39

(INVITED) Centro Nazionale HPC, BigData & Quantum Computing

Corresponding Author: ugo.becciani@oact.inaf.it

E-Infra / 40

(INVITED) USCVIII

Corresponding Author: possenti@oa-cagliari.inaf.it