Finanziato dall'Unione europea NextGenerationEU
Ministero dell'Università e della Ricerca
Italiadomani PIANO NAZIONALE DI RIPRESA E RESILIENZA
ICSC Centro Nazionale di Ricerca in HPC, Big Data and Quantum Computing

Mercury-Arxes
High-Performance Planet Formation

Diego Turrini (INAF-OATO), Romolo Politi (INAF-IAPS), Danae Polychroni (INAF-OATS)*

Spoke 3 Technical Workshop, Trieste October 9 / 11, 2023

Scientific Rationale

- New observations continuously increase the ***completeness*** of our knowledge of planetary systems, including our own Solar System.
- Observing ***young planets*** is now routinely possible, both before and just after the dispersal of their native protoplanetary disks.
- JWST and incoming space-based and ground-based facilities are projecting us into a new era of ***compositional characterization*** of planets.
- Italy has a leading role in the scientific design of the ***Ariel*** mission by ESA, which starting from 2029 will systematically characterize hundreds of exoplanets.

- To understand the roots of the great diversity of characteristics of the planetary systems we know, we need **planet formation codes** that are:
- ***Multi-physics*** and capable of tracking the multiple components of protoplanetary disks.
- ***HPC-capable*** to tackle the increasing complexity and details of the simulations.

Technical Objectives, Methodologies and Solutions

MERCURY ArXes

Mercury-Arχes (Turrini+2019,2021) is the multi-language n-body code of **Arχes**, the INAF suite of planet formation and astrochemistry codes:

- Builds on the Mercury n-body code (Chambers 1999; FORTRAN77) and the WHFAST library (Rein & Tamayo 2015; C99).
- Integrates planet formation libraries to track the growth, migration and interactions with the disk gas of planetary bodies (Turrini+2019,2021; Fortran2003).
- Parallelized and vectorized (SIMD) with OpenMP, porting designed so that the parallel and serial version produce the same results.

Technical Objective: *porting Mercury-Arxes to GPU computing*

Methodology: *pragma-based approach, Agile development process*

Adopted Solution: *OpenACC standard*

Timescale, Milestones and KPIs

- **Year 1:** computational bottleneck analysis, design analysis of bottlenecks, individual porting from OpenMP to OpenACC, performance analysis.

  Milestone: release of first OpenACC-capable version.

  KPI: application of first release to real scientific use cases with gain on OpenMP version.
- **Year 2:** design analysis, implementation of new solutions in order of increasing redesign needs, performance analysis on real scientific use cases.

  Milestone: release of second OpenACC-capable version.

  KPI: validation of gain w.r.t. to first release and OpenMP version in realistic scientific use cases.
- **Year 3:** design analysis, implementation of new solutions or evaluation of using OpenMP for GPU Computing in place of OpenACC.

  Milestone: release of final OpenACC-capable version.

  KPI: test of gain in the general-purpose use of final release in place of current OpenMP release.

Accomplished Work, Results

- Identification of the two **computational bottlenecks** (newtonian and non-gravitational forces).
- Porting, data movement optimization and test of the **first OpenACC-capable release**.
- **Test on real use case for Ariel mission** based on simulation setup from Turrini+2021 (2 massive bodies, 10^5 massless bodies, collisional system, newtonian + non-gravitational forces) using Tesla T4 on Genesis cluster at INAF-IAPS.

  Results: performance limited to 6-core equivalent.

  Outcome analysis: data movement accounts for more than 2/3 of the time within OpenACC.
- **Test on real use case from GAPS project** (1005 massive bodies, collisional system, newtonian forces) also using Tesla T4 on Genesis.

  Results: 3x initial speedup gain over 16-core OpenMP version, average performance equivalent to 21-core system (massive bodies collisionally decrease by a factor of 2 in the first ¼ of simulation).
- Identification of **alternative OpenACC design solutions** to be implemented.
- **Publication** of paper on V1298 Tau system's origins using OpenACC simulations (Turrini+2023).
- NVCC proves underperforming on serial and OpenMP codes w.r.t. Intel and GNU compilers.

Next Steps and Expected Results (by next checkpoint: April 2024)

- Merging of computation of newtonian and non-gravitational forces into single OpenACC suboutine to increase computational density per data transfer (goal for April 2024).
- Explore optimization/redesign of computation to maximise vectorization opportunities (optional goal for April 2024).
- Hiring of two researchers, one devoted to WP 1-2, to support team activities (goal for April 2024).
- Test performance on real scientific use cases (current baseline is Ariel and GAPS simulations for comparison with first release).
- Expand physical modelling in Mercury-Arχes with integration of GroMiT code (Polychroni+, in prep.).
- Explore impact of NVCC low performance on non-OpenACC code and possible mitigation strategies.