

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

Thursday, 15 June 2023 10:10 (15 minutes)

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 $[10,100]$ μ arcsec precision, the astrometric parameters of $\sim 10^8$ stars, the attitude and instrumental settings of the Gaia satellite, and the global parameter γ 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 $\sim 10^{11} \times 10^8$ 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 $\sim 14x$, $>9x$ 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.

Primary author: Dr CESARE, Valentina (Istituto Nazionale di Astrofisica (INAF))

Presenter: Dr CESARE, Valentina (Istituto Nazionale di Astrofisica (INAF))

Session Classification: Population Synthesis and Infrared-Ultraviolet

Track Classification: Astrophysics from Infrared to Ultraviolet