

MPI-multicore: Applications

Einstein Tk - Roberto De Pietri (INFN-Parma)

The Einstein Toolkit is an open software that provide the core computational tools needed by relativistic astrophysics, i.e., to solve the Einstein's equations coupled to matter and magnetic fields. In practice, the toolkit solves time-dependent partial differential equations on mesh refined three-dimensional grids. The code has been parallelized using MPI/OpenMP and is actually in production with simulation involving up to 256 cores on the PISA site.

- Refs: Einstein Toolkit

GAIA Mission - Ugo Becciani (INAF)

The parallel application is for the development and test of the core part of the AVU-GSR (Astrometric Verification Unit - Global Sphere Reconstruction) software developed for the ESA Gaia Mission. The main goal of this mission is the production of a microarcsecond-level 5 parameters astrometric catalog - i.e. including positions, parallaxes and the two components of the proper motions - of about 1 billion stars of our Galaxy, by means of high-precision astrometric measurements conducted by a satellite sweeping continuously the celestial sphere during its 5-years mission.

The memory request to solve the AVU-GSR module depends on the number of stars, the number of observations and the number of computing nodes available in the system. During the mission, the code will be used in a range of 300,000 to 50 million stars at most. The estimated memory requirements are between 5 GB up to 8 TByte of RAM. The parallel code uses MPI and openMP (where available) is characterized by an extremely low communication level between the processes, so that preliminary speed-up tests show a behavior close to the theoretical speed-up.

Since AVU-GSR is very demanding on hardware resources, the typical execution environment is provided by Supercomputers, but the resources provided by IGI are very attractive for debugging purpose and to explore the simulation behaviour for a limited number of stars.

NEMO - Massimiliano Drudi (INGV)

NEMO is an ocean modelling framework which is composed of "engines" nested in an "environment". The "engines" provide numerical solutions of ocean, sea-ice, tracers and biochemistry equations and their related physics. The "environment" consists of the pre- and post-processing tools, the interface to the other components of the Earth System, the user interface, the computer dependent functions and the documentation of the system. The NEMO 3.4 oofs2 multiscale simulation package has been ported in the Grid environment in order to run parallel calculations. The NEMO code has significant CPU and memory demand (from our calculations we estimated 1GB/core for a 8 cores simulation). The NEMO code can be used for production as well as for testing purposes by modifying the model (this means recompiling the source code) or varying the input parameters. The user community was interested in exploiting the Grid for the second use case (model testing and tuning) which implies that the package must be executed several times in a parameter sweeping approach. An additional benefit coming from this work is that scientist operating in the field of oceanographic who would be interested in the execution of the code, can share application and results using the Grid data management and sharing facilities.

Namd - Alessandro Venturini (ISOF-CNR-BO)

NAMD is a powerful parallel Molecular Mechanics(MM)/Molecular Dynamics(MD) code particularly suited for the study of large biomolecules. However, it is also compatible with different force fields, making possible the simulation of systems of quite different characteristics. NAMD can be efficiently used on large multi-core platforms and clusters.

Applications < MPI < TWiki

The NAMD use case was a simulation of a 36000 atoms lipid provided by a CNR-ISOF [21] group located in Bologna. To have a real-life use case the simulation had to be run for at least 25 nanoseconds of simulated time resulting on a wallclock time of about 40 days if run on a 8 cores machine.

To port NAMD to the Grid environment, the whole application was rebuilt on Scientific Linux 5 with OpenMPI libraries linked dynamically. Sites supporting the MPI-Start framework and OpenMPI were selected to run the jobs through JDL requirements.

The porting was challenging for two main reasons:

- a data management strategy was needed because we had to make available the output files to the ISOF researchers and the size of the output could not be easily handled via the WMS. This was obtained through pre-run and post-run scripts, both enabled via MPI-Start.
- the length of the simulation implied many computation checkpoints given the time limits on the batch system queues of the sites matching the requirements. We decided to split the simulation in 50 steps each 500 ps of simulated time long, allowing to complete each step without reaching the queues time limits.

Gaussian - Stefano Ottani (ISOF-CNR-BO)

provides state-of-the-art capabilities for electronic structure modeling can run on single CPU systems and in parallel on shared-memory multiprocessor systems Starting from the fundamental laws of quantum mechanics, Gaussian 09 predicts the energies, molecular structures, vibrational frequencies and molecular properties of molecules and reactions in a wide variety of chemical environments. Uses a custom system to parallelize the code or linda.

The use case foresees an umbrella sampling calculation with many short parallel simulations whose output is statistically analysed.

Globo - (ISAC-CNR-BO)

RegCM - Stefano Cozzini (SISSA)

RegCM is the first limited area model developed for long term regional climate simulation currently being developed at the Abdus Salam International Centre for Theoretical Physics (ICTP), Trieste, Italy. RegCM4 is a regional climate model based on the concept of one-way nesting, in which large scale meteorological fields from a Global Circulation Model (GCM) run provides initial and time-dependent meteorological boundary conditions for high resolution simulations on a specific region. The RegCM4 computational engine is CPU intensive using MPI parallel software.

Standard climate RegCM simulations require large dataset (ranging from a few gigabyte for small region up to hundreds of Gigabytes for the largest ones) to be downloaded on Grid and transferred back and forth several times during the model execution. There are however other kinds of computational experiments that can be conducted in a Grid environment: validation runs. This experiment requires to run many different short simulations with different initial conditions. This mixed HTC/HPC approach could be efficiently done on multiple SMP resources made available by Grid resources.

We therefore provide the possibility to run RegCM (or any other MPI parallel application actually) through a relocatable package approach. With this approach all the software needed, starting from an essential OpenMPI distribution is moved to the CEs by the job. All the libraries needed by the program have to be precompiled elsewhere and packaged for easy deployability on any architecture the job will land on. The main advantage of this solution is that it will run on almost every machine available on the Grid and the user will not even need to know what the GRID will have assigned to him. The code itself will need to be compiled with the same relocatable libraries and shipped to the CE by the job. This alternative approach allows a user to run a small RegCM simulation on any kind of SMP resource available to her, quite widely available

nowadays. The main drawback of this solution is that a precompiled MPI distribution will not take advantage of any high speed network available and will not be generally able to use more than one computing node.

Quantum Espresso - Stefano Cozzini (SISSA)

QUANTUM ESPRESSO (Q/E) is an integrated suite of computer codes for electronic-structure calculations and materials modeling, based on density-functional theory. The acronym ESPRESSO stands for opEn Source Package for Research in Electronic Structure, Simulation, and Optimization [23] The suite contains several heterogeneous codes with a wide range of simulation techniques in the area of quantum simulation for material science. Typical CPU and memory requirements for Q/E vary by orders of magnitude depending on the type of system and on the calculated physical property, but in general, both CPU and memory usage quickly increase with the number of atoms simulated. Only tightly-coupled MPI parallelization with memory distribution across processors allows to solve large problems, i.e. systems requiring a large number of atoms. The resulting MPI programs which composes the Q/E suite need fast communications and low latency requires the need to access via Grid HPC cluster resources. Our goal here is to check and evaluate which kind of highly intensive parallel production runs can be done on the top of the IGI MPI infrastructure.

* Refs: Calculation of Phonon Dispersions on the Grid Using Quantum ESPRESSO

Quantum Espresso - Cristian Degli Esposti (CNR-BO)

primo test su Napoli con inserimento di Degli Esposti nella VO unina.it

This topic: MPI > Applications

Topic revision: r24 - 2012-10-26 - DanieleCesini



Copyright © 2008-2023 by the contributing authors. All material on this collaboration platform is the property of the contributing authors.

Ideas, requests, problems regarding TWiki? Send feedback