Mem. S.A.It. Vol. 73, 23

Memorie della

HPC in Astronomy: overview and perspectives

G. Peres[1],

Dip. di Scienze Fisiche ed Astronomiche - sez. di Astronomia, Universitá di Palermo - Piazza Parlamento 1; 90134 Palermo e-mail: peres@astropa.unipa.it

Abstract. I present a personal overview of italian high performance computing in astronomy, as given with the presentation of this meeting, then suggest ways for a more fruitful interaction with the rest of Astronomy; finally and more important, I suggest a qualitative improvement of high performance computing in astronomy through the support of challenging and competitive key projects.

Key words. High performance computing - key projects

1. Introduction

I have been asked to make final comments on the ideas presented in this congress, on Astrophysical High Performance Computing (HPC) in Italy, and on possible perspectives; in some cases I will discuss, more in general, HPC. The task is not easy and many of my views may not be shared by all.

An overview of HPC in italian astronomy is rather difficult because many different fields of physics are addressed, sometimes even together within the same project; drawing perspectives is even more difficult because it is a field rapidly evolving under the effect of technological and scientific developments, of industrial choices and sometimes even governmental decisions.

Also, this meeting has presented us with many new other aspects of computing, like GRID computing. These are rapidly expanding fields and deserve lots of attention; on the other hand they may or may not be related to HPC. In any case, given the rapid pace of their evolution, I prefer not to discuss them otherwise my writing would be obsolete by the time it gets published.

Astrophysical HPC in Italy certainly is a rather active field: the range of areas of physics addressed include plasma physics, fluid dynamics, N-body simulations, relativistic physics, radiative processes and much more, with considerable ingenuity in tackling numerical techniques, computer science and applied mathematics. The relevant areas of Astronomy range from solar physics to stellar structure and evolution, to the formation, dynamics, interaction and merging of galaxies, to cosmology.

Each problem is typically very complex: often several physical effects are at work simultaneously, many of them are highly non-linear, and the physical behaviour is thus highly not-obvious; therefore one has to resort to a numerical solution of the problem translating it, along with relevant boundary and initial conditions, into a numerical code and then into a program to be run on a high performance parallel computer. As time progresses, the availability of more and more powerful computers allows scientists to address more and more challenging problems; but, similarly, it is typicall of all science to tackle the most challenging problem accessible with the tools (including conceptual ones) available.

Running the code typically produces a huge amount of numbers up to the range of terabytes. The task of interpreting results of present day astronomical HPC requires "ancillary" but fundamental activities: visualization and presentation of results, analysis and formatting of results in a way amenable of comparison with observations. The last task may include spectral synthesis from a plasma simulation or the "imaging" of a field of simulated galaxies; in any case this phase of the HPC is mandatory, given the complexity of the codes and of their results and given the need to compare models with observations. Without the last step HPC would hardly be useful for astrophysics.

2. INAF-HPC resources

Along the path to more and more poweful computing resources, this is a particular time for astronomical HPC in Italy, in the light of the new three-year contract between CINECA and INAF, an initiative which grants fresh resources. It is worth commenting on alternative solutions proposed before this decision.

The alternative choice of an INAF-owned supercomputer would have required the allocation of dedicated resources: a few dedicated rooms to create a center in one of its observatories, dedicated high speed connections and dedicated staff for hardware and software maintainance, help in program development, consultancy, 24 hours running and surveillance. The amount of human resources required would have been of the order of ten highly specialized persons: quite expensive and hardly achievable with present time shrinking budgets. Also this choice should have been tied to a long term strategy to maintain this computing center, since people in Italy is hired for ≈ 30 years while computers are obsolete after 3 years.

Another alternative proposal was to upgrade, maintain and make available to the whole INAF some facilities for HPC of intermediate size already available at some observatories: e.g. the Compaq machine in Palermo, the IBM machine in Catania, the Beowulf machines in Torino and probably others. While these machines are quite good for solving the intermediate-sized HPC in astrophysics for which they have been acquired, upgrading them to serve as multi-purpose machines would require significant investments and thus, for what said above, this choice would just multiply the personnel, room and resources problems that a single INAF-owned machine would have. This choice cannot satisfy national HPC needs for astronomy.

All considered the choice of the INAF-CINECA three years contract appears (fortunately!) as the most reasonable one at this time. However to maintain local facilities, good for specific set of problems is very important in order to maintain "bio-diversity" in a rapidly evolving and competitive environment, as HPC is. Also many medium-sized projects of astronomical HPC may be run at these facilities, in order to relieve the pressure on the limited resources available for astronomy at CINECA. The fact that most of these local machines are typically acquired with funds not from INAF is an additional point in their favour.

3. HPC and data analysis

Episodically, within this meeting and elsewhere, people has mentioned the possibility to use a High-Performance-computer such as the SP4-IBM for massive retrieval, handling and analysis of observational data. This is a rather delicate problem. For some areas, such as X-ray observations, expert scientists prefer powerful workstations, even when massive parallel machines are available. Other areas, e.g. the analysis of extensive optical data sets collected with wide-field CCDs, may indeed benefit from the use of HPC; however a back of envelope calculation invariably shows that the overall requirements of italian scientists involved in these observations easily overwhelms the limits of the INAF-Cineca three years contract. Probably an ad-hoc contract, or an ad-hoc machine, possibly a Beowulf array, could be more appropriate.

More in general, extensive data analysis have problems and methods radically different from HPC and require radically different technologies and computers, to the point that attempting to find a technical solution in common with HPC may lead to an unsatisfactory solution for all.

4. How is HPC considered in Astronomy?

Many of us often face nagging questions from our colleagues. Let me present here just a few representative examples:

How can I trust your code and your results?
This question reflects the feeling that a code is a complex rendering of a complex model of reality. The task of translating the physical theory into a numerical scheme and then into a numerical code is prone to numerical errors, numerical approximations etc., not to mention the possibility of programming mistakes or of misterious bugs. A possible way to answer this question is to validate a code with respect to experiments (see, for instance, Calder et al., 2002), with respect to well known solutions (as in Woodward and Colella, 1984) and/or to benchmark different codes, i.e. to run different codes on the same, specific, clearly stated, problem, then to evaluate and to compare the relevant results and, if needed, find a reason for possible differences.

Why should we spend money for super-computing?
The simple answer is "Because this is the only way to solve many fundamental problems of Astrophysics".

However the un-easiness that experimentalists, observers, some bosses etc. express when they ask this question, on one hand may be the small scale materialization of Italy's disregard for HPC (see next section), on the other may disappear by proving the usefulness of HPC to understand complex astrophysical phenomena or by finding a connection with the work of experimentalists, as discussed below. It is worth noting that the cost of HPC in astronomy is small relative to those of other big astronomical projects: as far as the INAF spending at CINECA is concerned the cost per person per year is rather small: 300 KEuro/y for a community of approximately 70 scientists amounts to just 4.3 KEuro/y per person.

Are your calculations useful for "real" astronomy?
or even tougher
Can your code explain these data?
I tend to consider these questions very important, having always interacted, since my thesis work, with experimentalists. So I believe that HPC astronomers cannot ignore the implicit request of contact with data and observations, but rather should seek the connection with observations, lab experiments and the like. Not only explaining why HPC is important for Astronomy but, more practically, trying to put HPC results in a context and format more familiar to other astronomers. An important aspect to put into light is the physical insight obtained through the computations and the diagnostic power gained when comparing the relevant results with observations. Observations are the templates we must compare our results with; they tell us which set of parameters lead to realistic simulations and which do not, i.e. observations can constrain our calculations.

Striving for the comparison requires specific ways for bringing code results closer to observations; an example is synthesizing emission from plasma simulations and taking into account instrument response to derive the focal plane data (Maggio et al., 1994); or, in the context of cosmological N-body simulations, generating galaxies fields from code results, and then compare simulated morphological characteristics with the observed ones to restrict the range of realistic values of simulations parameters (for extended treatments, see the books by Saslaw, 2000 and by Peacock, 1999).

5. Are italian HPC resources adequate to italian needs?

Simply, NO.

While most of this paper is devoted to astronomical HPC, in this section I prefer to consider italian HPC more globally. An objective way to answer the above question is to compare what various nations do, using both their Gross National Products (GNP) and their HPC resources. http://www.scaruffi.com/politics/gnp.html and http://www.top500.org/ provide info on worldwide GNP and HPC resources, respectively.

From the first site one discovers that the USA GNP ($10,533 billions) is slightly more than **eight times** the corresponding Italian GNP ($1,260 billions). But when one compares the HPC resources (for Italy almost entirely concentrated at CINECA's SP4) one easily sees that USA has a HPC power probably even **more than hundred times** the italian one (adding up all resources available in both countries).

An even more disconforting comparison can be made with Japan ($GNP =$ $4.852 billions, three times the italian one) which at the time of writing (November 2002) has the most powerful supercomputer of the world, the 5120-processor NEC "Earth simulator" in Yokohama (http://www.es.jamstec.go.jp/).

Why is it so? To begin with, USA (and most technologically advanced countries) allocate $\approx$ 3% of GNP to research, Italy only $\leq 1\%$ at present time, and this fact by itself widens the gap already present among GNPs; but, since the factor between HPC resources is significantly larger than twentyfive (the ratio of resources allocated to research) something else has to happen to increase even further the gap between USA HPC and the Italian one, most likely decisions on the policy of scientific and industrial development. However most other advanced countries consider HPC as an index and a tool of scientific and industrial development.

The list of the top 500 supercomputers in the world, and the above considerations clearly show that Italy as a whole has very few HPC resources, with respect to most industrialized countries. Therefore, Italy should increase the fraction of GNP assigned to scientific research and, within that fraction, should allocate more resources to develop other centers equivalent to CINECA. Furthermore, but at this point it is no surprise, Italy lacks national initiatives of the kind of the Accelerated Strategic Computing Initiative, sponsored by the Dept. of Energy in USA, to promote HPC.

6. Possible paths to improve HPC in Astronomy

Obviously, in order to progress, we need to keep up with technical improvements, look for more challenging problems, and try to develop and use more advanced codes. We can benefit from a better interaction with other scientists, not only with astronomers in Italy and abroad, but also with physicists, applied mathematicians and computer scientists.

As I discuss below, however, we may try a qualitative change of the way we select our problems and allocate our national HPC resources. The idea is somehow based on the way other nations have made HPC to progress.

As stated above, italian Astronomy has moderate HPC resources, nonetheless it has access to a big computer, namely the CINECA-SP4. In my opinion this is a good opportunity for this community to grow, in scope and quality of problems through the small scale version of the biggest HPC initiatives: we can set and tackle, within the moderate amount of resources available, bigger challenges in HPC by selecting and supporting **key projects**, i.e. important problems which can be tackled with our resources and which are interesting in the fore-front of astronomy. As stated in the Introduction, italian HPC includes several fields and each may give birth to more than one big key project.

These **key projects** should have a significant impact in the fore-front of astrophysics, could be executed only with the biggest machine available (presently the SP4 at CINECA) and should require a substantial fraction of the computer resources available to INAF. These requirements would make sure that italian astronomy can promote the "best" and most advanced projects and get a good return of the resources assigned to HPC by INAF. Computing resources should be allocated to key-projects in a competitive way in order to make sure that only significant and very good projects get supported. On one hand this would force astronomers to address big challenges and on the other to think hard to the best problems.

At the time of writing this idea has already materialized in the first call for proposal of "key projects" for HPC within INAF (the selection has not been completed yet) and we may expect that we have to keep improving on this idea.

Small projects should still be supported in order to allow the growth of new ideas but, most likely, the path for a real qualitative improvement in Astronomical High Performance Computing in Italy has to go through relatively few, important, competitive **key projects**.

Acknowledgements. Part of this work was supported by the MIUR funds for PRIN-COFIN

References

Calder, A.C. et al. 2002, ApJ, 143, 201

Maggio, A., Reale, F., Peres, G., Ciaravella, A., 1994 Comp. Phys. Comm, 81, 105

Peacock, J.A. 1999, Cosmological Physics, Cambridge University Press.

Saslaw, W. C., 2000, The distribution of the galaxies: Gravitational Clustering in Cosmology, Cambridge University Press.

Woodward, P., and Colella, P., 1984, J. Comput. Phys., 54, 115