Mem. S.A.It. Suppl. Vol. 3, 380

Memorie della
Supplementi

The TRIWULF Project: year one results ⋆

S. Zaggia, G. Taffoni, P. Bonifacio, P. Di Marcantonio, P. Molaro, and P. Santin

INAF - Osservatorio Astronomico di Trieste, via Tiepolo 11, 34131 Trieste

Abstract. The TRIWULF project started in 2002 at the OAT in collaboration with the Dip. di Astronomia of the Trieste Univ. aims at fully exploiting the capabilities offered by high performace computing *Beowulf* systems for the Italian astronomical community. In this brif review we summarize the first year of activities mainly focused on the porting and testing of the specific astronomical applications.

Key words. Wide field imaging – Parallel processing – Image analysis – Spectroscopy

Parallel processing of astronomical images through *Beowulf* systems is actually the most promising way (in terms of cost/performance) to handle, process and analyze big datasets of either images and spectra. The processing of wide field CCD mosaic images, like the ones coming from the ESO/WFI instrument, is intrinsically a parallel process since each CCD frame of the mosaic is independent from the others. Therefore, each frame can be fully processed by a *Beowulf* node, reducing substantially the processing time. Recent experiences demonstrate that in the case of an 8 CCD mosaic image a *Beowulf* cluster can attain an improvement of factor $4 \div 7$ compared to the single CPU processing time. Also for big datasets of fiber spectra, like the ones available with the ESO/FLAMES facility, the *Beowulf* can become a powerful tool in the data analysis.

Only in recent years *Beowulf* systems have been adopted by the Italian astronomical community and in many cases they are still experimental. In order to fully exploit the capabilities offered by this systems the Osservatorio Astronomico di Trieste (OAT) have started a pilot project in close connection with the Dipartimento di Astronomia dell'Università di Trieste and the Osservatorio di Capodimonte aiming at acquiring experience in the wide field image processing and in the multi-spectra analysis. This project started with the support of the 2001 COFIN/MIUR (Sedmak et al. 2003).

The TRIWULF (TRIeste beoWULF) cluster built at the OAT aims at addressing some of the open issue of the overall *Beowulf* system optimization, in connection with the development of a fully working parallel data reduction/analysis pipeline. Once the TRIWULF project will be completed it will be of invaluable help for the exploitation of the enormous capabilities offered by modern imaging instrumentation like WFI and the future VST, VISTA and

Send offprint requests to: S. Zaggia
⋆ http://wwwuser.oat.ts.astro.it/triwulf/
Correspondence to: via Tiepolo 11, 34131 Trieste

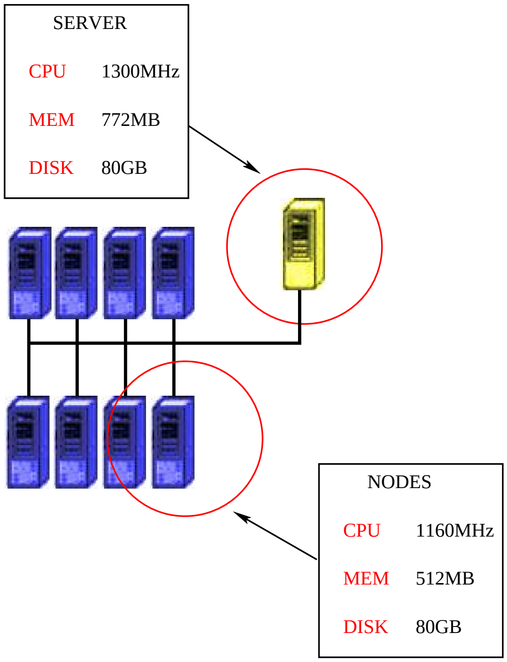

Fig. 1. Configuration of the TRIWULF cluster.

LBT cameras and also for multi-fibers spectrographs like the ESO/FLAMES facility.

1. The TRIWULF cluster

The TRIWULF cluster computing facility (Fig. 1) is a standard *Beowulf* cluster. It is based on a "Clients-Server" configuration, where a single computer ("Server") acts as firewall toward the outside network and also as disk and authentication server. The Server monitors the activity of the cluster itself using *Ganglia Monitor tools* and performs the the normal and extra-ordinary maintenance of the whole cluster.

The cluster topology requires that a full *Linux* operating system (Debian 3.0) is installed on the server while a minimal base *Linux* system is installed on the nodes. The whole cluster is totally homogeneous (all the Clients are created perfectly identical) and we assured that that the Clients and the Server are self-similar for the directories structure.

The hardware was chosen to assure good computational and network performances. The whole cluster is based on AMD Athlon XP+ CPUs (1300/1160 Mhz, top AMD CPUs in spring 2002) and a traditional 10/100 Ethernet network. The CPU performances is satisfying with a top computing power of 210 millions of floating point operations per seconds. The network allows communication bandwidth between nodes up to 12.5 Mbyte/sec.

To permit the development and the use of parallel software we installd the *Message Passing Interface (MPI)* standard libraries: a message-passing library specifically designed to provide access to advanced parallel hardware like a *Beowulf* cluster.

A whole set of compilers, scientific software and supporting libraries optimized for parallel processing has been installed. We installed Fortran77/90 and C/C++ compilers of different vendors (GNU, Intel, Vast) and tested their different performances on the application softwares. A special set of high performance parallel libraries to compute Fast Fourier Transform and Linear Algebra operation useful for data reduction/analysis has been installed.

2. WFI parallel reduction pipeline

In order to develop a modern, high capacity, fully parallel pipeline for the reduction/analysis of wide field images we have ported and adpted to the TRIWULF cluster the software for the un-supervised reductions of Optical/IR images developed by the *ESO Imaging Survey* Project (EIS, Vandame et al. 2003). The EIS software is based on a collection of C-based libraries integrated with the XML (Extensible Markup Language) technology and a host of new algorithms (mainly based on wavelet libraries). The code is extremely flexible and can deal with a variety of different observing strategies, either in the optical or in the infrared domain, as well with a variety of single- and multi-CCD instruments. Moreover it operates completely *in*

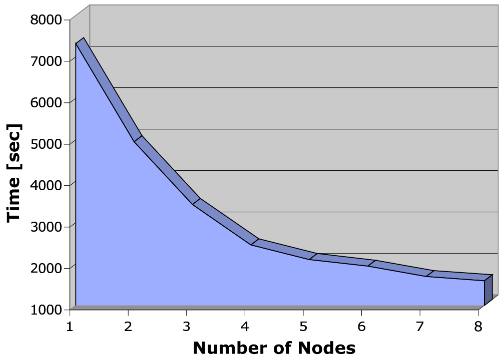

Fig. 2. Performance of the image "splitting" procedure varing the number of the nodes.

memory speeding up computations avoiding too much frequent disk acesses.

In order to test the EIS pipeline performances on the TRIWULF in comparison with its scalar version, we have selected 2 nights of ESO/WFI images (taken in July 2000), for a total of 293 images, or $\simeq 43.0$ Gb, of raw data which are part of the EIS *Pre-FLAMES* survey (Zaggia et al. 2001; Momany et al. 2001) and consists of observations of stellar fields for the support of the ESO/FLAMES facility. Our first goal was to optimize the data distribution over the cluster in order to reduce as much as possible network transfers. The first test we have done is the "splitting" procedure where the FITS multi-extension raw images are broken in the different CCDs and the header cleaned, rewritten and adapted to the pipeline. This is the more time consuming part of the code: on a single CPU it took $\sim$121 minutes. In the parallelized version we distributed the computing load over the whole cluster, assigning to each node a group of raw images and distributing the output of the splitted images over the cluster: i.e., CCD#1 goes to node #1, CCD#2 node #2, and so on. The results of this exercise is shown in Fig. 2 where the splitting time is plotted against the number of nodes used. With 8 nodes the "splitting" procedure takes $\sim$26 minutes with a gain of nearly a factor 4.5. As it can be seen from Fig. 2 the gain saturates after 5 nodes due to the limitations of the 10/100 Ethernet network. Since all the subsequent operations are done on a single CCD/node basis and the network communications are limited the prospected overall gain is around a factor $\simeq 7$. More tests are underway.

3. Analysis of FLAMES data

The TRIWULF cluster has also been used to analyze rapidly and efficiently the multi-fibers spectra delivered by the ESO/FLAMES facility. For this purpose we have developed an automatic software for the determination of abundances from stellar spectra Bonifacio & Caffau (2003). The code performs χ^2 fitting of synthetic spectra to selected stellar features, the minimum is sought using MINUIT (James 1998). The TRIWULF system is used, for this particular problem, in a "farm" configuration (each node is completely independent from the others). We developed a C wrapper using the MPI libraries which reads as input a multi-fibers spectra, splits it in N groups and send each group to one of the N available cluster nodes to perform the computation (χ^2 fitting).

In order to test the scalability and performances of our pipeline, we tested the code with the FLAMES Science Verification data of the globular cluster NGC 2808 (Primas 2003). We selected one of the frames in which 57 Giraffe spectra of RGB stars are available and run the code on different number N of nodes. The results are shown in Fig. 3 where the computational time is plotted as a function of the number N of nodes. The gain in performance is a factor of $\simeq 6.5$ with 8 nodes. As a consequence of the network performances the reduction of computational time flattens out when more than five nodes are in use.

To get an idea of the power of the system, we point out that a "traditional" determination of abundances analysis *line by line, star, by star* of 57 Giraffe spectra, which requires only $\simeq$ 69 seconds on

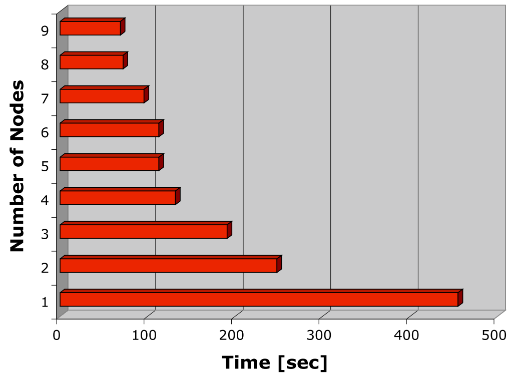

Fig. 3. Computational time of the FLAMES abundance analysis as a function of the number of TRIWULF nodes.

8 TRIWULF nodes, would require at least **57 days of human work**. With this computing power it is feasible to perform extensive and realistic Monte Carlo simulations (of the order of 1000 "events", where one "event" requires the analysis of ~ 400 simulated spectra) which allow reliable estimates of the errors associated with the abundances determination. The actual implementation of the code is extremely robust: we kept one such Monte Carlo parallel simulations running for three days in a row without encountering any problem.

4. Future Prospects

The actual implementation of the TRIWULF system at the OAT demonstrates that *Beowulf* systems can be extremely flexible and powerful for the development of dedicated pipelines for image and spectra data reduction/analysis. With a limited amount of manpower we have been able to develop specific parallel pipelines with an overall gain in computational time of at least a factor $\simeq 6$. This can be still improved with the possible upgrade of the internal network to a wide-band Ethernet connection (Gigabit).

Acknowledgements. We would like to thank Benoit Vandame and Luiz Da Costa for their kind support with the EIS pipeline.

References

Bonifacio P. & Caffau E. 2003, A&A, 399, 1183

James F., 1998 MINUIT, Reference Manual, Version 94.1, CERN, Geneva, Switzerland

Momany Y. et al. 2001, A&A, 379, 436

Primas F., 2003, The Messenger, 112, 3

Sedmak G., Cascone E., Russo G., & Zaggia S. 2003, Memorie della Societa Astronomica Italiana, 74, 415

Vandame B. et al. 2003, A&Ain preparation

Zaggia S. et al. 2001, The Messenger, 105, 25