Design of a Functional Architecture for the Management of Cluster Resources and Services through the Web

by Juan Antonio Ortega, Jorge Cantón, Ana Silva, David Bosque and Francisco Velasco

In 2006 Junta de Andalucía created the Andalusian Supercomputing Network (RASCI). RASCI consists of supercomputing nodes distributed geographically throughout Andalusia that provide the region with a large number of computing resources. Increased network bandwidth, more powerful computers and acceptance of the Internet have driven a growth in demand for new and better ways to utilize high-performance technical computing (HPTC) resources.

The Andalusian Scientific Computing Centre (CICA) is a supercomputing centre in Andalusia that provides supercomputing nodes and the Andalusian e-Science Portal to RASCI (see Figure 1).

CICA Infrastructure
Currently, CICA has a distributed cluster. It consists of a set of nodes with Intel Core 2 Duo or Intel Quad Core CPUs and 4 GB or 8 GB of RAM, connected through gigabit ethernet. The batch systems installed are the Sun Grid Engine (SGE) and Condor interface. At present, the Web interface application allows users to queue their jobs into Condor and SGE for subsequent execution by the cluster.

CICA promotes the use of open-source software, and the supercomputing nodes run under Linux, make use of open-source libraries and utilities, and use several OS.

The Andalusian Supercomputing Network will allow nodes located at different universities and research centres to be joined together in a unified computing system. To achieve this, a technical committee (for infrastructure management), an access committee (for access management) and a strategic committee have been established.

Other services are offered by CICA to users of the Andalusian Supercomputing Network, including assistance in parallelizing applications to take advantage of the Grid infrastructure. This involves providing programming courses, expert advice and consultancy for optimization and parallelization of applications and the use of message-passing libraries for the efficient use of distributed memory. With these services CICA intends to increase the number of users who can potentially take advantage of the super-
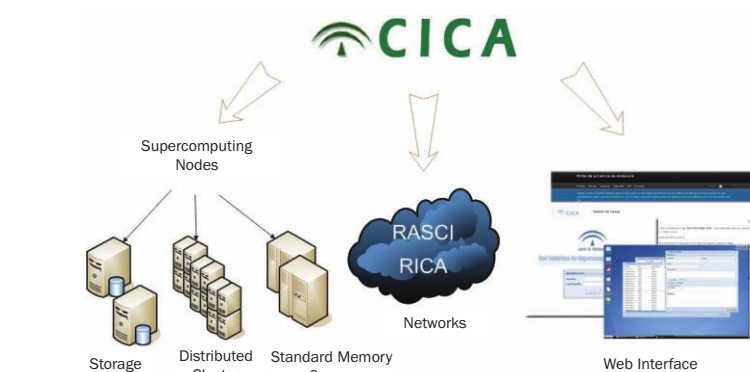

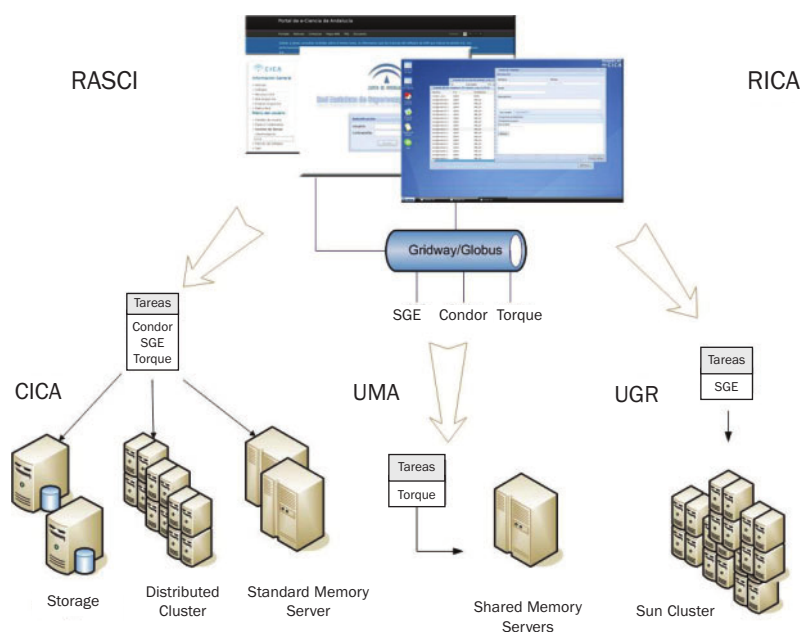*Figure 1: CICA connections.*

Figure 2: Supercomputing infrastructure.

computing infrastructure, rather than limiting its use to those who are already technically proficient (see Figure 2).

All of this makes it possible to bring supercomputing technology to a heterogeneous community, allowing for faster technological progress, and also to increase the demand for this type of service, due to its simplicity of use.

Web Interface
Rather than being graphical or accessible through a Web interface, cluster queue manager interfaces tend to be accessed through a command line interface from a Linux or Windows console.

Different interfaces have been developed for the Condor queue manager, each with a specific purpose depending

on who is making use of it. One of the services to which this queue manager's developers have contributed is the use of Web services (SOAP-WSDL). For its implementation they have developed the BirdBath package, which, by means of WSDL calls, allows jobs in the cluster to be submitted and managed. The main inconvenience presented by these calls is the ease of discovery and indexation by search engine Web crawlers, since they are accessible though URLs. They can thus be found and stored through search engines. Moreover, the management capabilities are limited to the functionalities implemented by BirdBath developers.

The architecture presented by the CICA-named Cluster Looking Class (CLG) allows the user to make use of resources through the Web, eliminating the need to enter commands through a command line console. This ease of use extends the reach of computing services to a greater number of users as a technical profile is no longer required. At present, this interface is operational at the Andalusian e-Science Portal, giving registered users access to computational resources connected to the Andalusian Computing scientific Network (RICA) and Andalusian Supercomputing Network (RASCI).

However, with the ability to access the system through the Web, security becomes a critical aspect. In a system such as this, the high level of computing power turns it into a very powerful tool for illicit uses such as brute-force password cracking. In order to avoid this problem, CLG defines both a logical and a physical architecture that allows the user to manage and submit jobs and to receive their results, uncoupling the physical architecture of the cluster from the Web interface and thus putting an end to security problems.

Further Work
In the future, the supercomputing resources of CICA will grow and will be included in this infrastructure. Our estimation is that this will lead to bigger massive storage, shared memory servers, more distributed nodes and new batch systems.

Links:
Andalusian e-Science portal:
https://eciencia.cica.es
Condor Project HTC:
http://www.cs.wisc.edu/condor

Please contact:
Juan Antonio Ortega
Centro Informático Científico de
Andalucía, Spain
E-mail: ortega@cica.es

Exploitation of Cell Multi-Processor Array in Solution of Spatio-Temporal Dynamics

by Zoltán Nagy, László Kék, Zoltán Kincses, András Kiss and Péter Szolgay

Array computers can be useful in the numerical solution of spatio-temporal problems. IBM has recently introduced a topographic array processor called the Cell Broadband Engine (Cell BE). Researchers at the Cellular Sensory Wave Computers Laboratory of SZTAKI in collaboration with the Department of Information Technology, Pázmány Péter Catholic University, Budapest and the Department of Image Processing and Neurocomputing, Pannon University, Veszprém, have implemented a Cellular Neural Network (CNN) simulation kernel on the Cell BE. The CNN simulation is optimized according to the special requirements of the Cell BE and implements both linear and nonlinear (piecewise linear) templates. We have used the CNN simulation kernel to solve Navier-Stokes partial differential equations (PDEs) on the Cell architecture.

Array processing is a good candidate for increasing computing power by using parallel computation. Additionally, it can help to solve architectural problems (eg distribution of control signals on a chip). The effectiveness of implementations is measured by a set of parameters, namely by the silicon area (A), execution time (T), dissipated power (P), consumed energy (E) and number of input/output pins (#I/O).

A number of different implementations of array processors are commercially available. In this project we have mainly concentrated on the topographic IBM Cell heterogeneous array processor architecture, because its development environment is open source and we wished to compare the results with those from our previous FPGA (field programmable gate array)-based implementations.

The Cell Broadband Engine Architecture (CBEA) is designed to achieve high-performance computing with better area/performance and power/performance ratios than the conventional multi-core architectures. The CBEA defines a heterogeneous multi-processor architecture where general-purpose processors called Power Processor Elements (PPE), and Single Instruction Multiple Data (SIMD) processors called Synergistic Processor Elements (SPE), are connected via a high-speed on-chip coherent bus called an Element Interconnect Bus (EIB). The CBEA architecture is flexible and the ratio of the different elements can be defined according to the requirements of different applications. The first implementation of the CBEA is the Cell Broadband Engine (Cell BE or informally Cell) designed for the Sony PlayStation 3 game console, which contains one PPE and eight SPEs.

In this work we have concentrated on an efficient CNN implementation on the Cell architecture. The basic CNN simulation kernel was successfully implemented on the Cell BE, and both linear and nonlinear CNN arrays can be simulated. The kernel was optimized according to the special requirements of