Thesis Overview

Methodology to evaluate performance of the I/O systems on High Performance Computers

Sandra Méndez

Computer Architecture and Operating Systems Department (CAOS)
Universitat Autònoma de Barcelona, Barcelona, Spain
Advisor: Dra. Dolores Rexachs
sandra.mendez@uab.es, dolores.rexachs@uab.es
Supported by the MICINN Spain, under contract TIN2007-64974 and TIN2011-24384.

The increase in computational power of processing units and the complexity of scientific applications which use high performance computing require more efficient Input/Output (I/O) systems. To use the I/O systems efficiently it is necessary to know its performance capacity to determine if it fulfills I/O requirements of the application. Evaluating the performance capacity of the I/O system is difficult due to the diversity of architectures and the complexity of its software stack. Furthermore, parallel scientific applications will have different behavior depending on their access patterns. Then, it is necessary to have some method to evaluate the I/O system performance taking into account the applications access patterns. We propose a methodology to evaluate performance of the Input/Output systems on High Performance Computers based on an I/O model for the parallel scientific applications [1] [2] [3]. The proposed methodology is composed of three stages: Characterization, I/O analysis and Evaluation. Figure 1 shows the proposed methodology.

CHARACTERIZATION The characterization is applied to the I/O system and parallel scientific application. This stage has two objectives: i) Extracting the I/O model of application; and ii)

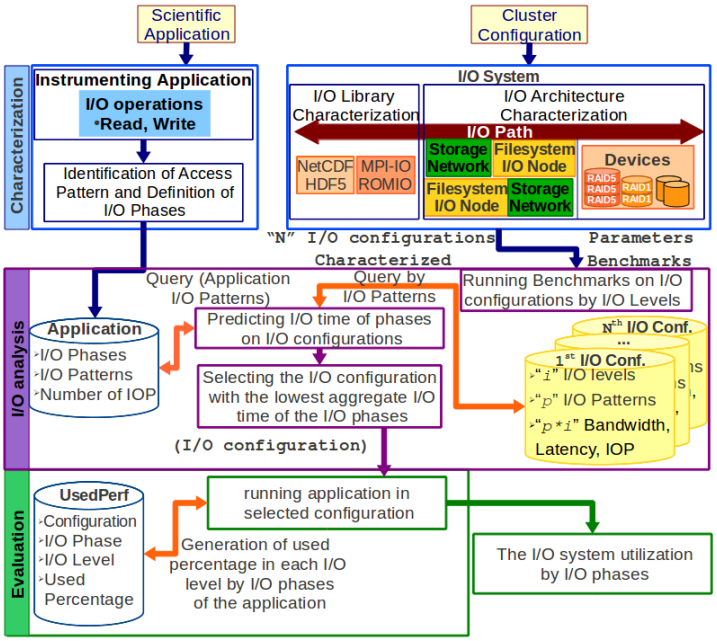

Figure 1: Methodology to evaluate performance of the Input/Output systems on High Performance Computers

Extracting the performance characteristics of the different I/O configurations of the I/O system. These activities are independent. The characterization of application is done off-line and the application I/O model can be applied to analyze different target systems. The application I/O model is defined by three characteristics: metadata, spatial global pattern and temporal global pattern. The I/O model of the application is expressed by I/O phases, where an I/O phase is a repetitive sequence of same pattern on a file for a number of processes of the parallel application. A phase will be significant depending on data transferred that represent its *weight*.

I/O ANALYSIS In this step we evaluate the I/O performance for the application I/O model for different I/O systems on the computer cluster. In this process it obtains a transfer rate for different access patterns that compose the I/O phase and it is posible estimate the I/O time for each I/O phase. Furthermore, the I/O model is used to set up the input parameters of the benchmark IOR [4] when there is not a previous characterization for a specific I/O pattern.

EVALUATION It evaluates the utilization of the I/O system by the relation between the bandwidth characterized at I/O devices level and bandwidth measured expressed as $SystemUsage(phase[i])$. When a phase has two or more I/O operations then the bandwidth characterized is calculated as the average of the bandwidth characterized of each I/O operation that composes the I/O phase.

OUR METHODOLOGY IN CLOUD PLATFORMS Cloud computing is gaining popularity in many areas, including High Performance Computing (HPC). Beside being CPU-intensive, HPC applications are sensitive to network bandwidth/latency as well as the performance of I/O storage systems, making them particularly challenging to run efficiently in clouds. Besides the well-known elasticity in resource acquisition, usage, and releasing, cloud platforms provide a level of congurability impossible with traditional HPC platforms. Unlike supercomputers or in-house clusters, which only offer fixed, one-size-fits-all configurations, clouds allow users to customize their own virtual cluster and perform reconfiguration at the user, application, or even job level [5]. In this context, our methodology can be applied in cloud platforms. The I/O model of a specific parallel application can be used to estimate the I/O time for different configurations and use this information to select the more convenient for the I/O characteristics. The setting up defined for IOR allows the user to apply a fast method for testing the I/O performance for an specific I/O model. Due to the I/O model is independent of the I/O platform, this can be used in different I/O systems configured by the user.

References

[1] S. Méndez, D. Rexachs, and E. Luque, "Methodology for Performance Evaluation of the Input/Output System on Computer Clusters," in *Workshop IASDS on Cluster Computing (CLUSTER), 2011 IEEE International Conference on*, sept. 2011, pp. 474 –483.

[2] S. Méndez, J. Panadero, A. Wong, D. Rexachs, and E. Luque, "A new approach for analyzing I/O in parallel scientific applications," in *CACIC 12, Congreso Argentino de Ciencias de la Computacion*, 2012, pp. 337–346.

[3] S. Méndez, D. Rexachs, and E. Luque, "Modeling Parallel Scientific Applications through their Input/Output Phases," in *CLUSTER Workshops'12, 2012 IEEE International Conference on*, sept. 2012, pp. 7–15.

[4] H. Shan and J. Shalf, "Using IOR to analyze the I/O performance for HPC platforms," LBNL Paper LBNL-62647, Tech. Rep., 2007. [Online]. Available: www.osti.gov/bridge/servlets/purl/923356-15FxGK/

[5] M. Liu, J. Zhai, Y. Zhai, X. Ma, and W. Chen, "One optimized I/O configuration per HPC application: leveraging the configurability of cloud," in *Proceedings of the Second Asia-Pacific Workshop on Systems*, ser. APSys '11. New York, NY, USA: ACM, 2011, pp. 15:1–15:5.