Jornadas Lip 2014

Lisbon, Portugal

Advanced Computing

V. Oliveira[1], A. Pina[1],J. Rufino[2], A. Alves[2]

{vspo, pina}@di.uminho.pt, {rufino, albano}@ipb.pt

[1] Departamento de Informática, Universidade do Minho,
Campus de Gualtar, Braga, Portugal

[2] Departamento de Informática e Comunicações, Instituto Politécnico Bragança
Bragança, Portugal

21th Mars 2014

Outline of the presentation

1. Introduction
2. Objectives
3. Research
4. Education

Overview

In most sciences the amount of both experimental and simulated data has been increasing because instruments are getting much better and cheaper and storage costs have been decreasing dramatically.

ATLAS

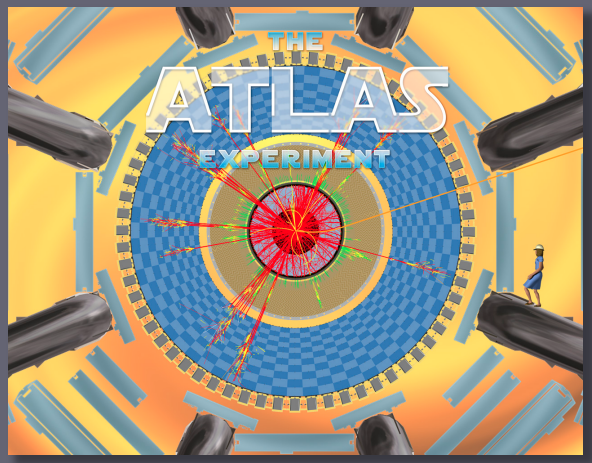

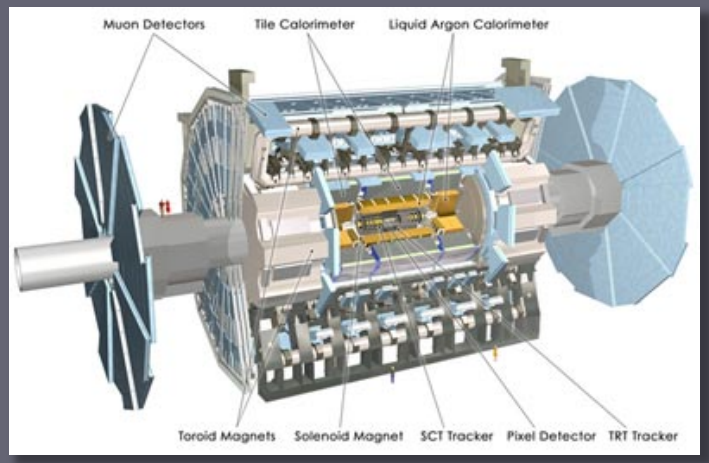

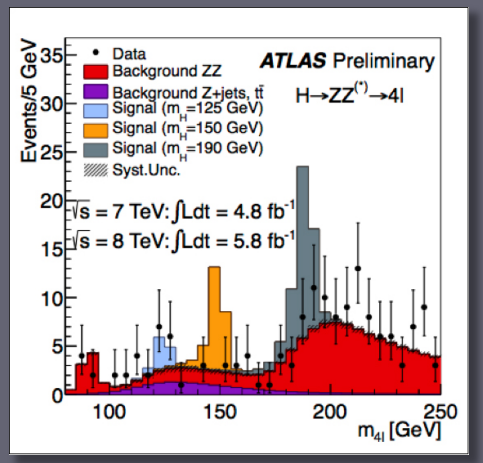

The Large Hadron Collider's ATLAS experiment generates a huge volume of information.

- The trigger system transforms 40M events/s in 200 events/s using 3 layers of specialized and general purpose hardware;
- Those 200 events/s require around 320 MB/s of storage bandwidth just for the raw data;
- Each year produces 3.2 PB of raw data, 2PB of reconstructed data and 200TB of physics data.

Analysis applications

Data analysis is performed by applications that:

- uses analysis frameworks based on ROOT;
- developed and tested during the course of several years.

In order to remain competitive in the ATLAS collaboration results must be delivered in useful time, but:

- the solution quality is dependent on the amount of time that is available to process events;
- and the amount of the information to analyse is very large…

Even bigger data

… but the ATLAS data set will grow one order of magnitude following the 2013-2014 machine operation conditions upgrade!

So a major challenge is now presented:

How to deal with a 10-fold increase in the data volume of each analysis?

- while supporting the increase in the complexity of the analysis applications,
- and reducing the turnover time of the results.

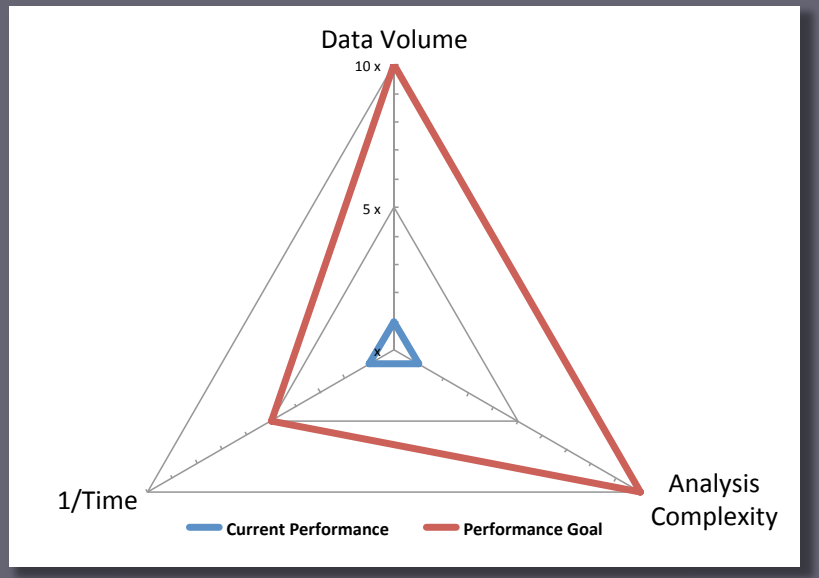

And all of this on a very tight budget!

LIP-Minho

The Advanced Computing team, integrated in LIP-Minho, since 2013, intends to:

- focus on advanced computing research on HEP related applications and on
- bridging the Computer Science academic courses available there with the LIP research interests.

Approach

The main issues on scientific computing are no longer confined to homogeneous parallelism and conventional job task scheduling/allocation.

Now nodes on emerging computer clusters include:

- several multi-core CPUs,
- accelerator devices with their own disjoint memory,
- several levels of hierarchy both in memory and in storage,
- with several types of networks connecting those highly heterogeneous components.

Software and hardware

The software and hardware configurations must be adjusted closely to the applications that will be used and not to the average case;

A thorough profile of the application execution becomes essential to restructure the system and the applications to obtain the desired performance benefits.

Previous work

The team has already explored application behaviour analysis, particularly using IO profiling, to optimize execution time of the analysis programs developed by the ATLAS team at LIP-Minho

- the work provided significant performance gains to the ATLAS team, and even suggested a more economic system architecture with nodes that can both compute and store data.
- these architecture can perform better than centralized systems of similar capacity, providing a smooth upgrade from current systems and address scalability, resource efficiency and cost concerns.

V. Oliveira, A. Pina, N. Castro, F. Veloso, A. Onofre, *Even Bigger Data: Preparing for the LHC/ATLAS Upgrade*, 6th Iberian Grid Infrastructure Conference (Ibergrid'2012), Lisboa

Directions

- to make systematic analysis of a range of HEP applications, resorting of complete execution profiles, both in terms of CPU and IO, to evaluate the optimization opportunities from the computational points of view;
- to use profilers to optimize the underlying execution system or to build new systems more adapted to the applications that are most widely used;
- towards data-location aware scientific computing, which adaptively improve execution efficiency by reducing highly expensive data migrations and contentions on their scheduler strategies, both across primary memory modules and storage units.

Data distribution

- data should be distributed and replicated between the available nodes, presenting users a unified view of the file structure;
- in complement there should be a scheduling strategy that places applications near the resources they effectively require by using adapted file systems or by using virtualization as a mechanism by which VMs are migrated to the nodes where data is being sourced;
- the local job scheduler should also be improved with storage related restrictions and limit concurrent accesses to levels that provides optimized storage performance.

Computer Scientist and Physicists

- a tighter cooperation between Computer Scientists and Physicists, with training and knowledge transfer, seems fundamental for the development of better models, techniques and tools that lead to faster advancement of science by promoting an effective use of scarse human and computing resources;
- CERN School of Computing has become a highly regarded reference, training physicists, computer scientists and engineers from all over the world in HEP applications.

CSC2014

- CSCs consist of a two-week period of theoretical and hands-on tuition for students coming from all over the world, featuring advanced academic topics on scientific computing;
- the school's main aim is to create and share a common culture in the field of scientific computing, which is a strategic necessity to promote mobility between scientific institutions, and to carry out large transnational computing projects;
- the 2014 edition, which will be hosted by the University of Minho and LIP at Braga, is a major achievement that may be further expanded with related activities in the following years. *https://indico.cern.ch/event/298406/*

The End

Any questions?

António Pina

pina@di.uminho.pt