

CERN-LHCC-2005-018  
ALICE TDR 012  
15 June 2005

# ALICE

## Technical Design Report

of the

## Computing

Cover design by CERN Desktop Publishing Service.  
Visualisation of ALICE Experiment as simulated by AliRoot.

Printed at CERN  
June 2005.

ISBN 92-9083-247-9

# ALICE Collaboration

---



# Contents

---

<b>Summary</b>	<b>0</b>
<b>References</b>	<b>3</b>



# Summary

---

The principal objective of this document is to present the computing model of ALICE (A Large Ion Collider Experiment) at the CERN [1] Large Hadron Collider [2] (LHC) and the current estimates of the computing resources needed to implement this model. The content of this document is the result of a long series of consultations and discussions within the ALICE Collaboration and its governing bodies. At the beginning of 2005, a LHCC review [3] examined the computing resources requested by the LHC experiments and found these requests reasonable. As expected and announced at the time of the review, the computing model and the projected computing needs have further evolved as a result of the additional experience gained in processing the data produced by the Data Challenges.

At the time of writing, we are slightly more than two years away from the first collisions at the LHC. This is, however, still a long lapse of time because of the fast pace of evolution in the field of Information Technology. In addition, the anticipated needs for LHC computing are very large. So is the complexity of the environment required to process the data and make optimal use of the available resources. Therefore, the deployment and organisation of the software, of the material and of the human resources needed have to be properly planned. This requirement is particularly critical, since the resources will be distributed in many centres around the world which will have to work together in a coherent way as a single entity.

In consideration of the above, this document contains the appropriate level of detail to support the ALICE requests and guide the collaboration in the implementation and deployment of the ALICE software and computing infrastructure, ideally without basing the ALICE computing strategy on elements that can and will still change in the course of the next few years.

The ALICE offline framework (AliRoot) has been under development since 1998. It has provided inputs for the Technical Design Reports of all ALICE detectors and for the performance and physics studies presented in the ALICE Physics Performance Report [4]. The AliRoot framework is based on Object-Oriented technology and depends on the ROOT framework. Although AliRoot already allows quite detailed and realistic studies of the detector, it is still under intense development.

Advanced code inspection tools have been developed in collaboration with computer science experts, and these have now been deployed in production. They play an important role in maintaining the quality and uniformity of the AliRoot code.

Simulation has so far been performed with the GEANT 3 Monte Carlo transport program through the use of a set of interfaces that allow the transparent implementation of other Monte Carlo transport programs. The interface to the FLUKA Monte Carlo program has also been validated, and it is being used in production. In the near future we plan to upgrade the existing interface to the GEANT 4 Monte Carlo transport program and therefore to discontinue the GEANT 3 program. The geometry is described via a modeller developed jointly by ALICE and the ROOT team.

A large amount of work has been dedicated to reconstruct trajectories and identify particles. These tasks are particularly difficult in the context of heavy-ion collisions, as we expect that such collisions will produce a number of tracks an order of magnitude larger than in proton-proton (pp) collisions and that the occupancy can be as high as 40% in some regions. The results obtained from the simulation in terms of efficiency and contamination are very close to the design parameters. More work is being carried out to consolidate these results in conjunction with the calibration and alignment algorithms, still under development.

The complete design of the condition infrastructure (calibration and alignment) is ready. Pilot implementations have already been achieved for a few detectors. Validation tests will be performed during the Physics Data Challenge 2005.

The development of the visualization application has just started and will be continued over the coming year.

The computing model applies to a 'Standard Data Taking Year' (SDTY). During a SDTY, ALICE

will take heavy-ion data for  $10^6$  effective seconds per year (one month), while for the rest of the time, when the accelerator is active,  $10^7$  effective seconds, ALICE will collect proton-proton data. During the initial phase, we assume that the effective beam time may be less, and increasing luminosity will progressively become available. However, the exact operation during the first three years, the so-called ‘initial running conditions’, is being periodically reassessed. Our model takes into account this period through a staging of the deployment of the computing resources, i.e. 20% to be available in 2007 for the first pp run, 40% for the first Pb–Pb pilot run in 2007, and 100% for the first full Pb–Pb run in 2008, even if the nominal luminosity is not yet available. This responds to the requirement to delay the acquisition of hardware as much as possible, every year bringing a reduction in cost that for CPU’s can reach of 30–40%.

The computing model for the pp data is similar to that of the other LHC experiments. Data are recorded at a rate of 100 MB/s. They are reconstructed quasi on-line at the CERN Tier 0 facility. In parallel, data are exported to the different Tier 1s outside CERN (hereafter ‘external Tier 1s’), to provide two copies of the raw data, one stored at the CERN Tier 0 and another copy shared by all the external Tier 1s. All Tier 1s will have collectively enough resources to perform a second and third reconstruction pass.

For heavy-ion data this model is not viable, as data are recorded at up to 1.25 GB/s. Such a data rate would require a prohibitive amount of resources for quasi real-time processing. ALICE therefore requires that heavy-ion data be reconstructed at the CERN Tier 0 and exported during a period of four months after data taking. Additional reconstruction passes will be performed at the Tier 1s.

It is customary to assume that scheduled analysis will be performed at Tier 1 centres, while unscheduled analysis and simulation will be performed at the Tier 2 centres. On the basis of the experience gained with the Physics Data Challenges, this hierarchical model, based on the MONARC [5] work, may be progressively replaced by a more ‘symmetric’ model often referred to as the ‘cloud model’. In the latter model, the only distinctive features of the Tier 1s, apart from size, are service levels and the commitment to store the data safely, most likely on mass storage systems.

The choice of the model finally adopted will also depend on the functionality and reliability of the Grid middleware. Should the middleware have a limited functionality in deciding where to perform the calculations and where to direct the data, a hierarchical model will be useful in organizing ‘by hand’ the computing activity. A middleware implementing a set of functionalities closer to the ‘Grid vision’ could benefit from some more freedom of choice, leading to a usage pattern of the resources similar to the one predicted by the cloud model.

At the time of writing, the functionality of the Grid middleware that will be installed on the LHC Computing Grid [6] (LCG) resources is still evolving. The elaboration of the ALICE computing model has implicitly assumed that a functional middleware will exist, optimizing to some extent the storage and workload distribution. Based on the experience gained with the ALICE-developed AliEn [7] system, it is believed that the application of the cloud model is technically possible. Currently, it is planned to provide the required Grid functionality via a combination of the common Grid services offered on the LCG resources and the ALICE-specific services from AliEn.

To ease the estimation of required resources, each task has been assigned to a specific Tier, in accordance with the MONARC model. Throughout this document the MONARC terminology will be used to discuss the different elements.

Finally, it is important to note that all the information contained in this document is provided to the best of our knowledge. The contents of this document depend on a number of human and technological factors that are in rapid evolution. We anticipate a qualitative as well as quantitative evolution of the ALICE computing model.

The document is organized as follows. Chapter 1 contains a description of the data acquisition system (DAQ) and of the basic parameters of the raw data. These are fundamental inputs for the computing infrastructure and the computing model. Chapter 2 provides an overview of the computing framework together with the condition infrastructure. Chapter 3 describes the ALICE distributed computing en-

vironment and the ALICE experience with the Data Challenges. Chapter 4 describes the simulation infrastructure. Chapter 5 illustrates the reconstruction strategy and the current status of the performance of the algorithms. Chapter 6 contains our current plans for the development of the analysis framework and some prototype implementation. Chapter 7 describes the ALICE computing model and the projected computing needs. Chapter 8 presents the organization and funding structure of the ALICE Computing Project and lists the major milestones of the project.

# References

---

## Chapter 0

- [1] <http://www.cern.ch>.
- [2] <http://lhc-new-homepage.web.cern.ch/lhc-new-homepage>.
- [3] LHCC Review of Computing Resources for the LHC Experiments,  
CERN-C-RRB-2005-006;  
CERN-LHCC-2005-006;  
LHCC-G-089.
  
- [4] CERN/LHCC 2003-049, ALICE Physics Performance Report, Volume 1 (7 November 2003);  
ALICE Collaboration: F. Carminati *et al.*, J. Phys. G: Nucl. Part. Phys. **30** (2004) 1517–1763.
- [5] <http://monarc.web.cern.ch/MONARC>.
- [6] <http://lcg.web.cern.ch/LCG/>.
- [7] <http://alien.cern.ch>.

