A Statistical Analysis of Job Performance on LCG Grid

David Colling, Olivier van der Aa, Mona Aggarwal, Gidon Moont

(Imperial College, London)

GridPP
UK Computing for Particle Physics

Introduction

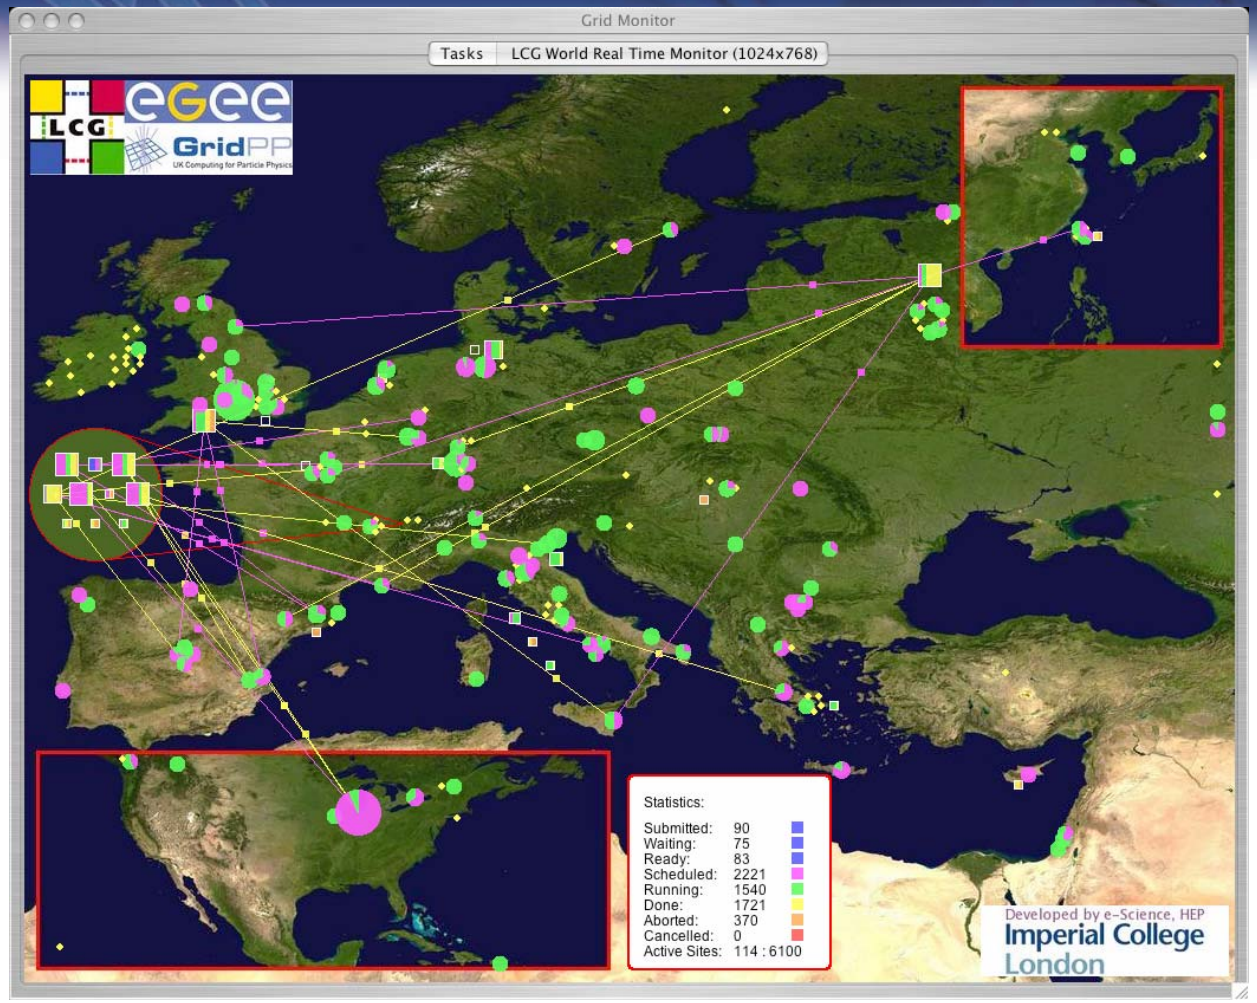

http://gridportal.imperial .ac.uk

Introduction

We decided to keep the data that we gather and to perform some statistical analysis on it. In this talk I will briefly discuss..

- What it can tell us about the different usage of the system by the different VOs
- What it can tell us about the performance of the individual components and the system as a whole
- We now produce daily reports (available from the website)
- In general I will just describe what we see rather than trying to interpret it. That is the next step.

This is still very much work in progress

GridPP
UK Computing for Particle Physics

Caveats

- This view of the LCG is that of the RBs (well actually the LBs)
- We don't see any jobs that are submitted by local users
- We don't see any any grid jobs that are submitted via RBs to which we do not have access (small effect)
- We do not see grid jobs submitted by directly not using an LCG RB. Specifically we do not see jobs submitted by Rod Walker's CondorG submission system.
- These stats are only for the last quarter.

The system

- The LCG is an operational Grid currently running over 200 sites in 36 countries, offering its users access to nearly 14,000 CPUs and approximately 8PB of storage.
- Defining meaningful metrics and monitoring the performance of such a system is challenging exercise but important for successful operation.
- Primary motivation for this research is to analyze LCG performance through a statistical analysis of the lifecycles of all jobs on the grid.
- In this paper we define metrics that describe typical job lifecycles. The statistical analysis of these metrics enables us to gain insight into the work load management characteristics of the LCG Grid [2]. Finally we will show how those metrics can be used to spot Grid failures by identifying statistical changes over time in the monitored metrics.

Analysis Dataset

- The dataset is obtained by
  - the information published by the about 28 Grid Resource Brokers (RBs) across the EGEE grid.
  - Job lifecycle obtained through RBs log files.
  - Dataset are taken from Sept 2005 –Jan 2006
  - More than 3 million jobs.
- The performance metrics are measured for main four LHC VO's:
  - ALICE
  - ATLAS
  - LHCB
  - CMS
- Metrics are defined to measure performance and effectiveness from three perspectives:
  - User
  - Resource
  - Grid

So what can see?

- Number of Active Users in a system at a given time.

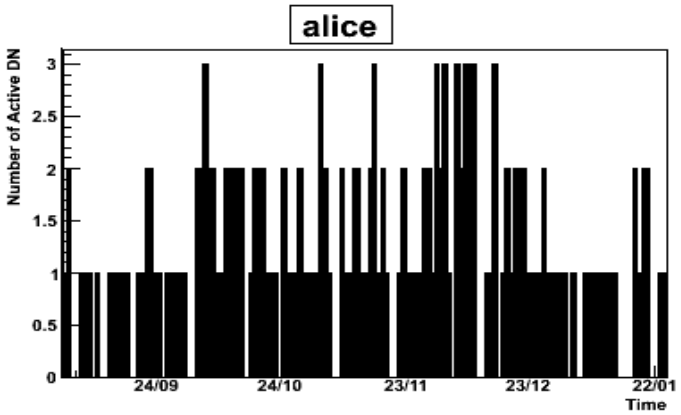

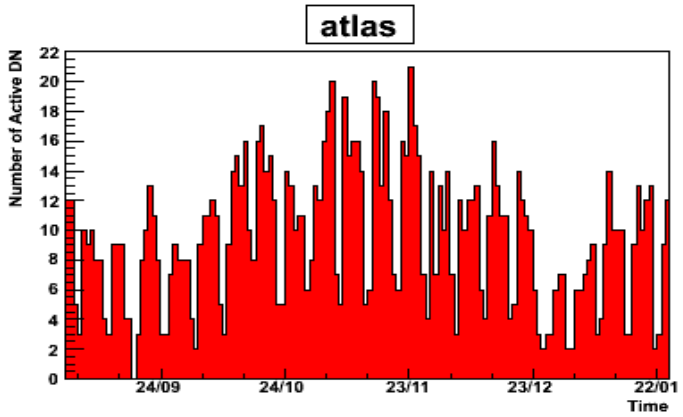

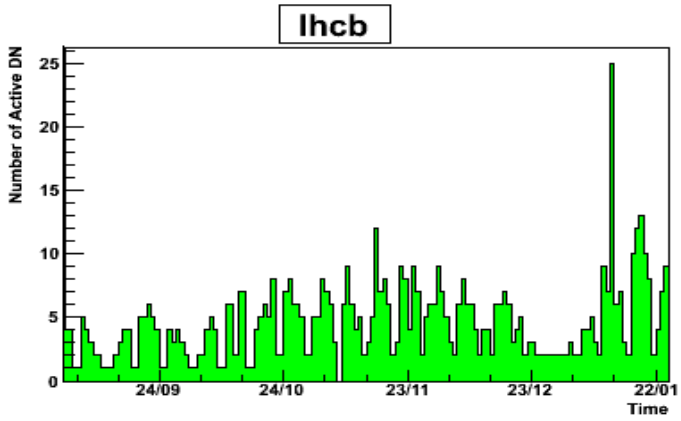

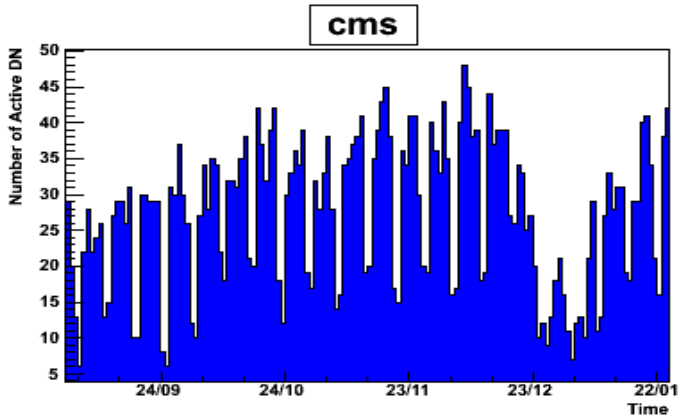

- Distribution of Job Run Time(h) for the LHC VO.

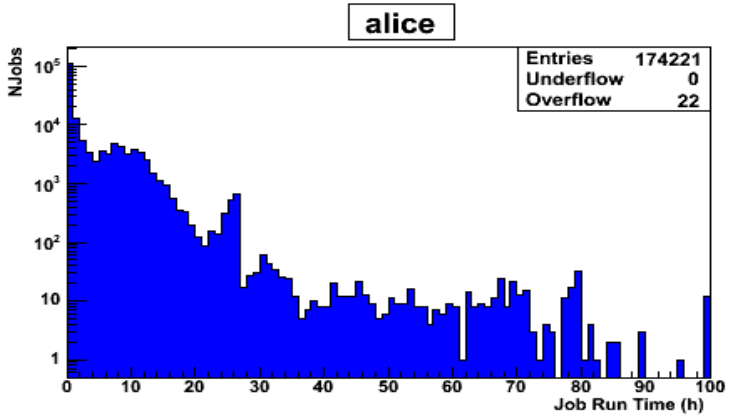

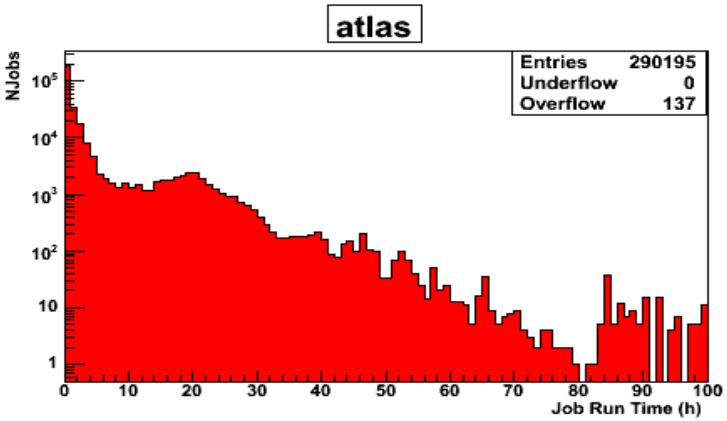

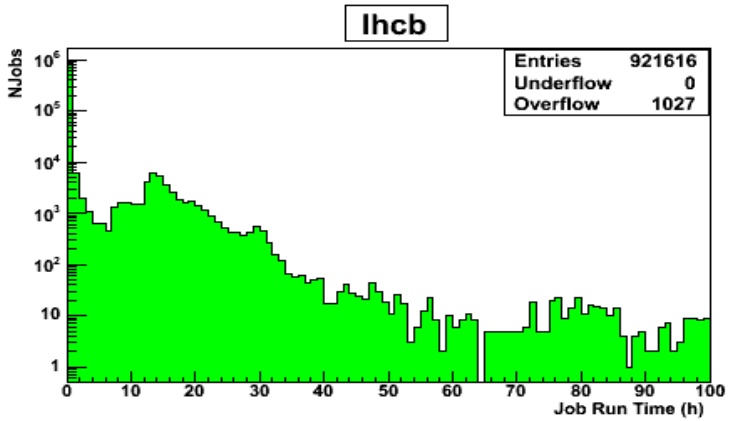

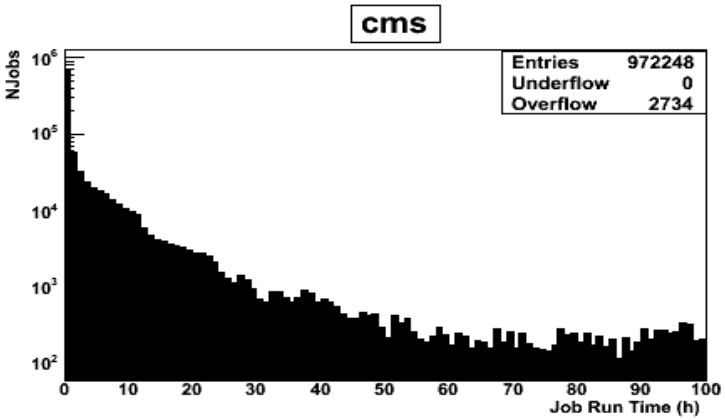

- Distribution of Job Run Time(h) weighted by Job Run Time (h).

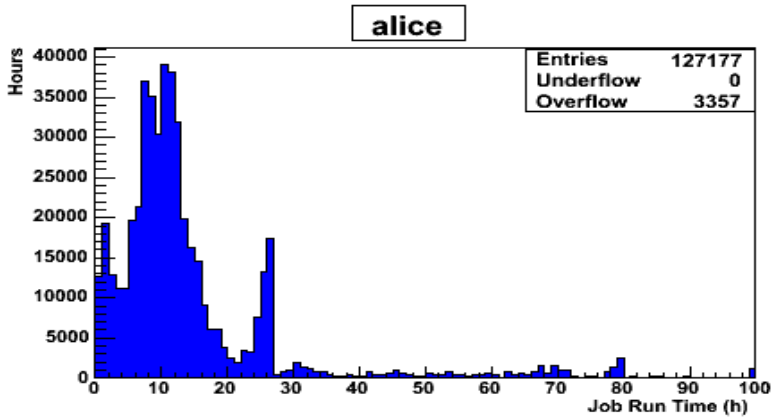

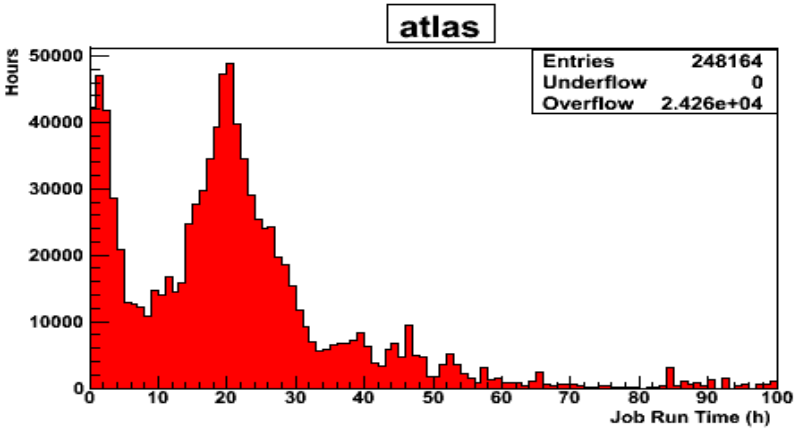

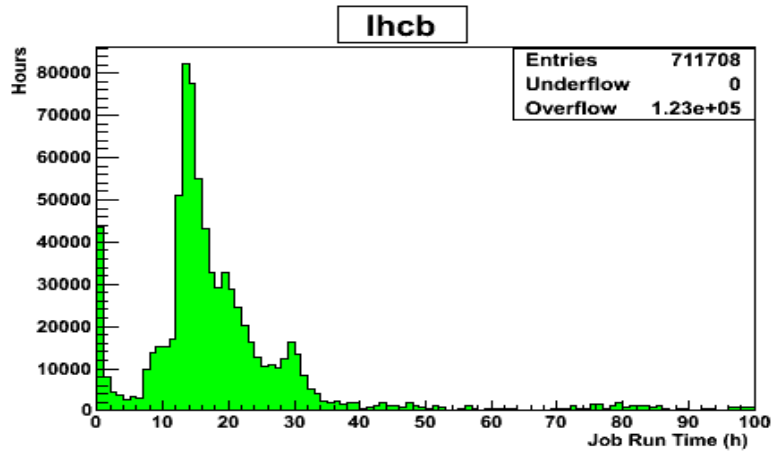

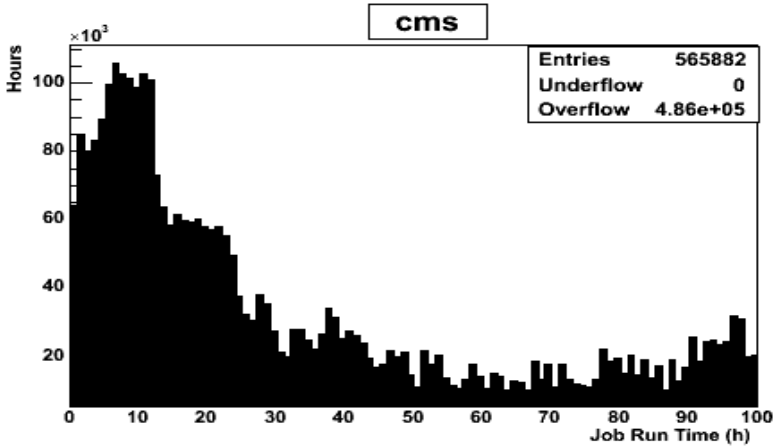

- Distribution of Job Efficiency for each LHC VO
  (efficiency=Time spent running successfully/total time in system)

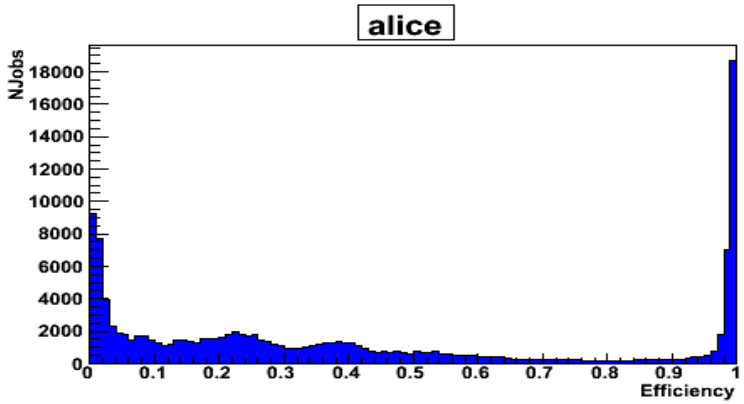

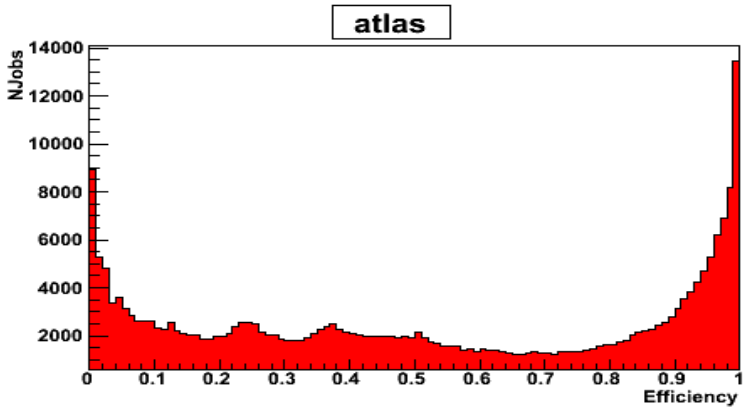

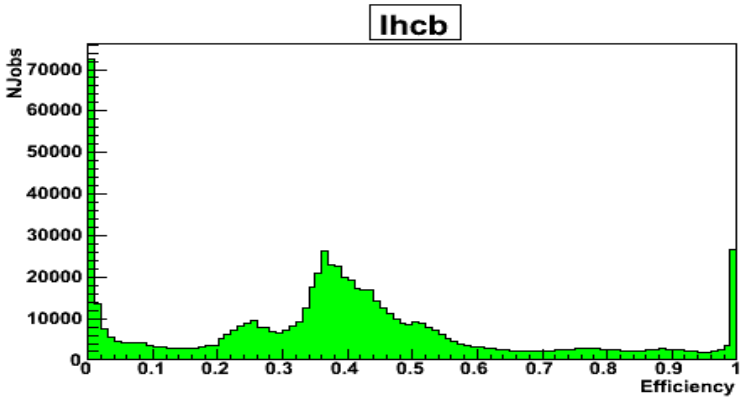

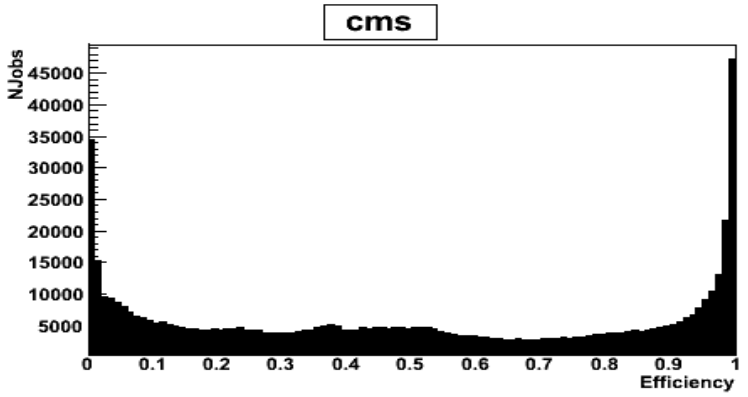

- Job Efficiency versus Job Run Time (h).

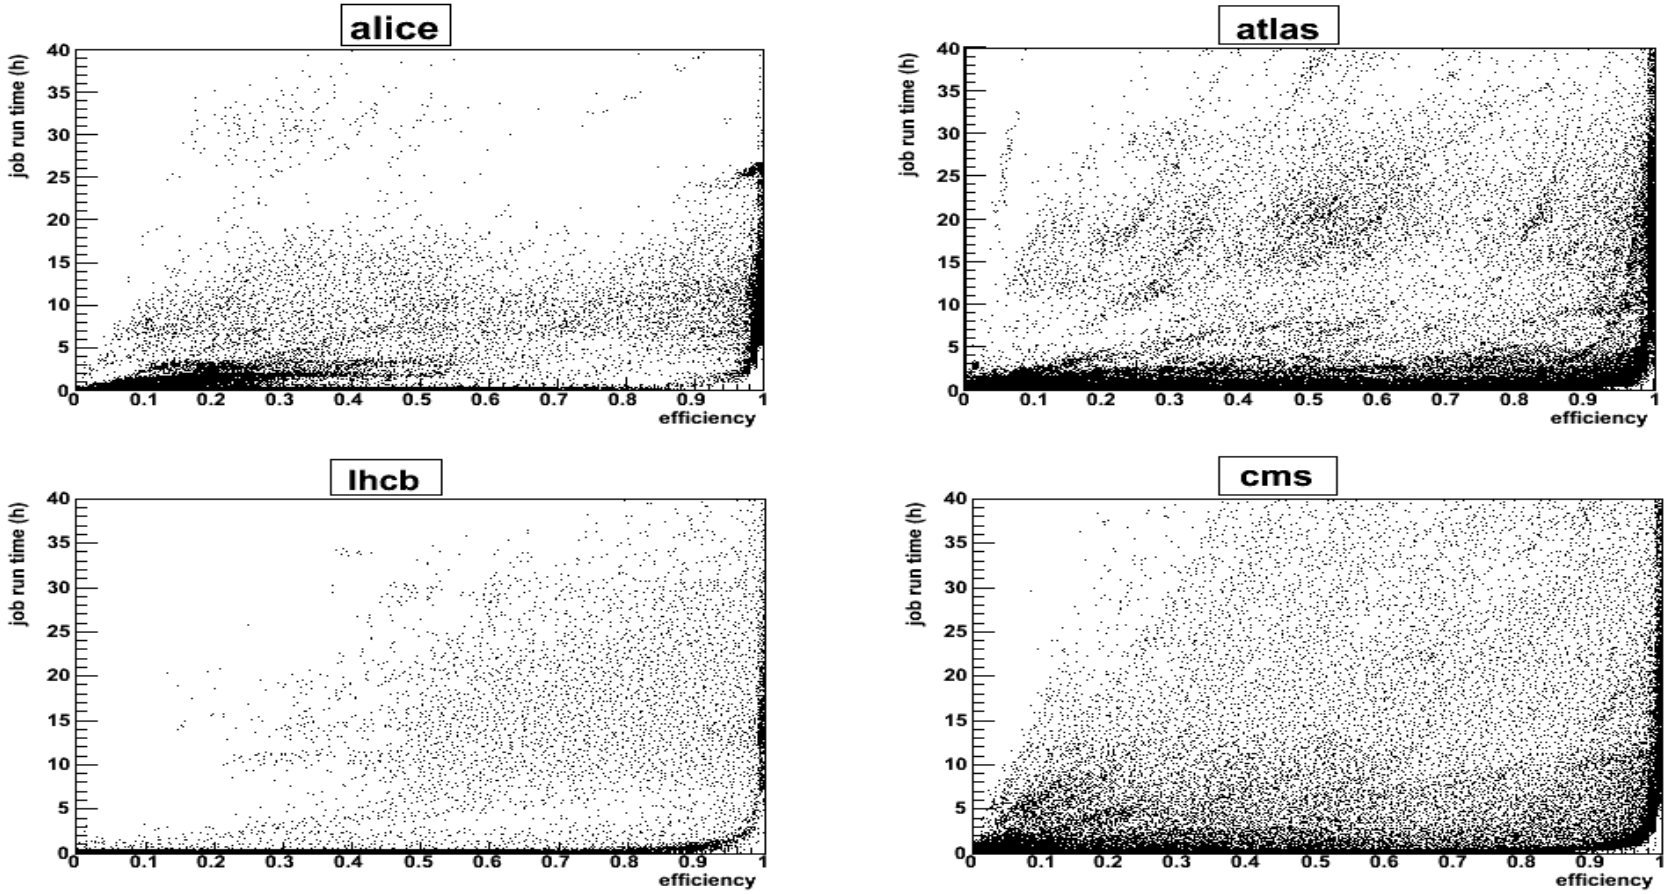

RB Load

- Number of Jobs on a given RB.

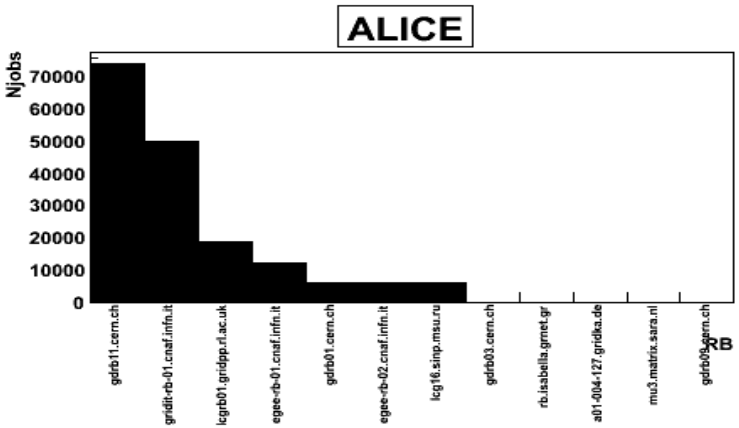

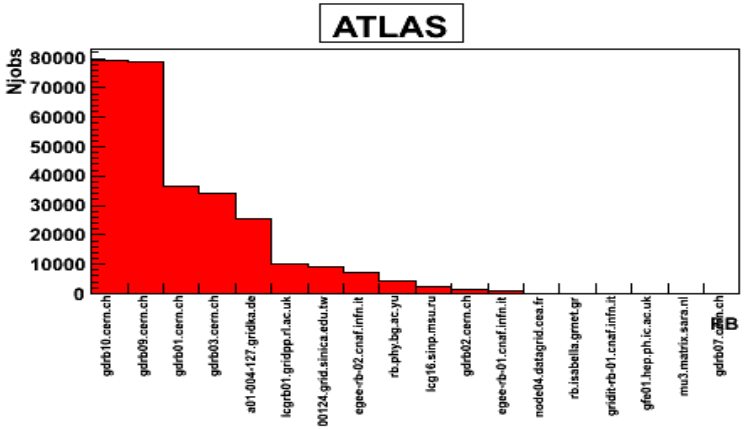

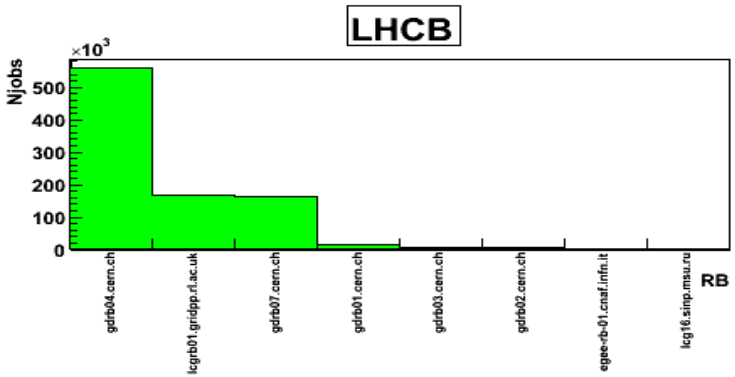

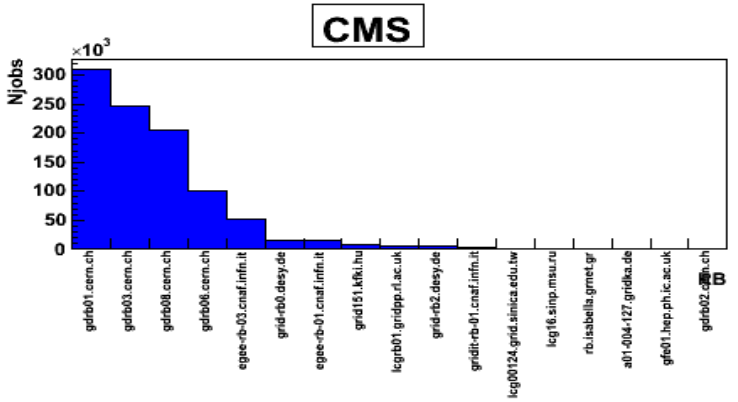

CE Jobs Distribution

Number of Jobs

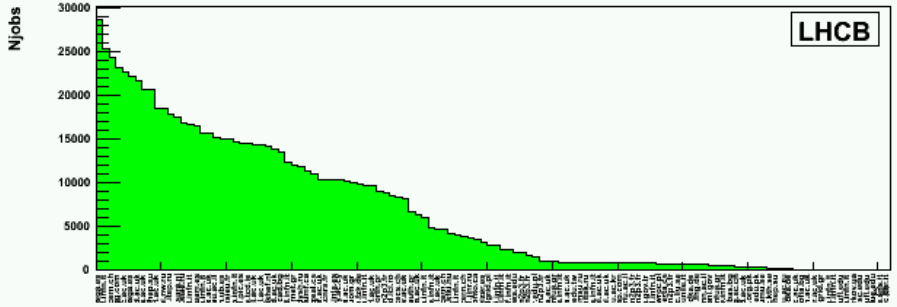

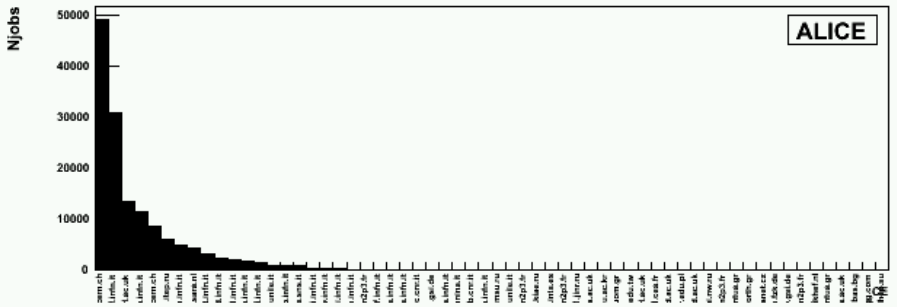

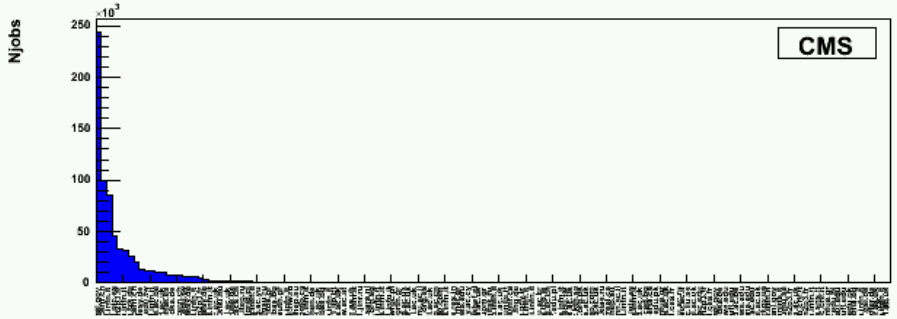

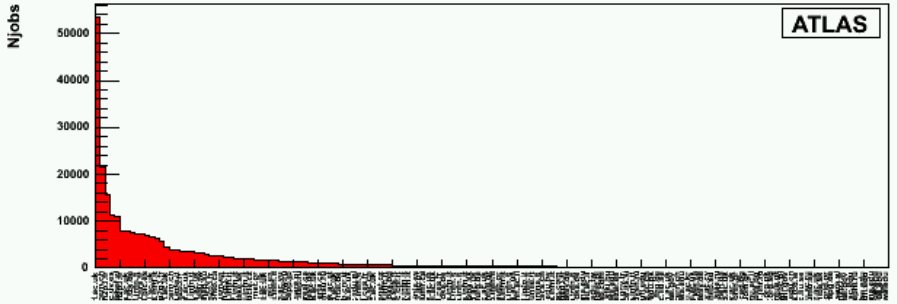

CE Hours distribution

GridPP
UK Computing for Particle Physics

Job Hours

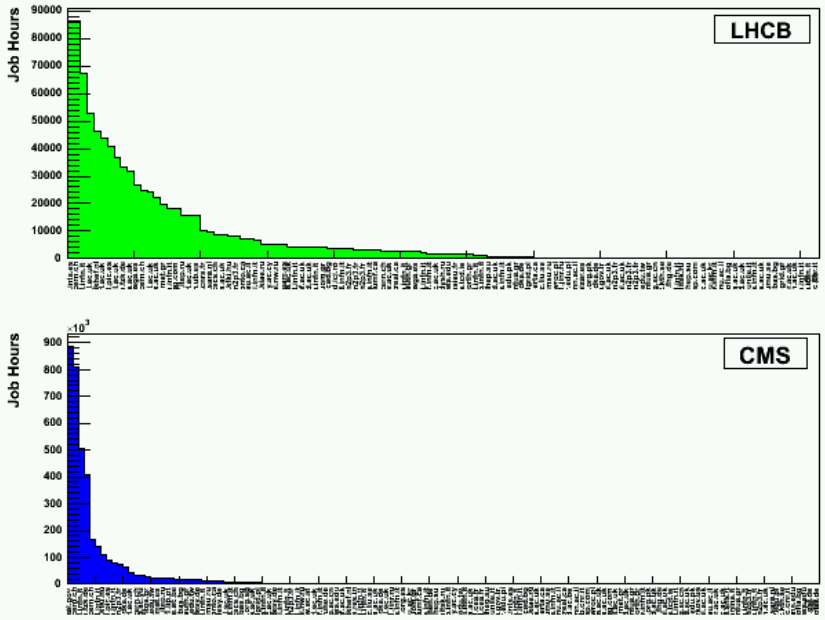

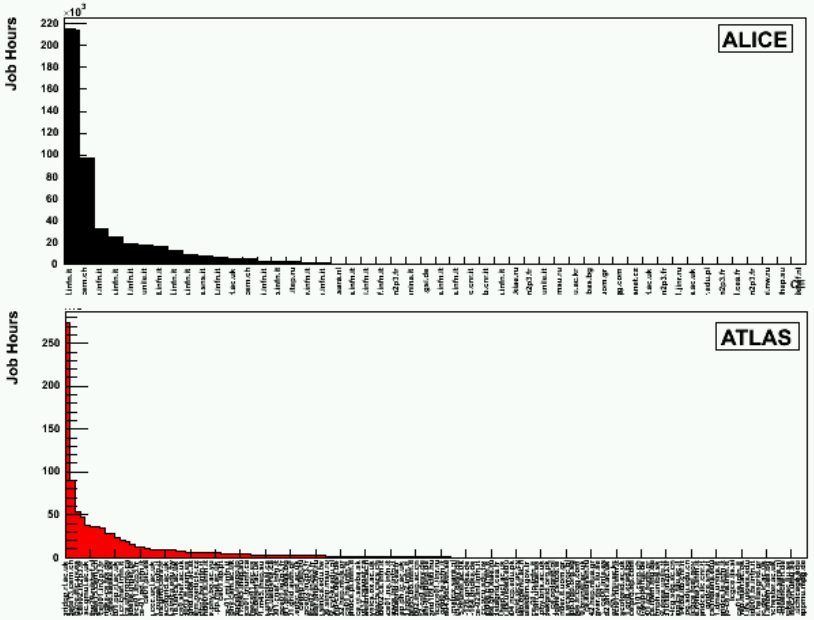

Number of Jobs LCG

Efficiency=N Success/N total

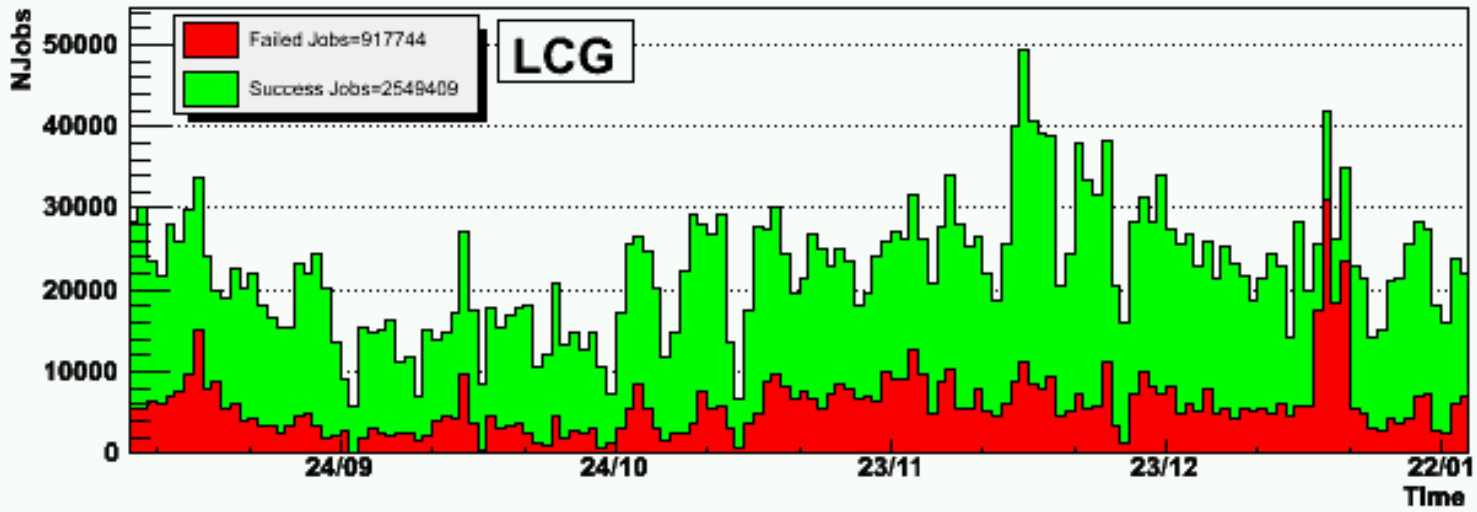

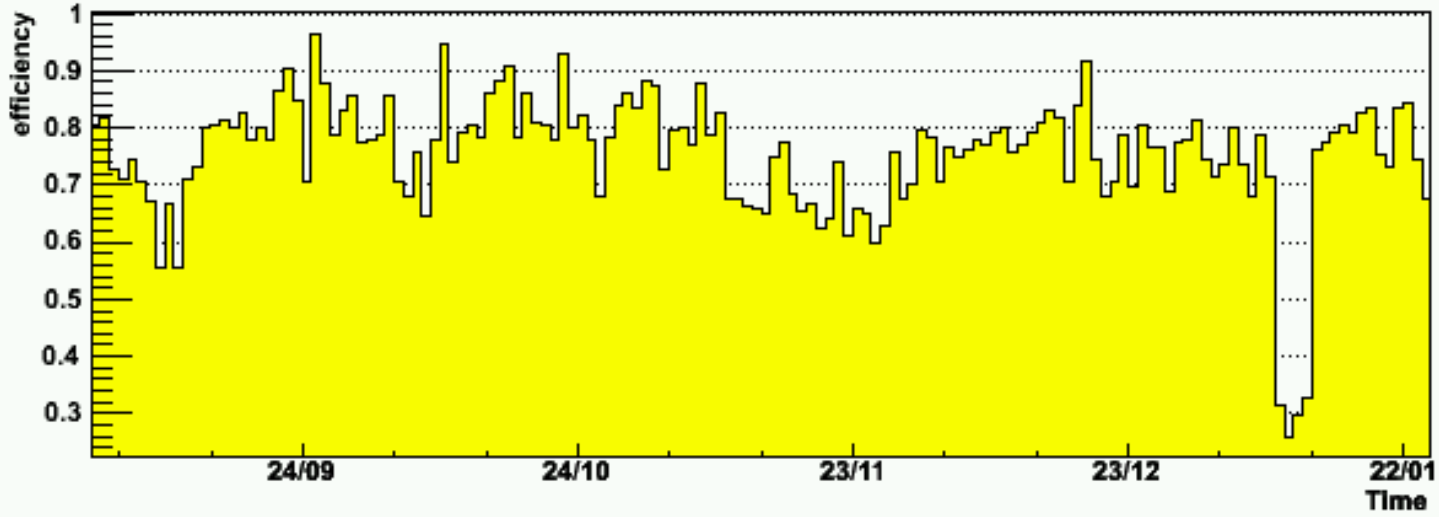

Number of Jobs Alice

Efficiency=N Success/N total

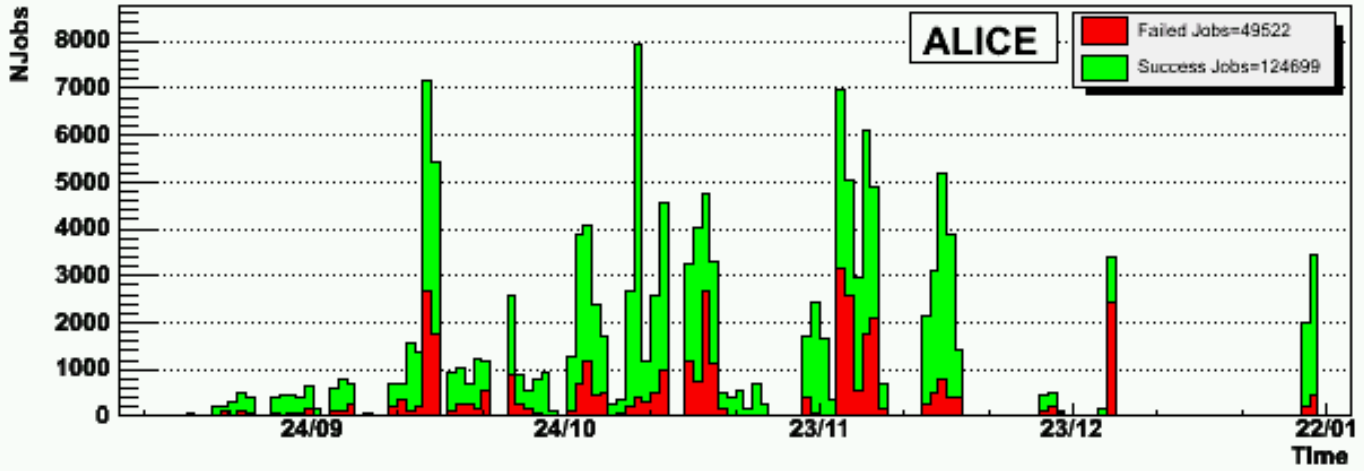

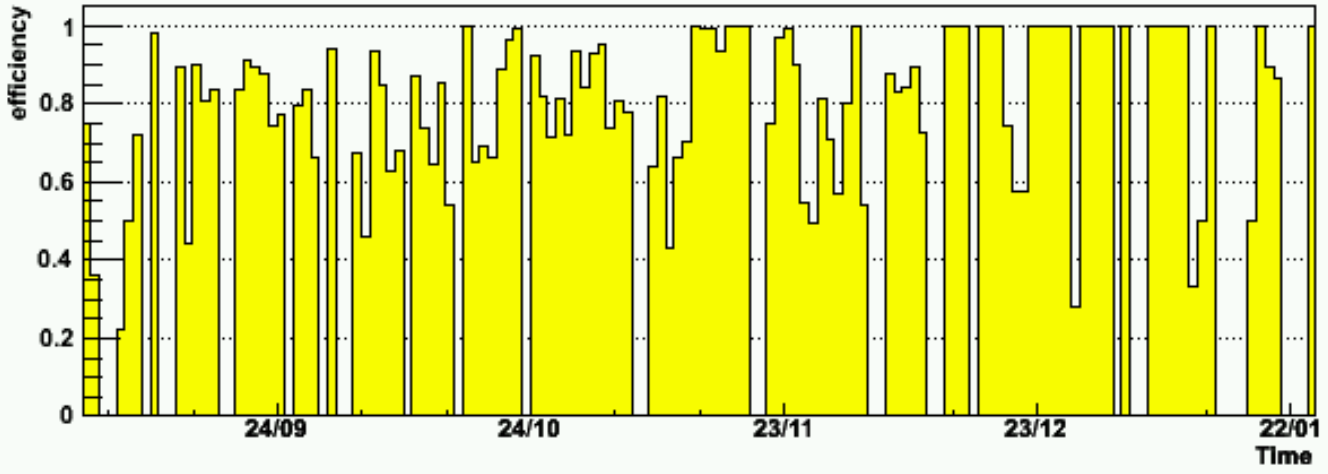

GridPP
UK Computing for Particle Physics

Number of Jobs Atlas

Efficiency=N Success/N total

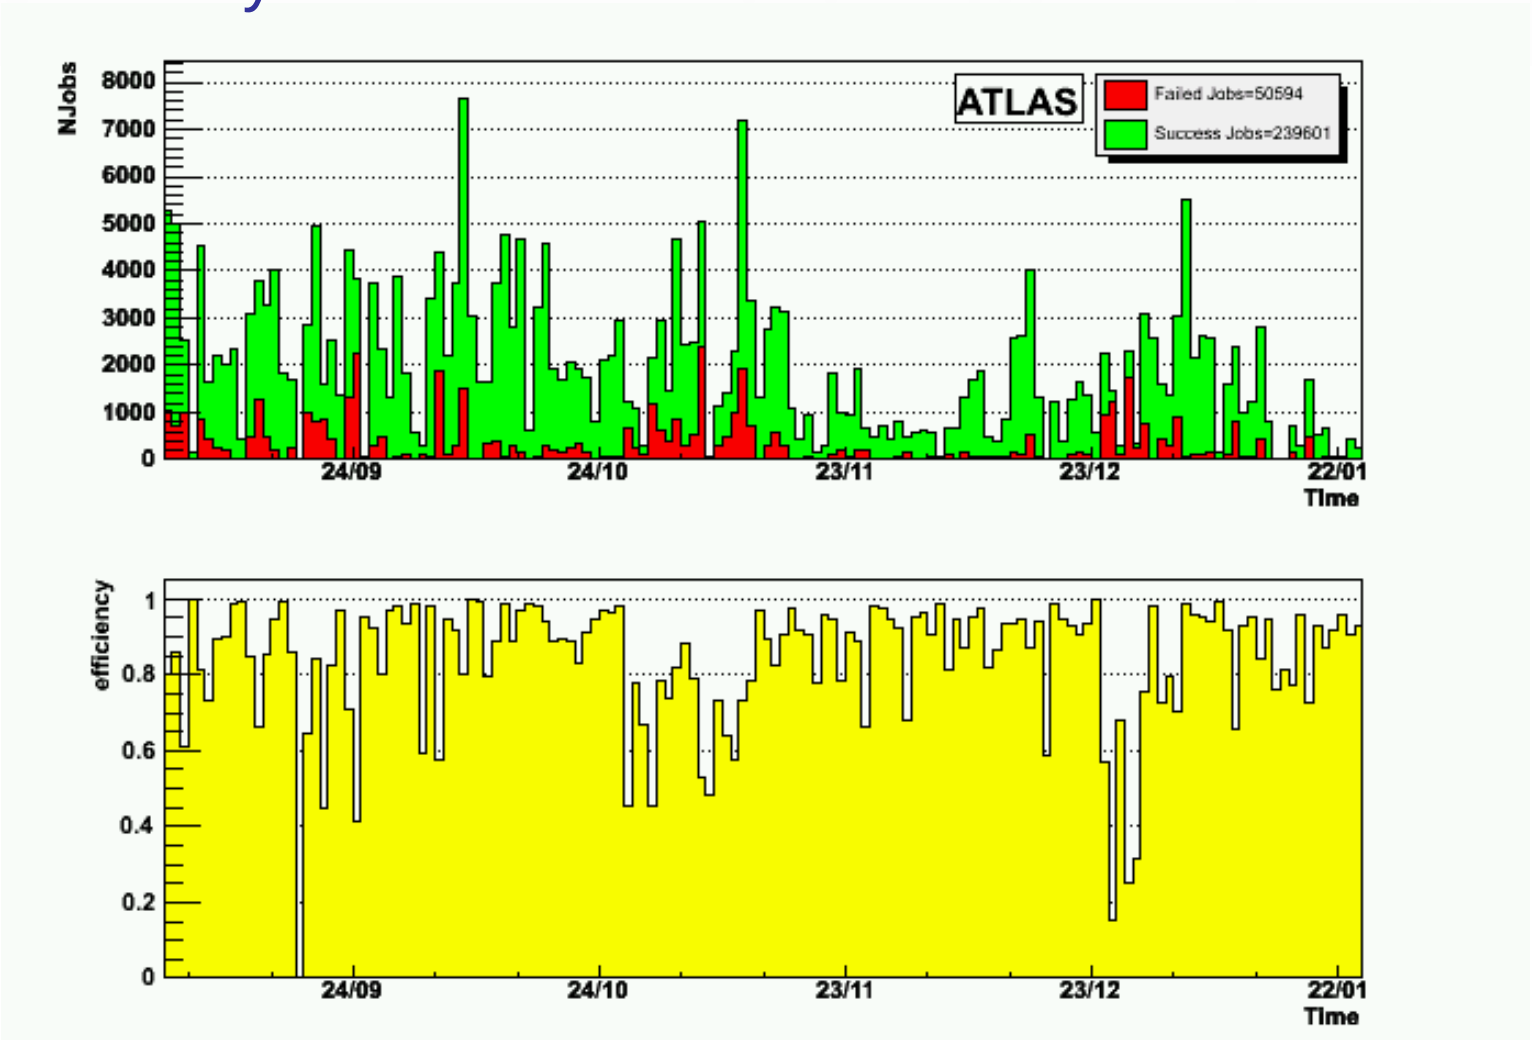

Number of Jobs CMS

Efficiency=N Success/N total

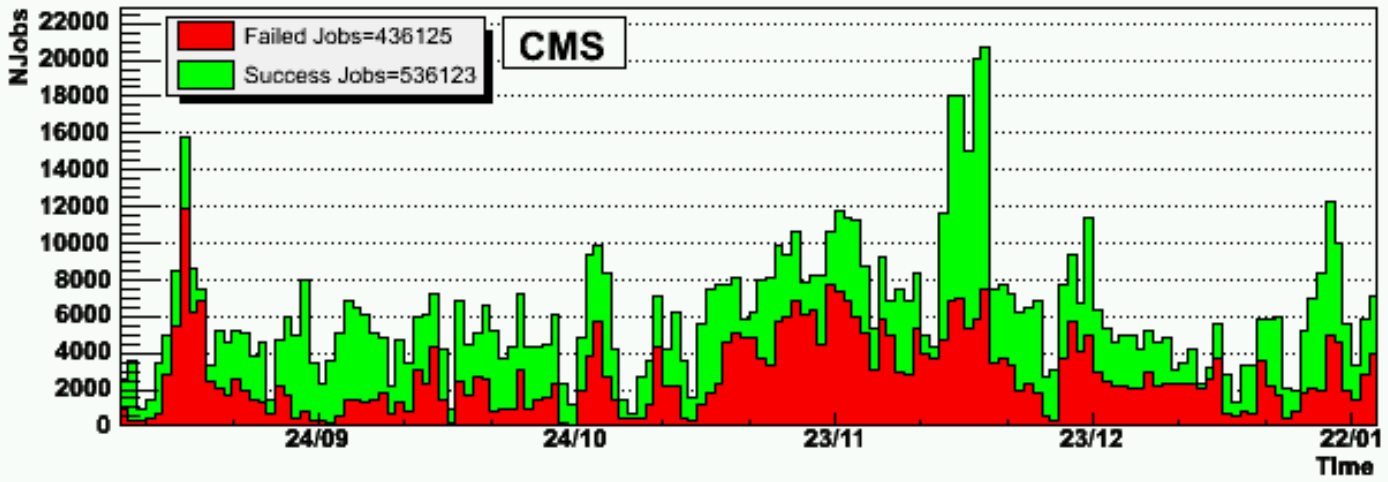

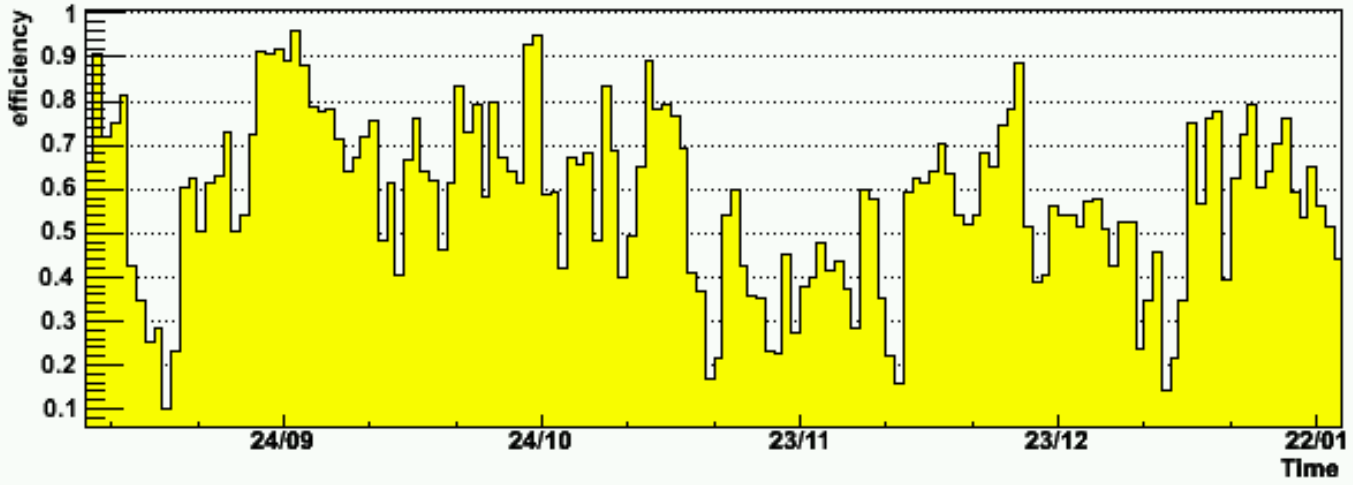

Number of Jobs LHCB

Efficiency=N Success/N total

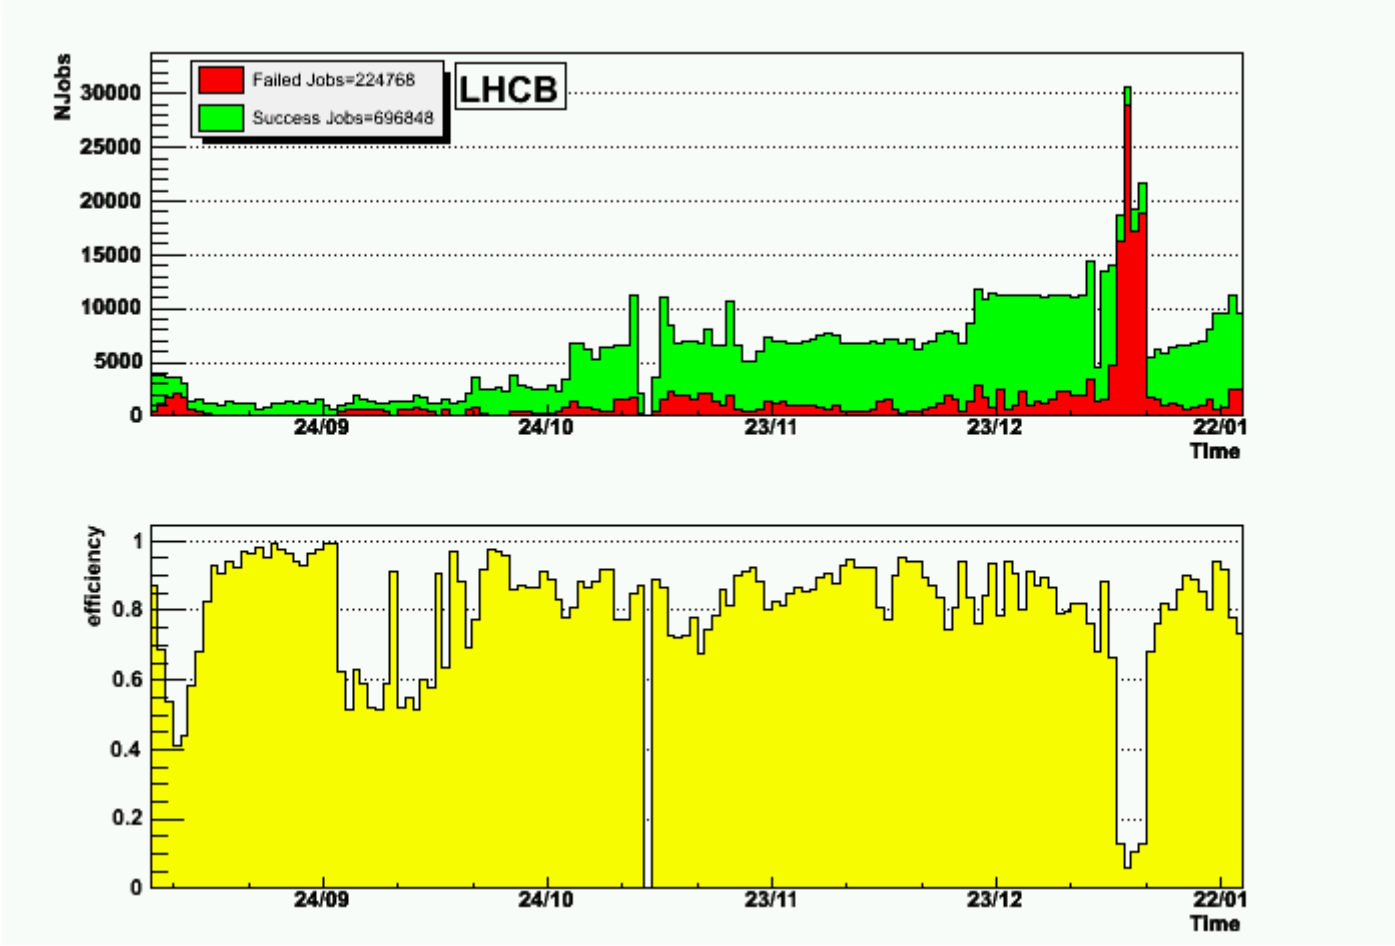

GridPP
UK Computing for Particle Physics

Grid Load

- Number of Job Hours submitted at a given time

ALICE

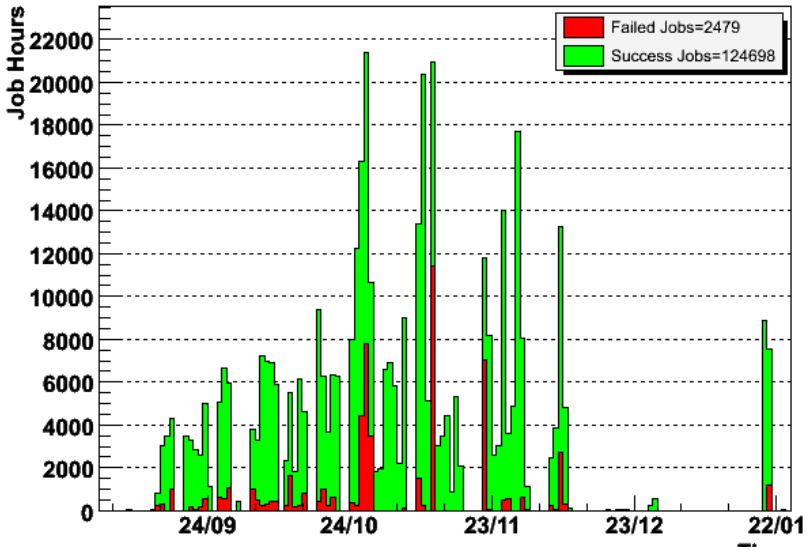

ATLAS

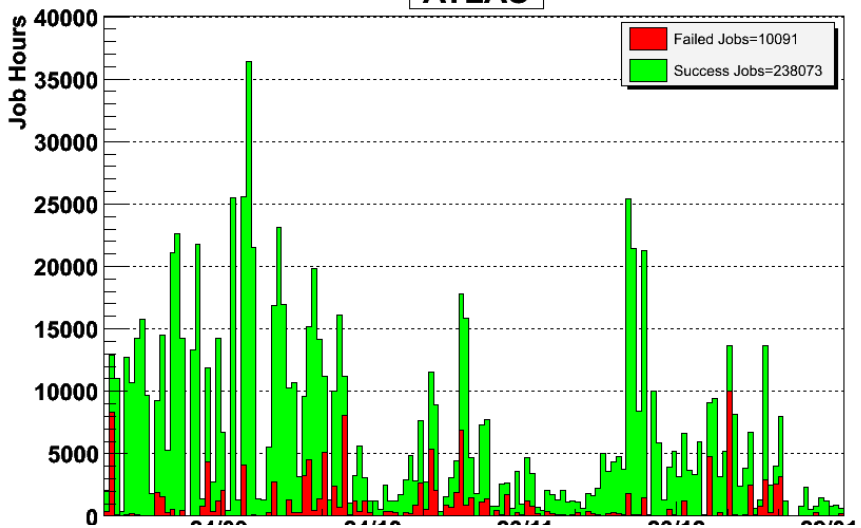

LHCB

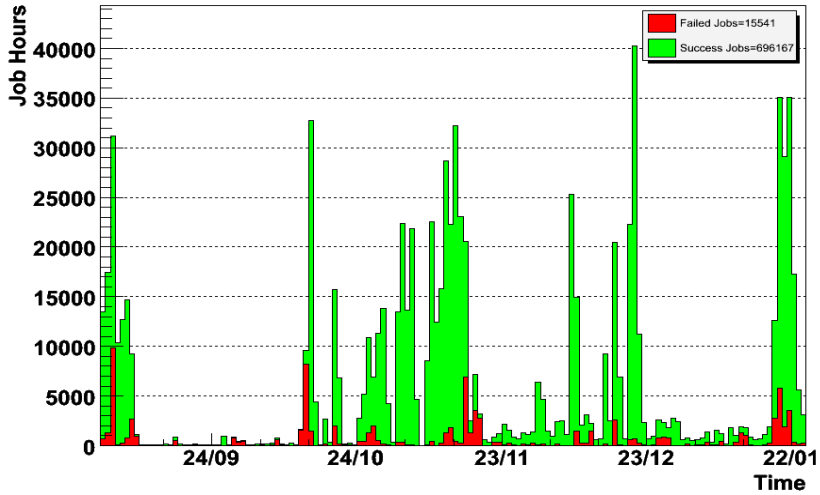

CMS

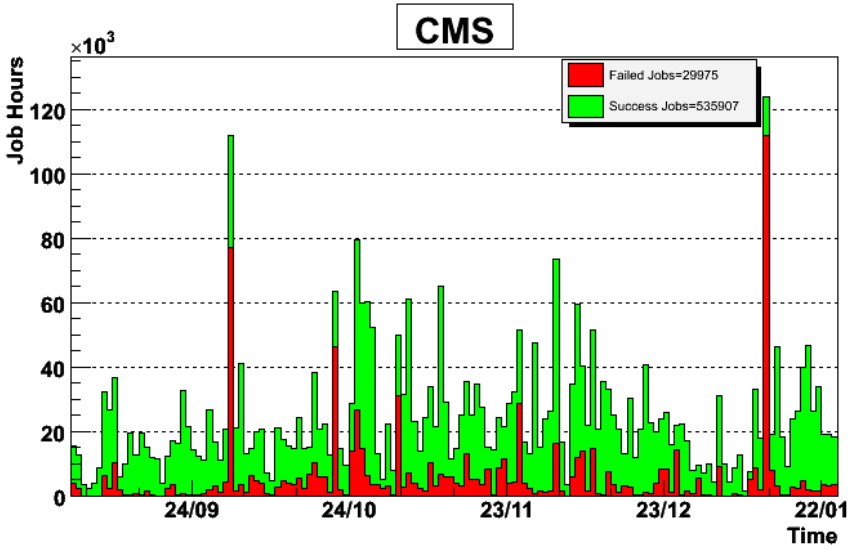

Efficiency CMS and Atlas

- Efficiency=Total Succ Hours/Total Hours

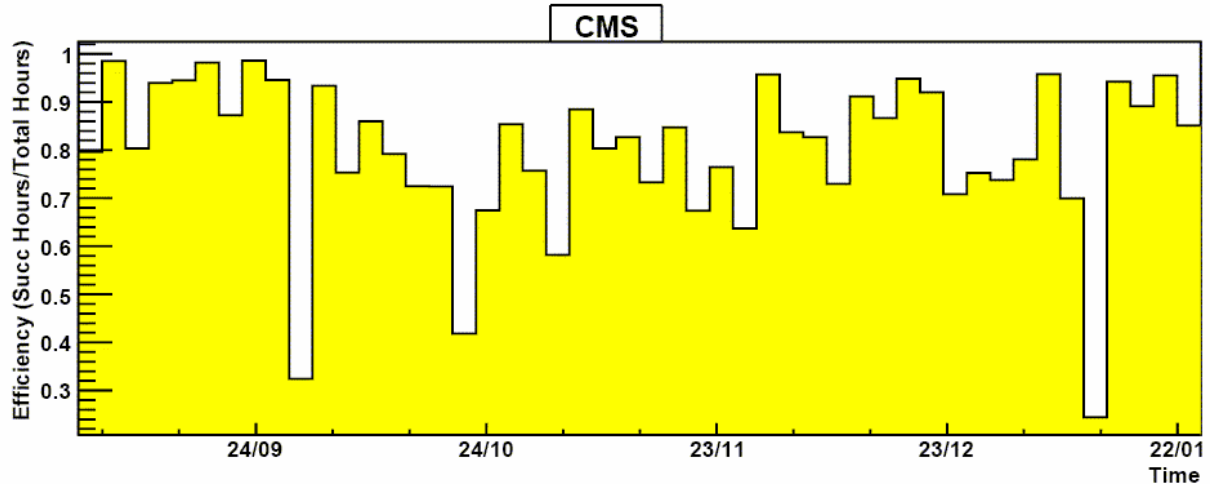

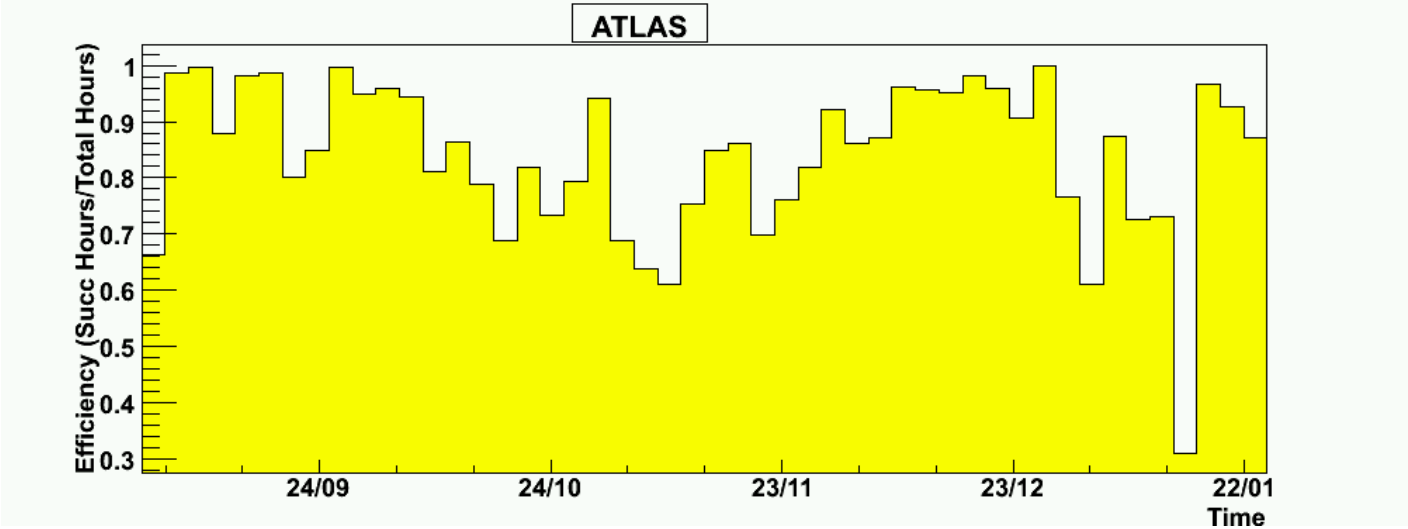

Efficiency LHCB & Alice

- Efficiency=Total Succ Hours/Total Hours

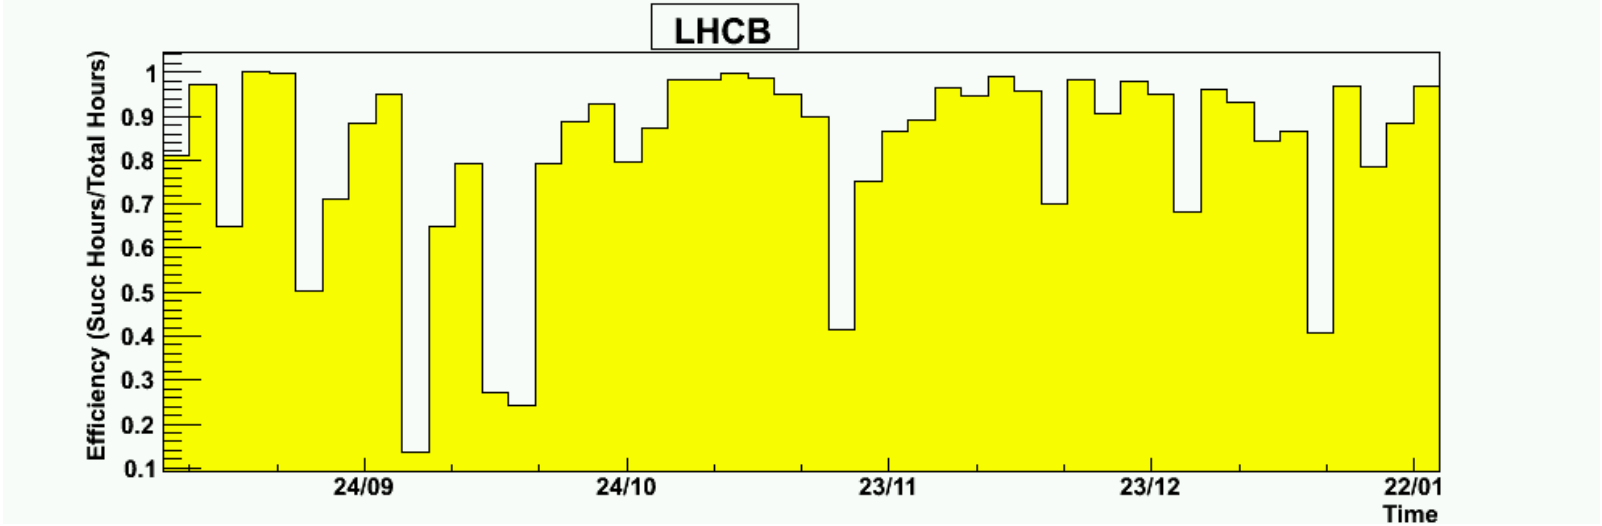

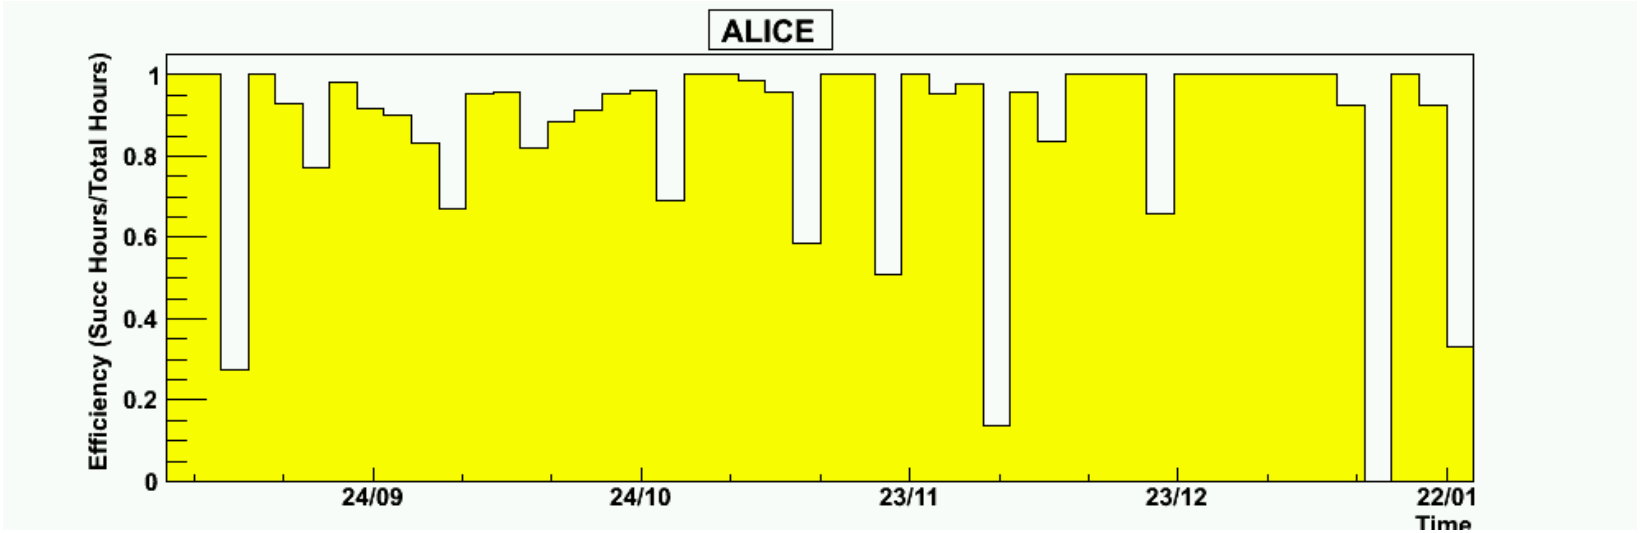

Grid Load

- Number of jobs in the system at a given time.

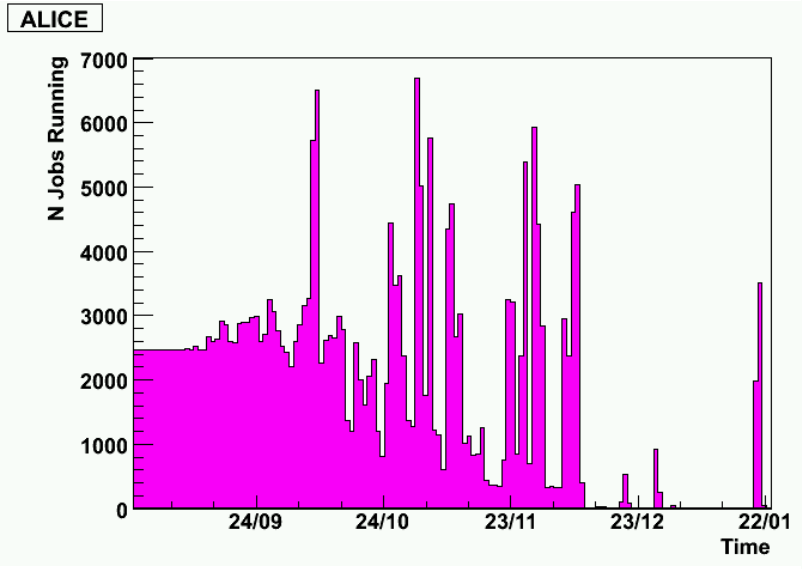

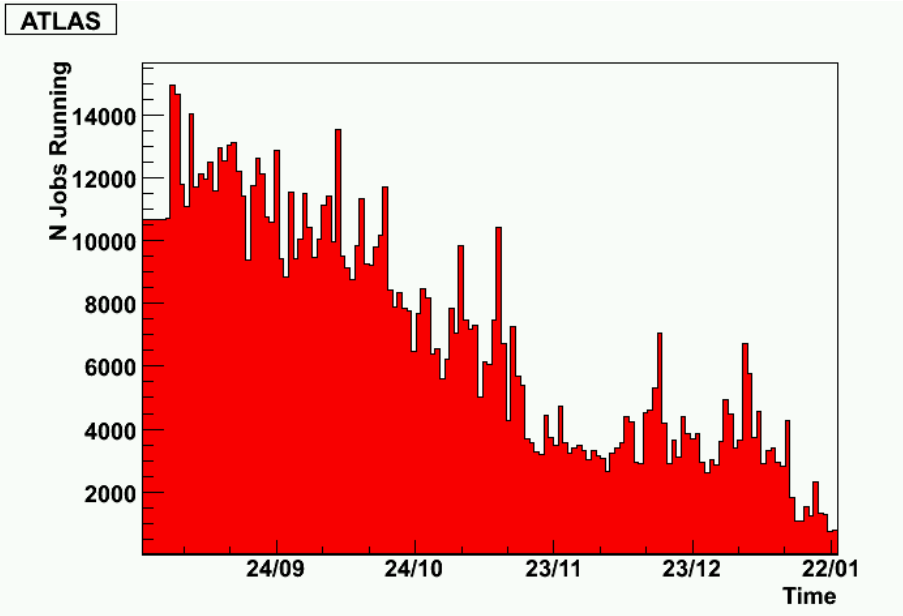

Grid Load

- Number of jobs in the system at a given time.

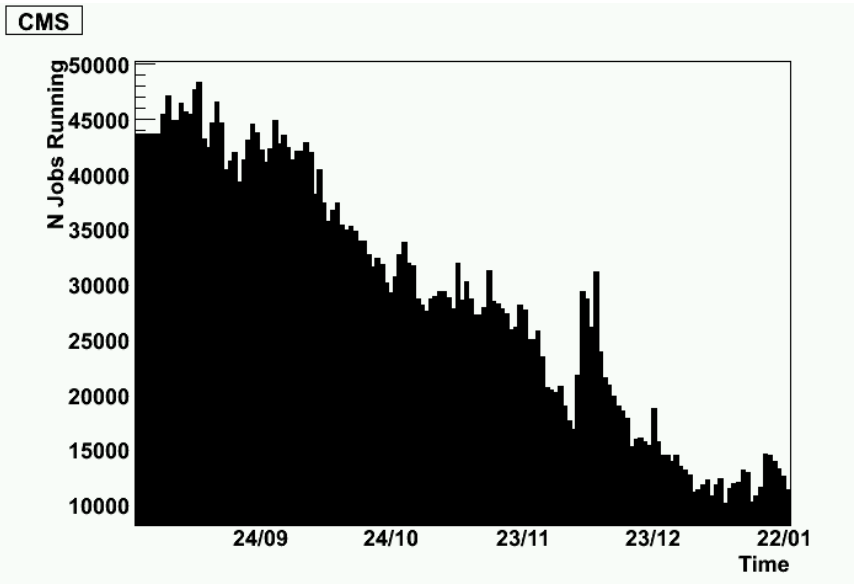

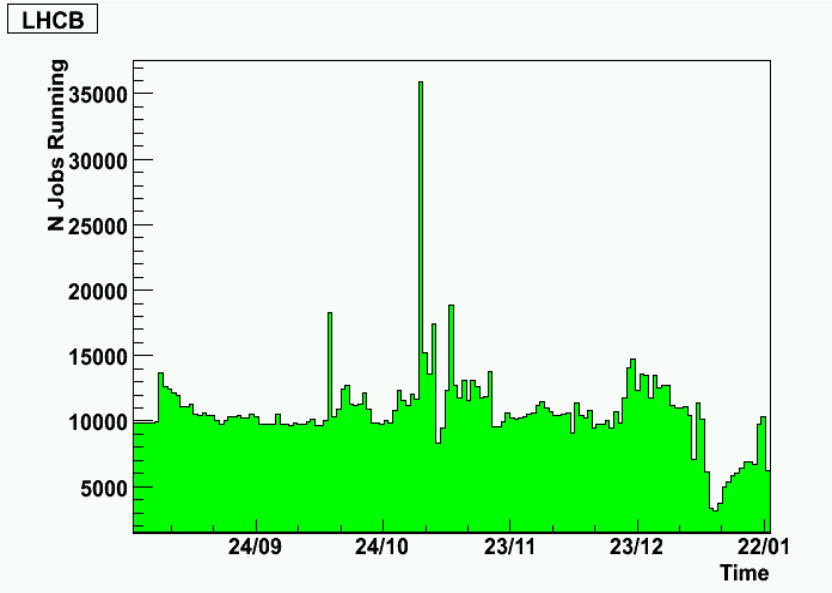

Grid Load

- Number of jobs in the system at a given time.

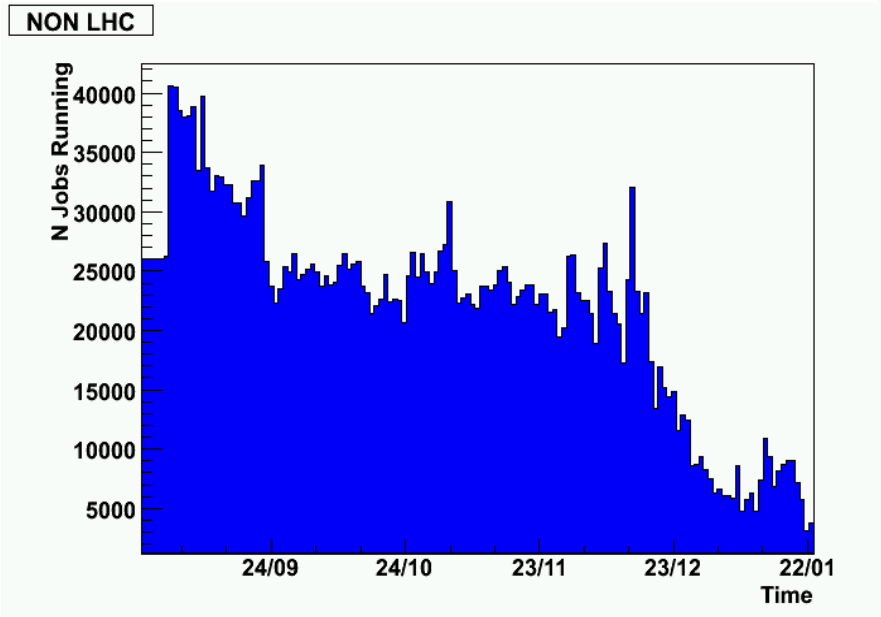

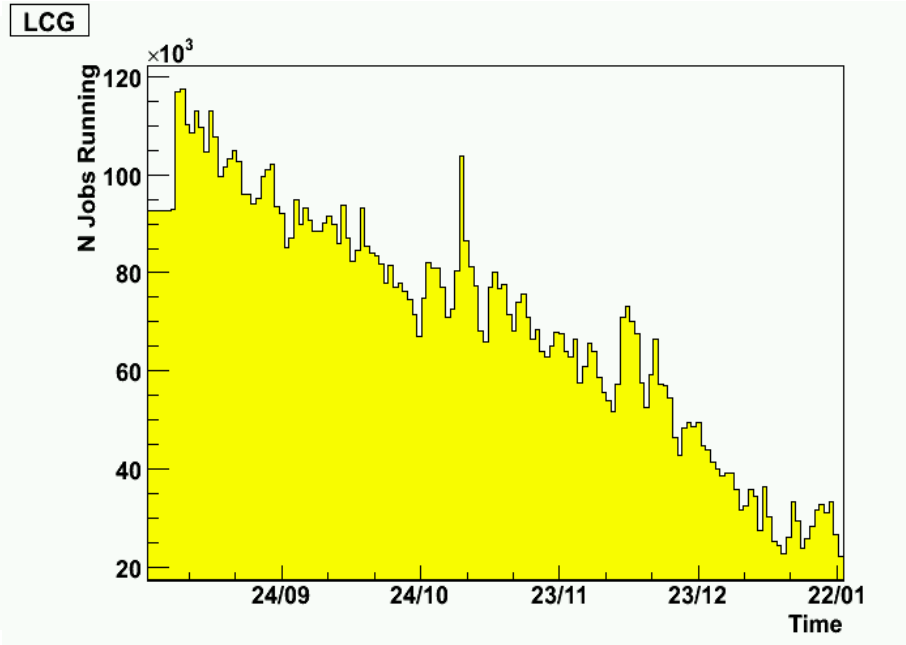

RB Load

- Job scheduling (Match Time) versus load (mean number of jobs/sec during the matching)

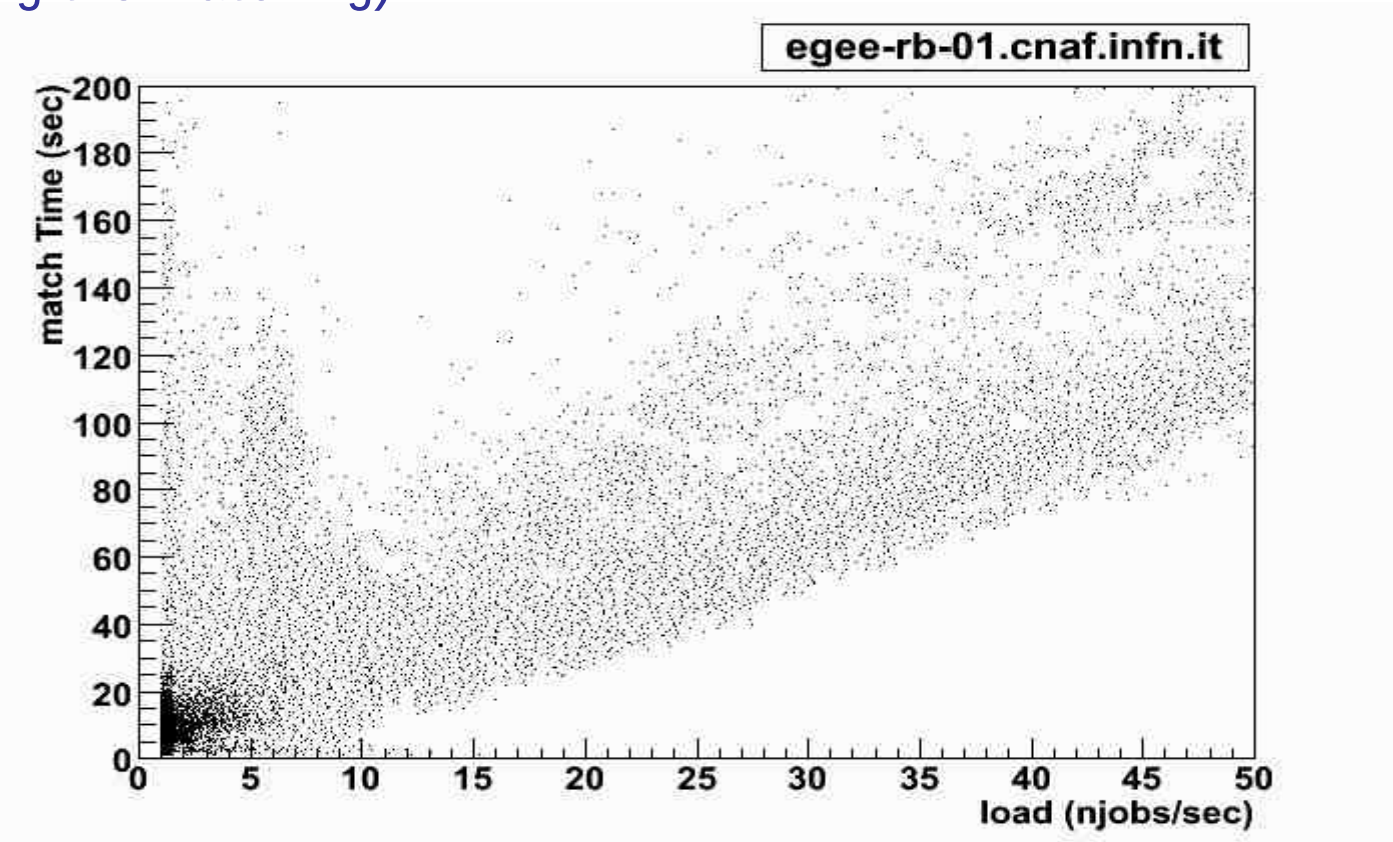

RB Load

- Job scheduling (Match Time) versus load (mean number of jobs/sec during the matching)

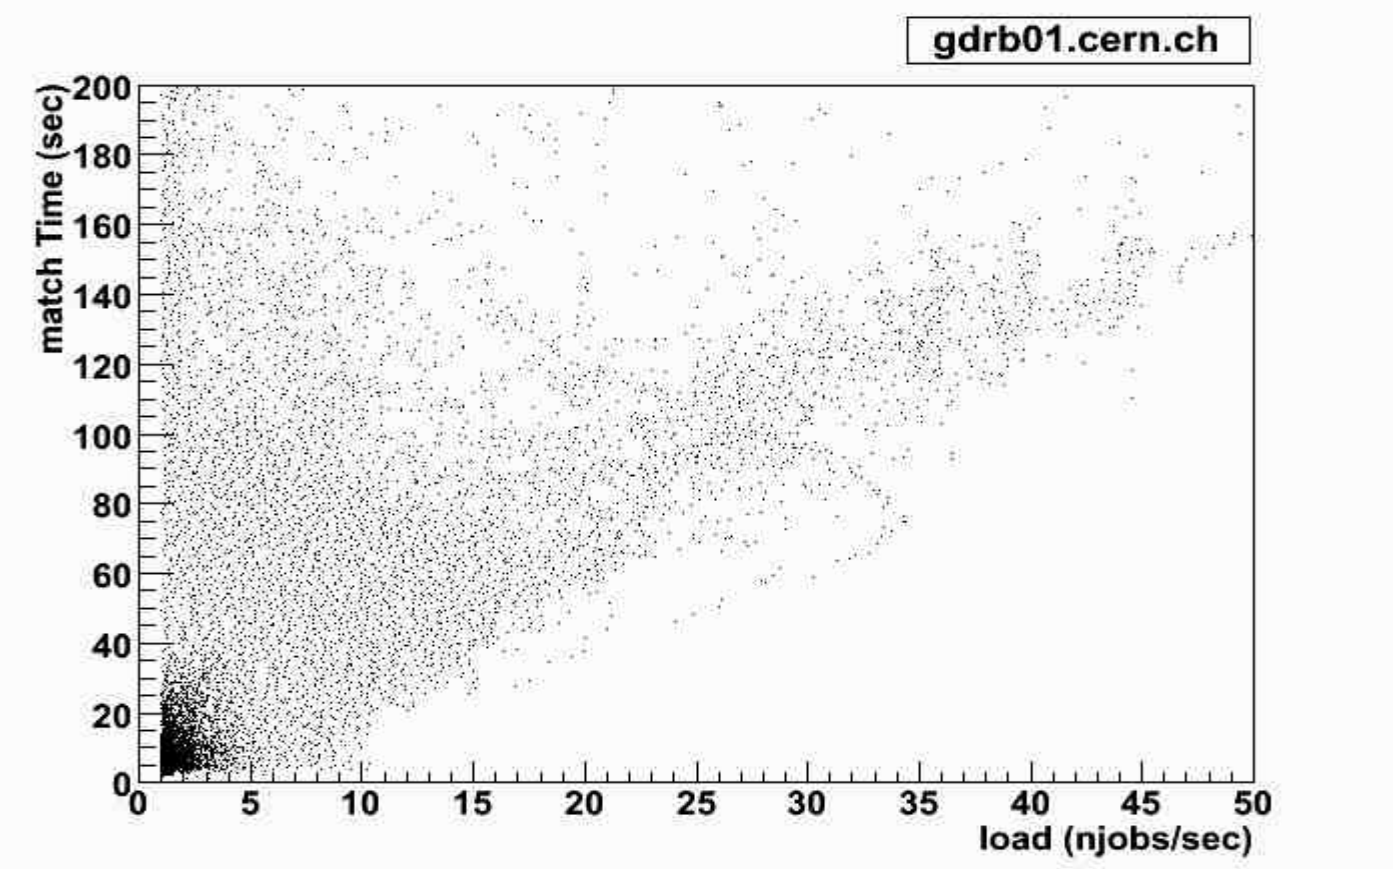

RB Load

- Job scheduling (Match Time) versus load (mean number of jobs/sec during the matching)
- RB.(gdrb04)

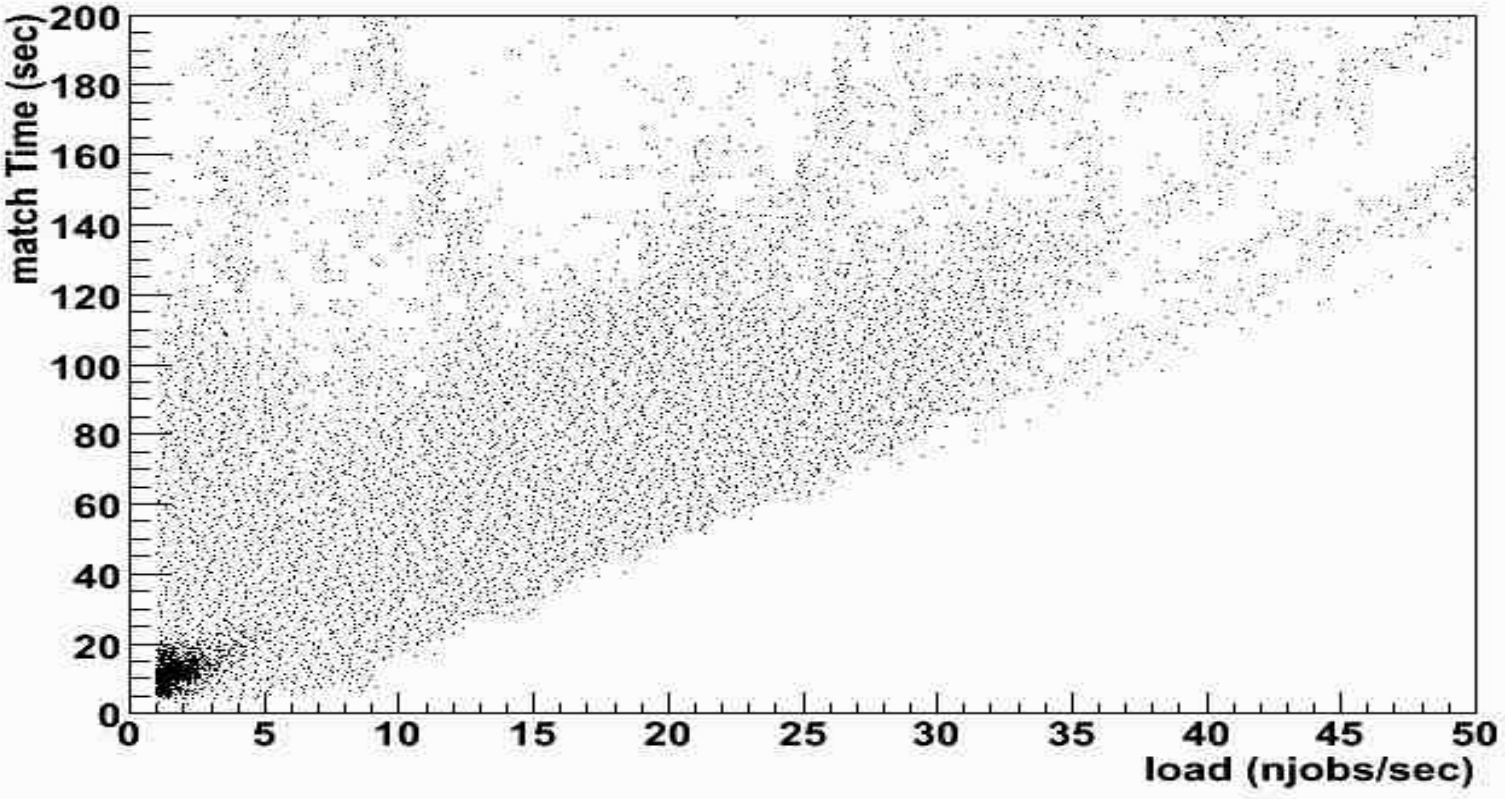

Conclusions

- We have started to analyse the distribution of jobs submitted to the LCG

- Distinct usage patterns are beginning to emerge for each VO

- These uasge patterns have different efficiencies

- There are many more plots that I could have shown and there is a lot more work to do to try to understand what we see

References

- [1] GridPP-UK Computing for Particle Physics: http://www.gridpp.ac.uk/
- [2] Crosby P, Colling D, Waters D, **Efficiency of resource brokering in grids for high-energy physics computing**, IEEE Transactions on Nuclear Science, 2004, vol: 51, Pages: 884 - 891, ISSN: 0018-9499

Backup slides

Grid Load

- Number of Job Hours submitted at a given time

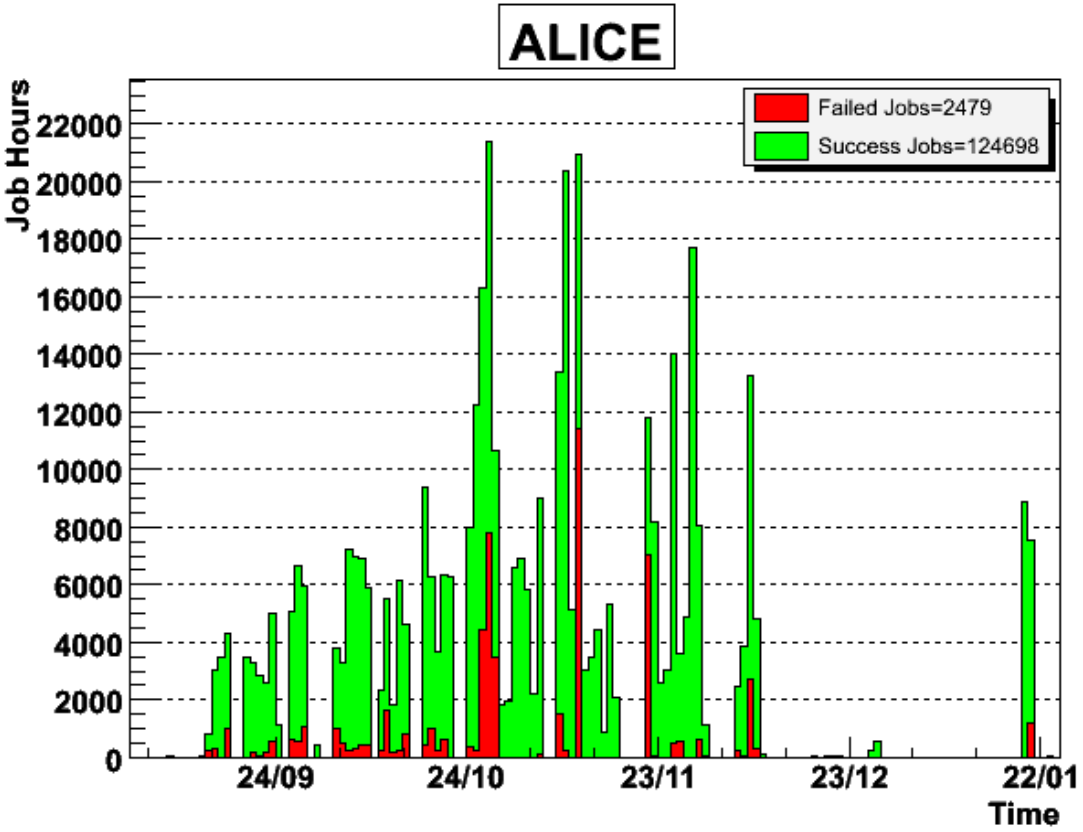

Grid Load

- Number of Job Hours submitted at a given time

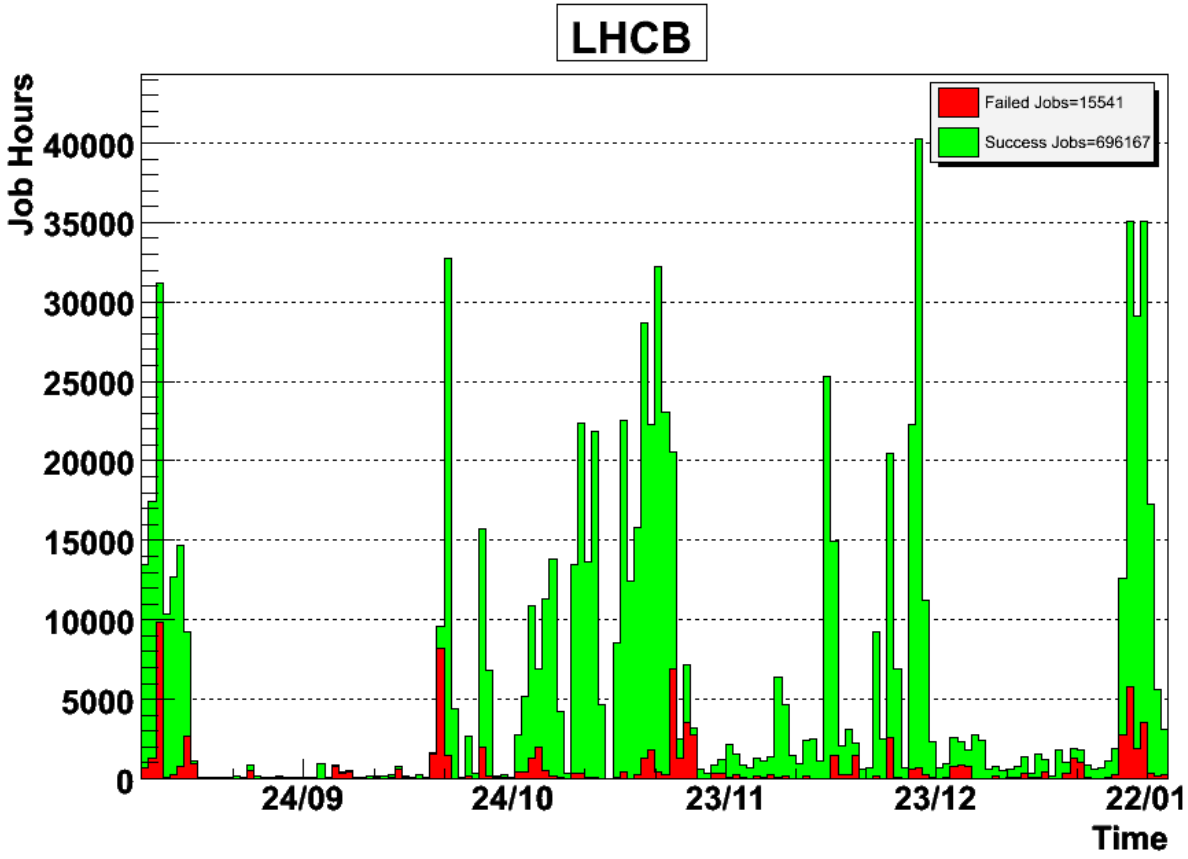

Grid Load

- Number of Job hours submitted at a given time

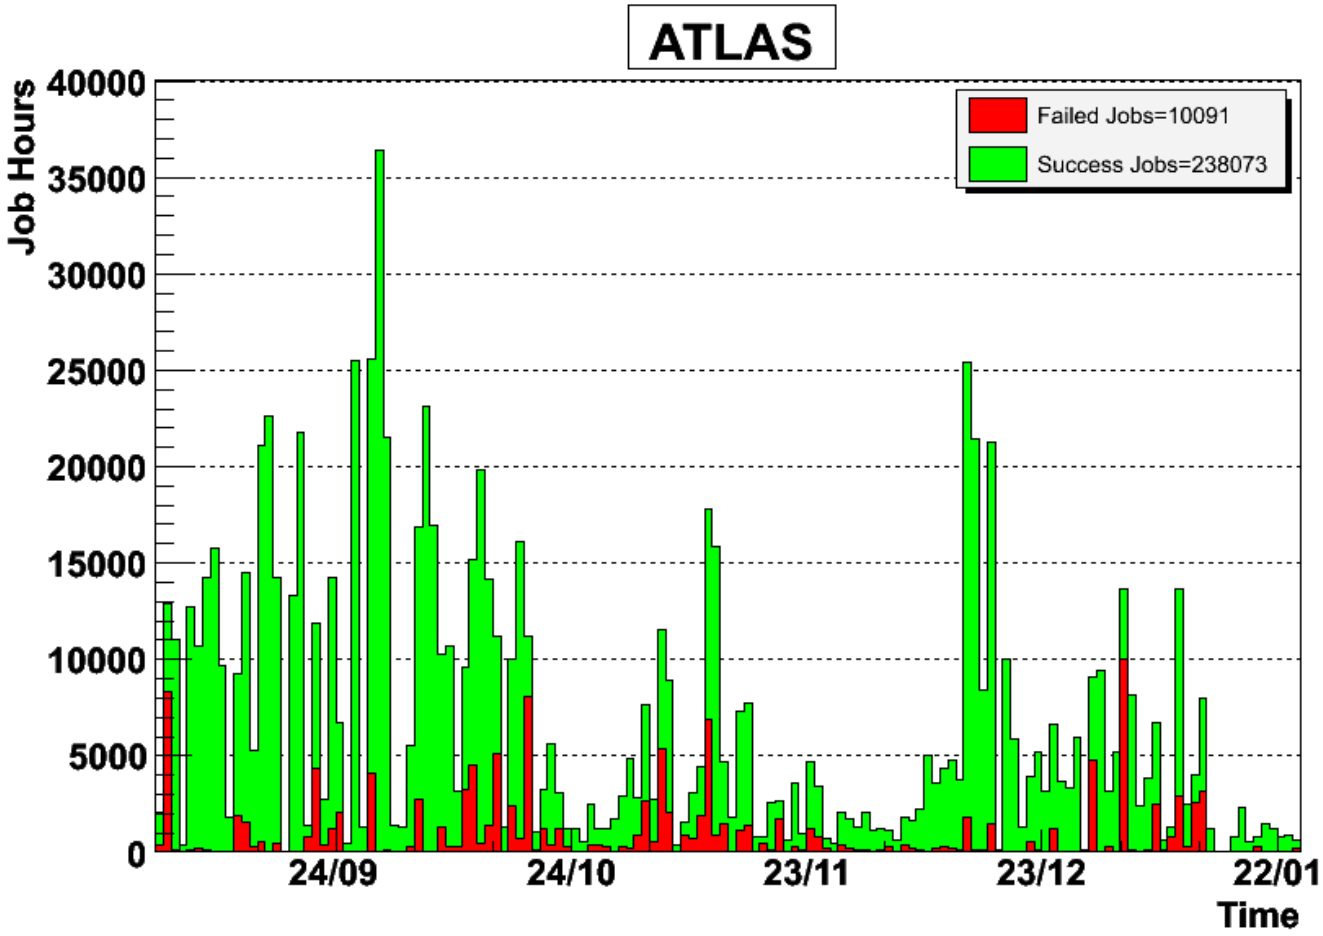

Grid Load

- Number of Job hours submitted at a given time