Assessment of ZEUS Offline Computing Beyond the End of Data-Taking

Rainer Mankel

DESY

4-October-2005

1. Reasoning

The ZEUS Offline Computing strategy has been defined in a previous document for the years 2004-07 [1]. This document is a first attempt to assess the model for offline computing in the years beyond the end of data-taking, 2008-11.

The recent survey of ZEUS manpower estimates has also addressed the number of persons doing analysis per year until 2010. According to this query, the number of analysis users will have decreased relative to the present level by about 25% in 2007, by 50% in 2008 and by 70% in 2010. The analysis volume however, which determines the workload for analysis computing, is governed not only by the number of users but also by the amount of data under analysis. The expected development of the HERA-II real data volume is displayed in fig. 1, it is expected to increase by 600% by 2007 and by 800% by 2008, to a final number of about 27 TB at MDST level. For comparison, the total amount of HERA-I data comprises only around 5 TB.

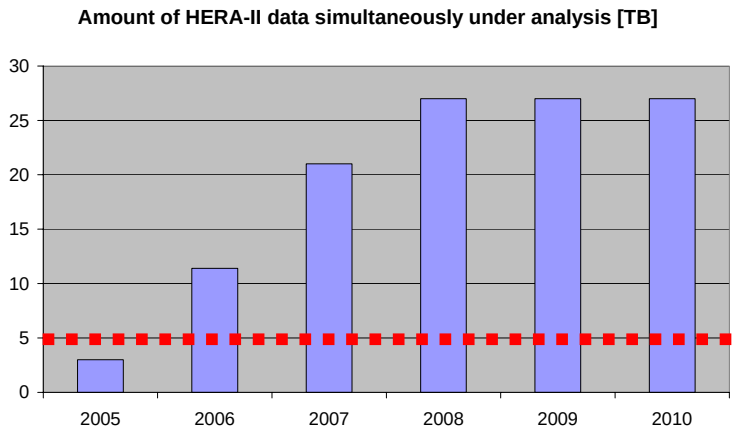

Fig.1: Cumulative amount of real-data under analysis in given year

The resulting analysis volume of real data (in TB * user) is displayed in fig. 2. After a steep increase, it reaches a flat peak through 2007/08 whereupon it decreases, albeit with a substantial residual in 2010. It should be noted that the number of users determined in the manpower questionnaire does not assess the intensity or fraction of time that each user is investing on average. As time moves on, people may become more and more occupied with other projects. On the other hand, being relieved of detector chores may give many people the leeway to become more deeply involved in analysis than before, and quantum leaps in the understanding of the new data, perhaps manifesting themselves in a grand reprocessing of the HERA-II events, may kindle additional enthusiasm that would increase the turnaround of analysis.

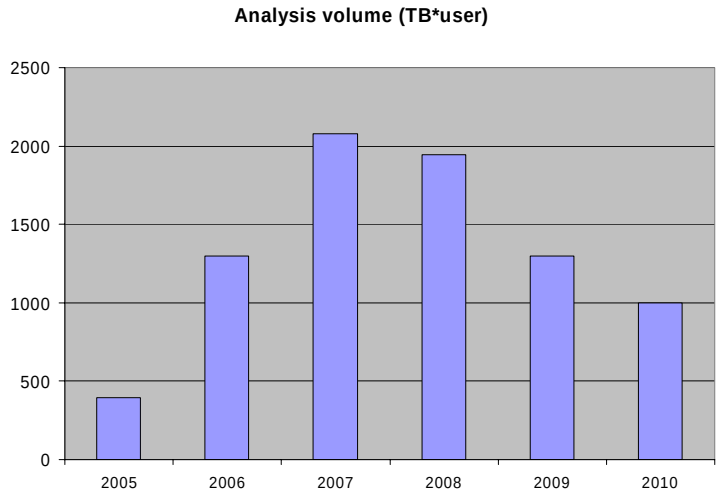

Fig. 2: Analysis volume of real data

Besides real-data, also the analysis of Monte Carlo events must be considered. According to experience, production & thereby analysis of Monte Carlo events is performed with a certain time offset with respect to the real data. Assuming a relative aggressive schedule, the total analysis volume displayed in fig.3 is obtained, where this offset has been set to only one year and the ratio of Monte Carlo to real data events has been set to five. The result is a strong peak in 2008, and a falloff towards 2010 that is less steep than before.

2. Expected Reconstruction Operations

Starting from the year 2008, there will be no more “A” production of fresh data to perform, assuming that subtleties as processing of data of a possible F_L run have already been taken care of in 2007. There should be the standard reprocessing of the 2007 data. Then the issue of a general “grand reprocessing” of all HERA-II data will arise probably in 2009, since it will be the only way of bringing all data on the same footing in terms of

reconstruction algorithms and the calibration and alignment of the detector. This grand reprocessing would certainly be ideal to have as early as possible, but according to experience a substantial amount of time will be required to finalize code and constants on all relevant issues.

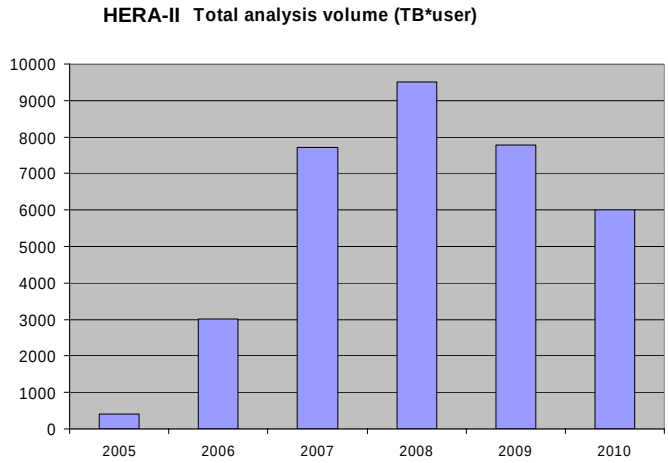

Fig. 3: Total analysis volume (real data and Monte Carlo)

3. Expected Analysis Operations

The year 2008 will mark a milestone since for the first time the full and final HERA-II statistics will become available for analysis. For this reason, the analysis volume is expected to reach an unprecedented peak in spite of the decrease of number of analysts. Also Monte Carlo production can be expected to reach a sustained all-time high. In the following years 2009/10, the analysis volume would normally decrease, but it is well possible that improvements through the grand reprocessing will lead to another boost of analysis and simulation activity.

During these years, it should be conceivable to expect that analysis will move more and more from MDST format to final n-tuples (see also section 5). A reasonable target may be to have the bulk of analyses moved to n-tuples by end of 2010. (On this point however, consensus should first be reached within the collaboration, in particular with the physics groups.) However, there may still be occasional need to access MDST or even RAW format data, and there will also still fresh Monte Carlo continue to be produced. For these reasons, it is advisable to keep central ZARAH services available for at least one year longer, i.e. until end of 2011.

4. Monte Carlo and Grid

Grid technology is already being used for about 2/3 of our current Monte Carlo production, and this fraction is expected to increase further, reducing our dependency on the traditional funnel clusters. Recently a new project has been started with the aim to

investigate the possibility of ZEUS data analysis on the grid. Since analysis is more data-intensive, it requires considerably more local resources on a grid site in comparison to Monte Carlo production, so that only specially selected sites may be ready and able to contribute. Nevertheless, it is conceivable that towards end of 2006, a certain amount of ZEUS data analysis can be performed on the grid. It is presently not planned to move event reconstruction to the grid.

5. Long-Term Conservation of ZEUS Data

It is unlikely that the present MDST-based analysis model (Fortran/Adamo/Zebra) can be kept working for much longer than ~2011. If the collaboration intends to keep the data usable for a much longer time, which would be natural in view of the unique role of the HERA data, only relatively simple formats of high abstraction level have a chance to survive. Such a format should be based on the level of physical quantities rather than hardware-related measures, and reside in an encoding that can be easily converted from one contemporary standard to the next. PAW, ROOT or XML may be choices to be considered. Such n-tuples in the most general sense could make us finally independent of particular mass storage installations, but would require a consensus of how to package events sufficiently dense to allow storage on less sophisticated media.

6. Cost Estimate for ZARAH Upgrade

According to the presently valid computing strategy, we will keep upgrading the capacity of the system regularly until including 2007, which is an important asset also for the years beyond. It is therefore possible to reduce the investments gradually and focus on replacements of older farm nodes and disk servers, with the intention of conserving our computing ability in accordance with the development of the analysis volume as outlined earlier. The estimated numbers are listed in table 1.

In case the ZARAH facility will continue to be used beyond 2011, additional resources may become necessary.

Year	ZARAH Upgrade [kEUR]	Computing running cost [kEUR]
2008	90	140
2009	80	140
2010	60	130
2011	40	110

Table 1: Cost estimates

7. ZARAH Operations

As long as any ZEUS-specific central computing services are required at DESY, also funding of personnel must be foreseen, i.e. one position for ZARAH operations.

8. Estimate for Computing Running Cost

From 2008 on we will not be writing any more raw data to tape, so that the corresponding part of the media cost will disappear. We will however still need media for reconstruction output of real data, as well as for freshly produced Monte Carlo.

The biggest fraction of computing running cost is the share of maintenance cost for the mass storage infrastructure at DESY, in particular tape robots and attached high end server machines. This share depends on our relative fraction of tape slots in use, and will not change very much by the end of data-taking, but may be affected by changes of the user community at DESY. The computing running cost cannot be predicted very reliably, and only rough estimates have been entered in table 2. (In the end it will be fixed by the DESY directorate.)

Acknowledgements

I would like to thank all my ZEUS colleagues who have given me their views and opinions in this matter, in particular A. Geiser, T. Haas, U. Stösslein, K. Wrona and W. Zeuner.

References

[1] J. Cole, T. Doyle, T. Haas, R. Mankel, W. Smith: ZEUS Computing Strategy: an Update for the years 2004 – 2007; ZEUS Note 03-024

[2] R. Yoshida: ZEUS Manpower; presented at PCOOR meeting on 20-Sep-2005