Summary of Atlas Database Meeting 24 February 1998

We had 5 presentations and a discussion at the end:

- **RD Schaffer: general points and status of Event access**
- **Christian Arnault: Identifier and range classes**
- **Martin Schaller: work on clustering strategies**
- **Simona Rolli: work with Objy on run I CDF data**
- **David Malon: points of interest from Grand Challenge work**
- **K. Sliwa: discussion on PAP for LCB project on data management and computing using distributed architectures.**

RD Schaffer
RD.Schaffer@cern.ch
4 March 1998

RD Schaffer

Objj V5:

mid-March to mid-April

Atlas database server:

will soon have a server machine

Detector Description Requirements:

finished review with a good round of discussions in Jan.

General planning for 1998:

Geant 3 digits for all detector systems in db	January 1998
first version of a detector description database available	June 1998
Geant 4 event data available for parts of the ATLAS detector	December 1998

Access to digits:

now ready to provide access events with digits in Inner Detector from Zebra tapes.

Have added to the model:

identifiers, detector **hierarchy, visitors** to collect detectors and digits, and a simple **Event** interface.

Christian Arnault: Identifiers

Discussed design of the Identifier and Range classes

=> these can be used to select/access data corresponding to detector elements

e.g. event raw data and detector description information

Martin Schaller: Clustering Strategies

He discussed some new work on a simple scheme for clustering events for overlapping subsets.

Simona Rolli: CDF Objy performance

Discussed performance questions for alternate scenarios:

1) One object per YBOS bank

2) Put same banks into Varrays of floats

Found that 2) is significantly better => moving away from concept of banks in DB.

RD Schaffer
RD.Schaffer@cern.ch

David Malon: DOE Grand Challenge

Working on optimising storage and access using Objy/ HPSS, primarily for RHIC.

Different parts of work may be of interest to Atlas:

Query Estimator - get number of events satisfying a query, estimate retrieval time

Storage Manager - can prefetch dbs from tape

Order Optimized Iterator - provides an iterator which is “correctly ordered” according to how events are stored. Allows easy parallelisation.

RD Schaffer
RD.Schaffer@cern.ch

Laura Perini: Computing Model analysis results

With M. Campenella, she has resumed the study of

where to put the cpu

where to put the data

(questions of networking have not yet been included.)

They have taken the analysis scenarios from the CTP and added variations.

They find some variation on the CPU requirements according to the different analysis scenarios for Central/Regional Centre/Local models.

=> further understanding of the analysis models, and inclusion of networking are next steps.

Krzysztof Sliwa: PAP for LCB project on data management and computing using distributed data architectures

It is generally agreed that an LCB project would be worthwhile.

The important questions now are:

defining well the breadth of the project, and

how it is coordinated with RD45 and the Technology Tracking Team of IT.

It was agreed to “take one’s time” so that things are well defined.

=> submission of the PAP may only be in the June LCB instead of April

Next steps:

First draft of PAP before Barcelona workshop,

General meeting at CERN of interested parties after the Barcelona workshop