Simulation of Grid Computations

Kathrin.Paschen@cern.ch*

June 23rd, 2003

*Financed by BMBF (the German Ministry for Education and Research)

Goal: evaluate variations of grid architecture with respect to

- performance and
- stability

in the context of the LHC experiments.

The idea is to find answers by simulating jobs in a grid environment.

Structure of a job:

1. request data (query by trigger, run period, properties of the event, etc.)

2. process data

3. return result (in the form of histograms for example) or write output to file.

4. goto 1.

Relevant for simulation:

- Size of return set of initial query
- Time complexity of initial query
- Common queries
- Processing time for 'typical' query
- Size of output

Jobs compete for access to resources such as CPU, storage and the network.

For the purposes of simulation, these resources can usually be viewed as black boxes.

- identify intelligent scheduling and replication strategies

- identify bottlenecks and points of failure

Which strategies qualify as intelligent depends on details of the physics jobs. For example, we could ask

- When should replicas of data be created or destroyed?

To answer this, we need to know how to predict future use of a piece of data based on its past use.

Also: estimate load on replica manager, define what happens when it fails.

Another example: scheduling.

- When several jobs need access to the same data, which is available only in one place, is it better to replicate the data or to have some of the jobs wait?

- What about jobs requiring more than one input file?

- Is advance reservation of resources worth the trouble? Is it needed?

A final example: graceful degradation

- How should the system detect and respond to errors?
- What is the system's expected average performance given realistic assumptions on failure rates?
- Can the programming mistake of a single user affect system performance globally?

Current work: designing and implementing experiments to answer these (and other) questions.

In order for our answers to be verifiably correct, it would be very useful to have a better understanding of the analysis process. It would also be useful to have access to monitoring data on the performance of test installations.