Physics results from the STAR experiment at RHIC benefit from production Grid data services.

News Update – 19 Mar 2004

Triggered Correlation Studies from Datasets on the Grid[1]

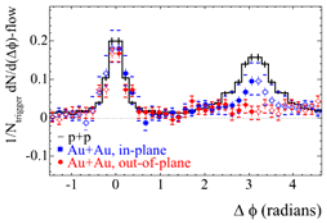

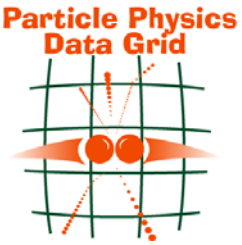

Physics on a Data Grid

A wealth of physics results have emerged from the RHIC[2] program and STAR[3] has published 26 papers in refereed journals, 12 of which were published since the start of 2003. The first publication from the d+Au run in the spring of 2003 was in print in a record four months after the completion of data taking. This paper, along with those of the three other RHIC experiments and previous results, provided strong evidence that the density of the medium generated in central Au+Au collisions is many times that of cold nuclear matter. The latest findings from STAR were covered in 14 talks at the Quark Matter 2004 conference. The first direct measurement[4] of open charm production at RHIC (example below) provides the first salvo of a full program of measurements utilizing the large datasets recorded in the current and future runs at RHIC.

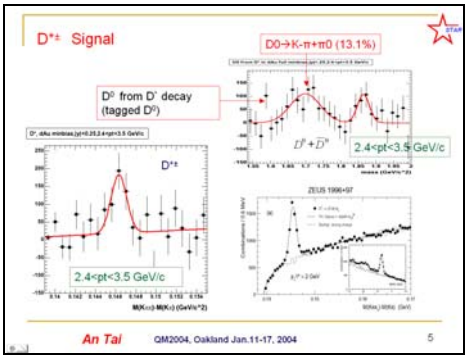

Data Grid Performance

The STAR computing groups at BNL and LBNL have operated a production Grid data movement service for nearly two years. In collaboration with the LBNL/SDM computer science group, they have hardened and extended the Storage Resource Management and GridFTP implementations to achieve sustained and robust automated data transfers of up to 5 TB a week in both directions between the HPSS mass storage systems at BNL and LBNL. This allows "next day" access to fresh data for analysis and physicists using the facilities at LBNL and BNL are able to collaborate more effectively on the analyses that have led to the recent physics results. Performance is measured regularly and monitoring is used to spot and identify problems in the data transfer. Peak data transfer rates from BNL to NERSC of up to 30 MB/sec are achievable with about 50 streams. Network engineers from RCF and ITD at BNL as well as NERSC are looking into issues that cause the single TCP stream performance to be less than one would expect. The figure below illustrates the robustness of the end-to-end system even though the time trace for individual files shows delays due to various component errors or down times.

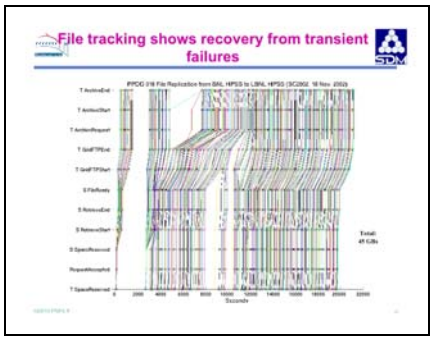

Improvements Due to Grid Technologies

In the BG age (before grid) transferring 100's of Gigabytes of data and 1000's of files was tedious, time consuming, error prone and insecure. Typically a request for data would be made (by email) that amounted to identifying a directory path in the source HPSS system. The Data Transfer Technician (DTT) would convert this request into a list of files by making a directory listing using ftp. This list would be divided into "convenient size" chunks (able to fit on disk) for extraction from the source HPSS using one set of scripts. Another set of scripts would perform the clear text ftp across the wide area network, and a third set of scripts would then copy files into the target HPSS system. This would be repeated until all the "chunks" of files were transferred. Each stage was prone to errors, network glitches, machine response timeouts, disks filling up, etc., and a "chunk" of perhaps 100 files would fail somewhere. The DTT would have to check the logs, determine what had succeeded and what had failed, and do the transfer stage again with the remaining file list. Overall this process would take up to 10 days for < 1 TB of data and would take a large fraction of an FTE, and still result in about 1% inconsistency in the number of files in a dataset between the two sites.

Now in the AG age (after grid) we have end-to-end data transfer using grid middleware along with catalog components such that replicated data are automatically registered after successful replication. The course of metadata and file replication is indicated in the diagram below.

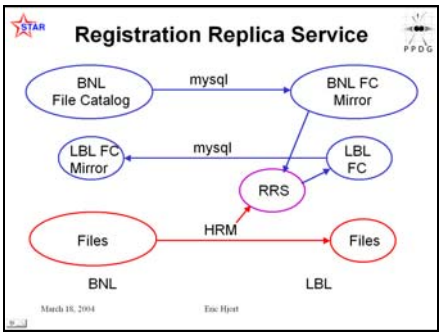

Transfer of TB datasets is accomplished with a single command and the average throughput is 10 times greater than BG (mostly due to improved operational efficiencies). Unsuccessful transfers due to transient problems in the end-to-end system, though rare, may be corrected easily by comparing the BNL and NERSC file catalogs to generate a secondary list of files to be transferred. A file discrepancy rate of 0.02% (50 times less than BG) or less is now obtained with the AG infrastructure.

Reliable file transfer and catalog registration not only benefits the scientists and the experiment's physics throughput but also considerably reduces the tedious everyday task of the DTT, allowing for more productive or new developments.

---

[1] K. Schweda, "STAR Experiment Highlights", Quark Matter 2004, http://qm2004.lbl.gov/

[2] www.bnl.gov/RHIC

[3] www.star.bnl.gov

[4] An Tai, "STAR Measurements of Open Charm", Quark Matter 2004, http://qm2004.lbl.gov/