ATLAS Detector Data Processing on the Grid

Alexandre V. Vaniachine on behalf of the ATLAS Collaboration

***Abstract*—The ATLAS detector is in the second year of continuous LHC running. A starting point for ATLAS physics analysis is data reconstruction. Following the prompt reconstruction, the ATLAS data are reprocessed, which allows reconstruction of the data with updated software and calibrations providing coherence and improving the quality of the reconstructed data for physics analysis.**

The large-scale data reprocessing campaigns are conducted on the Grid. Computing centers around the world participate in reprocessing providing tens of thousands of CPU-cores for a faster throughput. Reprocessing relies upon underlying ATLAS technologies providing reproducibility of results, scalable database access, orchestrated workflow and performance monitoring, dynamic workload sharing, and petascale data integrity control. These technologies are also empowering ATLAS physics and subsystem groups in further data processing steps on the Grid.

We present the experience of large-scale data reprocessing campaigns and group data processing on the Grid.

I. Introduction

THE "raw" data from the ATLAS detector (Fig. 1) are processed to produce the reconstructed data for physics analysis. During reconstruction ATLAS applications are processing raw detector data with sophisticated algorithms to identify and reconstruct physics objects such as charged particle tracks. Fig. 2 shows data processing flow of raw event and conditions/calibrations data used in reconstruction operations. The distributed multi-tier data processing architecture handles petascale data flow. Since the detector data are comprised of independent events, massively parallel reconstruction applications process one event at a time. Events taken during few minutes are collected in one raw file. Files with events that are close in time are collected in one dataset.

TABLE I. ATLAS Computing Resources

Center	Number	Cores (10^3)	Role
CERN	1	10	Data recording Calibration and prompt processing Data distribution
Tier-1	10	35	Reprocessing Group data processing Permanent storage
Tier-2	70	65	Simulation End-user analysis Storage

Manuscript received November 21, 2011. This work was supported in part by the U.S. Department of Energy, Division of High Energy Physics, under Contract DE-AC02-06CH11357.

A. V. Vaniachine is with the Division of High Energy Physics, Argonne National Laboratory, Argonne, IL 60439 USA (telephone: 630-252-1878; e-mail: vanachine@anl.gov).

The "first-pass" processing of the raw event data at the ATLAS Tier-0 site provides promptly the data for quality assessment and analysis. Later, the quality of the reconstructed data is improved by optimizing further software algorithms and conditions/calibrations data. A recent reprocessing of 2011 proton-proton collisions events recovered more than 8% of good quality data for physics analysis. For data processing with improved software and/or conditions and calibrations (reprocessing) we use distributed computing resources.

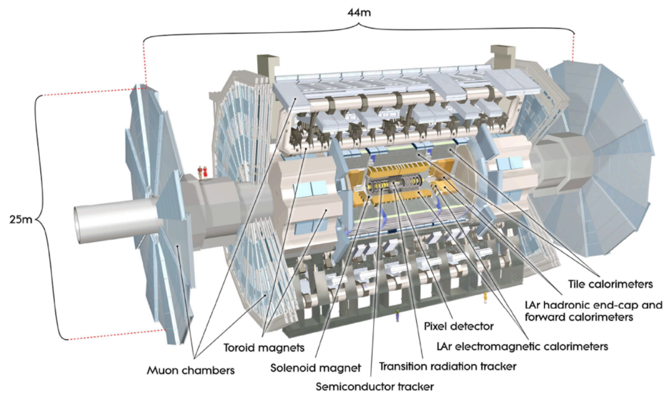

Fig. 1. The data for reconstruction come from the ATLAS detector [1].

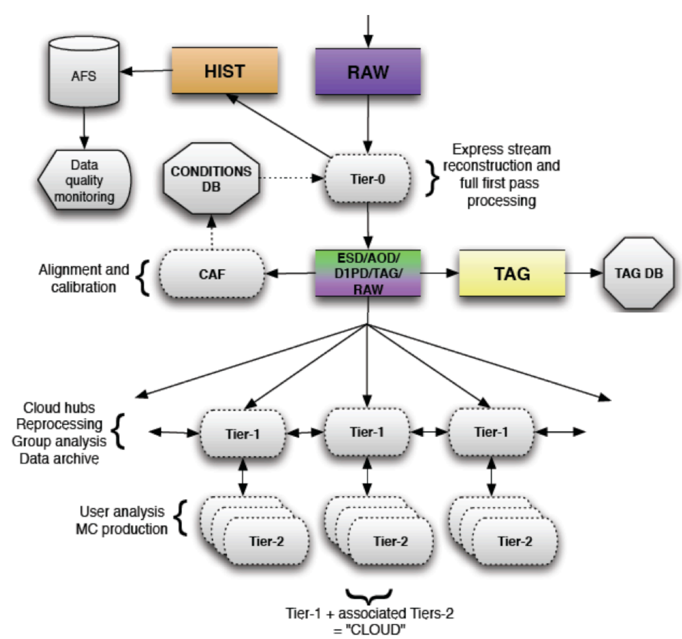

Fig. 2. Simplified flow of raw event and conditions/calibrations data used in reconstruction at the Tier 0 site at CERN (top) and on the Grid at the Tier-1 sites (bottom).

II. Requirements and Technologies

Table I shows that ATLAS distributed computing resources are an order of magnitude larger than the resources at the CERN alone. ATLAS uses Grids with three different interfaces split in ten "clouds" organized as large computing centres with tape data storage (Tier-1 sites) each associated with 5-6 other computing centres (Tier-2 sites). Plus more than a hundred of ATLAS Tier-3 sites used for the physics analysis. All these computing resources are used in further data processing steps following the ATLAS detector data reconstruction.

A. *Reproducibility of Results*

For Grid Data Processing (GDP) we developed Sites Validation tools that ensure that all Grid sites produce numerically identical outputs for the same raw data inputs. Sites Validation ensures that the same ATLAS software and conditions/calibrations versions are used worldwide and excludes discrepancies due to heterogeneity of the Grid such as site-specific CPU (AMD/Intel), system libraries, batch systems, etc. Sites Validation tools also enable software release validation with a large-scale data sample, which reveals rare software bugs.

B. *Scalable Database Access*

To prevent scalability problems in database access on the Grid we developed a GDP technology for access to the conditions/calibrations data, which is similar to the raw event data and software release distribution on the Grid. Matching the underlying architecture of the data-processing Grid, the Database Release technology integrates in a single dataset all conditions/calibrations data required for reconstruction [2].

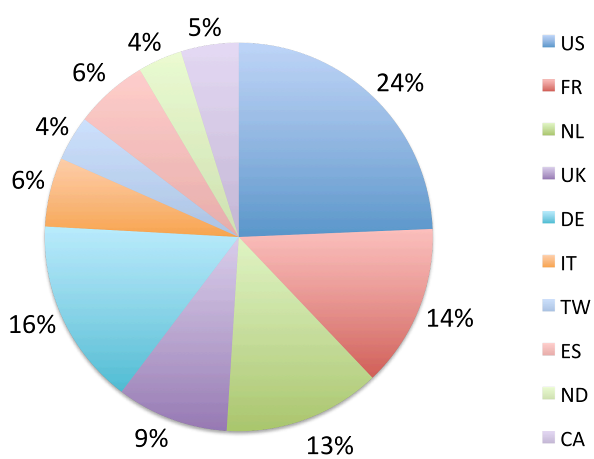

Fig. 3. During data taking, ATLAS data management algorithm placed two copies of each raw dataset at different Tier-1 sites according to the agreed shares; one copy was to disk, another copy to tape. ATLAS workload management system PanDA matched data processing tasks to the Tier-1 sites to minimize the reprocessing duration. Regardless whether the copy was on disk or on tape, PanDA algorithm brokered each raw reconstruction task[2] avoiding busy sites, which resulted in the reconstruction jobs[1] shares shown on the pie chart. (Dynamic workload sharing causes slight deviations from the agreed data shares.)

C. *Workflow Orchestration*

Leveraging the underlying workload management system PanDA [3], our stable but flexible GDP meta-application framework orchestrates ATLAS data processing applications to ensure efficient usage of tens of thousands of CPU-cores. Designed for generic applications, the system has a pull scheduling implemented with "pilot" agents, an approach described earlier in [4]. During reprocessing the system monitors site performance and supports dynamic workload sharing minimizing the reprocessing duration (Fig. 3). In addition, the fault-aware GDP framework allows sophisticated management of jobs[1] and tasks[2] to support fault management and system resilience.

D. *Data Integrity*

To facilitate physics discoveries, the reprocessing must minimize event losses. This is assured in GDP by automated resubmission of the failed data processing jobs, which excludes transient failures. The events that cannot be reconstructed during the reprocessing campaign are recovered promptly in a dedicated post-processing step using an updated software release and/or conditions and calibrations.

III. Petascale Data Processing Experience

It takes about three million core-hours to processes one petabyte of ATLAS data. During processing of one petabyte of data on the Grid, the peak throughput achieved twenty thousand CPU-cores. Since transient job failures and retries delay the reprocessing duration, the average reconstruction throughput, defined as the total number of core-hours divided by the duration of the reconstruction, serves as a performance benchmark. Optimization of our MapReduce-like workflow and other improvements increased the average reconstruction throughput from 3.6 10^3 cores in 2010 to 5.0 10^3 cores in 2011. This halved the duration of the petabyte-scale reprocessing on the Grid from almost two months in 2010 to less than four weeks in 2011 (Fig. 4).

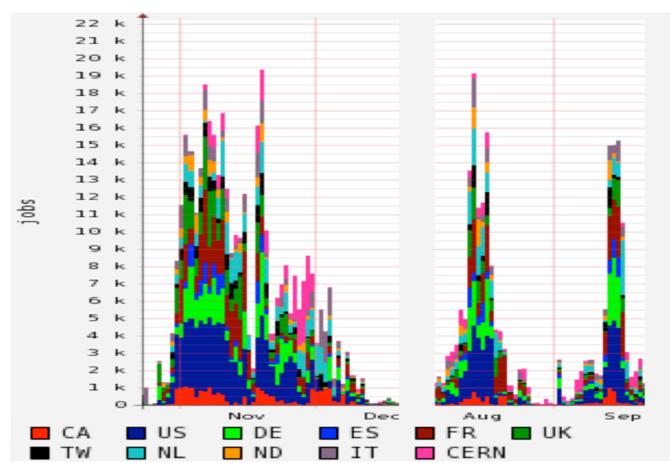

Fig. 4. Number of jobs running concurrently at the Tier-1 sites and CERN during petascale reprocessing campaigns in November-December, 2010 and in August-September, 2011.

A. *Six Sigma Quality*

Thanks to the technologies described in the previous section, no events were lost during the main ATLAS

[1] One reconstruction job is a 32-bit Linux Python/C++ data processing application using up to 2 GB of memory on a single core for 8 to 12 hours to process few GB of input data and producing output data of a similar volume.

[2] One task is a collection of jobs that process the same dataset.

reprocessing campaign of the 2010 data that reconstructed on the Grid more than 1 PB of data with 0.9 10^9 events. Attesting to high ATLAS software quality, in a recent 2011 data reprocessing only two collision events out of 0.9 10^9 events total could not be reconstructed. These events were reprocessed later in a dedicated data recovery step.

Later, silent data corruption[1] was detected in six events from the reprocessed 2010 data and in one case of five adjacent events from the 2011 reprocessed data. Corresponding to event losses below the 10^{-8} level, this demonstrates the "six sigma quality" performance sustained during one year.

GDP experience shows that correcting silent data corruption in a distributed petascale event store is prohibitively costly. To assure scalability, the data corruption must be detected *in situ*[2]. In a petascale event store, every layer of services should not assume that the underlying layer never provide corrupted or inconsistent data. We must have redundancy in order to detect and recover from data corruption errors.

During the "first-pass" processing, the event losses have been at the 10^{-4} level, which were tolerated to deliver the reconstructed data promptly [5]. In reprocessing, GDP achieved a reduction in the event losses by four orders of magnitude at the expense of the core-hours used to recover transient failures.

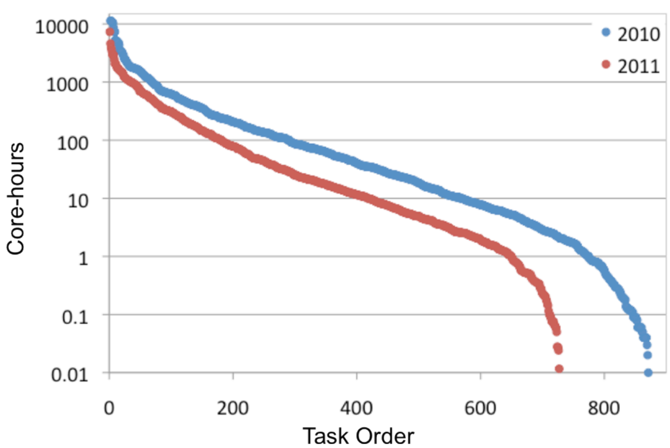

Fig. 5. Tasks ordered by core-hours used to recover transient failures.

B. *Performance Monitoring*

Job resubmission avoids data loss at the expense of core-hours used by the failed jobs. Fig. 5 shows that the distribution of tasks ordered by core-hours used to recover transient failures is not uniform: most of core-hours required for recovery were used in a small fraction of tasks. In 2010 reprocessing, the core-hours used to recover transient failures were 6% of the total core-hours used for reconstruction. In 2011 reprocessing, the core-hours used to recover transient failures were reduced to 4% of the total core-hours used for the reconstruction.

IV. Group Data Processing

The reprocessing technologies empowered further data processing steps on the Grid performed by dozens of ATLAS physics groups with coordinated access to computing resources worldwide. Unlike major reprocessing campaigns that are conducted only few times per year, the centrally managed production for physics groups process the whole available dataset once every few weeks, providing further improvements in the data used for ATLAS physics analysis shortly after the reprocessing or data taking. In 2010 group data processing consumed most of the ATLAS Grid computing resources. In 2011 group data processing achieved peak consumption of allocated Grid computing resources at the 100% level. The GDP technologies were also adopted for the trigger reprocessing, which is performed to validate new trigger menus and/or software releases during data taking.

V. Conclusion

In 2011 the ATLAS detector continues to perform extremely well recording petabytes of good-quality data for physics analysis. ATLAS technologies for data processing on the Grid coped well with the first petabytes of Large Hadron Collider data. Since the start of LHC data taking ATLAS successfully completed nine reprocessing campaigns on the Grid delivering sustained "six sigma quality" performance in ATLAS distributed computing operations. Thanks to workflow optimization, a recent petascale reprocessing campaign was completed within weeks providing a coherent 2011 dataset for physics analysis.

Peaking at twenty thousand cores during reprocessing of one petabyte of data on the Grid, an average reconstruction throughput increased from 3.6 10^3 cores in 2010 to 5.0 10^3 cores in 2011. In comparison with the 2010 reprocessing, the fraction of core-hours used to recover transient failures was reduced by a factor 1.5 to 4% in the 2011 reprocessing.

The workload management dispatches tasks according to their data location and resource load, tracks job status and avoids data loss. Further data processing steps performed by dozens of ATLAS physics groups are empowered with coordinated access to computing resources on the Grid. We are ready for the next petascale data processing challenges.

Acknowledgment

I wish to thank all my colleagues who contributed to ATLAS reprocessing activities.

References

[1] The ATLAS Collaboration, G Aad *et al*. (2008, August). "The ATLAS experiment at the CERN Large Hadron Collider," *J. Inst*., [Online]. Vol. 3, S08003. Available: http://iopscience.iop.org/1748-0221/3/08/S08003

[2] M. Borodin, P. Nevski, A. Vaniachine on behalf of the ATLAS Collaboration. (2011). "Scaling up ATLAS Database Release technology for the LHC long run," *J. Phys.: Conf. Ser.*, to be published.

[1] During job execution, the data corruption may be missed by the ATLAS software framework or Distributed Computing components. This is called silent data corruption (regardless whether it was logged or not).

[2] Further reduction in data corruption rates requires improvements in logging and log parsing during the job execution.

[3] T. Maeno. (2008). “PanDA: distributed production and distributed analysis system for ATLAS,” *J. Phys.: Conf. Ser.,* [Online]. Vol. 119, 062036. Available: http://iopscience.iop.org/1742-6596/119/6/062036

[4] P. Nevski, A. Vaniachine, T. Wenaus. (2001). “Bilevel Architecture for High-Throughput Computing” in *Proc. of the Int. Conf. on Computing in High Energy and Nuclear Physics* (CHEP 2001), Beijing, China, p. 696. Available: http://www.ihep.ac.cn/~chep01/paper/10-016.pdf

[5] S. Oda. (2011, Apr 19). “Prompt calibration, assessment of data quality and data processing.” Poster ATL-PHYS-SLIDE-2011-163. [Online]. Available: http://cdsweb.cern.ch/record/1345770