

A CASE FOR APPLICATION-AWARE GRID SERVICES

Gabriele Garzoglio^{*}, Andrew Baranovski, Parag Mhashilkar, FNAL, Batavia, IL 60510, USA
Ljubomir Perković, CTI, DePaul University, Chicago, IL
Anoop Rajendra, University of Texas at Arlington, Arlington, TX 76019, USA

Abstract

In 2005, the DZero Data Reconstruction project processed 250 tera-bytes of data on the Grid, using 1,600 CPU-years of computing cycles in 6 months. The large computational task required a high-level of refinement of the SAM-Grid system, the integrated data, job, and information management infrastructure of the RunII experiments at Fermilab. The success of the project was in part due to the ability of the SAM-Grid to adapt to the local configuration of the resources and services at the participating sites. A key aspect of such adaptation was coordinating the resource usage in order to optimize the typical access patterns of the DZero reprocessing application. Examples of such optimizations include database access, data storage access, and worker nodes allocation and utilization.

A popular approach to implement resource coordination on the grid is developing services that understand application requirements and preferences in terms of abstract quantities e.g. required CPU cycles or data access pattern characteristics. On the other hand, as of today, it is still difficult to implement real-life resource optimizations using such level of abstraction. First, this approach assumes maximum knowledge of the resource/service interfaces from the users and the applications. Second, it requires a high level of maturity for the grid interfaces. To overcome these difficulties, the SAM-Grid provides resource optimization implementing application-aware grid services. For a known application, such services can act in concert maximizing the efficiency of the resource usage. We describe what optimizations the SAM-Grid framework had to provide to serve the DZero reconstruction and Monte Carlo production. We also show how application-aware grid services fulfill the task.

INTRODUCTION

High energy physics applications have different resource utilization requirements. As the LHC experiment will start taking data in 2007, competition for computing resources for high energy physics will skyrocket. In this scenario, the optimization of resources becomes of crucial importance to achieve the maximum amount of computing possible for the field. Studying the characteristics of the typical high energy physics applications is the key to understand how to implement these optimizations.

We have implemented resource optimization for the DZero applications within the context of the SAM-Grid system [1]. The SAM-Grid is a meta-computing infrastructure for job, data, and information management.

In this paper we present the characteristics of three major computing activities, data analysis, data reconstruction, and Monte Carlo simulation, referring to our experience with the DZero experiment. We then present four examples where the knowledge of the application helps the grid optimize the resource utilization. We use these examples to show that application-specific knowledge helps grid services optimize resources and run grid jobs efficiently.

THE OPTIMIZATION PROBLEM

To optimize the resource usage for the DZero experiment at Fermilab, we have considered the characteristics of three major types of applications: data filtering, also called data "reconstruction", the production of simulated events, and data analysis. This third activity broadly consists of the selection and the statistical study of particles with certain characteristics, with the goal of achieving physics measurements. It should be noted that the first two activities are indispensable for the third one. During data reconstruction, the binary format of events from the detector is transformed into a format that more easily maps to abstract physics concepts, such as particle identification, energy, momentum, and direction measurements. The original format, instead, is very closely dependent on the hardware layout of the detector, in order to guarantee the performance of the data acquisition system, and is not suitable for data analysis. The production of simulated events, also called "Monte Carlo" production, is necessary to understand the characteristics of the detector either related to the hardware, such as the particle detection efficiency, or to physics phenomena, such as signal to background discrimination.

These three typical activities, which ultimately correspond to software application families, differ among themselves principally, but not uniquely, by the usage of the computing resources. The communities that run the software range from a handful of almost dedicated experts in the case of reconstruction and Monte Carlo activities, to potentially the whole physics community in the case of data analysis. The typical duration of a single reconstruction or Monte Carlo job is dozens of hours, while data analysis ranges from a few minutes to days, depending on the problem studied. All the activities are

^{*}garzoglio@fnal.gov

CPU intensive, but while both reconstruction and analysis are highly I/O intensive, Monte Carlo is not. In fact, Monte Carlo almost never requires any input data, while for reconstruction and, especially, analysis the input ranges from a GB to hundreds of GB. In addition, while the data access pattern of reconstruction is highly predictable, since all the "raw" data have to be filtered a few times throughout the lifetime of the experiment, the data access patterns of data analysis varies widely, as a few datasets can be accessed over and over again, while others may be almost neglected. All three activities can be run trivially in parallel because of the independent nature of particle physics events. On the other hand, while Monte Carlo and reconstruction are purely "batch" activities, analysis is run in both interactive and "batch" modes.

Table 1 summarizes the order of magnitudes of different characteristics of reconstruction, Monte Carlo and data analysis for the DZero experiment.

Activity	Description	Community	Load	time/job
Reconstruction	data filtering	Small	CPU & I/O	10 hours
Montecarlo	data simulation	Small	CPU	10 hours
Analysis	data mining	Large	CPU & I/O	hours to days

Activity	Input/Job	Output/Job	Input/Year	Output/Year
Reconstruction	GB	GB	100 TB	100 TB
Montecarlo	None	10 GB	None	TB
Analysis	100 GB	GB	varies	varies

Table 1: Comparison of different characteristics of three typical computation activities of the DZero experiment. The bottom table focuses on the input/output data size. The numbers represent the order of magnitude.

Even restricting our system to manage resources for data reconstruction [2] and Monte Carlo generation [3] only, it was still a challenge to run efficiently jobs with such different characteristics. In order to let the grid organize the usage of the resources efficiently, we decided to expose details of the applications to the grid.

DATABASE ACCESS PROBLEM

Grid jobs submitted to an execution site are split into multiple parallel instances of the same application by the SAM-Grid grid-to-fabric interface [4]. This typically results in dozens to hundreds of jobs starting approximately at the same time and, therefore, accessing key resources essentially concurrently.

In practice, not all the services have the same degree of accessibility. In particular for Monte Carlo generation, the parameters describing what type of physics to generate were accessed from a central database, which initially responded with a "denial of service" to 40% of the jobs. Introducing retrieval with randomized exponential back off reduced the final job failure rate to 5%. Despite the reduced failure rate, grid jobs and their retrials increased the load of the database to a point where interactive access was extremely inconvenient (minutes per query).

This problem was properly solved by informing the grid of the database access characteristics of the Monte Carlo application. All the hundreds of jobs submitted by the grid, in fact, were parallel replicas of a single grid job and, therefore, required access to the same input parameters from the database.

The grid-to-fabric interface was enhanced to perform a single database access per grid job, when the job entered the site. The information was saved and redistributed to the parallel jobs by internal cluster transport mechanisms. This solution reduced the "denial of service" failure rate to essentially 0% and still maintained a high availability for interactive database accesses.

In conclusion, access to a grid resource (the database) was optimized by instructing grid components (the grid-to-fabric interface) of the characteristics of the application (parallel jobs requiring the same input parameters).

DATA STORAGE ACCESS PROBLEM

Different applications have different typical input data access patterns. For DZero, data reconstruction applications begin data processing when a single input file, typically 1 Gigabyte in size, is delivered to the worker node. Instead, data merging applications, used in production operations to concatenate files typically 200 Megabytes in size, begin processing when multiple "small" input files are delivered to the worker node. Optimizing access to the storage resources with such different regimes is a concern.

In the SAM-Grid, applications transfer files from storage services that maintain queues of data access requests. The storage services, in fact, control their load by granting access to the data transfer servers a few requests at the time. Access to a transfer server is granted in the order in which the access request is submitted. When reconstruction and merging applications use the same data queue to access their input, transfer requests for the various input files are interleaved. This leads to inefficiencies, because in real life, on a cluster, reconstruction jobs are one or two order of magnitude more abundant than merging jobs. This means that requests for each input file of a merging application are interleaved with a dozen input files of reconstruction applications. Therefore, before starting processing data, a merging application often needs to wait for these multiple reconstruction transfers to occur, thus substantially increasing its idle time. For a cluster of 900 CPU, this idle time is between one and two hours.

This inefficiency can be reduced by instructing the grid of the input data access characteristics of the application. Knowing the number of required input data files, the data storage service can organize requests from reconstruction applications in a queue different from the requests from merging applications. Thus, a few merging applications only compete amongst themselves for file access, drastically reducing their idle time.

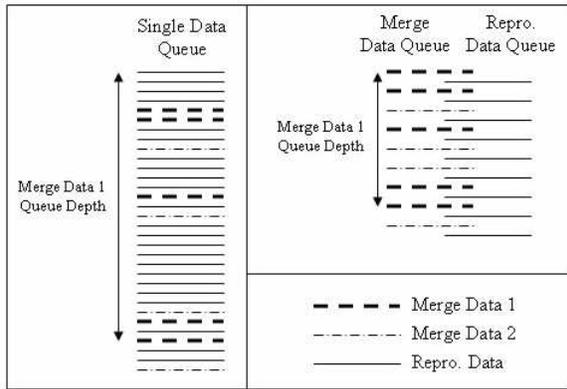


Figure 1: a diagram representing queues of requests for file access. On the left, a single queue manages requests from reprocessing jobs (straight lines), and merging jobs (dashed and dashed-dotted lines). Reprocessing jobs are two orders of magnitude more abundant than merging jobs. Merging job 1 needs to access five input files before it can start running (dashed lines, bold for clarity). On the right, requests from merging and reprocessing jobs are managed by two different queues. If access requests are granted one at the time, the queue depth for merging job 1 is much shorter than in the case of the single queue (left diagram). If the data storage server knows the typical data access pattern of the jobs, it can optimize access to the data. The SAM-Grid storage elements have knowledge of the typical data access patterns of each application.

In conclusion, as in the previous example, access to grid resources (data files) was optimized by instructing grid components (the storage service) of the characteristics of the application (multiple or single input data requirement).

WORKER NODES ALLOCATION PROBLEM

The grid-to-fabric interface of the SAM-Grid submits multiple batch jobs for every grid job entering the site. How many worker nodes should be allocated for a given application? In general, to accomplish the same amount of computation for a grid job, the fewer batch jobs are submitted, the longer each job runs, and vice versa. There is an acceptable range for the running time of a job. Batch jobs should not run too long to minimize the probability of termination before completion. Jobs are typically terminated because they run beyond the maximum wall-clock time allowed by the local scheduler, or because they are evicted due to a higher-priority job entering the scheduler, or because of hardware failures. On the other hand, batch jobs should not run too short in order to maximize the ratio between running time and setup time i.e. the time needed to prepare the job environment (in the SAM-Grid typically around half an hour).

The “suitable” expected running time is managed by the grid controlling the number of worker nodes allocated for running the application. It should be noted that applications may have additional constraints on the

number of jobs. These constraints are dictated by considerations on ease of bookkeeping and of recovery after failures. In any case, the number of worker nodes to allocate depends on the type of application. For reconstruction applications, the grid-to-fabric interface allocates a worker node for every file in the dataset specified for the grid job. Given the computational requirements of the reconstruction application, this approach gives an acceptable running time of a few hours on a modern CPU and eases bookkeeping and recovery operations. For Monte Carlo applications, the interface computes the number of worker nodes to be allocated by dividing the total number of events to be produced as specified for the grid job by the “optimal” number of events per job. The “optimal” number of events is a parameter configured at the site, considering the speed of the average CPU at the site, the computational requirements of the Monte Carlo application, and other scheduler constraints (maximum allowed wall clock time, etc.).

In conclusion, as in the previous examples, allocation of grid resources (worker nodes) is optimized by instructing grid components (the grid-to-fabric interface) of the characteristics of the application (computational requirements of the application and other constraints).

MINIMAL RESOURCE IDLE TIME PROBLEM

Grid jobs are often internally composed of interdependent tasks. We let the grid manage the order of execution of each internal task/job automatically. This automation minimizes the idle time between job submissions, thus minimizing the idle time of the resources.

In order to decide whether to submit a job, the grid must be able to determine whether the jobs on which it depends were successfully executed. In general, determining the success of a job is a difficult task. In case of complex computational activities, success is generally never defined only by the exit status of the job. To determine whether a Monte Carlo generation job was successful, for example, the grid has to check the number of events produced by the job by querying a bookkeeping database and compare this number with the number of events originally requested. Success is determined by policy: typically, if more than 90% of the events have been produced, the job is successful. For reconstruction applications the success policy is defined differently: typically a job is successful only if it has reconstructed all the input files, unless subsequent recovery jobs fail twice on the same event with the same error, thus exposing a corrupted input file. At any rate, having the grid determine the success of a job is an application-specific task.

In conclusion, as in the previous examples, the idle time of grid resources is minimized by instructing grid components (the job management component) of the

characteristics of the application (policy defining the success condition).

CONCLUSIONS

Application-specific knowledge is important in the optimization of grid resources. Two approaches are possible:

1) Applications communicate their requirements and preferences in terms of abstract resource/service-specific quantities. This is difficult to achieve as it requires a very high level of maturity of the grid interfaces and a thorough understanding of application requirements.

2) Applications rely on Application-Aware Grid Services for resource optimizations. This is less general but easier to implement and extend.

The SAM-Grid uses successfully Application-Aware Grid Services for grid resource optimization.

REFERENCES

- [1] I. Terekhov et al., "Meta-Computing at D0", Nuclear Instruments and Methods in Physics Research, Section A, NIMA14225, vol. 502/2-3 pp 402 - 406.
- [2] J. Snow, D. Wicke, M. Diesburg, G. Garzoglio, G. Davies, "DØ Data Reprocessing with SAM-Grid", in Proceedings of Computing in High Energy Physics 2006 (CHEP06), Mumbai, India, Feb 2006
- [3] G. Garzoglio, I. Terekhov, J. Snow, S. Jain, A. Nishandar, "Experience producing simulated events for the DZero experiment on the SAM-Grid", in Proceedings of Computing in High Energy and Nuclear Physics (CHEP04), Interlaken, Switzerland, Sep 2004.
- [4] G. Garzoglio, I. Terekhov, A. Baranovski, S. Veseli, L. Lueking, P. Mhashilkar, V. Murthi, "The SAM-Grid Fabric services", Nuclear Instruments and Methods in Physics Research, Section A, 534:33-37,2004