

Building Efficient Data Planner for Peta-scale Science

Thursday, February 25, 2010 2:25 PM (25 minutes)

Unprecedented data challenges both in terms of Peta-scale volume and concurrent distributed computing have seen birth with the rise of statistically driven experiments such as the ones represented by the high-energy and nuclear physics community. Distributed computing strategies, heavily relying on the presence of data at the proper place and time, have further raised demands for coordination of data movement on the road onwards achieving high performance. Massive data processing will be hardly “fair” to users and hardly using network bandwidth efficiently whenever diverse usage patterns and priorities will be involved unless we address and deal with planning and reasoning about data

movement and placement. Although there exist several sophisticated and efficient point-to-point data transfer tools, the lack of global planners and decision makers, answering questions such as “How to bring the required dataset to the user?” or “From which sources to grab the replicated data”, is for most part lacking.

We present our work and a status of the development of an automated data planning and scheduling system, ensuring fairness and efficiency of data movement by focusing on the minimal time to realize data movement (delegating the data transfer itself to existing transfer tools). Its principal keystones are self-adaptation to the network/service alteration, optimal selection of transfer channels, bottlenecks avoidance and user fair-share preservation. The planning mechanism is built on constraint based model, reflecting the restrictions from reality by mathematical constraints, using Constraint Programming and Mixed Integer Programming techniques. In this presentation, we will concentrate on

clarifying the overall system from a software engineer’s point of view and present the general architecture and interconnection between centralized and distributed components of the system. While the framework is evolving toward implementing more constraints (such as CPU availability versus storage for a better planing of massive analysis and data production), the current state of our implementation in use for STAR within multi-user environment between multiple sites and services will be presented and the benefit and consequences summarized.

Primary authors: Dr LAURET, Jerome (BROOKHAVEN NATIONAL LABORATORY); Dr SUMBERA, Michal (Nuclear Physics Inst., Academy of Sciences, Praha); Mr ZEROLA, Michal (Nuclear Physics Inst., Academy of Sciences); Prof. BARTAK, Roman (Faculty of Mathematics and Physics, Charles University, Praha)

Presenter: Mr ZEROLA, Michal (Nuclear Physics Inst., Academy of Sciences)

Session Classification: Thursday, 25 February - Computing Technology for Physics Research

Track Classification: Computing Technology for Physics Research