Contribution ID: **69** Type: **oral presentation**

Resource Predictors in HEP Applications

Thursday, September 30, 2004 2:00 PM (20 minutes)

The ATLAS experiment uses a tiered data Grid architecture that enables possibly overlapping subsets, or replicas, of the original set to be located across the ATLAS collaboration. The full set of experiment data is located at a single Tier 0 site, and then subsets of the data are located at national Tier 1 sites, smaller subsets at smaller regional Tier 2 sites, and so on. In order to understand the data needs, both in terms of access, replication policy, and storage capacity, we need good estimations of resource needs for data manipulation. Specifically, we envision a time when a user will want to determine which is more expedient, downloading a replica from a site or recreating it from scratch.

This paper presents our technique to predict the behavior of ATLAS applications, and then to combine this information with Internet link bandwidth estimation to improve resource usage in the ATLAS Grid environment. We studied the parameters that affect the execution time performance of event generation, detector simulation, and event reconstruction. Our results show that we can achieve predictions within 10-40% of the execution time (depending on the application), better than many other pragmatic prediction techniques. We implemented a software package to provide data transfer bandwidth estimation and execution time prediction that can be used with the Chimera software to aid in managing application execution and to improve resource usage for ATLAS.

Primary authors: HUTH, J. (HARVARD UNIVERSITY); SCHOPF, J. (Argonne National Lab); GRINSTEIN, S. (HARVARD UNIVERSITY)

Session Classification: Distributed Computing Services

Track Classification: Track 4 - Distributed Computing Services