The Farming of the Data of HyperCP Experiment

Lanchun Lu
University of Virginia

(April 24, 2001)

Abstract

The main goal of HyperCP(E871) experiment at Fermilab is to search for CP violation in Ξ and Λ decays at the $\sim 10^{-4}$ level. This level of precision dictates a data sample of over a billion events. The experiment collected about 240 billion raw events on about 30,000 5-GB tapes in ten months of running in 1997 and 1999. In order to analyze this huge amount of data, the collaboration has reconstructed the events on a farm of 55 dual-processor Linux-based PCs at Fermilab. A set of farm tools has been written by the collaboration to automate much of the farming and to allow non-expert farm shifters to submit and monitor jobs through a web-based interface. Special care has been taken to produce a robust system which facilitates easy recovery from errors. The code has provisions for extensive monitoring of the data on a spill-by-spill basis, as is required by the need to minimize potential systematic errors. Over 16 million plots of various parameters produced from the farm analysis can be accessed through a data management system. The entire data set was farmed in eleven months, or about the same time that was taken acquiring the data. We will describe the architecture of the farm, our experience in operating it, and show some results from the farm analysis.