# 1 Introduction

KTeV is a fixed target experiment at FNAL. It has two distinct operating modes making it two experiments. E832 is a high precision measurement of $\epsilon'/\epsilon$, the direct CP violation parameter. E799 is an experiment to study rare kaon decays, and has about 15 distinct triggering conditions, each one aiming at a different rare kaon decay mode. Thus the E799 data is easy to divide up for separate analyses, whereas the E832 data must all be analyzed in order to get a result.

Table 1 summarizes the data collected by KTeV in the 1996/1997 fixed target run.

Table 1: KTeV Data Summary

| Experiment | # Data Tapes | Data (TB) | # Calibration Tapes | Calibration (TB) |
|---|---|---|---|---|
| E832 | 2650 | 33 | 340 | 2 |
| E799 | 870 | 12 | 120 | 0.6 |
| Total | 3520 | 45 | 460 | 2.6 |

# 2 Computing Parameters

Table 2 summarizes the relevant parameters for the KTeV data analysis and Monte Carlo generation.

Table 2: KTeV Computing Parameters

| Rec. Time | Ana. Time | MC Gen. Time | Read Event Time | Event Size | MC Event Size |
|---|---|---|---|---|---|
| 3.3 | 5 ms (440MHz DEC 8400) | 14.3 ms | 4 ms (from DLT) | 7 kB | 2 kB |
| 0.93 mip-seconds | 1.4 mip-seconds | 4.0 mip-seconds | | | |

# 3 Collaboration Makeup

KTeV consists of about 80 collaborators from 12 different institutions, 10 US institutions and 1 Japanese institution. Table 3 lists the institions, the aproximate number of people from each institution computing (when relevant) at the institution, and the approximate fraction of data analyzed at each insitution (typically this represents a rare kaon decay mode). An N/A in the table means the institution does not run any ktev analysis or MC jobs (although they may use the computers for ntuple analysis).

Table 3: Summary of KTeV Computing at Remote Institutions

| Institution | No. of Collab. | Total CPU (MIPs) | Disk (GB) | Tape Drives | % total analysis |
|---|---|---|---|---|---|
| U of Arizona | 3 | 640 (DEC) | 120 | 1 DLT4000 | 2.5 |
| UCLA | 4 | N/A | N/A | N/A | 0 |
| UCSD | 2 | N/A | N/A | N/A | 0 |
| U of Chicago | 17 | 820 (SGI) | 250 | 6 DLT4000 | 1 |
| U of Colorado, Boulder | 6 | 1700 (SGI) | 200 | 5 DLT4000 + 2 DLT4700 | 5.5 |
| Elmhurst College | 1 | N/A | N/A | N/A | 0 |
| Osaka Unversity | 6 | 1200 | 250 | 8 DLT4000 + 2 DLT4700 | 1 |
| Rice Unversity | 2 | ? | ? | ? | 1 |
| Rutgers Unversity | 11 | 3500 (DEC) | 320 | 12 DLT4000 | 5 |
| U of Virginia | 9 | 1400 (DEC) | 150 | 4 DLT4000 | 5 |
| U of Wisconsin | 3 | N/A | N/A | N/A | 0 |

# 4 Central Computing Architecture

Most of the computing for KTeV is done at Fermilab. Table 4 shows the computing at FNAL dedicated to KTeV. The DEC machines are the central machines for offline analysis and the SGI machines are the DAQ processors which have been used for offline analysis (mainly generating MC) since the end of the data taking run.

## 4.1 Data Storage

The original DLT tapes are archived at FCC. A copy has also been made, rearranging the data into a more usable form (splitting by trigger type). This copy is also archived in FCC. Then the split tapes were crunched, using much harder cuts than were made online with final calibrations. The typical crunch reduction in data volume is 90%. Typically an analysis runs from the ”crunch” tapes, or if there is a problem from the ”split” tapes. The original tapes are rarely accessed.

KTeV has allocated 6 TB of data storage space in an IBM STK robot running HPPS. The central analysis machines are connected to the tape robot via FDDI. This space is used for crunch data output, Monte Carlo and calibration data.

Table 4: Summary of KTeV Computing at FNAL

| Flavor | Total CPU (MIPs) | Disk (GB) | Tape Drives | Mass Storage |
|---|---|---|---|---|
| DEC Alpha (UNIX) | 5600 | 1100 | 28 DLT4000 + 4 DLT4700 | 6 TB |
| SGI Challenges | 3900 | 120 | 11 DLT4000 | N/A |

Figure 4.1 shows the central computing at FNAL.

# 5 Networking

# 6 Other Considerations for Distributed Computing

## 6.1 Database needs

KTeV has two databases. One is made by the SlowDA which consists of spill by spill information of the # of protons on target, temperatures and humidities in the calorimeter ”block house”, high voltages, etc.

The other database is made by the Fast DA and consists of calibration constants (made by online then refined offline), run parameters, etc. This database also contains the geometry of the experiment, magnet mappings and other ”stationary” parameters.

The database is maintained on the FNAL central computers and is mirrored via ftp to remote institutions. The databases themselves are a few hundred MB in size, so this is not a huge demand in bandwidth.