VI. OFFLINE COMPUTING - FROM THE DRAFT MINOS MOU

OVERVIEW

MINOS Computing on FermiGrid

This document summarizes MINOS computing needs which have direct implications for processing on the Fermilab Grid, formerly the FBS Batch System on the farms. At the present time (June, 2007) we are able to use only the GPFarm, but development is underway to enable us to expand our operations to FermiGrid, perhaps exclusive of the CMSGrid. We anticipate this development will mature during the fall of 2007. The numbers presented here are drawn from and agree with the MINOS/Fermilab MOU.

We are at this time almost completely using Sequential Access through Metadata (SAM) for data handling, developed for the Run II experiments. The SAM-server is part of the MINOS cluster and has no impact on the grid system other than the need for a client installation which is now in place. We are also using one of the GPGrid head nodes to serve our mysql database and we expect to continue to do so throughout our data processing. During the past year, at the request of the Grid group we moved mysql from a dedicated worker node (where it was placed when the old SGI FGS master was unable to cope) to the head node, and we have noted that this configuration is completely satisfactory. We will continue to use AFS disk for some user data, but by using the /grid/data area for 'staging,' we are now able to access the AFS from non-grid systems where /grid/data is mounted, relieving us of the necessity for kerberized access to AFS directly from the Grid.. The overall needs for hardware are summarized in Table 1. The CPU needs for data processing and Monte Carlo include the requirements for reprocessing which are somewhat uncertain but we have assumed one reprocessing of all prior data per year. <u>We assume all the Monte Carlo generation needs will be met from offsite resources.</u> This is the bulk of the CPU needs for Monte Carlo. The reconstruction takes about 13% of the total CPU and we assume we will do that at Fermilab.

Year	2007	2008	2009
Tape (TB)	81	85	86
Disk (TB)	19	21	22
CPU, Reco (GHz)	150	60	0
CPU, MC (GHz)	1370	2250	3090
CPU MC (GHz) at FNAL	178	293	402
CPU, analysis (GHz)	269	108	108

Table 1 Summary of MINOS hardware needs for the last 3 years of data taking.

MINOS 5 YEAR RUN PLAN

For background, we present our run plan, which is directly proportional to the rate at which data is collected. The experiment has presented a request to the laboratory management for the number of protons on target per year (POT/year). This request is summarized in Table 2.

Year	POT x 10^{20}	Protons per pulse x 10^{13}
2005	1.0	2.5
2006	3	3.5
2007	3	3.5
2008	3	3.5
2009	3	3.5

Table 2 Requested number of protons for the MINOS run plan

The rate and raw data volumes for the Far detector are shown in Table 3 and for the Near Detector in Table 4. The rate from the far detector is dominated by 0.5 Hz of cosmic ray μ interactions, which will be used both for calibration and for cosmic ray and atmospheric neutrino physics studies. The far detector also has about 3.5 Hz of noise triggers, which are small events and contribute to the raw data size but are eliminated after reconstruction. There are also pedestal runs, light injection etc which are listed in the “Other” category. The neutrino interaction rate for the near detector will vary depending on the number of protons on target as will the spill size. The spill size is quoted for the lowest number of protons on target and is scaled for the higher intensities. The near detector DAQ system is capable of recording the full 250 Hz of cosmic rays μ seen by the near detector but it is expected that we will record only a fraction of these calibration purposes, we have assumed 28 Hz, and this is reflected in the numbers in Table 4. Note that we only process 1 in 8 of all the cosmic rays and this is reflected in the CPU needs and reconstructed data storage. The far detector assumes 3×10^7 seconds in one year (cosmic rays are always there) and the near detector assumes an effective year of 1.7×10^7 seconds for beam and 3×10^7 seconds for the cosmic rays. For simplicity we assume that 1Kbyte $\equiv$ 1000 bytes.

Sample	Rate/second (Hz)	Events/year	Raw Event Size (Kbytes)	Data Volume /year (GB)
Cosmic ray μ	0.5	1.5×10^7	1.1	17
Noise	3.5	1.05×10^8	0.2	21
Other				312
Total				350

Table 3 Event rates and raw data volumes for the Far detector

		Raw Data Volume per year (GB)				
Sample	Spill Size (Kbytes)	2005	2006	2007	2008	2009
Beam	17	68	200	200	200	200
Cosmic Rays	0.8	672	672	672	672	672
Total		740	872	872	872	872

Table 4 Raw Data volumes for the Near detector

EVENT RECONSTRUCTION AND STORAGE

The event reconstruction for both detectors is being done at Fermilab. A summary of the processing needs is given in Table 5 for steady state, which will keep up with the data taking, and in Table 6 for reprocessing. These numbers are based on the performance of the existing MINOS C++ reconstruction code. Both detector numbers are taken from real data. The number of neutrino events in the Far detector is small but we apply a generous window around the spill time to make sure that we do not miss any events. This is the Spill Sample line in the Far detector tables. Also we decided to perform a blind analysis of the Far detector data so a fraction of the data is placed in separate files called the Open Spill Sample. The processing time per event or spill will be given in GHz-seconds per event and the CPU requirements will be given in GHz. The Near detector processing times are increased to account for the higher beam intensity over time. We have assumed that in the years 2004 and 2005 we will do one complete reprocessing pass of the Far detector data per year that will be completed in 3 months. For the Near detector we assume one reprocessing per year in 2005 and 2006, again taking 3 months. In 2009 we assume one pass of all the data taking 6 months. The reprocessing numbers include processing for all the data taken up until that time. We assume a farm efficiency of 70%.

		GHz per year					
	Execution time	2004	2005	2006	2007	2008	2009
Beam data (Near)	40 GHz-sec/spill		7	23	23	23	23
Cosmic Rays (Near)	0.34 GHz-sec/event		1.7	1.7	1.7	1.7	1.7
Cosmic Rays (Far)	2 GHz-sec/event	1.4	1.4	1.4	1.4	1.4	1.4
Far Spill	0.6 GHz-sec/spill		0.2	0.2	0.2	0.2	0.2
Total		1.4	10	26	26	26	26

Table 5 Steady state event reconstruction needs for the Near and Far detectors

	GHz per reconstruction pass					
	2004	2005	2006	2007	2008	2009
Beam Data (Near)		15	76	213	259	175
Cosmic Rays (Near)		3.4	10	17	24	15
Cosmic Rays (Far)	15	21	27	33	38	22
Far Spill		0.4	1	2	3	2
Total	15	40	114	265	324	214

Table 6 Reprocessing needs for Near and Far detectors

MONTE CARLO PROCESSING

There are two types of Monte Carlo required for MINOS, simulation of neutrino interactions in the detector for oscillation measurements/conventional neutrino physics and simulation of the neutrino beam to understand features of the beam such as beam profiles, flux etc. In both cases the requirements are not precisely known so the numbers here are based on assumptions. We give the total requirements here but about 80% of CPU required for the Monte Carlo is provided by offsite resources at collaborating institutions. The files are transferred to Fermilab for reconstruction (which is 20% of the CPU) and storage.

Physics Monte Carlo

For studies of cosmic ray and atmospheric neutrino events in the Far detector we assume a Monte Carlo sample equal to the data. For the Near detector we have made a similar assumption. The event size is larger due to storage of the “truth” information” for the event. The needs per year are summarized in Table 11.

	Year					
	2004	2005	2006	2007	2008	2009
Far detector Cosmic Rays						
Events ($\times 10^7$)	1.5	1.5	1.5	1.5	1.5	1.5
CPU time/event (GHz-sec)	16	16	16	16	16	16
CPU for 2 months (GHz)	70	70	70	70	70	70
Far detector Beam						
Events ($\times 10^6$)	2	2	2	2	2	2
CPU time/event (GHz-sec)	16	16	16	16	16	16
CPU for 1 month (GHz)	17	17	17	17	17	17
Near detector Beam						
Spills ($\times 10^6$)		4	8.5	8.5	8.5	8.5
CPU time/spill (GHz-sec)		264	264	264	264	264
CPU for 6 months (GHz)		97	292	292	292	292

Table 11 Physics Monte Carlo off-site generation needs per year

We expect to reprocess the Monte Carlo data with the same frequency as the data. These requirements are shown in Table 12 for CPU .

	GHz per reconstruction pass					
	2004	2005	2006	2007	2008	2009
Near Beam MC		145	567	1850	4000	7000
Far Cosmic MC	34	101	169	237	304	372
Far Beam MC	9	26	44	62	79	97
Total	43	272	780	2149	4383	7469

Table 12 CPU for reprocessing for the Far and Near Monte Carlo

SUMMARY

Year	2007	2008	2009
CPU, Reco steady state (GHz)	0	0	0
CPU, Reco reproc (GHz)	151	600	0
CPU, MC (GHz)	0	0	0
CPU, MC reproc (GHz)	1370	2250	3090
CPU Speed (GHz)	4	4.6	5
Duals, Reco steady state	0	0	0
Duals, Reco reproc	19	7	0
Duals MC reproc	190	313	309

Table 21 MINOS needs for processing CPU for the last 3 years of data taking.

Table 21 shows the resources required per year for the last 3 years of data taking. This table shows the actual equipment that must be purchased each year to meet the requirements described in the previous sections. This takes into account equipment purchased in previous years. It also assumes a 3 year replacement cycle for CPU. The CPU for reprocessing is the amount required over a 3 month period for the years 2007-2008 and a 6 month period in 2009. Changing the assumptions about reprocessing will change these amounts. Table 22 shows the costs associated with the needs in Table 21. The following assumptions have been made. A dual CPU farm node costs $2200 independent of CPU speed. The costs for the Monte Carlo CPU are 20% of the total requirement as explained above.

Year	2007	2008	2009
Duals Reco steady state	14000		
Duals Reco reproc			94000
Duals MC reproc			18000
Duals, analysis	64600	22300	21000
Total	78600	22300	133000

Table 22 MINOS costs in $ per year for the last 3 years of data taking