Computing for CLEO III

The CLEO III Computing Committee for the Collaboration

December 11, 2000

1 Introduction

The CLEO III detector is now collecting physics-quality data and CESR has achieved B-factory luminosity of 1.2×10^{33} cm^{-2}s^{-1}. To meet the challenge of extracting physics from this data in a timely way, the CLEO collaboration and the Laboratory of Nuclear Studies must acquire and operate computing systems that are far more extensive than those that were used for CLEO II. This document describes the collaboration's computing plan, and provides quantitative assessments of the resources required for all aspects of CLEO III data processing and analysis.

The computing tasks for CLEO III fall into four broad categories:

1. On-line data acquisition, experiment control and data-quality monitoring.
2. Off-line reconstruction, including precise calibration of all subsystems and pattern-recognition tasks (track reconstruction, calorimeter cluster finding, RICH reconstruction, lepton identification, etc.).
3. Generation of simulated data samples (Monte Carlo) for algorithm design, efficiency determinations and background studies. Both generic and analysis-specific simulations are required.
4. End-user analysis tasks for specific physics studies. This includes the "skimming" of subsamples of events (data and Monte Carlo), the processing of these subsamples, and the extraction of results with PAW and other tools.

The first category, on-line computing, has been addressed through the CLEO III detector project. The system is currently functioning very well. The data acquisition system has been built with principal support from the DOE, and it is anticipated that there will be some continuing need for support to maintain and upgrade this system. On-line processing and data storage facilities will also require upgrades, but the requirements will be modest. LNS has principal responsibility for these on-line computing systems.

During the construction of the CLEO III detector the bulk of our manpower was dedicated to hardware tasks. A core group of physicists and software professionals were responsible for the development of the framework for CLEO III computing. With very limited human resources, they were able to deliver a system that was nearly complete by the time construction was finished. Detector-specific components have subsequently been developed by susbsystem experts freed up by the end of construction, so that CLEO III now has a functioning and nearly complete off-line reconstruction package running on data as it is collected. The Monte Carlo simulation framework is ready and nearly all subsystem simulations have been demonstrated.

During the CLEO III software development period, limited computing facilities for CLEO III off-line tasks were acquired with LNS operating funds. With the onset of data taking, aggressive steps were taken to meet the initial needs of routine data reduction. This has been a great success. The system will continue to evolve in the future, and this document provides an assessment of the overall requirements.

CLEO is a strong and very experienced collaboration, with human and technical resources both on site at LNS and at collaborators' campus locations. Success for CLEO III demands that all of these resources be tapped. This will require a robust, stable and well-supported framework for CLEO III analysis. It must provide both on-site and off-site users with access to reconstructed events and Monte Carlo samples, as well as the necessary support to accomplish analysis tasks locally or with central facilities. There must be coordination in the selection of hardware, including computer configurations and storage systems.

We have developed a model for how computing for CLEO III will be done. The central repository for data and expertise must be at LNS, maintained by a combination of CLEO members and laboratory support staff. LNS facilities must provide for all on-line and routine off-line data-processing tasks, as well as much of the analysis capacity. The entire data sample will be "spinning" on disk at LNS at all times, providing local and remote users with an extremely efficient and stable analysis environment. It follows from this description that the CLEO III investment in computing hardware and software will be heavily weighted toward LNS.

The one major component of CLEO III computing that does not need to be lab-centered is the generation of Monte Carlo samples. CLEO II has very successfully used large computer farms at the University of Florida and the University of Minnesota, and smaller farms at several other sites, to meet the great majority of Monte Carlo needs. Samples of several hundred million fully GEANT-simulated events have been produced. For CLEO III, it is essential to continue in this mode of operation, with farms for generating both "generic" and analysis-specific "signal" Monte Carlo samples.

Analysis tasks will be carried out both at LNS and at other CLEO institutions. Since complete CLEO III data and Monte Carlo samples will exist only at Cornell, LNS must provide sufficient CPU and storage capacity for all collaborators' skim jobs. Once in hand, these smaller skim samples may be processed on central facilities, on desktop systems or on other computer systems at the home site.

CLEO groups carrying out analyses at their home institutions have several options for the level of local CLEO III computer support they would like to provide. Monte Carlo farm sites and other large groups will require full CLEO III analysis capabilities, and will require staff and facilities to support this. Some other groups may elect only partial support to reduce their overhead, providing adequate facilities to process skims locally, while not supporting full analysis or Monte Carlo. Small groups, or those whose personnel are mostly stationed at Cornell, may prefer to concentrate analysis activities and facilities there, with at-home usage limited to windowing into LNS and the interactive final stages of analysis.

The remainder of this report is organized as follows: Section 2 describes the assumptions we have made in assessing CLEO III computing needs. Section 3 describes the specific computing requirements for Cornell LNS. Plans and specifications for CLEO III simulation farms are detailed in Section 4. The analysis needs of CLEO III collaborators, for both on-site and off-site personnel, are described in Section 5. Section 6 summarizes the plan and CLEO III needs.

2 Assumptions: Data Samples and Computer Performance

The original plan for CLEO III envisioned that all running would be at energies on and just below the $\Upsilon(4S)$ resonance for intensive studies of b quarks. The impressive performance of the PEP-II and KEK-B B-factories has changed the terrain for b studies. The CLEO collaboration and Cornell LNS are now assessing their physics potential both at the $\Upsilon(4S)$ and for CESR operation at other energies. Possibilities include the narrow upsilon resonances and the region just above charm threshold. Clearly such a change of plan would affect the details of our computer needs. While the basic framework for the readout and analysis of the experiment and the organization of analysis and simulations would be unaltered, the details of trigger rates, event size and characteristics, and backgrounds would depend on accelerator running conditions and specific physics processes studied. We have chosen to base our initial estimates in this document on the original $\Upsilon(4S)$ program, recognizing that reassessment will be required as plans evolve. It is essential that we proceed with acquisitions of facilities to ensure that CLEO III data is analyzed in a timely and effective manner.

To estimate CLEO III computing needs we must begin with estimates of the CLEO III data sample. The time frame for our analysis is the next three LNS fiscal years (November-to-November). As of the start of CLEO III physics data last summer, CESR reported a peak luminosity of approximately 0.7×10^{33}. An aggressive luminosity improvement program and a correction to the normalization of the CESR luminosity measurement, have resulted in a current peak luminosity of 1.2×10^{33}. Further improvements are anticipated in 2000-1, with no major accelerator upgrades or downtime required. The addition of more superconducting RF cavities and other improvements are expected to give 2.2×10^{33} and 3.0×10^{33} in 2001-2 and 2002-3, respectively. The conversion of instantaneous luminosity to expected integrated luminosity is based on a 50%

effective efficiency over the entire year, consistent with past CESR experience. The full three-year sample is thus estimated to be about 120 fb^{-1}.

Running at the $\Upsilon(4S)$ with these luminosities, the CLEO III rate of collecting Bhabha-scattering and hadronic events will increase from 5 Hz now to about 30 Hz in 2002-3. (This is the rate to tape. Including events taken for calibration, the trigger rate is approximately 30-40 Hz.) The read-out raw-data size of a hadronic event is currently 80 kB. As sparsification and other techniques mature, this is expected to decrease to 50 kB. Likewise the analyzed data size for a hadronic event is expected to decrease from 100 kB to 50 kB. Monte Carlo events will require this 100 kB for raw and analyzed information, plus an additional 10% for Monte Carlo-specific information. The standard CLEO practice, based on a great deal of relevant experience, is to generate generic Monte Carlo samples ($B\bar{B}$, continuum) with five times as many events as we have in our data. This has proved essential for understanding backgrounds to rare B decays and other processes.

To estimate CPU requirements, we start from these assumptions and add our predictions of the performance and price of currently available computer systems. CLEO III code is standardized on Sun/SPARC computers running the Solaris operating system. Because of our large installed base from the CLEO II era, we also support Compaq/Alpha processors running the Tru64 operating system. Since neither of these systems provides the best available price/performance, and for compatibility with other HEP experiments, we also plan to support Intel x86 (and compatible) processors running Linux. It is expected that most collaborators will adopt Intel/Linux from the start and that LNS will migrate from Sun to Intel/Linux by 2001-2.

The current CLEO III hadronic event processing time (on a 440 MHz Sun Netra) is 5.0 s. This is expected to decrease to 3.0 s as software matures in the next year. Beyond that we allow for an increase by 30% per year as capabilities are added.

The speed of available processors will also increase. We expect a 30% boost in the performance of our Sun platform in the next year, with the price stable at about $3K. The currently available Intel/Linux systems with the best price/performance characteristics (800 MHz, 256 MB memory, 20 GB disk) are estimated to be ~ 1.4 times faster than a Netra/440 for a bit less than half the cost. By 2001-2 we conservatively expect another 30% improvement in speed with no increase in cost.

On-line storage of CLEO III data will require very large hard disk capacities and automated tape storage. This area is particularly problematic since there are considerable technical uncertainties. SCSI disk prices have fallen consistently in past years, and we expect this trend to continue. We assume a reduction from about $17 per GB now (75 GB disks) to $4.2 per GB in 2002-3 (170 GB). RAID systems using inexpensive IDE disks are a recent innovation. These are a viable alternative for our Intel/Linux systems. IDE disks currently offer a cost advantage relative to SCSI of at least a factor of 3.

By 2003, the CLEO III data and Monte Carlo samples will have grown to more than 500 TB, and there is no alternative to high-density magnetic tapes and sophisticated robot libraries with hierarchical storage management (HSM) software. In contrast with disks, the rate of technical progress and cost reduction for tape storage has been very slow. With the advent of multiple new technologies, it is hoped that this situation will change. CLEO III is currently using AIT-2 tapes (50 GB), and we expect to move next year to AIT-3 (100 GB) Beyond that we anticipate switching to the new LTO-1 (100 GB) and LTO-2 (200 GB) technology. As is described in Section 3, tape drives, robots and media represent by far the largest expense in the CLEO III computing plan.

3 Computing Requirements for Cornell LNS

During the next year, CLEO II data analysis will consume virtually all of the Tru64 Alpha computing capacity at Cornell. Also, during this period it is unrealistic to assume that the full package of analysis software will have been converted to run on any currently unsupported platform. Therefore the initial CLEO III computing plan assumes that we use Sun computers running under Solaris. Special pricing and an expected upgrade of our standard Sun Netra T1 platform from 440 MHz to 600 MHz make this a viable strategy through 2001. As described in Section 2, however, by 2001-2 we plan to have migrated our off-line software to more cost-effective Intel (or compatible) processors running Linux.

The projections described in this report are based primarily on scaling arguments applied to our early experience with the CLEO III software.

Tables of projected needs and costs are provided for each of the several required facilities. Entries for numbers additional nodes, tape libraries and other components represent the incremental investment from the installed base at the beginning of each LNS grant year. Future capacity requirements are given in units (CPUs, tape libraries) that are our current standards. The projected acquisitions, however, are based on the more powerful and cost-effective products that are expected to be available at the time of purchase.

FULL RECONSTRUCTION PROCESSING FACILITY - Table 1 shows the projected needs for processing nodes, disk space and servers to reconstruct CLEO III raw data. The indicated processor capacity allows us to keep pace with ongoing data collection (under the luminosity projections of Section 2), and to reanalyze all previous years' CLEO III data at a rate of once per year.

As faster production nodes become available, it will be possible to meet our needs with smaller numbers of processors. Thus we anticipate adding only about 20 nodes each year, resulting in a total of fewer than 80 nodes by November, 2003.

The associated costs of short-term disk storage are modest. We start with 1 TB of disk cache, and provide storage for one hour's analyzed events for each node added to the facility. Efficient use of this increasingly complex system will require refinement of our management and scheduling procedures.

Table 1: Full reconstruction processing facility.

Item	Year		
	2000-1	2001-2	2002-3
Processor capacity (Netra/440's) for new data	23	40	65
Processor capacity (Netra/440's) for reprocessing	16	40	80
Additional processor nodes required	18	19	23
Cost of additional processor nodes	\$58K	\$33K	\$38K
Disk cache (1 hour's data per node +30%) (GB)	80	101	154
Cost of additional disk cache	\$1K	\$1K	\$1K
Cost of additional server(s) (\$3K/server)	\$3K	\$3K	\$3K
Total cost	\$69K	\$42K	\$54K

CALIBRATION FACILITY - Table 2 shows the projected needs and plan for the CLEO III Calibration Facility. To avoid significant delay in analyzing data, we are taking an aggressive approach to providing dedicated calibration facilities. Since only a subsample of the data is analyzed, and the analysis is an abbreviated version of full reconstruction, we project that only one fourth of the processing power of the full reconstruction facility is needed. We have already acquired 1.0 TB of disk space for the initial implementation of this facility. With increasing luminosity, we will add sufficient disk to accommodate three weeks of Bhabha events and a tenth of the hadronic events. The total cost per year to maintain the calibration facility is no more than \$20K.

Table 2: Calibration Facility.

Item	Year		
	2000-1	2001-2	2002-3
Processor capacity (Netra/440's)	6	10	16
Additional production nodes	5	2	3
Cost of additional production nodes	\$15K	\$3K	\$4K
Calibration disk (GB) (bhabhas + hadronic)	1269	1345	1459
Additional disk (GB)	269	75	115
Cost of additional disk	\$3K	\$1K	\$0
Total cost	\$17K	\$4K	\$4K

MASS STORAGE - Because the amount of data to be stored is so large, mass storage is the most challenging and costly part of the off-line computing facility. It would be completely impractical to analyze CLEO III data by manually loading tapes, as we have done for CLEO II. Our plan is to utilize robotic tape libraries to store all raw and reconstructed data plus reconstructed Monte Carlo events, making them accessible through server nodes and gigabit ethernet to the rest of the computing facility. Table 3 shows the amount of data, the size and number of robots required, and the costs.

The specifics are heavily dependent on tape technology which is advancing slowly. As described in Section 2, we are using 50 GB AIT-2 tape now and anticipate adding robots using 100 GB AIT-3 drives during 2001-2 and 2001-2. For 2002-3, we base our projection on 200 GB LTO-2 technology. This is an aggressive introduction of new technology to keep costs down, assuming a schedule for roll-out of higher density tapes and drives that we hope is realistic. AIT-3 has been demonstrated already, and uses the same media as AIT-2, so this is a safe next step. Switching to LTO (or the similar Super DLT) in 2002-3 is based on our expectation that very high capacities will be more readily achievable with LTO (or S-DLT) because of the four-times-larger tape area. Since the tape robot libraries operate largely independently of each other, there is no obstacle to adopting the most cost-effective technology at the time we add each new library.

Table 3: Mass storage.

Item	Year		
	2000-1	2001-2	2002-3
Additional off-line data storage (TB)	31	36	46
Additional reconstructed Monte Carlo event storage (TB)	40	52	85
Total additional mass storage needs (TB)	70	88	131
Cost of tape for raw Monte Carlo events	$28K	$24K	$24K
Media capacity (GB/tape)	100	100	200
Tapes per robot library	600	600	1000
Capacity of robot library (TB)	60	100	200
Number of robot libraries needed	1.2	1.5	0.7
Cost per robot library	$180K	$140K	$250K
Cost of media per library	$60K	$42K	$100K
Cost of HSM software per library	$6K	$6K	$6K
Cost of two servers per library	$6K	$6K	$6K
Cache disk per library (GB)	800	1000	1500
Cost of cache disk and network per library	$14K	$11K	$10K
Cost of fully equipped robot library	$266K	$205K	$372K
Total cost	$339K	$323K	$268K

SKIM AND GENERIC MONTE CARLO ANALYSIS FACILITIES - As described in Section 5, the collaboration will make roughly 100 passes through the entire data set per year to make skims for physics analyses. This will only be practical if the data sample to be skimmed is resident on disk, and we must therefore strictly limit the size of the events made available for analysis. We plan to provide a disk-based skim facility at LNS for use by the entire collaboration. Table 4 shows the expected requirements for compute nodes and SCSI disk space for three years of CLEO III operation. As for the full reconstruction facility, compute nodes are Sun Netras at first and Intel/Linux in subsequent years. In projecting disk needs, we have estimated that only 7 KB per event will be needed for hadronic and tau skims.

The generic Monte Carlo analysis facility (Table 5) is modeled after the skim facility. We assume the same event compression, but scale by a factor of five Monte Carlo events for each real event.

Table 4: Skim Analysis Facility.

Item	Year		
	2000-1	2001-2	2002-3
Additional processor capacity (Netra/440's)	10	12	18
Cost of additional processor nodes	\$25K	\$10K	\$11K
Additional data skim disk (TB)	2.0	3.0	4.2
Cost of additional data skim disk	\$21K	\$20K	\$18K
Total cost	\$46K	\$30K	\$28K

Table 5: Generic Monte Carlo Analysis Facility

Item	Year		
	2000-1	2001-2	2002-3
Additional processor capacity (Netra/440's)	20	24	36
Cost of additional processor nodes	\$51K	\$19K	\$21K
Additional Monte Carlo skim disk (TB)	7.8	12.2	16.6
Cost of additional M. C. skim disk	\$83K	\$82K	\$70K
Total cost	\$133K	\$101K	\$91K

4 Monte Carlo Farms for CLEO III

The CLEO II experience has repeatedly demonstrated that very large simulated event samples are essential for doing precision physics. Understanding backgrounds requires very large samples of "generic" events, with generators that are continually tested and tuned with data, and with full GEANT simulation. For CLEO II, simulation farms at Florida and Minnesota have produced generic $B\bar{B}$ and continuum ($udsc$) Monte Carlo samples of more than 200 million events. These samples have been used in nearly all CLEO analyses.

Specific analyses also rely on dedicated samples of simulated signal events for design of algorithms, determination of efficiencies and assessment of systematic errors. While "reweighting" and other economies allow the sample size for most analyses to be kept fairly modest (from tens of thousands to several million), there are many analyses. Based on a survey of CLEO II groups we have estimated that the sum of all generated special-purpose samples has been in excess of 100 million events. These events, comparable to either the $B\bar{B}$ or continuum generic samples, have been produced during idle periods with the Florida and Minnesota farms, on smaller facilities at Ohio State, Caltech, Syracuse, and elsewhere, and by "stealing" cycles on desktop machines.

We base our CLEO III requirements on the CLEO II experience and the assumptions in Section 2. An important lesson has been that running a simulations farm is a major continuing responsibility, requiring high-level support of software, coordination with analysis management, set-up of farm runs, operator tasks during runs, and transfer of event samples to LNS. Operation of each major farm represents an ongoing "service task" for one graduate student, as well as requiring supervisory oversight and staff support.

Under the assumption of $\Upsilon(4S)$ running, we base CLEO III generic Monte Carlo needs on roughly equal-sized $B\bar{B}$ and continuum samples. Simulating five times the expected $B\bar{B}$ sample in three years of running (120 fb^{-1}) requires 600 million events. At 5 (Netra) seconds per event for generation and 3 seconds per event for reconstruction, this demands 4.8 billion CPU seconds. To complete the simulation in three years demands a minimum of 51 CPU's.

To meet CLEO III needs for generic and signal Monte Carlo samples, including tau and two-photon physics in addition to continuum and $B\bar{B}$, we propose to establish two generic production facilities and three signal Monte Carlo farms. The latter would be considerably more labor-intensive to operate on a per-event basis, because they would deal with smaller samples and would need to coordinate closely with individual consumers. Florida and Minnesota have agreed to continue to operate the generic farms, while faculty at Southern Methodist University, the University of Pittsburgh and Carnegie Mellon University have proposed to operate signal farms. We expect that the farm facilities would be funded by DOE or NSF, with significant matching contributions from the host universities.

While specific hardware choices will be deferred to ensure optimal price/performance at the time of acquisition, we have sketched out a prototype generic Monte Carlo farm in Table 6. This two-year phase-in would be appropriate for either Florida or Minnesota, and could appropriately be halved for any of the signal farms. The minimal requirement calculated above has been escalated by a factor of 1.7, primarily to account for inefficiencies in operation and regeneration of early samples as the simulation is refined.

Table 6: Example generic Monte Carlo farm for CLEO III.

	2000-1	2001-2
Compute Nodes		
CPU	800MHz PIII or Athlon	1200MHz PIII or Athlon
Memory	256MB	256MB
Disk	45 GB IDE	45 GB IDE
OS	Linux	Linux
Unit Cost	$1.4K	$1.4K
Number of Computers	25 (35 Netra equivalents)	25 (53 Netra equivalents)
Total Cost	$35K	$35K
Miscellaneous Utilities		
Staging Disks	150 GB	300 GB
Disk Cost	$1K	$1K
KVM	$2K	$2K
Switch w. Gbit/s Uplink	$3K	$3K
Tot. Misc. Util.	$6K	$6K
Tape Storage		
AIT-2 drive w. 20-tape robot	$10K	$10K
Number of AIT-2 tapes	100	100
Cost of AIT-2 tapes	$5K	$5K
Tot. Tape Storage	$15K	$15K
Total cost	$56K	$56K

The mass-storage needs for Monte Carlo farms are very minor compared to those of the central archive, which were described in Section 3. The compute nodes as configured above include adequate local storage during generation and reconstruction. Generated event files will be congealed on a node equipped with somewhat more staging space and written to tape for transport to LNS, where they will be archived on the Monte Carlo Analysis Facility described in Section 3. The choice of tape technology is driven by compatibility with LNS. Transport of Monte Carlo event samples by tape can be accomplished even with the CLEO III first-generation AIT-2 drives. While 500 to 1000 tapes would be written each year, all tapes would subsequently be recopied at LNS. Therefore a relatively small tape supply could be continually recycled for the farms. Each farm will require a modest robot to allow unattended running over a period of a few days. This plan is reflected in Table 6.

Farm operation must be highly automated and fault-tolerant. For CLEO II each farm developed its own farm-management scripts, and for CLEO III we would like to develop a standard package to be used by all farms. Minnesota is working on a farm-management system based on the Condor package from the University of Wisconsin.

5 Analysis Computing Overview

CLEO historically has published 20-30 analyses per year, a rate of physics productivity that we expect to continue with CLEO III. This requires perhaps 75-100 full passes over the data per year, including development of analysis procedures, selection of specialized event skims, cross-checks of analyses, systematic studies, analyses not leading to publication, and mistakes. These passes of the full data sample, as well as of large generic Monte Carlo sets, will be carried out primarily on central facilities at LNS.

The usual procedure is to reduce the large signal and background samples first to skims and then to "Ntuples," a very efficient format for final-stage analysis. The reduced data samples are typically transferred from the central facility to individual users' dedicated computers, in offices at LNS or at home institutions.

While these dedicated analysis machines need not be as powerful as the central systems, the physicist's effectiveness is greatly enhanced by exclusive access to a modern workstation, with adequate memory, a large amount of local disk storage, and a high-bandwidth network connection. Powerful graphics accelerators and large high-resolution displays can be important enhancements for some analyses. These needs are equally applicable whether the physicist performing the analysis is located at LNS or at a remote site.

5.1 Analysis at Remote Sites

5.1.1 Full Analysis Requirements

Full analysis implies the ability to perform any CLEO III software task. In practice, the difficult ones involve reconstruction of data and require full access to all associated constants databases. While reconstruction of real physics data will take place only at LNS, reconstruction of simulated events will take place at several sites. Reconstruction of Monte Carlo is the prime reason to support full analysis capabilities away from LNS.

It is anticipated that most sites supporting full analysis will do so to support Monte Carlo farms. Reconstruction requires a similar amount of CPU to the generation of events; thus it is required that the farms do both, delivering fully processed Monte Carlo events to LNS for collaboration use.

Individual CLEO groups at their home university sites have very limited personnel resources. It is therefore essential that CLEO III software maintenance require substantially less than one full-time equivalent staff member. It must also be compatible with expertise that is typical of a physicist rather than of a computing professional. As a result of these constraints, there will be an unavoidable burden on the collaboration and the lab to maintain centrally a user-friendly and robust software infrastructure and distribution system. This must be developed with continual input from the outside groups, and those groups with access to professional computing support can play an especially valuable role.

Issues concerning licensing of the Objectivity database software to user groups have been successfully resolved by the computing staff at LNS. Objectivity data base maintenance will be a nontrivial component of the support of production facilities, requiring both manpower and dedicated computer nodes.

5.1.2 Partial Analysis Requirements

By partial analysis, we refer to the capability to transfer reconstructed events, for example skims produce for a particular analysis, and do further analysis work on them remotely. We plan to support a self-contained binary data format that is independent of the Objectivity database structure. It will also be independent of the constants databases. Since reconstruction is complete, access to calibration constants is not necessary for the analysis of these reduced data samples. Other database information external to events, such as beam spots, can be subsumed into the binary version of begin-run events.

We expect many sites to operate in this partial analysis mode. Since this environment does *not* allow for reconstruction of simulated events, users at these sites will rely on the signal Monte Carlo farms described in Section 4. Close coordination with the signal Monte Carlo farms will be required.

Ntuple-based analyses are by their nature portable. All information deemed to be necessary for the analysis is encoded in the Ntuple at the time of production. Invariably in the course of an Ntuple-driven analysis, new ideas, reviewer questions or discovery of bugs or omitted information will require additional information from the event database. That follow-up access can be driven very efficiently by event lists generated from the Ntuple, accessing only the required event information. While this necessitates running again at LNS, it is far more efficient than starting over from scratch.

5.2 Analysis at LNS

5.2.1 LNS support

One of the main ways to increase efficiency of physics analysis involves providing high-level event tagging information for public use. For example, most, if not all, users selecting leptons can agree on common criteria in advance, and the database can tag all events meeting such criteria. This allows the first stage of event skimming to be accomplished at the time of data reconstruction. Work is under way to identify the most useful classes of data for such tagging. The same can be done with the disk-resident Monte Carlo.

By providing access to all real and simulated data on disk at LNS, we expect to essentially eliminate individual users' needs for tapes and tape drives. This requires both sizable temporary disk on the central analysis farm and large network bandwidth to write information (skims, Ntuples) to user disks at LNS and remotely. The increasing capacity and falling cost of available disk drives makes it feasible to provide very large dedicated local storage on desktop machines.

5.2.2 Outside contributions

University groups generally have personnel both at LNS and the home institution. It will continue to be the responsibility of CLEO groups to provide dedicated facilities for their members who are pursuing CLEO analyses. As mentioned above, effective analysis requires one modern CPU with large local disk capacity ($\sim$ a few hundred gigabytes) for each analysis.

Our experience is that users are practical when it comes to sharing resources among themselves at LNS. Thus, the hardware located at the lab often benefits others, and is used very effectively. Having enough resources for everyone is crucial to maintaining a large, quality physics output.

6 Summary and Conclusion

The computing plan described in this document represents a realistic and comprehensive blueprint for all phases of CLEO III data processing and analysis. It is not cheap, with Cornell LNS investing more than three million dollars in equipment and software, including 3000 tapes' worth of robotic libraries. Although cost-reduction strategies will be sought, the data storage needs in particular appear very difficult to reduce. This very large investment in central computing facilities will be essential for the success of CLEO III.

While the effort and costs associated with remote Monte Carlo farms and analysis computing for individual groups are much smaller than those centered at LNS, they are critical for the success of the experiment. Individual groups will apply to their funding agencies (DOE or NSF), as well as to matching support from their universities, for these facilities. We preliminarily estimate agency requests of approximately \$200K for support of CLEO III Monte Carlo farms. For analysis, we anticipate costs of approximately \$3K for each physicist actively engaged in analysis.

Because CLEO III is collecting data now, we need to move quickly into the next phase of LNS acquisitions and farm proposals. Specific proposals will be forthcoming soon, with implementation anticipated next year.