

Review Report
of the
NSF Review Panel
for the
Laser Interferometer Gravitational Wave Observatory (LIGO)

**Performed for the
National Science Foundation**

**Conducted at Caltech in Pasadena, CA
on April 30 – May 2, 2001**

**Date of this Report
May 21, 2001**

LIGO-M010352-00-M

Table of Contents

Executive Summary	1
1. LIGO Data Analysis System Hardware Procurement	4
2. LIGO Data Analysis System Software Development and Preparations for Data Analysis with the LIGO Scientific Collaboration	8
3. Detector Installation and Commissioning	12

Appendices

A. Charge to the NSF Technical Review Panel	15
B. NSF Review Panel Membership	16

Executive Summary

The regular annual review of the Laser Interferometer Gravitational Wave Observatory (LIGO) was conducted by a review panel of experts on behalf of the National Science Foundation on April 30 – May 2, 2001 at the California Institute of Technology in Pasadena, CA. The focus of this review was on the technical aspects of the LIGO Project. Progress on detector installation and commissioning was reviewed along with the proposed plan for completion of the commissioning and the beginning of the science run for LIGO. The proposed procurement plan for the hardware necessary for the analysis of data from LIGO along with the proposed plan for software development and personnel support for the data archiving and analysis system were also reviewed. The NSF charge to the review panel is provided in Appendix A of this report. Members of the review panel and the observers from the NSF are listed in Appendix B.

Written material was provided by LIGO to the panel in advance of the meeting by setting up an access path via the web to the documents. The panel found this method convenient and would encourage its use in future reviews. The available documents were examined, oral presentations were heard and subgroups of the panel met with appropriate members of the LIGO team to explore details of the project. The format of the review followed the one used during recent reviews with short overview sessions followed by expanded breakout sessions with individual groups for in depth discussions. This arrangement provides a very effective format for the review of the project. Based on these evaluations, the panel discussed its findings in executive session and generated the written summary conclusions and observations given below. Details of the assessment of the LIGO project status are given in the full text of this report.

The LIGO Review Panel formulated the following summary observations:

- Final installation and commissioning of the detector systems at both the Hanford Observatory and the Livingston Observatory are going very well. Robust lock of the full configuration of the 2 kilometer interferometer at the Hanford Observatory has been achieved and first lock of the full configuration of the 4 kilometer interferometer at the Livingston Observatory is expected within two weeks.
- The LIGO team is to be commended for its effective response to many of the panel concerns expressed last year. The panel was favorably impressed with the activities in hardware testing, purchasing strategies and establishing an independent test environment for test data analysis.
- The LIGO team is to be congratulated on the rapid commissioning and initial operation of the 2 kilometer interferometer at Hanford. The application of the end to end model by Matt Evans to develop the lock sequence for the interferometer was essential to achieving lock. The initial attempts based on previous experience with the 40 meter facility and the locking of the input Michelson interferometer

were unsuccessful owing to the complexity of the control system. It is particularly satisfying to see that the environment created by LIGO enables an outstanding graduate student to carry out such an important development.

- The panel notes the important contributions of the LIGO Scientific Collaboration in the commissioning of the interferometers. They are developing software necessary for assessing the performance of the interferometers and members of the LSC are manning shifts during the engineering runs. The progress since the last review is impressive.
- In order to transition to science running as soon as possible, the panel encourages LIGO to continue to prioritize running time on the interferometers during the commissioning phase to increase their sensitivity and the reliability of their systems while providing early data samples to the LSC during the engineering runs.
- In the current environment of electronic communications, the issue of computer system security is an increasingly important and difficult problem. While it is desirable in a distributed collaboration to maintain as open a system as possible, this must be balanced against the risk of malicious damage to both the control system and the data stream. The panel recommends that a document be prepared within six months and presented to the NSF that assesses the risks and defines appropriate means to mitigate these risks. Among the items that need to be addressed are
 - Multiple, independently secure copies of data and metadata
 - Facilities augmentations at CACR and other sites appropriate to the needs of the LIGO production facility
 - Network vulnerability, data and system integrity, and robustness under partial system failures
 - Independent development, test and production systems
 - Protection against hackers, especially for critical systems.
- The panel recommends augmented efforts on the part of both LDAS and other members of LIGO in obtaining reasonably priced high-speed network connections between the Hanford, Louisiana and Pasadena sites in a timely manner. LIGO may benefit by seeking advice from the NSF and local universities.
- The panel recommends additional effort by the leadership and staff of LDAS in building direct working relationships with other groups and organizations in the high-performance computing, communications and data storage and retrieval communities including the HPSS consortium.

- The panel recommends that professional database administration be guaranteed by a dedicated person. A plan should be developed describing redundancy and how to perform hot backups, integrity checks and disaster recovery procedures and all of these should be regularly tested.
- The panel encourages the LSC to amplify its efforts in making collaborative connections with the astrophysical community. The makeup of the LSC Upper Limit Teams is a good beginning. Participation in topical meetings and workshops might well serve this function. This will be beneficial to the lab and will also keep the astrophysical community aware of the progress of LIGO.

In summary, LIGO is 97% complete. The LIGO team continues to make excellent progress towards the completion of commissioning of the three interferometers by December 2001. The LIGO Laboratory organizational structure has been changed from that of a construction project to that of an operational scientific laboratory. The Panel continues to be confident that the LIGO construction project will be completed within the project budget and the first LIGO science run will begin in early 2002.

1. LIGO Data Analysis System Hardware Procurement

Findings

LIGO staff reported on the present state of the hardware procurement for the LIGO Data Analysis System. Their progress is summarized below.

- The Phase I procurement was approved in February 2001 and pursuant to that plan they have purchased or are in the process of purchasing:
 - StorageTek 9310 Powderhorn Tape Silo(refurbished):
Purchased and installed
 - Beowulf clusters:
 - 18 Dell P4s for Hanford
 - 18 HP AMD K7s for Livingston
 - 4 Dell dual-processor SMPs for cluster masters (2 at each site)
 - 4 Dell quad-processor SMPs for data conditioning (2 at each site)
 - Sun Microsystems T3 RAID systems for all LDAS sites
28.3 TB total, divided among Caltech, MIT, Hanford, and Livingston.
- There is a detailed plan for Phase II and Phase III procurements to
 - Complete the HPSS data archive system by adding servers and tape drives and obtaining the necessary supply of tape media, plus a large ethernet switch
 - HPSS STK 9940 Tape drives
 - HPSS data movers: 1 metadata server and 2-4 data movers
 - Sun Solaris 8 port of HPSS
 - Beowulf computer clusters upgrades with additional nodes, data conditioning servers and Ethernet networking
 - Low-cost IDE/SCSI RAID at Caltech for science users
 - Peripherals and accessories.
- That plan has been determined through
 - Substantial functionality, performance and reliability testing
 - Determination of needs
 - Consultation with experts external to LIGO
 - Price negotiations.
- Investigations are underway to provide network connections to each of the sites at bandwidths sufficient to relay the full data stream from each of the interferometers to the data archive at CACR.

Comments

HARDWARE SYSTEMS

Supported platforms

The panel notes that LIGO has addressed the issue of limiting supported platforms and has decided on two, precisely defined platforms for the LDAS data processing. This seems appropriate.

This is one example of a commendable effort to make best use of limited human resources. Additionally, the limited resources available within the LDAS effort have correctly tempered acquisition choices against component assembly and self-maintenance.

Hardware testing programs

The panel commends LIGO on its substantial effort in testing hardware options before making key decisions that depend on several aspects of hardware performance as well as on price. This testing has properly been tailored to LDAS's specific needs and has correctly used actual LDAS and LCS codes.

Costing and price negotiations for hardware

The prices obtained for the various purchases already made appeared to make sense and to exhibit sensible decision-making. Their long-standing policy of postponing acquisition until necessary so as to take advantage of generally decreasing costs is paying off. The processes being used for future decisions, including hardware comparison testing and comparison pricing, as well as price negotiations, seem sound and are an improvement over the processes that appeared to be in place a year ago.

The large anticipated volume of stored data has correctly caused them to focus their tape choice on the upcoming STK 9940B drives. The panel encourages their efforts to negotiate a low-cost "bridging" alternative to serve until those drives become available.

SECURITY AND RELIABILITY

Data Redundancy and Protection

The panel reiterates the importance of maintaining a minimum of two complete copies of the data at distinct locations sufficiently distant from each other that they are highly unlikely to be subject to the same potential natural or man-made disasters. It may not be necessary for both copies to be immediately or automatically accessible.

The easiest and least expensive way to accomplish this is probably to keep a complete copy of the data safely stored at each of the observatories, while the data are also being constantly transmitted to the data center at Cal Tech. However, the desire for compatibility argues for writing second copy tapes at CACR provided that the HPSS now supports "shelf tape".

It is important to treat the possibility of storing all the data in two distinct striping modes for efficiency in searching as separate from the redundancy issue. Even if this is done

(through DTF or some other affordable program), this should be a separate effort and the second (backup) copy should be continuously maintained and sequestered.

The panel suggests the investigation of the potential benefits of additional facility safety measures, appropriate to a production shop, which might not otherwise be appropriate for a research shop such as CACR, including automatic injection of Halon into tape silos in case of fire.

Security, robustness and system integrity

The proper balance between tightly secured computer systems and ease of use within a geographically spread collaboration must be found and kept. However the panel feels the measures in place are not adequate once scientific data taking starts. The exposure of resources to the public internet needs protections against disruptions and damage which may occur by malicious hacker attacks. Special care should be taken for systems of critical importance to the success of the scientific mission of LIGO. The control of systems representing a major capital investment and crucial parts of the detectors must be designed to be protected against any malicious attacks.

The panel discussed and questioned LIGO project members about plans and strategy for security during several sessions. As a result of those discussions, the panel would like to see LIGO present NSF with a comprehensive analysis of potential security risks both in terms of

- possible network break-ins
- data and system integrity
- Ability to continue critical services given component system failures.

A particular data integrity issue that comes to mind is that of database recovery in case of hardware failure, which can be handled by, among other options, maintaining a mirror database.

Development and Test Systems

The panel commends LDAS for initiating the use of a separate test system in support of the LSC. This is one example of an important general principle. For example, it is not uncommon for operations depending upon substantial databases to maintain separate systems for development, testing and production. This arrangement coincidentally provides a useful backup position should the production system fail. In cases where newer generations of equipment are acquired to provide greater capacity, such as is the case for LDAS, the older systems can readily be used for development and testing.

NETWORKING

Local Area Networking

The panel commends LDAS for moving to Gigabit Ethernet and other, more readily supportable networking media, for the majority of local area networking in accordance with the panel's recommendations of 2000.

High speed data connection for the laboratories

The panel commends LDAS for working to implement network connections sufficient to support routine data transfers from the observatories as recommended last year. The panel feels that obtaining access to an OC3 connection in both observatories is important, and encourage the project to enter into ongoing and well-informed negotiations in both Washington State and Louisiana to get favorable rates. The situation in Washington sounds quite promising. There is more background research and more negotiating to be done in Louisiana.

Timely, favorable resolution of this issue is important to facilitate network data transfer as soon as possible both to allow work as soon as possible on the proper "production" data transfer mode and to minimize the need for additional hardware of only short-term benefit at the observatories. Because the acquisition of network connections at favorable pricing is not only a technical issue, help may be appropriate from outside the LDAS group.

EXTERNAL RELATIONSHIPS

Technical community relationships

The panel continues to encourage the LIGO project managers to build working relationships with key technical communities, including mass store, database management and others. For issues ranging from questions about existing HPSS units with Sun Solaris front ends to those dealing with the future of mass store silo and tape technologies to those relating to negotiations of networking solutions, the panel recommends not only calling Bernie O'Lear but also sending relevant people to a few key meetings where they can develop strategic relationships. The panel specifically suggests that LIGO make a point of sending their HPSS FTE to the next IEEE mass store meeting or its equivalent and of requesting permission for him to attend HPSS meetings along with the CACR representative.

Such relationships can provide inside information about hardware and software directions and pricing, about little publicized problems with either and about likely future developments. For example, people within LIGO need to track mass store technologies and HPSS. Such relationships also provide important professional development for LDAS staff.

The scope and importance of LIGO warrant a stronger attempt to build links with the high-performance computing and communications community, members of which share many of the challenges confronting LDAS. In addition to those meetings mentioned

above, attendance at more general meetings such as the Super Computer Conference or Computing in High Energy Physics would likely provide benefit to the LDAS group.

Recommendations

1. Security Issues

The panel recommends that a document be prepared within six months and presented to the NSF that assesses the risks and defines appropriate means to mitigate these risks. Among the items that need to be addressed are

- Multiple, independently secure copies of data and metadata
- Facilities augmentations at CACR and other sites appropriate to the needs of the LIGO production facility
- Network vulnerability, data and system integrity, and robustness under partial system failures
- Independent development, test and production systems
- Protection against hackers, especially for critical systems.

2. Networking

The panel recommends augmented efforts on the part of both LDAS and other members of LIGO in obtaining reasonably priced high-speed network connections between the Hanford, Louisiana and Pasadena sites in a timely manner. LIGO may benefit by seeking advice from the NSF and local universities.

3. External relationships

The panel recommends additional effort by the leadership and staff of LDAS in building direct working relationships with other groups and organizations in the high-performance computing, communications and data storage and retrieval communities including the HPSS consortium.

2. LDAS Software Development and Preparations for Data Analysis with the LIGO Scientific Collaboration

Findings

Software is being developed in collaboration between the LIGO laboratory (LDAS) and the LSC (LAL) with the overall effort adding up to about 20 FTE's. Responsibilities are defined and documented in MoU's. Regular software meetings between all partners participating in the software process are organized. Institutions take responsibility for

long term maintenance and support of the software that they contributed. The analysis software is used for data from the engineering runs and the organized and well planned Mock Data Challenges are important milestones. The analysis system critically depends on the availability of the metadata database, which is implemented using the DB2 product.

Reference platforms for LDAS software development are Solaris and Intel based Linux systems. Software is supported and tested on these reference platforms. The operating system versions are adjusted according to the availability and needs of the hardware used. LAL software development in the LSC is sometimes done on other platforms. It will then be ported to the reference platforms where it will be maintained and supported. The "conditions API" is currently the only module which exploits SMP machines when several user processes are run simultaneously. It performs equally well on a single CPU system running the users' code sequentially.

The LDAS software design has always contained a component known as the "dataCache API" which will provide transparent access to all LIGO data. This particular component is one of the last components to be fully scoped. This was a conscious decision based on the coupling of the storage technology to be used by LIGO with the software tools necessary to interface with the data.

The LDAS software has limited computer security and user authentication methods built in. TCP wrappers control the use of resources and allow access from a defined list of IP-domains only. Only LIGO specific IP-ports are enabled to receive remote requests for data processing and analysis; a shared username and password is used for these requests. Beowulf clusters are hidden behind gateway machines in order to minimize exposure to the general internet.

LIGO is one of the four experimental activities in the GriPhyN Collaboration. They are part of the two proposals submitted for performing R&D on the middleware for distributed data analysis and to implement two Tier 2 centers. Penn State University and University of Wisconsin-Milwaukee are the proposed sites for the Tier 2 centers. In addition, the proposal asks for approximately five FTE's for software R&D for LIGO. The current version of LDAS must be upgraded to make it fully Grid-aware. Caltech is the natural place to play the role of the Tier 1 center.

Comments

The software development process is well organized and is able to address the analysis needs of the collaboration. The institutional responsibilities are essential to guarantee a long and successful scientific run. The proper support and maintenance of the database is crucial and must be high on their priority list.

The strategy to restrict the supported platforms and the choice of platforms is necessary and well done. It limits the support effort which may otherwise not be manageable.

The panel suggests that LIGO Laboratory perform a survey of products addressing the interface to mass storage systems. This may result in a solution of how to address the needs of the "diskcache API" without having to write it by adopting or adapting existing products to LIGO needs. The Storage Resource Broker is an example of such a product that has not been reviewed by LIGO but should be considered.

LIGO 1 will benefit from the additional resources provided by GriPhyN to perform their data analysis and to do science. The modularized structure of LDAS means changes will be needed in only a few places. These enhancements will need to be made compatible with the GriPhyN computer security strategies. The resources for this LDAS implementation and the Tier 1 data service to Tier 2 centers need to be identified.

The recently completed MDC's and the formation of the upper limit groups both provide an excellent venue for obtaining astrophysical input into the LIGO program, as well as exercising the full LDAS/LAL combination. The data analysts in these groups will identify the noise sources which are most troublesome for the science goals, while the experimentalists will be able to provide a reality check and communication with the Lab about the progress. The goal of producing scientific papers is excellent. It will motivate these groups to keep up with the rapidly developing astrophysical insights during the LIGO science run. The team building and collaboration between the LSC software development and LIGO laboratory is to be applauded. The panel is encouraged by this development and recognizes the rapid progress that has occurred since the last review.

The recent "opening" of the LSC to participants who do not necessarily provide a concrete deliverable is a positive development that will allow for more participation by astrophysicists who will bring their insights to the table. The panel encourages this to continue and that the LSC do all it can to involve young people, especially graduate students.

Though not presented in detail, the panel was encouraged to learn about the very strong coupling of the GEO 600 group with LIGO both in terms of experimental overlap as well as on the sharing of data and data analysis efforts.

The panel learned that the LSC data analysis efforts appear to be funded at a level which allows them to participate fully as scientific partners with the Lab as the science run approaches. They also seem to have found the hardware that they need for software development through alternate NSF programs such as the MRI program. Their experience in setting up a Beowulf cluster was clearly of great benefit to the LIGO laboratory's planning for constructing similar clusters.

The panel sees a need for the LIGO laboratory to establish a mechanism by which the scientific impact of engineering improvements during the science run are assessed and acted upon. The Upper Limit LSC groups will presumably evolve into science groups which, by analyzing the data differently, will be advising the lab about those noise sources which are the most troublesome for their scientific goals. The desired improvements will sometimes be of a conflicting nature. When asked about this, the LSC

response is that they are prepared to make clear the scientific tradeoffs of various types of improvements, as this is a capability they have developed for advanced LIGO planning.

A possible example would be to consider how science would drive the LIGO laboratory's response to a situation where one interferometer ends up being much better than another. Some science goals (unexplained bursts, for example) depend critically on a coincidence, whereas others (e.g. continuous wave sources) can be done with a single interferometer. It seems to the panel that the mechanism for this input from the LSC is in place, but that the LIGO laboratory should begin to develop an internal mechanism to provide expertise in astrophysics. One mechanism for this might be adding astrophysicists as members of the Program Advisory Committee. These astrophysicists could both help the LIGO Director resolve conflicting needs of the science groups and provide more continuous astrophysical input that would enhance the scientific reach of LIGO.

The science targets of the upper limit groups are excellent, well-chosen and cover the most often discussed sources. The obviously missing search was for r-modes, either from newly born neutron stars (as are being indicated by optical supernovae searches, such as those being led by Prof. Filippenko at Berkeley) or in X-ray binaries in our galaxy (such as Sco X-1). The panel appreciates that these searches are difficult to define, but would encourage the LSC to continue its investigation of just how such a search could be done. Such a definition clearly needs to be made for the science run, so it seems to us that it should be an important part of the upper limits run as well.

The panel encourages the LSC to continue its efforts in making collaborative connections with the astrophysical community. The makeup of the LSC Upper Limit Teams is a good beginning. Participation in topical meetings and workshops might well serve this function. This will be beneficial to the lab and will also keep the astrophysical community aware of the progress of LIGO. Progress on the electromagnetic front on the properties of compact objects as well as their birth events has been dramatic in recent times. This is due to the successful launches of XMM-Newton and Chandra and the rapid developments in our understanding of gamma-ray bursts, radio pulsars and magnetars. Observations of black holes in binaries is now commonplace, masses have been measured and spin-rates inferred. Upcoming radio pulsar searches might also find deficits of ms pulsars that would point to the importance of gravitational wave emission (either via r-modes or crustal quadrupoles) during the accretion spin-up. It is important to develop a mechanism for these rapidly developing astrophysical insights to influence the science plans for LIGO and possibly LIGO II development.

Recommendations

Professional database administration must be guaranteed by a dedicated person. A plan must be developed describing redundancy and how to perform hot backups, integrity checks and disaster recovery procedures; all of these must be regularly tested.

3. Detector Installation and Commissioning

Findings

Progress on the installation and commissioning of the LIGO interferometers was presented to the Committee. The 2 kilometer interferometer at the Hanford Observatory has been locked for extended periods in its full configuration. While the power level in the arms was limited due to sensitivity of the photo detectors used in the alignment system to the laser beam, the control and data acquisition system worked well. The flexibility of the digital approach to the control and data acquisition system enabled a straightforward work around of the problem with the alignment system by using the optical levers in place of the local optical sensors on the mirrors.

One of the key elements in obtaining initial lock of the full interferometer was the end to end model of the interferometer. The model provided the necessary sequencing of programming the gain and phase of the many digital feedback loops which was essential to achieve lock. Before the end to end model was used, there it was not possible to lock the interferometer.

Significant progress has been made in identifying the sources of excess noise in the gravity wave output channel. The initial engineering runs indicate excessive sensitivity to acoustic noise generated by tank firing, aircraft, automobile traffic, etc. This was traced to phase noise generated in the prestabilized laser system due to vibration of various optical elements. Unfortunately, the Olympia earthquake caused sufficient ground motion that some of the mirrors were damaged and required entrance into the vacuum system to remove and refurbish the damaged optical elements. Operations are expected to begin again at Hanford in two months. During this access, the photodetectors on the mirror alignment system will be replaced with units that will have a factor of 10^3 reduction in sensitivity to the laser beam.

At the Livingston Observatory, the input Michelson and the individual arms of the 4 kilometer interferometer have been locked. Lock of the full interferometer is expected within the next couple of weeks. The commissioning of the interferometer has gone very rapidly. The lessons learned for the commissioning of the 2 kilometer interferometer at the Hanford Observatory have reduced this projected effort by nearly a factor of ten. Unfortunately, the same photodetectors that are sensitive to the laser beam are installed in the mirror alignment system at Livingston. After the new detectors are validated on the 2 kilometer interferometer at Hanford, it will be necessary to replace the detectors at Livingston.

Comments

The LIGO team is to be congratulated on the rapid commissioning and initial operation of the 2 kilometer interferometer at Hanford. The application of the end to end model by Matt Evans to develop the lock sequence for the interferometer was essential to achieving lock. The initial attempts based on previous experience with the 40 meter facility and the

locking of the input Michelson interferometer were unsuccessful due to the complexity of the control system. It is particularly satisfying to see that the environment created by LIGO enables an outstanding graduate student to carry out such an important development.

At the Livingston Observatory, the commissioning of the 4 kilometer interferometer is progressing much faster than estimated. The lessons learned at Hanford were much more effective than previously thought in facilitating this process. It is expected that full system lock will be achieved within two weeks. The panel looks forward to the achievement of this important milestone.

The Diagnostics Test Tool (DTT) and Data Monitoring Tool (DMT) work effectively to provide information about the state of the detectors and their environments and to identify noise sources and malfunctioning components. While these two software frameworks were developed by members of the LIGO Laboratory, the software tools which operate within them are being developed by members of the LSC. At last year's review, the panel was concerned about the effectiveness of this approach. The very significant progress in this area is to be commended.

Preparations for the next engineering runs appear to be under control. The Upper Limit Teams will provide realistic tests of end-to-end hardware and software systems and simultaneously yield important science results.

The LSC is thoroughly integrated into the engineering runs with both experienced users and trainees participating in the acquisition of information about the state of the detector and environment. The size of this community should continue to grow and thus be adequate for LIGO I science runs.

The number of detector wizards and graduate student wizards-in-training is adequate now but may be a limiting factor when 3 detectors are operational.

While the current noise limits are a factor of 10^5 above the design levels, there is a consensus that all but the final factor of 10 or so should be comparatively easy to achieve. The final factor of 3 or so will require understanding of novel noise sources and will be very difficult to achieve. The progress made to this point in understanding the detector and in overcoming currently known obstacles provides some confidence that timely progress will be made in meeting the design goals.

Overall, the panel was very impressed with the rapid progress in transforming a large scale and technically very challenging construction project into an extremely sensitive set of scientific instruments.

As the effort changes from engineering to science, there will be conflicts between continued running of the interferometers and the essential development time needed to further reduce noise and improve sensitivity. The LIGO Scientific Collaboration has as part of its charter, the responsibility to recommend to the director the priorities for

various improvements and upgrades to the system performance. The director has the ultimate responsibility for these critical decisions. In conversations with the director and the spokesperson for the LSC, the panel is satisfied that the present approach is adequate for the initial engineering runs. As the sensitivity of the interferometers improves, it will be important to seek advice from a broader cross section of the astrophysics community in arriving at the appropriate balance.

Recommendations

None

Appendix A

**Charge
NSF Panel to Review
The Installation and Commissioning of LIGO,
and
The LIGO Data Analysis System**

April 30 – May 2, 2001

Installation of the LIGO interferometers is nearing completion and commissioning of the 2 km detector at Hanford is well under way. The 4km detector at Livingston is now installed and commissioning is just beginning. The last major acquisitions under the LIGO construction project are the staging buildings at Livingston and Hanford, and the procurement of the hardware for the LIGO Data Analysis System (LDAS).

1. Review progress on the LDAS hardware procurement plan. NSF has approved Phase I of the LDAS procurement plan. Comment on the schedule of the Phase I procurements and the plan (including cost) for Phase II and III procurements.
2. Review and comment on progress of LDAS software and preparation by LIGO Lab and LSC for data analysis.
3. Review and comment on the installation and commissioning of LIGO interferometers at Hanford and Livingston.

The final report of the review by the committee should be available to the Foundation by May 14, 2001.

Appendix B

Members of the LIGO Technical Review Panel and NSF Staff April 30 – May 2, 2001

Consultants

Prof. Beverly Berger
Department of Physics
Oakland University
Rochester, MI 48309
Office #: (248) 370-3422
Email: berger@oakland.edu
Fax #: (248) 370-3416

Prof. Lars Bildsten
Institute for Theoretical Physics
UC, Santa Barbara
Santa Barbara, CA 93106
Office #: (805) 893-3979
Email: bildsten@itp.ucsb.edu

Prof. Donald L. Hartill (Chair)
234 Newman Laboratory
Cornell University
Ithaca, NY 14853
Office#: (607) 255-8787
Email: dlh@lns.cornell.edu
Fax#: (607) 255-8062

Matthias Kasemann
C1 CD-COMPUTING DIVISION OFFICE
Mail Station: 370 (FCC 1W)
Fermilab
P.O.Box 500
Batavia, IL 60510-0500
Office# 630 840 6387
Email: kasemann@fnal.gov
Fax#

Michael J Levine [Fellow]
5822 Marlborough Ave
Pittsburgh PA 15217
Phone (412) 268-6488
Email: levine@psc.edu
Fax# (412) 268-5832

Barbara Mihalas
5239 Hoag Lane
Fayetteville NY 13066
Phone 315 234 2414
Email: bmihalas@cginetwork.com

NSF Staff

Dr. Victor Cook
Department of Physics,
PAB Rm. B201
University of Washington
Seattle, WA 98195-1560
Office#: (206) 543-4185
Email: cook@u.washington.edu

Dr. Richard Isaacson
National Science Foundation
4201 Wilson Blvd.
Arlington, VA 22230
Office#: (703) 306-1899
Email: ri@einstein.mps.nsf.gov
Fax#: (703) 306-0566