REPORT

OF THE

NSF REVIEW PANEL

FOR THE

LASER INTERFEROMETER GRAVITATIONAL WAVE OBSERVATORY (LIGO)

Performed for the National Science Foundation

Conducted at the California Institute of Technology in Pasadena, CA
November 9-11, 2005

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

November 9-11, 2005

Table of Contents

Attachments

EXECUTIVE SUMMARY

The committee to review the LIGO project met at Caltech on November 9-11, 2005. The charge, the members of the committee, and the agenda are listed in attachments to this report. Here we repeat the charge and provide our assessment of each of the issues raised. The remainder of our report amplifies this brief summary and provides some recommendations.

Review the current state of detector performance and progress towards achieving design sensitivity at all frequencies.

All of the three interferometers, the 4 km and 2 km at Hanford and the 4 km at Livingston, have been operated in coincidence for an extended period of time. *All three instruments have achieved, and slightly surpassed, the design requirement for strain sensitivities of* 10^{-21} *RMS integrated over a 100 Hz bandwidth centered around maximum sensitivity*, as defined in the document LIGO-E950018-02-E, "LIGO Science Requirement Document (SRD)", dated 25-03-1996. The Review Panel congratulates LIGO for this remarkable milestone achievement. LIGO is now ready for, and in fact has just started, an extended observing run with the goal of obtaining a full year's worth of coincident data at design sensitivity.

The computing, networking, data archiving, and data processing infrastructure has been developed to a sufficient level of maturity to start the efficient reduction of the huge volume of data now beginning to flow from the present major science run (S5).

The data analysis tools and search algorithms for the four major astrophysical signals expected in LIGO have been developed, and results from previous short science runs have been published in peer-reviewed journals. The collaboration is now ready to analyze the higher sensitivity data coming from present and future science runs.

The computing solutions implemented by LIGO are based on sound modern tools and techniques. There should be adequate resources to support the key inspiral analyses of the current and projected S5 data sets. The full scientific reach, however, will likely be limited by the computing resources now dedicated to LIGO. This gap could be mitigated by effective use of emergent grid computing, and the collaboration is committed to exploring and possibly exploiting these resources.

Review progress in Advanced LIGO R&D, especially the recent results in high power laser development, the mechanical loss properties of optical coatings, and the recently raised issue of parametric instabilities.

Test mass down-selection to fused silica has been accomplished in a timely fashion. Improvement of silica performance has been steady and there is no real compromise in this selection. Progress in mirror coatings has resulted in a titania-doped tantala composition that is sufficient to meet requirements (i.e., is no longer a show-stopper), but not optimal in terms of mechanical dissipation. Because additional reduction of thermal noise from the coatings translates directly to additional range of Advanced LIGO, aggressive continuation of the coating development program is desirable to ensure that the instrument achieves the highest practical sensitivity.

Input optics, including the pre-stabilized laser, temperature-compensated Faraday isolator, and Mach-Zender electro-optic modulator, have progressed well toward achieving high power capabilities. Further development of the modulator over the next year is planned and necessary to achieve design specifications.

Preliminary work on parametric instabilities arising from acoustic deformations of the Fabry-Perot mirror surfaces suggests that instability is unlikely or avoidable and, if present, may be controllable. Work is continuing and is likely to provide a straight-forward solution to mitigate the impact in the Advanced LIGO instrument.

The committee finds that the LIGO collaboration has an active and effective program to meet the seismic isolation requirements of Advanced LIGO. The HEPI system is a success and good progress has been made on the in-vacuum isolator components. There is concern that the procurement time for the two-stage seismic isolators may become a critical path item for Advanced LIGO.

Concurrent R&D for initial LIGO and Advanced LIGO has been extremely valuable and has proved essential for achieving high sensitivity in initial LIGO. This parallel R&D approach mitigates risk and demonstrates that the LIGO team can meet the significant challenges for the many Advanced LIGO systems. Early incorporation of HEPI technology in Livingston has resulted in a significant improvement in duty cycle. Where timing, technology, and budget permit, further incorporation of other Advanced LIGO innovations in initial LIGO can improve performance and accelerate readiness.

Review the two-year work plan and budget for the extension of the Cooperative Agreement.

The primary task for the next two years is the operation of initial LIGO at the design sensitivity in order to obtain and analyze a year’s worth of coincident data. Other tasks include continued R&D in preparation for Advanced LIGO and education and outreach activities focused in the communities near Hanford and Livingston. The budgeted amounts seem appropriate given that there is now an established basis for understanding the operational needs. Because of the importance of realizing the full scientific potential of initial LIGO, we recommend that the requested 2-year extension of the current cooperative agreement be granted at the requested budget level.

According to the work plan, R&D spending will be phased out by the end of FY 2008 because it is assumed that funding for the construction of Advanced LIGO will be authorized by that time. If for any reason this is not the case, the resulting disruption to the continuity of the project will cause very significant delays in the completion of Advanced LIGO and quite probably the loss of key personnel. NSF should make every effort to avoid such disruption.

Assess Caltech's and MIT's management and operation of the LIGO Laboratory. In making your assessment, please consider: 1) their roles and effectiveness, including their joint oversight function, in maximizing the Laboratory's scientific and educational impact; 2) Caltech's and MIT's vision for the future of LIGO, including Advanced LIGO; 3) Caltech's and MIT's institutional commitments to LIGO, including the appropriateness of the number of faculty and staff involved in exploiting LIGO for science and education considering the scale of the project; and 4) Any other unique considerations that Caltech and MIT management provides to maximize LIGO's science and educational impact.

Caltech and MIT have sustained a multi-decade commitment to gravitational wave physics in general and to LIGO specifically. Both provide a rich intellectual environment, including faculty members and research staff who are international leaders in research in this field, and support for outstanding post docs, graduate students, and undergraduates. Both maintain experimental capabilities that enable full-scale prototyping of LIGO subsystems, and such prototyping has proved to be key to reducing risk. MIT has recently hired two new faculty members in this field. Caltech will continue to seek suitable faculty candidates, and indeed has recently made two offers, both ultimately and unfortunately unsuccessful. Both institutions have provided substantial amounts of space, and Caltech has committed to a plan that would make it possible to co-locate most of the LIGO team.

Recent changes in the structure of the LIGO Directorate and its oversight committee, along with a provision for visiting appointments at Caltech, have made it possible for members of the LIGO Science Collaboration, independent of their institutional affiliation, to participate fully in all aspects of the LIGO project from governance and policy making to leading subsystem efforts.

The National Science Board (NSB) has concurred that planning for Advanced LIGO is sufficiently advanced and the intellectual value of the project sufficiently well demonstrated to justify consideration by the Acting Director and the National Science board for funding in FY 2007 or a future NSF budget request. The Board approved the resolution with the understanding that the existing LIGO Program will collect at least a year's data of coincident operation at the science goal sensitivity before initiating facility upgrades to the new Advanced LIGO technology. There is extremely close coupling and sharing of expert personnel between the operation of the initial LIGO and the design, prototyping, construction, installation, and commissioning of Advanced LIGO. Because of this close linkage, we do not believe that it is cost effective or technically desirable to separate the management of ongoing operations from the management of the Advanced LIGO project. Therefore, while we endorse the general principle of recompetition of cooperative agreements, we recommend that the cooperative agreement with Caltech covering the operation of the initial LIGO be continued through the construction of Advanced LIGO. We also recommend that the Advanced LIGO Cooperative Agreement be structured in such a way as not to create artificial barriers to recompetition when Advanced LIGO reaches stable operational status.

Assess Caltech's progress and plans toward appointing a new LIGO Laboratory Director.

The search process, which involved representatives of the key stakeholders in LIGO, has led to the recommendation of a single candidate who has the strong support of both the LIGO Lab staff and members of the LSC. Negotiations between Caltech and the candidate are currently underway. While the formalities may take several months to complete, it is likely to be known whether or not the candidate will accept an offer before our report is submitted.

THE REPORT

I. MANAGEMENT

Findings

The management of the LIGO project has demonstrated its ability to provide the administrative and technical leadership required to execute effectively this important and extremely challenging national project. The LIGO Directorate, which has evolved over time and now includes the Spokesperson of the LIGO Scientific Collaboration (LSC) as a member, provides an excellent vehicle for detailed coordination and oversight of work undertaken both at the LIGO Laboratory and by members of the broader collaboration. The total effort is very broad, including as it does, facility operations, R&D and design activities, experimentation, and data analysis.

The External Oversight Committee, which reports to the Caltech and MIT administrations, and the LIGO Management Council, which reports to the LIGO Directorate, provide critical multi-institutional oversight and technical/administrative support for LIGO.

Comments

The review committee applauds the demonstrated strong commitment of the Caltech and MIT administrations to the LIGO project in terms of administrative support and other resources. Specifically, the committee is pleased to note the recent actions at MIT in the appointment of junior faculty and provision of new space for LIGO activities. The committee was also pleased to hear a firm statement of intention from the Caltech administration to continue to search for faculty candidates who work in gravitational wave physics and a commitment to provide contiguous space once a new astrophysics building is complete. The appointment of new faculty whose research interests intersect strongly with LIGO will ensure that Caltech continues to nurture LIGO intellectually as it has since its inception. Co-location of the LIGO team will become even more important as the Advanced LIGO project gets underway.

The 40-m interferometer at Caltech and the LASTI facility at MIT have proven their utility in Advanced LIGO R&D and improvements made to initial LIGO. The LIGO Lab should, however, continue to evaluate the need for offline R&D capability against other competing resources.

The committee finds that the budget planning process and the administrative and project controls implemented by the LIGO Directorate in order to execute the project at several sites have worked effectively, and the LSC and LIGO Laboratory activities are well integrated.

II. THE COOPERATIVE AGREEMENT

It is the Committee's understanding that two actions are required in order to continue the LIGO program. First, the project has requested a two-year extension of the existing cooperative agreement to cover the period of time until the construction of Advanced LIGO begins. Second, a new cooperative agreement will be required for the continued operation of initial LIGO after FY 2008. This latter agreement must be considered in the context of the proposal to build Advanced LIGO, which has been endorsed by the NSB for funding. In the sections that follow, we discuss these issues separately.

1. EXTENSION FOR 2 YEARS

Findings

At the start of FY2006, LIGO will complete commissioning of the detectors in preparation for the S5 Science Run. LIGO expects to initiate coincident operations for the S5 run in December 2005. The run is planned for approximately 18 months and is intended to produce one year of coincident data at design sensitivity. Some short-term commissioning tasks may be warranted by the detector performance. A continuation of Advanced LIGO R&D is also planned.

The primary objectives of the work plan defined in the proposal for FY2007 and FY2008 include:

> Operations of the LIGO laboratory including scientific observations with the initial LIGO detectors and the completion of the fifth science run (S5) with the goal of collecting one year of coincidence data near design sensitivity;
>
> Data analysis and management in support of scientific observations with the goal of publishing results from S5;
>
> Continued development and fabrication of full-scale Advanced LIGO prototypes;
>
> Pursuit of forward-looking research related to upgrades of Advanced LIGO and more advanced detectors; and
>
> Outreach to schools and the general public, primarily focused on the regions surrounding the two observatories.

The requested budget for the work plan for FY2007-2008 is $33 million/year. The estimate of the work and the corresponding budget is an extrapolation from the current experience and includes Advanced LIGO R&D activities at a level consistent with the LIGO plan to start construction of Advanced LIGO during FY2008. The annual budget for LIGO activities defined in the work plan is dominated by labor, with roughly two-thirds of the total for salaries and associated costs. The LIGO Lab takes a project approach to budgeting using a work breakdown structure and change control process.

The operations of the LIGO Observatories at Hanford and Livingston are now in steady state, and the planned changes in total costs are driven by inflation. Adjustments have been made in the budgets to accommodate changes in the mix of planned activities (Advanced LIGO R&D vs. commissioning, etc), the plans for the year-long science run (S5), and the completion of preparations for the start of Advanced LIGO.

Comments

Successful completion of S5 is essential in order to realize the full scientific potential of the investment made to date in LIGO. Successful completion is defined to be the collection of 12-months of coincident data at the design sensitivities.

The work plan is consistent with this and the other stated priorities of the LIGO Scientific Collaboration. The proposed budgets seem adequate to support the planned work, and there appears to be an appropriate balance between staffing levels and the other expenses. We note, however, that the operational experience in previous years was dominated by commissioning activities, and a transition to more steady state or routine operations is planned for FY2006-2007. The experience with the S5 scientific run may result in the need to revisit the FY2007-2008 work plans to ensure that needs and priorities match. LIGO Lab may find that the requirements for extended operations (18 months for S5 in order to achieve 12 months of coincident data) are somewhat different from the requirements for the shorter science runs completed up until this point (typically 1-2 months). There may be a need to have more spare parts on hand to meet the duty factor goals, or additional operations staff may be required. The level of effort sustainable over a month may not scale to 18 months. Extended working hours can be required for a month but not indefinitely, and adequate provision for vacation and other absences will be required.

The low level of effort to investigate improvements to initial LIGO detectors is appropriate given the current emphasis on a sustained scientific run. There is always the potential that difficulties with the S5 run will prompt an increased interest in expanding the investment in improvements at the expense of run time. The LIGO leadership understands the external expectations for a long, productive data-taking run and is organized in a way that should allow trade-off decisions to be made in a transparent and efficient manner.

There is little doubt that the S5 run will be a learning experience and will likely result in a better understanding of system performance and reliability issues. In addition the data handling and data analysis requirements will increase dramatically and will require the active engagement of the full LSC for a sustained period of time. LIGO should be prepared to reassess the budget allocations based on some operational experience in extended science mode in the latter part of S5 or in preparation for a possible S6 following additional instrumentation improvements. However, these are nice challenges to face.

The investment in Advanced LIGO R&D during the proposed 2-year extension is appropriate and necessary to prepare for the Advanced LIGO construction project that is scheduled to begin in FY2008.

There is some flexibility to reprogram items in the course of FY2007-2008 but far less than existed in previous years. LIGO has experienced essentially flat funding over the past five years, with a reduction of $1M in 2005 and 2006 to a total of $32M per year. The LIGO team has accommodated the cumulative effects of inflation against static funding using a multi-year budget approach and careful timing of planned expenditures. Despite this, it was necessary to reduce staff by twelve in 2005 and eliminate some unfilled positions. The proposed extension restores funding to the $33M per year level and foresees a shift in effort as the initiation of Advanced LIGO is approached. Under this plan no additional staff reductions are foreseen. This is prudent given the anticipated staffing *increases* that will be required to execute Advanced LIGO.

If Advanced LIGO does not begin on the currently foreseen (2008) schedule, further extension of the ongoing LIGO Operations at the current funding level (implying flat funding over more than 7 years) would force additional staff reductions--jeopardizing the ability of the LIGO Lab to begin ramping up staffing for the new project as quickly as is required for maximum cost effectiveness.

A change to the projected start date for Advanced LIGO will therefore have a significant impact on work plans. NSF should notify LIGO leadership immediately if there is a potential change to the Advanced LIGO start date in order to ensure accurate impact assessments to inform decisions and to enable LIGO to adjust work plans in a timely way, if necessary.

The committee was not presented with LIGO Operations budgets for the period of time during which the current LIGO will be operating in parallel with the Advanced LIGO Project although there is an expectation that personnel will shift from operations to construction in accordance with the planned work scope.

Recommendations

Because of the importance of realizing the full scientific potential of initial LIGO, the current cooperative agreement should be extended for two years at the requested level.

LIGO should continue to place the highest priority on completing the sustained scientific run (S5) and meeting the stated goal of achieving one year of coincident data at design sensitivity.

The LIGO Operations budgets for the period of time when LIGO operates in parallel with the Advanced LIGO construction project should be reviewed as part of the Advanced LIGO baseline review in order to ensure that there is a consistent understanding of the entire effort and the constraints that will exist in the two budget categories.

2. RENEWAL/RECOMPETITION IN FY 2008

Findings

The President's FY 2006 budget included Advanced LIGO for an FY 2008 start. The construction plan relies on some of the same personnel as have been involved in the design, installation, commissioning, and operation of initial LIGO. Advanced LIGO would be installed in the same buildings as initial LIGO, and very close coordination of scheduling between Advanced LIGO and ongoing activities at the sites, including possible science runs and prototyping will be required. Arrangements for maintaining operation of at least two gravitational wave detectors through collaborations with foreign collaborators may prove desirable.

It is the normal practice of the NSF to recompete cooperative agreements when they come up for renewal (see the attachment). Since the award for Advanced LIGO has already been approved on the basis of a submitted proposal, this portion of the future work on LIGO cannot be competed. Therefore, it appears to the Committee that the only question is whether to recompete the continued operation of initial LIGO.

Comments

The Review Committee agrees with the general principle of re-competition for facilities that are in a stable or slowly evolving operational phase. However, LIGO is unique in that the plan is to carry out a very challenging development project while maintaining some continuity in the operation of existing facilities. The present arrangement, with Caltech as the awardee institution, and MIT and Caltech sharing the responsibility for oversight and implementation of LIGO, is working very effectively. The achievements to date are truly remarkable and fully justify the trust placed in these two institutions. The technical infrastructure and intellectual leadership to execute the next phase of the project clearly continue to reside at Caltech and MIT. The decision has already been made to proceed with Advanced LIGO. Because of the strong coupling between initial LIGO activities already underway and the upcoming construction project, seamless coordination will be a great strength for both. It is the strong view of the Review Committee that recompetition would seriously delay and jeopardize the timely completion of Advanced LIGO within the projected budget.

Recommendations

Due to the very close linkage between initial LIGO and Advanced LIGO in terms of personnel, facilities, scheduling, and agreements with foreign partners, we do not recommend recompeting LIGO at this time.

We do recommend that future co-operative agreements be structured in such a way as not to create artificial barriers to recompetition when Advanced LIGO reaches ongoing operational status.

III. CURRENT LIGO PERFORMANCE

1. LIGO SENSITIVITY

Findings

The strain sensitivities of all three interferometers, the 4 km and the 2 km at Hanford (H1 and H2) and the 4 km at Livingston (L1) have seen substantial improvements this past year. The design strain sensitivities of the spectrometers have been defined for the 4 km interferometers to be 10^{-21} RMS, integrated over a 100 Hz bandwidth at the peak sensitivity, in document LIGO –E950018-02E "LIGO Science Requirement Document (SRD)", dated 25-03-1996. All three interferometers have achieved, and slightly surpassed, this design sensitivity, reaching an RMS strain of 0.4×10^{-21} integrated over a 100 Hz bandwidth near peak sensitivity. (Of course, the sensitivity of the 2 km interferometer has to be scaled by the inverse of the length). LIGO plans to continue efforts to improve this sensitivity further during the upcoming science run S5.

Comments

LIGO should be congratulated on this remarkable milestone of achieving design sensitivity as well as the impressive improvement in the interferometer uptime fractions.

Recommendations

The vigorous efforts to improve further the strain sensitivities of the interferometers, especially at the lower frequencies below 100 Hz, should continue. For example, we encourage the project and NSF to explore the feasibility of early installation of HEPI devices at the Hanford site. However, the priority of these improvement efforts should be carefully balanced with the high priority of completing the year-long observing run S5.

2. LIGO DUTY CYCLE

Findings

In the recently completed S4 science run, the duty cycle of S4 improved to an average of 80.5% on-time for H1, 81.4% for H2, and 74.5% for L1, respectively, up from the corresponding numbers in the previous run S3 of 70%, 63%, and 22%. The fraction of time with all three interferometers operating at the available sensitivity simultaneously, i.e., triple coincidence, was 57% in S4, compared to 16% in run S3. For the S5 run that has just commenced, the targets are 85% for each individual interferometer and 70% for triple coincidence duty cycle. These duty cycles include scheduled downtime and clearly offer a duty cycle that enables the detection of gravitational signals. The committee has reviewed, as an example, the downtime of H1 during the S4 run. The fraction of downtime specifically due to wind, seismic disturbances, and calibrations was reasonable. Improvements have been made to increase the duty cycle for the S5 run, and the opportunities for the largest future improvements seem to be in the areas of retaining lock and faster lock acquisition.

Recommendations

The panel fully supports the ongoing efforts of the LIGO Lab to improve the lock robustness and reduce the time to re-acquire lock, including broad thinking about different control and lock acquisition strategies. New strategies that could lead LIGO and Advanced LIGO to more robust locking and >95% duty cycles have a clear value for the future science of these detectors.

IV. COMPUTING

Findings

The computing and data analysis collaborators and staff are top-notch and have based their solutions on sound modern tools and techniques. The scope of the computing challenge is about 10% of an LHC (ATLAS,CMS) collider experiment, and the currently deployed resources scale reasonably to this well documented reference. The LIGO project has a challenging cyber-security problem spanning from highly critical systems at the interferometer sites to the non-trivial authentication issues of grid computing. A cyber-security plan exists, was reviewed previously, and is appropriate for the challenges particular to LIGO.

The LIGO laboratory and the LSC have together developed a grid computing environment (LDG, LIGO Data Grid) based on centers associated with the laboratory and collaborating institutions. The LDG is built from standard grid tools and commodity hardware and is well suited to the challenges of LIGO. Caltech, MIT, and several institutions of the LSC have made substantial in-kind contributions to the LDG by providing space, power, and cooling--which together are becoming a leading cost driver in grid computing. Searching for pulsars in the LIGO data set is very computation-intensive, far outstripping the capacity of the LDG. The proponents have answered this challenge by developing "Einstein@home", modeled after "SETI@home". Currently 60,000 subscribers participate in Einstein@home delivering about 4x the capacity of the LDG. In addition to genuine scientific productivity, Einstein@home is an impressive vehicle for outreach.

The LDG and Einstein@home resources have been adequate to serve the demand of the collaboration to date. It is recognized that these resources will be outstripped as the S5 data set grows. In the near term, analysis efforts will have to be managed, and there are now administrative mechanisms in place to establish and implement priority access to computing resources. For the longer term, the LIGO project is working to federate with the "Open Science Grid" (OSG) in order to gain opportunistic access to a much larger pool of computing resources. Effective participation of LIGO in the OSG has been difficult to date, and the particular need of LIGO to move relatively large amounts of data from their mass data stores to OSG compute resources has led to inefficiencies. The promise of genuine returned value from OSG has yet to be realized. The LIGO proponents and OSG consortium are working together to overcome these difficulties.

Comments

The LIGO collaboration is now transitioning from a long and successful R&D period to an operational phase. Computing support for the collaboration has been provided on a best-effort basis, which may become a stress on the computing staff as demands increase through the S5 period. The laboratory and LSC should consider developing a sustainable expectation for computing support, perhaps outlined in MOUs.

The computing staff is now refining a computing model based on the early experience of S5 data handling and processing. This model should be used to guide and motivate future proposed upgrades to the LDG and to understand what the best strategy is to interact with larger grid consortia such as OSG.

Recommendations

The LIGO project should develop, document, and qualify a computing model for the experiment. This model should be expressed in a living document that serves the computing proponents, collaboration, and funding agencies in determining the data handling, production, and analysis requirements of the LIGO data set.

The LIGO project should continue working with the Open Science Grid to develop effective use of resources outside of the LIGO Data Grid.

V. ASTROPHYSICS AND DATA ANALYSIS

Findings

The LIGO team has published a number of papers with upper limits on gravitational waves from binaries, burst sources, continuous wave emitters, and stochastic backgrounds. Although none of the upper limits yet constrain plausible astrophysical models, the progress has been impressively rapid. For example, the upper limit to the background energy density in gravitational waves has improved by nearly five orders of magnitude. With LIGO now at the promised sensitivity for the upcoming S5 science run, detections of binary neutron star inspirals are possible (although still requiring good fortune), and the limits placed on stochastic backgrounds by nondetections will for the first time improve on existing limits, in this case from Big Bang nucleosynthesis constraints. LIGO is therefore poised to start contributing to astrophysics.

Comments

The statistical techniques employed to search for all four categories of sources (binaries, burst, continuous, and stochastic) continue to develop well. This work has been facilitated by free exchange of ideas within the LSC as well as by communication with statisticians. In addition, the community has established websites to which numerical modelers can upload waveforms from binary inspirals, supernovae, or other sources, which are then used by the data analysis teams to cross-correlate with data. Currently, the inspiral portion of binary coalescence has been the primary focus, due to astrophysical expectations and to the availability of second-order post-Newtonian waveforms. Work is in progress on templates for ringdown waveforms. There has been exciting recent progress in numerical relativity supported by the NSF, and Caltech and Cornell have augmented this with their own initiatives. This strong effort, at Caltech and other institutions, has already produced the first crude merger waveforms. There is, however, a great deal of work yet required to yield reliable coalescence simulations and waveforms before Advanced LIGO is on line.

Recommendations

The LIGO data analysis group has displayed a commendable thoroughness and flexibility in its approach and has drawn from a considerable diversity of expertise in time series analysis. It is important, however, as the experiment gets closer to detections, that the group make its presence felt more frequently at mainstream astronomy conferences. This will enrich the community of gravitational wave astrophysicists, lay the groundwork for fruitful interactions, and perhaps as a dividend, yield suggestions for currently unexplored sources.

We also note the need for strengthened investment in numerical relativity, particularly simulations of mergers between two black holes or a black hole and a neutron star. These events involve the extremes of strong gravity and will be the best tests of general relativity. Extraction of the most exciting science from gravitational waves is dependent on having the numerical infrastructure needed to interpret merger waveforms. NSF in reviewing its priorities should consider additional investment in this area.

VI. ADVANCED LIGO

1. SEISMIC ISOLATION AND SUSPENSION

Findings

To achieve the level of isolation from environmental disturbances that is required by Advanced LIGO, a three tier isolation strategy has been proposed: HEPI platforms for the BSC chambers, a two stage vibration isolation unit inside the BSC chambers, and a quadruple pendulum suspended from the isolation unit.

To mitigate unanticipated seismic noise from logging near the Livingston site, development of the HEPI system was accelerated, leading to the successful deployment of the pre-isolators on L1. We applaud the LIGO collaboration for this plan of action and their success in the implementation of the plan. The installed HEPI system not only permits L1 to operate at high efficiency, but provides important input data for the design of the in-vacuum isolation systems.

It was reported that the active in-vacuum seismic isolators, because of their cost and importance to Advanced LIGO, represent perhaps the most serious potential bottleneck for Advanced LIGO. We find that good progress has been made toward the realization of a seismic isolation system. A pre-prototype isolation unit has been built and tested, and its performance met the specifications for Advanced LIGO over much of the required frequency range. Construction of a full scale prototype is underway, to be installed and tested at the LASTI facility. We find the LIGO collaboration R & D plan for seismic isolation is sound and on track.

The decision to employ a quad pendulum mass suspension system for the final layer of seismic isolation appears to be sound, and good progress has been made toward its implementation. The design was made in conjunction with UK collaborators who have experience with multiple pendulum suspension systems used by the GEO project. A prototype quad mass unit is being installed at LASTI where tests of the three sensing and control components will be made. A collaborative effort is underway to study a lingering issue about potential problems associated with charge accumulation on the test masses.

Comments

In summary, we find that the LIGO collaboration has an active and effective program to meet the seismic isolation requirements of Advanced LIGO. Performance of the HEPI system is a success and will probably be improved further as the system is fine-tuned under operating conditions. Problems may arise with the seismic isolation units and quad masses, but there is time to solve such problems, and the LIGO collaboration has proven itself to be up to the task of identifying and implementing such solutions.

The work of the LIGO collaboration on seismic isolation is ground breaking and is likely to continue to lead to technology transfer to other areas of science.

Recommendations

We recommend that the LIGO collaboration and NSF explore avenues through which the HEPI system may be implemented at the Hanford site to further improve the operation of H1 and H2 in advance of the commissioning of Advanced LIGO.

2. PRE-STABLIZED LASER

Findings

A key component for Advanced LIGO is a 20 times higher laser power. In the past year, the collaboration between the Albert Einstein Institute, Laser Zentrum Hannover, and the LIGO Lab has demonstrated a power of 195 W. This exceeds the 180 W Advanced LIGO power requirement. The Advanced LIGO PSL recommended that this design proceed to preliminary design. High power photodetectors are needed to utilize the higher power fully, and preliminary tests are showing no problems at this time.

3. FARADAY ISOLATORS

Findings

A terbium gallium garnet Faraday isolator has been proven in. 40 dB isolation and $\lambda/4$ passive thermal lens compensation have been achieved at 100 W power. It is believed that $\lambda/20$ is possible with a better matched waveplate.

4. ELECTROOPTIC MACH-ZEHNDER MODULATOR

Findings

High-power electrooptic modulators based on rubidium titanyl phosphate or rubidium titanyl arsenate are used in a parallel modulation Mach-Zehnder configuration to eliminate the problem of side bands on the side bands that occurs with serial modulation. A prototype modulator is locked and undergoing characterization. No damage was seen after >400 hours of irradiation at 85 W, which is greater than would be experienced in Advanced LIGO. Given the laser frequency stability requirements, the finesse gives the requirements for the Mach-Zehnder modulator phase noise as $<7 \times 10^{-6}$ rad/Hz$^{-½}$ at 30 Hz. An optimized design that will meet this specification is estimated to be a year away.

Comments

Another possible method to generate a carrier with sidebands at two frequencies without sidebands of sidebands is to phase modulate at both sideband frequencies and then amplitude modulate at the sum and difference frequencies. This method is naturally immune from introducing phase variations to the carrier.

Recommendations

This is a significant development program worthy of the effort being put into it. It should remain a high priority to keep the one-year time schedule. The ongoing work to understand the Advanced LIGO sensitivity to noise from the Mach-Zehnder configuration is important, and we support continued discussions of alternative solutions.

5. TEST MASSES

Findings

The down-selection of the test masses has been resolved with the choice of silica. Over time, the performance of silica has steadily improved to the point where there is no real compromise in performance. Although sapphire still performs better at longer wavelengths, silica was calculated to perform at least equally in detecting three out of four types of events modeled. Sapphire is estimated to be significantly better only for LMXBs at much higher frequencies. It is certain that silica can be produced in the quality necessary to meet design specifications. Sapphire continues to have very significant technical risks, though it is still being pursued as a contingency. If work on optical absorption and thermal noise in sapphire is successful, the suspensions have been designed for dual use so they could be retrofitted for sapphire with only slight modifications.

Absorption problems with existing test masses degraded the performance in the Hanford H1 interferometer. One test mass was replaced with a fully characterized spare that was available, and the other was successfully cleaned in situ. Forensic examination of the test mass that was removed revealed point defects were the problem. Their exact nature has not yet been identified.

Silica ribbons have been chosen for the test mass suspension as they are more stable than fibers. However, there are still a large number of different low frequency resonant modes from 0.3-4 Hz that must be damped. A CO_2 laser-heated ribbon pulling system has been developed as commercial processes cannot achieve the necessary uniformity. The laser ribbon puller was found to be preferable to a flame puller.

6. COATINGS

Findings

Based on previously known values of the mechanical dissipation of the optical coatings deposited onto the fused silica mirror substrates, the ultimate sensitivity of the Advanced LIGO detectors could be limited over a significant frequency band by thermally-generated position noise of the mirror surfaces. LIGO's collaborative experiments with other groups have shown that the tantalum-dioxide layers in the optical coating were responsible for the majority of the mechanical dissipation, and hence thermal noise generation. The source of a substantial fraction of the mechanical loss has been identified as oxygen vacancies and reduction in tantala layers in the mirror coatings, though loss in the silica layer has not been ruled out.

Work has been ongoing to reduce coating thermal noise below the optical noise. The recently-completed TiO_2 doping experiments reveal that the coating mechanical loss angle can be reduced from $\sim 2.6 \times 10^{-4}$ to $\sim 1.5 \times 10^{-4}$ for TiO_2 concentrations of >3% in the Ta_2O_5 layer. Additional doping up to 50% did not further reduce the loss angle indicating that this correction has saturated.

Anecdotal results suggest that cobalt doping in tantala may have achieved lower mechanical losses, but had excessive optical absorptions. One ppm optical absorption translates to 400 mW absorption in the coating in this high power application. This can result in substantial heating in the test mass. The cobalt optical absorptions were a factor of 10^4 above the acceptable level, so this is not a viable option.

Previous results with hafnia-silica coatings did not have significantly better mechanical losses and displayed some scattering suggestive of crystallite formation in the hafnia. The team is hopeful that doping might reduce this tendency to crystallize.

There is an extensive experimental plan going forward to test other doping/coating configurations for low mechanical and optical loss. Lutetium doped tantala coatings are being tested now. Unfortunately there is only a relatively small body of measurement data on mechanical loss to guide the experimental program.

Coated optics will be tested in LASTI in April 2006 using the best available coating. Ideally this will be the same coating as will be used in Advanced LIGO, but it is not crucial. In 2007 a baseline coating will be determined from LASTI experience and current research. The Advanced LIGO test masses are scheduled to be coated in fall 2008, but there is some room for delay if promising research is nearing completion.

Comments

It is considered possible that additional doping experiments can find a sweet spot combining both low optical and mechanical losses. If no further improvement is made, the mechanical loss of the present titania-doped tantala coating formulation is sufficient that adequate performance may be achieved in Advanced LIGO. However, this is an important result since it has long been supposed that the mirror thermal noise would set the achievable sensitivity, whereas it now is clear that, at the best achieved loss values, the range of the detector can be increased by re-optimizing the recycling and other adjustments. Still at present, until the mechanical loss is reduced below the optical noise, improvements in mechanical loss will directly translate to improved sensitivity and range of Advanced LIGO.

It is most likely that doping with titania provides stoichiometry compensation to reduce dangling bonds from oxygen vacancies. However there may still be some reduced Ta^{4+} that absorbs both optically and mechanically. Cobalt is a compliant dopant with multiple valence states that is more easily reduced than tantalum. Therefore it could buffer the coating against both Ta^{4+} and oxygen vacancies over the entire growth, cooling, fabrication, and coating ranges. However, like other transition metals, the various valence states have significant broad absorption peaks as was observed.

Recommendations

The LIGO team and this panel are hopeful that properly amorphous and oxidized hafnia would have little mechanical loss since it is not subject to reduction. The present coating vendors appear to have relatively little experience in coating with hafnia. It is suggested that hafnia-silica coating experiments by more knowledgeable vendors would help to define this problem fully.

There are other compliant dopants besides cobalt oxide. Other transition metal oxides such as manganese oxide would also have significant absorption at the system wavelength. However it is suggested that the team investigate doping the tantala layer with other compliant oxides. Ytterbium and europium oxides both are somewhat compliant dopants and have a limited number of sharp absorption lines.

7. LARGER MODE-SIZE BEAMS TO REDUCE THERMAL NOISE

Findings

The LIGO team has made some interesting experiments using large-diameter laser beams, with the objective of obtaining enhanced averaging of the thermal position noise of the mirror surface. This represents a second approach, beyond the interaction with the coating group at Lyon to use TiO_2-doping of the Ta_2O_5 coating layer to reduce its mechanical losses, and hence the amplitude of the thermally-induced surface motion. The approach of using larger than usual laser beams leads to substantially more strict requirements on the mirror surface figure, in that the reduced angular content of the larger beam diffracts less and so can probe a smaller area on the remote mirror. So now the wavefront matching on the mirror no longer generally involves spherical wavefronts, but rather ones with richer phase and amplitude dependences spatially. The LIGO team has used rotating masks during the coating to allow depositing of axially-symmetric, but not spherical, reflecting surfaces. Experiments showed that field distributions approximating the expected "Mexican-Hat" mode-shapes could indeed be supported with the specially-fabricated mirrors.

Comments

For the LIGO application, this approach to thermal noise reduction might well be painful because of the added system requirements; the mode location on the mirror surface now contains another constraint beyond the usual low-noise condition to have mode symmetry around the mirrors' axes of rotation and rocking. Another predictable difficulty would be to achieve dark fringes in the input beamsplitter's output port, as the mode phase and size properties in the long arms will depend in much more detail on the mirror coating profiles. Probably this mode mismatch can be trimmed with their usual mirror local thermal-expansion tricks, for example by using an absorbed auxiliary laser beam to "paint" hotter spots onto the mirror's surface. It seems likely that modeling could offer insight into larger mode spatial distributions that may be less critically dependent on the correct mirror surface profile. Certainly modeling should be helpful in developing the proper masking strategies to use during the coating process.

Recommendations

The LIGO team is encouraged to keep up a low-level effort on the use of specialized non-Gaussian modes to attain lower rms levels of mirror surface position noise.

8. THERMAL COMPENSATION

Findings

It has been shown in Initial LIGO that thermal aberrations in the test masses can be successfully controlled using adaptive corrective heating with a CO_2 laser from outside the vacuum system. Using this correction, operation at the full design power of 6 W is possible in Initial LIGO. Heating to a radius of curvature approximately half the length of the arm is required for stability. Advanced LIGO with 120 W of power will require an even higher degree of control. In addition to the CO_2 heat projectors already used, incandescent ring heaters will be suspended in the vacuum. This compensation will be axisymmetric as will be required to compensate homogeneous coating or substrate absorption in the test masses. This has been tested on a bench top. Feedback from wavefront sensors will be required to be fully optimized at such high powers.

9. PARAMETRIC INSTABILITIES

Findings

Braginsky has pointed out a parametric instability that couples acoustic deformations of the Fabry-Perot mirror surfaces to the excitation of higher transverse Fabry-Perot modes. The targeted circulating power in Advanced LIGO of 0.86 MW is sufficient to excite this instability. The LIGO Lab is analyzing the likelihood of, and possible solutions to, this instability. The work to date shows that only the lowest four to five modes can contribute because higher modes are strongly attenuated by diffraction losses. Given this small number of modes, a coincident resonance between the acoustic modes of the mirror substrate and the transverse cavity mode splitting seems unlikely. The acoustic modes are very narrow, of order 1 mHz, and their frequency detuning is shifted by small changes in the mirror radii. Work is ongoing to confirm the current expectations. The panel agrees that a coincidence seems unlikely and that the ongoing work will likely avoid any coincidence in Advanced LIGO and also be able to find a reasonable and straight-forward solution to mitigate their impact in the Advanced LIGO instrument.

10. ADVANCED LIGO CONTROLS

Findings

One of the important issues for Advanced LIGO is the control system to acquire and maintain a robust lock. The 40-m Caltech prototype interferometer has recently shown major advances in the cavity locking and gravitational signal recovery strategies for the optics configuration of Advanced LIGO. This strategy allows lock and alignment signals to be generated into substantially independent channels, independent of departures from perfect locking conditions. This is a very important improvement over the previously-used quadrature detection scheme in which the orthogonalization matrix values were fundamentally dependent upon the state-of-lock, making the approach to full lock slower and unpredictable. The new capability is enabled by an additive scheme in which two electrooptic phase modulators used within a Mach-Zehnder interferometer discussed above can produce two sets of independent modulation sidebands, without cross-terms, thus enabling a data recovery at difference and sum frequencies, as well as the modulation frequencies individually. The increased number of uniquely-coded signals greatly improves signal identification and association to a particular misaligned optic. This improvement should allow greater lock stability and faster lock acquisition in Advanced LIGO, leading to a high duty cycle. This work appears to be on a sufficiently advanced schedule so that any important interactions between the Advanced LIGO Core Optics Components and Interferometer Sensing and Control groups can occur. Work on the recent lock demonstration will continue with e2e modeling for Advanced LIGO. The panel is pleased by the progress on this important task.

Comments

The main challenge for lock acquisition in initial LIGO is bringing the long Fabry-Perot cavity arms to within a small fraction of an optical wavelength so that the control signals are linear. For the higher power levels of Advanced LIGO, thermal effects and their settling times are likely to be much more important than in initial LIGO. One possibility to facilitate rapid lock re-acquisition, much faster than any long thermal time constant, is to augment the lock acquisition system by locking additional low-power lasers to each Fabry-Perot cavity arm. This control might allow the Fabry-Perot cavity arms to be initially controlled very close to resonance to allow a faster cross-over to lock with the more traditional controls.

Recommendations

We strongly support the ongoing pursuit of control enhancements to improve lock robustness and reduce lock acquisition times for Advanced LIGO.

LIGO 2005 Annual Review
- Conflict of Interest Matters -

- **The P.I.**

  Barry Barish is a member of the National Science Board. He is thereby not allowed by regulation to represent himself or Caltech in negotiations or any related dealings with the NSF on any proposal, cooperative agreement, etc.

- **The Panelists**

  a) Read and sign Form 1230P.

  b) In general, reviewers should not have any "direct and predictable financial interest" in the outcome of the proposal. Any present or planned future collaboration, employment, or contracts with any part of the LIGO operation, should be disclosed to the NSF staff.

LIGO 2005 Annual Review
- Charge to the Review Panel -

1. **Review the current state of detector performance and progress towards achieving design sensitivity at all frequencies.**
2. **Review progress in AdvLIGO R&D, especially the recent results in high power laser development, the mechanical loss properties of optical coatings, and the question recently raised issue of parametric instabilities.**
3. **Review the two-year work plan and budget for the extension of the Cooperative Agreement.**
4. **Assess Caltech's management of the LIGO Laboratory. In making your assessment, please consider:**
   - **Caltech's role and effectiveness, including oversight functions, in maximizing the Laboratory's scientific and educational impact;**
   - **Caltech's vision for the future of LIGO, including Advanced LIGO;**
   - **Caltech's institutional commitment to LIGO, including the appropriateness of the number of faculty and staff involved in exploiting LIGO for science and education, considering the scale of the Project; and**
   - **Any other unique considerations that Caltech management provides to maximize LIGO's science and educational impact.**
5. **Assess Caltech's progress and plans towards appointing a new LIGO Laboratory Director.**

LIGO REVIEW
NOVEMBER 9-11, 2005
PANELISTS

Dr. Charles Baltay
Department of Physics
Yale University
New Haven, CT 06511
charles.baltay@yale.edu

Ron Davidson
MS17 C-Site L240
Princeton Plasma Physics Laboratory
P. O. Box 451
Princeton, NJ
rdavidson@pppl.gov

Vincent Fratello
Integrated Photonics, Inc.
132 Stryker Lane
Hillsborough, NG 08844
vjfratello@integratedphotonics.com

Kurt Gibble
0104 Davey Laboratory
PennState University
University Park PA 16802
kgibble@psu.edu

John Hall
JILA
University of Colorado
440 UCB
Boulder, CO
jhall@jila.colorado.edu

Blayne Heckel
Physics Department
University of Washington
Seattle, WA 98195
heckel@phys.washington.edu

Thom Mason
Oak Ridge National Laboratory
P. O. Box 2008
Oak Ridge, TN 37831
masont@ornl.gov

Cole Miller
Department of Astronomy
University of Maryland
College Park, MD 20742-2421
miller@astro.umd.edu

Bob Tschirhart
CD-Computing Division Office
Mail Station: 370 (FCC 1W)
Fermilab
P. O. Box 500
Batavia, IL 60510-0500
tsch@fnal.gov

Sidney Wolff
National Optical Astronomy
Observatory
P. O. Box 26732
Tucson, AZ 85726
swolff@noao.edu

Jim Yeck
IceCube
222 W. Washington Ave.
Madison, WI 53703
jim.yeck@icecube.wisc.edu

LIGO 2005 Annual Review
- Agenda -

Wednesday, November 9, 2005	Thursday, November 10, 2005
08:00 Continental breakfast 08:30 Panel Executive Session (closed) 08:55 Introduction and Charge to Panel 09:00 Welcome — Tom Tombrello (Physics, Math, & Astronomy Division Chair) 09:10 LIGO Overview & Status (Stan Whitcomb) 10:00 LIGO Commissioning and Run Plan (Peter Fritschel) 11:00 break 11:15 LIGO Science Results, Analysis Plans (Patrick Brady) 12:15 Lunch 13:15 Panel Executive Session with Emlyn Hughes and other members of the Oversight Committee 14:15 LIGO Lab/LSC Integration and Management (Peter Saulson) 15:15 break 15:30 Panel Executive Session with Peter Saulson and selected LSC members 16:30 Panel Executive Session (closed) 17:00 Panel presents questions to LIGO for 2nd day sessions 19:00 Dinner	08:00 Continental breakfast 08:15 Panel Executive session (closed) 09:00 Advanced LIGO Status (Carol Wilkinson 10:00 Break-out sessions (as appropriate) 12:30 Executive Lunch 13:30 Executive session: Meeting with Prof. Jennings (Provost) and Prof. Tombrello 14:30 Additional Breakout Sessions (if needed) TBD Panel executive session to write report, etc. **Friday, November 11, 2005** 08:30 Panel Executive Session (closed) 11:00 NSF Closeout Presentation 12:00 Meeting ends

<u>Appendix D to NSB-97-241</u>

<u>NSB-97-224</u>

RESOLUTION APPROVED BY THE NATIONAL SCIENCE BOARD
AT IT'S 346TH MEETING, NOVEMBER 13, 1997 CONCERNING

COMPETITION, RECOMPETITION AND RENEWAL OF NSF AWARDS

Whereas the Committee on Programs and Plans has outlined, at its meeting in November 13, 1997, the major principles and key issues in a report "Competition, Recompetition and Renewal of NSF Awards: (NSB 97-216) in the context of the various types of NSF Awards; and

Whereas the Committee on Education and Human Resources concurs in the principles articulated in the report;

Now, therefore, be it RESOLVED, that the National Science Board:

Affirms its strong support for the principle that expiring awards are to be recompeted unless it is judged to be in the best interest of U.S. science and engineering not to do so. This position is based on the conviction that peer-reviewed competition and recompetition is the process most likely to assure the best use of NSF funds for supporting research and education. And

Requests that the Director, NSF, take such steps necessary to ensure that NSF practices embody this principle.