THOUGHTS ON LIGO PHASE A OPERATIONS SCENARIO

Peter Saulson December 14, 1988

Assume around 12 "scientists", plus graduate students (around 8). In addition, we have LIGO engineering staff of 6, 4 technicians at home institutions, and staff of 5 (?) at each site.

Baseline plan allows about two years after design freeze (how specific is design and how frozen?) before first facility "available", I guess meaning ready to move in. We need to consider what sort of facility problems will have to get shaken down along with receiver and data system problems.

Another one year before second site available. We will hope that work at second site goes substantially faster because of everything learned at the first site.

There are a number of aspects of division of labor to consider. First, two institutions will be well-coordinated, and the lag between first and second site availability means there will be every incentive for whole team to work on that one site (to the extent that people work on site). We need to distinguish in our planning between work where engineers take lead responsibility (site, buildings and power, vacuum system, data hardware) and work for which scientists take lead (receivers, data software). We need to figure out how much work can realistically be done at home, and what will need time and manpower on site. Presumably, responsibility for testing which can be done at home will be divided between the two institutions according to facilities and personnel available. Design, installation, acceptance, debugging, running, and data analysis will not have any institutional identity. The division between home and site work will determine whether home technicians or site technicians are involved in particular tasks. There will probably be a lot of travel for scientists and staff engineers.

I think most people will agree that the more work which can be done at the home institutions, the better. A corollary to this is, unfortunately, that a substantial investment will have to be made to construct laboratory test-beds which are as identical as possible to the facilities on-site. This includes, at a minimum, tanks which are identical to the LIGO tanks. (How many copies of corner, instrumentation, and splitter tanks required?) It probably also means a computer system which at least looks the same to the user as what is installed at the sites. Is there any aspect of this which is a luxury, or is postponeable to "Phase B"? I don't know.

Assuming that such facilities exist at the home institutions, there is a large number of tests of receivers which, because of the ease of working, can best be done at home. These include tests of lasers, mode cleaners, fore-optics, the beam-splitter optics, and all other features of the optics which don't actually require the long arms themselves. Also some aspects of suspension performance, damping servos, and alignment systems (again to the extent that long arms aren't crucial to the tests) can be studied and shaken down. We will also be able to verify some of the vacuum properties of the design (outgassing rates, thermal behavior, and the like) of the receiver parts. We will also, in the course of doing the above tasks, be validating the designs of the mechanical, electrical, and optical interfaces.

Even with all of the above tasks performed at home, there are clearly a number of key tasks which can only be performed at the LIGO sites. After "installation", we will still need to check the optics of the long cavities (which of course is related to the rest of the optics as well), test the alignment systems, lock the interferometer, measure the noise, understand the noise by cross-correlating interferometer signals and signals from other sensors (and by doing special experiments). All this will gradually shade into "running", but even as we start archiving data there will almost certainly be an ongoing effort to understand the receiver performance as we accumulate more information.

How many people on site on average, and what is the duty cycle likely to be? First, let's consider the resident site staff. If we want an operator on duty 24 hours a day, that means 3 shifts or 4 operators to cover weekends and holidays. Perhaps one of these four is actually the site manager. (is this practical?) We probably also need a resident electronics technician. We also will need a technician for vacuum and mechanical work. Do we need in addition a computer operator or system manager, or will operators take care of that?

For scientists, the question is less clear-cut. The difficulty is that at different stages of our work, the load will be quite different. The heaviest concentration of scientists will probably come during the installation of the first receiver, as then there will be the most to do and the least known about both the receiver and the facilities. At that stage, I imagine we would have 5 or 6 scientists on site for some extended period of time.(It is hard to imagine making effective use of more people than that at one time.) Whether that means actually living at or near the site or commuting depends on the site, as well as other factors I'm not sure of right now. If the site is rather remote, it is likely we would want some sort of bunkhouse at least. The rather austere living quarters at Kitt Peak or other telescopes are an appealing model to me, but even that may be too elaborate here.

At other stages of our work, there will almost certainly be fewer scientists on site. However, I don't believe there will be many periods without any visitors. This is based on a guess that either we will always be working on either diagnostic tests of installed receivers or else on installing something new. Diagnostic work probably would only take teams of two or three, installation work probably seldom as many as the 5 or 6 involved in getting the very first receiver in.

It is probably easier to imagine life-cycle of engineering staff than of scientific staff. Civil and mechanical engineering will build to a peak of activity from design through construction and acceptance of the two sites, then will taper off dramatically. Vacuum engineering will (hopefully) follow a similar course, maybe extending some more into active lifetime of facilities. On the other hand, there will likely be ongoing need for electronic engineering at some moderate level, as well as computer hardware and software work. Continued receiver development will likely call for

ongoing mechanical engineering help for receiver (especially suspension) parts, from someone who enjoys working on small mechanisms.
If we employ an optics engineer (as distinct from scientists), the need for his/her services is probably ongoing.

Again, we can imagine several different aspects of on-site engineering work. The acceptance and first receiver installation stages will demand many staff engineers on site, but maybe only for repeated short visits. An exception is probably for the vacuum engineer, who will probably have plenty to do which will require extended tests. Another exception may be for the person responsible for the data system.
In the long term, site visits by engineering staff may become rare, unless the make-up of the engineering staff changes into a reciever support team, involved in designing and installing "engineered" versions of validated receiver designs. In that case, some members of that team will spend extended periods of time installing and debugging (from the mechanical end, at least, and maybe in terms of noise performance as well) the "engineered" receivers.

Summing up the numbers mentioned in the previous few paragraphs, it looks like there will be a time when as many as eight or nine scientists and staff engineers will be on-site. On many occasions, four to five people will be visiting.

The organization of the scientific work takes more thought. I agree with the assumption that it is hard enough to get one receiver (plus the facilities) going. But I want to include in planning for that first receiver the construction of a half-length version early on to help in the shakedown process. I think it is a valuable part of the arsenal of diagnostics.

To help think about the work more clearly, we have to think about where we are right now in the receiver design business. We have paper designs for a first LIGO receiver, variants of the designs in the appendices of the Dec 1987 proposal. There are a number of untested features which are part of our baseline design: recombined FP, recycling, optics at high power, optics at long baseline, vacuum-ready suspension of all but most rudimentary type. It wouldn't be fair to say data recording or analysis are untested, but there is plenty of work to do there too. We need to guess how much progress we will have made by proposal submission, by design freeze, by first site availability date. My guess is that at best we will have accomplished at each stage just enough to get by, i.e. by proposal we'll look convincing but know of many unsolved technical problems, by design freeze we'll specify something which has some substantial untested aspects, and by move-in we won't know the biggest problems we'll face when we first turn things on. Probably we should also assume that we'll be wrestling with some facility problems, too.

Another thing to keep in mind is that once LIGO receiver hardware is being built there will be substantial supervisory roles for staff scientists, not completely different from now but with higher stakes in terms of cost, and more work done outside. Also, there will continue to be a need for scientists to help in things which are primarily the engineers' province.

One way to approach this planning exercise is to make a list of subsystems or tasks which are major enough to warrant the better part of one person's time. Here is a try at such a list. (Some things may not take full time, some things may occupy more than one scientist.)

- Lasers
- Mirrors
- Other optics
- Suspensions
- Alignment
- Diagnostic Measurements
- Computer System
- Data Analysis (Algorithms and/or special hardware)
- Electronic Hardware for above

I'll guess that we'll be lucky to have full confidence in a receiver design by the time the second facility is available, even with everyone working full blast. If we took this fear seriously, we have to consider how many parts we order while doing the shakedown in the first facility.

I think we need to make it an absolute rule that a receiver design is understood and functioning before switching all but a small amount of effort to the next generation.

Another ground rule has to be that if, in spite of our best design efforts, there turns out to be interference between development and running, we have to guarantee our sponsor that substantial dedicated data-collection runs will be carried out between rounds of improving the receiver. If we are successful in using the elevator tanks to minimize this problem, we still need a ground rule, as has been discussed already, to limit the daily duty cycle of work in and around the tanks, say to one shift per day.

The natural place for installing the very first receiver is the corner tank. It is my guess that throughout the lifetime of the LIGO facilities the corner tank will be the place where most developmental work goes on, because of its extra flexibility, and because hairier, less well engineered things will function there better. For this reason, I would dearly wish that we would return to the idea that there would be a large flange at the bottom of the corner tank, allowing us (maybe not very rapidly) to completely expose the insides. I think the provision of a two foot wide walkway around the allowed working space, while it solves some problems, would seem very confining. This is not just a psychological statement, but a feeling based on some experience worrying about establishing the alignment of an optical system, for example. Setting up an optical lever or a transit is much easier in a less-confined space. Also, I think installation of receiver components will be unnecessarily encumbered by the restricted access of the present fixed design. (Installation will surely be awkward in any case, since we will probably be filling a large fraction of the available footprint with components, and we have to figure out a way to reach things in the middle.)

It is an important near-term task for someone to produce a drawing of what the first receiver (with full-scale components) might look like in the corner tank, to make sure that the space we are planning is sufficient. The question of whether components for full- and half-length interferometers can fit there is also very important.

Similarly, there ought to be a drawing of how similar components would be installed in a three-tank arrangement, again to verify that there is room.

Let's try to fill out a timeline for the first receiver.
The schedule above has us installing, understanding, and modifying the first receiver in the first facility for the whole year before the second site is available, and it may take longer than that.
It might be that around that time the first
receiver design is validated, and forwarded to our engineering staff to be neatened up and prepared for installation in the forward tanks at both sites. Then, the question arises, how best to use two available developmental facilities, that is, the corner tanks at the two sites.
We could
imagine several different scenarios. One is that we make copies of a developmental version of Mark II for both sites. Another is to pursue two different new designs at the two different places, with the idea that both would be eventually installed in engineered versions at both sites. Exactly how that works depends on many things, among them how many sets of tanks are installed in Phase A. A third possibility is that development work will go on at one site only, which we might choose to do if people and money were tight. I hesitate to mention this, but it might be that if resources looked like they would be scarce for the foreseeable future, we might end up designating one sight as the primary place for developmental work, with the other getting copies of validated interferometers only. This could only work, of course, if the process of validating an interferometer was something that could be sensibly done at a single site alone. It is not clear that this operations scenario would have much, if any, impact on how the facilities are designed. It might only affect the number and frequency of visits by staff people.

One other feature of two-site operation that we need to give some thought to is what sort of communications to have between the two sites. Some sort of wide-band dedicated link would facilitate coincidence (cross-correlation) studies, both by exchange of data in real-time and by the coordination of activities which it would allow. I think this will be of key importance in understanding the performance of the receivers.
If practical, the
concept might be extended to links between the remote sites and home institutions. We might even imagine approaching the status of "remote observing" which is starting to be popular in optical astronomy. I have no idea how costly this might be, but it would be a desirable feature. Even if we think we would have to postpone implementing an elaborate version until Phase B, a clear commitment to the concept might influence how we instrument the control panel, for example.