APOLLO's APO Interface: Preliminary Design

T. W. Murphy, L. Carey, S. Fisher

February 28, 2003

1 Preliminaries

This document contains an integrated response to the twenty issues raised by APO staff that need to be addressed before proceeding with the implementation of APOLLO. The sections that follow each address one or more of these points. Some sections are more complete than others, but an attempt has been made at fully covering the issues of immediate concern. The sections do not precisely follow the order of the presented list of issues, but are roughly ordered along a time-line of immediate to long-term issues.

In what follows, we will often refer to the "laser bench enclosure" and the "laser box". The laser components are all contained within an aluminum box built onto the laser bench. This is what we mean by *laser box*. The entire bench & laser box are themselves housed in an insulated enclosure, referred to as the *laser bench enclosure*.

2 Proposed Equipment Locations

APOLLO is to be a distributed system, with some components on the telescope's primary mirror cell (PMC), some components on the observing floor level, and some immediately below in the intermediate level (IL).

Because much of our equipment requires a minimum operating temperature of 10°C (50°F), all must be placed in insulated enclosures to protect against the very cold temperatures experienced in the dome environment. The thermal requirements and control are discussed in Section 3, and the impact to observers is covered in Section 17. In order to minimize the number of enclosures needed, we have consolidated all equipment into two volumes—one on the PMC, and one split between the intermediate and observing levels.

In arriving at this arrangement, we sought to simultaneously achieve:

- minimum weight on the PMC
- co-location of all equipment requiring a high degree of thermal stability
- consolidation of primary timing system (APD/CAMAC/clock) in one location
- minimum thermal loading on the dome environment and on the PMC
- minimum cable-lengths as required by certain equipment
- simplest optical path from laser to telescope
- maintain high degree of laser beam safety

Together, these goals strongly favor the back port of the telescope as the natural location for the laser. Rigidity of the relative alignment between the transmit and receive paths strongly favors placing the receiver on the laser bench. Considering the goal of consolidating the timing system, this puts the laser, receiver optics, APD detector, CAMAC crate, and GPS clock all on the telescope PMC. All the rest can go in a separate enclosure mechanically divorced from the telescope. (See the "option A" drawing previously FAXed).

Unfortunately, a number of cable length restrictions prevent us from simply putting all remaining equipment in the intermediate level. Three separate items drive this restriction:

- The two capacitor units have a maximum allowable cable length of 15 ft, translating to about **13 ft** from the laser-bench-enclosure cable-entry point.

- The CCD control box has a maximum separation from the CCD head of 19 ft. Because the CCD head is high in the laser bench enclosure, this translates to a roughly **14 ft** separation of the CCD control box from the laser-bench-enclosure cable-entry point.
- The picomotors that actuate optics on the bench have a maximum allowable separation from the motor controller of 20 ft. As these share the approximate position of the CCD, the separation from the laser-bench-enclosure cable-entry point is roughly **15 ft**.

Given that the back port reaches a position 11 ft above the observing floor when the telescope is pointed at the horizon, these three cable restrictions mean either placing these items on the PMC or placing them high enough above the observing floor to allow a short cable run at all telescope orientations. A cabinet 4 ft high behind the telescope just meets this criterion.

We should point out that even if these (four) items were to be placed on the PMC, we would still have to route some cables from the intermediate level directly up to the laser bench enclosure through a hole in the observing level floor. The limitations here are:

- The CPU–CAMAC distance is restricted to ¡ 10 m (**33 ft**—and cables *this* long are anecdotally not guaranteed to work).
- The CPU–GPS Clock distance is theoretically restricted to the 4 m (13 ft) GPIB limit, though 8 m (**26 ft**) GPIB cables are available, and it is this longer cable that we will try to use with APOLLO.

If these items were to be placed on the intermediate level below the back port, a conduit would need to breach the floor and carry the cables up to the laser bench enclosure from a position behind the back port to keep the run short enough. It was this line of thought that ultimately resulted in the cabinet idea: simply take that conduit and enlarge it to cabinet dimensions so that it can accommodate the capacitors, CPU, etc.

Table 1 lists the various electrical components and their various properties and requirements as relevant to positioning. PMC indicates a position inside the laser bench enclosure on the back port. CAB means a location in the cabinet behind the telescope. IL means a location in the intermediate level—on a platform above the steel beams, just below the back port, and situated directly below the cabinet. Table 2 lists the weights of the passive (mechanical) equipment to go on the PMC.

The total PMC weight in Table 2 is larger than we anticipated—partly because we were conservative with our guesses, and partly because it's the first time we've deliberately added it all up. Previous estimates ranged between 400–500 pounds. We are aware that this is a potentially prohibitive amount of weight, and we should devote a good deal of discussion to this issue during the March design review. We need to be explicit about what the limitations are and why. If the elevation drive motors can not handle this load (plus counterbalance), let's look at the possibility of new elevation motors, and the expense of such a switch. If changes of this sort *must* be done to accommodate APOLLO, then APOLLO is willing to contribute substantially to the cost of the effort.

3 Thermal Properties and Control

We have carefully considered the thermal requirements both of APOLLO and of the 3.5 m telescope/environment. Our goals can be stated as:

- Maintain laser bench at nearly constant set-point temperature.
- Maintain CAMAC crate and GPS clock at stable temperature during operation.
- Keep auxiliary electronics within specified operating ranges (usually $> 10°$C).
- Limit heat flux into dome to < 50 W both during off-state and warm-up for external temperatures down to $-15°$C (5°F).
- Limit heat flux into dome to < 100 W during active lunar ranging operation.
- Limit heat flux directly into PMC to < 10 W in all operating stages.

Table 1: Properties and Locations of APOLLO Electronics Components

Component	Size (in) H×W×D	Approx Weight (lbs)	Min.Temp. (°C/°F)	Location	heat load (W)	Elec. Sup.
Laser Bench	14 × 36 × 48	250	10/50 stable	PMC	100–200	—
CAMAC Crate	13 × 13 × 16	45	10/50 v. stable	PMC	125	115 VAC
GPS Clock	2 × 17 × 11	10	10/50 stable	PMC	25	115 VAC
APD Array	8 × 8 × 3	5	10/50 stable	PMC	< 10	(DC)
CCD head	$\phi 3.5 \times 2$	1	−15/5 (?)	PMC	< 10	—
CCD Controller	4 × 12 × 10	10	−10/15 (?)	CAB	7	115 VAC
CPU (`houston`)	15 × 9 × 20	30	5/40 (?)	CAB	75	115 VAC
Picomotor Driver	4 × 9 × 13	8	10/50	CAB	< 25	115 VAC
T/R Motor	4 × 4 × 6	8	−10/15	PMC	40	48 VDC @1A
DC Power Supplies (5)	4 × 8 × 9	7 ea.	0/30 (?)	CAB	< 5 ea.	115 VAC
Laser Capacitor Banks (2)	5 × 18 × 27	30 ea.	5/40 (?)	CAB	< 50 ea. (?)	208 VAC/3ϕ
Laser Electronics Rack	29 × 25 × 28	200	10/50 (?)	IL	~ 1 kW (?)	208 VAC/3ϕ
Laser Chiller	27 × 17 × 24	200 (ship)	10/50	IL	2.5 kW	200–230 VAC
Bench Chiller	23 × 13 × 21	140 (ship)	10/50	IL	1.5 kW	115 VAC
Heat Exchangers (2)	6 × 23 × 4	9 ea.	−20/5 (?)	PMC	−200 W ea.	115 VAC

Table 2: Weights of the components slated for the PMC

Component	Approx Weight (lbs)
Laser Bench & Laser	300
Electronics	60
Optics	25
Heat Exchangers	20
Structural Metal	90
Insulation	40
T/R motor	15
Corner Cubes	10
Total	560

- Dump operating heat into intermediate level with good mixing properties.

To motivate our desire to maintain a stable temperature, consider Figure 1. APOLLO works largely in a *differential* mode, meaning that by comparing the "range" to the corner-cube in the telescope (Section 13) to the lunar range, any systematic errors are canceled. Most temperature variations don't have an impact, then, unless the variations are faster than the 2.5 second round trip travel time to the moon and back. The time-to-digital converter (TDC)—located in the CAMAC crate—is the *one* unit in the APOLLO scheme that is not purely differential, and therefore is susceptible to temperature excursions.

The TDC converts the time interval between START/STOP events spanning 0–100 ns into a 12-bit integer between 0–4095. Assuming for the moment that the TDC is purely linear, the mapping between the reported integer (bin) and the actual time interval is characterized by a slope and an offset (intercept). Figure 1 shows how these parameters vary with temperature. To interpret this, let's say the corner cube measurement just happens to be reporting intervals like 40 ns on the TDC, and that the lunar returns are reporting 60 ns typically. The 20 ns$\approx$ 400 bins between the two can't be taken at face value, given the temperature-dependent mapping from bins to nanoseconds. The net result is that a gain error of 0.014 bins/ns translates to an interpreted lunar range error of one millimeter. This corresponds to about 2°C in temperature, indicating that we would like to hold the temperature of the TDC to within ±1°C. Note that the warm-up of the unit suggests a 30-minute requirement to reach this sort of stability (lower right in figure).

3.1 APO Temperature Records

The temperature record at APO provides the constraints against which we have formulated our thermal design. Figure 2 shows the minimum and maximum daily temperatures both outside and inside the dome enclosure. The sinusoidal fits to the enclosure temperature (in °F) are:

$$T_{\rm min} = 42.92 + 15.16 \sin\left[2\pi(y - 0.3004)\right],$$

$$T_{\rm max} = 53.22 + 16.66 \sin\left[2\pi(y - 0.2886)\right],$$

meaning that the overall average enclosure temperature is around 48°F, or 9°C (and indicating a transition at 0.3 yr = 108 d = Apr 17). Because many of our electronics components have a minimum operational temperature rating of 50°F (10°C), we have adopted this as a reasonable set-point for much of our equipment. Note that the maximum enclosure temperature is reported as 79.8°F on 6/28/1998, and the minimum enclosure temperature is 13.3°F, on 1/18/2001. The reported outside extremes are 88.8°F on 6/27/1998 and −5.4°F on 12/17/1996.

3.2 Thermal Leakage During Off-State

If we want to keep our equipment at 10°C in an environment that may get down to −15°C (5°F), and keep the heat flux below a certain level, P, then we need substantial insulation around such equipment. Using

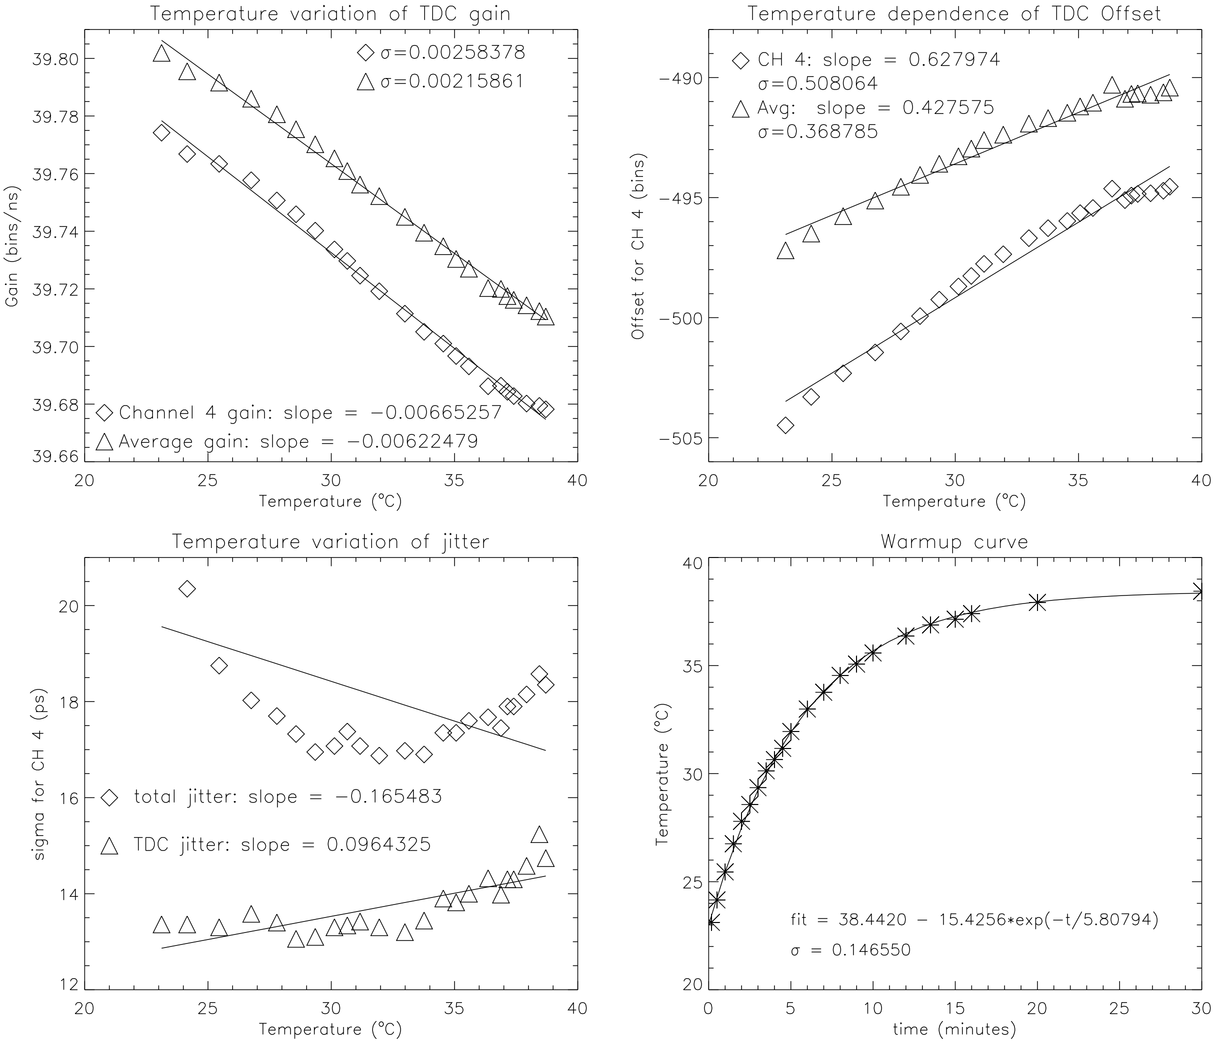

Figure 1: Temperature response of the TDC during warm-up. For this test, the temperature sensor was closest to channel 4. The two parameters characterizing TDC calibration—slope and offset—clearly are strong functions of temperature.

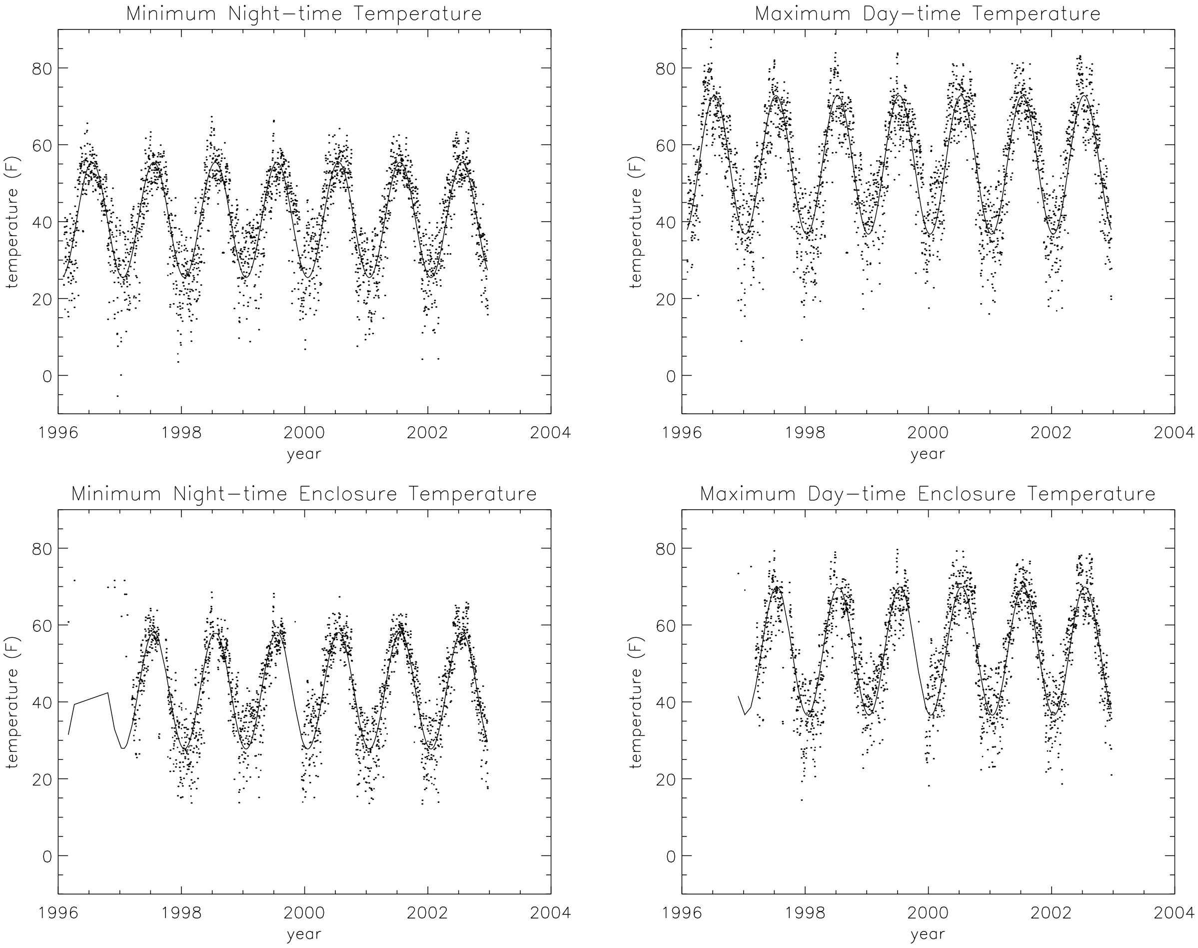

Figure 2: Temperature record for the APO 3.5 m telescope for the last several years.

the commonly-encountered Imperial units for insulation, we can express the *R-value* needed to limit the flux to P by: $R = 5.678A\Delta T/P$, where A is area in m^2, ΔT is the interior/exterior differential in °C, and P is in Watts. Given the area of our laser bench enclosure (using the mean of the interior area and the exterior area) of 7.2 m^2, we need $R \approx 20$ to limit P to 50 W, which is achieved by 3-inch Thermax insulation ($R = 24.4$). A small "space" heater in the laser bench enclosure can maintain the interior temperature at 10°C, on average consuming less than 50 W even on the coldest days.

The thermal leakage from the cabinet items is also an issue. Also maintained at or above 10°C, if this enclosure were not separately vented (as described below) the insulation would need an R-value of about 12 (area is 4.3 m^2)—corresponding to 2-inch thickness—to hold below 50 W. But now we have two 50 W sources—the laser bench enclosure *and* the cabinet. The insulation necessary to limit both enclosures to a sum heat loss less than 50 W would require thicker insulation than is available in single sheets—not to mention being more cumbersome than we desire.

The cabinet issue is solved by a double-walled enclosure. The outer wall is a durable metal structure, with an air gap between it and the insulated box within. Air is pulled from fans in the intermediate level through the air gap, with inlets at the top of the cabinet to admit ambient dome air. In this way, the outer surface of the insulation is held at ambient temperature, eliminating radiative coupling to the dome. All heat flux is then convected into the intermediate level. The residual leakage into the dome may be expected to be less than 10 W.

The cabinet enclosure effectively extends into the intermediate level, forming a contiguous volume with the enclosure containing the laser electronics rack and chillers. Here the surface area becomes rather large, as the chillers need ample room in which to vent. 2-inch insulation with an R-value of 14 and an area of 20 m^2 would pass about 200 W into the intermediate level when $\Delta T = 25$°C. Again, a heater can maintain the set-point, with enough internal circulation provided to assure proper mixing.

It is worth pointing out that the total heat load is maintained below about 50 W at temperatures 10°F *below* the *coldest temperatures yet recorded for the telescope enclosure.* At a more typical cold night of 20°F, the heat leakage is reduced by a factor of 2/3, so that our 50 W figure is reduced to 33 W.

There is also an issue of direct thermal coupling to the PMC through the supports that mechanically interface the laser bench to the telescope. The stainless steel flexure legs we have constructed only conduct 0.01 W per °C differential, so that the maximum heat load from six such legs would amount to 1.5 W at the coldest temperature.

Also worth exploring is the radiative coupling between the bottom of the laser bench enclosure and the PMC. Given the area overlap of roughly 1.5 m^2, roughly $(1.5/7.2) \times 50\,\mathrm{W} \approx 10$ W emerges from this wall in convective and radiative forms. At these temperatures and differentials, free convection tends to be less efficient than radiation, so one might expect the full 10 W to be coupled directly to the PMC at high surface emissivities. Making the bottom of the laser bench enclosure shiny reduces this number by at least a factor of three to < 3 W.

3.3 Thermal Leakage During Warm-up

During the roughly 30-minute warm-up phase prior to lunar range operation, we want to maintain aggressive control of thermal leakage so that no ill effects impact other observers. For us, warm-up means that *everything* is on, except for the transmit/receive (T/R) motor. During warm-up, the laser rods are flashing, green light is being generated, the CAMAC crate is on, and the chillers are cranking out heat.

Almost all of the ~ 1500 W of power generated by the laser rods is pulled away by the laser cooling water. The residual 20–40 W diffuses into the optical bench and into the air within the laser. The auxiliary electronics components that are part of the laser additionally contribute about 20–30 W. Also within the laser bench enclosure is the CAMAC crate, GPS clock, and APD arrays, expending a total of about 150 W. All of this heat is to be carried away by cooling water circulating to a pair of heat exchangers within the laser bench enclosure, so that no effect should be seen on the part of observers.

The ~ 5 kW of heat generated in the intermediate level is clearly a concern. The enclosure here will be vented such that the internal temperature will not rise (cold night scenario). Thus the cabinet leakage into the dome is not changed, and is still kept low by the double-walled cabinet ventilation scheme. But the heat that *is* vented into the intermediate level could potentially warm the air there enough to affect the dome environment. Mark Klaene performed a calculation that indicated a 2.5°F rise in the intermediate level temperature during one hour of operation, but that recovery would be very fast—within 5–10 minutes after shut-down. In order to realize this performance, we *do* need to ensure adequate mixing of the vented air so that we don't end up with a hot-spot under the observing floor.

3.4 Thermal Leakage During Full Operation

This scenario is really no different from the warm-up aside from the influence of the T/R motor. This motor is currently intended to sit outside the laser bench enclosure, exposed to ambient air. The motor should be able to operate at the coldest dome temperatures. The motor puts off 40 W when spinning the T/R mirror at 20 Hz. Since we are the only observers using the telescope at this time, we will live with whatever negative consequences this brings. If we suspect this is causing trouble, we can annex the motor into the laser bench enclosure with a small extension (also allowing the motor shaft to shorten significantly).

The heat load directly into the PMC through the motor mounting bracket can be reduced to a negligible value. Rough calculations indicate a heat load of about 2 W given a stainless steel bracket having a 0.5-inch by 4-inch cross section, and a motor that is 30°C above ambient (consistent with our experience). A clever bracket design would reduce the cross-sectional area well below the figure above. There is no reason this heat flux term could not be reduced to less than 1 W. There is still the matter of radiative coupling. A worst-case calculation indicates a total radiation of 15 W, at most one third of which would impinge on the PMC. Therefore, less than 5 W would find its way to the PMC. Using reasonable guesses for emissivities, ($\varepsilon = 0.8$ for both the motor and the telescope), this 5 W is reduced to 3 W of coupled power.

3.5 Heat Loss in Hoses

One element that has not been treated in the above analysis is the heat load imposed by the coolant hoses to the laser bench. These consist of two sets (See also Section 4):

- a coolant loop to the laser heads containing de-ionized water, operating at 29.5°C (85°F) in warm-up/operation, and 10°C (50°F) during quiescent times,
- a coolant loop to the laser bench heat exchangers containing a 50/50 solution of ethylene glycol and water, running at 0–10°C (30–50°F).

If uninsulated, these hoses could give off almost 3 W per °C differential apiece over 5 m lengths. This could reach a number as high as 135 W apiece for the warm hoses during warm-up/operation, and 75 W during the off-state! Simply adding pipe insulation helps, but only by a little more than a factor of two. The best solution is to place the hoses in a ventilated tube such that the outsides of the (moderately insulated) hoses are held at the ambient temperature and all convected heat is carried away. Since these hoses could emerge from the cabinet, and the cabinet wants to pull in ambient air, these two functions can be combined so that the cabinet air draw comes through these ducts. Calculation of the efficiency of this ventilation will have to wait until we have selected the ducting scheme and airflow rate.

3.6 Thermal Control of Laser Bench Enclosure

Our scheme assumes that we have the ability to maintain a constant temperature within the laser bench enclosure even when the laser and other electronics are generating heat. This certainly is not a trivial matter, but we believe we can effectively carry this out. During quiescent times, it is rather easy to maintain temperature—provided a set-point warmer than ambient—simply by including a small heater on a thermostat, and ensuring adequate circulation throughout the volume.

As mentioned before, the bulk of the heat generated by the laser is carried away by cooling water. This water is run through a heat exchanger in the laser electronics rack which itself is hooked to a closed-cycle chiller (Section 4). Both of these units are on the intermediate level. The remaining 190–220 W of heat generated/left-over in the laser bench enclosure will be carried away by additional coolant lines. We have selected small forced-air heat exchangers made by Noren Products, Inc. to handle the laser bench cooling. These devices provide 25 W of cooling for every 1°C difference between ambient and coolant temperature. There is no stated limit to how far this can be pushed, but 200 W per unit is certainly within their range of capability. We expect to run at a temperature differential of $-5°C$, or a fluid set-point of $\sim 5°C$. At a flow rate of 1 gallon-per-minute, the exit temperature of the fluid is negligibly different from the input.

A serial interface to the cold-loop chiller will enable us to dynamically control the fluid set-point temperature so that we can vary the amount of cooling necessary to maintain a fixed temperature inside the enclosure. More information on the chillers can be found in Section 4.

There is a subtlety to this scheme in that we want the GPS clock to run all the time (at 25 W). In the absence of heat flow out of the box, this would raise the internal temperature by about 0.6°C/hr. If no attempt were made to remove this heat from the enclosure, it would drive a 15°C differential between internal and ambient temperatures. On very cold nights this is fine since we need to supply heat to maintain temperature anyway, but on mild nights ($-5°C < T < 10°C$), we don't want to have to run the laser bench chiller (at 1.5 kW consumption) just to take care of this paltry 25 W source. An alternative solution is to put a small pump in-line with the laser bench cooling system and to keep the intermediate level enclosure a little colder (by just 1°C) than the laser bench so that we send slightly chilled coolant to the laser bench whenever this auxiliary pump is on. If we then put the Noren heat exchanger fans on a thermostat, they only activate when a bit of cooling is needed. The 1°C differential on two Noren coolers is enough to remove 50 W if the fans were on full-time. Of course we may end up tuning many of these set-points for best performance and minimal energy expenditure, but this at least illustrates the idea.

One other point to mention here is that we want to keep a constant, light flow of water in the loop that cools the laser rods. If left stagnant, the portion of these hoses between the cabinet and the laser bench enclosure could freeze. We could either accomplish this with an auxiliary pump as for the Noren loop, or by letting the laser cooling unit handle this function.

3.7 Thermal Control of Cabinet/Intermediate Enclosure

Because the equipment in the cabinet and in the intermediate level is not particularly temperature sensitive—aside from having minimum operating temperatures—we don't have to be as careful about controlling the temperature in this enclosure. Specifically, we don't have to worry about active cooling. The computer is the only device that is always on, running at 75 W. If the temperature inside the enclosure drops to less than the ambient temperature, we simply need to run a heater (a few-hundred Watts capacity) to maintain our minimum temperature. Any time the temperature is above our desired set-point (roughly 10°C), we simply open large louvers to let the internal components warm up to ambient conditions.

When in (identical) warm-up and full-run modes, we open these louvers wide, and let as much heat out as we can. It is probable that the 5 kW generated inside this enclosure will be enough to keep all the equipment warm even on very cold nights. If this turns out to be less than true, we simply open the louvers less widely to achieve the right (approximate) balance. We only need to worry about staying above the minimum (10°C), but don't need to dial in any exact temperature. More detail on this scheme is presented in Section 7.

3.8 Cool-down Phase

We will probably also want a cool-down phase in which only the laser bench chiller is left to run, thus continuing to pull heat out of the laser bench components. Some of these components warm up for operation, and we would like to extract this heat rather than permit it to diffuse into the rest of the laser bench, thus raising the temperature above our desired set-point. One-half hour will probably be sufficient time for cool-down. This phase probably isn't necessary on the colder nights, when we can rely on thermal loss to carry the excess heat away (always less then 50 W, of course).

3.9 Multiple Set-points

This general scheme works well enough for half of the year, when the dome temperature is less than 10°C. But if we wanted to maintain this set-point on the laser bench during warmer periods, we would have to actively *chill* during a large fraction of the time. Rather than do this, we propose a dual set-point operation, with a warm-month set-point of 24°C (75°F). Doing so would require a re-alignment of the laser, but this should probably happen at least twice per year anyway. In terms of thermal leakage, our ability to withstand

$\Delta T = -25°\text{C}$ differentials translates to being able to cope with freezing temperatures during the warm months.

A look at the temperature record indicates that the switch-overs would typically happen around mid-April and mid-October. Over five years of data, only five nights plummet below freezing after the sensibly timed switch to the higher set-point. The winter is a bit crueler, in the sense that daily maximums in the dome frequently exceed 10°C, even in the dead of winter. But they hardly ever exceed 13°C during times when the set-point would be at 10°C.

A simpler scheme would of course always have the laser at roughly room temperature (which would make the laser company very happy), but this is much harder to handle during the cold months, resulting in much thicker insulation, difficulty keeping the hoses from releasing heat, and a higher heating bill to keep both laser bench and intermediate level enclosures toasty warm.

Another disadvantage to the single warm set-point scheme is that when one wanted to perform work on the laser, there could be a very large temperature differential between the interior and ambient environments. Opening the enclosure could thermally shock the optics bench, likely making it impossible to do any effective tuning.

3.10 Remaining Flexible

Managing the thermal state of this complex, distributed system in an environment whose temperature spans such a large range is a very difficult challenge, especially if we wanted to maintain tight tolerances. In truth, our system will not be perfect. Especially tricky are those days in which the temperature swings up to near or above our set-point. Because we don't want to constantly run the cooling loops (and in fact will have *no* active cooling capability in the IL enclosure), we must simply live with such excursions.

One way to imagine handling these scenarios is to simply take stock of the thermal state at the time of turn-on, and proclaim that as our set-point for that particular run. For such small excursions, we can be reasonably sure that the laser alignment is still good (it is built to withstand normal lab temperature variations), and we can calibrate our timing electronics to compensate. This doesn't mean we can be happy letting the temperature run wild, but we *can* tolerate the occasional departure.

4 Plumbing Considerations

We are forced to use water to cool the laser flashlamps/rods, which together generate nearly 1500 W of power when operating. It would be nice to utilize the same water cooling circuit to manage the thermal state of the laser bench itself, but the higher set-point of this circuit precludes such a combination. We explored air cooling schemes, but found this to be cumbersome, and additionally had concerns about dust, moths, and transporting the warm air to the IL without bleeding heat into the dome. So we end up with two plumbing loops going to the telescope.

During operation, the laser rods/flashlamps are cooled by circulating de-ionized water (DI water), with a set-point temperature of 29.5°C (85°F) through the laser flashlamp/rod assemblies. This temperature offset choice is made for thermal equilibrium of the laser while operating. It is this equilibrium condition that drives the warm-up period of the laser. We tried to get the laser company to agree to a lower set-point (some reasonable ΔT above the laser bench set-point), but they convinced us that the optical properties of the laser head *really do* need this absolute temperature. The DI laser-rod closed-loop is referred to here as the *hot loop*, for obvious reasons.

Our proposed solution employs two heat exchangers in the laser bench enclosure, fed by chilled coolant. Using two such units, we can achieve 400 W of heat removal using water chilled to about 8°C below the laser bench set-point. Because this temperature is dangerously close to freezing, and because we would like to cease circulation of this fluid during cold and inactive periods, we have chosen to use a 50/50 solution of ethylene glycol and water. This loop is called the *cold loop*.

A third cooling loop—the *dump loop*—takes the heat away from the laser DI (hot) loop via the Neslab Merlin M-75 chiller. This exchange takes place in the lowest unit in the laser electronics rack, referred to as the "cooling group" (CG-604). Figure 3 shows the general plumbing layout for APOLLO. For maximum efficiency of the M-75 chiller, the dump loop will be run at 20°C (68°F).

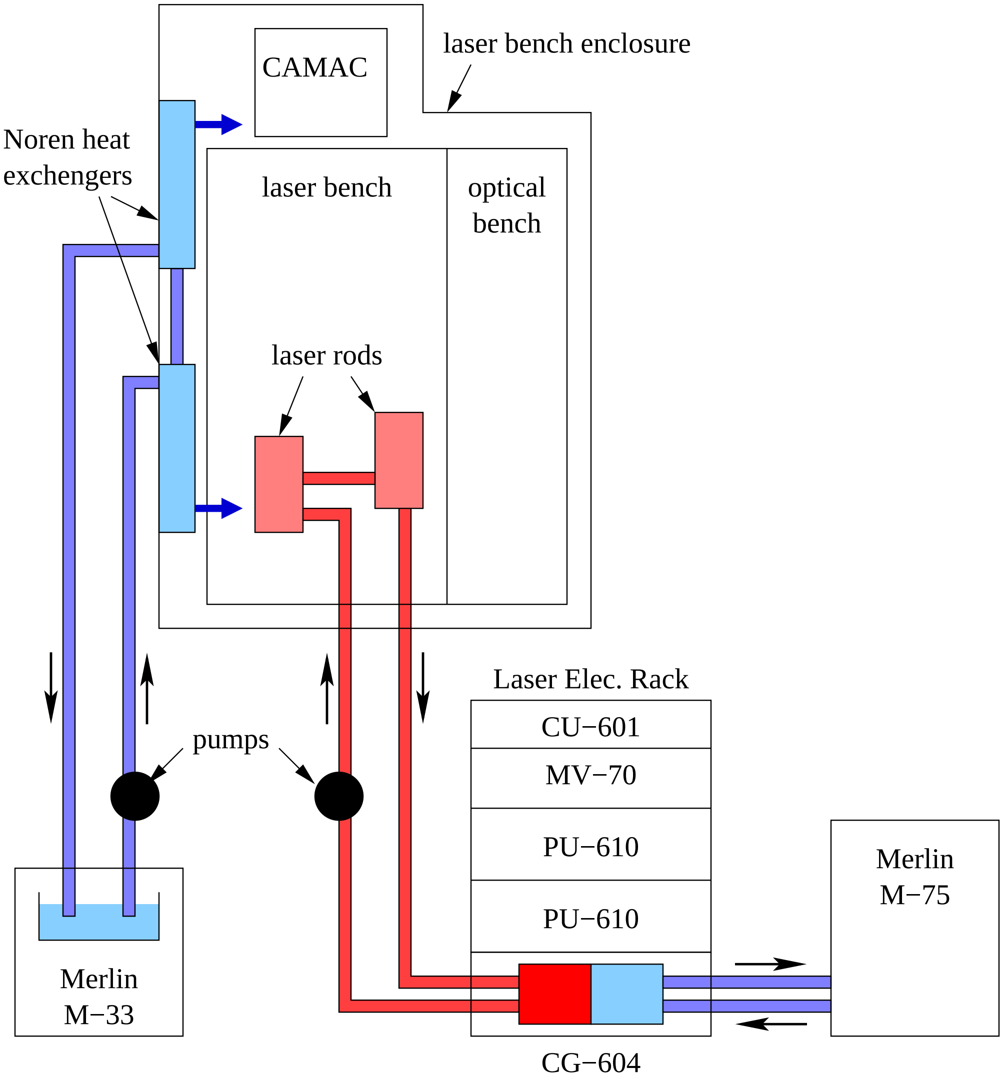

Figure 3: The APOLLO plumbing scheme. The three closed-loops are the ethylene-glycol "cold" loop to cool the laser bench enclosure, the DI water "hot" loop to take heat from the laser rods, and the "dump" loop to remove the heat from the hot loop.

We use two Merlin series recirculating chillers built by Neslab. The M-33 used for the cold loop has a maximum cooling capacity of 1250 W, though at altitude we can expect only 400 W for the roughly 0°C cold-set-point regime (cold-loop has to be colder than the 10°C laser bench set-point), and 750 W for the warm-set-point regime. The M-75 unit used for the hot loop should provide 1700 W of cooling capacity at a 20°C set-point. These figures assume a 50/50 ethylene glycol solution is used. If we were to use water in the M-75 unit, we would see greater than 1850 W capacity, at the risk of freezing in case of heater failure. The Noren heat exchangers are simply finned aluminum structures through which air is passed. A single 5-inch fan pushes the chilled air into the enclosure.

The two Neslab Merlin chillers have their own pumps, as does the laser cooling group. The auxiliary pumps pictured in Figure 3 perform the following two functions. The hot-loop pump maintains a slow circulation at all times, so that the temperature of the fluid in this loop is maintained near the temperature of the cooling group's reservoir. This not only prevents freezing in the lines, but also minimizes thermal shock to the laser rods when the system is first turned on. The second cold-loop pump is intended to remove excess heat from the laser bench enclosure—largely deposited by the always-running GPS clock. By keeping the M-33 reservoir slightly cooler than the laser bench enclosure, we get the requisite amount of cooling simply by forcing circulation. Another way to put this is that provided cold loop circulation, we can expect a differential between the laser bench enclosure and the IL enclosure of no more than about 1°C.

Because there is a heat load issue from the hot loop, we want to keep the exposed length short. Thus we prefer a routing similar to that of the cabling (Section 9): an umbilical emerging from the top of the cabinet and feeding into the lower side of the laser bench enclosure. The cold-loop, however, runs colder, is not subject to freezing, and is mostly off during the really cold times—when the laser bench enclosure requires no cooling. So this may be routed up the fork tines and around the primary mirror cell in (mostly) rigid tubing. It would be a greater tragedy to have the chemically corrosive cold loop burst than to have the plain-water hot-loop burst. Fortunately, the temperatures favor this more robust arrangement for the cold loop. The dump loop is a very short loop, and presents no obvious danger to the telescope should a leak or burst occur.

Addressing the possibility of a leak/burst, we point out that the reservoir volumes of the M-33, M-75, and CG-604 are 0.5 gallons, 0.5 gallons, and < 2 gallons, respectively. Additionally, a 25 ft hose with 0.5 inch I.D. contains about 0.25 gal. So the total fluid volume in all loops is roughly 4 gal. Messes are possible, but catastrophes unlikely. Freezing of the DI loop is our worst enemy, and we need to pay more attention to this than we have thus far.

We may want to outfit a number of our cooling loops with flow interlocks/alarms. The laser hot loop is already well outfitted in this regard, shutting off the laser when flow is interrupted.

5 Electrical Power Requirements

Table 3 lists the APOLLO electrical equipment requiring site power, and the type of service needed. Note that the current requested is not necessarily the current used, but the current that should be made available to each device. The usage times are denoted WU for warm-up, LLR for active ranging, and CD for cool-down. Location abbreviations follow the same convention as used in Table 1. The 1% and 2% duty cycles were derived from an estimated 30 min and 60 min run-time every other day.

By consolidating the capacitor banks into the cabinet/intermediate level enclosure, we have limited the 208 VAC run to a small region. Also, most of the 115 VAC is required in this structure, with only three (as yet identified) items on the PMC requiring AC power.

6 Mechanical Mounting Issues

This section only treats the mounting of equipment on the PMC. The mounting of the equipment in the cabinet/intermediate level is not critical, and can be accommodated in a variety of ways.

A chief problem in mounting the laser bench to the PMC is the variation in temperature between the PMC and the laser bench. As a result of this, there is a reasonably large differential contraction between the two. As the laser bench remains at a constant temperature, it stays the same length while the PMC changes scale beneath it. For an extreme 30°C differential, the roughly 4 ft span of the PMC under the laser

Table 3: Electrical Power Requirements for APOLLO

Component	Voltage	Phases	Current	Duty Cyc.	Runs When:	Loc.
Laser Elec. Rack	208 VAC	3	6–10 A/phase	2%	WU, LLR	IL
Laser Cap. Banks	208 VAC	3	¡1A/phase (?)	2%	WU, LLR	CAB
M-75 Laser Chiller	200–230 VAC	1 (?)	12 A	2%	WU, LLR	IL
M-33 Bench Chiller	115 VAC	1	11 A	3%	WU, LLR,CD	IL
Circulation fans (4)	115 VAC	1	0.25 A ea. (?)	0–100%	as needed	IL/CAB
Auxilliary Pumps (2)	115 VAC	1	0.5 A ea. (?)	10% (?)	as needed	IL
CPU	115 VAC	1	2 A	100%	full-time	CAB
Picomotor Driver	115 VAC	1	0.5 A	1%	LLR	CAB
CCD Controller	115 VAC	1	0.5 A	1%	LLR	CAB
T/R motor P.S.	115 VAC	1	0.75 A	1%	LLR	CAB
DC Power Supp. (5)	115 VAC	1	0.25 A ea.	2%	WU, LLR	CAB
Noren x-changers (2)	115 VAC	1	0.25 A ea.	2%	WU, LLR	PMC
CAMAC Crate	115 VAC	1	1.5 A	2%	WU, LLR	PMC
GPS Clock	115 VAC	1	0.5 A	100%	full-time	PMC

bench changes by 0.017 inches (about half-a-millimeter). A rigid mount would therefore impose stress in the laser bench, possibly affecting alignment. For this and other reasons, we have elected to employ a kinematic mount so that the bench may live a stress-free life.

The laser bench will be kinematically assembled onto a triangular mounting frame. This sub-assembly can then be mounted as a unit onto the back face of the primary mirror cell. The mounting system between the bench and the mounting frame is a 6-flexure design. The flexures are orthogonally oriented in a 3, 2, 1 arrangement, with the set of three flexures mounted perpendicular to the face of the bench, (its least rigid axis). Two links will be oriented parallel to the longest axis of the bench and one parallel to the shorter axis of the bench face. The five load bearing flexures (the sets of 3 and of 2) are all designed to be loaded in compression and each flexure will be enclosed inside a “fail-safe” mounting. This mounting consists of two concentric tubes—one nested inside the other—with one tube extending from each end of the flexure. The tubes serve the dual function of preventing over-travel (and consequent over-stressing) of the flexure perpendicular to its axis, as well as providing a backup support system for the bench if one or more flexures were to fail. (Note: This tube cavity is recognized as a potential moth trap.) The bench mounted tubes can also provide a fixed surface around which to seal the insulating box.

7 Ventilation and Moth Control

The laser bench enclosure itself will have no intentional airflow. This keeps dust levels on the optics down, which is a desirable property of a high-power laser enclosure. Clearly there must be breaches in the insulating shield to allow:

- mechanical mounting structures to access telescope,
- light to get into/out of enclosure,
- hoses and cables to pass in/out.

The mounting structures will penetrate through minimal diameter holes in the insulating shroud, and batting will be stuffed into the gap in such a way as to prevent significant airflow. This should also hopefully keep the moths at bay. The light will emerge through a small (2-inch) window in the side of the box. The hoses and cables can be treated in a way similar to the mechanical mounting structures.

We may have problems with moths getting splattered against the quaternary mirror (M4), but we will have to simply keep an eye on this and hope for the best. A second window can always be mounted between this mirror and the tertiary, allowing M4 to sit in a sealed space.

For the intermediate level, the enclosure has three modes: *sealed*, *venting*, and *wide-open*. The sealed mode is used when the external temperature is less than 0°C. Below this temperature, the always-running CPU at 75 W is not sufficient to maintain temperature, and a separate heater will maintain the internal temperature at about 10°C. The venting mode is to be used when the ambient temperature is between 0°C and 10°C, in which case the CPU heat would tend to warm the interior of the box above the set-point. For this, a fan would draw in ambient air, creating positive pressure inside the box, forcing air out through whatever cracks exist. The wide-open mode is used either when the ambient temperature has exceeded the desired set-point (rare if we manage it right), or when the laser is in operation. This is necessary to keep the 5 kW of generated heat from warming the interior of the box excessively.

In all cases, there will be small fans inside to keep air circulating within the box, so that no hot-spots develop. This is especially important for the "chimney" (cabinet) poking into the observing level, where heat would naturally collect. During laser operation, larger fans can affect aggressive circulation so that the heat is indeed flushed out into the intermediate level.

It is not clear yet whether the wide-open mode should be accommodated by louvers on the walls of the box, or by leaving the front doors open, or some combination. The disadvantage of louvers is that they are poor insulators when we want to operate in sealed mode. Perhaps the best technique would be to have the side-walls hinged on the top so they can be propped out by the desired amount. The "front" doors may either hinge the same way, or hinge on the sides like normal doors. In either case, the ventilation would be through large gaps produced by the hinged doors/walls.

None of these cases provides any protection against moths. Dare we wait and see if this causes a problem?

8 Enclosure Dimensions

See attached drawings (For now, refer to previously FAXed drawings).

9 Electronics Interfaces

APOLLO consists of a relatively complex arrangement of electronics, using a variety of interfaces (see Figure 4). This section describes a number of the electronics interfaces on APOLLO.

A brief word about proposed routing. Clearly a number of cables will run from the cabinet directly to the laser bench enclosure, since the purpose of the cabinet is to accommodate a variety of short cable runs. Two electrical umbilicals are envisioned: a first umbilical designated **U1** that carries data interface cables, signal cables, and DC power cables. The latter will be separated into its own shielded run designated **U1B**, with **U1A** carrying all interface and signal cables. A second umbilical, **U2**, carries all cables to the laser. By keeping these runs separate, we can guard against crosstalk when the laser fires. A similar umbilical will carry the hot-loop hoses, but this is discussed elsewhere (Section 4).

9.1 Data and Control Interfaces

Table 4 lists the interface cables associated with the computer called `houston`.

Many of the devices in Table 4 have been introduced elsewhere. New items are:

- the CU-601, which is the control center at the top of the laser electronics rack,
- the Dataq DI-720, a data acquisition module hooked to `houston`'s parallel port,
- and an ethernet hub of unknown nature or location.

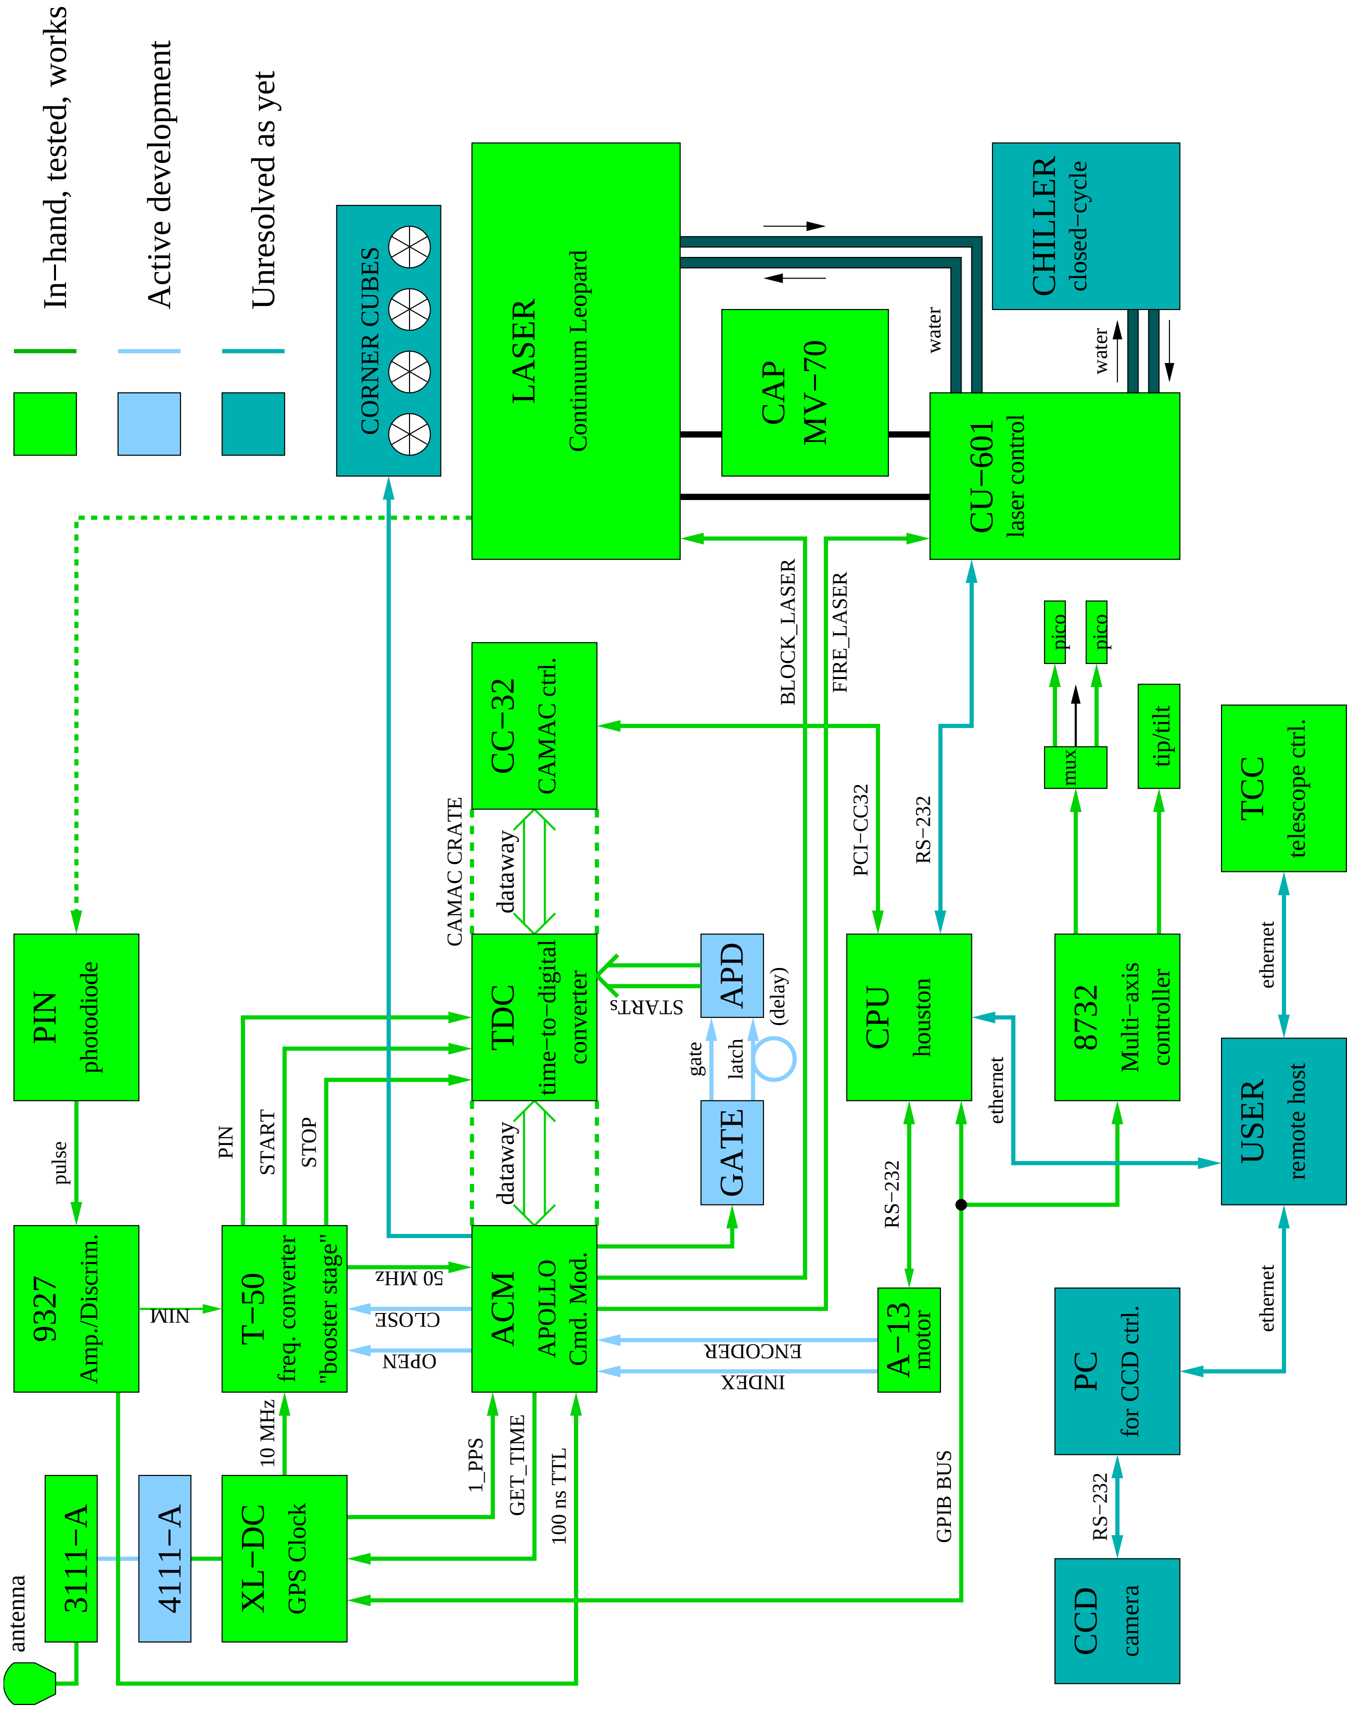

Figure 4: APOLLO Electronics/Timing System

Table 4: Interface cable types and routings for APOLLO

From	To	Interface	Length	Cable	Purpose	Routing
CPU	CAMAC	custom	8 m	SCSI-2	data, control	U1A
CPU	Picomotor	GPIB	2 m	IEEE-488	control	CAB
Picomotor	GPS Clock	GPIB	8 m	IEEE-488	data, control	U1A
CPU	CCD Ctrl	RS-232	2 m	multi-cond.	data, control	CAB
CPU	T/R motor	RS-232	8 m	multi-cond.	control	U1A-PMC
CPU	CAMAC	RS-232	8 m	multi-cond.	program	U1A
CPU	CU-601	RS-232	4 m	multi-cond	control	CAB-IL
CPU	M-25	RS-232	4 m	multi-cond	control	CAB-IL
CPU	DI-720	parport	1 m	IEEE-1284	data, control	CAB
CPU	HUB	ethernet	?	CAT 100	data, control	CAB-?
CCD control	CCD head	custom	19 ft	multi-cond.	data, control	U1A
Picomotor	motors	custom	20 ft (×2)	IDC mod.	control	U1A

Table 5: Signal lines and routing

From	To	Cable	Length	Purpose	Routing
T/R motor	CAMAC	RG-58	2 m	encoder	makeshift
T/R motor	CAMAC	RG-58	2 m	index	makeshift
Control Room	4111A	fiber	100 m	GPS feed	U1
CAMAC	CU-601	RG-58	10 m	charge	U1A
CAMAC	CU-601	RG-58	10 m	fire request	U1A
CAMAC	corner cubes	RG-58	5 m (×2)	deploy/stow	makeshift
DI-720	laser bench	multi-cond.	8 m	temperature	U1B
DI-720	points within	multi-cond.	4 m	temperature	CAB/IL
DI-720	laser bench	multi-cond	8 m	DAC out	U1B

9.2 DC Power Supplies

Many of the electronics for APOLLO require DC power, provided by linear AC-DC modules integrated into small box-packages and located in the cabinet near the computer. Rather than list all the DC lines here, it will just be pointed out that most of these lines will be routed through the U1B umbilical to the laser bench, and others will power small opto-isolators within the cabinet. The DC lines will be shrouded in a well-shielded braid to avoid unwanted pickup from nearby signal lines.

9.3 Signal Lines

A number of signal lines within the apparatus are used to request actions or change states, etc. Most of these lines are self-contained within the laser bench enclosure, and therefore unimportant in the context of routing. Table 5 summarizes the external lines, some of which coordinate the actions of the APOLLO apparatus, and some of which are used for temperature sensors. Table 6 lists the intended positions and functions for an array of miniature RTD temperature sensors. The RTD readings are made available to `houston` through the DI-720 data acquisition unit.

10 Control Software and Observatory Interface

As can be guessed from Table 4, the computer, `houston`, is the seat of control for APOLLO. The tasks that `houston` will take care of include:

Table 6: Positioning of the RTD temperature sensors

Sensor Location	Used for:
CAMAC/TDC air intake	Control of M-33 set-point, Noren Fans
TDC Internal #1	Calibration of TDC/Control of set-point
TDC Internal #2	Calibration of TDC/Control of set-point
TDC Internal #3	Calibration of TDC/Control of set-point
CAMAC/TDC exhaust	Calibration/Diagnostic/Control
Upper (CAMAC) Noren exhaust	Control of M-33 set-point
Noren intake air	Diagnostic/Control of set-point
Internal to laser box (air)	Diagnostic
Top surface of laser bench	Diagnostic/Control of set-point
Ambient dome air	Diagnostic
Upper Cabinet/Capacitors	Diagnostic
IL enclosure/laser elec. rack	Diagnostic

• Gathering data from CAMAC (TDC, ACM)	0.025
• Calculating arrival time of next photon & informing ACM	0.050
• Getting time from GPS Clock	1.00
• Gathering temperature information	1.00
• Sending requests to TCC for guiding and focus corrections	5–30
• Commanding T/R motor to spin & controlling speed	~ 60
• Controlling M-33 chiller for thermal control of laser bench	~ 60
• Gathering pressure information	~ 60
• Adjusting optics through picomotor control	~ 300
• Controlling laser state	~ 300
• Controlling CCD camera and occasionally collecting images	~ 300
• Occasionally re-programing ACM	—

These tasks are listed roughly in order of frequency, the approximate periods of activity in seconds listed to the right. For instance, the first happens 40 times per second, the next 20 times per second, the next two once per second, etc. The last one may happen only once or twice per year (if that) once stabilized.

Most of the low-level communications and interface is programmed in C, and a Python/Tkinter wrapper presents a graphical control interface granting access to most of the functionality listed above.

The exact scheme has not yet been worked out, but ideally, `houston` will not itself run the graphical user interface. Rather, a remote machine will run the user interface, sending command requests to `houston`, and listening to the data that `houston` is churning out.

In terms of direct observatory interface, `houston` will compute an appropriate `track` and `offset arc` command-set that will cause the telescope to track our target, and will send these to the TCC. This command sequence has been tested on the 3.5 m, demonstrating the ability to maintain one-arcsecond tracking over ten minutes (during which time the moon moved over 300 arcseconds). Once tracking, the user can make manual nudges via the graphical interface. The result of each such action will be a requested TCC action on the part of `houston`. `houston` may also evaluate return signatures and generate requests automatically, as part of a search/optimization strategy.

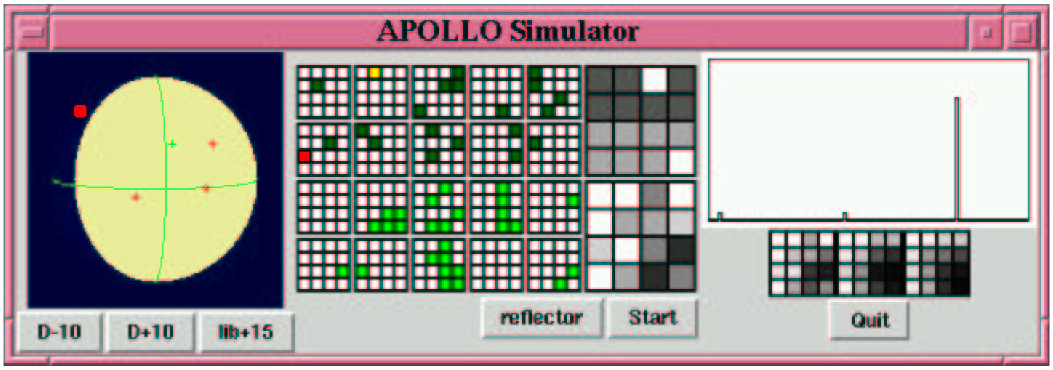

Figure 5: Early look at APOLLO status screen, produced from an APOLLO data/background simulator. Each 4 × 4 grid represents the APD array output. The top 10 frames represent the corner-cube calibration signal, and the lower 10 are lunar returns. At 20 shots per second, this represents 0.5 seconds-worth, so a frame like this would update twice per second. At lower right are the 1, 5, and 10 second sums of the lunar returns. From this you can clearly see that the return is not centered—a state that would trigger corrective action to the guiding/alignment. The time histogram at right shows the calibration return.

The operator will have access to the graphical interface, which should contain enough informative displays that assessing the current status and evaluating the need to initiate any corrective actions should be obvious (see Figure 5 for an example). In addition, the video output from the STV CCD camera will be routed to the control room for a real-time guiding display. The details of this routing have not been worked out, though it is very likely a solved problem waiting for a purchase order.

11 Auxiliary Interfaces

Unfortunately, millimetric lunar ranging isn't a simple point & click operation. If we simply pointed the telescope at the moon and acquired range signals, we would know very precisely the light travel-time from the fiducial point on the telescope (Section 14) to the particular reflector on the moon. But to convert this into a range measurement from the center of the earth to the center of the moon, one must know both the instantaneous crustal deformation and the atmospheric propagation delay. The former varies by about 0.35 m, and the latter is about 1.6 m at zenith, scaling with airmass.

11.1 Crustal Displacement

We seek to understand crustal displacement through a combination of gravimetry and GPS measurements. The superconducting gravimeter is an impressively sensitive device capable of sensing sub-millimeter displacements relative to the earth's center of mass simply because gravity gets weaker the farther away one travels from the center of earth! The gravimeter achieves this precision on very short timescales of roughly one second. But the system is not perfect. Atmospheric pressure and ground water—both of which can deflect the crust by a few millimeters—also exert direct gravitational influences on the gravimeter, thus masking the true deflection in a way that may or may not be easy to model. Additionally, gravimeters are completely insensitive to horizontal motions due either to loading phenomena or continental drift. For this, GPS can save the day, with new installations capable of 0.5 mm horizontal determination and 2.5 mm vertical determination over 24 hr averages. So the GPS is not capable of the short timescale tracking that the gravimeter enables, but has excellent long-term stability, allowing us to look for gross trends that we may not see or believe in the gravimeter data alone.

11.1.1 Gravimeter Installation

Funding for the gravimeter is not yet secure, but it would ultimately have its own data acquisition system hooked up to the Internet. It can be viewed as a standalone addition to APOLLO at some later point in

Table 7: Requisite meteorological precision for 1 mm delay determination

Parameter	zenith	60°	30°	20°
Pressure (mb)	0.43	0.38	0.23	0.14
Temperature (°C)	125°	70°	14°	4.3°
Rel. Humitidy	20%	19%	10%	8%

time. In truth, it doesn't have to be located on the APO grounds. Anywhere nearby that shares the same bedrock (the whole mountain?) would work. Of greatest relevance is a quiet site away from human traffic.

11.1.2 GPS Installation

The GPS *may* come for free. I have two separate contacts who may want to place a precision GPS installation at Apache Point if we provide the power/caretaking. They aren't that expensive in any case (compared to the gravimeter), so we can expect this to happen regardless of the funding situation. Like the gravimeter, this can be viewed as a standalone installation, divorced from the rest of the APOLLO system.

11.2 Pressure Sensors and Other Meteorology

The relevant factor in assessing the refractive delay through the atmosphere is the density of air along the line of site of the telescope. We have no means of measuring this directly, but if hydrostatic equilibrium of the atmosphere is believed, a pressure measurement at the site is directly proportional to the overhead (zenith) integrated density, since this is just a measurement of the mass of overlying air. If we further assume no horizontal pressure gradients, we can estimate the integrated density along oblique lines-of-sight as well.

Temperature and relative humidity also play a role, but as Table 7 indicates, the pressure is the quantity of greatest relevance. The entries of this table can be interpreted as the degree of precision to which the associated parameter needs to be measured to achieve 1 mm precision in the atmospheric delay determination (for various zenith angles). Temperature accuracy is clearly not important, and we can certainly rely on APO's current measurements of this. Likewise for the relative humidity, though caution is warranted, as the surface humidity may not be indicative of the tropospheric volume. An attractive possibility for assessing humidity profiles is through analysis of the infrared all-sky camera images already in use at APO.

But clearly the pressure is critical. We should have as a goal 0.1 mbar accuracy. Such transducers can be bought for about \$2,000 each. One such unit can be placed in the dome, and its data read by `houston`. But we may want to go one step further. Not willing to trust hydrostatic equilibrium and zero-gradient conditions, we would actually do much better to place a local *array* of transducers so that we could sense local gradients. The moon is not always straight overhead, so the pressure measurement at the dome doesn't tell us much about the integrated path if we're looking across a regional pressure gradient. Since the scale height of the atmosphere is about 7 km, and APO sits near 3 km, an array of roughly 5–10 km scale is sensible. Because this is a long way to run signals (by cable, fiber, or RF), we probably would want a junker PC at each location to record the data. Obviously much needs to be worked out here, but the need for such an array is some time off yet.

In addition to large-scale meteorological pressure gradients, local terrain could produce gradients, as wind flows over the ridge and produces dynamic Bernoulli effects. For this, one might want a smaller grid of transducers. Short of this, wind speed and direction could potentially be turned into a modeled correction factor. As a guide, the ram pressure associated with air at 20 mph comes out to about 0.4 mbar at APO. Bernoulli effects would be some fraction of this, perhaps about a quarter. So it's not wholly unreasonable to make relevant (and small) corrections to our pressure measurements by using wind information in this way.

12 GPS Antenna Feed

We have purchased an optical fiber link to send the GPS signal to our clock in the dome via a 325 ft single-mode optical fiber. The 3111A transmitter and 4111A receiver is made by Agere (formerly Ortel), and will allow us to place the antenna on top of the control room, running the fiber along the connecting corridor and into the dome. The only lightning threat then comes from the short run of coax cable from the antenna on the roof to the 3111A transmitter somewhere inside. We can install a lightning arrester at the entry point into the building. Though these are not 100% guaranteed to prevent problems, the added protection is a good precaution. The 3111A is powered by a wall-plug transformer, so it may be that the 3111A and transformer together take the heat of an unarrested—or partially arrested—strike, leaving the rest of the local environment unharmed.

13 Corner-Cube Implementation

A crucial part of our differential measurement of the range to the moon is the small corner-cube prism located in the exit aperture of the telescope. The corner-cube intercepts a small portion of the light emerging from the primary mirror, and sends this light back through the optical train in the direction to which the 3.5 m telescope is more accustomed. This light is attenuated heavily on its way to the APD receiver, ultimately knocked down to a couple of photon detections per pulse—similar to the lunar return strength. By comparing the arrival time of the corner-cube pulse to the lunar arrival time, one in essence has a measurement of the range from the telescope corner-cube to the corner-cube array on the moon. Ultimately, one wants to add to this value the distance from the corner-cube to the fiducial fixed-point on the telescope (intersection of axes), thus giving a measurement from a "fixed" point on the earth to a "fixed" point on the moon. Section 14 addresses the determination of the distance of the corner-cube from this fiducial point.

The requirements for the corner-cube are not very hard to achieve. The mounting simply has to be stable (relative to the PMC) to some fraction of a millimeter (0.005 in is more than good enough). It is desirable to place the corner cube as close as possible to the elevation axis, so that the distance to this reference point is minimally affected by thermal expansion of the telescope structure. The aperture of the corner-cube will most likely be about 1.5 inches. Alignment is not at all critical (within a few degrees).

A second function is performed by the corner-cube: if two (or more) corner-cubes are placed around the periphery of the mirror cell, a focal plane image of the returns from these corner-cubes will perfectly coincide only if the outgoing laser beam is collimated (assuming the focal plane is at infinite focus). So a pair of corner-cubes provides a handy beam collimation diagnostic. In truth, the diffractive spread from the corner cubes will be about 7 arcsec FWHM, so that it will be hard to verify collimation to better than about 1–2 arcsec in this way. But this technique can at least get us close enough to see light return from the moon, at which point we can refine the collimation (via picomotor actuation) so as to optimize return signal strength.

From the timing point of view, one would not want to have two permanent corner-cubes at approximately the same place with respect to the fiducial point on the telescope, because the returns from the two would overlap and smear out the apparent pulse shape, reducing both temporal precision and interfering with a useful pulse diagnostic. However, two corner-cubes staggered at different levels (separated by a few inches, and on opposite sides of the mirror cell) *would* be separately distinguishable, would provide a redundant sanity-check, and would produce more calibration returns per pulse than would be achieved by one alone.

An alternative to rigid mounts would be deployable/stowable corner-cubes so that the obstructions could be pulled out of the way of other observers (small though they are), and one would also have the capability of "blinking" the two corner cubes to better assess the degree of overlap of the two diffraction patterns by alternately comparing positions. The corner cubes would be signaled to move in/out by a TTL-level signal. The actuation mechanism is at present undecided, though a pneumatic actuation has been suggested.

Which of the two approaches we decide to take is something of a toss-up at the moment. With so many other jobs on the plate, the simpler permanent mount is attractive. One factor driving the decision is the rapidity at which the picomotors can drive changes in the collimation. The picomotors are rather slow, so waiting for the 7 arcsec diffraction patterns to separate appreciably could be very painful, so that in practice blinking becomes more desirable. But this could still be accomplished with a permanent mount and a flip-cover that would optionally obscure each of the corner-cubes.

14 Reference Axis Survey Plan

Setting aside the notion that there is no such thing as a fixed point on this floppy earth, we *do* need to reference our lunar measurements to *something* that approximates a stationary reference. Other ranging systems (e.g., satellite ranging systems that achieve millimeter resolution) use the intersection of telescope axes as such a point. In principle, the azimuth axis of the 3.5 m telescope intersects the elevation axis at a position near the surface of the tertiary mirror (on the optical axis). This point does experience slight motions as the telescope grows and shrinks thermally (referenced to base of azimuth cone near ground level), but remains fixed with respect to telescope pointing. The corner-cube we use as a fiducial does *not* remain fixed in space as the telescope moves. Therefore, it is necessary to carefully measure the distance from the vertex of the corner-cube to this center of telescope rotation.

Before going on, it should be noted that the two axes of the telescope may well *not* intersect to millimeter precision. In fact, we should assume that they don't. In order to accommodate this, we could simply include a horizontal offset between the two as a parameter in our model. If we know the azimuth of the telescope for a given range measurement, we allow the elevation axis to be displaced by some fit-for amount in the direction of that azimuth. Our survey measurement may in fact yield insight into the size (or limit thereof) of any such displacement.

One possible survey strategy is this: use a precision theodolite stopped down to a small aperture (better definition of nearby points) and get as far away as the dome enclosure will allow. Without moving the azimuth drive (and therefore dome), move the telescope in elevation while taking measurements of the vertex of the corner-cube (both, if you can see both of them). Along with these measurements is a measurement to understand the distance to the corner-cube from the theodolite. But also, it should be well known what the elevation angle of the physical mount is for each measurement. The sequence of theodolite measurements produces a three-dimensional arc in space. The physical mount angles sets the scale (confirmed by distance to corner-cube). The error of the measured points away from a theoretical arc indicates the measurement uncertainty. The radius of the arc is the sought-after measurement of the physical distance from the vertex of the corner-cube to the elevation axis. The measurement can be repeated from multiple positions on both sides of the telescope.

The best way I can think of to directly measure the intersection of axes is to set up on a stable position outside the dome, look through the open enclosure with a theodolite at a small, spherical target (e.g., ball bearing) mounted near the top end of the front-end. Ideally, one would use two such spheres at a known separation (a foot or so apart) mounted on the sky-side of the secondary support. Rotate in azimuth so that you can just see the two targets (sidelong) through a fair bit of the elevation rotation range, and survey an arc of elevation motion. Now flip to the other side in azimuth, so you are looking from the other side of the telescope and perform similar measurements. Each sequence (again, with knowledge of physical mount angles) gives a 3-d arc in space (the two targets help define the distance-to-target). Each arc has a 3-d center point. The degree to which the centers coincide is the degree to which the axes intersect.

How accurate can we expect these measurements to be? A typical quality theodolite is good to a couple of arcseconds. At distances like 30 m, this translates to 0.3 mm. Provided the targets themselves can be unambiguously centered in the theodolite to comparable precision, the measurement should be robust. If our target is a 3 mm ball bearing, such centering (1/10 the diameter) should be straightforward. It may be trickier to unambiguously center on the same point on the corner cube as the viewing angle changes slightly, but this is probably easily worked out.

15 Scheduling Plan

15.1 Requisite Observing Length

We envision typical APOLLO observing runs lasting 30 minutes. This number is dictated by the total number of photons received per target. The ultimate precision is given by $\sigma_{\mathrm{tot}} = \sigma_{\mathrm{indiv}}/\sqrt{N}$, where σ_{indiv} is the RMS error of the range as measured by a *single* photon, and N is the number of photons collected. For the Apollo 11 and Apollo 14 arrays, the maximum single-photon range error (due to lunar libration) is 25 mm. For Apollo 15, it's 46 mm. Thus we need about 600 photons from the smaller Apollo 11 and Apollo 14 arrays, and 2500 photons from Apollo 15—which is 3 times larger in area than the other two. At an average return rate of one photon per pulse (for Apollo 15), and a pulse rate of 20 Hz, we achieve our goal in 125 seconds, or about two minutes. For the smaller arrays, assuming one photon every three pulses, we get to one millimeter in 90 seconds.

Based on the assumption that our efficiency will not be perfect, and that more is better, and that once you've spent two minutes, you might as well spend five, a nominal plan of five minutes per target seems reasonable. There are four lunar retroreflectors that will be part of our usual suite: the three Apollo arrays, and one Soviet-landed array. The latter is of poor thermal design, so we may only choose to range to this one during its lunar night (waning phases). In any case, 3–4 targets at 5 minutes each, plus 5 minutes for initial acquisition and 5 minutes for shutdown looks like 30 minute runs. We may ultimately prove more efficient, in which case we will have to decide whether we want to simply use up 30 minute slots to maximal advantage, or shorten the slots to 15 or 20 minutes.

15.2 Fitting into the Schedule

Even though APOLLO time will be scheduled by the normal time allocation process, if we were to grab a thirty-minute slot at 9:30PM in the middle of what would otherwise be a contiguous observing session, we will likely be viewed as an unwelcome interruption. In poor conditions, this may not chafe, but if the seeing is very good, the intrusion could well be greeted with some bitterness. APOLLO aims to be a campaign of many years' duration, so we would like to have the smoothest possible interface with the observatory and the observing community.

APOLLO's science goals are accomplished by measuring the *shape* of the lunar orbit throughout the month. One of our key goals—measurement of the equivalence principle—is most sensitive in comparison of new-moon range to full-moon range. But more generally, a uniform distribution of measurements throughout the month makes for the cleanest extraction of science signals. Roughly speaking, about 10 measurements per month would establish a nice sequence, corresponding to a range measurement every 2–3 days.

From a background point of view, APOLLO could happily operate during the daytime. In an ideal world, APOLLO would snag a lunar measurement when the moon crosses the meridian, thus cutting through a minimal airmass. Continuing this idealization, APOLLO would also gather measurements a few hours before and a few hours after meridian crossing in order to tie down the erratic earth rotation (the meridian measurements alone say nothing about this). For now sticking with the meridian observations, we have the problem that the moon swings right past the sun once per month. Ranging to a slim crescent on the meridian means that the sun is high in the sky, not far away, and beaming its light onto the telescope (even the primary mirror!). Clearly this doesn't work, and we must range to crescent moons only after the sun is below the horizon. New moon ranges are out of the question, and we are limited to about 2–3 days on either side of new moon. This means no dark-time activity except during twilight.

On the full-side of quarter moon, the moon is more than 90° away from the sun, and daytime operations become feasible. When close to quarter moon, we will have to be careful to understand operating limits with respect to dome/sun azimuth, etc. so that we may avoid any unwanted heating of the telescope structure or enclosure.

In keeping with the current half-night-oriented operating scheme at APO, it seems natural that we should try to fit APOLLO's time blocks into the mid-night switch slot, (bright) twilight slots, and daytime. Twilight slots are *necessary* to acquire crescent-moon ranges. Midnight works well for full moon and gibbous phases. If we use the lunar phase angle from the sun, D, where $D = 0^\circ$ corresponds to new-moon, and $D = 180^\circ$ to full, we can arrange nominal observations as follows:

- $30^\circ < D < 90^\circ$: Evening twilight slot
- $90^\circ < D < 135^\circ$: Evening twilight slot with occasional afternoon supplements
- $135^\circ < D < 150^\circ$: Midnight slot with occasional afternoon supplements
- $150^\circ < D < 210^\circ$: Midnight slot; twilight supplement if necessary
- $210^\circ < D < 225^\circ$: Midnight slot with occasional morning supplements
- $225^\circ < D < 270^\circ$: Morning twilight slot with occasional morning supplements

- $270° < D < 330°$: Morning twilight slot

"Supplements" are occasionally sought to pin down earth rotation, or may sometimes be used when a scheduling difficulty precludes operation during the preferred time slot. The break-points in the list above are chosen such that of the two slots available—twilight and midnight—the moon is higher in the chosen slot.

Midnight slots are far less disruptive to observers than twilight slots, since twilight is often used for calibration frames (flat fields, etc.). There is no immediate and simple solution to this potential conflict. Some ideas include:

- Observers state preference for access to early twilight or late twilight, and we take the other.
- A central APO-observers stockpile archives all calibration frames. From this, the expected variability can be studied, and it may be that observers are happy using flats & biases from previous nights and previous observers. If the archive becomes cumbersome, only the last several months need be kept.
- Use the daytime observation to hit all four targets, and just come in for 10 minutes during twilight to pick up the more optimal (higher in sky) datapoint to a single reflector.

16 Operations Plan

As identified previously, there are three operational modes for APOLLO: warm-up, ranging, and cool-down. Each will nominally last about 30 minutes.

At the start of operations, the telescope assistant (or other staff, e.g., during daytime) goes up to the intermediate level to power on the chillers, and laser, and open the doors. A good design will permit this to be done from a standing position on the IL floor. Next the electronics within the laser bench enclosure are turned on from within the dome. Probably at the same time, a few buttons on the laser remote control box (within the cabinet behind the telescope) will be pressed to begin the laser warm-up. If not done manually from the remote-control box, this step can be initiated through software a few minutes later.

As an alternative to the above, switches could be installed in the control room for activating the various equipment.

The APOLLO software is started, and set into warm-up mode, during which it monitors the equipment state, runs calibrations, and awaits the command to start. Once the observations are to begin, the observer (be it a UW scientist or an APO staff member) selects a target, and the APOLLO software issues the telescope track commands to follow the designated target. The video feed is used to verify pointing. In all likelihood, the system is well-enough optimized that initial acquisition is very straightforward. The APOLLO software is envisioned as being capable of executing a simple search strategy. Once a signal is received, a quick optimization sequence maximizes the signal return rate. In the event that the apparatus is out-of-whack, the observer can check various diagnostics such as beam collimation, focus, receiver alignment, etc. In truth, we won't know until we try how robust and automated the APOLLO sequence can be.

Assuming smooth operation, the observer will monitor the signal rates, pointing, air traffic (means as yet undetermined), etc. At the end of the session, a summary report will be e-mailed to collaborators, and any anomalous problems will be registered by the observer.

The mode of observation is far from fleshed out at present. At first, operations will be on-site with APOLLO's builders present. The transition away from this is less certain. It is not clear whether we want to push in the direction of remote observing (how do we handle the video feed?) or leave the matter in the capable hands at APO. Whichever the case, we would like to have a very strong and positive relationship with the APO staff, considering them a part of the APOLLO team—and sharing in its success.

17 Thermal Impact to Other Observations

The thermal impact to observers has in many ways been addressed already (exhaustively?) in Section 3. Here, we summarize the likely impacts due to our "hot" boxes in the dome, and also address thermal impacts to optics. Because of our active cooling scheme, we expect no significant difference to the observer during the off-state and during warm-up.

Before launching into individual details of the thermal emissions from each component, it should be noted that such emissions always originate from behind the telescope, and not directly in the line-of-sight of the observer. It is also true that an escape path for warm bubbles of air generated by our equipment exists straight up over the equipment, through the open enclosure shutter. Nonetheless, our design should always keep thermal leakages into the dome below the 50 W level.

17.1 Laser Bench Enclosure

Given 3-inch Thermax sheathing surrounding our laser bench, with its R-value of 24 and effective area of 7.2 m^2, we anticipate a heat load into the observatory of $P = 1.70\Delta T$ Watts, where ΔT is in °C. This exceeds 50 W for a temperature differential of 29°C, or 53°F. At a set-point of 10°C (50°F), this corresponds to a dome temperature of approximately -20°C (-3°F)—colder than any yet-recorded enclosure temperature. When our set-point is raised to $\sim$ 24°C, the 50 W limit is reached at an enclosure temperature of -5°C (23°F). These temperatures are seldom, if ever, reached during the “warm” season, when we would be using the higher set-point.

Mechanical conduction into the PMC is expected to be less than $P = 0.1\Delta T$ Watts, for a maximum of 3 W. Radiative loading of the PMC from the underbelly of the laser bench enclosure is estimated to be less than 3 W given a metallic bottom surface for low emissivity.

17.2 Cabinet on Observing Level

Our double-walled, constantly vented cabinet design is intended to carry virtually all of the heat leaking out of the “chimney” of the cabinet immediately down into the intermediate level. If we use 1-inch insulation for this upper section ($R = 7.2$), then the $3.39\Delta T$ Watts of leakage, if removed at 95% efficiency, would leave $P = 0.17\Delta T$ Watts finding its way into the dome.

17.3 Intermediate Level Enclosure

Though the IL enclosure leaks $P = 8\Delta T$ Watts (at 2-inch insulation: $R = 14.4$), this heat is of little concern to the observer, as it is easily removed by the ventilation systems in place. When in warm-up (and cool-down), we generate nearly 5 kW of heat in the IL. This is expected to raise the temperature slightly (2.5°F), but only for short times, and with a fast recovery timescale. It is not anticipated that this will have any noticeable impact on observers.

17.4 Hoses and Cables

Hose issues were addressed in Section 3.5, and is an unresolved issue worth attention. If we are not careful, we could easily match or exceed the laser bench enclosure thermal output with the heat given off by the hoses. For now, if we assume that the hoses in the so-called hot loop are housed in vented tubes, then we may guess that the hoses contribute something like 20 W at the coldest temperatures, or $P \approx 0.8\Delta T$ Watts.

So far untreated are the cable contributions. We have cables diving into and out of warm boxes. Some heat will be conducted out through these paths. At first blush, the relatively small cross-sectional area of the aggregate bunch of cables (maybe 0.005 m^2) pales in comparison to the many square meters of box/cabinet surface areas. But the thermal conductivity is very much higher. A quick-and-dirty estimate might say that we have 1 square-inch of copper conductor. If there is a $\Delta T = -25$°C, and the temperature of the cables makes the full transition over about a foot, then the heat flux (from each box) would be about 15 W, for a total of 30 W. We could surely reduce this by using some insulation, but probably not easily by more than a factor of two. Splitting the difference, we might say that $P \approx 0.9\Delta T$ Watts.

17.5 Adding It Up

If we add the heat into the dome as contributed by the laser bench enclosure, the cabinet, the hoses, and the cables, we get a total of: $P = (1.70 + 0.1 + 0.17 + 0.8 + 0.9)\Delta T = 3.67\Delta T$. This sum figure exceeds 50 W at a $\Delta T > 14$°C. While this is not as thermally clean as we had hoped to achieve, records indicate

Table 8: Tertiary Mirror Properties

Property	symbol	value
mass density	ρ	2200 kg m^{-3}
heat capacity	c_p	745 J kg^{-1}K^{-1}
thermal conductivity	κ	1.38 W m^{-1}K^{-1}
thermal expansion coefficient	α	2.8×10^{-6} °C^{-1}
front plate thickness	t_1	0.0107 m
rear plate thickness	t_2	0.0127 m
rib length	l	0.1031 m
rib thickness	t	0.005 m

that typically only about 20 nights per year get this cold. And on most of these nights, the leakage goes up to a number like 60 W instead of 50 W.

The surprise contributions are the hoses and the cables. But it is worth bearing in mind that these two are the least well studied in our system, and A) may not be as bad as calculated and B) may be easily improved by a bit of engineering.

17.6 Optical Heating

Another issue concerns the optical heating of the telescope by the laser power itself. At an average power of 2.3 W, the laser delivers non-negligible heat loads directly to the parts of the telescope that matter the most—the optics.

The tertiary is the hardest-hit, having a surface illumination intensity four times higher than that of the secondary, and 92 times that of the primary. Thus we will only consider the effect on the tertiary mirror here.

The laser beam is expanded to fill the entire 3.5 m aperture of the telescope, meaning that it emerges from the laser bench enclosure in an $f/10$ beam. The tertiary is 3.07 m from focus, so that the footprint on the tertiary is an ellipse 0.307×0.434 m^2. The area is about 0.10 m^2, so that the incident power intensity is 23 W m^{-2}. Assuming a terribly inefficient tertiary that absorbs 20% of the incident light (more likely *absorbs* only 8–9%, and dust may *scatter* additional light), we can get a handle on the worst-case heat loads. The resultant heat flux into the mirror is $q = 0.46$ W, or $q\prime\prime = 4.6$ W m^{-2}.

Using the values for borosilicate glass and the geometry of the tertiary in Table 8, we can calculate the warming of the front plate after some time Δt to be $\Delta T = q\Delta t/\rho c_p V$, where V is the volume heated. If we assume no heat diffuses into the ribs or out into the plate, the volume is small (0.00107 m^3), and $\Delta T = 0.47$ °C after 1800 s (half-hour) of continuous exposure. If we let the heat diffuse into the bulk of the mirror (total $V = 0.018$ m^3), then $\Delta T = 0.028$ °C.

These simple calculations at least set the scale of the problem. In truth, heat is taken away from the surface of the heated glass. A finite-element analysis of the front plate (no ribs) indicates that the temperature rise is more like 0.37°C after 1800 s, 0.60°C after 3600 s, equilibrating to 0.94°C eventually.

Let's build a few worst-case scenarios for deformation of optics. One mode of deformation would have the ribs immediately take heat away from the part of the plate immediately above them, leaving heat in the middle of each cell to be pulled away by convection/radiation. In this worst case, the edges of the cells would have a ΔT of 0°C, while the middles of the cells would be sitting at roughly 1°C above ambient (after prolonged exposure). Given the thermal expansion coefficient, and assuming equal deformation on the top and bottom of the thin plate, one would get $\delta t = \Delta T \alpha t_1 = 3.0 \times 10^{-8}$ m, or 30 nm, or about $\lambda/20$ for optical wavelengths, λ. One should not worry about such small deformations (and keep in mind that this is after a few hours of exposure, using a whopping 20% absorption).

The next worst-case is to assume the ribs take the heat. Though no finite-element analysis has been performed for the ribs, it is assumed that the total ΔT would be smaller for these, given their large surface-area-to-volume ratio. One might assume that given the ribs as a conduction path, the equilibrium temperature of the front plate is about half its no-rib value, or about 0.5°C, so that the average temperature of the rib is half of this, being wedged between the heated top plate and the ambient back plate. If so, the rib grows by $\delta l = \Delta T \alpha l = 7.2 \times 10^{-8}$ m, or 72 nm, or about $\lambda/8$. This is starting to be a distortion.

To some extent, the back flexes too. Being 20% thicker, and bending being proportional to the cube of thickness, one would expect the top to bend about 45 of the 72 nm, for a distortion of about $\lambda/12$. So what does this distortion look like? The central ribs are hotter than the outer ribs, partly because only the center of the tertiary is illuminated, but also because the Gaussian beam intensity puts more power in the center. Although the illuminated pattern is elongated into an ellipse, from the vantage of the observer, it is a circularly symmetric distortion—just as the illuminating beam is circular as viewed along the optical axis. So this is in essence a spherical aberration mode, looking much like focus. The worst case focus shift is about 40 μm, which doesn't sound at all troubling. It is hard to imagine ever discerning any image degradation from this small amount of focus aberration, especially given the atmosphere's several wavelengths of low-order distortions (focus, astigmatism, etc.).

Once we terminate lasing activity, the recovery timescale of the mirror ought to be relatively fast—especially considering that there may be no recovery necessary to begin with. Again, these calculations assumed a 20% absorption factor, so that when more realistic values of $< 10\%$ are considered, the effects that were already marginal all but disappear.

18 Safety Plan

As exciting as the APOLLO science is, and as eager as we are to launch our program, we must not allow harm to come to anyone, be they APOLLO scientists, APO personnel, or an unknown pilot. So we must carefully integrate safety into our scheme.

18.1 Laser Hazard

The laser is our chief concern. A direct intra-beam exposure to the unexpanded beam would surely result in blindness. Skin exposure will cause a light burn, leaving a temporary small red mark behind. Once the beam is expanded to fill the primary mirror, however, there is *no* danger to the skin, and actually little danger to the eye. The latter is a bit surprising, so following is a bit more information on this point.

The 3.5 m outgoing beam has an intensity of ~ 0.2 W m^{-2}. This is small compared to the sun's 120 W m^{-2} (in the eye's photopic band), though the laser is concentrated onto an individual resolution element, unlike the extended sun. Coincidentally, they work out to about the same retinal intensity—the laser as viewed from in front of the telescope would have the same apparent surface brightness as the sun, but is much smaller. You still don't want to look at it—it would certainly form afterimages. Even worse is that a dark adapted eye is much wider than a sunny-day eye (e.g., 7 mm vs. 2 mm), so the actual retinal intensity is likely to be 10 times higher for our laser. It is for this reason that we exceed the ANSI standard for intra-beam eye safety (by about a factor of 20).

But the more dangerous part of the beam is the part where the beam is small, i.e., within the APOLLO instrument. The small section of collimated beam within the apparatus is clearly dangerous. But what about the $f/10$ expanding beam emerging from the laser bench enclosure? Because of its divergence, doesn't this appear to be extended? If the viewer's eyes are focused at infinity, then yes. But the portion of the beam sampled by the pupil is only mildly divergent. It works out that the retinal intensity is independent of distance along the expanding beam (up to about 25 m away): the farther one is, the less light is accumulated, but the less divergent the sampled bundle. The intensity turns out to be 20 times solar, independent of pupil size. So this is bad. What's worse is that the eye could focus this to a tighter illumination, so that the 20× solar value is a *minimum* intensity. So we cannot permit intra-beam exposure—period.

In its stable operation regime, the laser beam is never exposed until the run from the back-port hole to the tertiary. So the only risk of exposure would be to someone in the light-gathering beam of the telescope. The one possible exception is the potential over-filling of the primary mirror, such that a look at the secondary mirror from the edge of the PMC would put light in the eye. This may not even be possible given the sizing of the secondary. But in any case it would not be difficult to mask out.

The real danger comes in regard to the installation/testing phase, in which case we will likely have the beam traversing the dome at times. Safety glasses are available for Nd:YAG lasers, protecting simultaneously against the green and infrared wavelengths. Our laser *does* produce infrared light, though it is terminated within the laser box. Still, for the purposes of working with the laser, we will get the (more expensive) glasses that offer protection at both wavelengths.

We can implement a number of interlocks and safeties, but no system is foolproof. In fact, since we *do* need to sometimes be in the dome when the laser runs, we can't implement a system that absolutely precludes presence in the dome with the laser running. So we will have to settle for something moderately restrictive, so that only authorized individuals can gain access to the dome when the laser runs. A suggested scheme is a door lock that activates when the laser is on, and is defeated only by a key (or from inside the dome) that authorized persons have access to. Perhaps this key could be stowed in a box containing safety glasses, so that even authorized individuals are reminded to don their glasses before entering.

18.2 Laser Shut-off

There will be at least two shutters in the laser. The first is inserted into the cavity, suppressing any lasing activity. This shutter can be software-controlled, push-button controlled on the laser remote control box, and is also tripped by various interlocks built into the laser system. A second shutter on the exit aperture of the laser box blocks the green light that would otherwise emerge from the laser box. The laser box is designed to be light-tight, and we will maintain this quality throughout any modifications we make.

During warm-up, we want the laser to be producing green light, so the internal shutter will be open and the exit shutter closed. We imagine the exit shutter spring-loaded into the closed position, defeated only by a positive voltage deliberately applied to the retraction mechanism. We can combine multiple logic signals in an AND sense to require that all enablers are simultaneously happy to let the laser light out.

There will certainly be a button in the control room that can shutter the laser off immediately. It is also worth thinking about incorporating a shutoff when any emergency stop button is pressed in the observatory.

As for external safety, we can help ensure no danger to site personnel/visitors by not allowing the laser to emit light below some minimum elevation—maybe 10°. Since the laser rotates in elevation, this is relatively easy to incorporate in a robust way. An inclinometer on the laser bench can generate one of the signals sent to the AND input of the exit shutter actuator.

18.3 Safety Organization

We ought to interface with safety offices at UW or NMSU in the near future. We have interacted with the UW safety office in the past with regard to this project, and now need to re-establish connection.

We will likely want to designate a laser safety officer at APO to maintain safe operating conditions for our laser. During our March visit, we should spend some time to draft a number of ideas for implementing various safety schemes.