APOLLO Systems Manual

August 29, 2005

Contents

Preface

This manual is intended to serve a broad audience. Whether a student new to the project, a telescope observing specialist, or a long-time member of the APOLLO collaboration, this document is intended to provide the necessary background to understand the construction and functioning of the APOLLO apparatus.

The focus of this manual is the hardware. The nitty-gritty details are mostly kept out, with the intention that this manual is actually readable. Non-essential technical information is relegated to other documentation. A software manual is separately under construction.

Chapter 1

System Overview

1.1 The Lunar Ranging Technique

APOLLO (Apache Point Observatory Lunar Laser-ranging Operation) is a system for acquiring the range to the retroreflectors on the lunar surface. Lunar Laser Ranging (LLR) involves launching a very brief laser pulse at one of several retroreflector arrays placed on the moon by the Apollo astronauts (or by an unmanned Soviet rover, in one case). The round-trip travel-time of the pulse is measured to high precision and converted to a one-way range measurement. This time ranges between 2.34–2.71 seconds, depending on how far the moon is.

The signal loss is staggering, so that we lose all but about 10^{-8} of the launched photons on the up-link, and another comparable factor on the down-link. Almost all of this loss is due to divergence of the beam: owing to atmospheric turbulence on the way up and corner-cube diffraction on the way back. The net effect is a signal only a few photons-strong upon receipt.

This tiny signal (19th magnitude for APOLLO) is picked out against the bright moon (magnitude -13 for full moon) by a variety of filtering techniques:

- we only accept green photons at the laser wavelength (factor of 200 background suppression)
- we only look at the tiny region around the reflector (a few arcseconds for 10^6 suppression)
- we only open our detector for about 100 ns around the return time ($> 10^6$ suppression)

In the end, the laser photons returning from the moon outshine the bright lunar background. This means that APOLLO can operate in high background conditions, including full moon and daylight.

How many photons we expect to receive per pulse depends on: the number of photons launched (pulse energy), the one-way system optical efficiency, the narrow-band filter throughput, the detector quantum efficiency, the size of the reflector array (assuming 3.8 cm diameter corner cubes), the divergence imposed by the atmosphere, the diffractive spread at the corner cube, and the distance to the reflector array. All of these factors are represented in the equation below, in the order introduced above.

$$N_{\mathrm{received}} = 5.4\left(\frac{E}{115\,\mathrm{mJ}}\right)\left(\frac{\eta}{0.4}\right)^2\left(\frac{f}{0.25}\right)\left(\frac{Q}{0.3}\right)\left(\frac{n_{\mathrm{refl}}}{100}\right)\left(\frac{1\,\mathrm{arcsec}}{\phi}\right)^2\left(\frac{10\,\mathrm{arcsec}}{\Phi}\right)^2\left(\frac{385000\,\mathrm{km}}{r}\right)^4.$$

If the various factors take on the values indicated in the parentheses, we should expect to detect about 5 photons per pulse. If ϕ, the atmospheric seeing, becomes 2 arcseconds, then this term causes the photon rate to drop to 1.35 photons per pulse, since it is squared and in the denominator. The throughput and efficiency numbers are a bit pessimistic (especially the 40% optical efficiency!), so that we may even approach some approximation of reality. By contrast, the full moon background contributes less than one photon per detection window, and these are spread across the 100 ns window, so that they do not confuse the narrowly confined lunar return.

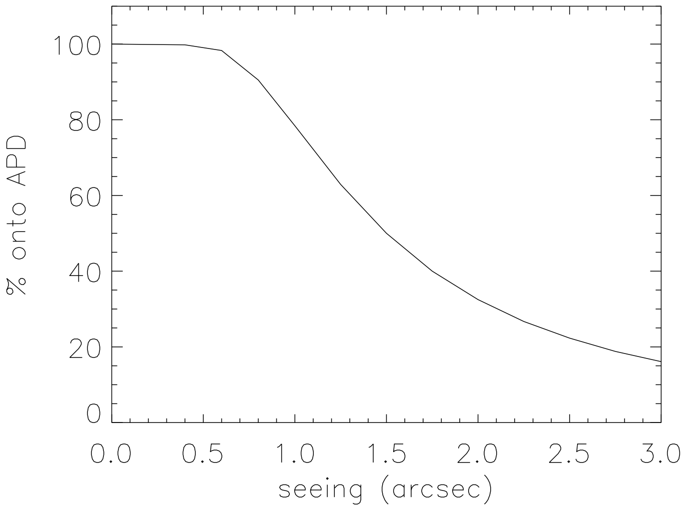

Figure 1.1: APD coverage as a function of atmospheric seeing

The above relation leaves out one critical piece. The image size of the return depends on the seeing. After all, the lunar array is essentially a point source on the lunar surface, much as a star is. So the return photons are blurred into a Gaussian-like distribution at the focal plane just as a star image would be. Our detector is small—about 1.3 arcsec on a side—so under bad seeing conditions, our detector looses light around the edges. Figure 1.1 shows the fractional coverage as a function of seeing (full width at half-maximum of the best-focus image).

The time-of-flight measurement is performed in a way that is as close to a *differential* measurement as possible. A corner-cube prism attached to the telescope's secondary mirror intercepts some small fraction of the laser pulse on its way to the moon. This light is returned and attenuated to the single-photon level so that the APOLLO detector array detects and time-tags these *fiducial* photons in exactly the same way as it processes returning lunar photons. By subtracting the return time of the fiducial photons from the return time of the lunar photons, a differential measurement is performed, resulting in a range measurement between the telescope's corner cube and the corner cube array on the moon. A fast-photodiode improves the fiducial measurement accuracy, described in detail in Section 7.4).

1.2 APOLLO Apparatus Outline

The system, in outline, consists of:

- a short-pulse, high-power laser,
- a receiver employing avalanche photodiode (APD) arrays,
- a rotating beam switch for transmit/receive modes,

- a high precision (and relatively accurate) GPS-slaved clock,
- a timing system composed of 25 picosecond resolution (100 ns range) time-to-digital-converters (TDCs) and a timing/counting module,
- a CCD acquisition and guiding camera,
- a thermal control and health-monitoring system
- and a computer to control the operation.

This chapter briefly describes each of these subsystems and their interconnection, with subsequent chapters describing individual subsystems in greater detail.

1.3 Physical Layout

The APOLLO equipment is distributed in three main volumes; the Utah box, the cabinet, and the ILE (discussed in Section 1.14). Figure 1.2 shows the general arrangement of the components. The contents of each of the volumes are as follows.

1.3.1 The Utah Box

- Laser bench
- Receiver tube and all optics besides the quaternary mirror (just outside box)
- APD array/electronics
- T/R mirror and motor
- CAMAC crate (ACM, Booster, TDC)
- STV CCD camera head
- XL-DC clock (in Utah at the moment)
- Power meter head
- Fast photodiode and Ortec 9327 timing discriminator
- 100 W heater
- Two Noren air-to-water heat exchangers (connected to M33 chiller)

1.3.2 The Cabinet

- `houston`, the control computer
- DC power supplies (6)
- RTD signal conditioner
- Power distribution logic hub

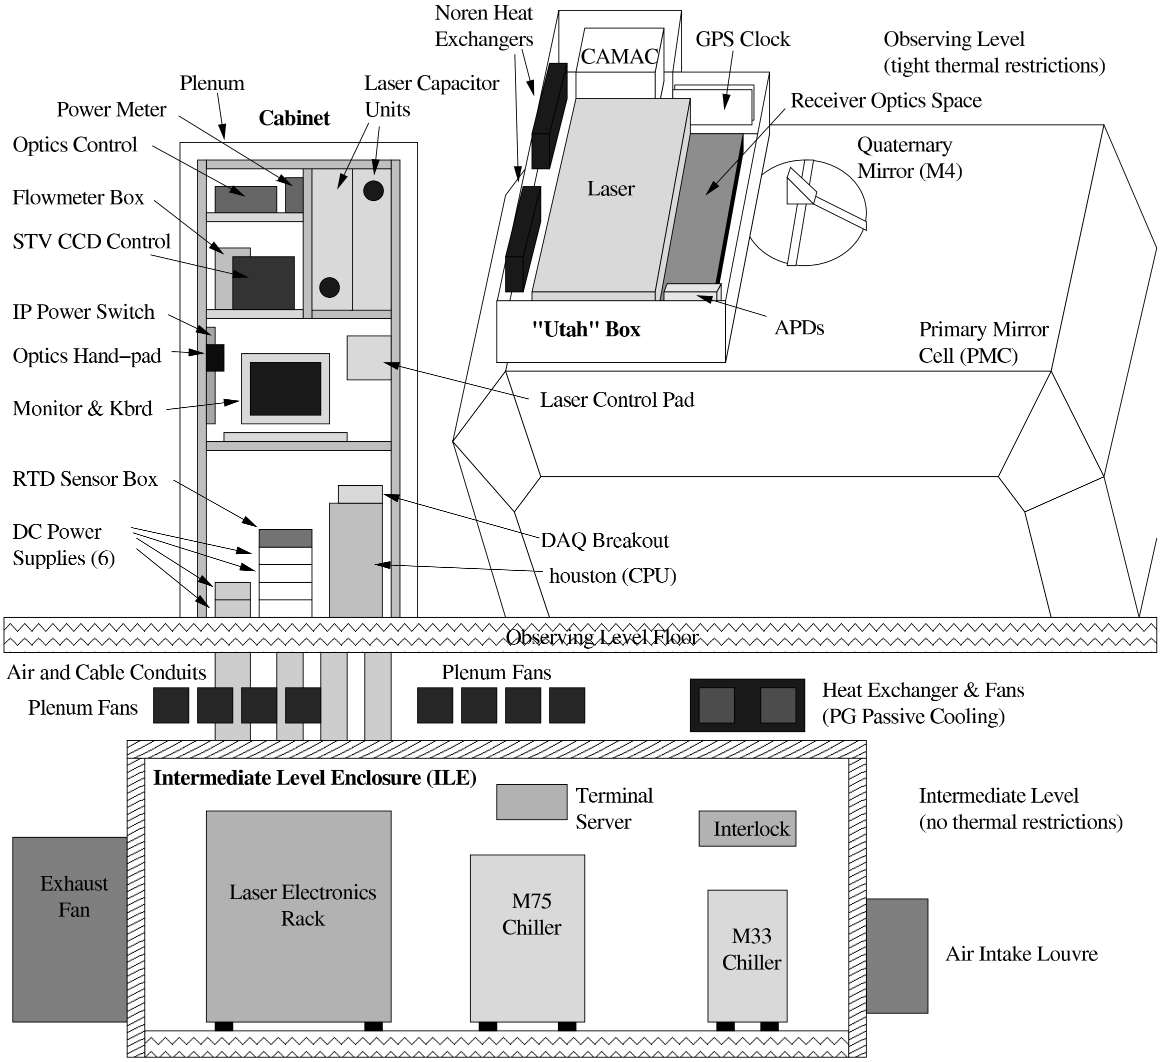

Figure 1.2: Layout of the APOLLO enclosures (**bold**) and main components

- Laser control box
- Laser capacitor units
- Optics actuation controller
- STV CCD controller
- Power meter analyzer
- Flowmeter signal conditioner
- IP power switch

1.3.3 Intermediate Level Enclosure (ILE)

- Laser electronics rack

- M75 chiller (for laser heat extraction)
- M33 chiller (for extracting heat from Utah)
- Terminal server
- IP switch/hub
- Interlock controller
- Cabinet/ILE circulation fan
- 400 W heater

1.3.4 Outside the ILE

- ILE exhaust fan
- Plenum fan bank (8 fans)
- Auxiliary pump and heat exchanger for passive Utah cooling

1.4 The Laser

The laser used for APOLLO is a *Continuum Leopard* Nd:YAG pulsed laser, frequency-doubled to 532 nm wavelength (green), with a pulse width of ≈ 100 ps, a pulse energy of about 115 mJ (3.08×10^{17} photons per pulse), and a repetition rate of 20 Hz. The average power is then 2.3 W, though the peak power is in the neighborhood of 10^9 W (and a corresponding duty cycle in the range of 2×10^{-9}). The output beam is about 9 mm in diameter, with a roughly Gaussian intensity profile.

The pulse full-width at half-maximum (FWHM) of 100 ps translates to a root-mean-square (RMS) width closer to 50 ps, which corresponds to about 7 mm in round-trip range to the moon. Thus an individual photon making the round trip has this unavoidable uncertainty due to the laser pulse width. Many other components in the system add to this uncertainty (in an independent, and thus quadrature-sum way), so that the resultant uncertainty in an individual photon's round-trip travel time is closer to 200 ps, corresponding to about 30 mm of range uncertainty. In order to achieve the APOLLO goal of millimeter range precision, we need approximately $30^2 = 900$ return photons in order to statistically reduce the ensemble uncertainty to 1 mm. This puts a premium on the number of photons detected, which sets the demand for the most powerful laser we can reasonably afford that has a short pulse width. Nd:YAG is very mature technology, with product-line support for this kind of application.

Associated with the laser is an electronics rack, detached capacitor units, and a separate closed-cycle water chiller. The placement of the laser electronics and chiller is such that we can confine the heat load to the observatory's intermediate level, so that we aren't dumping heat into the telescope enclosure. We are asked by the observatory to restrict our total thermal emissions into the dome to less than 50 Watts. Water carries most of the heat away from the laser bench itself, so that the heat load we need to handle in the dome is minimized. Our heat load into the intermediate level is effectively unrestricted.

1.5 The APD Array Detector

The detector for APOLLO must be capable of single-photon detection with a high degree of time-resolution. Photomultiplier tubes (PMTs), microchannel plates (MCPs) and avalanche photodiodes (APDs) fit the bill, though PMTs by their larger geometry lead to more dispersion (jitter) in the reported time of photon arrival. APDs have the highest quantum efficiency of the lot, and if thinned, require very little bias voltage (25–30 V) and have timing resolution < 100 ps. APOLLO uses APD array devices fabricated Lincoln Labs, which allow us to multiplex our measurement (multiple photons per pulse), as well as preserve spatial information to aid in acquiring/tracking the signal. The current arrays in hand are 4×4 arrangements with 100 μm element spacing and either 20 μm or 30 μm diameter active areas. The 30 μm device is the one used in the APOLLO system. The less-than unity fill-factor means we need to place a lenslet array in front of the detector to recover full areal coverage.

1.6 The Optical System

The APOLLO optical system provides a transmitter beam path, a receiver beam path, and a rotating optic to switch between these two. There are also secondary optical paths, such as that to the CCD camera, beam dumps, and to the laser-monitoring photodiodes. In brief, the transmitter beam path consists of a two-lens beam expander, a lens to diverge the beam to match the $f/10$ beam of the telescope, and a variety of flat turning mirrors. The receiver shares the $f/10$ transmitter lens (now serving to collimate the incoming light), has a variety of turning mirrors, a narrow-band filter to eliminate all but 532 nm light, a two-lens spatial filter for eliminating light from all but the 3 arcsec region around the target, a rotating set of attenuators and diffuser (for the fiducial return), a lens to produce a focus on the APD's lenslet array, and the lenslet array itself. The CCD camera gets 8% of the light entering the receiver tube (just before the narrow-band filter), and produces a focus with a 100 mm focal length, $f/2.3$ camera lens. More detail on the optical layout can be found in Chapter X.

1.7 The Transmit/Receive Beam Switch

If we want to use the full aperture of the telescope to transmit the laser beam, while obviously also using the full aperture to scrape up return photons, it is inevitable that we have some optical element periodically switching between these two modes. For this, we use a rotating optic that is mostly transmissive, but has a patch of dielectric coating with high reflective efficiency at 532 nm (99.9% for S-polarization). By spinning this optic at 20 Hz and slaving the laser-fire to the optic's rotational phase, we can guarantee that the patch is in the correct place when the laser fires, so that the emitted blast is sent to the telescope. At all other times, the receiver "looks" through the clear optic awaiting return photons from the sky. We can adjust the rotational frequency slightly to ensure that when the lunar photons return, the reflective patch will be well out of the way—say 180° out of phase.

The rotating optic is driven by a stepper/servo motor with on-board smarts allowing versatile control over the drive frequency as well as start-up and slow-down acceleration curves. An encoder on-board the motor sends an index pulse and 2000 encoder pulses per revolution to a delay generator, which counts encoder pulses and issues a fire request to the laser at the appropriate time.

1.8 The System Clock

The clock forms the backbone time reference against which we measure the travel time of our photons. As such, it must provide a low-jitter (< 10 ps) measurement over periods of 2.5 s, and additionally keep a consistent idea of what 2.5 seconds *means* over month timescales. The reason for this last statement is that if the Equivalence Principle is violated, the lunar distance at new moon is different than at full moon—and this is 15 days different in time. We'd better not have our yardstick change length in the interim.

For our clock, we are using a TrueTime XL-DC GPS-disciplined ovenized quartz oscillator, which combines the unsurpassed short-term stability of quartz with the long-term stability inherent in the atomic-clock-based GPS constellation. It is this latter stability that ensures a constant-length 2.5 s yardstick. One way to think of it is this: our quartz oscillator produces a 10 MHz signal, and it is forced to make 10 million oscillations in a one second interval, with the second being defined by GPS. In reality, there is a complex filtering function with GPS control coming in at > 10 s timescales, but the idea that the GPS is dictating the length of a second gets the idea.

Because we want to synchronize our APD gate control to the clock reference, and because we would like to have greater control over the time interval measured by the time-to-digital converter (TDC), we are multiplying the 10 MHz signal by 5 to get 50 MHz. We don't use the raw 10 MHz for anything. This multiplication gives us 20 ns resolution in positioning our gate and in choosing a TDC time interval. The resultant RMS jitter between 50 MHz clock edges 2.5 s apart is something like 7 ps.

1.9 The Timing System

Like the hands on a clock–having a second hand, a minute hand, and an hour hand—APOLLO's timing scheme is a three-tiered system. APOLLO has a 25 ps hand (the TDC), a 20 ns hand (the ACM) and a second hand (the GPS clock). The basic idea is that the time interval between an APD-generated photon event and a selected 50 MHz clock pulse (roughly 30–50 ns later) is measured to 25 ps resolution with a Phillips Scientific 7186H 16-channel time-to-digital converter (TDC), whose intrinsic jitter is below 15 ps, and range is 100 ns. Performing this action on both *fiducial* photons (returning from the telescope exit aperture via a small retroreflector affixed to the secondary mirror) and *lunar* photons references these events to the 50 MHz clock train.

As a next step, the APOLLO system needs a timer (in the ACM) to keep track of the clock pulses against which these short-time measurements are made. This functionality is incorporated into the APOLLO Command Module, discussed below.

The timer is responsible for counting every 50 MHz clock pulse that goes by, and recording the time when critical events happen. In addition, the timer tells the APD array when to turn on in anticipation of the lunar return. The timer, in essence, coordinates the activities of the APD array, clock, and TDC, while providing the temporal framework to patch these things together.

The timer contains numerous counters clocked by the 50 MHz input, keeping track of: number of seconds elapsed, number of 50 MHz pulses per second (should always be 50 million), the time of a particular event within the second, how long the APD has been on, and one counter that runs free. The last counter is used to determine when the APD gate should be opened for a lunar return. If we know where this counter was when the laser pulse was fired, we can calculate the lunar distance and know what the count should be when the photon returns. By comparing this calculated value against the counter, we can take action (open the gate) when the counter reaches this magic number. Control of the APD gate width works in a similar way, shutting off after the gate counter reaches a prescribed value.

1.10 The APOLLO Command Module (ACM)

The aptly named APOLLO Command Module (ACM) is a custom single-wide CAMAC module based on Altera programmable logic devices (PLDs). The ACM is responsible for timing, coordination, and various device controls. The ACM is based on two Altera chips, one of which interfaces to the CAMAC dataway and interprets NAF CAMAC commands and generates Look-At-Me (LAM) requests. The other Altera chip contains the timer, APD control, and laser fire control.

By implementing the ACM as a CAMAC module, we gain full computer control of ACM activities, and place this timing device right next-door to the critical TDC. The ACM also acts as the primary hardware interface to most of the rest of the system, so that it controls laser fire requests, laser safety, APD array protection, rotating optic phase, and diffuser/attenuator phase.

The ACM also provides a calibration service, so that the TDC may be checked with START/STOP pulses arriving precise multiples of 20 ns apart. The generation of START and STOP pulses is carried out by the "Booster" board.

1.11 The "Booster" Board

At present, the Booster board sits in the CAMAC crate (though not in any way connected to the CAMAC dataway), and performs the following functions:

- Multiplies the 10 MHz sine-wave output of the GPS clock by 5 to deliver a 50 MHz sine wave frequency reference
- At the request of the ACM, slices individual START and STOP pulses out of the 50 MHz train in the form of ECL signals
- Provides a 50 MHz TTL reference to the ACM as the ACM's clock
- Converts the fast-photodiode's NIM signal (out of the Ortec 9327) to ECL

The Booster is currently under redesign, so that it will no longer itself provide the $5\times$ frequency multiplication, receiving and dealing only with a 50 MHz square wave. Every STOP pulse used in the APOLLO timing scheme is produced by the Booster board. It does so via a comparator fed by the 50 MHz clock, and enabled only during a 20 ns interval as specified by the ACM. Appropriate delay cables ensure that the output ECL pulse is centered within the 20 ns request signal. For calibration, a START pulse can be generated in a similar way, resulting in a calibration source comprised of START/STOP pairs an integral number of 20 ns apart.

1.12 The Control Computer

The APOLLO mission control center is a Linux box named `houston`. This is a 1 GHz Pentium III (runs cooler than Pentium IV or Athlon) with 512 MB RAM and mirrored 120 GB disks. `houston` contains a GPIB interface card for access of the GPS clock, and control of the moving optical elements (via New Focus motion controller), and a PCI-CAMAC interface for control of the CAMAC crate. `houston` also controls the rotating optic motor and communicates with the laser and chillers via serial interfaces.

It is `houston`'s responsibility to coordinate laser ranging activities and store the gathered data. Through the CAMAC interface to the ACM, in conjunction with the other hardware interfaces, `houston` sets the hardware state of the system and performs the juggling necessary to range to the moon at 20 Hz (~ 50 pulses en-route at any given time).

1.13 The Guide Camera

A CCD camera using light reflected off of a glass window on the front of the receiver enables field identification, rudimentary tracking, and provides the only mechanism for evaluating laser beam collimation (via light returned from corner cubes in telescope). The SBIG STV cameraoffers a video mode with up to 16 frames/sec video output, or a more traditional CCD integration mode with serial image transfer (at 14 sec per frame). If there are no adjustable optics between the receiver and the guide camera, the same pixel(s) on the CCD always correspond to the location of the APD array, making the act of placing a target onto the APD relatively easy.

The field of view on the 656×480-pixel CCD camera is about 35 arcsec, the field being limited by optical vignetting rather than by the chip size. In the "Normal" image or focus mode (as opposed to "Zoom" or "Wide", the CCD is binned 2×2, delivering a 320×200 image with a plate scale of 0.17 arcsec per pixel (5.8 pixels per arcsec).

1.14 Thermal Measurement and Control Scheme

APOLLO is a complex set of instruments distributed across a variety of enclosures, each with varying temperature requirements. The external environment ranges from $-15°$C to $28°$C. It is important to keep the laser and the timing electronics (especially the TDC) at a stable temperature (ideally within $1°$C). The GPS clock can track against temperature changes, but only if they are slow. The chillers and various other electronics devices are not meant to operate at temperatures below $10°$C, but otherwise are not sensitive to temperature fluctuations. We need to be very careful that we do not allow the de-ionized water that runs through the laser to freeze during cold months.

So the laser, APDs and CAMAC crate are bundled in the "Utah" box—a thick-walled insulated enclosure with a 100 W heater and air-to-water heat exchangers. The clock is currently also in the Utah box, but will soon be moved to its own dedicated enclosure because it cannot tolerate the changing orientation of the telescope. The rest of the equipment sits in the cabinet and the intermediate level enclosure (ILE). These two volumes are contiguous, and exchange air at a rate of 90 cubic feet per minute. The ILE has a 400 W heater, and an exhaust fan rated at 1600 cubic feet per minute. The ILE and cabinet cannot be cooled below the ambient temperature.

A series of RTD temperature sensors are deployed around the equipment to monitor the thermal state at one-minute intervals. We can implement a flexible set of rules for when to activate heaters, chillers, heat exchanger fans. In this way, we can control the temperature of the Utah box. And we can prevent the cabinet/ILE from overheating.

1.14.1 Plumbing Systems

There are three closed-loop plumbing systems in APOLLO. There is a laser rod cooling loop, called the Deionized loop (or DI loop), a laser chiller with associated loop, and a Utah enclosure cooling loop, carrying a propylene-glycol solution and called the PG loop.

The laser rod DI loop (whose input and output to the laser bench can be seen in Figure 2.1) is responsible for carrying heat away from the laser rods and flashlamp cavity. In normal operation, the temperature of this water is maintained at $85°$F. Because this water forms part of the optical system (laser light passes through it) relatively pure water is required, so no anti-freeze components may be added. We use distilled water, and the system contains a de-ionizing filter. This water must run through the cold dome environment, so circulation must be maintained in cold months so that there is no risk of freezing. A reservoir sits in the bottom of the laser electronics rack in a unit called the "cooling group." We have modified the cooling group so that if enabled (via front panel switch), an auxilliary pump activates when the laser electronics rack is powered off. A flowmeter monitors the flow through this loop.

The Neslab Merlin M75 chiller in the intermediate level enclosure extracts heat from the DI loop through a heat exchanger. The laser cooling group regulates the flow from the chiller so that cooling is only provided when needed to maintain the DI loop at $85°$F. In practice, not much coolant flow is needed unless the flashlamps are flashing. This plumbing loop is very short, stretching only from the chiller to the adjacent laser electronics rack.

The PG loop is driven by the Neslab Merlin M33 chiller, and this loop flows into the two (series-connected) Noren air-to-water heat exhchangers within the Utah box. These units are rated to extract 25 W of power per °C differential between air and water. An auxilliary pump and heat exchanger are rigged in parallel with the M33 chiller so that one can still provide cooling (sinking Utah toward ambient temperature) even without operating the M33. This capability is especially useful to remove constant sources of power from the Utah box, such as the XL-DC clock, or circulation fans that are always on. The propylene glycol solution will protect this line from freezing even in the coldest conditions.

1.15 Safety Systems

The APOLLO laser is not eye-safe. The collimated 9 mm beam emerging from the laser can blister skin at full power, and would without question cause serious permanent eye damage to anyone catching the direct beam in the eye. We have an evolving safety system to protect APOLLO collaborators, APO staff, visitors, pilots, and spacecraft from harm. When the laser lid is on, the Utah box closed, and the quaternary mirror in place, there is no accessible collimated laser beam except for in the column in front of the primary mirror. The beam emerging from the primary is nearly eyesafe, so that safety glasses provide adequate protection from this light level.

We have established a safety protocol that pertains to human presence in the dome when the laser is firing. This information is not repeated here, but can be summarized by saying that only authorized personnel may be in the dome when the laser is firing at full power, and all such persons must be wearing designated safety glasses with a 532 nm blocking factor of $> 10^7$. All personnel in the dome must also be made aware when the laser shutter is being opened or closed. The stairwell door into the dome must remain locked during these operations.

For external safety, we have an interlock system that controls a shutter at the output aperture of the laser enclosure. At present, two aircraft spotters on the external catwalk can control the state of this shutter, as can a person at the APOLLO cabinet. Ultimately, this shutter will be controlled by a sophisticated network of interlocks and switches. Inputs will include an infrared sky monitor to detect aircraft, and perhaps a transponder-listening device or a kill signal based on a radar feed. Also, we can interlock the stairwell door lock, activate warning signs, etc. with this system. We notify and comply with the FAA and U.S. Space Command authorities in order to avoid damage or harm to persons or equipment along our line of sight.

1.16 Interconnection

Many systems have to be married into a single functional unit for APOLLO to work. Here is an attempt to verbalize the various interconnections.

The laser is bolted to the telescope and fires its beam through an expander, the focus of which is set by the New Focus controller, via GPIB computer control. The laser itself is asked to fire either by its internal electronics (in test/setup/alignment modes) or by the rotating T/R optic through the ACM, with the control computer coordinating activities, and talking to the laser controller via a serial interface (actually emulating the remote keypad box that came with the laser).

A fast photodiode just outside of the laser box senses the green light, sending a fast negative signal to the ORTEC 9327 amplifier and timing discriminator. This device performs a zero-crossing constant-fraction measurement on the photodiode pulse, the output NIM signal of which is converted to an ECL pulse and sent to one channel of the TDC. The transmission is delayed enough so that the same STOP used for the APD fires can be used. The 100 ns TTL pulse from this discriminator is sent to the ACM to be one-shotted to 20 ns to act as both the fast-photodiode signal and the fiducial start request (thus turning on the APD gate).

The APDs have a constant 23–24 V DC power supply, to which a gate pulse is added for the purpose of biasing the APD above breakdown. This pulse is typically about 6 V in amplitude, and about 100 ns in duration. The ACM supplies the logic for this drive. The APD output is sensed by a fast comparator (AD96685), which sends an ECL pulse to the corresponding channel on the TDC.

The True-Time XL-DC sends one of its 10 MHz sine-wave signals to a 5× frequency multiplier (from Wenzel Associates) located in what is known as the “Booster board.” The emerging 50 MHz sine wave is converted to a square wave and routed to the ACM to clock all the various counters. The 50 MHz output also feeds two AD96685 comparators. Most of the time these comparators are disabled, and the ECL output is the quiescent state. At signals from the ACM corresponding to the start and end of the APD gate, these comparators are alternately enabled thereby each producing a single ECL pulse tightly associated with a given 50 MHz zero-crossing. The latter of these acts as a COMMON STOP pulse for the TDC. The other one can be used as a START pulse for TDC calibration, being a precise multiple of 20 ns before the STOP pulse. The XL-DC also sends its (10’s of microseconds long) 1 pulse-per-second (PPS) to the ACM for use in keeping track of seconds, time of event within second, and 50 MHz counts per second. Once a second, the control computer can grab the full time string from the GPS via GPIB, taking an estimated 1 ms to perform.

The T/R motor, controlled via serial interface from the computer, sends its encoder index and positional pulses to the ACM, which issues flashlamp charge and fire commands to the laser electronics rack (in the intermediate level enclosure).

The computer is connected to the New Focus multi-axis driver by way of GPIB, and has control over laser beam expansion, overall focus (common to both laser and receiver), and a steering optic that changes receiver pointing relative to transmit. This steering is especially necessary if ranging to satellites, with perhaps 10 arcsec of transmit/receive misalignment due to relativistic velocity aberration. Even for the moon, this is 0.75 arcsec, and therefore relevant compared to our 1.5 arcsec APD field of view. As long as this optic is upstream of both the guide camera and the receiver, the two will track each other, so that a given pixel on the CCD corresponds to the APD field center. For satellite tracking, this steering mirror must be dynamically adjusted as the apparent direction of the satellite changes with respect to the alt-az frame of the camera/APD.

Chapter 2

The Laser

The function and basic properties of the laser are described on page 11.

To repeat the basic parameters, the laser is a *Continuum Leopard* Nd:YAG pulsed laser, frequency-doubled to 532 nm wavelength (green), with a pulse width of ≈ 100 ps, a pulse energy of about 115 mJ (3.08×10^{17} photons per pulse), and a repetition rate of 20 Hz. The average power is then 2.3 W. The output beam is about 9 mm in diameter, with a roughly Gaussian intensity profile. Nd:YAG lasers have a fundamental wavelength of 1064 nm (infrared), and laser energy at this wavelength *does* exist within the laser enclosure. We frequency double (combine two infrared photons to make one green one) for two main reasons: silicon detectors are only a few per-cent efficient at 1064 nm (and the time-response of our APDs would additionally be horrible), and visible lasers are inherently much safer than comparably-powerful infrared lasers because there is a natural avoidance mechanism.

The Continuum Leopard laser (Figure 2.1) has been built with a few customizations to suit our needs. The laser table itself has been extended by a foot to give us a 1×4 square foot area of breadboard on which to mount our T/R switch and receiver. This way we have a rigid structure on which the transmitter and receiver are co-located, and therefore not easily misaligned. The laser electronics have been split up to allow the capacitor units to be located near the laser itself, while the heat load from the majority of the laser electronics is generated (and disposed of) in the intermediate level. We have also taken pains to limit the amount of electromagnetic interference stemming from the Marx bank—an array of capacitors that rapidly switch 4000 V onto the Pockels cell in order to dump the cavity. Most of the electromagnetic noise from the laser originates at the Marx bank, so we have shielded this component, and are working on a fully-shielded, fiber-fed version.

The laser has two shutters. One is provided with the laser (labeled G in Figure 2.1), and sits in the laser cavity. When this is in place, no lasing whatsoever happens. This shutter is controlled either from the laser remot control box or via serial interface to the laser. An additional shutter is placed at the exit aperture of the laser, and is under control of the interlock system (including spotters, ACM, operator control, etc.) This shutter merely blocks the full-power laser light from exiting the laser enclosure (i.e., doesn't reach beam expander, T/R mirror, etc.) This distinction is important: it should be understood that the laser can still be popping away at full power when the interlock shutter is closed. This has certain advantages for warmup and power checks, but if the laser is to be off for more than a minute, the internal cavity shutter should be activated rather than the exit shutter.

The laser is contained within a heavily insulated box so that we may both keep its temperature relatively constant and also minimize heat loss into the dome. This thermal shroud provides the secondary advantage that we are somewhat immune to moths and dust on our optics.

The laser bench is mounted on the primary mirror cell via a system of six mechanical flexures, critically-(but not over-) constraining it in its six degrees of freedom. This system makes the laser bench immune to

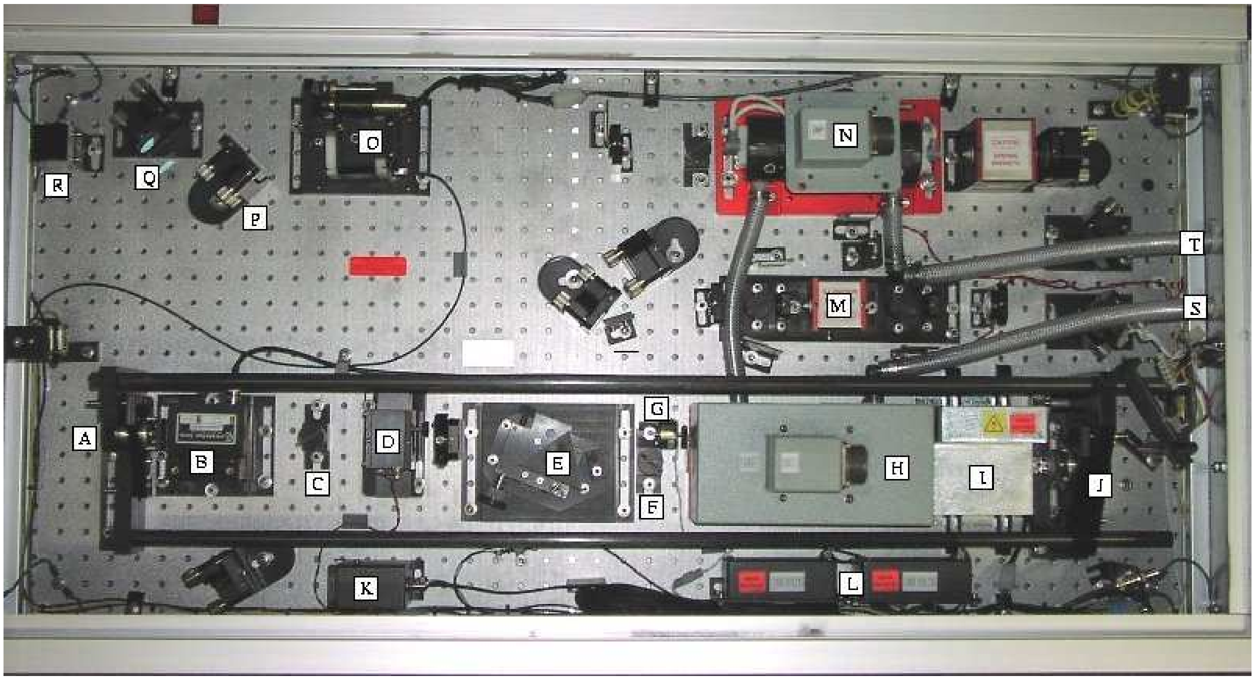

Figure 2.1: Laser assembly as delivered by Continuum. The amplifier flashlamp cover is off so that one may see the laser rod housing. Missing are the flip-dichroic and power meter, and the second output dichroic is facing the wrong way (now sends light outside). Legend: A) cavity end-mirror; B) AOML (acousto-optic mode locker); C) photodiode split-off; D) GaAs pulse-shaper; E) delay prism; F) cavity dump split-off; G) shutter; H) oscillator flashlamp/rod; I) Marx bank/Pockels cell shield; J) cavity end-mirror; K) photodiode; L) 100 V and 300 V power supplies; M) Faraday isolator; N) amplifier flashlamp/rod; O) second harmonic generator (SHG); P) output dichroic; Q) photodiode dichroic (residual); R) infrared beam dump; S) coolant intake; T) coolant output.

stresses that would arise from differential thermal contraction between the laser and telescope if mounted rigidly.

2.1 Basic Operation

At present, we turn on the laser manually at the breaker on the lower right of the electronics rack. This will ultimately be under computer control via relay. After turning the laser on, wait about one minute for the coolant flow to stabilize and flush out air. Then turn the keyswitch to the on position (and toggle off then on if already in the on position). The keyswitch too will ultimately be under computer control. Now turn on the M75 chiller that sits to the left of the laser electronics. If you forget to turn the chiller on, the laser will quickly overheat. And yes, once the other processes are automated, the chiller will also be turned on via serial interface.

At the control pad in the cabinet (to be replaced with—you guessed it—computer control), you can hit the AUTO/MANUAL button, then the START button to begin flashing. No lasing is taking place yet. The laser should flash for about 25 minutes before attempting to lase. The default startup program (PGM1) as seen on the control pad does not allow the cavity shutter to be opened, so no lasing activity can start until selecting PGM2. To do this, hit the PROGRAM UP button (upper left) *twice* so that you see PGM2 on the display, then press ACTIVATE within 5 seconds to finalize the selection. If already flashing, the laser will continue to do so. Otherwise hit START again (and obviously STOP ceases flash activity). Now the SHUTTER button will work, and lasing can begin. But always think before opening the shutter. What is the state of the interlock shutter? Where is the power meter dichroic? Is the manual shutter on the laser enclosure open or closed? How is the T/R mirror oriented? In short, what will happen to the beam if lasing commences?

At the level of this manual, no detailed information on laser tuning is offered. Only one or two select individuals are authorized to tune the laser. But there *are* some adjustments that are relevant. Almost all laser adjustments reference the output power. There is a power meter permanently mounted within the laser enclosure, and a flip-in dichroic that serves to direct all green light to the power meter, passing the infrared along to its beam dump. The first rule is not to flip the dichroic in or out *while the shutter is open*. Make sure the cavity shutter is closed first. Otherwise the optical mount can get hit with full laser power and scorch/ablate the surface. For similar reasons, do not put full laser power (> 2 W) onto the power meter for more than a few minutes at a time. The power meter is also ablated by the laser. The flip-dichroic is actuated either by the hand paddle on the top wall of the laser enclosure (difficult but not impossible to reach with the Utah upper lid on), or by computer control, provided the CAMAC crate is on.

The only relevant optical tuning accessible to anyone is the orientation of the second-harmonic generator (SHG), which is responsible for converting infrared into green light. This has a known dependence on telescope orientation (gravity direction), and is likely to need tweaks. The two upper buttons on the right column of the control box actuate this rotation in either direction.

To turn down laser power, adjust the delay knob on the left-most of the two capacitor units in the cabinet (CB 632C). Full power is at 50 μs, and by about 150 μs, the power is down to tens of mW. For alignment purposes, something in the neighborhood of 30 mW works well. If aiming for full power, be sure the flashlamps have been going for at least 30 minutes. Nominal full power is 2.3 W. Anything above 2.0 W should be considered to be full power. The laser is also sensitive to ambient temperature. If the insulating enclosure has been open for a while, the laser bench is likely out of equilibrium. Do not expect full power and robust performance under such conditions.

Also related to the thermal stability is the limitation on repetition rate. Because the laser rods deform under thermal load, their ends behave like lenses in the optical system. This system is optimized for operation at 20 Hz. Deviating from this by more than 15% (3 Hz) is likely to result in unstable, sub-optimal performance. Fortunately, even though we vary the frequency to ensure lunar photons do not return while the laser is firing, we do not have to vary the frequency by more than about 2%.

Chapter 3

The APD Array Detector

The function and basic properties of the APD array are described on page 11.

Avalanche Photodiodes (APDs) are semiconductor *p-n* junctions reverse-biased to the point where the electric field across the junction is so high that electrons running across the junction impact-ionize lattice atoms, creating more electron-hole pairs that in turn beget even more. This runaway process is the avalanche. In our scheme, the avalanche continues in a steady-state until we deliberately lower the bias voltage below the breakdown level. For our devices, the breakdown voltage is approximately 24.2 V. Biased below this, an avalanche will die on its own within nanoseconds, as the electric field is no longer strong enough for a self-sustaining runaway process.

When we "turn on" the gate, we are effectively changing the voltage from sub-breakdown (perhaps 24 V) to several volts above breakdown. The first electron to cross the junction will then probably start an avalanche. This lucky electron could stem from any of the following processes:

- a photon impinges on the silicon, creating an electron-hole pair (desired signal);
- a thermal fluctuation pops an electron-hole pair into existence;
- an electron left over from an incident photon that arrived before gate turn-on;
- an avalanche-generated photon (we believe) arrives from a neighboring element.

These four mechanisms correspond to detection of light, "dark current," "pre-pulsing," and cross-talk, respectively. In the 30 μm detectors, the dark rate is about 100 kHz, meaning that a given element has a 1% chance of experiencing a "dark event" during a 100 ns gate. The crosstalk is a strong function of distance between elements, so that the four nearest neighbors are by far the biggest contributors, with the other four "corner neighbors" being the only other noticeable participants. Given that element A is avalanching, there is a X% probability that the nearest neighbor will avalanche over a 100 ns period. Pre-pulsing is an issue for the fiducial gate events, because the laser has just fired and dumped quadrillions of photons into the enclosure. Of these that are lucky enough to find the well-shielded APD to create photoelectrons, most are promptly swept out of the junction by the electric field (which is still present with 24 V of bias, but not enough to trigger avalanche). But some may still be kicking around when the avalanche field is established, and can result in an immediate avalanche. Even innocent dark or background (light) electrons can be present when the gate comes on, so it is not unusual to see an over-abundance of events at the beginning of the gate.

The APD array is packaged in a 40-pin DIP, which is itself mounted in a socket on the "motherboard" (Figure 3.1). On the back of the motherboard, 16 individual coaxial cables route the APD signals to 16 "daughter" boards, each responsible for turning an avalanche event into an ECL logic signal that is sent to the TDC as a START signal.

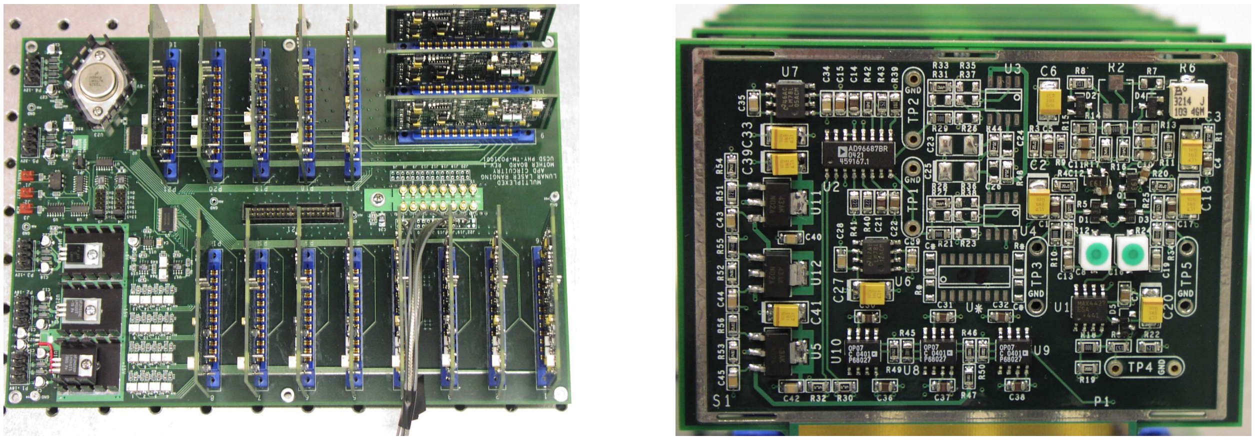

Figure 3.1: Pictures of the APD motherboard (left) and daughterboard (right). The APD itself sits on the opposite side of the motherboard behind the array of gold trim-caps on the right side.

A logic signal from the ACM commands the APDs to turn on. This signal is processed by logic on the motherboard before being sent to each of the 16 daughter boards where individual APD elements are biased above breakdown. Each element's bias voltage (in the "on" state) is individually set-able, typically sitting around 5 volts above breakdown.

In order to recover the full fill-factor at the APD array focal plane, it is necessary to place a lenslet array in front of the APD array. The APD array elements are 100 μm apart, and for the 30 μm devices, we need to concentrate the light into, say, a 25 μm diameter. Optical raytracing or application of the conservation of the $A\Omega$ product leads us to conclude that our only degree of freedom is the focal length of the lenslet arrays, and this has a direct bearing on the "pixel" scale obtained. The apparent size on the sky of a single element, size $a = 100\,\mu$m, is

$$\sigma = \frac{sa}{Df},$$

where $s = 25\,\mu$m is the concentrated light diameter, $D = 3.5$ m is the telescope aperture, and f is the lenslet focal length. Thus an off-the shelf lenslet array with 100 μm pitch and a 0.5 mm focal length gives 0.29 arcsec pixels, or 1.26×1.26 arcsec^2 field of view. A custom 0.4 mm focal length array yields a more reasonable 1.47×1.47 arcsec^2 field.

Chapter 4

Optical System

Figure 4.1 shows the optical layout on the 1-foot by 4-foot section of optical bench beside the laser. The orientation is rotated counterclockwise 90° from the “normal” orientation as seen on the telescope, such that the left of the figure is “up” in the real system. The following paragraphs walk out along the transmitter beam, and then in along the receiver beam.

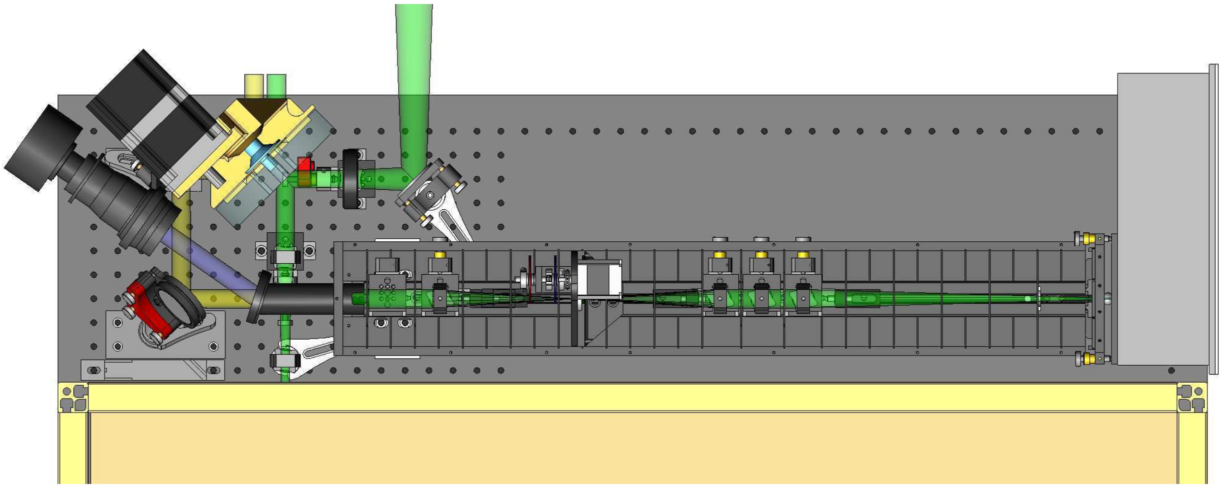

Figure 4.1: SolidWorks model of optical bench

The laser beam emerges in Figure 4.1 as a thin green beam at lower left, heading up. Immediately outside the laser enclosure, the beam passes through a negative lens (L1), diverging until it hits a positive lens (L2) that re-collimates the beam, now expanded. The second lens in the beam expander is on a motor-actuated stage for beam divergence control. Next, the expanded beam strikes the T/R mirror, reflecting to the right. The next lens (L3) is the $f/10$ expander, causing the beam to begin expanding to fill the telescope aperture. A final turning mirror (M5) sends the beam out of the Utah box through a pair of windows (not shown) with dry air sealed between them. The two windows are AR coated, and separated by air to provide thermal insulation between inside air and outside air. Outside the window is the quaternary mirror (M4). Both M5 and M4 have high-power dielectric coatings that reflect very efficiently at 532 nm, in a passband about XX nm wide. One result of this is that all light entering the system from the outside has a definite green appearance. The entire transmit beam is situated in a plane 61 mm off of the optical table.

Tracking the receive beam in Figure 4.1, we start at top center, where the converging cone of light enters the scene. This light has reflected off of M4, and entered through the dual-pane window. It reflects to the left off of M5, and would come to focus somewhere near the right edge of the optical bench but for L3 re-collimating the beam to 19.2 mm diameter. This beam passes through the T/R mirror, either unattenuated (for lunar), or suffering an attenuation factor of 10^6 if the reflective patch is in place (for fiducial returns). Hidden under the T/R motor is M6, a turning mirror (broadband coating) turned to 45° and then tipped 21.3° forcing the beam to climb above the 61 mm level experienced until now. The beam at this point in the diagram is yellow. A second broadband turning mirror, M7, is next in line, also tipped 21.3°. This mirror is on a three-axis actuated (red) mount that controls the relative pointing of the transmit and receive beams. After this lens, the beam is again level (at 4.50 in) and headed straight at the receiver assembly (long gray rectangular box). At the front of the receiver snout, the beam encounters a flat piece of uncoated glass tilted at a 17.5° angle, so that 8% of the light is diverted to the CCD camera (blue beam). At the rear of the snout sits the narrow-band filter with a 1.5 nm bandpass. The green beam depicted inside the receiver tube makes its way to the APD at far right.

The path through the receiver is as follows. The light is made to go through a spatial filter in the form of a lens (R1), a pinhole, and a second identical lens (R2). After R2, the beam is again collimated, but only light within a 3 arcsec circle on the sky has survived the pinhole. A mask (R3) trims the light pattern (to match the entrance pupil and central obstruction) so that a tidy pupil image forms on the APD. A final lens (R4) brings the light to a focus on the APD/lenslet.

Around the location of the pinhole, where the beam narrows considerably, sit three rotating glass disks—two upstream of the pinhole and one downstream of the pinhole. This is called the diffuser assembly, and is driven by a stepper motor synchronized to the T/R motor. The role of the diffuser is to both provide an additional 10^3–10^4 attenuation (variable) for the fiducial photons, and to diffuse the fiducial photons to illuminate the APDs uniformly as do the lunar photons. The mask's (R3's) main role is to shape this diffuse pattern so that the illumination at the APD is identical for both lunar and fiducial photons.

Tying up a few loose ends, when the T/R mirror is rotated such that its mirror intercepts the outgoing laser beam, the receiver is "looking" at the back surface of the mirror, and thus is staring at the side of the Utah box. This path is indicated as the yellow beam emerging from the T/R mount. Likewise, though almost all the laser light is reflected by the T/R mirror, one millionth of the light (still 3×10^{11} photons per pulse!) will pass through both front and back reflectors on the mirror. This path is indicated by the green beam emerging from the T/R mount. The T/R mirror is as thick as it is to prevent the residual laser light from being "seen" by the APDs during a laser pulse. Carefully designed beam dumps in *both* locations ensure that the receiver is looking into a dark hole, and the laser light largely disappears into its beam dump, incapable of contaminating the receiver's dark vista.

One last issue related to the optical train: at times we place temporary corner cubes around the periphery of the primary mirror, jutting into the beam so that we may check beam alignment. This reassures us that the beam emerges from the telescope collimated, and that the laser is pointing in the same direction that the receiver is looking. This operation is typically performed at low laser power, and with the T/R mirror oriented such that the laser mostly passes through, a tiny fraction reflecting off the AR-coated front surface.

Chapter 5

The Transmit/Receive Beam Switch

The function and basic properties of the transmit/receive switch are described on page 12. The T/R switch performs a variety of services:

- Allows use of full aperture for both transmit and receive modes;
- Via encoder pulses tells laser when to fire;
- Through its arrangement of coatings, protects receiver from laser light, while allowing enough through for fiducial corner-cube measurement;
- Drives diffuser disk in receiver at (typically) half the speed of the T/R mirror.

Behind the highly-reflective dielectric patch on the optic's front surface, there is a similar, larger, patch on the rear surface. Laser light scattered back toward the receiver from either dusty optics or air is mostly reflected back toward the laser by the front patch (which only moves 2 μm during the light's round-trip through telescope). Though the outgoing laser light is S-polarized (parallel to mirror face), scattering events can change the polarization. For S-polarized light, these patches reflect 99.9% of the light. For P-polarization, the reflected fraction is 99.X%. So for the fiducial corner cube—with its reflective coating that preserves polarization), this pair of patches passes 10^{-6} of the potential photons streaming back toward the receiver. This should also eliminate most sources of scatter, and help to attenuate the corner cube return to the single-photon level at the APD.

The rest of the T/R optic is anti-reflection (AR) coated. At 532 nm and 45° incidence, these surfaces reflect only 0.08% of the incident light. For looking visually at light returning from temporary corner cubes placed in front of the telescope primary mirror, it is best to let the (low power!) beam reflect off of this surface: you will get a factor of 1200 attenuation, but this is better than a factor of one-million if staring through the two reflective patches together.

Because the spinning T/R mirror directs the laser beam out of the telescope, we cannot afford to have wobble in this mirror. In general, a misalignment of the optic normal to the spin axis by an angle ϕ means a full swing (at, say, the bottom of the mirror) of 2ϕ, which then steers the beam by 4ϕ, since angles are doubled on reflection. The beam seen by the T/R mirror is a 20 mm collimated section, which means a telescopic magnification of (3.5 m)/(20 mm) = 175. Thus one arcsecond on the sky looks like 175 arcseconds inside. Then in order to keep the beam steady to 0.5 arcsec outside, we need to be steady to about 90 arcsec at the T/R mirror. This implies a limit on ϕ of 23 arcsec, which is 8 μm (0.00033 inch) at the T/R mirror radius of $\sim$75 mm (3 inches). This is a very tight tolerance, but in truth we don't need to hold this so tightly. The reason is that the outgoing beam always strikes the same point on the T/R mirror with very high precision. So the direction should be the same as long as ϕ is fixed, however large. Because the use of temporary corner cubes in front of the primary mirror to check alignment uses a *different* part of the T/R mirror (a non-reflective part), we can't ignore this requirement entirely, the net result being that we need to ensure wobble at a level smaller than 0.001 inches at the outer radius.

The rotation rate of the T/R mirror can be tweaked so that the lunar returns fall well away from the reflective patch. Since the round-trip time to the moon is $2r/c$, where r is the one-way distance, and c is the speed of light, the T/R mirror spinning at $\nu = 20$ Hz would turn $N_{\rm rev} = 2r\nu/c$ revolutions. We can adjust ν slightly so that $N_{\rm rev}$ is an odd-multiple of $\frac{1}{2}$, like 50.5, for example. Since the nearest such number is never more than half a revolution away, and the lunar round-trip is typically 50 revolutions at 20 Hz, one need adjust the frequency by only 1% maximum, so that we may favorably place the lunar returns with T/R mirror rotation rates between 19.8 and 20.2 Hz.

The motor used as the T/R drive is a "brushless servo" motor; really a stepper motor with lots of built-in smarts. All we tell the motor to do is spin at a certain angular velocity, and it manages the acceleration curve, tracking the performance relative to the requested motion, and various filtering parameters governing the control loop. We communicate with the motor via serial interface. Though we do not much exceed 1200 rpm, the motor is capable of 2000 rpm motion. Its torque capability of 1.8 N-m (250 oz-in) at speed means that it can spin the 4.2×10^{-3} kg-m^2 disk at an angular acceleration of 430 rad/s^2, or 70 revolutions per second per second. This means we could go from a stand-still to 20 Hz rotation in a mere 0.3 seconds. We never need to be this extreme, and can accommodate a 19.8 Hz to 20.2 Hz transition in a single 50 ms cycle period using less than 15% of the torque capability.

The motor also allows us to position the mirror at a select phase under remote control. The motor provides two encoder outputs at 1000 cycles per second, 90° out of phase, plus an index pulse once per revolution. We run the two encoder outputs through an XOR logic gate to derive a 2000 pulse-per-revolution encoder, and use this in conjunction with the index to fire the laser and control the diffuser disk rotation (both via ACM control).

Chapter 6

The System Clock

The function and basic properties of the System Clock are described on page 12.

The XL-DC provides the 10 MHz frequency standard locked to a true second as defined by the GPS frame. A 1-pulse-per-second output is also provided, phased with the 10 MHz signal. We communicate with the clock via a GPIB interface.

TrueTime uses a stable ovenized quartz oscillator from Wenzel Associates (same people who make our 5× multiplier) as their central oscillator, providing the 10 MHz output. The phase noise reported for this configuration is:

Offset Frequency (Hz)	SSB Phase Noise (dBc)
1	-100
10	-130
100	-145
1000	-151
10000	-153

Using these phase noise measurements to estimate clock jitter, we find that the XL-DC 10 MHz output has a jitter of 2.8, 3.0, 4.0 and 12.1 ps over time intervals of 0.1, 1.0, 2.5, and 10.0 s, respectively. The quantity of interest then is 4.0 ps. Multiplying a frequency by five introduces an unavoidable phase noise increase of $20 \log 5 = 14$ dBc, which means that the jitter for our ultimate 50 MHz reference becomes 6.2, 6.3, 6.9, and 13.3 ps for 0.1, 1.0, 2.5, and 10.0 s intervals. Here, the relevant number is 6.9 ps.

An interpretation of these statements about jitter is as follows: for the 50 MHz source, two clock edges (zero-crossings) are separated by exactly the amount expected for a perfect 50 MHz source, with 7 ps RMS (Gaussian) uncertainty added to this. Therefore for the APOLLO timing scheme, our STOP pulses—which are derived straight from the clock—are true over the 2.5 s interval to an RMS precision of 7 ps.

The XL-DC "listens" to the GPS satellites (via a 325 ft fiber link to the RF antenna) to determine a position and time solution, and compares this GPS-derived time to that of its internal ovenized oscillator. After filtering with a 2100 s time constant, any frequency offset between the oscillator and the GPS time signal is turned into a 16-bit digital number that turns into an analog voltage (via a digital-to-analog converter, or DAC) that is applied to the crystal oscillator in order to change its frequency. A new DAC value is calculated and applied every ten seconds. The DAC value (between $-16,384$ and 16,384) is made available so that we may monitor the voltage control being applied to the oscillator. The result of this feedback loop is that as long as the clock's environment changes slowly compared to the 2100 s filter timescale, the oscillator will stay locked to the GPS reference, and the frequency is reliable.

One step of the DAC corresponds to about 15 ps of change per second, turning into almost 50 ps of change over the ~ 2.5 second round-trip to the moon. This corresponds to 7 mm of one-way range between the earth and the moon. The net effect is that whenever the DAC changes value, we may see the moon appear to lurch by 7 mm. Because of this, we will monitor the DAC every 10 seconds (at all times) so that we know when these jumps occur, and can fit a smooth function to the generally monotonic trend of the DAC. In normal circumstances, the monotonic drift of the oscillator results in about ten DAC steps per day.

Much to our surprise, we found that when we located the XL-DC within the Utah box and it began riding around on the telescope, the DAC value began to vary by huge amounts (a few hundred DAC counts, corresponding to more than a meter of effective one-way range to the moon). The orientation apparently couples into the control loop and disrupts the oscillator frequency. For this reason, we must move the clock out of the Utah box, which should happen during the fall of 2005.

Chapter 7

The Timing System

A brief description of the timing system can be found on page 13. In summary, we use the 50 MHz clock to establish a yardstick against which to measure $\sim$2.5 s photon travel times. In essence, we measure the time interval between a photon event and the nearest 50 MHz clock tick using the TDC, providing 25 ps resolution. We use a counter to count the number of 20 ns clock pulses (ticks) that occur in the interim. The measurement becomes a differential measurement between the time that the laser pulse strikes the fiducial corner cube (on its way out of the telescope) and the time the pulse struck the retroreflector array on the moon. Any changes on our end (cable changes, temperature effects) are for the most-part common to both measurements, so that the difference is insensitive to such changes.

7.1 Starting the Cycle

The 50 MHz clock is the heartbeat of the timing system, and all commanded actions are synchronized to this, including the command to fire the laser. But the laser does not fire for about 300 μs after the requestto fire, and this value has a shot-to-shot variation on the order of 1.5 μs. So the laser ends up being random relative to the 50 MHz clock. This also means that we cannot base the fiducial gate of the APD off of a priori knowledge of when the laser *should* fire. We have no choice but to detect the laser fire and swing into action as quickly as we can thereafter.

The solution is a fast-photodiode receiving a small fraction of the green laser light. The dichroic mirror that directs the green light out of the laser enclosure leaks about 0.2% of the green light hitting it (see Figure 7.1). A second dichroic intercepts this residual, sending it also out of the laser enclosure. A piece of uncoated microscope slide then diverts 8% of this light to the fast photodiode, which has four sheets of paper in front of it to diffuse and attenuate the light. The roughly 150 mV signal from the photodiode is sensed by the Ortec 9327 amplifier and timing discriminator, which triggers on a fixed fraction of the peak pulse height, so that it has immunity from variable pulse strength. The 9327 puts out a 100 ns TTL pulse and also a fast NIM pulse announcing its trigger. The TTL pulse alerts the ACM to the laser fire, which responds by commanding the APD to open a fiducial gate.

7.2 The TDC

The Phillis Scientific 7186H time-to-digital converter (TDC) performs the highest-resolution time measurement in the system. This unit measures the time interval between START and STOP events by charging a capacitor at a known current for the duration of the time interval being measured. The resultant charge (voltage) on the capacitor is measured with an analog-to-digital converter (ADC) to 12-bit precision. Thus the time interval is converted to a number between 0 and 4095. The TDC has several jumper-settable ranges

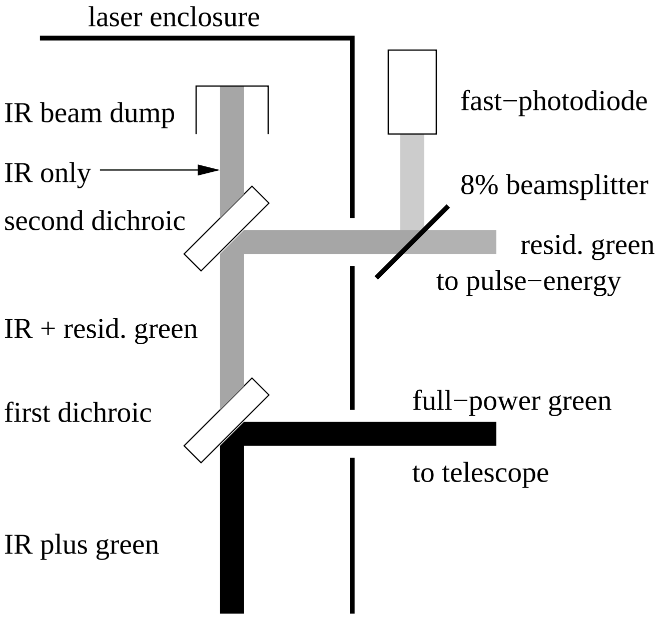

Figure 7.1: Arragement of dichroic splitters at laser output.

(100, 200, 400, 800 ns, and so on), but we use the smallest range for the highest precision. 4096 divisions over 100 ns means each “bin” is about 25 ps wide. But the TDC is actually capable of measuring time intervals to a jitter (variation) of about 13 ps RMS (about 30 ps FWHM), so that the time domain is somewhat undersampled.

We run our TDCs in COMMON STOP mode, meaning that there is only one stop pulse applied to all 16 input channels. Thus individual channels present START pulses, and all channels share the same STOP pulse. For us, START pulses are associated with photon events, and the STOP pulse is derived from the 50 MHz clock. There is also a TEST input on the TDC, which allows a single signal to be distributed to all 16 input channels simultaneously. We use this for calibration purposes, sending a START derived from the 50 MHz clock into the TEST input, followed by a STOP also derived from the 50 MHz clock. The result is a pair of pulses separated by an integral multiple of 20.00ns. The ACM and Booster together coordinate this calibration mode.

The TDCs form a weak-link in our differential measurement scheme, because the “gain” of the TDC (picoseconds per bin) varies with temperature, by about 150 parts-per-million per °C. Because of this, if our fiducial return comes in around 60 ns on the TDC, while the lunar return comes in around 30 ns, the TDC will report a difference that is temperature-dependent. If the temperature changes by 10°C, this 30 ns difference appears to change by about 45 ps, which looks like 6.5 mm of one-way lunar range. For this reason, it is important to control the TDC temperature as closely as we can. But we can also ensure that the placement of photon events within the gates is such that the TDC measures the same range for both fiducial and lunar events, so that our measurement again becomes differential. And even so, our calibration technique (STARTs and STOPs $n \times 20.00$ ns apart) allows us to monitor any departure of the TDC gain.

7.3 The Timing Sequence

A typical sequence of events, as pertaining to the timing, goes as follows. Refer to Figure 7.2, in which the circled numbers refer to the list below.

1. The ACM requests a laser fire, based on the T/R motor encoder count.
2. The laser fires, sensed by the fast-photodiode and Ortec 9327 discriminator.
3. The ACM responds to the photodiode pulse by turning on the APD gate.
4. The photodiode triggers a START on one TDC channel (channel 15, presently).
5. Photons from the fiducial corner cube hit the APD array, triggering TDC STARTs.
6. The APD gate shuts down, and a STOP pulse is requested from the 50 MHz source.
7. At this same time, the time-within-second (TWS: 50 MHz count into current second) is recorded. Also, the count on a free-running counter (FRC—never reset) is recorded. If a new second has dawned, the number of 50 MHz clock counts in that second is recorded, as is the n umber of seconds elapsed, and the time is grabbed from the XL-DC to the nearest microsecond.
8. About 25 ms later, the free-running counter reaches a value predicting the arrival of lunar photons from a previous shot (50 or so ago), and the APD is again turned on.
9. Lunar photons strike the APD array, creating TDC START pulses.
10. Steps 6 and 7 are repeated, and 25 ms later, another laser fire is requested, beginning the cycle again.

The TDC takes care of the high-resolution START–STOP measurement, and the TWS and FRC together (redundantly) tag the clock-pulse number of the STOP tick. The TWS value recorded for the outgoing laser pulse (fiducial, or FID) is used to calculate the round-trip time for that particular pulse, with that particular earth orientation, etc. The delay is converted to a multiple of 20 ns, and added to the value reported by the FRC at the time of the FID event (less a few pulses to account for gate width, etc., and wrapped if the new count exceeds the counter's maximum number). This value is loaded into a software first-in-first-out (FIFO) circular queue, which, about 50 pulses later, finds itself being sent to the ACM for comparison with the FRC. When the FRC does reach this count, the APD gate is opened for a lunar (LUN) event.

A gate-width counter is reset on LUN or FID gate-open signals, and counts to a preset value (7, for instance) upon which the APD gate is requested off. This sets the gate width to some multiple of 20 ns.

7.4 The Role of the Fast-Photodiode

In addition to serving as Paul Revere, the fast-photodiode in the laser plays a critical timing role. First, a word about why we want few photons in a fiducial event.

In order to have a fair comparison for our differential measurement, we want the fiducial photons to look no different and be treated no different than the lunar return—the same optical path, the same signal level, the same electronic processing, etc. The lunar return is expected to be 2–5 photons per pulse (though highly variable). Any more than this, and our 16 APD array elements are inadequate for handling the return. The reason is that if 16 out of 16 APD elements were hit, we would be forced to conclude that the return was very strong, and we probably only caught the first 16 out of many more detectable photons. By catching the first photons, we're biasing our measurement. On the other extreme, if we only see one photon every 6 shots, with 16 elements, the inferred probability is 1 in 96, or about 1%. We would conclude that we're not missing any photons, and our time measurement is reliable (if not dreadfully slow to accumulate statistics). We'd like to be somewhere in between—collecting photons at a reasonable rate, without irrevocably biasing our result to early times.

If, on average, we concluded that we got 1.6 photons per pulse into 16 APD elements, we would say that our hit rate was 10%. In this case, it's then true that about 10% of *those hits* are effectively double-counts. So there is a slight bias to early time measurements, but it should be correctable if we know the parent distribution (pulse shape). Then the trick is to have enough photons to get a high signal rate, but few enough that you can actually estimate the hit rate, and therefore the bias. Given the highly variable lunar return per pulse (due to speckle interference), this has to be done on a shot-by-shot basis. Probably the magic number is somewhere around 4–5 elements out of 16. Any stronger and we'll probably have to throw the return out.

To the extent that the fiducial is comparable to the lunar return in strength, any applied bias correction will in effect cancel between the two, and that's a desirable thing. So we're talking about fiducial signals in the 2–5 photon range as well—preferably at the lower end of this range.

So what happens when you get a 2-photon fiducial return? The laser-plus-APD-plus-electronics introduces a time uncertainty per photon of about 100 ps RMS. If I have two photons, do I really know where the pulse was centered? In general, I only know to about $100/\sqrt{2} = 70$ ps. Sure, the lunar returns are subject to the same uncertainty, but why add to that? Surely we can tell when the laser fired to better precision. But not with the APDs operating in the single-photon regime. That's where the photodiode comes in.

By using a high signal-to-noise-ratio measure of the laser fire time, as reported by the photodiode and constant-fraction discriminator, we have a good estimate of the laser fire time (to maybe 20 ps precision). If we send this photodiode signal to one channel of the TDC (using the same STOP pulse as for the FID photons), we get a good measure of when this pulse occurred. A cable delay can even put this in the same ballpark of the TDC. Now we have a reference by which to compare the few FID photons we've time-tagged. We can say, for instance, "this photon was 73 ps late, and that one was 15 ps early." Over the course of several minutes, we build up enough FID photons to determine the offset between the photodiode report and the average FID report. Even if this number changes (due to temperature fluctuations or whatever else), as long as the timescale for such change is longer than 5 minutes, we can use the photodiode as an **anchor** for the center of the FID photon events.

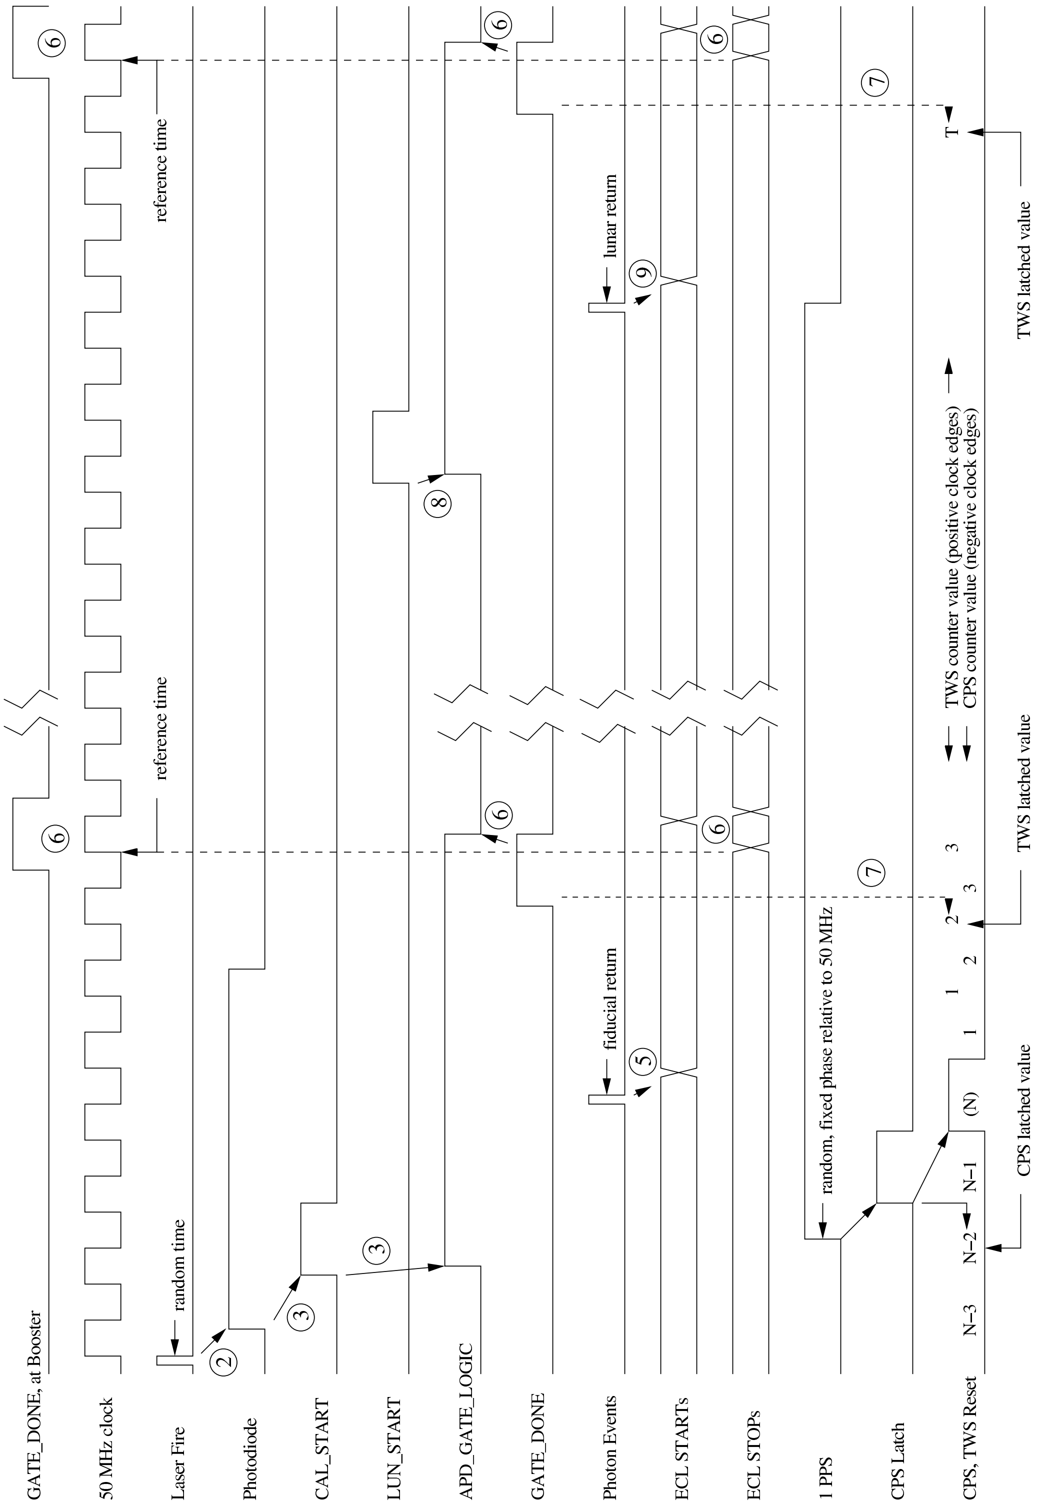

Figure 7.2: APOLLO Timing Diagram. Circled numbers refer to the list in Section 7.3.

Chapter 8

The APOLLO Command Module

The function and basic properties of the APOLLO command module are described on page 13.

8.1 ACM Description

The APOLLO Command Module (ACM) is a highly versatile customized piece of equipment providing much of the smarts in the apparatus. The ACM is a single-slot CAMAC module based on a pair of Altera 7000AE series programmable logic devices (PLDs). These chips are separated into two functions. One handles the CAMAC dataway, interpreting addresses, function calls, read and write data, and sends LAM ("Look At Me") notices to the CAMAC controller. This chip can afford to be relatively slow, and has pretty simple innards, demanding only about 35% of the smaller 7256 chip (though larger physically, with 208 pins). The timing functions are contained on the "TIMER" PLD, clocked by the 50 MHz TTL from the Booster. The TIMER chip is full of many-bit counters, and as such demands the larger 7512 chip, though fewer I/O pins are needed. The TIMER PLD is capable of being clocked at speeds of < 116 MHz, and has an end-to-end propagation delay of ~ 7 ns. The slightly slower CAMAC chip is still high speed, with clocking up to 95 MHz and 10 ns propagation delay.

The Altera 7000 PLDs are re-programmable on the circuit board, so that we can in principle actually change the internal logic and functionality with the board in place in the CAMAC crate! To the extent that the hardware pinout and external connectivity is not changed, this allows versatility in the interface and internal logic scheme.

8.2 ACM Central Concept

The ACM employs a counter/register/comparator scheme to schedule gate events. In other words, the value of a counter, clocked at 50 MHz, is compared against a pre-loaded register containing some data value. When these match, an action is taken. For example, in scheduling the lunar gate, the free-running counter (FRC) is compared against a value—set on each shot by software—that we know the FRC will reach when the lunar photons are about to return. This value has been loaded into the ACM, having been derived from a polynomial prediction using the time of laser fire as the time input to the polynomial. When the FRC reaches the prescribed value, the APD gate is turned on. The mechanism for setting the width of the APD gate works in a similar way. When the gate is turned on, a counter is reset. When this counter reaches a count that equals a preset and pre-loaded data value (the width parameter), the gate is commanded to turn off. At each gate closing, the FRC value is latched so that it may be read out (though the FRC counter keeps counting uninterrupted). Likewise, the time-within-second counter—which is reset once per second—is latched so that the time of the gate may be known to 20 ns precision.

Stare mode refers to a mode wherein the FRC is reset after it reaches the FRC delay set (thus opening a LUN-type gate). Because it is reset, it will reach the delay count again in the same amount of time, repeating the cycle. So this becomes a way to run a series of gates at a selectable period, making this a convenient way to “stare” at the sky, a star, etc. No FID-type gates are generated, only LUNs.

8.3 ACM Front Panel

The ACM front panel takes a number of inputs and provides a number of outputs, described here. The single-width panel accommodates 22 LEMO connectors at a 0.30-inch spacing. LEDs to the sides of the connectors indicate functionality.

The bottom five connectors on the ACM are labeled as input/output lines. These are jumper-configurable to act either as outputs or inputs to the Altera TIMER chip. At present, the Altera chip is programmed to have all five act as outputs. The first of these has been dedicated to serve as the diffuser stepper motor step request. The other five are configured so that CAMAC commands can set them to a TTL high state or low state (default).

8.3.1 Inputs

Logic Name	Front Panel	Signal Type	Connects To:	Action
CLOCK	CLK	50 MHz square	Booster Clock	ACM time ref.
PPS	1 PPS	20 μs TTL	XL-DC 1 PPS	alerts new second
PHOTODIODE	PIN	100 ns TTL	9327 TTL out	alerts laser fire
DIFF_INDEX	CAL	~ 4 kHz TTL	diffuser opto-isolator	1 per revolution
TR_INDEX	IDX	$\sim 50\,\mu$s TTL	T/R opto-iso index	1 per revolution
TR_ENCODER	ENC	~ 40 kHz TTL	T/R opto-iso encoder	2000 per revolution

The first three inputs are terminated in 50 Ω, while the last three are terminated in 4 kΩ.

8.3.2 Outputs

Logic Name	Front Panel	Signal Type	Connects To:	Action
LUN_START	LUN ST	20 ns TTL	Booster START	slices START pulse
CAL_START	CAL ST	20 ns TTL	—	on FID gate open
GATE_START	OPN	20 ns TTL	—	on any gate open
GATE_DONE	CLS	20 ns TTL	—	on gate close
STOP	STOP	20 ns TTL	Booster STOP	slices STOP pulse
GET_TIME	TME	160 ns TTL	XL-DC trig in	latch GPS time
APD_GATE_LOGIC	GTE	$N \times 20$ ns TTL	—	mimics gate
EN_APD_GATE	APD	$N \times 20$ ns TTL	APD gate input	turns on APD
LASER_CHARGE	LSR	DC TTL	laser “externals”	tells caps to charge
LASER_FIRE	ZAP	40 μs TTL	laser “externals”	tells laser to fire
LASER_BLOCK	BLK	DC TTL	interlock	shutters laser output
DIFFUSER_STEP	I/O 1	~ 4 kHz TTL	diffuser opto-iso	drives diffuser motor
I/O 2	I/O 2	DC TTL	—	—
I/O 3	I/O 3	DC TTL	—	—
I/O 4	I/O 4	DC TTL	—	—
POWER_METER	I/O 5	DC TTL	flip dichroic	illuminates pwr meter

All outputs are capable of driving 50 Ω termination. The LUN_START occurs only at the beginning of a lunar gate, and it is this that requests a START clock pulse from the Booster when the calibration mode is enabled. The STOP signal and GATE_DONE signal are identical: both occur at the end of a gate. The APD_GATE_LOGIC and EN_APD_GATE are effectively identical, though the former activates for every gate event, while the latter is onnly activated if the APD drive is enabled.

In addition to these outputs, various LEDs indicate the status:

- An orange "N" light blinks when the module is being addressed via CAMAC commands.
- A green LED blinks at one pulse per second to indicate clock heartbeat.
- A blue-white LED blinks with the APD gate logic to indicate gate action.
- A yellow LED indicates the APD drive is enabled, and thus follows the gate logic.
- A green LED indicates that the laser is enabled.
- A blue-white LED flashes when a laser fire request is issued.
- A red LED indicates that the laser is unblocked and able to fire at the sky.

8.4 ACM Commands

The ACM commands take the form of CAMAC NAF commands (N = slot number; A = address, F = function), and follows standard CAMAC numbering conventions. A complete list of the commands is included in Appendix B. Here, we summarize the commands in an abbreviated form, the intention being to flesh out the functionality of the ACM. In the following, FRC means free-running counter, TWS means the time within second counter, CPS is the counts-per-second counter, and a LAM is a look-at-me interrupt request.

8.4.1 Reading Data

Upon the GATE_DONE signal, the FRC and TWS counters are latched, and each of these can be read out. Upon the 1 PPS signal, the CPS and accumulated second counters are latched, and these can be read out. The T/R encoder counts between the T/R index and diffuser index indicates the diffuser motor phase, and can be read out. The value set fo the gate width and the value stored as the FRC delay can be read.

8.4.2 Writing Data

The following data can be written to the ACM: delay setting (comparison trigger value) for the FRC; gate width parameter; T/R mirror phase delay (at what phase mirror fires); laser fire period for internally generated firing; phase adjustment parameters for modifying the diffuser phase.

8.4.3 Setting States

Various commands exist for clearing registers and LAMs. The 50 MHz counters can all be reset. The accumulated seconds counter can be reset. The APD gate can be commanded to the off position. In addition, the following state changes can be made. In what follows, the default state (happens on CAMAC initialization command) is the first listed, the alternate state indicated after the slash.

Stare mode can be disabled/enabled. The APD drive can be disabled/enabled. All 50 MHz counters can be disabled/enabled. The laser drive can be disabled/endabled. The laser shutter can be blocked/unblocked.

The LUN_START output can be disabled/enabled—when enabled, the Booster processes START pulses that are used in TDC calibration. The laser fire mode can be switched to internal pulse-train generation or T/R encoder drive. Each of the I/O outputs (2–5) can be set to low/high states.

LAMs can be configured to fire on lunar gates, fiducial gates, GATE_DONE (all gates) via a LAM mask. The LAM status can be read, as well as LAM request (those LAMs enabled by the mask). LAMs can be software-generated, which can be useful in verifying LAM operation and response.

Chapter 9

The Control Computer

The function and basic properties of the hardwre control computer are described on page 14. The APOLLO computing tasks can be broken into three main categories: hardware control; interface to the observatory; graphical monitoring of data and issuing of commands. These tasks are broken out to different computer functions known as `houston`, the ICC (instrument control computer), and TUI (telescope user interface). At present, the ICC is a program that runs on `houston` itself, but may later occupy its own machine. We will talk about the ICC as if it is a seperate machine. `houston` talks to the ICC directly, and the ICC communicates data to the “hub,” which then disseminates information out to all TUIs (multiple machines can be running TUI).

The tasks of the hardware control computer (named `houston`, since it serves as APOLLO mission control) are:

- communicate with the CAMAC crate to get TDC and ACM data, and set the hardware state through the ACM,
- communicate with the XL-DC to get GPS-accurate time strings,
- communicate with the T/R motor to control spin rate and accelerations,
- communicate with the optics controller to set focus, beam collimation, and transmit/receiver offset,
- communicate with laser control unit to enable and activate laser,
- communicate with chillers to turn on/off, set temperature,
- compute the target delay for each outgoing pulse,
- write collected data to disk, and make available to separate user computer,
- monitor temperatures and activate thermal control mechanisms,
- accept parameter and control commands from the user computer.

The majority of the laser ranging is therefore coordinated by `houston`.

The ICC retrieves real-time data from `houston` and packages this data for transmission to the hub. The ICC will ultimately issue telescope move commands for tracking and acquisition, and both issue and pass-on commands to `houston` to control its operation and states. The hub passes data to the TUI in keyword=value format, so that the TUI may present a graphical display of the APD array, histograms of the

TDC hits, etc. Commands may be sent to `houston` from the TUI, which may be interpreted by the ICC and passed on to `houston` in a way that it can understand.

The bulk of the software on `houston` is in the C language, as this is the language most thoroughly supported for GPIB, CAMAC, and serial interfacing. The `houston` software is a complex code incorporating real-time scheduling (at the millisecond level) in a multi-thread, multi-prioriy, resource-locking way. Most of the internal loop of the laser ranging operation consists of CAMAC commands, with occasional GPIB, serial, and computational tasks.

The user interface is written in Python, running in conjunction with TUI, and using Hippodraw as its graphical resource. This provides a reasonably slick graphical interface, and is straightforward to integrate into the APO environment.

Chapter 10

Modes of Operation

Having described the major components of the APOLLO system, this chapter describes the actions taken by the control program on `houston` for the various states of operation. This program is called `housctl`, and comprises several source files (`housctl.c` and `housutil.c`), the latter of which performs many of the background functions like monitoring temperature and activating thermal controls. The following is organized by operating state, of which `housctl` has many—numbered 0–7. Entering state 0 performs a clean exit of `housctl`, turning off all equipment and closing files, etc.

10.1 IDLE State (`state =` 1)

In IDLE state, nothing but "background" tasks are active. These include monitoring the temperatures, flowmeters, and GPS clock (especially the DAC). The following actions are taken upon entering the IDLE state:

- Set GPS clock statistics interval to 10 seconds
- Set temperature log interval to 60 seconds
- Set flowmeter log interval to 120 seconds
- Begin loop to check scheduled tasks and look for ICC commands

This last entry is performed in most of the other states, and is referred to as "background" tasks. In the IDLE state, nothing happens between checks—just a wait. But in the STARE state, for instance, a gate/data cycle happens between each background check.

10.2 WARMUP State (`state =` 2)

This state is not yet fully functional, but will become so when we need an automated 30-minute warmup cycle. When this happens, all equipment will be turned on to establish a stable thermal state, the laser flashlamps will run, and lasing will commence about 5 minutes before show time. At present, the WARMUP state performs the following:

- Set GPS clock statistics interval to 10 seconds
- Set temperature log interval to 10 seconds

- Open the data file (will be changed to be associated with reflector)
- Read polynomial coefficients
- Power on the CAMAC Crate and the APD cooling fan
- Begin loop to check scheduled tasks and look for ICC commands

10.3 RUNNING State (`state =` 3)

This is the main state for acquiring lunar range data. It *should* only be entered from the WARMUP state (and soon also STANDBY state). Once in RUNNING state, the following setup actions are taken:

- Open connections to the CAMAC crate and the XL-DC clock
- Set the XL-DC into hardware trigger mode
- Initialize CAMAC Crate
- Disable ACM laser drive, block laser beam (ACM input to interlock)
- Set ACM T/R mirror phase for laser fire (170 counts after index, currently)
- Set ACM gate width (to 180 ns, currently)
- Set ACM LAM mask to listen to GATE_DONE type (all gates generate LAM)
- Clear the ACM LAM status
- Reset the ACM accumulated seconds counter
- Based on polynomial prediction's round trip time, set T/R motor speed for interleaved return
- Wait 6 seconds for T/R motor to spin up
- Set ACM laser drive to follow T/R encoder
- Enable ACM laser drive
- Unblock ACM hold on laser exit shutter
- Set TDC lower and upper thresholds (to 2 and 4093, respectively)
- Set up CAMAC LAM mask to pay attention to ACM and TDC; Clear LAMs; Enable interrupts
- Enable ACM APD drive
- Enable ACM 50 MHz counters
- Reset TDC (clear data, LAMs, make ready for START/STOP cycle)

Now the ACM is out of the blocks, and begins firing the laser, responding to the photodiode START signal by opening fiducial gates. But we have to refrain from taking the fiducial gates seriously until the GPS clock triggers are synchronized to the operation. The first time a new second rolls over, the GPS clock's triggered time is read, but the first one is a vestigial from some earlier, irrelevant trigger. The fiducial gates leading up to this first rollover are labled as shot number -2. During the next second, the fiducial shot numbers are all -1. As long as the shot number is negative, no lunar gates are scheduled, because we do not yet have an accurate GPS time on which to base the prediction. As soon as the second GPS trigger-time is read, the next shot is numbered 1, and begins to increment by 1 with each shot. Now that there is a valid GPS trigger time, the following sequence ensues:

- The computer waits for a LAM interrupt, signalling the end of a gate
- The accumulated seconds counter, the FRC counter, and the TWS counter are read from the ACM
- TDC data is read, employing the sparse-read technique
- The TDC is reset, and LAMs are cleared
- If a second has wrapped (as judged by TWS), get the latest GPS time
- Check CPS counter to verify that no 50 MHz pulses were missed (or added) in last second
- Combine GPS time and TWS to make 20 ns-accurate gate closing time
- If an integral multiple of 10 seconds has elapsed, get the clock statistics (e.g., DAC)
- If the previous gate was a fiducial type, calculate the round-trip time from the prediction polynomial using the combined GPS and TWS time stamp
- Convert this prediction to an FRC value, make appropriate adjustments, and wrap to 28 bits
- Queue the prediction for later use (go ahead and send to ACM if this is the first queue entry)
- If the gate type was lunar, pop the oldese queue entry into the ACM FRC delay set
- Write the data for this event to disk and send to ICC
- Check to see if T/R motor speed needs to be changes to appropriately phase fiducials and lunars
- Perform background functions, such as temperature and flow checks
- Return to top of list, waiting for next LAM

When the shot count reaches its prescribed maximum, perform the following actions:

- Disable the ACM laser drive
- Disable the ACM APD drive
- Disable the 50 MHz counters
- Reset the CAMAC LAM mask back to its previous state
- Close the CAMAC crate
- Stop the T/R motor
- Set state to enter COOLDOWN (will be changed to STANDBY)

10.4 COOLDOWN State (`state =` 4)

The COOLDOWN state is intended to give the equipment time to cool down while still under active thermal control (e.g., chiller on). At the moment, this is a very short-lived state automatically entered and exited on finishing the RUNNING state. We will soon change to enter a STANDBY state right after the RUNNING state, and make COOLDOWN more fully functional (actually turning off equipment, etc.). But for now, COOLDOWN simply does:

- Set dwell time of COOLDOWN (now 5 seconds: will be about 30 minutes)
- Set temperature log interval to 10 seconds
- Set GPS statistics interval to 10 seconds
- Run background tasks
- Before exiting, close data file from previous RUNNING state

10.5 STARE State (`state =` 5)

Stare mode (described on page 38) enables us to look at the flux impinging on the APD array from a star, the moon, or whatever. To enter STARE mode, we:

- Set the temperature log interval to 60 seconds
- Set the GPS statistics interval to 10 seconds
- Set the flowmeter log interval to 120 seconds
- Open a data file
- Power on the CAMAC crate and APD cooling fan
- wait five seconds for CAMAC warmup
- Open a connection to the CAMAC crate
- Initialize the CAMAC crate
- Set the APD gate width (currently 180 ns)
- Set the ACM LAM mask to GATE_CLOSE (all gate events)
- Clear ACM LAMs, reset accumulated seconds counter
- Enable ACM Stare mode at a specified rate (currently 1000 gates per second)
- Enable TDC LAMs and set TDC thresholds
- Set up CAMAC LAM mask to pay attention to both ACM and TDC LAMs
- Enable LAM interrupts
- Enable ACM 50 MHz counters

- Enable ACM APD drive
- Reset TDC (clear data, LAMs, make ready for START/STOP cycle)

Now that the setup is done, enter the loop:

- Wait for LAM event
- Read ACM accumulated seconds, FRC, and TWS counter values
- Read the TDC values via sparse read technique
- Reset TDC (clear data, LAMs, make ready for START/STOP cycle)
- When a preset number of gates have transpired (now 500), write data to file and output to ICC
- If a new second has elapsed, check the CPS counter for extra/missing 50 MHz pulses
- Reset CAMAC LAMs
- Check to see if background tasks are due & do them if necessary

Stare mode is exited by an ICC-initiated change of state. When this happens, the following actions are executed:

- Disable ACM 50 MHz counters
- Disable ACM APD drive
- Disable ACM Stare mode
- Reset CAMAC lam mask to original state
- Close CAMAC crate and data file

10.6 FIDLUN State (`state` = 6)

The FIDLUN state is meant as an in-dome test mode that mimics the RUNNING state almost entirely, with the following differences:

- Setup: Power on the CAMAC Crate and the APD cooling fan & wait 5 seconds
- Setup: Write static FRC delay of 1,000,000 (20 ms delay)
- Setup: No T/R phase, no T/R laser drive enable, no T/R speed calculation or spin-up
- Setup: Enable ACM internal laser drive at prescribed frequency (default = 20 Hz)
- Running: The real difference is no polynomials, no predictions, no queuing of FRC delays: the lunar gates simply arrive 20 ms after the fiducials in a static way.

The FIDLUN state is useful for testing aspects of the fiducial gates: the photodiode start pulse, the acquisition of fiducial photon events, etc. The lunar gates provide a handy reference as to the behavior of the APDs in the absence of laser light. No WARMUP state is required.

10.7 STANDBY State (`state` = 7)

Though not currently implemented, the STANDBY state will allow us to enter the RUNNING state multiple times in an observing session without cycling through WARMUP and COOLDOWN states. If several minutes go by, the laser flashlamps will begin to pulse again to maintain the thermal state. If more than 15 minutes goes by, COOLDOWN will be entered by default.

Appendix A: Acronyms

ACM	Apollo Command Module: controls timing from the CAMAC crate (pp. 13, 37)
APD	Avalanche Photodiode (pp. 11, 23)
APOLLO	Apache Point Observatory Lunar Laser-ranging Operation
AR	Anti-Reflection
CAMAC	Computer Automated Measurement and Control
CCD	Charge-Coupled Device (electronic camera: p. 14)
CPS	Counts Per Second (ACM data value)
DAC	Digital to Analog Converter
DAQ	Data Acquisition
DIP	Dual In-line Package
ECL	Emitter-Coupled Logic: differential logic scheme with low at -1.6 V, high at -0.8 V
FID	not an acronym, but stands for fiducial (as in fiducial corner-cube return)
FRC	Free-Running Counter (ACM data value)
FWHM	Full-Width at Half-Maximum
GPIB	General Purpose Interface Bus
GPS	Global Positioning System
ICC	Instrument Control Computer (Chapter 9)
ILE	Intermediate Level Enclosure (Section 1.3)
LAM	Look At Me: an interrupt request from one of the CAMAC modules
LLR	Lunar Laser Ranging
LUN	not an acronym, but stands for lunar (as in lunar-type return)
NAF	CAMAC command structure: slot N; address A; function F
NIM	Nuclear Instrumentation Methods (also fast, -1 V timing signal)

PLD	Programmable Logic Device
PPS	Pulse-Per-Second (ACM data value)
RF	Radio Frequency
RMS	Root-Mean-Square
RTD	Resistance Temperature Detector
SHG	Second-Harmonic Generator (in laser: Chapter 2)
STV	model name of our CCD camera (p. 14)
TCC	Telescope Control Computer (Chapter 9)
TDC	Time-to-Digital Converter: measures intervals < 100 ns to 25 ps resolution
T/R	Transmit/Receive (pp. 12, 27)
TUI	Telescope User Interface (Chapter 9)
TWS	Time Within Second (ACM data value)
XL-DC	model name of the GPS-disciplined clock

Appendix B: ACM Commands

Following is a list of the CAMAC-actuated ACM commands in A(x)F(x) form, with brief descriptions.

A	F	Function Description
0	0	Read lower 16 bits of free-running counter (FRC)
1	0	Read upper 12 bits of free-running counter (FRC)
2	0	Read lower 16 bits of time-within-second counter (TWS)
3	0	Read upper 10 bits of time-within-second counter (TWS)
4	0	Read lower 16 bits of counts per second (CPS)
5	0	Read upper 10 bits of counts per second (CPS)
6	0	Read accumulated seconds counter
7	0	Read diffuser motor phase
0	1	Read lower 16 bits of FRC delay set
1	1	Read upper 12 bits of FRC delay set
2	1	Read gate width set value
12	1	Read LAM status (1=FID, 2=GATE_DONE, 4=LUN), enabled or not
13	1	Read LAM mask (1=FID, 2=GATE_DONE, 4=LUN)
14	1	Read LAM request (1=FID, 2=GATE_DONE, 4=LUN), live LAM
15	1	Read Gate type: 0=FID, 1=LUN
15	8	Test LAM (returns status on Q: Q=0 if no LAM, Q=1 if LAM set)
0	9*	Clear group 1 registers (data values: FRC, TWS, CPS, PIN)
0	10	Clear LAM status register
0	11*	Clear group 2 registers (set values: FRC delay, gate width)
1	11*	Clear groups 1 & 2 registers
2	11	Clear LAM status register; Clear group 1 registers

3	11	Clear LAM status register; Clear groups 1 & 2 registers
0	12	Reset accumulated seconds counter
1	12	Reset all 50 MHz counters
2	12	Reset accumulated seconds and 50 MHz counters
3	12*	Disable Stare mode (enable normal ranging mode)
4	12	Enable Stare mode (for looking at stars, background)
0	13	Request APD gate OFF
1	13*	Disable APD drive (logic still goes, but APD drive disabled)
2	13	Enable APD drive
3	13*	Disable 50 MHz clock
4	13	Enable 50_MHz clock
5	13*	Disable Laser drive (ACM can't fire laser)
6	13	Enable Laser drive (ACM allowed to fire laser)
7	13*	Block laser beam
8	13	Unblock laser beam
9	13*	Disable START pulse request via LUN_START output
10	13	Enable START pulse request via LUN_START output
11	13*	Switch laser fire mode to internal 20 Hz generator
12	13	Switch laser fire mode to T/R switch drive
13	13*	Disable external CAL_START input: Go with PIN for gate control
14	13	Enable external CAL_START input for gate turn-on
2	14*	Set I/O 2 output low
3	14	Set I/O 2 output high
4	14*	Set I/O 3 output low
5	14	Set I/O 3 output high
6	14*	Set I/O 4 output low
7	14	Set I/O 4 output high
8	14*	Set I/O 5 output low
9	14	Set I/O 5 output high

0	17	*data*	Write lower 16 bits of delay data for FRC
1	17	*data*	Write upper 12 bits of delay data for FRC
2	17	*data*	Write gate width set data
3	17	*data*	Write mirror phase delay
4	17	*data*	Write lower 16 bits of internal laser fire period
5	17	*data*	Write upper 10 bits of internal laser fire period
6	17	*data*	Write diffuser phase adjust command
12	19	*mask*	Set LAM status (1=FID, 2=GATE_DONE, 4=LUN) simulate LAM
13	19	*mask*	Set LAM mask (1=FID, 2=GATE_DONE, 4=LUN)
0	20		Strobe FRC delay write (transfer all 28 bits to timer chip)
1	20		Strobe internal laser fire period
2	20		Reserved Strobe
3	20		Reserved Strobe
12	23	*mask*	Clear LAM status
13	23	*mask*	Clear LAM mask
*			Happens on Initialize (Z) command