

The Apache Point Observatory Lunar Laser-ranging Operation: Instrument Description and First Detections

T. W. MURPHY, JR.,¹ E. G. ADELBERGER,² J. B. R. BATTAT,³ L. N. CAREY,⁴ C. D. HOYLE,⁵ P. LEBLANC,¹ E. L. MICHELSEN,¹
K. NORDTVEDT,⁶ A. E. ORIN,¹ J. D. STRASBURG,⁷ C. W. STUBBS,⁸ H. E. SWANSON,² AND E. WILLIAMS¹

Received 2007 August 23; accepted 2007 November 13; published 2008 January 22

ABSTRACT. A next-generation lunar laser-ranging apparatus using the 3.5 m telescope at the Apache Point Observatory in southern New Mexico has begun science operation. The Apache Point Observatory Lunar Laser-ranging Operation (APOLLO) has achieved 1 mm range precision to the moon, which should lead to approximately 1 order-of-magnitude improvements in several tests of fundamental properties of gravity. We briefly outline the scientific goals, and then give a detailed discussion of the APOLLO instrumentation.

1. INTRODUCTION

1.1. Scientific Motivation

A variety of observations and theoretical explorations—including the apparent acceleration of the expansion of the universe (Perlmutter et al. 1999; Riess et al. 1998), the possible existence of extra dimensions (Arkani-Hamed et al. 1998), and attempts to reconcile quantum mechanics and gravity—provide motivation for improved tests of the fundamental aspects of gravity.

Lunar laser ranging (LLR) currently provides the best tests of a number of gravitational phenomena (Williams et al. 1996, 2004) such as:

1. The strong equivalence principle (SEP): $\eta \approx 5 \times 10^{-4}$ sensitivity
2. Time rate of change of the gravitational constant: $\dot{G}/G < 10^{-12} \text{ yr}^{-1}$
3. Geodetic precession: 0.6% precision confirmation
4. Deviations from the $1/r^2$ force law: $\sim 10^{-10}$ times the strength of gravity at 10^8 meter scales

LLR also tests other gravitational and mechanical phenomena, including, for example, gravitomagnetism (Murphy et al. 2007), preferred frame effects (Müller et al. 1996; Nordtvedt 1987), and Newton’s third law (Nordtvedt 2001). LLR may also provide a window into the possible existence of extra dimen-

sions via cosmological dilution of gravity (Lue & Starkman 2003; Dvali et al. 2003). In addition to the SEP, LLR tests the weak equivalence principle (WEP) at the level of $\Delta a/a < 1.3 \times 10^{-13}$, but the LLR constraint is not competitive with laboratory tests. In addition, LLR is used to define coordinate systems, probe the lunar interior, and study geodynamics (Dickey et al. 1994).

These constraints on gravity are based on about 35 years of LLR data, although the precision is dominated by approximately the last 15 years of data at 1–3 cm precision. APOLLO aims to improve tests of fundamental gravity by approximately an order of magnitude by producing range points accurate at the 1 mm level.

1.2. A Brief History of LLR

The first accurate laser ranges to the moon followed the landing of the first retroreflector array on the Apollo 11 mission by less than two weeks (1969 August 1). These were performed on the 3.0 m telescope at the Lick Observatory. One month later, a second station using the 2.7 m telescope at the McDonald Observatory began ranging to the moon (Bender et al. 1973). The operation at the Lick Observatory was designed for demonstration of initial acquisition, so that the scientifically relevant observations over the next decade came from the McDonald station, which used a ruby laser with 4 ns pulse width, firing at a repetition rate of about 0.3 Hz and $\sim 3 \text{ J pulse}^{-1}$. This station routinely achieved 20 cm range precision, with a photon return rate as high as 0.2 photons per pulse, or $0.06 \text{ photons s}^{-1}$. A typical “normal point”—a representative measurement for a run typically lasting tens of minutes—was constructed from approximately 20 photon returns.

In the mid 1980s, the McDonald operation was transferred to a dedicated 0.76 m telescope (also used for satellite laser ranging) with a 200 ps Nd:YAG laser operating at 10 Hz and $150 \text{ mJ pulse}^{-1}$. This station is referred to as the McDonald

¹Department of Physics, University of California, San Diego, La Jolla, CA; tmurphy@physics.ucsd.edu.

²Department of Physics, University of Washington, Seattle, WA.

³Harvard-Smithsonian Center for Astrophysics, Cambridge, MA.

⁴Department of Astronomy, University of Washington, Seattle, WA.

⁵Department of Physics and Astronomy, Humboldt State University, Arcata, CA.

⁶Northwest Analysis, Bozeman, MT.

⁷Pacific Northwest National Labs, Richland, WA.

⁸Department of Physics, Harvard University, Cambridge, MA.

Laser Ranging System (MLRS; Shelus 1985). At about the same time, a new station began operating in France at the Observatoire de la Côte d’Azur (OCA) (Samain et al. 1998). Using a 1.5 m telescope and a 70 ps Nd:YAG laser firing at 10 Hz and 75 mJ pulse⁻¹, this became the premier lunar ranging station in the world. In recent years, the MLRS and OCA stations have been the only contributors to lunar range data with typical return rates of 0.002 and 0.01 photons per pulse, respectively. Typical normal points from the two stations consist of 15 and 40 photons, respectively.

Other efforts in LLR are described in Williams et al. (2005), and more detailed histories may be found in the preceding reference as well as in Bender et al. (1973) and Dickey et al. (1994).

1.3. Millimeter Requirements

The dominant source of random uncertainty in modern laser ranging systems has little to do with the system components, but rather comes from the varying orientation of the lunar retroreflector arrays. Although the arrays are nominally pointed within 1° of the mean earth position, variations in the lunar orientation—called libration—produce misalignments as large as 10°, and typically around 7°. This means the ranges between the earth and the individual array elements typically have a root mean square (rms) spread of 15–36 mm, corresponding to about 100–240 ps of round-trip travel time. This dominates over uncertainties associated with the laser pulse width, and with jitter in the detector and timing electronics. A typical normal point containing 16 photons will therefore be limited to 4–9 mm range precision by the array orientation alone, though range residuals reported by analysis at the Jet Propulsion Laboratory tend to be larger than this.

Reaching the 1 mm precision goal demands at a minimum the collection of enough photons to achieve the appropriate statistical reduction. Assuming an ability to identify the centroid of N measurements—each with uncertainty σ —to a level of $\sigma_{\text{net}} = \sigma/\sqrt{N}$, the uncertainty stemming from the retroreflector array orientation typically demands 225–1300 photons in the normal point to reach the 1 mm mark. Worst-case orientations push the individual photon uncertainty to 50 mm, demanding 2500 photons. This is far outside of the capabilities of the aforementioned LLR stations. We point out that any constant range bias is accommodated in the analysis, so that only *variations* in the range are important to the experiment.

While adequate photon number is sufficient to reduce statistical uncertainty to the 1 mm level, other sources of error could potentially limit the ultimate scientific capacity of LLR. Most importantly, the gravitational physics is sensitive to the center-of-mass separations of Earth and Moon, while one measures the distance between a telescope and reflectors that are confined to the body surfaces. The earth’s surface in particular has a rich dynamic—experiencing diurnal solid-earth tides of 350 mm peak-to-peak amplitude, plus crustal loading from oceans, at-

mosphere, and ground water that can be several millimeters in amplitude. Moreover, the Earth atmosphere imposes a propagation delay on the laser pulse, amounting to ~1.5 m of zenith delay at high-altitude sites. Satellite laser ranging, very long baseline interferometry, and other geodetic efforts must collectively contend with these same issues, for which accurate models have been produced. A good summary of these models is published by the International Earth Rotation and Reference Systems Service (IERS; McCarthy & Petit 2003).

As an example of the state of these models, the long-standing atmospheric model by Marini & Murray (1973) has recently been replaced by a more accurate model (Mendes et al. 2002; Mendes & Pavlis 2004). The model differences for a high-altitude site are no more than 2 mm for sky elevation angles greater than 40°—providing an indicative scale for the model accuracy. The primary input for this model is the atmospheric pressure at the site, as this represents a vertical integration of atmospheric density, which in turn is proportional to the deviation of the refractive index, n , from unity. Thus the zenith path delay, being an integration of $n - 1$ along the path, is proportional to surface pressure under conditions of hydrostatic equilibrium. A mapping function translates zenith delay to delays for other sky angles. Measuring pressure to 1 part in 2000 (0.5 mbar) should therefore be sufficient to characterize the 1.5 m zenith delay at the 1 mm level. Our experiment records atmospheric pressure to an accuracy of 0.1 mbar.

The principal science signals from LLR appear at well-defined frequencies. For example, the equivalence principle signal is at the synodic period of 29.53 days, and even secular effects (\dot{G} , precession) are seen via the comparative phases between periodic (monthly) components in the lunar orbit. Because many of the effects discussed in the preceding paragraphs are aperiodic, they will not mimic new physics. To the extent that these effects are not adequately modeled, they contribute either broadband noise or discrete “signals” at separable frequencies.

The science output from APOLLO may be initially limited by model deficiencies. But APOLLO’s substantial improvement in LLR precision, together with a high data rate that facilitates deliberate tests of the models, is likely to expose the nature of these deficiencies and therefore propel model development—as has been historically true for the LLR enterprise. Ultimately, we plan to supplement our LLR measurement with site displacement measurements from a superconducting gravimeter (not yet installed) in conjunction with a precision global positioning system installation as part of the EarthScope Plate Boundary Observatory (installed 2007 February as station P027).

1.4. The APOLLO Contribution

APOLLO—operating at the Apache Point Observatory (APO)—provides a major improvement in lunar ranging capability. The combination of a 3.5 m aperture and 1.1” median image quality near zenith translates to a high photon return rate.

Using a 90 ps full-width at half-maximum (FWHM) Nd:YAG laser operating at 20 Hz and 115 mJ pulse⁻¹, APOLLO obtains photon return rates around one photon per pulse, so that the requisite number of photons for 1 mm normal points may be collected on few-minute timescales. To date, the best performance has been approximately 8900 return photons from the Apollo 15 array in a period of 8 minutes. The average photon return rate for this period is about 0.9 photons per shot, with peak rates of 2.5 photons per pulse. Approximately 80% of these photons arrived in multi-photon bundles, the largest containing 10 photons. APOLLO brings LLR solidly into the multi-photon regime for the first time.

This paper describes the physical implementation of the APOLLO apparatus, including descriptions of the optical and mechanical design, the electronics implementation, and system-level design. For early reports on APOLLO, see Murphy et al. (2000, 2004a, 2003, 2004b). A list of acronyms commonly used in this paper appears in Table 2.

2. SYSTEM OVERVIEW

2.1. Overall Requirements and Differential Timing Scheme

Obtaining accurate laser ranges between a specific point on the earth's surface and a specific point on the lunar surface requires that one be able to establish both the departure and arrival times of the laser pulse, referred to an accurate clock. One must also identify a spatial location on the earth's surface to which the measurements are referenced. For APOLLO, this corresponds to the intersection of azimuth and altitude axes of the telescope, located near the tertiary mirror. Light travels 1 mm in 3.3 ps, so that determining the one-way lunar range to 1 mm precision requires round-trip timing that is accurate to the level of 6.7 ps. These considerations together place stringent demands on the performance of the laser, clock, and timing electronics.

Because the performance of the electronics can be a strong function of temperature, and issues of mechanical flexure and thermal expansion become relevant at the 1 mm level, it is highly desirable to implement a *differential* measurement scheme. As with other LLR stations, APOLLO has, in the exit path of the laser beam, a small corner cube that intercepts a small fraction of the outgoing pulse and directs it back toward the receiver. These *fiducial* photons follow the same optical path as the lunar return photons—including all the same optical elements—but are attenuated by a factor of $\sim 10^{10}$ by thin reflective coatings that are rotated into place by spinning optics. The fiducial return has a photon intensity similar to that of the return from the moon, and is processed by the detector and electronics in exactly the same manner as the lunar return. By adjusting the fiducial rate to be between one and two photons per pulse, it is possible to guarantee that a majority of outgoing pulses have an associated fiducial measurement.

Determining the start time via a single photon from a 90 ps FWHM laser pulse introduces an unnecessary uncertainty in the measurement of the round-trip travel time. A higher signal-to-noise ratio measurement can determine the laser fire time to higher precision (~ 15 ps), though doing so disrupts the differential nature of the measurement. It is, however, possible to accomplish both goals at once: each laser pulse triggers a fast photodiode (see § 5.2.3), producing a robust measurement of the start time to 10–20 ps precision. Though the comparative measurement of the start time as measured by the single-photon detector and by the fast photodiode may vary with time (as a function of temperature, for instance), this variation is expected to be slow. Thus the photodiode measurement can act as a timing “anchor” at a high signal-to-noise ratio for *every shot*, and the single-photon measurement of the start time can be used to “calibrate” the timing anchor over few-minute timescales.

2.2. Chief Components and Layout

Figure 1 provides a schematic view of the APOLLO apparatus. The APOLLO system consists of the following primary subsystems:

1. Laser
2. Optical system, including beam-switching optic
3. Avalanche photodiode (APD) array detector
4. Timing electronics (clock, counters, time-to-digital converter)
5. Environmental monitoring and thermal regulation system

The laser, optical system, detector, and most of the timing electronics are all affixed to the telescope. These components move with the telescope and therefore experience a different gravity vector as the telescope rotates about the elevation axis. The only discernible impact that tilt has on our apparatus is a need to adjust the second-harmonic generator crystal orientation in the laser, which is actuated.

Figure 2 shows the distribution of these pieces of equipment. Four separate umbilical groups connect to the moving telescope. Three of these groups run from the “cabinet” to the laser enclosure. Of these three, one carries laser power and signals (including large cables from the capacitor banks to the flashlamps); one carries signal cables for device communications; and one carries deionized water to the laser heads for cooling. A fourth grouping that is well separated from the laser and signal cables—connecting to the primary mirror cell near the elevation axis—carries DC power and also propylene-glycol coolant for the “Utah” enclosure; more discussion of thermal control may be found in § 6.

3. OPTICAL DESIGN

Figure 3 presents the layout of the APOLLO optical system. The view is rotated $\sim 90^\circ$ clockwise relative to the orientation

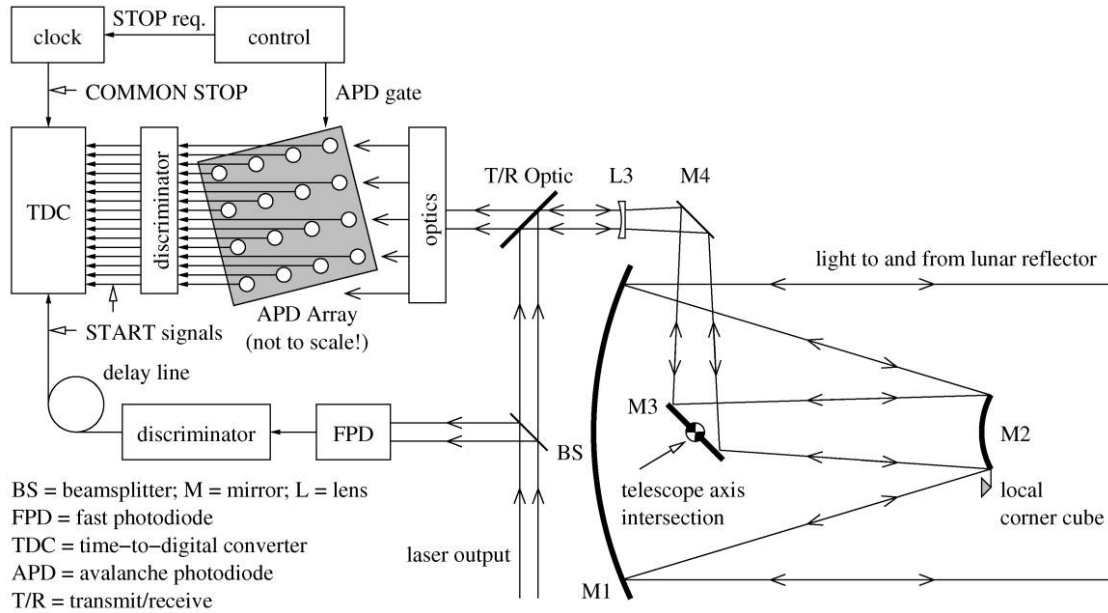


FIG. 1.—Overview of the chief components of the APOLLO system, showing the telescope, local corner cube, detectors, and timing components. Solid arrowheads denote electronic signal paths, whereas open arrowheads denote optical paths. Optics labels follow the convention of the more detailed optical layout in Fig. 3. The box labeled “control” represents the hardware control computer as well as the Apollo command module. The box labeled “clock” consists of the actual clock, clock multiplier, and clock slicer components. The timing system components are discussed in § 5.2.

in Figure 2. The following sections describe the optical components.

3.1. Laser and Laser Power Monitor

The laser is a Leopard solid-state picosecond product from Continuum Lasers. Pumped by flash lamps at 20 Hz, 1.064 μm laser emission from the Nd:YAG rod in the oscillator is shaped into ≈120 ps pulses (FWHM) via the combined efforts of an acousto-optical mode locker, a solid-state saturable absorber, and a GaAs wafer to clamp the pulse energy. A double-pass amplifier boosts the cavity-dumped 1 mJ pulse to ~250 mJ, after which a second-harmonic generator crystal frequency-doubles the light at approximately 50% efficiency to produce 115 mJ pulses of 532 nm light with pulse widths of about 90 ps (FWHM).

Heat from the laser rods is carried away by deionized water flowing at a rate of approximately 6 liters minute⁻¹, taking away ~1300 Watts of thermal power. An auxiliary pump maintains flow in this loop when the laser is not on and the ambient temperature is near or below freezing.

Remote setting of the laser output power is provided by digitally controlled potentiometers placed in the electronics units that control the amplifier flash lamp delay and the voltage applied to the oscillator flash lamp. The latter allows us to monitor the oscillator laser threshold and adjust for optimal laser operation on a routine basis.

A single “output” dichroic (D1 in Fig. 3) sends the green light out of the laser enclosure, past an output shutter actuated and controlled by a custom interlock system. The infrared light passes through to a beam dump. An actuated dichroic (not shown in Fig. 3) located before D1 alternately sends the green beam to a bolometric power meter, so that we may periodically check the laser power. Another fixed dichroic (D2) intercepts residual green light leaking through D1, sending this to two detectors. The first is a photodiode (Hamamatsu S2281; labeled PEPD in Fig. 3) that is used to integrate the pulse energy, presenting the result as an analog output that is read and reset after every pulse. The second is the fast photodiode (FPD) used in the differential timing scheme. The function of this photodiode is twofold: (1) establish a time reference of laser fire to ~15 ps precision (§ 5.2.3); and (2) alert the timing system that the laser has fired, initiating the data acquisition cycle.

3.2. Optical Train

The optical train (Fig. 3) has a transmit path and a receive path that share the full aperture of the telescope. The system must switch between transmit and receive modes (and back) 20 times each second. This is accomplished by the transmit/receive optic (T/R optic). The T/R optic is a 30 mm-thick, 150 mm-diameter glass disk antireflection (AR) coated (0.08% reflection) for 532 nm light incident at 45° across most of the disk. A small sector on the front of the disk has a multi-layer dielectric coating for 99.90% reflection of S-polarized




FIG. 2.—Schematic Layout of the APOLLO apparatus. The heavily insulated Utah-shaped box on the telescope contains the laser, APD detectors, and timing electronics. The insulated cabinet contains electronics that need to be close to the telescope but not necessarily on the telescope. The large enclosure at bottom houses the high power-load equipment of APOLLO, situated in an air volume separate from that of the telescope so that we may dissipate heat into the surroundings. The GPS clock is also on this level, in its own insulated box.

light at 532 nm. The disk is rotated at about 20 Hz by a stepper-servo motor coupled via belt drive. The angle encoder and index from the motor are processed by the APOLLO command module (§ 5.2.2), which directs the laser to fire based on the position of the optic. The laser pulse—slaved in this way to the T/R optic rotation—is arranged to strike the highly reflective patch on the front of the T/R optic so that it is sent to the telescope (transmit mode). Most of the time, the T/R optic presents a clear path from the telescope aperture to the receiver (receive mode).

3.2.1. Transmit Path

Starting at the laser output, the green pulse emerges as a ~ 7 mm diameter ($1/e^2$) beam with an approximately Gaussian

profile, centered 61 mm off the optical bench. A biconcave AR-coated lens (TL1: -74 mm f.l.) followed by a plano-convex AR-coated lens (TL2: 168 mm f.l.)—both spherical—expand the beam to 16 mm diameter prior to the T/R optic. Prior to 2006 June, we used different lenses that presented a 13 mm beam—and thus underfilled the telescope aperture. Following the T/R optic, the beam encounters a plano-concave lens (L3: -198 mm f.l.) that introduces a roughly $f/10$ divergence to the beam so that it may nearly fill the telescope aperture. The virtual focus of this lens is coincident with the telescope focus. After L3, the beam experiences two 90° turns on M5 and M4—both of which are multilayer dielectric coatings for high-efficiency reflection at 532 nm. After this are the telescope's aluminum-coated tertiary, secondary, and primary mirrors




FIG. 3.—APOLLO optical layout. Components are labeled according to the code table on the left. Beam directions are indicated. The T/R Optic rotates at ~ 20 Hz, and the laser fires when the high-reflectivity dielectric patch on the front surface rotates into position. The rear coating on the T/R Optic, identical to the front coating except for size, provides additional attenuation for both the local corner-cube return and the laser beam dump. The receiver beam dump provides a dark target for the receiver during the laser flash and coupling between the laser beam dump and receiver dump is minimized.

(M3, M2, M1). The beam emerges from the primary mirror collimated to well below 0.5 arcsec, as discussed in § 3.4.

3.2.2. Receiver Path

Light from the telescope is brought toward a focus, following the inverse path of the transmit beam, becoming collimated at L3. From here, the path through the T/R optic experiences two 90° turns on M6 and M7, in the process being elevated to ~ 115 mm off the optical bench so that it may cross the transmit path. M7 is tip-tilt actuated so that the receiver may be aligned relative to the transmit beam direction (§ 3.4). The collimated beam enters the receiver tube via an uncoated glass window, tilted to send the reflected light toward a charge-coupled device (CCD) camera that aids acquisition and alignment. The clear aperture up to this window is maintained to be at least 35 mm so that a 40 arcsec field of view is preserved for the CCD camera. Past this window, the optics are 25 mm in diameter, which is suitable for the very small field of the avalanche photodiode detector array.

A narrow passband filter sits at the front of the receiver tube, with a 2.1 nm FWHM passband centered at 532 nm, and 95% transmission at the center wavelength. Prior to 2007 June, we used a filter having 70% peak transmission and 1.5 nm band-pass. Beyond this, a doublet lens (RL1: 147 mm f.l.) concentrates the collimated beam to a focus, where a pinhole is placed to act as a spatial filter. The $400 \mu\text{m}$ hole corresponds to $3''$ on

the sky. An identical lens (RL2) is placed opposite the pinhole, re-forming the collimated beam. An optional mask in the collimated beam blocks light originating outside the telescope aperture, but also serves an important purpose for the fiducial return (discussed in § 3.2.5). A final lens (RL3: 347 mm f.l.) focuses the light onto the detector at the end of the receiver tube. The receiver tube is closely baffled at 50 mm intervals along its entire length so that scattered light from the laser fire is unlikely to survive a trip to the single-photon-sensitive detector.

3.2.3. Detector Configuration and Microlens Array

The detector (discussed in more detail in § 3.3) is a 4×4 array device with $30 \mu\text{m}$ diameter active elements in a square array on $100 \mu\text{m}$ centers. Thus the fill factor is low, at around 7%. A microlens (lenslet) array is placed in front of the detector so that the full fill factor is recovered. The microlens array is an epoxy replica on a 1 mm-thick glass substrate—each lens having a $500 \mu\text{m}$ focal length and covering a $100 \mu\text{m}$ square. The microlens array occupies the focal plane of RL3, where an image of the far field is formed. Each microlens element, or pixel, spans $0.35''$ of angle on the sky, so that the focal plane is oversampled in typical seeing conditions of $1.0''$. At the detector, pupil images—effectively images of the primary mirror—are formed on each $30 \mu\text{m}$ detector element (Fig. 4). This means that the lunar return photons illuminate an annulus on the detector with a central obstruction due to the secondary




FIG. 4.—Microlens array, showing three field points as dashed, solid, and dotted lines, within a single “pixel” of the lens array. A pupil image is formed at the detector element.

mirror and an outer edge determined by the outer radius of the primary mirror.

3.2.4. Fiducial Attenuation

An important ingredient in performing a differential measurement is to establish identical conditions for the measurement of the lunar photon returns and the fiducial returns from the local corner cube. By identical conditions we include the same optical path, the same detector illumination, the same signal level, and the same electronics.

In order to achieve the same signal level, we must heavily attenuate the light returning from the corner cube. Of the 3×10^{17} photons in the outgoing pulse, about 10^{13} strike the 25.4 mm-diameter corner cube. This return has a diffractive spread that overfills the detector array, and some attenuation is present in the form of narrowband filter throughput and detector efficiency (common to the lunar signal). It is nonetheless necessary to provide about 10 orders of magnitude of additional attenuation for the fiducial return.

Six orders of magnitude of attenuation—three from the front surface coating and three from an identical rear-surface coating—are provided by the T/R optic, which is virtually stationary during the ~ 100 ns round-trip time to the corner cube. The thickness of the T/R optic is determined by the requirement that reflections within the T/R optic do not result in overlap between the primary beam and secondary “ghost” beams. Also, the receiver, when “looking” off the reflective back side of the T/R optic, should not see an overlapping patch of illumination from the laser, since about 3×10^{11} photons (10^{-6} of pulse output) in the outgoing pulse are transmitted through *both* front and rear reflective patches on the T/R optic. Well-engineered beam dumps with 10^5 suppression are placed side by side where the laser beam dumps and the receiver looks to eliminate crosstalk between these paths (see Fig. 3).

The remaining 4 orders of magnitude of necessary fiducial return attenuation are achieved in the receiver tube near the spatial filter pinhole. Here, a series of rotating disks with custom coatings present an alternating pattern of attenuation and clear

path to the fiducial and lunar return photons, respectively. Two disks are placed before the pinhole, and a third disk immediately behind. The choice to place the disks here was dictated by the diffusing function, discussed in § 3.2.5 below. The first disk alternates between an optical density (OD) of 0.0 and 2.0. The second disk alternates between ODs of 0.0 and 0.25–1.75. The range indicates that this optic has an azimuthally varying (ramped) attenuation centered around OD 1.0, but spanning ± 0.75 OD on either side. This ramp allows tuning of the fiducial attenuation by setting the phase of the rotating disks relative to the T/R optic. The final disk in the series alternates between OD 0.0 and 1.15—rounding out the targeted total of OD ~ 4 . The rotation of the disks is accomplished by a stepper motor slaved to the T/R optic motor, so that relative phasing is stable and controllable.

3.2.5. The Fiducial Diffuser

Each of the three disks in the arrangement described above is coated in quadrants and spun at half the speed of the T/R optic (i.e., around 10 Hz). Opposite quadrants are either clear (AR-coated) or attenuating, giving a bow-tie appearance to each disk, seen in the insets in Figure 3. The third disk is different from the first two in that opposite the OD 1.15 attenuator is not an identical attenuator but a quadrant of ground glass to act as a diffuser. The purpose of the diffuser is to spread the fiducial photons so that the detector is illuminated in exactly the same manner as it is by the lunar photons.

As described in § 3.2.3, an image of the primary mirror is formed on each of the detector elements. Because the response time of the detector depends on the position of the incident photon, a true differential measurement requires identical spatial illumination of the detector in both circumstances. But the fiducial photons all come from a small corner cube located near the secondary mirror. The image of this corner cube on the detector would then be small, looking nothing like the illumination pattern from the lunar photons. The diffuser located adjacent to the spatial filter pinhole spreads the fiducial photons out, affecting a uniform illumination of the detector element. The optional mask between RL2 and RL3 (not yet installed) would allow one to impose precisely the same illumination pattern on both lunar and diffused fiducial returns, complete with the central obstruction. See § 3.3 for an estimate of the effect the mask would have.

The spatial dependence of the detector response time means that one sacrifices temporal clarity when diffusing the fiducial photons. Any comparative bias has been removed, but at the expense of a temporal spread. Taking full advantage of the timing system to characterize the laser pulse width and detector response, the quadrant opposite the diffuser is a simple attenuator, so that the illumination of the detector is still confined, and thus suffers less temporal spread. In this way, it is still possible to carry out system diagnoses with high temporal resolution. Additionally, one may directly measure the detector’s temporal bias by comparing the time offset of

each type of fiducial return (diffuser vs. attenuator) against the fast-photodiode timing “anchor” for each pulse. In practice, we find the average offset between the two to be about 10 ps.

3.3. Detector Array

The APOLLO detector is a 4×4 avalanche photodiode array fabricated at Lincoln Lab (Fig. 5) (Strasburg et al. 2002). The elements are $30 \mu\text{m}$ in diameter on $100 \mu\text{m}$ centers in a square pattern. The material is lightly doped p -type silicon, epitaxially grown onto a p^+ substrate. A buried p^+ layer is implanted about $1.0 \mu\text{m}$ deep, acting as the multiplication region. The front surface is heavily doped via diffusion to make an n^+ layer and thus the p - n junction. The breakdown voltage is about 25 V, and a typical bias voltage is about 30 V, placing the device in avalanche Geiger mode. The detector spends most of its time biased just below the breakdown voltage, and biased above breakdown for only a 180 ns window, or gate, around the time of an expected photon arrival. The timing window spans 100 ns within the larger APD window.

The device is front-illuminated, so that most photoelectrons arising from 532 nm photons—with a penetration depth of about $1.0 \mu\text{m}$ —very quickly avalanche without spending time in a slow diffusion state. The depletion region in the overbiased state extends a few microns into the device, meaning that a driving electric field exists in this volume. Photoelectrons created within the depletion region are driven at the saturation velocity of $0.1 \mu\text{m}/\text{ps}$ toward the avalanche region. Photons that happen to penetrate below the depletion region generate photoelectrons that may wander for tens of nanoseconds before encountering the depletion region, after which the avalanche is prompt.

One aspect of the avalanche process that impacts our measurement is the lateral propagation of the avalanche microplasma within the thin, disk-shaped multiplication region. Starting in a spot close to where the photon enters the detector,

the boundary of this microplasma expands in a roughly linear fashion at a speed close to the thermal velocity of electrons in silicon. Since the avalanche current is proportional to the area of the avalanching region, the current grows quadratically in time until one edge of the microplasma encounters an edge of the multiplication region—or active region of the detector element. Ultimately, full avalanche current is reached when the entire multiplication region is participating in the avalanche. Depending on the current level to which the trigger electronics are sensitive, this phenomenon could result in a spatially dependent delay of trigger. If the electronics are capable of triggering at a low threshold (but therefore more susceptible to noise triggers), the avalanche does not have to be very large to trigger, and the delay between photon arrival and electronic trigger will be relatively insensitive to the location of the incident photon until this location approaches the edge of the multiplication region. If, on the other hand, the trigger level is set to a high value such that the majority of the active area must be in avalanche, a photon landing in the center will achieve this state long before a photon impinging near the outer edge—by as much as a few hundred picoseconds (Strasburg 2004). The intermediate case would see time-delay insensitivity for a circular region around the center but a spatial dependence outside of this region. It is for this reason that APOLLO employs a diffuser for the corner-cube fiducial returns, as discussed in § 3.2.5. Calculations indicate a reduction of bias from about 100 ps to about 5 ps when a diffuser is employed for high trigger levels.

Because we set the avalanche trigger at a low level—about 30 mV compared to a ~ 150 mV full avalanche—the spatial dependence of the temporal response is somewhat suppressed. We estimate the effect that the optional mask would have on the timing to be about 2 ps, based on the 10 ps scale of the difference between diffused and nondiffused fiducial timing. It should be noted that while attempts to reduce known sources of bias in the differential measurement are warranted, *static* biases will not compromise APOLLO’s science goals.

Crosstalk has been extensively characterized for these arrays. The energetic avalanche process occasionally gives rise to photon emission, which can propagate through the silicon to be absorbed in the active region of a nearby otherwise quiescent element, causing it to undergo avalanche. A simple model of this phenomenon suggests a $1/r^4$ probability for crosstalk, where r is the separation between elements. This relationship is confirmed by experiment. Nearest neighbors ($100 \mu\text{m}$ separation) experience a steady-state crosstalk rate of 0.001 events per nanosecond while the avalanching element is held in a full avalanche current regime. Therefore, a photon-induced avalanche midway into the ~ 100 ns gate results in a 5% avalanche probability due to crosstalk in each nearest neighbor, down by a factor of 4 in the next-nearest neighbors.

One surprising aspect of the crosstalk phenomenon is that it only slowly reaches a steady-state rate. If we cause element A to avalanche, and sustain the avalanche current for > 100 ns, while




FIG. 5.—The 4×4 APD array, with $30 \mu\text{m}$ elements on $100 \mu\text{m}$ centers.

monitoring element B for a correlated event, we find that the crosstalk rate exponentially approaches its steady-state value with a time constant of about 30 ns. There is no sign of prompt crosstalk avalanche events. Practically, this means that since all of the lunar photons return within 1 ns of each other, the crosstalk phenomenon is not yet operative on this timescale so that we do not have to worry about crosstalk events masquerading as genuine lunar returns.

3.4. Acquisition and Alignment

Closing the laser link to the lunar reflectors is a demanding task, requiring simultaneous satisfaction of six variable parameters. One is the outgoing beam divergence, closely related to telescope focus. One is the timing of APD activation—a ~ 100 ns window that must be precisely positioned in time. Two relate to the telescope pointing, simultaneously affecting both the outgoing beam direction and the direction in which the receiver (APD array) looks. Another two parameters describe the pointing offset between the outgoing beam and the receiver direction. A return will only be strong enough to be readily identified if the well-concentrated beam illuminates the reflector, the receiver is aligned to collect the returning photons, and the APD is turned on at the appropriate time. Figure 6 illustrates five of these six parameters.

The relative alignment between the outgoing beam and the receiver may not be an obviously variable parameter. But one must intentionally point ahead of the lunar reflector—to where it *will* be in 1.25 s—while looking behind the “current” position

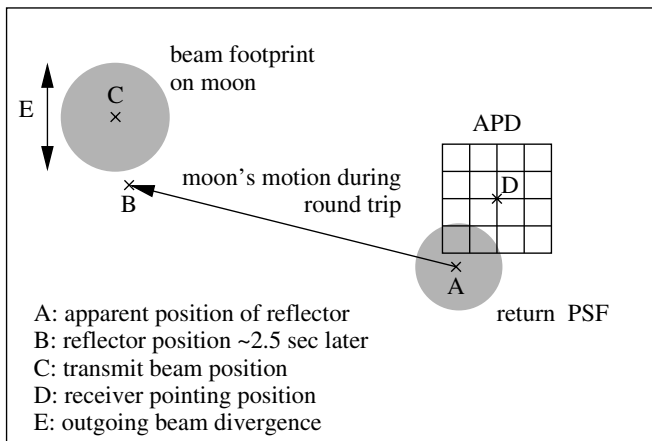


FIG. 6.—Depiction of pointing parameters. The lunar return is seen at position A, where the moon was 1.25 seconds ago, while the outgoing laser beam must point at B, where the moon *will* be in 1.25 seconds. Meanwhile, the receiver (APD) is pointed imperfectly at D, and the laser is striking off-target at C. The size of the return point-spread function (PSF) is determined by atmospheric seeing, while the footprint of the beam on the moon is governed both by seeing and improper focus or collimation. Ideally, D is coincident with A, C is coincident with B, and the lunar footprint is no larger than the seeing limit. The offset between D and C is under hardware control, intended to match the vector from A to B.

of the reflector—to where it *was* 1.25 s ago. At the transverse velocity of the moon ($\sim 1,000$ m s $^{-1}$), this translates to approximately 1.4" of intentional misalignment between transmitter and receiver. But because the telescope mount is driven on altitude and azimuth axes, the offset direction rotates relative to the instrument depending on where the moon is in the sky. Additionally, the earth rotation (~ 400 m s $^{-1}$) changes the magnitude of the necessary offset. Because these effects are comparable in magnitude to the divergence of the beam ($\sim 1''$) and to the field of view of the receiver (1.4"), they must be accommodated in an adjustable manner.

Acquisition and alignment are both greatly aided by the CCD camera (SBIG STV model), which picks off 7% of the incoming light at the front of the receiver tube. Because there are no moving optics between the CCD camera and the APD array, a particular pixel on the CCD may be identified with the center of the APD field. It is therefore straightforward to ensure receiver alignment to a star, or to a crater or other identifiable lunar landmark. A narrowband notch filter in an actuated filter wheel within the CCD prevents green laser light from saturating the CCD so that we may continue to view the target while the laser is flashing.

The CCD is also used to verify the transmitter pointing direction by looking at the return from the fiducial corner cube on the CCD. This is done with the laser turned to very low power, the T/R optic in a transmissive orientation, and an uncoated piece of glass replacing the notch filter within the CCD. The $\sim 4''$ diffraction spot from the corner cube is easily visible on the CCD camera, and its position compared to the known APD position. When the transmit beam and receiver are precisely co-aligned, the corner cube diffraction spot appears centered on the APD position. The relative transmit/receive alignment is adjusted via M7, which is in a tip-tilt actuated mirror mount. The intentional transmit/receive offset angle is likewise confirmed using this system. But because we perform the fiducial alignment check with the T/R optic in a clear position, its phase is rotated relative to the actual transmit phase during ranging. Any misalignment between the surface normal and the rotation axis manifests itself as a wobble in the T/R optic, potentially invalidating the alignment. We can map this wobble in a straightforward manner—just looking at the apparent fiducial return position as a function of T/R optic phase—and correct for it. Indeed, we must offset the beam by 2" when the T/R optic phase is in the nominal clear position to accommodate this misalignment. A second method for verifying transmitter pointing via the CCD works with the laser at full power and the T/R optic in its nominal transmit position. While this method is not sensitive to misalignment of the T/R optic, the dielectric coatings on the T/R optic complicate the transmitted diffraction pattern by preferentially passing P-polarized light, leading to a less distinct—though centered—7-lobe pattern.

Divergence of the beam is checked in two ways. First, a shear plate in a section of collimated beam ensures that the wave front

has a very large radius of curvature (e.g., near planar). This is sensitive at a level corresponding to $0.05''$ of divergence outside the telescope. To the extent that the telescope—when properly focused—delivers a collimated beam after L3, this measurement ensures that the reverse path (collimated beam sent into L3) will deliver a collimated beam at the output of the telescope. Because the CCD and associated lens, and the APD together with lens RL3, have both been set to focus at infinity, the apparent best-focus as determined by either device achieves proper collimation of the output beam. In the second method, a corner cube is inserted at various points around the periphery of the outgoing beam. The position of the corner-cube return on the CCD should not change as a function of which side of the beam the corner cube is placed. This test is sensitive at the $\sim 0.5''$ level—comparable to the precision with which telescope focus is determined in realistic scenarios.

The gate timing is based on predictions generated from the Jet Propulsion Laboratory's DE403 lunar ephemeris. Appropriate corrections for relativistic time delay, atmospheric propagation delay, polar wander, and earth rotation are incorporated. For the duration of the run, the predicted round-trip travel time to each of the four available reflectors is calculated at five minute intervals. A polynomial of the order of 8–12 is computed to fit the predictions to roughly subpicosecond precision. The time of laser fire is latched to submicrosecond precision, and used to compute the polynomial delay. Thus far, we have always seen the return signal arrive within 10 ns of the prediction, and suspect that we can refine our software and site coordinates to achieve 1 ns precision.

If the timing, beam divergence, and transmit/receive offset are properly established, the only remaining uncertainty is the telescope pointing. The APOLLO lunar pointing model is based on the selenographic coordinates of each feature, with libration, parallax, etc. all computed at the time the telescope track command is issued. Thus, once corrective pointing offsets are determined for one feature, any other feature may be acquired with little pointing error if the corrective offset is maintained from one pointing to the next. We can typically make moves across the moon's face without accumulating more than $1''$ of pointing error. Because none of the reflectors are directly visible from the ground, we must rely on local topographical features whose centers are well-defined. Near each reflector (within about $100''$) we have identified reference features that are compact and visible at all solar angles. When acquiring the reflector, we center on the reference crater, then issue the command to track the reflector. We often find the return signal at this nominal pointing position, though it is straightforward to execute a $1''$ spiral raster pattern around the nominal position to pick up the signal. Once the signal is established, we may try to optimize the telescope pointing, the transmit/receive offset, and telescope focus.

4. MECHANICAL CONSIDERATIONS

The laser is built onto a ~ 100 mm-thick honeycomb optical bench, occupying a 610×1220 mm² section of the 915×1220 mm² bench. The T/R optical switch, receiver tube, and CCD camera are all mounted to the remaining section of the bench—a 305×1220 mm² strip beside the laser (Fig. 3). Thus the transmitter and receiver share the same rigid platform, so that their relative alignment (and intentional offset) is preserved at all telescope orientations. Any flexure of the optical bench relative to the telescope would appear as a pointing error in the telescope. The optical bench is affixed to the telescope via a system of six flexures, each stiff in one dimension (for both tension and compression) and flexible in the other two. The flexures are arranged to critically constrain the optical bench relative to the telescope structure in its 6° of translational and rotational freedom. These six flexures contact the telescope at three locations in a triangular pattern, each near structural webs within the steel primary mirror cell frame. The purpose of the flexures is to allow thermal expansion between the thermally stabilized optical bench and the ambient-exposed primary mirror cell. In extreme conditions, this may reach 1 mm of differential contraction. The flexure system accommodates this motion without imposing stresses on the optical bench.

5. ELECTRONICS IMPLEMENTATION

The most important electronics in the APOLLO system are those responsible for the time measurement corresponding to the lunar range. Other systems monitor temperatures, fluid flow rates, laser power, laser pulse energy, telescope tilt, and actuate optics, adjust laser set points, and control power to the various devices in the system. Another system provides interlock control of the laser shutter for aircraft avoidance and in-dome safety. This section concentrates on the timing system implementation, as the rest—while important for the operation of the system—is not critical to the scientific quality of the data.

5.1. Detector Electronics

The 4×4 detector array is packaged in a 40-pin dual-inline package, which is situated in a socket on a multilayer electronics board. Directly behind the APD is an array of variable capacitors whose capacitance (0.4–2.5 pF) roughly matches that expected from the APD elements themselves. Each capacitor and APD element is connected via low-capacitance coaxial cable to an individual, modular “daughter” board inserted into one of 16 sockets on the main board. Figure 7 shows the differential APD triggering electronics as implemented on the daughter boards. The scheme closely follows that presented in Lacaita et al. (1995).

In the “gated off” state, the “gate pulse” voltage is held about one diode-drop below ground, so that diodes D1a and D1b conduct (D2a and D2b are reverse-biased), and voltage applied to the APD cathode and to the “dummy” capacitor is ~ 0.0 V. The

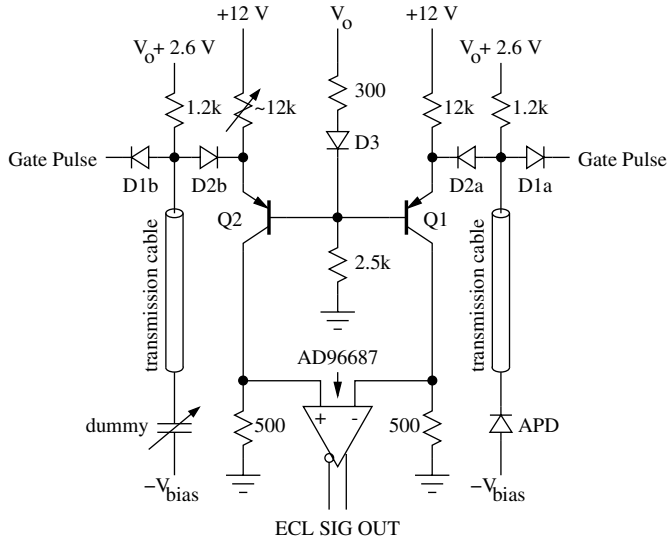


FIG. 7.—APD daughter board electronic readout scheme. The APD is biased with $-V_{\text{bias}} \approx -24$ V, about 1 V short of breakdown. V_0 is typically about 5 V, and controls how much excess voltage is applied to the APD cathode when the gate pulse is brought to a positive voltage (about 7 V). The APD is mirrored by a “dummy” capacitor acting like a nonavalanching APD. A fast comparator senses current draw of the APD when it undergoes avalanche.

APD anode is held at -24 V, which is just below the breakdown voltage of the device. In the “gated on” state, the gate pulse is brought to ~ 7 V, so that diodes D2a and D2b conduct (D1a, D1b are reverse-biased) so that a positive voltage near $V_0 \approx 5$ V is applied to both the APD cathode and the matching capacitor. The total voltage seen by each APD element (and capacitor) exceeds the breakdown voltage, placing the APD into avalanche Geiger mode.

The readout scheme directs current from the APD or capacitor across 500Ω resistors isolated from the APD or capacitor by the transistor. The voltage increase on the sensing resistors associated with turning on the gate is common to both comparator inputs. A single photon entering the APD causes it to avalanche—an action not mimicked by the capacitor—reducing the current through Q1 and triggering the fast comparator (AD96687) to produce an emitter-coupled logic (ECL) output signal. The comparator’s output signal is routed from the daughter board through the main board to the high-resolution time-to-digital converter (§ 5.2.3). Returning to the “gated off” state quenches the avalanche in preparation for the next event.

5.2. Timing System

Figure 8 displays the sequence of events in one cycle of the APOLLO lunar range measurement. The timing system, in brief, is built from the following hierarchical pieces:

1. A high-accuracy clock and associated high-quality frequency standard at 50 MHz;

2. A system of counters to track every clock pulse (thus 20 ns resolution);
3. A 12-bit time-to-digital converter (TDC) with 25 ps resolution and 100 ns range.

The following three sections treat each of these stages. Aside from the clock, the timing electronics are located in a Computer Automated Measurement and Control (CAMAC: IEEE 583 standard) crate situated in the optical bench enclosure. In addition to the crate controller (WIENER PCI-CC32), the inhabitants of the CAMAC crate (Fig. 2) are

1. A custom clock distribution board, referred to as the “clock slicer” (though with no CAMAC connections);
2. A custom counter and state machine board with CAMAC interface;
3. A commercial TDC (Phillips Scientific 7186H; 16 channels, common STOP) with CAMAC interface.

Figure 9 depicts the core elements of the timing system and their means of interconnection.

5.2.1. System Clock

The clock is an XL-DC model from Symmetricom (formerly TrueTime) that is GPS-disciplined to maintain time in accordance with atomic standards to approximately 100 ns. To achieve 1 mm range precision, one must know absolute time to better than one microsecond, as earth rotation and lunar orbit velocities are $\sim 400 \text{ km s}^{-1}$ and $\sim 1000 \text{ km s}^{-1}$, respectively. In the APOLLO implementation, the antenna is about 100 m away from the clock, separated by a single-mode optical fiber with RF converters on either end. A tracking loop with a slow time constant (≈ 2000 s) maintains synchronization of the clock’s internal ovenized quartz oscillator by applying a control voltage. The control voltage is converted from a digital solution to an analog voltage via a digital-to-analog converter (DAC). The DAC value is updated every 10 s and is accessible at all times via computer interface. A single step of the DAC corresponds to a frequency change of 1.5×10^{-11} , translating to about 40 ps (≈ 6 mm) in the 2.5 s round-trip lunar range measurement. This step size is large compared with APOLLO’s target precision. As such, we monitor the clock’s DAC value every 10 s whether actively ranging or not. By fitting a smooth function to the DAC, unphysical steps in the lunar range may be compensated. But even without performing this step, the frequency offset would average away over several-hour timescales.

The direct output of the clock is a 10 MHz 1 V rms sinusoid into 50Ω . The harmonic quality of the sinusoid is very good, with higher harmonics down by at least 30 dBc. The phase noise is also exceptionally low, with a high-frequency noise floor of -154 dBc, translating to a zero-crossing jitter of 2.5 ps. Over a 2.5 s period, the zero-crossing jitter is ~ 5 ps.

The 10 MHz clock signal is frequency-multiplied by 5 to become a 50 MHz sinusoid. The multiplication introduces an

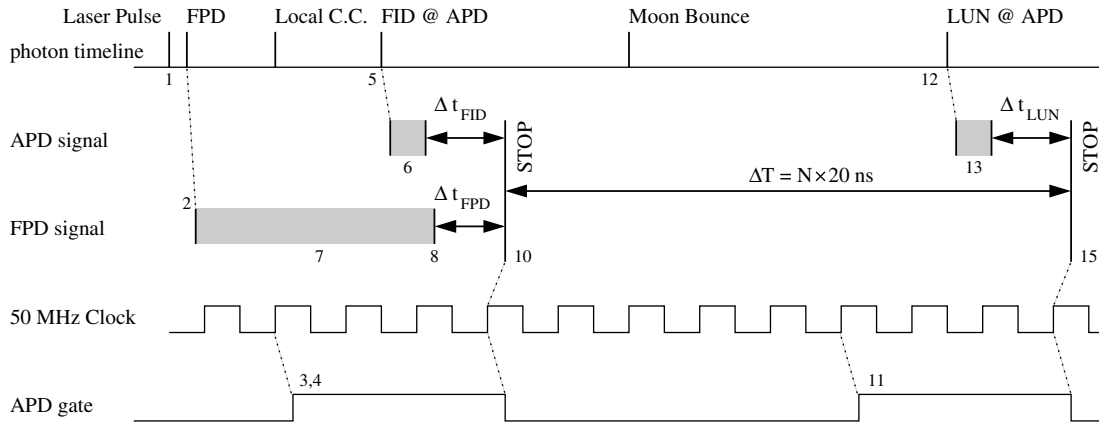


FIG. 8.—Sequence of events in the APOLLO measurement. The photon time line starts with the laser pulse. A small fraction of the light hits the fast photodiode to generate a high-precision measurement of launch time, and another small fraction is intercepted by the local corner cube, to be sent back for detection by the APD as a fiducial (FID) event. The light that bounces off the lunar reflector is detected by the APD as a lunar (LUN) event, the APD being gated on for <200 ns around the time of arrival of each event. Causally related events are indicated by dotted lines, and cable delays are indicated by shaded regions. The time-to-digital converter (TDC) measures the Δt values indicated in the figure to 25 ps resolution, referenced to a STOP pulse sliced from a low-noise 50 MHz clock pulse train. Counters track the number of clock pulses, N (8 in this case), between STOP pulses, measuring ΔT . The fundamental measurement is the round-trip travel time from the local corner cube to the lunar reflector and back, given by $\tau = \Delta T + \Delta t_{\text{FID}} - \Delta t_{\text{LUN}}$. The time axis is not to scale. Numbered events correspond to numbered signals in Fig. 9.

unavoidable $1 + \log 5 = 1.7$ degradation in phase noise so that the edge jitter over 2.5 seconds becomes ~ 8 ps. The clock is housed away from the rest of the timing electronics in a temperature-controlled enclosure. It was not viable to place the clock with the rest of the timing electronics on the telescope due to a pronounced tilt-sensitive clock frequency offset. Initially, the 10 MHz sinusoid was routed directly to the timing electronics crate via an 11 m RG58 coaxial cable. After 2006 July, the $5\times$ multiplier was repackaged to sit next to the clock, now sending a 50 MHz ECL (square) clock signal via a Belden 9207

shielded twisted pair cable to the electronics crate. No obvious degradation accompanied the temporary arrangement, though the decision to route the 50 MHz ECL rather than the 10 MHz sinusoid was largely based on a desire to accomplish greater noise immunity.

5.2.2. Clock Slicer and APOLLO Command Module

The 50 MHz clock signal arrives at the clock slicer board in the CAMAC crate, where it is distributed by a low-jitter, low-skew distribution chip to a number of comparators. The version

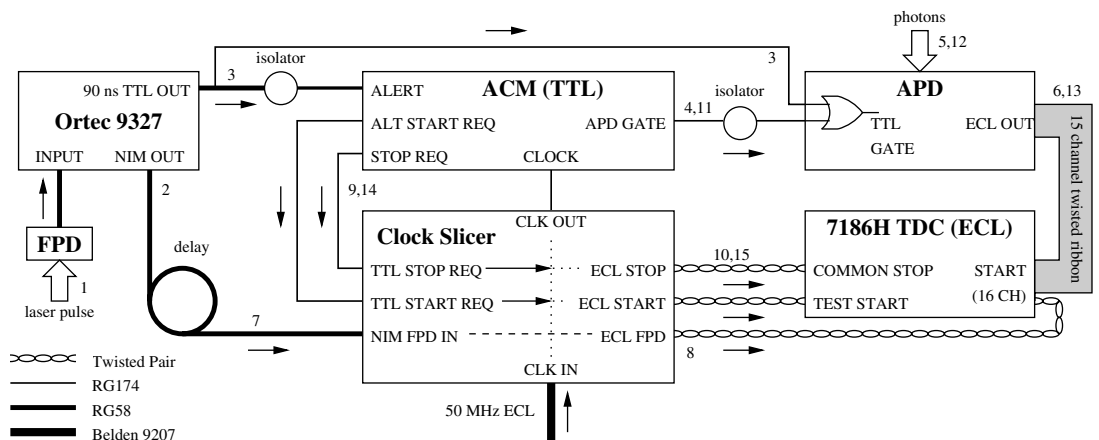


FIG. 9.—Core elements of the timing system. The ACM, TDC, and clock slicer all reside in the CAMAC crate in the “Utah” box on the telescope. The STOP and START pulses out of the clock slicer are derived directly from the high-quality 50 MHz clock source, latch-enabled by the TTL requests from the ACM. One input channel of the TDC is used for the appropriately delayed FPD “timing anchor” pulse. The TTL output of the Ortec 9327 is split so that it may both alert the ACM to a laser pulse and initiate the APD gate. Some signals are ground-isolated by the use of isolation transformers. Numbers relate to the sequence of events for a fiducial-lunar pair. For calibration, the TEST input of the TDC fans out to act as an input for all 16 channels.

of this board used prior to 2006 August also performed the $5\times$ clock multiplication, and distributed the 50 MHz sine wave via a transformer distribution system. This was more susceptible to noise than the current distribution scheme. Among the comparators in the Clock Slicer, one converts the ECL clock into a TTL square wave, which becomes the timing reference for the APOLLO command module (ACM). The ACM, under the direction of the hardware control computer, performs a variety of functions:

1. Counts clock pulses to achieve 20 ns resolution on event times, establishing the time within each second by reference to the 1 pulse-per-second signal from the XL-DC clock
2. Generates STOP (and START) requests for the clock slicer
3. Responds to the fast-photodiode signal from laser fire to initiate fiducial gate (APD turn-on)
4. Coordinates the APD gate event to catch the returning lunar photons
5. Latches counters at key events (APD gate events)
6. Requests latched time from the XL-DC clock to microsecond precision corresponding to gate events
7. Coordinates laser fire with respect to the rotating T/R optic
8. Coordinates the diffuser-attenuator disk rotation and phase relative to the T/R optic

Each gate event produces latched values of various 50 MHz counters, so that, for example, the time within the second of each gate is captured to 20 ns resolution. The latched time actually corresponds to the *end* of the gate event, and this same end-of-gate event is used to generate a STOP pulse in the high-resolution TDC (§ 5.2.3). A free-running 28-bit counter—wrapping each 5.4 s—is also latched at each gate event, and it is upon this counter that the scheduling of the lunar gate is based, in conjunction with the prediction polynomial discussed earlier. The time-within-second and free-running counters provide pulse-count redundancy, which confirm accurate counting with no exceptions to date.

5.2.3. High-Resolution Timing

At the end of each gate event, the ACM sends a 20 ns TTL signal to the adjacent clock slicer board, which unlatches (enables) for exactly one clock cycle a fast comparator (AD96687) with the 50 MHz ECL clock as its input. A single ECL clock pulse is thereby sliced out of the 50 MHz clock train, whose edge is a low-jitter derivative of the clock signal. This pulse serves as a common STOP pulse to the high-resolution TDC. The START pulses are generated by photon events at the APD detector array. The TDC generates a 12-bit digital number corresponding to a time ranging from 10 to 110 ns, corresponding to about 25 ps per resolution element. The actual jitter of the TDC is somewhat better, at <15 ps rms.

A key part of the high-resolution timing system is the fast photodiode (Hamamatsu G4176 with 20 ps rise time coupled to a Picosecond Pulse Labs model 5545 bias tee with 12 ps rise

time) that also serves to alert the system to laser fire. The FPD signal is received by the Ortec 9327 1 GHz amplifier and timing discriminator. The discriminator uses a constant-fraction technique to produce a nuclear instrumentation module (NIM) output having <10 ps jitter relative to the input signal. The combination of FPD, amplifier-discriminator, and TDC have been verified in the laboratory to respond to a short laser pulse with a timing jitter less than 20 ps. The FPD signal therefore provides a high-precision reference to the laser fire time. While we ultimately rely on the individual photon returns from the fiducial corner cube to establish a differential range measurement to the moon, the FPD “anchor” provides a low-jitter indication of laser fire *for every single pulse*. We assume that the time offset between FPD detection and fiducial detections at the APD—which involves cable delays and varying electronics response times and therefore may change with temperature—is relatively constant over the 5–10 minute timescales of a ranging run.

5.2.4. Calibration of the TDC

A possible nondifferential aspect in APOLLO’s ranging scheme stems from nonlinear properties of the TDC. Some function relates digital output of the TDC to the actual START-STOP time difference. Unless the TDC is used in precisely the same range for the fiducial and lunar gates—a condition that is deliberately arranged via our gate timing—imperfect knowledge of this function will contribute to a systematic error in the way that event timing is calculated. We know that the linear gain (e.g., digital counts per nanosecond) varies with temperature by approximately 160 parts per million (ppm) per °C. If the times measured by the TDC for fiducial and lunar events is allowed to be different by as much as 50 ns, 160 ppm per °C translates into 8 ps per °C, which is ~ 1 mm of one-way range.

The ACM, together with the clock slicer, is capable of establishing a calibration of the TDC. In much the same way that the ECL STOP pulse is created by the clock slicer at the request of the ACM, the ACM may also request an ECL START pulse from an earlier 50 MHz clock transition. The low jitter of the clock signal therefore guarantees that the separation of the START and STOP pulses is an integer multiple of 20 ns to ~ 10 ps precision. In this way, START/STOP pairs are created at 20, 40, 60, 80, and 100 ns separation. One thousand pairs of each type are sent to the TDC at a rate of 1 kHz, so that the calibration procedure is carried out in about 5 s. The START pulse is applied to the TEST input of the TDC, which distributes the ECL signal simultaneously to all 16 channels of the TDC. The calibration is carried out at the beginning and end of each ranging sequence, bracketing the roughly 5–10 minute runs.

In addition to determining the gain, offset, and approximate nonlinearity of each TDC channel, the calibration procedure provides a means to measure both the clock jitter (at high frequencies) and the TDC jitter. For each START/STOP pair

separation (i.e., 1000 events), each channel will possess an associated mean of the digital delay-number reported by the TDC. For each event pair, one may then compute the offset between the reported delay and the mean delay. If there is a systematic offset (across the 16 channels) from the mean for a given event pair, it may be concluded that the time separation of the externally generated pair was itself systematically off. In this way, one may separate the distribution of external offsets from internal TDC jitter. Typical results indicate an external (clock, multiplier, and clock slicer electronics) jitter of 10 ps and a TDC jitter of 15 ps. The separation cannot be without uncertainty: 16 channels of TDC values at 15 ps jitter may have a random collective offset from the mean of about 4 ps. But compared to 10 ps in a quadrature sense, this is a small influence.

5.2.5. Error Budget

Based on measurements in the laboratory of the timing performance of the individual components comprising the APOLLO system, we derive the rms error budget per photon presented in Table 1. The timing errors are presented as round-trip measurement errors, while the errors expressed in millimeters represent one-way errors at a conversion of 0.15 mm ps^{-1} . The APOLLO system, according to this tabulation, has a system random uncertainty of 93 ps, or about 14 mm. In this case, the typical uncertainty arising from the tilted retro-reflector array dominates the error budget, so that the required photon number is effectively determined by the lunar libration angle, and ranges from about $20^2 = 400$ to about $47^2 \approx 2200$.

In practice, the APOLLO system does not achieve this performance, most likely due to electromagnetic noise present at the time of laser fire. Two tests have supported this statement. First, two fast-photodiode units coupled to Ortec 9327 discriminators were triggered from the same laser pulse, and their times as registered on the TDC unit compared. In the laboratory, using a diode pulse laser, this measurement yields a 20–25 ps rms error in the comparative timing. In the implementation at the telescope, the same test produces a comparative timing with ~55 ps rms variation. Second, we can perform the calibration tests discussed in § 5.2.4 synchronized to the laser fire, such that any impact of laser noise on the clock slicer and TDC units can

be seen by comparing to the performance when the laser is not firing. These tests show an intrinsic TDC jitter going from 15 to 25 ps, and jitter on the external signals (arriving into the TDC) going from < 10 to 30 ps.

In looking at the spread of the fiducial corner-cube measurement—which should come out near the 93 ps estimate of Table 1—we find typical rms errors of 135 and 120 ps with and without the diffuser in place, respectively. These spreads were 180 and 168 ps prior to a laser pulse-width improvement in 2007 September. Because these numbers are considerably larger than the observed degradation of the TDC and fast-photodiode measurements discussed above, we conclude that the electromagnetic interference from the laser has a significant influence on the timing measurement of the APD signals. At present, APOLLO is still usually dominated by the reflector tilt, even though the system is not as precise as intended.

Most of the electromagnetic noise from the laser stems from the Marx bank generating a ~4000 V pulse with a ~1 ns rise time to dump the pulse out of the oscillator. We are in the process of designing a shield for the Marx bank and associated Pockels cell such that the light signal that triggers the switch is fed into the shield via optical fiber, and only heavily filtered DC power lines penetrate the shield. Initial laboratory tests indicate that substantial shielding is possible, so that APOLLO’s intrinsic per-photon error budget may approach the design goal in the near future.

6. THERMAL CONTROL AND MONITORING

Despite the differential measurement mode employed by APOLLO, thermal control is important. The differential measurement technique assumes that the system performance has not changed during the 2.5 s round-trip travel time of the pulse. But to take maximum advantage of the timing “anchor” discussed in § 2.1, we want the system to be stable for longer periods of time—over several minutes. If we tune the fiducial return to be about one photon per pulse, and allow a single-photon measurement uncertainty of 100 ps, reaching a 5 ps goal requires 400 photons in each channel. This translates to 6000 photons across the array, and thus about 6000 shots, taking 5 minutes.

The chief systematic thermal vulnerability in APOLLO is the variation of gain (picoseconds per bin) in the TDC as a function of temperature. At roughly 160 ppm per °C, a measurement in the middle of the TDC range (50 ns) translates to 8 ps per °C, which corresponds to nearly 1 mm of one-way range per °C. One solution is to arrange the timing of the lunar gate such that the TDC measurements span the same range for both fiducial and lunar photons. Doing this to even 5 ns precision reduces thermal coupling to the level of 0.5 ps per °C—a tolerable level.

All the same, it is desirable to regulate the temperature of the apparatus. Among other things, this promotes stability in the operation of the laser. Our goal, then, is to regulate the thermal environment at the ~1 °C level. We have deployed surface-

TABLE 1
APOLLO RANDOM ERROR BUDGET PER PHOTON.

Error Source	rms Error (ps)	rms Error (mm)
APD illumination	60	9
APD intrinsic	<50	<7.5
Laser pulse	45	7
Timing electronics	20	3
GPS clock	7	1
Total APOLLO	93	14
Retroreflector array	100–300	15–45
Total random uncertainty	136–314	20–47

mount resistive temperature devices (RTDs) throughout the APOLLO apparatus to monitor the thermal state. The hardware control computer periodically reads the temperatures and effects thermal controls to maintain temperature. Because the thermal state of the TDC is especially important, we have five RTDs distributed within the TDC, arranged in a vertical column at the positions of the charging capacitors whose charge is proportional to the measured time interval. The gradient within the device allows us to interpolate to intermediate channels for an estimate of temperature channel by channel.

In addition to the desire to maintain a steady thermal state, the observatory requires us keep thermal emission into the telescope enclosure below 50 W to avoid creating local turbulence along the telescope beam path that affects image quality (both for ourselves and for subsequent observers). The result is a thickly insulated laser enclosure employing polyisocyanurate at a thermal conductivity of $0.02 \text{ W m}^{-1} \text{ }^\circ\text{K}^{-1}$ and a thickness of 8.5 cm. This enclosure maintains a temperature difference, ΔT , at a power loss of $\approx 1.5\Delta T \text{ W}$. The enclosure is maintained at 20°C – 21°C . A 100 W heater maintains states of positive ΔT when the system is off, and a closed-cycle chiller coupled to heat exchangers within the enclosure maintains the thermal state when ΔT is negative, and also when the equipment is on ($\sim 250 \text{ W}$ total power).

7. SOFTWARE CONTROL

The overall control scheme consists of a hardware control computer (HCC), an instrument control computer (ICC), and a telescope-user interface (TUI).⁹ The TUI is a highly capable and versatile platform developed for the operation and control of telescopes at APO. The TUI was specifically engineered to be augmented for control of various observatory instruments, and APOLLO utilized this framework for its interface. The TUI expects to communicate with an ICC, but we did not want to burden the actual HCC with the ICC communications tasks, so placed a separate ICC machine between the TUI and the HCC. The ICC is then free to perform additional tasks including web service and video streaming without compromising the near real-time performance of the HCC.

7.1. Hardware Control

The HCC is implemented as a Pentium III machine running the Redhat 7.2 distribution of Linux. This machine hosts CAMAC, GPIB, and RS-232 serial interfaces, and incorporates a National Instruments 6031E 64-channel 16-bit analog-to-digital converter and 8-channel digital I/O card. The CAMAC interface is responsible for communication with the TDC and the ACM. The GPIB interface communicates with the XL-DC GPS-disciplined clock and the New Focus 8731 optics actuation controller. Serial interfaces are used to command the

laser electronics rack, control the T/R optic drive motor, read the laser power meter, and interface to the SBIG STV CCD controller. Analog inputs are used to read the RTD temperatures, flow meters and flow alarms, laser pulse energy, and telescope tilt-meter. The eight digital outputs are used to activate a configurable set of power relays for turning equipment on and off. A terminal server provides additional RS-232 interfaces to the chillers, a few programmable resistors for remote actuation of potentiometer “knobs,” and also provides control of additional power relays for device activation.

Because the HCC has access to all temperature and flow data, as well as command of the power states of the various APOLLO devices, the HCC controls the temperature of the system. This allows a “smart” approach to thermal control: upper and lower set points; programmable hysteresis, and anticipatory action when powering up for a run. Also running in the background is a check on the XL-DC oscillator statistics—updated every 10 s. Tracking this information allows a reconstruction of the discrete steps in frequency introduced to keep the average oscillator rate in agreement with the GPS time reference. All such background data appends a log that is archived and renewed daily.

The primary function of the HCC is to coordinate ranging activity. Various operational states are defined in the HCC control software. For example, in the RUN state, the HCC commands the T/R mirror to spin-up, performs the TDC calibration, then configures the ACM to flash the laser according to the T/R motor’s encoder/index pulses. Once the laser is firing, CAMAC interrupts alert the HCC that new data is available, at which point the TDC and ACM counter values are read, the pulse energy is read, and the gate time for the associated lunar return is queued. On a CAMAC interrupt associated with a lunar event, the TDC and ACM values are recorded, and the next queued gate time is deployed to the ACM. The primary timing events thus are read at a 40 Hz rate. Slower and less time-critical activities—such as serial and GPIB communications to devices—are threaded so that they do not block the primary activity. At the end of the RUN sequence—usually a fixed number of laser shots—the HCC performs a final TDC calibration.

Other states are defined, including those to warm up or cool down the apparatus, perform in-dome ranging simulations, stare at a star with the APD, obtain dark or flat fields from the APD, calibrate the TDC both with and without the laser firing, and measure the average laser power. Each state creates a data file containing a summary of the device configuration and all associated data, including background environmental data that also populate the log file.

7.2. User Interface

The HCC permits users to log in via a local telnet session so that HCC commands may be entered directly. In addition, the ICC may establish a similar connection, passing information

⁹ See R. Owen’s documentation for TUI at http://www.apo.nmsu.edu/35m_operations/TUI/.

from the HCC to the TUI session(s), and also passing commands from the TUI to the HCC. Multiple TUI sessions may be active at any given time, allowing APOLLO participants to control or monitor the full instrument activity from any location having a TCP/IP connection. A Python-based APOLLO extension to TUI provides a graphical interface to the APOLLO apparatus, with tabbed control over basic HCC operation, the STV CCD camera, the laser, device power, and a lunar pointer tool.

The APOLLO TUI extension employs HippoDraw to graph incoming data in real time, including histograms, stripcharts, and APD spatial response. The HippoDraw plots become our

primary feedback for signal acquisition and optimization. The user interface, augmented by streaming video from the CCD camera, permits operation of the apparatus at any remote location with high-speed Internet access.

8. PERFORMANCE SUMMARY

In the summer of 2005, small amounts of telescope time were made available for system engineering and shakedown, during which we made several attempts to detect a return signal. The instrument was completed with the installation of the microlens array in 2005 October, just before the first scheduled telescope

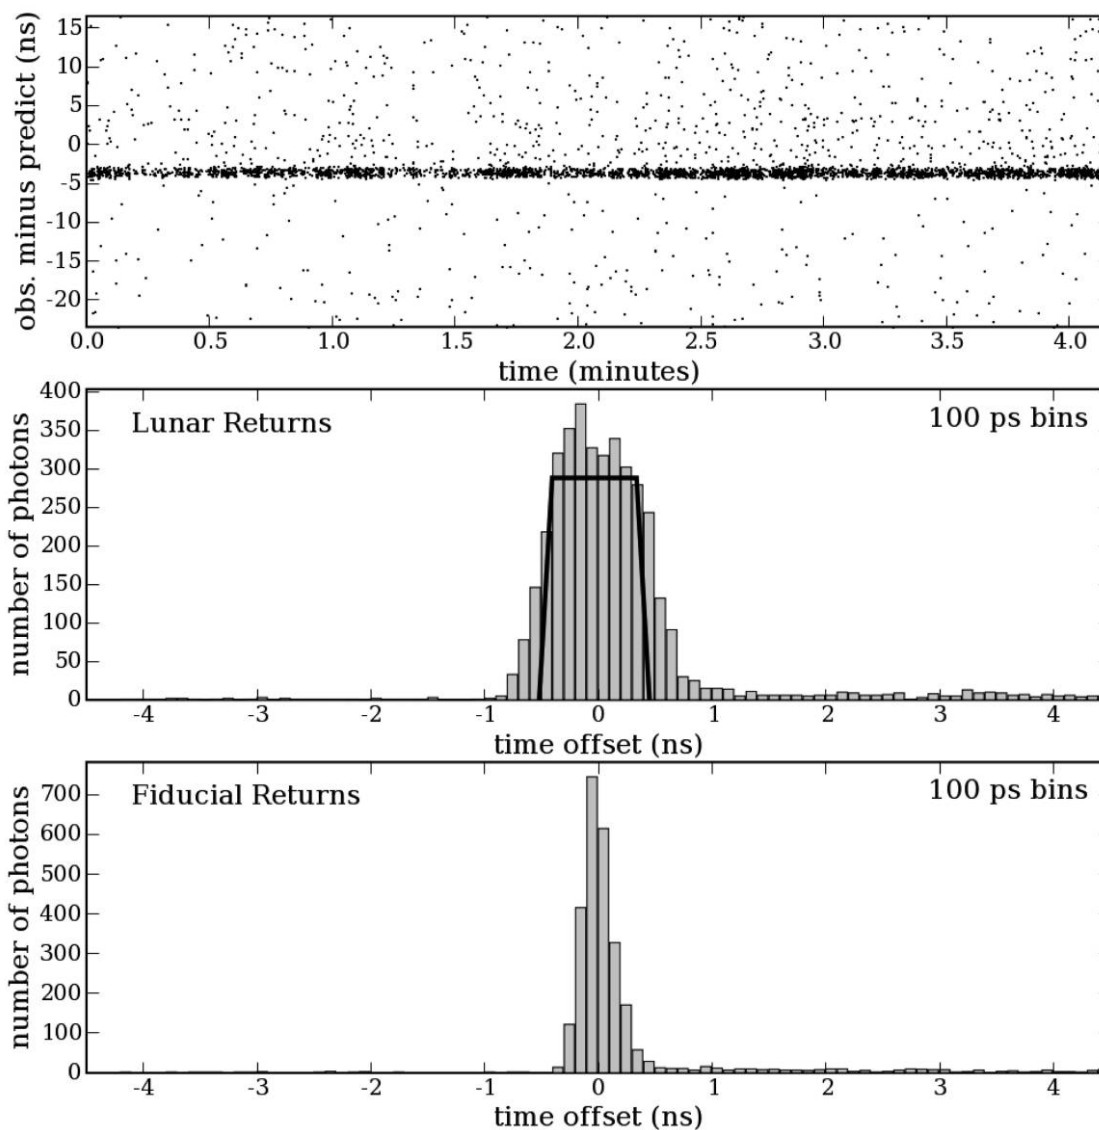


FIG. 10.—Example return from the Apollo 15 array on 2007 November 19, in which 3733 photons were collected in a 5000 shot run. The reflector return is clearly visible against the background. The asymmetry in the background—seen in both lunar and fiducial histograms—results from diffusion of some photoelectrons in the APD. The trapezoidal overplot represents the libration-induced temporal spread on this particular night, which is convolved with the system response (*lower panel*) to produce the lunar distribution. The estimated 1σ error in determining the lunar range from this run is 1.1 mm.

time for APOLLO. In this first observation period in October, we achieved record returns—the best of which garnered ~675 photons in 5000 shots (250 s) for a 0.135 photon-per-shot average. In subsequent months, we saw sustained rates of 0.25 photons per shot, occasionally peaking (for 15 s periods) at 0.6 photons per shot. After the narrowband filter replacement in 2007 July, a laser improvement in 2007 September, and careful alignment in 2007 October, we began seeing sustained rates exceeding one photon per shot, peaking at rates around 2.5. These numbers refer to the Apollo 15 array, which is 3 times larger than the Apollo 11 and Apollo 14 arrays. Figure 10 shows an example return from the Apollo 15 array.

Comparison of the lunar return to the fiducial return in Figure 10 shows that the fiducials are nearly background-free relative to the lunar return, both share the asymmetric tail, and the fiducial return is markedly narrower than the lunar return. The difference in background is due to the substantial attenuation employed for the fiducial return, blocking any background light from the moon. The asymmetric tail is due to photoelectrons created beyond the depletion region in the Geiger-mode APD, discussed in § 3.3. The difference in temporal width is a result of the tilted lunar reflector array, the prediction for which is represented by the trapezoidal curve. Using a correlation technique to compare the timing between fiducial and lunar returns results in a 1.1 mm rms range uncertainty for this run.

The return photon count rate has a high degree of shot-to-shot variability due to speckle structure and other “seeing” effects imposed on the outgoing and return beams. For example, in a 10,000-shot run on Apollo 15 where no changes to telescope pointing or device configuration were made, 1239 single-photon events were recorded, 319 two-photon events, 79, 29, 8, 4, 1, and 1 events with 3, 4, 5, 6, 7, and 8 photons, respectively, for a total of 2303 photons. Thus, 46% of the returning photons were in multiple-photon bundles. Strict binomial statistics at an average rate of 0.23 photons per pulse would suggest that <20% of the returns would be in multiple-photon bundles, with no events having more than four photons in a 10,000-shot run.

APOLLO is highly sensitive to atmospheric seeing conditions. Not only does the flux placed on the lunar reflector scale as the seeing-limited divergence angle squared, but the 1.4” receiver also begins to lose flux when the seeing is significantly worse than 1 arcsec. For large values of the seeing parameter, the total system response scales as the seeing parameter to the fourth power. Thus 3” seeing is 81 times harder than 1” seeing. Operationally, we find it too difficult to work when the seeing is worse than about 2.5”. Zenith seeing at APO exceeds this less than 10% of the time, though at high zenith angles where we must often work, the fraction of unusable time is somewhat higher.

We have not yet seen returns at the rate anticipated by a careful link analysis, which we defer to a future paper. However,

APOLLO's observed return rates provide sufficient statistics for 1 mm range precision on timescales less than 10 minutes.

We are indebted to many people for the success of the APOLLO project. Ed Turner and Suzanne Hawley—successive directors of the Astrophysical Research Consortium (ARC)—have generously allocated director’s discretionary observing time to the project on a continuing basis, and have also facilitated observatory support for our installation and operation procedures. Brian Aull, Bernie Kosicki, Richard Marino, and Robert Reich of MIT Lincoln Lab contributed in a crucial manner by allowing us to characterize and use their APD array technology for APOLLO. Bruce Gillespie of the Apache Point Observatory (APO) facilitated APOLLO’s interface with the observatory and pioneered our interface with Space Command, the Federal Aviation Administration, and local military authorities. Mark Klaene of APO was instrumental in coordinating site activities and suggesting approaches to APOLLO instrumentation that would achieve both project and observatory goals. Other APO staff, especially Jon Brinkmann, Jon Davis, Craig Loomis, Fritz Stauffer, and Dave Woods, have been helpful in getting APOLLO off the ground. Russell Owen wrote the ICC software, for which we are entirely grateful. Russet McMillan has performed many of the APOLLO observations. Jeff Morgan assisted with early project definition and telescope interface issues. Sterling Fisher helped explore and characterize early design ideas. Many undergraduate students have assisted our development: chiefly Jesse Angle, who developed the T/R motor control and Evan Million, who produced our prediction software. Daniel Miller, Justin Ryser, and Aimee Vu also contributed to electronic, software, and hardware development, respectively. Jonathon Driscoll at University of California, San Diego performed some timing experiments and further developed the T/R motor interface. Tim van Wechel at UW mastered

TABLE 2

THE FOLLOWING IS A LIST OF ACRONYMS FREQUENTLY APPEARING IN THE TEXT.

Term	Definition
ACM	APOLLO command module
APD	avalanche photodiode
APO	Apache Point Observatory
APOLLO	Apache Point Observatory Lunar Laser-ranging Operation
AR	antireflection
CAMAC	Computer Automated Measurement and Control
CCD	charge-coupled device
ECL	emitter-coupled logic
FPD	fast photodiode
GPS	Global Positioning System
LLR	lunar laser ranging
Nd:YAG	neodymium-doped yttrium-aluminum-garnet
RTD	resistive temperature device
TDC	time-to-digital converter
T/R	transmit/receive
TTL	transistor-transistor logic
TUI	telescope-user interface

the implementation of the ACM, built by Allan Myers. Allen White at University of California at San Diego contributed substantially to APOLLO electronics development, with critical help from George Kassabian and Mike Rezin—especially on the APD electronics, clock multiplier, and clock slicer. James MacArthur and William Walker at the Harvard University Instrument Design Lab contributed to electronics development. We also thank Mike Vinton of the University of Washington physics machine shop, and Ken Duff, Tom Maggard, and Dave Malmberg of the Scripps Institute for Oceanography machine shop for considerable help in fabrication and design. Jim

Williams of the Jet Propulsion Laboratory, and Randy Ricklefs, Judit Ries, Pete Shelus, and Jerry Wiant of the University of Texas at Austin all provided valuable advice about the LLR technique, as did Eric Silverberg. Jim Williams and Dale Boggs of the Jet Propulsion Laboratory performed frequent verification of our prediction quality. Finally, we express gratitude for the funding sources that enabled APOLLO. These include NASA NAG8-1756, NASA NNG04GD48G, and the National Science Foundation Gravitational Physics program (PHY-0245061 and PHY0602507), in addition to discretionary funds from C. W. S. and T. W. M.

REFERENCES

- Arkani-Hamed, N., Dimopoulos, S., & Dvali, G. 1998, *Phys. Lett. B*, 429, 263
- Bender, P. L., Currie, D. G., Dicke, R. H., Eckhardt, D. H., Faller, J. E., Kaula, W. M., Mullholland, J. D., Plotkin, H. H., Poultney, S. K., Silverberg, E. C., Wilkinson, D. T., Williams, J. G., & Alley, C. O. 1973, *Science*, 182, 229
- Dickey, J. O., Bender, P. L., Faller, J. E., Newhall, X. X., Ricklefs, R. L., Ries, J. G., Shelus, P. J., Veillet, C., Whipple, A. L., Wiant, J. R., Williams, J. G., & Yoder, C. F. 1994, *Science*, 265, 482
- Dvali, G., Gruzinov, A., & Zaldarriaga, M. 2003, *Phys. Rev. D*, 68, 024012
- Lacaita, A., Cova, S., Samori, C., & Ghioni, M. 1995, *Rev. Sci. Instrum.*, 66, 4289
- Lue, A., & Starkman, G. 2003, *Phys. Rev. D*, 67, 064002
- Marini, J. W., & Murray, C. W. 1973, NASA Technical Report, X-591-73-351
- Mendes, V. B., & Pavlis, E. C. 2004, *Geophys. Res. Lett.*, 31, L14602
- Mendes, V. B., Prates, G., Pavlis, E. C., Pavlis, D. E., & Langley, R. B. 2002, *Geophys. Res. Lett.*, 29, 1414
- McCarthy, D. D., & Petit, G. 2003, IERS Technical Note 32, <http://www.iers.org/>
- Müller, J., Nordtvedt, K., & Vokrouhlicky, D. 1996, *Phys. Rev. D*, 54, R5927
- Murphy, T. W., Adelberger, E. G., Strasburg, J. D., & Stubbs, C. W. 2003, in *Proceedings of the 13th International Laser Ranging Workshop*, ed. R. Noomen, S. Klosko, C. Noll, & M. Pearlman (NASA CP-2003-212248; Washington, DC: NASA)
- Murphy, T. W., Adelberger, E. G., Strasburg, J. D., Stubbs, C. W., & Nordtvedt, K. 2004a, *Nucl. Phys. B—Proc. Suppl.*, 134, 155
- Murphy, T. W., Nordtvedt, K., & Turyshev, S. G. 2007, *Phys. Rev. Lett.*, 98, 071102
- Murphy, T. W., et al. 2000, in *Proceedings of the 12th International Workshop on Laser Ranging*, ed. G. Bianco, V. Lucer (Matera, Italy), cddis.gsfc.nasa.gov/lw12/index.html
- Murphy, T. W., et al. 2004b, *Proceeding of the 14th International Workshop on Laser Ranging*, ed. J. Garate, J. M. Davila, C. Noll, & M. Pearlman (Boletin ROA No. 5/2005; San Fernando, Spain: Ministerio de Defensa), 165
- Nordtvedt, K. 1987, *ApJ*, 320, 871
- Nordtvedt, K. 2001, *Classical Quantum Gravity*, 18, L133
- Perlmutter, S., et al. 1999, *ApJ*, 517, 565
- Riess, A., et al. 1998, *AJ*, 116, 1009
- Samain, E., et al. 1998, *A&AS*, 130, 235
- Shelus, P. J. 1985, *IEEE Trans. Geosci. Remote Sens.*, GE-23, 385
- Strasburg, J. D. 2004, Ph.D. dissertation, University of Washington
- Strasburg, J. D., Murphy, T. W., Stubbs, C. W., Adelberger, E. G., Miller, D. W., & Angle, J. I. 2002, *Proc. SPIE*, 4836, 387
- Williams, J. G., Newhall, X. X., & Dickey, J. O. 1996, *Phys. Rev. D*, 53, 6730
- Williams, J. G., Turyshev, S. G., & Boggs, D. H. 2004, *Phys. Rev. Lett.*, 93, 261101
- Williams, J. G., Turyshev, S. G., & Boggs, D. H. 2005, preprint (gr-qc/0507083)