

An absolute calibration system for millimeter-accuracy APOLLO measurements

E G Adelberger¹, J B R Battat² , K J Birkmeier³,
N R Colmenares⁴, R Davis⁴, C D Hoyle⁵, L R Huang²,
R J McMillan⁶, T W Murphy Jr⁴ , E Schlerman², C Skrobol³,
C W Stubbs⁷ and A Zach³

¹ Center for Experimental Nuclear Physics and Astrophysics, Box 354290, University of Washington, Seattle, WA 98195-4290, United States of America

² Department of Physics, Wellesley College, 106 Central St, Wellesley, MA 02481, United States of America

³ TOPTICA Photonics AG, Lochhamer Schlag 19, 82166 Graefelfing/Munich, Germany

⁴ Center for Astrophysics and Space Sciences, University of California, San Diego, 9500 Gilman Drive, La Jolla, CA 92093-0424, United States of America

⁵ Department of Physics and Astronomy, Humboldt State University, One Harpst St, Arcata, CA 95521-8299, United States of America

⁶ Apache Point Observatory, 2001 Apache Point Rd, Sunspot, NM 88349-0059, United States of America

⁷ Department of Physics, Harvard University, 17 Oxford St, Cambridge, MA 02318, United States of America

E-mail: tmurphy@physics.ucsd.edu and jbattat@wellesley.edu

Received 23 June 2017, revised 21 September 2017

Accepted for publication 23 October 2017

Published 22 November 2017



CrossMark

Abstract

Lunar laser ranging provides a number of leading experimental tests of gravitation—important in our quest to unify general relativity and the standard model of physics. The apache point observatory lunar laser-ranging operation (APOLLO) has for years achieved median range precision at the ~ 2 mm level. Yet residuals in model-measurement comparisons are an order-of-magnitude larger, raising the question of whether the ranging data are not nearly as *accurate* as they are precise, or if the models are incomplete or ill-conditioned. This paper describes a new absolute calibration system (ACS) intended both as a tool for exposing and eliminating sources of systematic error, and also as a means to directly calibrate ranging data *in situ*. The system consists of a high-repetition-rate (80 MHz) laser emitting short (< 10 ps) pulses that are locked to a cesium clock. In essence, the ACS delivers photons to the APOLLO detector at exquisitely well-defined time intervals as

a ‘truth’ input against which APOLLO’s timing performance may be judged and corrected. Preliminary analysis indicates no inaccuracies in APOLLO data beyond the ~ 3 mm level, suggesting that historical APOLLO data are of high quality and motivating continued work on model capabilities. The ACS provides the means to deliver APOLLO data both accurate and precise below the 2 mm level.

Keywords: lunar laser ranging, experimental gravity, calibration, frequency standard, SESAM laser

(Some figures may appear in colour only in the online journal)

1. Introduction

Lunar laser ranging (LLR) has long produced superlative tests of general relativity (GR), using the solar system as a dynamical laboratory against which the theory—or parameterized variants thereof—may be tested [1–3]. The technique currently provides our best tests of the strong equivalence principle (SEP), time-rate-of-change of the gravitational constant, geodetic precession, gravitomagnetism, and the inverse-square law, among others. A recent review article summarizes the science case for LLR [4].

For scale, post-Newtonian effects on the Earth–Moon separation, as evaluated in the solar system barycenter (SSB) frame, appear at the 10 m level. Likewise, a complete violation of the SEP (e.g. if Earth’s gravitational self-energy did not itself experience gravitational acceleration) would result in a 13 m amplitude range signal [5, 6].

Early LLR efforts achieved ~ 200 mm measurement uncertainty [2], thereby constituting tests of GR effects at the few-percent level. The period from 1985–2005 saw an improvement in measurement uncertainty to the ~ 20 mm level, permitting $\sim 0.1\%$ checks on relativistic gravity. Beginning in 2006, the next level of LLR precision was achieved by the apache point observatory lunar laser-ranging operation (APOLLO [7]), working down to the ~ 2 mm precision level [8, 9]. Since coming online, APOLLO has produced the majority of LLR measurements worldwide, and also habitually acquires ranges to 4–5 reflectors in each $\lesssim 1$ h observing session. In principle, this significant improvement in measurement precision should push tests of relativistic gravity by a similar factor, down to the $\sim 0.01\%$ level in short order.

However, LLR science relies on an intricate and specialized model capable of reproducing all physical effects that can influence the measurement from a telescope on Earth’s surface to a reflector placed on the lunar surface. Few such models exist in the world [10–14]. The basic scheme performs a numerical integration of the solar system from a set of initial conditions (parameters in the model), following a relativistic equation of motion that may also be parameterized to explore deviations from GR—applied to the entire solar system in a fully self-consistent manner. Torques from non-spherical mass distributions are included, and a host of ancillary effects are layered on top, such as tidal distortion, Earth orientation, plate motion, and crustal loading. The difference between measured ranges and computed ranges form ‘residuals’, which are iteratively minimized in a fully-covariant least-squares process by simultaneously adjusting the unknown model parameters (sometimes within known bounds).

If measurement uncertainties are properly assessed, and the model contains all requisite physics (and is correctly coded), then the residuals should distribute around zero according to the uncertainties such that a reduced chi-squared measure would come out near unity. Another way to express this is that the weighted standard deviation (root-mean-square, or RMS) of the residuals should be commensurate with the measurement uncertainty.

This is not the case for any of the current LLR models. Residuals for APOLLO's ~ 2 mm scale measurements are characterized by a spread ranging from about 15 mm in the best model to about twice this in others [9, 15–17]. No obvious similar patterns emerge when comparing residuals from different models. Thus, introduction of LLR data an order-of-magnitude more precise than was previously available has not yet significantly moved the needle in testing gravity.

The situation raises questions such as: Are the data truly improved? Are the models capable of performing at the millimeter level? Has the technique run into fundamental modeling limitations? Does the past data collection effort of APOLLO represent ‘money in the bank’, establishing a baseline of high-quality measurements while awaiting ultimately successful model improvements—or does it indicate that resources were expended based on imagined gains?

Efforts to improve models are ongoing, but years have passed without a clear ruling on whether the large residuals indicate problems on the model side or data side. Evidence is mixed. Performing several measurement cycles around the lunar reflectors in the course of an hour shows a consistency in APOLLO data points commensurate with estimated uncertainties [8] (while often showing separations by reflector as an indication that the model was not getting lunar orientation right), but assessing possible longer-term trends has proven difficult. Meanwhile, measurements of a second corner cube prism at the telescope exit aperture—in addition to the ever-present fiducial corner cube prism—indicates measurement offsets that depended on axial location, but at least in a seemingly static way across months and years. We suspect that the local corner cube (fiducial) measurement may be impacted by the ring-down of the laser fire, which involves switching ~ 3000 V in a matter of a few nanoseconds and thus creating electromagnetic interference (EMI) that can impact the nearby detector and timing electronics.

In light of these questions and concerns, we devised a calibration scheme that allows us to independently assess APOLLO's data accuracy. The initial idea centered on injecting very short pulses onto the detector at well-controlled intervals as a way to investigate APOLLO's timing performance and track down sources of systematic error. The concept evolved into a high-repetition-rate fiber laser locked to a cesium clock, whose pulses could be selected at will and delivered to the APOLLO system by a long fiber, thus isolating the system from the APOLLO-generated EMI.

Most powerfully, we can overlay these calibration photons atop the lunar range measurements, acting as a calibrated ‘optical ruler’ delivering photon tick marks to the detector alongside lunar return photons. Thus, we are able to calibrate the measurements directly, even if we never manage to find and eliminate sources of systematic error. In so doing, we transform APOLLO from a few-millimeter *precision* apparatus to a few-millimeter *accurate* operation.

This paper describes the absolute calibration system (ACS), and what we have learned thus far about APOLLO data quality. The conclusion is that we see unappreciated influences at the ~ 3 mm level, but nothing that would account for the $\gtrsim 15$ mm scale of model residuals. The implication is the most beneficial to the scientific community: past APOLLO data are of high quality, so that as soon as model capabilities catch up, this long baseline of few-millimeter-accurate data can be leveraged for improved constraints on gravitational physics.

2. ACS overview and design

The original concept for the ACS involved a pulse-on-demand laser system capable of delivering photons to the APOLLO detector at prescribed times, in pulses that would have a high degree of relative accuracy (in the time interval between pulses). Such a setup would permit

a comparison of APOLLO-determined time intervals to the ‘truth’ data from the ACS pulse pairs. Any offset could be studied in a variety of conditions, including various delays after the main APOLLO laser fires. Equipment tunings, module swaps, shielding, temperature effects, etc could all be quantified in terms of their impacts on APOLLO accuracy.

We had difficulty identifying a pulse-on-demand laser operating at or near the APOLLO wavelength of 532 nm having a pulse width narrower than about 90 ps. Moreover, commercial delay generators that could produce pairs of triggers several seconds apart (simulating the lunar round trip travel time) could not provide better than 50 ps jitter. Finally, these devices were not easily configurable to multiplex interleaved pulse pairs in order to faithfully mimic LLR measurements at high rates. While some of these barriers may have yielded to solutions, the combination of challenges motivated us to explore alternate ideas, culminating in the more powerful ACS design described here.

Fiber-cavity lasers employing a saturable absorber mirror naturally produce picosecond pulses at high repetition rates. We adopted a laser design that generates 1064 nm \sim 10 ps pulses at 80 MHz—and most importantly, is capable of modulating the cavity length in a phase-locked loop referenced to a frequency standard such that the period could be stabilized. Intrinsic stabilization can result in pulse jitter well below one picosecond. For reference, 1 mm of one-way distance translates to 6.7 ps of round-trip time. When paired with a Cs standard, the interval of time Δt between pulses could be made extremely regular. This scheme does not support placement of a pulse at arbitrary time, but the tradeoff is that a pair of pulses separated by several seconds will have a reliable Δt limited only by the clock jitter. For a high-performance Cs standard, this jitter is typically better than 10 ps: roughly an order of magnitude better than the pulse-on-demand timing uncertainty. As a result, we rely much less on statistical averaging to get below the millimeter limit, so that in principle the requisite statistics may be accumulated in a few seconds, and the system timing stability can be readily investigated.

This scheme also lends itself well to the second, potentially far more powerful application of the ACS: we can inject calibration photons during lunar ranging operations, knowing that all pulses are tightly connected to a high-performance clock. Thus we deliver an optical clock to the avalanche photodiode (APD) array detector, overlaying a ‘ruler’ of ‘tick marks’ simultaneous with the lunar range measurement. If the fiducial corner cube measurement is being skewed by electromagnetic noise following the APOLLO laser fire, for instance, the ACS calibration will tell us by how much, effectively in real-time. In this way, it would matter less if we never successfully delivered on the first function of the ACS: finding and eliminating sources of systematic error. Naturally, it is important that we make every attempt to do so, but we can still realize a transition from precision to accuracy (in a more guaranteed manner, in fact) by employing this overlay technique—now standard practice in APOLLO data acquisition.

As a concise system overview, then, APOLLO delivers energetic laser pulses (\sim 100 mJ at 20 Hz) to the Moon, typically detecting \lesssim 1 return photo-electron per pulse. APOLLO turns on the APD twice per cycle: once for the local fiducial corner cube return and once for the expected lunar return. The ACS, meanwhile, delivers select pulses out of the Cs-locked 80 MHz train at \sim 1 photon per pulse directly to the APD (no lunar trip). Only those ACS pulses that coincide with expected fiducial and lunar return photon times are permitted to illuminate the APD.

2.1. ACS implementation

The ACS has three main sub-systems: (a) the ‘ACS Enclosure’ containing the 80 MHz 1064 nm pulsed laser, the Cs clock, and associated control systems, (b) a pulse processing system that selects and conditions specific pulses from the train and transmits them to (c) the

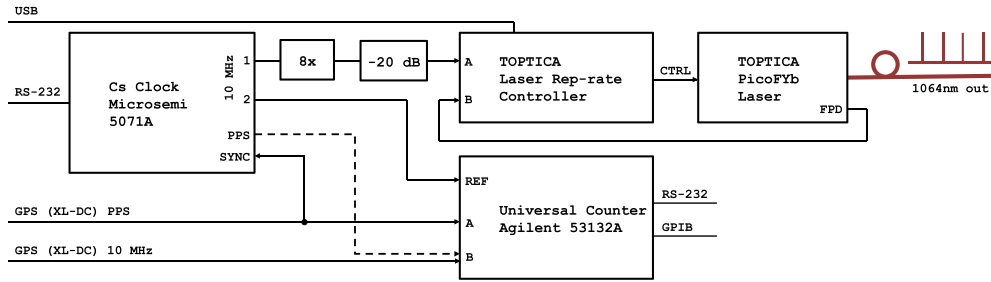


Figure 1. The ACS Enclosure contains a laser (TOPTICA PicoFYb) that delivers a series of 1064 nm 10 ps pulses into an optical fiber. A Laser Repetition-rate Controller (LRC) locks the fiber laser to a Cs frequency standard (Microsemi 5071A). A Universal Counter (UC: Agilent 53132A) enables a comparison between the Cs clock and a separate GPS-disciplined clock (TrueTime XL-DC). The dashed line represents occasional and brief reconfigurations to monitor the accumulating phase difference between the two clocks, as discussed in the text.

Table 1. Allan deviation of the 5071A Cs clock, as provided by the manufacturer.

Delay (s)	10^0	10^1	10^2	10^3	10^4
Allan Dev.	3.1×10^{-12}	2.1×10^{-12}	5.9×10^{-13}	2.0×10^{-13}	6.9×10^{-14}

APD delivery optics. The ACS Enclosure is located away from the main APOLLO system, and is electrically isolated from the rest of the ACS and APOLLO apparatus. The ACS laser pulse train is delivered to the pulse processing system via 12 m of polarization-maintaining single-mode optical fiber. The ACS Enclosure does not move in elevation with the telescope, unlike the APOLLO laser and electronics and the remainder of the ACS system. A pair of Raspberry Pi computers is used to manage the interface between the ACS and the rest of the APOLLO apparatus.

2.1.1. ACS laser and clock overview. The ACS enclosure (figure 1) houses the Cs frequency standard (Microsemi 5071A; stability in table 1), which provides a pair of equivalent 10 MHz 1 V RMS sinusoidal signals having low phase noise (< -155 dBc at high frequency). Phase jitter at the important 2.5 s lunar round-trip time computes to about 2.5 ps. One of the 10 MHz outputs from the Cs clock provides the 80 MHz reference for the ACS laser, after being frequency multiplied by a factor of 8 (Wenzel IFM-3R-10-8-13-13) and attenuated to match the input expectation of the loop controller. The frequency multiplication increases the phase noise by $20 \log 8 \approx 18$ dB, and the phase jitter at 2.5 s periods becomes about 6.3 ps.

The ACS laser is a turn-key system that can be operated remotely. It is a customized TOPTICA PicoFYb that is diode-pumped and uses an all-fiber (Ytterbium-doped) configuration based on a ring oscillator with a semiconductor saturable absorption mirror (SESAM). The laser pulse width at the output of the 12 m optical fiber is ~ 10 ps (full width at half-maximum: FWHM), delivering 30 mW average power within a wavelength envelope of 0.5 nm. Further details about the ACS laser are provided in section 2.2.

The relative frequency and phase difference between the Cs and GPS-disciplined (XL-DC) clocks are monitored by a universal counter (UC, Agilent 53132A). The Cs 10 MHz signal serves as the frequency reference for the UC, and the frequency of the XL-DC clock is measured relative to this reference to 1 μ Hz resolution in ~ 10 s gates. In an alternative configuration (requiring a cable exchange, see dashed line in figure 1), the time difference between the

Pulse-Per-Second (PPS) signals from the two clocks is measured to a resolution of 0.1 ns. The results of this comparison are summarized in section 3.2.

2.1.2. Pulse processing. Injecting calibration photons at critical times requires the ability to select individual pulses out of the 80 MHz pulse train. We accomplished this with a custom pulse processing system (see figure 2). The ‘pulse slicing’ is performed by an electro-optic modulator (EOM), or ‘pulse picker’, which blocks most pulses and passes only the desired ones onward to a temperature-controlled second harmonic generator (SHG). In the SHG, a non-linear optical element converts pairs of 1064 nm photons into single 532 nm photons. Next, the green laser pulses pass through a remote-controlled variable attenuator set to deliver ~ 1 photon per pulse to the APD array. The EOM is controlled by the laser slicer board (LSB), which monitors gate request signals from the APOLLO command module (ACM, not shown), and generates EOM-compatible windowing pulses with user-defined delays and widths.

Using on-board photodiodes, the LSB monitors the average laser power before and after the EOM (1064 nm: PD0 and PD1, respectively), and after the SHG (532 nm: PD2). The EOM contains an internal photodiode whose average photocurrent is also measured by the LSB. Potentiometers tune the sensitivity of each channel to effectively utilize the full range of the 12-bit analog-to-digital converter (ADC: AD5592R).

The LSB generates an 80 MHz clock signal to facilitate synchronized pulse-picking (described below) based on the PicoFYb laser pulses sensed by PD0 (see `Clock Gen.` in figure 2). The clock phase, relative to the optical pulses, is controlled by a PLL-based clock multiplier/divider with a programmable delay line (Texas Instruments CDCF5801A), described in greater detail below.

The EOM (Photline NIR-MX-LN-10-PD-P-P) is a fast (10 GHz) Mach–Zehnder interferometer in which the index of refraction of one arm, and therefore the relative phase of the light in the two interferometer arms, can be controlled by an external voltage. A change in bias voltage of $V_\pi \sim 4$ V will toggle the EOM from destructive to constructive interference, with an extinction ratio of 30 dB (on top of the ~ 5 dB insertion loss). The quadratic behavior of the SHG results in ~ 60 dB switching contrast in green light. Although the EOM could be used as a variable attenuator, we operate it as a two-state device: maximally opaque (‘blocked’) or maximally transmissive (‘open’). The LSB generates a DC bias for the EOM via a DAC (digital-to-analog converter in the AD5592R, followed by an amplifier), tuned to hold the EOM in the blocked state. The blocked-state DC bias voltage is measured experimentally by sweeping the DC bias voltage from 0 to 10V and monitoring the average laser power after the EOM (PD1). Figure 3 (top) shows an example sweep. The value of V_π , the voltage difference between the blocked and open states, can be verified from a DAC sweep. The DAC sweep also allows a measurement of the EOM extinction ratio ($\sim 10^3$) by comparing the average laser power in the blocked and open states (the ratio of the maximum and minimum readings of PD1).

The following paragraphs describe the scheme for clean selection of individual laser pulses. The overall idea starts with a logic signal initiating the request. Programmable delay chips define a user-controlled period during which 80 MHz clock pulses generated from the ACS laser pulse train may trigger a synchronous pulse request. To avoid the transmission of runt or double optical pulses, the gate pulses must bracket the optical laser pulses to ensure that the EOM is opened before the desired laser pulse arrives, and returns to the blocked state well in advance of the next pulse. The amplitude of this RF pulse is set to V_π (manual potentiometer on the LSB) in order to flip the EOM state from blocked to open for one or more pulses.

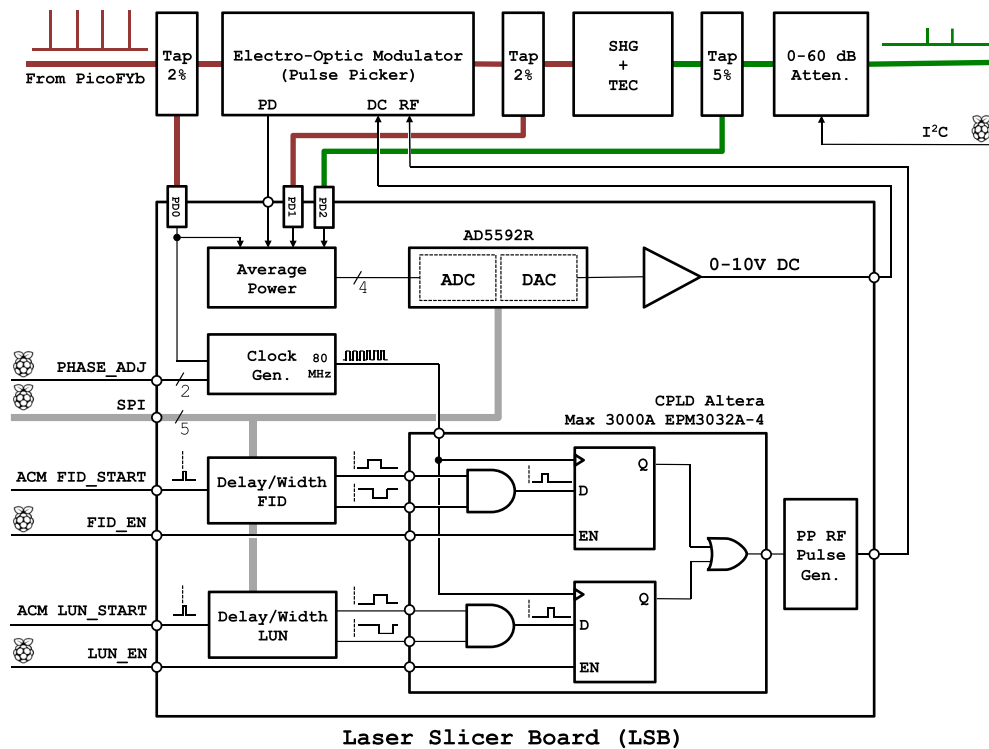


Figure 2. Schematic diagram of the pulse processing system that selects, conditions, and delivers specific ACS laser pulses to the APD injection optics. Thick red and green lines represent optical fibers, while the thin black lines indicate electrical connections. The ACS pulse train is modulated by an electro-optic modulator (EOM), which is normally held in the blocked ('off') state. The laser slicer board (LSB) generates RF gate pulses to make the EOM transmissive for one or more pulses. Details about the LSB and gate pulse generation are provided in the text. A second harmonic generator (SHG, thermally regulated by a thermo-electric controller, TEC) frequency-doubles the light to green (532 nm). A variable attenuator (0–60 dB) allows remote control of the pulse amplitude, generally tuned to deliver ~ 1 ACS photon per pulse to the APD array. Three optical taps enable the LSB to monitor the average laser power along the signal path, and also to generate an 80 MHz clock slaved to the pulse train. Signals labeled with indicate connections to a Raspberry Pi (GPIO or I^2C or SPI).

In more detail, the EOM gates that pick individual pulses are generated in response to a trigger from the ACM, which is part of the APOLLO timing system. Each firing of the main APOLLO laser generates two gated detector events: one for the local (fiducial) corner cube return and another for the eventual lunar return. We refer to these as FID and LUN gates, utilizing logic trigger pulses at the beginning of each gate type called FID_START and LUN_START, as labeled in figure 2. Considering a single gate type for what follows, a pulse-request window is created by first stretching the trigger pulse to ~ 120 ns, then passing to a pair of programmable delay chips (Maxim DS1023-50; 128 ns range in 0.5 ns steps), generating two differently delayed copies of the input pulse. An inverted version of the more-delayed signal is AND-combined with the earlier, non-inverted version to create a composite pulse whose beginning and end times are user-controlled at the 0.5 ns level for precise definition of the EOM window. This combined pulse serves as the input to a D-type flip-flop, so that if the logic signal is high when a positive-going clock edge arrives, a 12.5 ns pulse

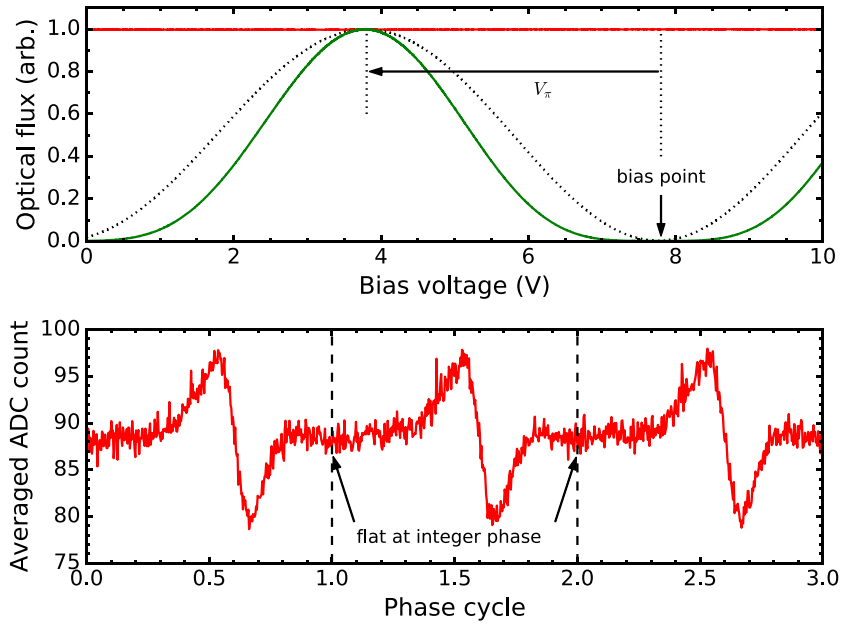


Figure 3. Routine sweeps establish the optimal DC bias of the EOM (top) and phasing so that selection edges are far from optical pulses at the EOM (bottom). At top is a sweep of the DAC output controlling the pulse-picker DC bias. The flat, solid (red) curve at top is the IR input (PD0, steady; unaffected by pulse picking); the black dotted curve is the IR pulse-picker output (PD1); solid (green) is the post-SHG output (PD2), essentially the square of its input. We typically bias the EOM near the 8V minimum and switch by amplitude V_π . At bottom is the PD1 signal during a sweep through three cycles of the LSB-generated 80 MHz clock phase. The hump and trough features arise from having the selection window edges coincident with optical pulses. To select pulses cleanly, we adjust the phase so that integer values find flat regions between the edge features.

request is exported. In this way, the gate request edges are guaranteed to be multiples of 12.5 ns apart, at a constant and controllable phase relative to the laser pulse train. The AND and flip-flop are implemented in a complex programmable logic device (CPLD, Altera Max 3000A EPM3032A-4), also employing an OR-combination of both FID and LUN pulse requests so that separate timing may be defined for each. Considering propagation delays, we are able to generate pulse requests as early as 27 ns after the initial trigger edge (FID_START or LUN_START) emerges from the ACM and as late as 155 ns, in 0.5 ns steps.

Correct phasing of the EOM requests relative to the optical pulses is facilitated by the CDCF5801A chip, which allows bi-directional control of the phase using two GPIO lines from the Raspberry Pi (see PHASE_ADJ in figure 2). The direction of the phase adjustment is determined by the high/low state of one GPIO line, while a phase step is triggered by a rising edge on the second GPIO line. After 1536 rising edges, the clock phase cycles by 360° (0.234° of phase per step). We get confirmation that the EOM-enable gate is properly phased relative to the laser pulse train by sweeping the phase and measuring the average transmitted laser power (PD1) while selecting pulses at a roughly 1 MHz rate. We identify and operate at a flat region in the output, away from fluctuations indicating runt or double pulses when the phase is such that the gate selection edges coincide with laser pulses at the EOM. A sample phase is sweep shown in figure 3 (bottom).




Figure 4. Optical delivery of ACS photons to the APD detector does not block lunar or fiducial photons from the telescope, as ACS photons are situated near the APOLLO receiver optical axis in the shadow of the telescope secondary mirror. Additional attenuation (presently 10^{-3}) at the receiver box entry, in tandem with controlled attenuation (figure 2) establishes single-photon illumination of the APD array. The alternate injection site is discussed in the text.

2.1.3. APD delivery optics. After pulse slicing, frequency doubling to 532 nm, and intensity adjustment, the ACS laser pulses must be delivered to the APD array without blocking the lunar signal photons. Figure 4 shows the optical system that does this. First, the light exits the fiber and a lens creates a free-space collimated beam. Next, a 45° dichroic mirror reflects the 532 nm light into the well-baffled APOLLO receiver box, eliminating $>96\%$ of any residual 1064 nm light that makes it through the 532 nm single mode fiber. A neutral-density filter at the receiver wall contributes to the overall attenuation of the system. The ACS photons are injected into one of two possible collimated sections in the receiver box. The central portion of this collimated section is effectively obstructed by the secondary mirror of the telescope, meaning that a small reflective patch placed in the center can couple the ACS photons into the APD without blocking any APOLLO signal photons. This is achieved by a 45° clear optic (anti-reflection coated for 532 nm), with a small silvered patch in the center.

The alternate injection site—seldom used—allows study of any impact to system timing imposed by additional optics and attenuators (not all shown) within the receiver. In normal operation, rotating attenuator disks near the pinhole impose a $\sim 10^5$ attenuation for FID returns compared to LUN returns, which would produce unequal ACS amplitudes for FID and LUN gates.

2.2. ACS laser

The setup of the TOPTICA PicoFYb laser is depicted in figure 5. A SESAM mode-locked fiber oscillator is the starting point. The repetition rate is controlled and stabilized by adjusting the cavity length with a piezo actuator. The complete fiber oscillator is temperature controlled by Peltier elements.

The laser emits pulses at a central wavelength of $\lambda = 1064.29$ nm, $\Delta\lambda \approx 0.191$ nm (FWHM) and a pulse width of ~ 10 ps—as measured via an autocorrelator using a sech^2 fit. An average power of about 35 mW is coupled out of the oscillator at a repetition rate of 80 MHz.

About 1% of the output power of the oscillator is split and detected by a photodiode in order to generate the electrical signal for the phase-locked loop (PLL) electronic control of the repetition rate (within the LRC in figure 1). The instantaneous phase difference between the Cs clock reference signal and the signal of the photodiode is measured by an analog phase detector. Its output voltage is proportional to the phase difference and is fed to a proportional-integral-derivative (PID) regulator as an error signal. The PID regulator minimizes the error signal by applying an electrical voltage to the piezo actuator and hence varying the optical length of the oscillator by displacing the SESAM along the beam axis. A software-controlled regulation of the oscillator temperature compensates for low-frequency drifts of the repetition




Figure 5. Schematic diagram of the ring-shaped fiber oscillator. See text for details, and [18] for patent information.

rate. The phase detector additionally delivers an RMS phase jitter signal that is used to determine the residual RMS timing jitter value.

A residual RMS timing jitter of 350 fs is measured in closed-loop action when integrated from 10 Hz to 100 kHz. This comparatively high jitter value is considered to be mainly caused by the type of the laser oscillator which has high anomalous net intracavity dispersion in order to operate in the soliton regime and is therefore strongly affected by Gordon-Haus timing jitter. Nevertheless, this oscillator type was chosen for stability reasons and the achieved RMS jitter value is far below the minimum requirement for the ACS of <2 ps.

The pulse picker introduces an optical loss of about 5 dB. Therefore a fundamental power of ~ 11 mW is used to seed the SHG unit. A temperature-stabilized 5 mm thick MgO:PPLN crystal is used for frequency doubling to $\lambda \approx 532$ nm (the spectrum is depicted in figure 6). The frequency-doubled pulses contain an energy of about 4 pJ. The low peak intensity of the fundamental pulse results in a rather low SHG efficiency of 3%, which is more than sufficient to deliver the desired one calibration photon per pulse to the APD array. It was not possible to measure the temporal pulse profile because of low pulse energy. Nevertheless, in theory the quadratic power dependency of the nonlinear process in the non-saturated case causes a reduction of the pulse length after the SHG to a pulse duration of <11 ps. A dichroic filter at the exit of the SHG aperture suppresses infrared transmission, letting only 0.02% through. Roughly half of the green light is coupled into the output PM480 fiber.

The pulse picker introduces an optical loss of about 5 dB. Therefore a fundamental power of ~ 11 mW is used to seed the SHG unit. A temperature-stabilized 5 mm thick MgO:PPLN crystal is used for frequency doubling to $\lambda \approx 532$ nm (the spectrum is depicted in figure 6). The frequency-doubled pulses contain an energy of about 4 pJ. The low peak intensity of the fundamental pulse results in a rather low SHG efficiency of 3%, which is more than sufficient to deliver the desired one calibration photon per pulse to the APD array. It was not possible to measure the temporal pulse profile because of low pulse energy. Nevertheless, in theory the quadratic power dependency of the nonlinear process in the non-saturated case causes a reduction of the pulse length after the SHG to a pulse duration of <11 ps. A dichroic filter at the exit of the SHG aperture suppresses infrared transmission, letting only 0.02% through. Roughly half of the green light is coupled into the output PM480 fiber.

As implemented in the ACS system, we find approximate laser power and pulse energy levels as detailed in table 2. The three taps shown in figure 2 deliver approximately 520, 220, and $9 \mu\text{W}$, respectively, if the EOM is held in its most transmissive (open) state. Individual green pulses emerging from the fiber contain approximately 6 million photons. The variable attenuator is typically set to -16 dB, and a fixed 3 ND filter is placed at the entrance to the optical receiver enclosure. Expanding the beam to uniformly cover the APD array adds another factor




Figure 6. Spectrum of the frequency-doubled pulses at the output of the SHG. Self-phase modulation in the 12 m fiber delivery causes a slight drop at the central wavelength.

of ten of effective attenuation. Together with other losses and detection efficiencies, we end up capturing approximately one photon per pulse across the 16 elements of the APD array. The contrast between green and infrared photons at the APD array is always high ($>10^5$ in the open state and $>10^2$ in the closed state) due to the combination of dichroic elements and fiber coupling/transmission.

3. ACS results

3.1. Clock comparison

A GPS-disciplined clock (TrueTime XL-DC model) previously served as the sole time base for APOLLO. The Cs clock recently added as part of the ACS provides superior frequency stability on intermediate timescales and can be used to characterize the GPS clock. Although the XL-DC clock is no longer the frequency standard for APOLLO ranging, it continues to operate as a time standard, referenced to universal time coordinated (UTC). A detailed comparison of the two clocks, along with a method for back-correcting a decade of APOLLO data based on recorded XL-DC clock statistics is described elsewhere [19].

On a roughly weekly schedule, we reconfigure cables into the UC for approximately one minute in order to check the accumulated phase drift between the two clocks. At the 2.8 km altitude of the site, we expect the free-running Cs clock to run fast by 3.0×10^{-13} due to gravitational redshift. Our Cs clock was measured by the manufacturer to have a static frequency offset of -1.3×10^{-13} relative to a hydrogen maser standard, for a net frequency offset from the UTC second of 1.7×10^{-13} . The resulting phase accumulation should be roughly 15 ns per day. Indeed, we measure a ~ 14 ns per day drift between the two clocks. The net frequency offset translates to a lunar ranging measurement error of less than 0.1 mm.

3.2. Example ACS run and products

This section illustrates the gains delivered by the ACS, while providing details on the techniques by which the results are realized. The next section provides a first-look into the longer-term behavior of the ACS calibration measures.

Table 2. Laser average power and pulse energy at various stages.

Stage	Power (mW)	Energy (pJ)
Out of 12 m fiber	30	375
After upstream tap	26	325
After (open) EOM	13	160
After downstream tap	11	140
SHG fiber output	0.18	2.2

We use for demonstration a 10,000-shot (500 second) observation of the Apollo 15 reflector on 2016-09-12—the first ever using the ACS. While this run is not fully representative in some ways (e.g. still using XL-DC clock as frequency standard rather than Cs clock), it is perhaps more instructive in that it illustrates the effects of clock drift and GPS disciplining. During this run, the fiducial gates and lunar gates each saw about 20 000 total photon events (for clarity, a ‘gate’ is a detector activation event associated with laser ‘shots’—a prompt one for the local corner cube and another delayed to catch the lunar return). Of these, we get roughly 11 000 useful ACS photons for each gate type, while about 5000 and 3000 photons contribute to the fiducial and lunar return peaks, respectively. Raw timing data are in the form of 12-bit integer counts (bins) from a time-to-digital converter (TDC) having 16 channels—representing pixels in the 4×4 APD array—and 25 ps resolution over 100 ns. We will focus on the lunar gate in what follows, but the fiducial gate is analogous.

Raw TDC data (corrected only for individual static channel offsets to improve visibility/alignment) appear in the left panel of figure 7, while the right panel shows the same data, but rearranged by subtracting the *predicted* TDC value that a lunar return would have *for the corresponding shot*—thus piling up lunar detections and smearing out ACS pulses. This is important to understand: based on laser fire time, we can predict to better than a nanosecond (40 TDC bins) when returning lunar photons will arrive. Because the return photons are asynchronous with respect to the 50 MHz APOLLO clock, they distribute uniformly over a 20 ns range in the TDC, which measures the time between the return photon (START) and a selected subsequent clock pulse (STOP). But we know *when*, in TDC-space, any lunar photon should appear. Those arriving at the expected time are colored red in figure 7.

Note that the ACS photons (yellow) in figure 7 form a ‘comb’ of stripes 2.5 ns (100 TDC bins) apart in the left panel, overlapping the (red) lunar photons reasonably well (the overlap is not perfect in this *first* on-sky run, but we have since adapted using fine control of ACS pulse positioning—as detailed in section 2.1.2). While ACS pulses are actually integer multiples of 12.5 ns apart, the 80 MHz pulse train has five possible positions relative to the 50 MHz clock pulse that forms the TDC STOP measure (see figure 8). The number of stripes appearing in figure 7 depends on the adjustable width of the pulse selection window as described in section 2.1.2. The drift rate of the XL-DC clock (GPS-disciplined) relative to the ACS clock (Cs) eventually spreads ACS photons across all TDC values. This drift disappears when the Cs clock is used as the APOLLO time base, as both the ACS laser pulses and TDC STOP pulses become synchronous.

We can play many games using identified ACS photons to learn more about APOLLO timing. The games we explore here *all start like this*: pick some shot: e.g. shot 1000. Consider each/any ACS photon identified for that shot in the fiducial gate, for which we record APD channel/pixel and TDC value. Now ask if there are any lunar-gate ACS photons in the *same* channel for the *same* shot number (or offset shot in some studies), roughly 2.5 s later, when the lunar return would arrive. Finally, ask: what time difference would APOLLO report between these two events, using our standard techniques for turning raw TDC measurements into time

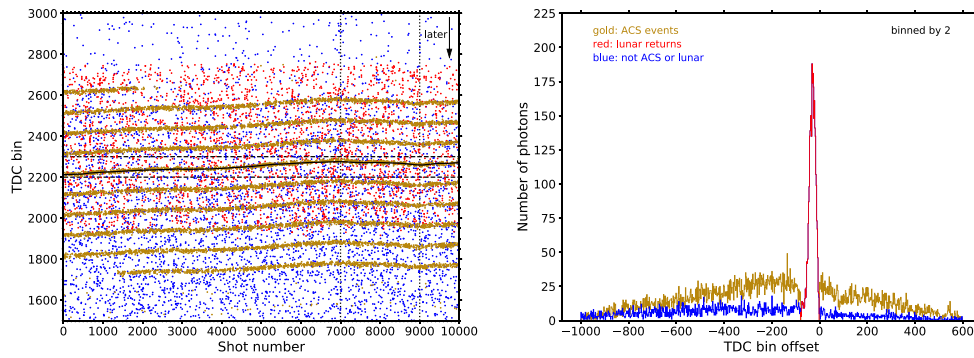


Figure 7. Left: raw data from the first LLR + ACS run obtained over a 500 s period on 2016-09-12. Each dot represents a photon detection during the lunar gate. TDC time measurements are 25 ps per bin; later photons appear lower in the plot. Yellow dots have been identified as ACS pulses based on nearly-static phase relative to the XL-DC clock. Red dots coincide with predicted lunar return times. Blue dots are the remainder—ones that cannot be confidently identified as either lunar or ACS—and largely represent background, slow avalanches due to carrier diffusion, or delayed crosstalk events in the APD (becoming more pronounced later/lower in the figure). The solid black line is constructed from the independent UC measurement of the APOLLO clock frequency referenced to the Cs clock. Horizontal dashed lines help to emphasize the degree to which the two clocks drift, and vertical dotted lines delimit the shot range during which the XL-DC clock steering reversed direction for 100 s. Right: histogram of the lunar-prediction-corrected TDC values, smearing the ACS spikes into a triangular distribution while pulling lunar returns into a high-visibility signal—even though weaker than the aggregate ACS signal and spread over approximately the same TDC region. Color assignments match those in the left panel. The masking by ACS photons in the left panel looks five times worse here than it really is, because only one in five stripes is ‘live’, or possible, for a given shot.




Figure 8. Relative timing between 50 MHz XL-DC clock and 80 MHz ACS laser pulses. Absent frequency shifts between the signals, there are 5 possible alignments of the 80 MHz laser pulse train relative to any given 50 MHz rising edge (then the pattern repeats). The resulting comb of laser spikes has a 2.5 ns spacing.

Table 3. APOLLO range errors are small, and can be determined precisely depending on how many channels and shots are considered.

Channel	δ_{shot}	N_{pairs}	Range error (mm)
16	0	74	-1.70 ± 3.58
16	5	730	-0.13 ± 1.09
1–16	0	826	-0.55 ± 1.01
1–16	5	8918	0.30 ± 0.32
1–16	9	15 234	0.47 ± 0.24

differences? How does this compare to an integer multiple of 12.5000 ns, assuming the ACS photons represent absolute truth in timing? Now repeat this timing error assessment for multiple shots, and possibly multiple channels, aggregating results into a single distribution for statistical analysis. If better statistics are wanted, one may also include shot numbers that are not exact matches, but are within $\pm\delta_{\text{shot}}$ of the target (so that $N_{\text{shot}} = 2\delta_{\text{shot}} + 1$ lunar gates are considered for each fiducial shot/gate).

The more channels we include, and the wider the shot window (δ_{shot}), the more photons we have for comparison. For example, using all 10 000 shots in the example LLR run, APD channel 16 exhibits 74 shots having ACS photons in the same shot for both gate types. Allowing $\delta_{\text{shot}} = 5$ shots (e.g. lunar gates for shots 995 to 1005 compared to the fiducial gate for shot 1000) results in 730 pairs to compare, and much better statistics. Opening up to all channels provides 826 pairs when exact shot/gate matches are required, and 8918 when $\delta_{\text{shot}} = 5$. Some results populate table 3. Note that timing errors are converted to range errors in millimeters in the round-trip sense, wherein 1 ps corresponds to 0.15 mm. Timing and range errors may be used interchangeably, in this sense.

The principal product of the tables in this section is the range error in the final column, obtained thusly. A Gaussian fit is applied to the histogram of offsets produced by accumulating APOLLO range errors reported for ACS photon pairs as described above (see figure 9). The range error is the Gaussian fit centroid, expressed in millimeters. The estimated error of the centroid, μ , is simply the Gaussian standard deviation σ divided by $\sqrt{N_{\text{pairs}}}$. As is apparent in figure 9, the fit is rather good. Reduced chi-squared (χ^2_{ν}) measures for all fits used in the present analysis are statistically probable. The spread seen in ACS time differences in figure 9, at 8 TDC bins (200 ps), indicates an intrinsic single-photon uncertainty down by $\sqrt{2}$ from this, or 140 ps (21 mm). This is consistent with other metrics for APOLLO timing performance (e.g. [7, 9]), and is likely principally due to APD detector jitter. For comparison, reflector tilt contributes typically 200 ps of RMS uncertainty per photon and dominates the error budget. The combined effect requires returns of ~ 1000 photons in order to reach millimeter-levels—which APOLLO can achieve. Before the ACS, we could not be sure how much of this spread was attributable to the APOLLO laser versus APD detection and other timing elements. Since ACS photons are generated in very short pulses, we now know that electronics influences dominate, motivating paths for future improvement.

Several key points should be made about the numbers in table 3. First, the range errors are all tolerably small. This suggests that APOLLO is not guilty of gross inaccuracy in its timing measurement. Second, the numbers are not in statistical tension with each other: relative to the weighted mean of range errors in table 3 of 0.355 mm, the χ^2 value would be larger for 19% of random realizations using these errors. Finally, when enough pairs are included in the comparison (all channels, all shots, $\delta_{\text{shot}} > 0$), a single ACS + LLR run is capable of sub-millimeter accuracy.

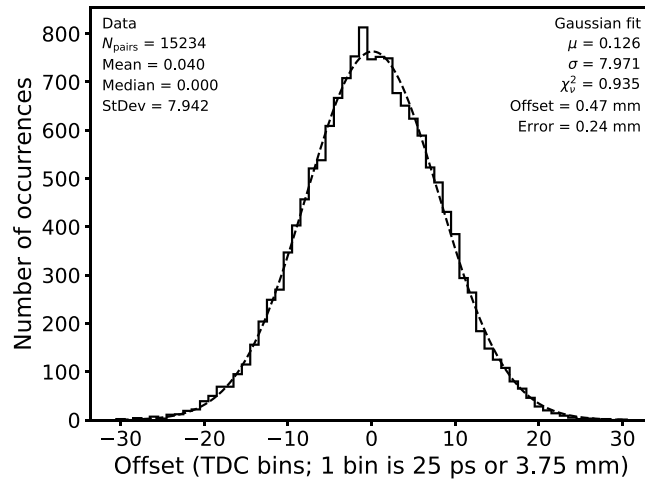


Figure 9. Example histogram corresponding to the final row in table 3. The result is well-described by a Gaussian distribution with mean $\mu = 0.126$ TDC bins (0.47 mm of one-way path), and standard deviation $\sigma = 7.94$ TDC bins. The error, computed from $\sigma/\sqrt{N_{\text{pairs}}}$, is 0.24 mm of one-way path.

In a second test, the ACS accurately captures XL-DC clock drift—a roughly 1.5 mm effect in this instance. In this game, we look at smaller shot ranges and ask if we can discern the clock rate shift in shots 7000–9000 compared to shots 1–7000, obvious in figure 7. We observe the first part of the run to accumulate 60 TDC bins of phase offset in 7000 shots, or 1.5 ns in 350 seconds, or -4.3×10^{-12} fractional frequency offset. Between shots 7000–9000, we see a reversal of 16 TDC bins, translating to a rate of 4×10^{-12} . Across the lunar round-trip time of 2.5 s, we might therefore expect a time measurement error of $-4.3 \times 10^{-12} \cdot 2.5 \cdot 0.15 \text{ mm ps}^{-1} = -1.6 \text{ mm}$ and 1.5 mm, respectively (a 3.1 mm difference). Table 4 presents the results. The first row indicates a negative offset for the first period, different from the latter period (second row) by about 3 mm—perfectly in line with clock drift expectations two sentences back. The two measures are about $4\text{-}\sigma$ discrepant.

A third test reveals that effects other than clock drift are small. The clock drift is due to a frequency offset, producing a range error only when developed over a time interval: e.g. 2.5 s. If we eliminate this time interval by comparing fiducial and lunar gates that are about 50 shots different (2.5 s times 20 shots s^{-1}), they are effectively simultaneous. In other words, fiducial shot 1000 will be compared to contemporaneous lunar return numbered 950 ($\pm \delta_{\text{shot}}$), exposing any range error that persists from influences other than clock drift. Table 5 indicates statistically consistent ‘other’ offsets with a weighted mean of 1.44 ± 0.34 mm. This error scale is in line with previous APOLLO assessments based on comparing the performance of different channels to each other, and may be due in part to EMI from the APOLLO laser firing. Comparing table 4 values to the weighted mean from table 5, we verify that the 1–7000 shot range has a clock-induced offset of -1.75 ± 0.51 mm and the 7000–9000 shot range has an offset of 1.25 ± 0.81 mm: again consistent with expectations from the independent measurement of clock frequency (-1.6 and 1.5 mm).

A final test shows that the ACS can nail down differential channel offsets in ten minutes, compared to months in the past. In a single ACS-LLR overlay run, we get an independent measure of two types of channel-specific timing offsets: large offsets (1 ns scale) common to both fiducial and lunar gates, due to essentially static delays that are different for each

Table 4. Clock drift is easily seen/recovered ($\delta_{\text{shot}} = 5$, all channels).

Shot range	N_{pairs}	Range error (mm)
1–7000	6219	-0.31 ± 0.38
7000–9000	1767	2.69 ± 0.73

Table 5. Neutralizing clock offsets by evaluating essentially simultaneous events, we see a consistent range error from sources other than the clock ($\delta_{\text{shot}} = 5$, all channels).

Shot range	N_{pairs}	Range error (mm)
1–7000	5992	1.65 ± 0.38
7000–9000	1749	0.68 ± 0.72

channel; and differential offset between gate types ($\lesssim 0.1$ ns level). The latter type of offset results in different lunar range measurements for each channel, for instance. We have routinely corrected for these offsets in the past, but adequate characterization required months of lunar data, relying on particularly strong returns in order to have enough ‘truth’ measurements from the Moon. Being able to measure the channel-dependent offsets in a single run opens the door to studying long-term evolution and causes, which will be the subject of future work.

Table 6 illustrates typical values of differential channel offsets from this run—each row representing a different channel. The offsets are sometimes substantial, wholly inconsistent with each other, and have decent precision (using $\delta_{\text{shot}} = 9$ and gate offset of -50 to neutralize clock drift). Within an observing session, we tend to see stable, self-consistent differential channel offsets, and little in the way of worrisome structure across longer timescales. In short, ACS delivers superior channel offset characterization, on *much* shorter timescales.

In summary, these example ACS data indicate that APOLLO inaccuracies are substantially lower than $\gtrsim 15$ mm model residuals; we can characterize inaccuracies at the sub-millimeter level; we no longer have to rely on lunar ranging for truth data, but can produce it when we want; and we can commence an effort to correct our now-characterized few-millimeter-scale errors.

3.3. Range error statistics

Section 3.2 highlighted key analysis schemes available using the ACS. A future paper will describe additional tests and measures, and elucidate the procedures by which ACS data are used in the APOLLO data reduction to produce accurate LLR normal points. For now, we apply the techniques from section 3.2 to all APOLLO operations since 12 September 2016 (through to 2 April 2017) to understand overall performance and stability.

Figure 10 summarizes the results. The top panel corresponds to the measure presented in table 3, using $\delta_{\text{shot}} = 5$ and all channels. The middle panel corresponds to the ‘simultaneous’ measure of table 5—eliminating the effect of clock drift (though not clock jitter). The top panel therefore represents all effects together, and the middle panel is everything but clock drift. The final panel performs a subtraction of the two to characterize the clock drift itself.

At the beginning of 2017, we switched the APOLLO frequency standard to the Cs clock (still using the GPS clock as an absolute time reference). We have therefore split the analysis into two periods corresponding to reliance on the two clocks. Besides reducing phase noise (thus jitter), using the Cs clock as the frequency standard guarantees no drift relative to the ACS laser pulses, given the phase lock to the same clock. Thus the Cs-based data in figure 10 appears to be much tighter.

Table 6. Differential channel offsets (partial set; channel 15 is not used).

Channel	N_{pairs}	Range error (mm)
11	1225	-0.42 ± 0.88
12	1382	-1.30 ± 0.76
13	1403	2.11 ± 0.81
14	998	6.60 ± 0.92
16	1174	3.75 ± 0.81




Figure 10. Channel-aggregated range offsets (0.5 mm bins); Blue/lighter: GPS-based; Red/darker: cesium-based data (darkest is overlap). The top histogram compares ACS fiducial to lunar photons having the same shot number and is an everything#157; measure, containing systematic contributions from timing electronics and clock effects; the middle histogram is the non-clock measure (lunar shot offset by -50 shots comparing effectively simultaneous events) and the last is the difference of the two previous distributions, describing the clock-only contribution to range errors. All cases use $\delta_{\text{shot}} = 5$ Weighted means are represented by dashed (GPS) and solid (Cs) vertical lines. Weighted standard deviations and their computed errors are shown for each. Reduced chi statistics describe how well the weighted mean fits the distribution relative to the offset errors. A factor of four reduction in clock-contributed range errors is apparent in the cesium-only spreads relative to GPS data.

Figure 10 demonstrates that the GPS clock contributes about 2.5 mm of spread in APOLLO range data. Previous estimates of clock-induced error were based on disciplining DAC steps resulting in 1.2×10^{-11} fractional frequency shifts, translating to 4.5 mm in one-way range. The stable DAC behavior suggested a roughly uniform offset distribution bounded by a 4.5 mm range, so that the RMS error would be a factor of $\sqrt{12}$ smaller, or roughly 1.3 mm. The ACS has revealed that the GPS clock drift is about twice what we anticipated—consistent with the clock comparison campaign alluded to in section 3.2 and treated extensively in a separate paper [19]. The reduced χ^2 measure for the Cs clock-drift-only distribution indicates no spread beyond the uncertainties contributed by the ACS statistics (corresponding to errors in table 3, for instance). Neither clock appears to contribute significant range bias (both zero-mean). We do see a several-millimeter offset in the collective statistics (best seen in middle panel of figure 10). We should not take this result literally just yet, as a deeper analysis of ACS data suggests that this is due to temporal structure likely stemming from the APOLLO laser EMI. A more thorough utilization of ACS data can correct for this as well, to be detailed elsewhere. We have also investigated the stability of channel offset measures, as explored in table 6 for the example run, finding the behavior to be stable over month-long periods.

The ACS has provided a powerful check on APOLLO performance. So far, the indications are reassuring. While it is true that historical APOLLO data are more impacted by clock drift than we had appreciated, the level is not severe (2.5 mm), and part of this is correctable [19]. We also see static offsets at the few-millimeter scale, but these may well disappear upon a more complete analysis, as ongoing work suggests.

4. Conclusions

The ACS has transformed the way APOLLO collects LLR data. While offering reassurance that historical APOLLO data appear to be free of large-scale systematic errors at the level of current model residuals ($\gtrsim 15$ mm), the ACS *has* exposed some systematic errors larger than previously appreciated. The GPS-disciplined clock contributed ~ 3 mm of random error—about double the previously-reasoned ~ 1.5 mm scale. Fortunately, the addition of a cesium clock (not to mention the ACS capability as a whole) renders this point moot going forward. A separate paper demonstrates a method by which we can reduce this clock error for historical APOLLO data [19]. We also see a static offset of several millimeters that was not previously measured, deserving further attention. While such an offset is habitually removed in the model via a constant range bias, it is preferable to understand and remove any such influences, especially if the offset is found to evolve over time. While not stressed as much here, the ACS also provides a powerful and fast diagnostic tool to help characterize and eliminate systematic error influences.

A likely outcome of this effort is increased pressure on model improvement, having largely settled the question of APOLLO data quality. In the past, APOLLO data were not as accurate as the estimated uncertainties indicated, although not by enough to explain model residuals. Sharpening up APOLLO data may more clearly expose systematic signatures in the model, suggesting routes to improvement.

In the end, we look forward to realizing the potential gains of millimeter-accurate LLR measurements. Order-of-magnitude gains in our understanding of fundamental gravity are a compelling reward.

Acknowledgments

We thank Ed Leon for performing periodic measurements of clock phase at the observatory. George Kassabian provided a design for producing well-formed RF pulses for the EOM with adjustable amplitude. Jim MacArthur and James Phillips provided assistance with other aspects of the LSB design. This work is based on access to and observations with the Apache Point Observatory 3.5 m telescope, which is owned and operated by the Astrophysical Research Consortium. This work was jointly funded by the National Science Foundation (PHY-1404491) and the National Aeronautics and Space Administration (NNX-15AC51G).

ORCID iDs

J B R Battat  <https://orcid.org/0000-0003-1236-1228>

T W Murphy  <https://orcid.org/0000-0003-1591-6647>

References

- [1] Bender P L *et al* 1973 *Science* **182** 229
- [2] Dickey J O *et al* 1994 *Science* **265** 482
- [3] Williams J G *et al* 1976 *Phys. Rev. Lett.* **36** 551
- [4] Murphy T W 2013 *Rep. Prog. Phys.* **76** 076901
- [5] Nordtvedt K 1968 *Phys. Rev.* **170** 1186
- [6] Damour T and Vokrouhlicky D 1996 *Phys. Rev. D* **53** 4177
- [7] Murphy T W Jr *et al* 2008 *Publ. Astron. Soc. Pac.* **120** 20
- [8] Battat J B R *et al* 2009 *Publ. Astron. Soc. Pac.* **121** 29
- [9] Murphy T W Jr, Adelberger E G, Battat J B R, Hoyle C D, Johnson N H, McMillan R J, Stubbs C W and Swanson H E 2012 *Class. Quantum Grav.* **29** 184005
- [10] Fienga A, Laskar J, Exertier P, Manche H and Gastineau M 2015 *Celest. Mech. Dyn. Astron.* **123** 325
- [11] Müller J, Hofmann F and Biskupek L 2012 *Class. Quantum Grav.* **29** 184006
- [12] Williams J G, Turyshv S G and Boggs D 2012 *Class. Quantum Grav.* **29** 184004
- [13] Pitjeva E V 2013 *Sol. Syst. Res.* **47** 386
- [14] Reasenberg R D, Chandler J F, Colmenares N R, Johnson N H, Murphy T W and Shapiro I I 2017 *Proc., 7th Meeting on CPT and Lorentz Symmetry (CPT 16)* (Bloomington, Indiana, United States, 20–24 June 2016) pp 137–40
- [15] Viswanathan V, Fienga A, Manche H, Courde C, Torre J, Exertier P and Laskar J 2016 *Proc., 20th Int. Workshop on Laser Ranging (Potsdam, Germany, 9–14 October 2016)*
- [16] Hofmann F and Müller J 2016 *Proc., 20th Int. Workshop on Laser Ranging (Potsdam, Germany, 9–14 October 2016)*
- [17] Yagudina E 2016 *Proc., 20th Int. Workshop on Laser Ranging (Potsdam, Germany, 9–14 October 2016)*
- [18] Haering R and Zach A 2013 Fiber laser comprising a ring-shaped resonator *US Patent* 8 457 164 B2
- [19] Liang Y, Murphy T W Jr, Colmenares N R and Battat J B R 2017 *Class. Quantum Grav.* **34** 245009