

Fifteen years of millimeter accuracy lunar laser ranging with APOLLO: data reduction and calibration

**N.R. Colmenares^{1,2,‡}, J.B.R. Battat³, D.P. Gonzales¹,
T.W. Murphy, Jr.¹, S. Sabhlok¹**

¹ Center for Astrophysics and Space Sciences, University of California, San Diego, 9500 Gilman Drive, La Jolla, CA 92093-0424, USA

² Geodesy and Geophysics Lab, NASA Goddard Space Flight Center, 8800 Greenbelt Rd., Greenbelt, MD, 20771, USA

³ Department of Physics, Wellesley College, 106 Central St, Wellesley, MA 02481, USA

E-mail: nicholas.r.colmenares@nasa.gov

Abstract. The Apache Point Lunar Laser-ranging Operation (APOLLO) has been collecting lunar range measurements for 15 years at millimeter accuracy. The median nightly range uncertainty since 2006 is 1.7 mm. A recently added Absolute Calibration System (ACS), providing an independent assessment of APOLLO system accuracy and the capability to correct lunar range data, revealed a $\sim 0.4\%$ systematic error in the calibration of one piece of hardware that has been present for the entire history of APOLLO. Application of ACS-based timing corrections suggests systematic errors are reduced to < 1 mm, such that overall data accuracy and precision are both ~ 1 mm. This paper describes the processing of APOLLO/ACS data that converts photon-by-photon range measurements into the aggregated normal points that are used for physics analyses. Additionally we present methodologies to estimate timing corrections for range data lacking contemporaneous ACS photons, including range data collected prior to installation of the ACS. We also provide access to the full 15-year archive of APOLLO normal points (2006-04-06 to 2020-12-27).

1. Introduction

Modern physics is built upon two primary pillars: general relativity and quantum mechanics, which describe nature at its largest and smallest scales, respectively. The standard model of particle physics, which describes three fundamental forces to a very high accuracy, omits gravity entirely. Of the two pillars comprising modern physics, gravity is less accurately tested than Standard Model forces. Experimental gravitational constraints are challenging to study due to gravity being about 39 orders-of-magnitude weaker than quantum mechanics, when comparing their scaled coupling constants [1]. Very massive objects are necessary for making meaningful measurements of certain aspects of gravitation, and as such, solar system objects have traditionally provided the most practical means of measurement to accomplish this. In our local region, the Earth-Moon system provides a convenient laboratory for probing gravity using lunar laser ranging (LLR). Many leading constraints on gravity (such as the strong equivalence principle[§]) have historically been tested in this natural laboratory, and the goal of modern-day LLR experiments is to improve upon these constraints [4, 5].

LLR began in 1969 after the installation of a retroreflector array (composed of fused silica corner cube reflectors, or CCRs) on the Moon by the Apollo 11 astronauts. Subsequent retroreflector arrays were placed on the Moon by Apollo missions 14 and 15 [6], as well as by remotely landed Soviet rovers Lunokhod 1 and 2 [7]. In LLR, a pulsed laser is fired from the Earth to one reflector array at a time, and the round trip time (RTT) of photon flight is recorded. Each nanosecond of RTT corresponds to 15 cm in one-way range; routine range precision and accuracy on the millimeter scale is achieved for the Apache Point Observatory Lunar Laser-ranging Operation (APOLLO) [8, 4, 9, 10]. Observations throughout the course of the Moon’s trajectory enable

[§] Recently, improved constraints on the strong equivalence principle were obtained from NASA’s MESSENGER mission data [2], and from a hierarchical stellar triple system consisting of a white dwarf and millisecond pulsar orbiting another white dwarf [3].

determination of the shape of the lunar orbit, allowing for judgement of the validity of opposing gravitational theories. For example, theories must account for ~ 10 m post-Newtonian (relativistic) effects on the Earth-Moon separation, as measured in the solar system barycenter (SSB) frame [4]. Gravitational tests aside, LLR is additionally sensitive to Earth orientation parameters [11], secular evolution of the Earth-Moon distance [12], and the physical properties of the Moon [13, 14]. A more detailed overview of LLR science deliverables can be found in Ref. [4].

APOLLO currently produces the highest-precision LLR measurements. The dominant source of temporal uncertainty for modern LLR stations such as APOLLO is due to reflector geometry being a flat tilted array of many small-aperture corner cubes. Several research groups are investigating next-generation single-element large-aperture fused silica corner cube reflectors ([15, 16, 17, 18, 19]) as well as novel bonding techniques for their high tensile strength and low coefficients of thermal expansion [20, 21, 18]. These next-generation corner cubes will largely address the geometry-induced uncertainty of the original reflector arrays, offering potentially significant improvements to range precision for an equivalent number of collected range photons.

APOLLO fires powerful laser pulses (~ 115 mJ, ~ 90 ps FWHM at 20 Hz, 532 nm) to retroreflector arrays on the Moon. A small portion of the outgoing light is intercepted by a local retroreflector attached to the secondary mirror of the 3.5 m telescope at the Apache Point Observatory (APO). The relative timing of the “fiducial” (FID) photons from the local corner cube and the returning “lunar” (LUN) photons yield a differential measurement, effectively determining the separation of the local and remote corner cubes using the same detector and timing electronics. The construction of a differential measurement is important to circumvent the difficulty in making a precise absolute timing measurement due primarily to the temperature dependence of the timing electronics. After each laser pulse (“shot”) the 4×4 avalanche photodiode (APD) detector array is turned on (in associated detector activation events called “gates”) once for the fiducial returns and once for the lunar returns, typically detecting $\lesssim 1$ lunar photon across the array. A custom Computer-Aided Measurement And Control (CAMAC) device called the APOLLO Command Module (ACM) is the primary state machine for APOLLO, scheduling gates using a series of counters, registers and comparators referenced against a 50 MHz clock train. A more detailed description of APOLLO instrumentation, operation and sequence of events can be found in Ref. [8].

More recently, an Absolute Calibration System (ACS) was added to APOLLO, to independently assess system accuracy and check for systematic errors [9]. The ACS utilizes a low-power semiconductor saturable absorber mirror (SESAM) fiber cavity laser (different from the main “range” laser) whose cavity length is thermally modulated by a phase locked loop (PLL) referenced to a cesium (Cs) time standard in order to stabilize the pulse repetition period. The laser generates 1064 nm light pulses of 10 ps duration at 80 MHz. Individual pulses are picked out of the pulse train to be delivered to the APD during the FID and LUN gates resulting in calibration photons overlaid onto range measurements that form a “ruler” of optical “tick marks.” We

can compare what APOLLO determines the timing separation of these tick marks to be against the known truth (a multiple of the 80 MHz pulse separation of 12.5 ns), given the exceptional regularity with which ACS pulses are delivered out of the laser. To distinguish range photons (non-ACS) from ACS photons, photon timestamps are compared against predicted range photon return times as well as expected ACS pulse timings to result in strongly peaked signals of range returns and ACS detections respectively. Details of the ACS apparatus and performance can be found in Ref [9].

Raw APOLLO data products, or “runs,” are the result of photon collection periods spanning 3–10 minutes (3000–10000 laser shots) on a single lunar reflector. A run typically contains a few hundred lunar range photons, and a single observing session spanning 1–1.5 hours typically contains 5–10 runs; further characterization of APOLLO operation can be found in Ref. [22]. During each run we record information about photon and gate timing, absolute timestamps provided by a GPS-disciplined quartz oscillator timing standard (XL-DC) and relevant system state parameters. APOLLO’s precision timer is a 12-bit 16-channel (one for each APD element) time-to-digital converter (TDC), configured to have 25 ps resolution. The data acquired during a run is processed into a final product called a “normal point.” A normal point provides a representative measurement of the round-trip-travel time from telescope to reflector and back at some specific epoch, and is the fundamental observable used in subsequent LLR-based analyses [23].

Utilization of the ACS allowed us to uncover a $\sim 0.4\%$ systematic error in the TDC calibration, which impacts APOLLO measurements at the millimeter scale; fortunately, we are able to apply timing corrections to account for this effect to even the earliest of APOLLO data. The rest of this paper describes how APOLLO’s calibrated normal points are generated. A full description of APOLLO’s performance using normal point statistics is presented in a companion paper [22].

2. Packaging of Range Photon Timing

To motivate the reduction procedure and application of ACS corrections, we first summarize the sequence of APOLLO events and how a range measurement is formed. It is critical to know to a high degree of precision when the range laser has fired, as scheduling of gates is referenced against this time. However, the firing time of the range laser cannot be controlled to enough precision given its large jitter (1.6 μ s). A timing “anchor” is used to obtain a low-jitter (~ 20 ps) measurement in the form of a fast photodiode (FPD), the signal from which is present for every laser shot. In contrast, the APD operates in single-photon mode, with ~ 100 ps of jitter per photon, and may not have detections for every laser shot. An unfortunate consequence of utilizing the FPD as a timing anchor is that it is not a differential measurement with respect to the returning lunar range photons. The FPD signal is handled on one reserved channel of the TDC that will have a different temperature dependence (and different electronic and optical paths to arrive there) compared to the other 15 TDC channels. The differences

between FPD and LUN signal handling would cause the timing between the two to drift over time; as such, we need to tie our timing “anchor” to the lunar photons on a run by run basis by relating the two to the fiducial photons, which take the same electronic and optical paths as the lunar photons. Relating the FPD and FID signals for each run is a relatively simple task given the presence of one or more photon events in most FID gates, providing ample statistics.

Figure 1 shows a crude timeline of APOLLO events for an idealized scenario of a single laser shot in which a fiducial photon and a corresponding lunar gate return photon are present in the same detector channel. The range laser fires, generating a FPD event a short time later as range laser photons leave the APOLLO apparatus. The telescope axis intersection is chosen as the fixed reference point to measure time between the outgoing and returning signal. This reference point is insensitive to telescope motion, and we know to a high degree of accuracy how long it takes light to travel along the optical axis between the axis intersection and the fiducial corner cube based on measurements using a theodolite.

The FPD signal (henceforth distinguished from the FPD itself by a different font: FPD) alerts the ACM that the laser has fired, thereby initiating the FID gate and scheduling the corresponding LUN gate about 2.5 seconds in the future. A cable-delayed version of the FPD is also fed to the TDC within the fiducial gate. A time T (the prediction time, *if* our prediction were perfect) after the transmitted laser pulse reaches the telescope axis intersection, range photons again cross the same reference point after returning from the Moon. Shortly thereafter the light reaches the APD which produces a **START** signal that is sent to the TDC.

Although the fiducial and lunar photon return times are asynchronous relative to the 50 MHz clock signal used to schedule APOLLO events, we are able to predict when those returns should arrive (based on the FPD), and schedule gates to the nearest clock pulse such that range returns uniformly populate a 20 ns interval in TDC-space, colloquially referred to as a “slug,” for both gate types. For each FID and LUN gate, we record the TDC reading for each channel that triggered, as well as the last clock pulse in the detection gate using a counter that tracks the number of clock pulses elapsed. Fiducial gate records contain information on when the FPD fired, and lunar gate records contain the predicted TDC timestamp for lunar photons based on the FPD timing from the associated fiducial gate.

During the reduction procedure, TDC calibrations obtained from a process called “CALTDC” — executed before and after each run — are applied to the raw data. During CALTDC, the TDC sends common **START** and **STOP** signals (derived from the 50 MHz clock train) to all channels 1000 times at 1 kHz each for 5 different **START–STOP** intervals in integer multiples of 20.00 ns. A quadratic function is fit to each channel’s results to provide a conversion between TDC number and an actual timestamp for any photon event—more information on this process can be found in Ref. [8]. Various data integrity checks are also performed, such as checking that counters are not missing any clock pulses. Below is a simplified description of the remaining reduction procedure, followed

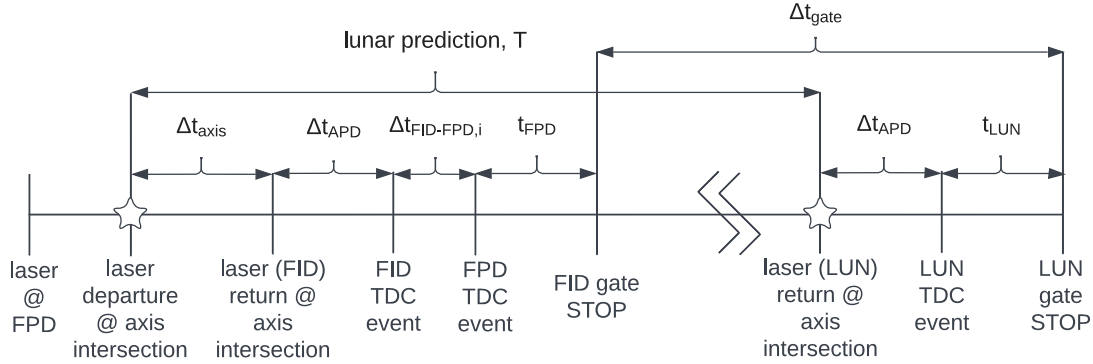


Figure 1: Rough sequence of notable APOLLO events for a particular laser shot fired with one fiducial and one lunar photon returned in their respective gate types, in the same detector channel. $\Delta t_{\text{FID-FPD},i}$ indicates an individual event that contributes to the aggregate fiducial signal $\Delta t_{\text{FID-FPD}}$ seen in equation 1 and discussed in the text. Δt_{APD} denotes the time between laser light returning through the telescope axis intersection and receipt of a photon event at the TDC; note this term cancels out in equation 1.

by a detailed description of each step.

- (i) Subtract the FPD timestamp from each FID event timestamp to produce a relative timing signal for each channel of the APD. Combine these into a single “aggregate” fiducial signal, thereby establishing channel-to-channel timing corrections as well as characterizing range laser and detector timing performance.
- (ii) Fit a physically motivated functional form to the “aggregate” fiducial signal, which is later used to inform the lunar functional fit.
- (iii) For each lunar event timestamp, form lunar “residuals” by applying the channel offset corrections determined in step i to each lunar event and subtracting the prediction TDC value for that laser shot.
- (iv) Remove any linear temporal trend in the lunar residuals across the run duration and form a histogram of the result.
- (v) Perform a one-parameter sliding fit between the lunar and fiducial histograms and record the result. The shape, size and background vertical offset are pre-determined from the FID signal treatment convolved with a trapezoid representing the libration-specific apparent reflector profile.
- (vi) Add back in the linear trend and prediction value at a calculated reference time.

Recall that most shots will not have a photon event in both fiducial and lunar gates for the same APD channel. To utilize all lunar events for our final data product, we form an aggregate fiducial signal to act as a timing reference for every lunar event. This aggregate fiducial signal is constructed by comparing the arrival times of each FID photon against the laser fire time (the FPD). The difference in FID and FPD arrival

times is effectively constant over the duration of a run, but slightly different for each channel (due to differences in electronic delays of each). Relative channel corrections δ_{ch} are determined on a per-run basis by correlating each channel against a reference signal (arbitrarily chosen to be the first TDC channel). We can then apply each channel’s timing correction and combine data from all channels to produce a strongly peaked signal (the “aggregate”). The width of the aggregate fiducial is set by the laser pulse width and detector/timing electronics jitter. It has an asymmetric tail due to photoelectron diffusion into the depletion region of the APD [4]. A function (gaussian with asymmetric tail) is then fit to the aggregate signal. From here, the RTT for an individual lunar photon is calculated from:

$$RTT = \Delta t_{\text{gate}} + t_{\text{FPD}} + \Delta t_{\text{FID-FPD}} - t_{\text{LUN}} - \delta_{\text{ch}} + \Delta t_{\text{axis}}, \quad (1)$$

where:

- Δt_{gate} is the number of clock oscillator periods between the ends of the FID and LUN gates for this laser shot (a multiple of 20.00 ns),
- t_{FPD} is the time of the FPD signal relative to the end of the FID gate,
- $\Delta t_{\text{FID-FPD}}$ is the fitted aggregate fiducial timing,
- t_{LUN} is the lunar photon return time relative to the end of the lunar gate,
- δ_{ch} is the relative channel correction,
- Δt_{axis} is the round trip light travel time from the axis intersection to the local CCR and back.

Lunar photon timing residuals are obtained by subtracting the predicted return time on a shot-by-shot basis from the calculated RTTs (with channel corrections applied), resulting in a strongly peaked signal similar to the aggregate fiducial, as seen in Figure 2. A line is fit to the residuals vs. time within the run to capture the offset and slope. Lunar returns that fall within a windowed region of the fitted line are tagged as “registered” lunar returns (red points in Figure 2, top), and a refined linear fit is performed on these registered photons. A representative time is formed for the run by finding the mean of registered lunar return times. The fitted line is subtracted from the residual signal to result in a new residual histogram. A one parameter fit is performed on the latest residual histogram by sliding the aggregate fiducial signal fit, convolved with the reflector tip/tilt profile at the time of ranging, until best-overlaid. During the fit we impose a constraint that the area under the fit function be matched to the area under a windowed region of the lunar histogram and an additive constant is set to match the measured background. The representative time of the registered lunar photons is used to calculate a predicted round-trip time for the run *as a whole*, and is combined with the centroid of the lunar histogram fit, plus the representative timing offset of the fitted trend line; the ensuing number is the fundamental lunar range measurement: a representative round-trip time for the *entire run*—the aforementioned “normal point.”




Figure 2: Top: lunar gate photon timing residuals versus time within run, for a run on the night of 2019 April 16, after subtracting the predicted round-trip time for each shot. The linear residual is represented by a blue line, and photons within some window surrounding the fit are colored red, indicating they are candidate lunar range photons. Middle: the aggregate fiducial signal with its fitted centroid subtracted. The asymmetric fit function is overlaid in red. Bottom: histogram of the difference between lunar events and the fitted blue line in the upper plot, with the retroreflector array timing profile for this observation shown in magenta. The red line is the fiducial fit convolved with the reflector profile and shifted horizontally to optimally align with the lunar histogram.

2.1. Clock corrections

Prior to installation of the ACS in 2016 August, APOLLO only used the GPS-disciplined Agilent XL-DC clock as its timebase. After the ACS was installed, APOLLO utilized either the newer Symmetricom cesium clock or the XL-DC clock as its timebase. Since the arrival of the cesium clock in 2016 February, a Universal Counter (UC) has monitored the relative frequency and phase differences between the two timebases. A methodology for determining corrections to round trip times informed by the XL-DC clock was developed based on the UC comparison data [24].

Runs that use the XL-DC as APOLLO's timebase without corresponding UC

comparison data are correctable in the manner described in [24], which gives a single RTT adjustment and uncertainty modification per normal point; this encompasses runs before 2016 February 12 and runs from 2017 December 21 to 2018 May 11, when the cesium clock underwent repairs. However, runs that use the XL-DC as APOLLO’s timebase but *do* have corresponding UC comparison data are eligible to receive slightly more direct clock corrections. The concept is effectively the same as in [24], but since we have UC comparison data every 10 seconds, we may apply clock corrections determined for each 10 second interval (using comparison data averaged over the closest three to four 10 second intervals, for smoothness) to individual lunar range photon round trip times instead of a singular correction to the normal point’s final representative round trip time and uncertainty. Runs from 2016 February 12 to 2017 January 4 and from 2017 February 3 to 2017 February 15 (during which we reverted to using the XL-DC while investigating an issue related to the cesium clock) are correctable using this slightly modified version of [24]’s methodology.

2.2. “Long-term” channel corrections

Raw APOLLO normal point uncertainties are based only on data from individual runs, but there may still exist systematic trends from RTT differences between channels. By comparing individual-channel RTTs against the weighted mean of RTTs over all channels for a given run, it is possible to assess how consistent and compatible a particular channel’s timing is relative to the overall collection. A single run will show some scatter among individual-channel RTT calculations that may or may not be statistically aligned to the whole, which becomes easier to discern for runs that acquire a large number of photons.

One run alone is not enough to elucidate systematic trends, so in practice we have analyzed these patterns over long timespans corresponding to distinct hardware periods. We restrict the study to include runs that accumulated at least 500 photons so that individual channels would have enough signal to produce useful statistics. Over the period of time represented in this analysis, we could assess for each channel whether any systematic offset existed, and if the excursions from the channel’s weighted mean exceed that expected from estimated uncertainties. In other words, a chi-squared measure of departures from the channel weighted mean indicates whether the scatter is consistent with or exceeds estimated uncertainties.

We may then use the results of these “long-term” analyses to apply corrections to individual channels equivalent to each channel’s deviation from the weighted mean of the whole; this has the effect of tightening-up the distributions of range returns such that the functional fit mentioned in the previous section performs better. Additionally, we can account for the disconnect between expected and observed scatter by applying a measure of excess spread to the estimated normal point uncertainties. These error inflations are computed in two ways: as a multiplicative scaling of the original errors, or (more appropriately, we think) as a root-sum-square (RSS) error term to be added in

quadrature to the raw normal point uncertainty. Note however that these error inflations are provided to analysts as suggestions; APOLLO normal point uncertainties have not been adjusted by these suggested amounts. Recommended error inflations are discussed in more detail in [22].

3. The ACS signal

A detailed description of the ACS hardware and integration into the rest of the APOLLO apparatus is provided in Ref. [9]. Here, we describe the salient characteristics necessary to understand the ACS-based timing calibration of APOLLO data.

The ACS laser repetition rate is 80 MHz, locked to a stable cesium frequency standard. ACS photons are coupled into the same APD detector array and timing system that is used for the lunar range measurement. Although ACS photon pulses are 12.500 ns apart, the 80 MHz pulse train has five possible positions relative to the 50 MHz clock pulse that forms the STOP signal for the TDC. These five possible positions produce an ACS “comb” in TDC-space, with comb “teeth” separated by $n \times 2.5$ ns, where n is an integer.

Even though the calibration and lunar range (or fiducial) photons may arrive in a single gate of the APD, timing within the gate can be used to tag detected photons as ACS or LUN or FID. An example of this is given in Figure 3. The top panel shows the FID and LUN photons (aggregated over all APD channels), while the middle panel includes the ACS photons. The ACS peaks are broad because timing corrections between different channels of the APD have not been accounted for. Once those corrections are applied, the ACS “comb” is clearly seen (such as in the lower panel of Figure 3). Comparing the measured locations of the peaks to the known intrinsic separation of $n \times 2.5$ ns provides a calibration of the APOLLO timing system.

The quality of the ACS calibration depends on the precision with which the centroid of each tooth is known. That, in turn, depends on the number of photons in the tooth. The ACS light level is limited to $\sim 1.5 \frac{\text{photons}}{\text{gate}}$ (we avoid higher rates because an APD element that has detected an ACS photon would be blind to a subsequent lunar/fiducial return signal in that gate). A typical lunar ranging run contains 5000 range laser shots (gates), spread across 15 active APD channels, and an ACS comb consists of ~ 10 teeth. So each tooth of an ACS comb for a given APD channel contains about 50 of the 7500 ACS photons. The ACS is often disabled during the first $\sim 20\%$ of a run to ensure proper lunar range signal acquisition, bringing the effective ACS shot duration down from 5000 to ~ 4000 , and the per-tooth per-channel ACS photon count down to 40. A typical comb tooth is $\sigma_\gamma \approx 130$ ps wide for data prior to August 2019, and $\sigma_\gamma \approx 80$ ps after (an upgrade to the APD in August 2019 improved the timing precision). The tooth centroid is therefore determined to a precision of $\frac{\sigma_\gamma}{\sqrt{N}} \approx 21$ ps (13 ps) before (after) the hardware upgrade when $N \approx 40$. Range photon timing corrections rely on the difference of tooth positions, which further degrades the calibration precision by a factor of $\sqrt{2}$ to 30 ps (18 ps), or 5 mm (3 mm) of one-way range. Given that the typical APOLLO normal




Figure 3: Histograms of time-converted range and ACS returns for the same run as in Figure 2. The top and middle plots do *not* include relative channel corrections, while the bottom *does*.

point uncertainty is 1–2 mm, and that the typical scale for ACS-based timing corrections of range data is 3–7 mm, we need a more precise method to calibrate APOLLO data.

To enhance the comb tooth centroid determination, one might consider aggregating ACS data taken across many runs, thereby increasing the number of total ACS photons in each comb tooth. Unfortunately, temporal drifts in the comb tooth location in TDC-space make this approach unviable. Instead, we accumulate measurements of the separation *between* comb teeth across runs. As shown in Figure 4, these separations are stable over timescales of years. Reduced chi-squared measures suggest the presence of extra structure and/or underestimated uncertainties (as seen for the FID result in Figure 4); subsequent investigations revealed no extra underlying structure, so uncertainties were inflated to result in reduced chi-squared measures near unity. As described in Section 4, we leverage this stability to arrive at an ACS-derived timing calibration with an accuracy of ~ 0.5 mm.

It is additionally worth noting that the weighted means of tooth separations for both gates in Figure 4 lie below the expected separation of 2.500 ns by approximately 0.4% (~ 10 ps), demonstrating that the TDC tends to undercount the time interval between this pair of teeth on average. This undercounting of time is a common feature across channels and tooth pairs, providing a clue that APOLLO may systematically count time too slowly in the TDC. This phenomenon will be seen repeatedly in successive sections

of this paper.




Figure 4: Tooth separations between an adjacent pair of teeth, for a single APD channel, over several years. Each gate type (FID or LUN) is plotted separately. The χ^2_ν , corresponding p-values*100 and the degrees of freedom (DOF) are shown for each gate, relative to the weighted mean which is displayed as colored solid horizontal lines. The horizontal black dash-dotted line shows the expected 2.500 ns tooth separation if APOLLO counted time perfectly. The weighted means for each gate lie approximately 10 ps below 2.500 ns, or $\sim 0.4\%$, which is common across channels and tooth pairs and indicates an undercounting of time in the TDC – a feature which will be discussed in subsequent sections of this paper. The two large gaps in the data (days 259 to 454 and 467 to 570) are due to a reference clock failure that halted ACS operations, and a scheduled observatory maintenance shutdown, respectively.

Any APOLLO run either has sufficient contemporaneous ACS data or it does not. In the former case, the ACS data can be used to calibrate the ranging data on a photon-by-photon basis using an approach described in Section 4. If ACS data is not available or inadequate for a given run, calibration options still exist, as described in Section 5.

4. Range timing calibration with contemporaneous ACS data

APOLLO runs containing sufficient contemporaneous ACS information allow us to calibrate individual FID and LUN photon timings. Photon timing corrections are determined by identifying the position of a range photon relative to a representative ACS comb (aggregated over years). This section describes the photon-by-photon calibration procedure in greater detail.

As suggested in section 3, we observe that ACS combs from one run to the next over timescales of years largely share the same pattern of tooth *separations* but suffer from temporal shifts in TDC-space. Aggregation of tooth separations relative to a

chosen reference tooth allow us to form high signal-to-noise representative ACS combs, or “template” combs (one for each channel/gate-type combination), while effectively decoupling the templates from TDC-space. Template combs are later matched to ACS combs from individual runs and photon-by-photon range timings may then be calibrated based on their position relative to the resulting comb. Occasionally APOLLO experiences hardware changes, which are typically accompanied by the formation of a new set of ACS template combs.

4.1. The Calibration Method

Having defined a template comb for each channel and gate type, we can form a calibration mapping to convert photon TDC records to time, distinct from the CALTDC-derived calibration functional form. We assert that the teeth of the template comb should have a fixed pitch of 2.5 ns, relative to a chosen reference tooth of the template comb (the same tooth for every channel/gate combination). The choice of reference tooth is kept fixed across hardware periods in order to connect ACS results in a consistent manner. For photon events that fall between template teeth, interpolation is used to permit calibration for any location within the span of the template comb. Any photon events that fall outside the template comb see an extrapolated portion of the calibration mapping, based on the template comb’s mean tooth separation. The final step is to determine, for runs with contemporaneous ACS data, where each run’s ACS combs *actually are* within TDC-space. We accomplish this by sliding each template comb until best overlaid with the corresponding comb in an individual run (the “daughter” comb). The alignment between a given template comb and a corresponding daughter comb is optimized via a one-parameter least squares fit. The resulting fitted comb position tells us where in TDC-space the previously determined calibration mapping should be placed. We refer to the fitted template combs as “idealized” combs to avoid confusion, given “daughter” combs are formed from contemporaneous ACS photons in a given run, and “template” combs are decoupled from TDC-space.

Now that we have a calibration mapping derived at the location of each template tooth, and a best-fit position for the *actual* comb (the “idealized” comb), timing records for range photons can be obtained. Each range photon has a measured position in TDC-space which can be used to compute the corresponding ACS-derived time calibration based on their proximity to idealized comb teeth. The data reduction process as described in Section 2 continues normally after range photons are time-calibrated.

One consequence of the ACS methodology described above that users of the data should be aware of is that normal point results will change slightly over time as ACS calibrations are improved from aggregation of new ACS data. ACS template combs are aggregated within time periods of fixed hardware configurations of APOLLO. Certain hardware changes to APOLLO may influence the structure of the ACS template combs, so “new” template combs are started any time these changes are expected or observed to impact ACS results.

4.2. Range Error Statistics

At the end of each run’s data processing, we seek to summarize the scale of timing inaccuracies that were corrected by using the ACS-derived calibration into a single representative differential estimate. We may then study how ACS corrections distribute themselves over time and/or collections of runs. Since using the CALTDC calibration is the default operation in absence of sufficient ACS information, we define an ACS timing correction for an individual FID or LUN range photon as the difference between its calibrated timestamp as determined by ACS and its calibrated timestamp as determined by CALTDC:

$$C_i \equiv t_{\text{ACS},i} - t_{\text{CALTDC},i}. \quad (2)$$

The mean of timing corrections to FID and LUN range photons are obtained separately for those range photons considered to be contributors to the actual signal, and differenced such that

$$C_{\text{differential}} \equiv C_{\text{FID}} - C_{\text{LUN}}, \quad (3)$$

where C_{FID} and C_{LUN} are the mean FID and LUN timing corrections, respectively^{||}.

We invite the reader to refer to Figure 5, which displays an example result of the process described earlier in this section, for an illustrative example. More template/idealized comb data points appear than the number of daughter comb teeth, due to slightly different regions of TDC-space being sampled by ACS photons over time as we learned how to properly utilize the ACS system. The slopes of interpolated portions of the correction mapping indicate how template/idealized *adjacent* teeth are spaced — a positive slope indicates the teeth are too close, while a negative slope indicates teeth are too far apart. Additionally, it is worth noticing the significance of the sign of corrections themselves, relative to what side of the reference tooth (where $x = 0$ in Figure 5) they fall on. Negative corrections on the left side of the reference tooth and positive on the right side of the reference tooth both correspond to comb teeth being separated less than expected, relative to the reference tooth, representing compressed time. Figure 5 suggests that, relative to a chosen reference tooth, all other teeth in this template comb are systematically spaced too closely compared to truth — a tendency that was first hinted at in Figure 4 due to the weighted means being systematically lower than 2.500 ns.

The gray dotted lines in Figure 5 indicate this channel counts time too slowly in the TDC part of the timing system — ~ 60 ps less than truth, over a 20 ns window in each gate type, as suggested by the fitted slopes. Similar to Figure 4’s weighted means being offset from 2.500 ns, the slope assessment for this particular channel provides another preview of what becomes a newly discovered overall 0.4% systematic inaccuracy in the TDC when analyzed for all channels.

^{||} Note: while in Ref [10] the discussion centers around timing *offsets* rather than timing *corrections* (one is the negative of the other), the corresponding differential quantity defined at the end of Section 5.3.1 of Ref [10] had an erroneous negative sign applied; as such, any references to timing offsets within or after Section 5.3.2 should be interpreted as timing *corrections* instead.

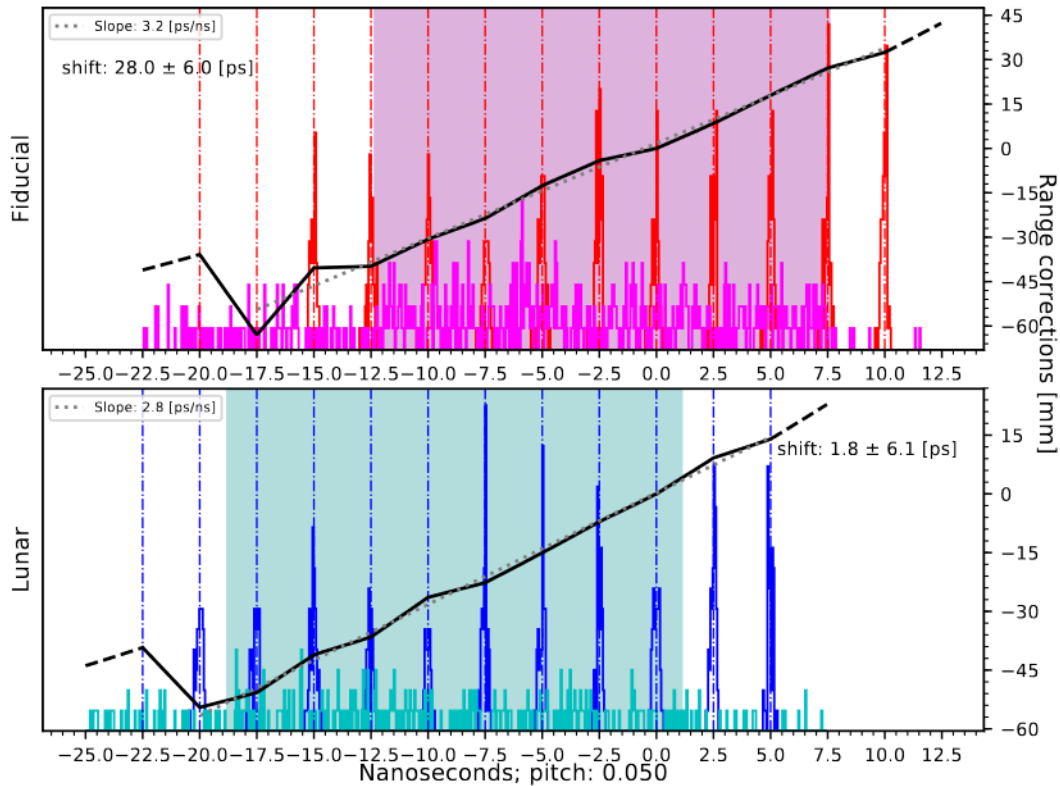


Figure 5: Actual data for the same run pictured in Figure 2 for the same channel as in Figure 4. Vertical dash-dot lines denote positions of template comb teeth which have been shifted to best align to the daughter comb shown as the red and blue spikes. This shifted comb is referred to as the “idealized” comb in the main text. Teeth on the ends of the template comb are sparsely populated, such that noise effects are more obvious. The dark sloping line indicates interpolated timing corrections, derived from the template comb as defined in Equation 2. Extrapolated portions of the timing corrections are indicated with dashed black lines. The timing corrections, whose scale in picoseconds appears on the right, demonstrate a mostly smooth structure that adds confidence to the interpolation approach. Dotted gray lines denote a linear fit to the timing corrections, excluding the first template comb tooth in each gate. The fitted slopes for each gate are $\sim 3 \frac{\text{ps}}{\text{ns}}$, or 0.3% — another preview of the eventual 0.4% systematic offset from truth discovered when averaged over all channels. The histograms shown in magenta and cyan are the raw FID and LUN range returns, respectively. Lightly shaded regions appearing in magenta and cyan atop the histograms denote the 20 ns span that registered range photons land within for the FID and LUN gates respectively; notice the 20 ns range return region is well-covered by the ACS comb for each gate type.

Figure 6 shows the distribution of ACS-based corrections and their uncertainties for APOLLO data that used the cesium clock between 2017 February 21 and 2019 August 14. The left histogram in Figure 6 suggests that APOLLO has a systematic tendency to underestimate the one-way range by ~ 4 mm, resulting from the TDC’s systematic undercounting of time. The right-side histogram indicates sub-mm accuracy of ACS corrections, and thus of APOLLO data once corrected. A more detailed discussion of this TDC systematic error effect and how it can be utilized to predict ACS-based timing corrections will be presented in Section 5.

Occasionally APOLLO experiences hardware and/or configuration changes which




Figure 6: Left: distribution of ACS corrections ($C_{\text{differential}}$) for runs using the cesium clock as the APOLLO timebase in the initial ACS data collection period, corresponding to 2017 February 21 through 2019 August 14. The range of ACS corrections is reflective of variations in the prediction used to acquire a range signal – a well-understood phenomenon that is expanded on in Section 5. Right: distribution of estimated ACS correction errors for the same dataset. A best-fit gaussian is overlaid, with a centroid of 0.4 mm.

have implications regarding range accuracy. In terms of ACS operation and usage, these hardware changes are usually accompanied by the formation of a new set of ACS template combs and systematic error statistics. More information regarding changes to APOLLO hardware and the system clock can be found in Ref [22].

5. Range timing calibration without contemporaneous ACS data

When the lunar range signal is weak, delivery of ACS light onto the APD is deliberately disabled. The addition of ACS photons interferes with telescope pointing feedback derived from APD event rates, and can effectively mask the range signal underneath a “floor” of detections when coupled with difficult observing conditions. Despite these runs lacking concomitant ACS data, we still wish to calibrate their lunar range measurements based on aggregate ACS information. Here we present a technique to utilize ACS correction trends to determine trustworthy corrections on a run-by-run basis for those runs without concurrent ACS photons, which we refer to as “non-ACS” runs.

Lacking ACS information for a run, a straightforward substitute correction could be obtained from the centroid of the distribution shown in Figure 6, which we label as the “overall” distribution. We observe, however, that we may be able to effectively

predict a run’s ACS correction by having noticed a tight correlation between derived ACS corrections and the positioning of LUN and FID photons in TDC-space. Imagine for a moment that the cyan histogram (and correspondingly the cyan shaded region) in the lower subplot in Figure 5 moves to the right while everything else in the entire figure remains fixed. The average value of ACS corrections in the lunar gate for this channel would therefore be different (more positive). Now imagine this were the case for every channel in this run. Consequently, the ACS correction for the entire run ($C_{\text{differential}}$) would be different.

In reality, the range photon distributions do move around in TDC-space somewhat over time. The prediction used to appropriately position the lunar return distribution in TDC-space is inherently imperfect, and can vary by a few nanoseconds from night to night; this causes the lunar distribution to “walk” by as much as ~ 150 TDC bins for collections of runs over multiple nights. The fiducial range photon distribution is better behaved, but can wander by ~ 7 TDC bins over time due to the FPD timing dependencies with temperature that cannot completely be eliminated within the temperature-controlled electronics/optics enclosure. We therefore seek to study the correlation between ACS corrections and the differences in TDC distributions of FID and LUN photons¶.

5.1. Corrections to “non-ACS” runs

During each run’s data processing, the positions of range photon distributions (one for each gate) are estimated by calculating the midpoint of the minimum and maximum TDC values of a windowed region over the registered range photons. We can plot ACS corrections for those runs having contemporaneous ACS photons against the FID - LUN shift and fit a line to it, as shown in the upper subplot of Figure 7. The fitted slope comes out to $0.58 \frac{\text{mm}}{\text{ns}}$, or $\sim 0.4\%$ when put in dimensionless units, which corroborates what was seen in Figures 4 and 5. It is additionally important to note that when the FID and LUN slugs are perfectly overlapped (a condition that can be contrived), APOLLO experiences a minimum of systematic error in the TDC.

The linear fit results can be used to produce a reasonable representation of what the ACS correction would have been for runs without contemporaneous ACS data, based on the run’s FID - LUN shift. An error contribution of $\sqrt{\sigma_{\text{residual}}^2 + \sigma_{\text{intrinsic}}^2}$ is root-sum-squared into the normal point uncertainty, where σ_{residual} is the model residual standard deviation (0.60 [mm]) which characterizes the scatter of points about the fit line and $\sigma_{\text{intrinsic}}$ is a measure of the intrinsic error in ACS corrections (0.40 [mm]). The lower subplot of Figure 7 displays a histogram of ACS correction residuals, obtained by subtracting the model line from the ACS corrections, with a gaussian fit function overlaid. The model residual standard deviation appears in the legend of Figure 7’s lower subplot; when compared against Figure 6, it demonstrates notable improvement

¶ Note: we did not appreciate this correlation at the time of publication of Ref [10]; a different methodology for correcting runs without contemporaneous ACS data appears there.

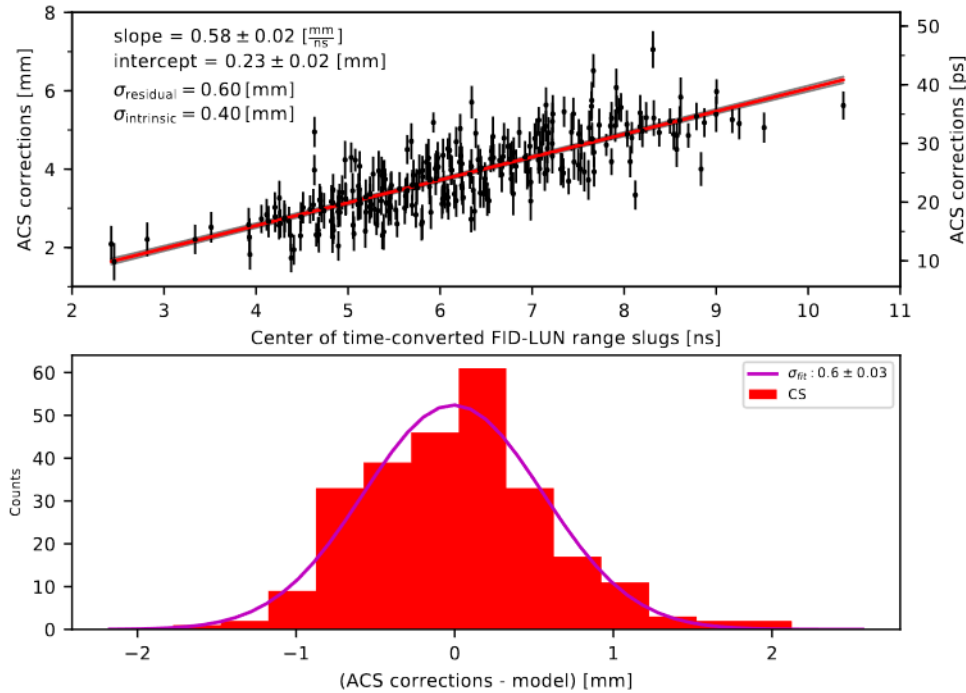


Figure 7: Top: ACS corrections versus differential (fiducial - lunar) position of range photon distributions in TDC-space, converted to time. A fitted line appears in red, with small gray shaded regions indicating the line model uncertainty. Bottom: histogram of the difference between ACS corrections and the fitted line from the upper subplot.

in the spread of ACS corrections.

5.2. Corrections to “pre-ACS” runs

We can also extend the usefulness of the non-ACS correction methodology to historic, or “pre-ACS” APOLLO runs — those that came before the installation of the ACS. The same TDC unit has been used for the entirety of the experiment, so pre-ACS corrections derived in a similar manner to non-ACS corrections should still be applicable to old APOLLO data, in principle.

The concept remains the same as in Section 5.1, except now we work with ACS correction data from the entire operation of the ACS instead of a specific period’s data (corresponding to a certain hardware configuration), such as that displayed in Figure 7. Figure 8 shows the result of this concept for all ACS runs spanning 12 September 2016 through the end of 2020; the upper subplot parallels Figure 6, while the middle and lower subplots parallel Figure 7. Four different groups of data appear, corresponding to three different hardware configurations; the earliest ACS configuration used either the cesium clock or the XL-DC clock as APOLLO’s timebase (distinct from the timebase used to send ACS pulses themselves, which is always the cesium clock) at various times,

while successive configurations used only the cesium clock as APOLLO's timebase.

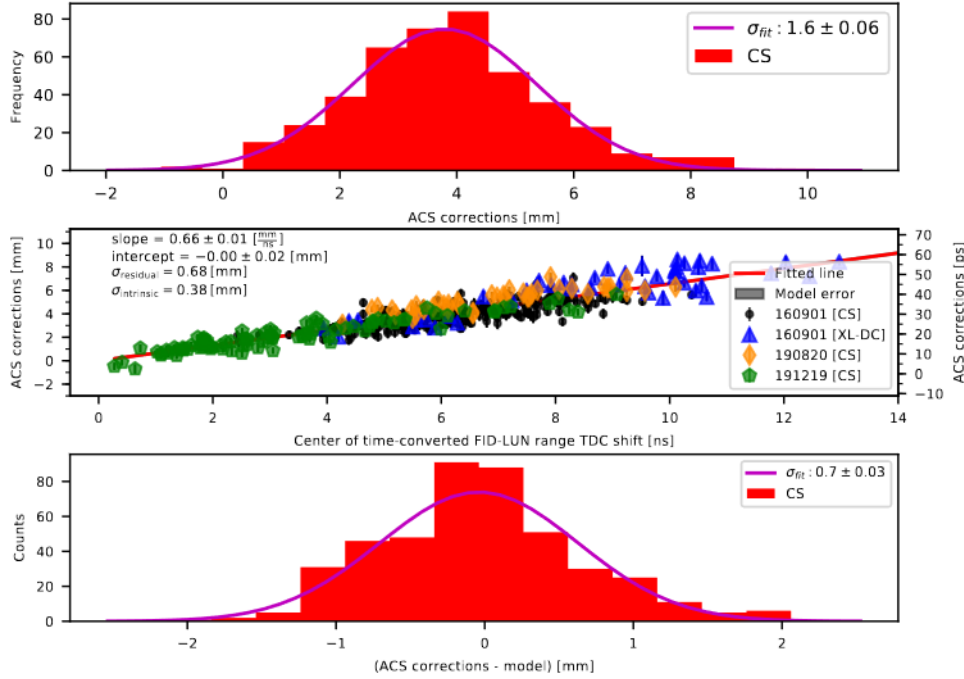


Figure 8: A version of the same plot as Figure 7, but for all ACS corrections from 2016 September 12 through 2020 December 27. Note that the black circular markers in the middle subplot are the same data points that appear in Figure 7's upper subplot.

Vertical shifts in the groups of data can be seen in the middle subplot of Figure 8, corresponding to slightly different systematic errors when the FID and LUN signals are perfectly aligned in TDC-space. A single linear fit is performed on the entire collection of data, and we can again estimate what the ACS correction would have been based on the run's FID - LUN shift. The error contribution to the NP is conceptually the same as for non-ACS runs, but the residual (0.68 [mm]) and intrinsic (0.38 [mm]) uncertainties are unique to this fitted line and dataset. The fitted slope of $0.66 \frac{\text{mm}}{\text{ns}}$, or $\sim 0.4\%$ again demonstrates the newly-discovered systematic undercounting of time by APOLLO, and further substantiates what was previously seen in Figures 4, 5 and 7.

5.3. Corrections to non-ACS and pre-ACS runs with weak lunar signals

Corrections to non- and pre-ACS runs rely on a reasonable estimation of where the fiducial and lunar signals fall in TDC-space. For runs with few lunar return photons, a reliable direct estimation of the lunar signal TDC position may not be possible. However, we may evaluate what a run's lunar signal position was expected to be, based on a few operational variables used in scheduling timing gates, plus a run's offset from the

predicted round trip time. Therefore, we wish to study the offset between directly-estimated lunar slug positions and the aforementioned operational variables for runs considered to have trustworthy measures of their lunar slug positions.




Figure 9: A plot of the offsets between directly-estimated lunar slug positions and indirectly-estimated lunar slug positions, using relevant gate positioning parameters, for runs between 2006 April 7 through 2020 December 27. Vertical shifts in groups of data points, separated by vertical blue dash-dot lines, can clearly be seen. Weighted means for each group appear as solid black lines while shaded gray regions around the solid black lines denote the weighted standard deviation for each group.

Figure 9 shows the difference between directly and indirectly estimated lunar slug positions for runs from the start of APOLLO through the end of 2020. Green circular markers denote runs whose direct estimate of the lunar slug are considered to be “good” based on certain user-defined thresholds of quantities such as estimated number of photons contributing to the lunar signal, the signal-to-noise ratio, how wide the directly-estimated lunar slug is, etc; red \times markers are runs that did not pass these cuts. Magenta diamonds are runs that did pass threshold cuts, but are far enough removed from the green points to not be considered part of them. Data point uncertainties incorporate the estimated error on the direct measure of lunar slug position as well as the prediction offset error (normal point error, pre-ACS-adjusted).

Vertical shifts in Figure 9 are obvious, and are due to (currently) unidentified/unrecorded changes to the APOLLO system that affected the lunar gate positioning such as temporary changes to the operational software, cable changes, biased channels,

etc; these less-than-obvious reasons may be further investigated in the future. However, within each group of points, the mapping between parameters used to indirectly estimate a lunar slug position and the direct measures of lunar slug positions is fairly tight and thus reliable. We know that a lunar slug offset corresponds to ACS corrections, as demonstrated by Figures 7 and 8. Therefore, the offset and spread of points in Figure 9, combined with a run’s operational variables relevant to gate positioning, are all we need to determine the lunar slug position for lower-signal runs.

Uncertainties pertinent to indirectly estimating a lunar slug value include the error on the weighted mean and the weighted standard deviation of the appropriate group of points in Figure 9, as well as the uncertainty of the normal point in question. For example, say we wish to look-up a lunar slug value for the last run performed on 2009 March 16; the relevant group in Figure 9 is the group that includes day 1000. The error on the weighted mean is ~ 13 ps, the weighted standard deviation is ~ 220 ps and the normal point uncertainty is ~ 38 ps to result in a lunar slug uncertainty of $\sqrt{13^2 + 220^2 + 38^2} \approx 224$ ps. Next, the looked-up lunar slug value and its uncertainty are used to estimate a differential slug position and a corresponding uncertainty. Finally, an ACS correction value is estimated as described in Section 5.1 (or Section 5.2 for pre-ACS runs) and the estimated ACS uncertainty is RSSd into the normal point uncertainty.

As mentioned earlier, normal points that participate in this estimation of a lunar slug position are chosen based on certain user-defined thresholds of various quantities like the estimated signal to noise ratio. However, even if a normal point is not flagged to use this logic based on threshold conditions, we will check what a looked-up lunar slug uncertainty would be and compare it against the normal point’s directly-estimated (the average TDC position of lunar range returns contributing to the signal) lunar slug error. If the looked-up uncertainty is less than the directly-estimated uncertainty, we will utilize the looked-up lunar slug information instead.

6. Conclusions

The Earth and Moon play a critical role in studying aspects of our most poorly-understood fundamental force, gravity. LLR has been able to take advantage of these local massive bodies to constrain gravitational parameters that may otherwise be difficult to investigate. Additionally, LLR is sensitive to geodetic studies that involve establishment of the lunar and terrestrial coordinate systems, as well as composition of the lunar interior. APOLLO has contributed high-caliber data to advance these aforementioned pursuits for 15 years using a robust method of data reduction, as described in Section 2.

In 2016, an Absolute Calibration System was added, allowing the independent study of APOLLO timing inaccuracies. The ACS is able to run *while collecting* range data, granting the powerful ability to correct timing inaccuracies of the lunar range data itself on a photon-by-photon basis. We are able to extend the ACS’s usefulness to those runs lacking in contemporaneous ACS data by referencing their FID - LUN shifts and predicting

appropriate timing corrections to apply to their final normal points. Finally, the original purpose of the ACS was fulfilled by identifying a significant source of systematic error in the TDC part of APOLLO’s timing system; the TDC’s systematic 0.4% inaccuracy is multiplied by how imperfectly overlapped the FID and LUN signals are to produce a RTT error.

Runs on or after September 12, 2016 with sufficient contemporaneous ACS data are correctable as described in Section 4. Runs on or after September 12, 2016 that lack ample ACS statistics and runs prior to September 12, 2016 are correctable based on the systematic error ACS uncovered, as described in Sections 5.1 and 5.2, respectively. The data presented in Figure 6 indicates APOLLO is accurate at the < 1 mm level. APOLLO normal point statistics are presented in a companion paper, [22]. That study indicates that APOLLO operates in the millimetric regime, with a median nightly range precision of 1.0 mm since the ACS was added in September 2016.

At the beginning of 2021, stewardship of APOLLO was transferred to the National Aeronautics and Space Administration’s (NASA) Space Geodesy Project (SGP) at Goddard Space Flight Center (GSFC). Lunar range data taken prior to 2021 will be presented in “JPL” (or “MINI”) format and Consolidated laser-Ranging Data (CRD) format on Zenodo [25]; CRD-formatted data will additionally appear on the International Laser Ranging Service’s (ILRS) global data centers: NASA’s Crustal Dynamics Data Information System (CDDIS) and the EUROLAS Data Center (EDC) [26].

Acknowledgments

We thank Russet J. McMillan for her continued support as APOLLO’s operator during observation, and handling of ACS functions during ranging. We thank Ed Leon for periodically performing measurements of clock phase at the observatory, and for occasional hardware troubleshooting. We thank the NASA Postdoctoral Program (NPP) for continuing to fund this study after the NASA transition took place; we thank Stephen M. Merkowitz of NASA GSFC’s Space Geodesy Project for his aid in facilitating the NPP appointment and support thereafter. We also thank Frank G. Lemoine of NASA GSFC for providing feedback on this publication prior to formal peer review.

This work is based on access to and observations with the Apache Point Observatory 3.5-meter telescope, which is owned and operated by the Astrophysical Research Consortium. This work was jointly funded by the National Science Foundation (PHY-0602507, PHY-1068879, PHY-1404491, PHY-1708215) and the National Aeronautics and Space Administration (NNG04GD48G, NAG81756, NNX12AE96G, NNX15AC51G, 80NSSC18K0482). JBRB acknowledges the support of the NASA Massachusetts Space Grant (NNX16AH49H).

References

- [1] Zyla P *et al.* (Particle Data Group) 2020 *PTEP* **2020** 083C01
- [2] Genova A, Mazarico E, Goossens S, Lemoine F G, Neumann G A, Smith D E and Zuber M T 2018 *Nature* **9**
- [3] Archibald A M, Gusinskaia N V, Hessels J W T, Deller A T, Kaplan D L, Lorimer D R, Lynch R S, Ransom S M and Stairs I H *Nature* **559**
- [4] Murphy T W 2013 *Reports on Progress in Physics* **76** 076901
- [5] Nordtvedt K 1968 *Phys. Rev.* **170** 1186–1187
- [6] Chang R F, Alley C O, Currie D G and Faller J E 1972 Optical properties of the Apollo laser ranging retro-reflector arrays. *Space Research Conference* vol 1 pp 247–259
- [7] Fournet M 1972 Le réflecteur laser de Lunokhod. *Space Research Conference* vol 1 pp 261–277
- [8] Murphy Jr T W, Adelberger E G, Battat J B R, Carey L N, Hoyle C D, LeBlanc P, Michelsen E L, Nordtvedt K, Orin A E, Strasburg J D, Stubbs C W, Swanson H E and Williams E 2008 *Publ. Astron. Soc. Pac.* **120** 20
- [9] Adelberger E G, Battat J B R, Birkmeier K J, Colmenares N R, Davis R, Hoyle C D, Ruixie L H, McMillan R J, Murphy Jr T W, Schlerman E, Skrobol C, Stubbs C W and Zach A 2017 *Class. Quantum Grav.* **34**
- [10] Colmenares N R 2020 *Calibrating the Apache Point Observatory Lunar Laser-ranging Operation (APOLLO) Apparatus* Ph.D. thesis University of California San Diego <https://escholarship.org/uc/item/5kw8935h>
- [11] Müller J, Biskupek L and Hoffman F 2014 Earth orientation and relativity parameters determined from llr data *Proceedings of the 19th international workshop on laser ranging* (Annapolis, MD, USA) pp 27–31
- [12] Williams J G, Boggs D H and Ratcliff J T 2016 Lunar tidal recession *Lunar and Planetary Science Conference* p 1096
- [13] Williams J G and Boggs D H 2015 *Journal of Geophysical Research: Planets* **120** 689–724
- [14] Williams J G, Boggs D H, Yoder C F, Ratcliff J T and Dickey J O 2001 *Journal of Geophysical Research: Planets* **106**
- [15] Currie D G, Dell’Agnello S, Delle Monache G O, Behr B and Williams J G 2013 *Nuclear Physics B - Proceedings Supplements* **243-244** 218 – 228 ISSN 0920-5632
- [16] Ciocci E, Martini M, Contessa S *et al.* 2017 *Advances in Space Research* **60** 1300 – 1306 ISSN 0273-1177
- [17] Turyshev S G, Williams J G, Folkner W M *et al.* 2013 *Experimental Astronomy* **36** 105–135
- [18] He Y, Liu Q, Duan H Z, He J J, Jiang Y Z and Yeh H C 2018 *RAA* **18** 136
- [19] Araki H, Kashima S and Noda H o 2016 *Earth, Planets and Space* **68**
- [20] Preston A M and Merkowitz S 2013 *Appl. Opt.* **52** 8676–8684
- [21] Preston A M and Merkowitz S 2014 *Optical Engineering* **53** 1 – 6
- [22] Battat J B R, Adelberger E, Colmenares N R, Farrah M, Gonzales D P, Hoyle C D, McMillan R J, Murphy, Jr T W, Sabhlok S and Stubbs C W 2023 *Publ. Aston. Soc. Pac.* **XX** XX
- [23] Müller J, Murphy T W, Schreiber U, Shelus P J, Torre J M, Williams J G, Boggs D H, Bouquillon S, Bourgoïn A and Hofmann F 2019 *Journal of Geodesy* **93**
- [24] Liang Y, Murphy Jr T W, Colmenares N R and Battat J B R 2017 *Class. Quantum Grav.* **34**
- [25] Colmenares N R, Battat J B R, Gonzales D P, Murphy T W J and Sabhlok S 2023 Apache Point Observatory Lunar Laser-ranging Operation (APOLLO) normal point data: 2006 through 2020 URL <https://doi.org/10.5281/zenodo.7818557>
- [26] Noll C E, Ricklefs R, Horvath J, Mueller H, Schwatke C and Torrence M 2019 *Journal of Geodesy* **93**