

# UC San Diego

## UC San Diego Electronic Theses and Dissertations

### Title

Calibrating the Apache Point Observatory Lunar Laser-ranging Operation (APOLLO) Apparatus

### Permalink

<https://escholarship.org/uc/item/5kw8935h>

### Author

Colmenares, Nicholas Raymond

### Publication Date

2020

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA SAN DIEGO

**Calibrating the Apache Point Observatory Lunar Laser-ranging Operation (APOLLO)  
Apparatus**

A dissertation submitted in partial satisfaction of the  
requirements for the degree  
Doctor of Philosophy

in

Physics

by

Nicholas Raymond Colmenares

Committee in charge:

Professor Thomas W. Murphy, Jr., Chair  
Professor William A. Coles  
Professor Frederick C. Driscoll  
Professor Dusan Keres  
Professor Shelley A. Wright

2020

Copyright  
Nicholas Raymond Colmenares, 2020  
All rights reserved.

The dissertation of Nicholas Raymond Colmenares is approved, and it is acceptable in quality and form for publication on microfilm and electronically:

---

---

---

---

---

Chair

University of California San Diego

2020

## DEDICATION

To my marked academic mentors through the years for all their amazing help, effort, mindful attention and caring: Tracy Hanson, first grade teacher; Marie Jordan, third grade teacher; Linda Huttzell, ninth & eleventh grade Mathematics teacher; Susan James, tenth & twelfth grade English teacher; Scott Wissink, undergraduate research advisor; Tom Murphy, graduate research advisor.

To my parents Patricia A. Scheer and Thomas J. Colmenares, and my Aunt, Shirley J. Colmenares, for their unyielding love, faith, encouragement, and willingness to listen.

To my feline son, Juno, with me through thick and thin since the early days of undergraduate study.

To Meg Wiley, for having a significant positive impact on my personal development.

## EPIGRAPH

*If we define Futurism as an exploration beyond accepted limits,  
then the nature of limiting systems becomes the first object of exploration.*

—Frank Herbert

## TABLE OF CONTENTS

Signature Page . . . . .	iii
Dedication . . . . .	iv
Epigraph . . . . .	v
Table of Contents . . . . .	vi
List of Figures . . . . .	viii
List of Tables . . . . .	x
Acknowledgements . . . . .	xi
Vita . . . . .	xii
Abstract of the Dissertation . . . . .	xiii
Chapter 1	
Introduction . . . . .	1
1.1 Science Motivation . . . . .	1
1.1.1 Brief History of LLR . . . . .	2
1.1.2 Science Deliverables . . . . .	4
1.2 Challenges to Science Deliverables . . . . .	6
Chapter 2	
Meet and Greet: An Introduction to APOLLO . . . . .	9
2.1 APOLLO Overview . . . . .	9
2.2 APOLLO Tip/Tilt Mirror . . . . .	11
Chapter 3	
Proving APOLLO’s Worth: The Absolute Calibration System (ACS) . . . . .	16
3.1 ACS Concept . . . . .	16
3.2 ACS Overview . . . . .	18
3.2.1 ACS laser and clock overview . . . . .	18
3.2.2 ACS Instrument Monitoring and Communications . . . . .	20
3.3 Pulse Processing . . . . .	22
3.4 APD Delivery Optics . . . . .	26
3.5 ACS Laser . . . . .	28
Chapter 4	
Early ACS Results . . . . .	32
4.1 Clock Comparison . . . . .	32
4.2 Example ACS run and products . . . . .	34
4.3 Range error statistics . . . . .	44
4.4 Initial Conclusions . . . . .	46

Chapter 5	ACS Corrections to Range Data . . . . .	48
	5.1 ACS Corrections Concept . . . . .	49
	5.1.1 APOLLO Range Photon Timing . . . . .	49
	5.1.2 Correction Scheme . . . . .	53
	5.1.3 Establishing Trust in the Interpolated Mappings . . . . .	58
	5.2 Aggregating the ACS Corrections to Meaningful Levels . . . . .	64
	5.2.1 Channel Offsets and Tooth Labelling . . . . .	64
	5.2.2 Aggregation of Tooth Separations . . . . .	68
	5.2.3 Processing the Aggregate Tooth Separations . . . . .	71
	5.3 Correction Scheme Application . . . . .	78
	5.3.1 ACS-enabled Run Corrections . . . . .	78
	5.3.2 Updated Range Error Statistics . . . . .	80
	5.3.3 Corrections to ACS-disabled Runs . . . . .	82
	5.3.4 ACS Offsets Over Time . . . . .	89
	5.4 Conclusions and Future Work . . . . .	91
Appendix A	Final notes . . . . .	93
Bibliography	. . . . .	94



## LIST OF FIGURES

Figure 2.1:	Layout of main APOLLO optical bench components . . . . .	12
Figure 2.2:	Design assembly of the new RX mirror setup . . . . .	13
Figure 2.3:	RX assembly after fabrication and installation . . . . .	15
Figure 3.1:	ACS enclosure schematic diagram . . . . .	19
Figure 3.2:	ACS pulse processing system schematic diagram . . . . .	21
Figure 3.3:	ACS optical pulse train conditioning . . . . .	24
Figure 3.4:	ACS optical delivery setup . . . . .	26
Figure 3.5:	ACS receiver/optic mount design . . . . .	27
Figure 3.6:	The custom ACS optic and mount after fabrication and installation . . . . .	28
Figure 3.7:	Schematic diagram of the ring-shaped fiber oscillator. See text for details, and Ref. [1] for patent information. . . . .	29
Figure 3.8:	Spectrum of the frequency-doubled pulses at the output of the SHG. . . . .	30
Figure 4.1:	Clock Allan deviation measurements . . . . .	33
Figure 4.2:	TDC-space data visualization for the first on-sky ACS run . . . . .	35
Figure 4.3:	Relative timing between 50 MHz XL-DC clock and 80 MHz ACS laser pulses	36
Figure 4.4:	TDC-space data visualization of the <i>true</i> ACS comb spacing . . . . .	42
Figure 4.5:	Example histogram of ACS-determined APOLLO range errors . . . . .	43
Figure 4.6:	Channel-aggregated histograms of run range errors for early ACS results .	45
Figure 5.1:	APOLLO event timeline with ACS pulse train overlaid . . . . .	51
Figure 5.2:	Data visualization of range photon timing before and after ACS photon identification and sequestration . . . . .	52
Figure 5.3:	Raw (except for channel aggregation) TDC-space histogram visualization of range and ACS photon timing separately . . . . .	54
Figure 5.4:	Cartoon demonstration of ACS timing offset determination (“interpolation” correction method concept) . . . . .	62
Figure 5.5:	Time-aggregate timing offset map example for single run, channel . . . . .	63
Figure 5.6:	Relative channel offset data product visualization for single run, channel . .	65
Figure 5.7:	Channel offsets visualization across time for single channel . . . . .	67
Figure 5.8:	Example plot of ACS tooth timing separations for particular pair of teeth, across time . . . . .	69
Figure 5.9:	Example plot of ACS tooth timing separations for particular pair of teeth, across time, after uncertainty inflation . . . . .	70
Figure 5.10:	Attempt at tooth separation timing aggregation across gates but not channels	71
Figure 5.11:	Attempt at tooth separation timing aggregation across channels but not gates	72
Figure 5.12:	Demonstration of timing alignment between particular ACS parent comb and daughter comb . . . . .	76
Figure 5.13:	Interpolated line slope distribution for use in timing map extrapolation regions	77

Figure 5.14: Updated channel-aggregated histogram of run range timing errors and associated round trip time $\chi^2$ measure . . . . .	81
Figure 5.15: Distributions of ACS offsets for each run, each session, and runs compared to sessions . . . . .	84
Figure 5.16: Cumulative distribution functions of the <code>session wMean</code> relative distribution and its randomly sampled (from the <code>overall</code> distribution) analog . . . . .	86
Figure 5.17: Demonstration of the use of a joint probability distribution in correcting runs without ACS data present . . . . .	87
Figure 5.18: Scatter plot across time (runs) of run range timing errors . . . . .	90

## LIST OF TABLES

Table 3.1:	Laser average power and pulse energy at various stages. . . . .	31
Table 4.1:	Example APOLLO range errors depending on how many channels and shots are considered . . . . .	38
Table 4.2:	Clock drift timing error demonstration . . . . .	39
Table 4.3:	Timing error demonstration after clock drift removal . . . . .	40
Table 4.4:	Differential channel offsets (partial set; channel 15 is not used). . . . .	40
Table 5.1:	Example calculation of timing offsets from fitted teeth centroids. . . . .	55
Table 5.2:	Scaling of ACS-disabled run correction uncertainties, as a function of number of ACS runs in the corresponding session . . . . .	88
Table 5.3:	<i>Adjusted</i> scaling of ACS-disabled run correction uncertainties, as a function of number of ACS runs in the corresponding session. . . . .	89

## ACKNOWLEDGEMENTS

I would like to acknowledge Professor Tom Murphy for his years of patient mentorship and guidance, poignant insights on a wide array of professional, personal and societal subjects, as well as general kindness and understanding.

I also appreciate the significant collaborative effort and acumen of Professor James Battat, of Wellesly College.

Additionally, Nathan Johnson, former graduate student to Tom Murphy, is due gratitude for patiently fielding many questions and his general wisdom as I familiarized myself in the research group.

Russet McMillan, of Apache Point Observatory, is thanked for her years of dedication in running observations. Ed Leon, also of Apache Point Observatory, assisted in various engineering efforts.

Finally Brad Hanson, in charge of the UCSD Physics Department Machine Shop, provided much advice towards my hardware fabrication efforts.

Section 1.2 and chapters 3 and 4 are excerpts with slight modifications and additions from the material appearing in *Classical and Quantum Gravity*, 2017. 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.; Zach, A.

Figure 4.1 was excerpted from the material as presented in *Classical and Quantum Gravity*, 2017. Liang, Y.; Murphy, Jr., T. W.; Colmenares, N. R.; Battat, J. B. R.

## VITA

- 2010-2012 Student Researcher, Indiana University, Bloomington; Advisor: Scott Wissink
- 2011 B. S. in Physics, Indiana University, Bloomington
- 2011 B. S. in Astronomy/Astrophysics, Indiana University, Bloomington
- 2012-2014 Graduate Teaching Assistant, University of California San Diego
- 2014 M. S. in Physics, University of California San Diego
- 2014-2020 Graduate Student Researcher, University of California San Diego; Advisor: Tom Murphy
- 2020 Ph. D. in Physics, University of California San Diego

ABSTRACT OF THE DISSERTATION

**Calibrating the Apache Point Observatory Lunar Laser-ranging Operation (APOLLO)  
Apparatus**

by

Nicholas Raymond Colmenares

Doctor of Philosophy in Physics

University of California San Diego, 2020

Professor Thomas W. Murphy, Jr., Chair

Lunar laser ranging has long provided precision tests of gravitational effects including, but not limited to, the strong equivalence principle (SEP), geodetic precession, secular evolution of the gravitational constant, gravitomagnetism, and the inverse square law. The most recent ranging station, the Apache Point Lunar Laser-ranging Operation (APOLLO), achieves median lunar range precision to the  $\sim 2$  mm level. However, the best model residuals (observed - predicted) are an order of magnitude larger than that, throwing into question whether APOLLO has gross systematic inaccuracies in its data collection, or the models are incomplete in some manner. The recent addition of an Absolute Calibration System (ACS), consisting of a high-repetition-rate, low

jitter calibration laser slaved to a cesium clock, allows for independent assessment of APOLLO system accuracy and the ability to correct ranging data in-situ by delivering “truth” photons to the APOLLO detector at well-known time intervals. Initial ACS analysis suggested APOLLO is accurate to the  $\sim 2 - 3$  mm level, pushing primary suspicion of model residuals onto the models themselves. Further study, presented in the later chapters of this dissertation, has corroborated this story. A newer and more robust method of correcting those data runs containing ACS information has been developed, demonstrating the ability to remove range bias *without significantly affecting the uncertainty of the range result*. Additionally, this correction method has been extended to permit bias corrections to data runs *without ACS* information present, based on trends in ACS corrections overall and on a per-observation-night basis, given the presence of similar scale ( $\sim 2 - 3$  mm) variations in range bias over the course of a year.

# Chapter 1

## Introduction

### 1.1 Science Motivation

Modern physics is primarily built upon two pillars: general relativity and quantum mechanics, which describe nature at its largest and smallest scales, respectively. Relativity contains singularities in which infinitely high densities result from matter being compressed to very small scales, and cannot further describe physical interactions in such regions. One might then naturally expect quantum theory to play a key role in extending our understanding of singularities. Unfortunately, the standard model of particle physics, which describes three fundamental forces to a very high accuracy, utterly fails to unify gravity with quantum theory, resulting in a variety of unphysical expectations and inconsistent descriptions.

Even though gravity is arguably the most commonly experienced force on a day-to-day basis, it is also the weakest of the fundamental forces by a wide margin. e.g., if we scale the coupling constant of the strong force to be 1, electromagnetism sits at  $\frac{1}{137}$ , the weak force at  $1 \times 10^{-6}$ , while gravity's scaled coupling constant is a paltry  $1 \times 10^{-39}$  [2]! Due to this, gravity is also the most poorly tested of the fundamental forces by modern-day experiments. Thus, of the two pillars of modern physics, gravity is considered to be less completely described than quantum



mechanics. Very massive objects are the greatest boon towards meaningful measurements of certain aspects of gravitation, and as such, solar system objects traditionally have provided the easiest measurements to accomplish this. However the scale of interest extends widely, from pulsars and the production of gravitational waves to small laboratory scales. 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, or SEP) come from this natural laboratory, and the goal of modern-day LLR experiments is to improve upon these constraints [3, 4].

### **1.1.1 Brief History of LLR**

LLR began in 1969 after the installation of a retroreflector array on the Moon by the astronauts of Apollo 11. Retroreflectors are composed of fused silica corner cubes operating via total internal reflection. Subsequent retroreflector arrays were placed on the Moon by Apollo missions 14 and 15, as well as by remotely landed Soviet rovers Lunokhod 1 and 2. In LLR, a pulsed laser is fired from the Earth to one reflector array, and the round trip time (RTT) of photon flight is recorded. Round trip times are commonly expressed as a one-way distance ( $D$ ) for a more intuitive description of lunar range measurements; routine range precision on the millimeter scale is achieved in modern times, with further improvements expected in the future. Observations throughout the course of the Moon's orbit enables the formation of "maps" of the shape of the lunar orbit, allowing for judgement of the validity of opposing gravitational theories; For scale,  $\sim 10$  m post-Newtonian (relativistic) effects on the Earth-Moon separation  $D$  are present, as measured in the solar system barycenter (SSB) frame. Higher precision tests of gravity may soon observe deviations at the quantum scale, potentially leading to a more complete understanding and modified theory [3].

The 3.1 meter telescope at the Lick Observatory began ranging to the Moon shortly after the placement of the first reflector array in August of 1969, primarily to demonstrate initial data

acquisition. Shortly thereafter a second campaign began at the McDonald Observatory using the 2.7 meter aperture (shared) telescope [5]. It utilized a nanosecond pulse width ruby laser firing at about 0.3 Hz at around  $3 \frac{\text{J}}{\text{pulse}}$ , thereby achieving 20 cm range precision, collecting photons at a rate as high as  $0.2 \frac{\text{photon(s)}}{\text{pulse}}$ . This supplied all of the meaningful measurements for the next decade. Range results are packaged into a representative measurement called a “normal point (NP),” typically built from around 20 photon returns over data collection periods, or “runs”, of tens of minutes [3].

LLR then saw a large improvement in data acquisition starting in the mid-1980s. A French station opened in 1984 using a 1.5 m telescope at the Observatoire de la Cote d’Azur (OCA) producing pulses of 75 mJ, 70 ps widths from a Nd:YAG laser firing at 10 Hz [6]. A station in Haleakala, Hawaii contributed range results for a brief stint between 1984 and 1990. In 1985 the McDonald operation went through an upgrade and relocation on-site, now called the McDonald Laser Ranging System (MLRS) [7]. It moved to a dedicated 0.76 m aperture telescope also being used for satellite laser ranging, with a 150 mJ pulse strength, 200 ps width, 10 Hz rep rate Nd:YAG laser. OCA and MLRS were the main work horses for data acquisition from 1990 to 2006, with typical normal points consisting of  $\sim 40$  and 15 photons, respectively, achieving range precision at the 1 to 2 cm level. More history and other efforts can be found elsewhere [5, 8, 9].

In 2006, the Apache Point Observatory Lunar Laser-ranging Operation (APOLLO) came online. Using a 90 ps wide, 115 mJ Nd:YAG laser pulse at 20 Hz, APOLLO normal points typically consist of a few hundred photons. Range precision of a few millimeters was historically achieved by APOLLO, and a major hardware change in 2013 resulted in range precision between 1 and 2 mm. While two other stations (OCA, MLRS) routinely contribute data (as well as the Matera and Wetzell stations occasionally contributing data) at the centimeter level, APOLLO has been leading the charge in data density, data quality and variety of reflector utilization since coming online [3].

## 1.1.2 Science Deliverables

Traditionally the most significant measurement of gravity that LLR has contributed to is the equivalence principle, one of the foundational principles of general relativity. The equivalence principle is commonly described in a few different forms, each with successively more stringent conditions. The key concept involved is that physical laws do not differ between inertial frames, and as such, all forms of mass-energy experience the same acceleration under a gravitational potential.

The weak equivalence principle (WEP) asserts equality in the outcome from two inertial frames for any local non-gravitational experiment involving masses with negligible gravitational binding energy. The masses may differ in internal structure and composition; In other words it argues for the equivalence of inertial mass with the gravitational mass of an object. Current limits on the differential acceleration seen from potential violations of the WEP are at the level of  $\frac{\Delta a}{a} < 1.3 \times 10^{-13}$  [9, 10]. Even though LLR is sensitive to the WEP, given that the Earth has a massive iron nickel core and the Moon does not, laboratory torsion balance experiments are typically more sensitive. While they are forced into using significantly smaller test masses, they can optimize selection of materials to maximize intrinsic differences in internal composition to be greater than that of the Earth-Moon system [3].

The Einstein equivalence principle (EEP) extends the WEP to include experiments not at rest in an inertial frame, so that the outcome is independent of its location within space-time. Any inequalities seen in these outcomes could imply differing values of important dimensionless physical parameters such as the fine structure constant, vacuum permittivity, the speed of light, the gravitational constant  $G$ , etc.

Finally the SEP, the most inclusive and demanding formulation of the equivalence principle, extends the EEP to include local gravitational experiments and gravitational binding energy. In order to be sensitive to SEP violations to a high degree of confidence, measurements *must* involve objects with a significant amount of gravitational binding energy, and thus must involve a

massive body. Naturally, laboratory experiments are unable to compete against astronomical-scale objects in realizing large amounts of gravitational binding energy. The fractional contribution of Earth's gravitational binding energy to its overall gravitational energy is about  $4.6 \times 10^{-10}$ ,  $\sim 20$  times larger than that of the Moon's. This easily allows for testing of SEP violations by looking for differential accelerations giving an expected amplitude of 13 m for a full violation of the SEP. The current limit on the Nordvet parameter  $\eta$ , which characterizes the degree of violation of the SEP, is  $3 \times 10^{-4}$ . Millimeter-scale lunar ranging should bring this limit down by another order of magnitude [3].

LLR can test several other gravitational science deliverables. Current limits on the secular evolution of the gravitational constant constrained by the long time span of LLR measurements is  $7 - 9 \times 10^{-13}$  per year [11, 12]. Gravitomagnetism, the gravitational analog to Maxwell's equations, is constrained by LLR to  $< 0.2\%$  [13, 14]. Geodetic precession is also known to the 0.3% level by examination of the orientation of lunar orbit [15].

LLR is also sensitive to a variety of non-gravitational science deliverables. The Moon's orientation, orbit and tidal deformations have some dependence on the interaction of its internal mass structure with imposed forces and torques. Range results from the five different reflectors can be used to aid in understanding of the lunar interior. Range results are also, of course, sensitive to the three-dimensional orientation of the Earth, at the time of ranging. In a similar manner, LLR can contribute to knowledge of the physical processes affecting Earth's orientation, and add to measurements of the relative orientation of the ecliptic and equatorial planes [3].

Further in-depth explanation of science deliverables can be found elsewhere. What follows will serve as an introduction to APOLLO, and motivations for independently assessing system accuracy given model inaccuracy.

## 1.2 Challenges to Science Deliverables

Since coming online, APOLLO has produced the majority of LLR measurements worldwide, and also habitually acquires ranges to 4—5 reflectors in each 1 hr observing session. In principle, this significant improvement in measurement precision should push tests of relativistic gravity by a similar factor, down to the 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 [9, 10, 16, 17, 18].

The basic scheme involves 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 [19].

This is not the case for *any* of the current LLR models. Residuals for APOLLO’s  $\sim 2$  mm scale measurements are characterized by a spread ranging from about 15 mm in the best model to about twice this in others [20, 21, 22, 23, 24]. No obviously 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 [19]?

Efforts to improve models are ongoing, but years have passed without a clear indication 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 [25] (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 fiducial corner cube measurement may be impacted by the ring-down of the laser fire, which involves switching  $\sim 3000$  V in a matter of a few nanoseconds and thus creating electromagnetic interference (EMI) that can impact the nearby detector and timing electronics [19].

In light of these questions and concerns, we devised the ACS 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 precise apparatus to a few-millimeter precise *and* accurate operation [19].

This dissertation ultimately seeks to describe the analysis process of utilizing ACS data, and what results have come from it. A robust method of systematically correcting APOLLO range data has been developed, incorporating various trends seen over long timespans, as quantified by the ACS. The conclusion is that we see unappreciated influences at the  $\sim 2 - 3$  mm level, but nothing that would account for the 15 mm-scale of model residuals. The beneficial implication is that 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.

This section, in part, has been excerpted from the material as presented in *Classical and Quantum Gravity*, 2017. 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.; Zach, A. [19]

# Chapter 2

## Meet and Greet: An Introduction to APOLLO

Before proceeding, it would behoove us to issue a concise APOLLO system overview, to introduce required and/or useful knowledge to further motivate discussion of the ACS system, and to highlight certain pre-ACS efforts.

### 2.1 APOLLO Overview

APOLLO delivers energetic laser pulses ( $\sim 115$  mJ at 20 Hz) to the Moon, typically detecting  $\lesssim 1$  return photo-electron per pulse. The detector is a Lincoln Lab  $4 \times 4$  avalanche photodiode array (APD) with  $\lesssim 100$  ps resolution. APOLLO turns on the APD twice per cycle: once for the local fiducial corner cube return and once for the expected lunar return [19].

Fine-scale time measurement is governed by a 16 channel (corresponding to the 16 APD elements) time-to-digital converter (TDC) that receives START signals from detector photon events. STOP pulses, common to all channels, are sliced out of the 50 MHz low-phase-noise clock train from the recently added cesium clock (Microsemi 5071A) for use with the ACS system.



TDC timing results, via measuring the voltage across a capacitor charged at a well-known current, are characterized by 15 ps jitter, 25 ps resolution and 100 ns range.

Intermediate-scale tracking of time is handled by the APOLLO Command Module (ACM), which utilizes the 50 MHz clock train to give us 20 ns resolution between events. The ACM relies on a series of counters, registers, and comparators to schedule various gate events (for clarity, a “gate” is a detector activation event associated with laser “shots”). The ACM is the primary hardware interface for the rest of APOLLO, overseeing timing, coordination and assorted device controls. The cesium clock additionally sends its 1 pulse-per-second (PPS) signal to the ACM to allow tracking of seconds, times of events within each second, and a free-running (resets every  $\sim 5.4$  s) counter for the number of 50 MHz clock pulses elapsed (as a consistency check).

It is important to note that the fiducial signal—photon returns from a local corner cube at the secondary mirror which intercepts and reflects a small portion of outgoing light, and is heavily attenuated down to the  $\sim 1$  photon level upon receipt at the detector—forms the basis of our differential measurement. The large ( $\sim 1.5 \mu\text{s}$ ) jitter of the range laser makes it impossible to predict fiducial and lunar returns to the necessary accuracy ( $\sim 20$  ns) for proper gate positioning. Given this difficulty, we rely on detection of outgoing laser light from the laser housing by a high-fidelity fast photodiode (FPD). Channel 15 of the TDC is reserved for cable-delayed timing of the FPD signal, within the fiducial gate. Through clever use of various cable delays, well-known per-channel timing offsets, and the time between fiducial and lunar signals, the differential measurement times the crossing of outgoing light at the telescope axis intersection to the return light at the same point. The number of clock pulses between the fiducial STOP signal and the lunar STOP signal are counted for use in the differential measurement. Further discussion of the differential timing scheme is deferred until section 5.1.1.

Shortly after joining APOLLO in the summer of 2014, I was handed the responsibility of maintaining and modifying the software suites for both APOLLO operation and data reduction. I participated in various efforts, including assisting in detector characterization, studying data


reduction byproducts, and heading-up replacement of an important optical component (tip/tilt mirror), before the ACS project began. A summary of the tip/tilt mirror effort follows.

## 2.2 APOLLO Tip/Tilt Mirror

As a short aside, the tip/tilt mirror is a critical portion of APOLLO's optical system that allows us to fire the range laser towards where the Moon *will* be, while directing the detector's attention to where the Moon *was* at reflection, by the time the light returns. In the latter part of 2014 I was put in charge of designing, fabricating (the add-on mount portion) and implementing a replacement for the primary optic mount responsible for offsetting the transmit and receive beam paths: what we call the "RX" mirror, which can be seen in the lower left-hand portion of Figure 2.1. This *was* an all-in-one New Focus (now Newport) model 8852, a remotely operable mirror/mount capable of three-dimensional pointing offsets along the tip/tilt and translation axes. The stock 8852 actuators were replaced with piezoelectric pico motor actuators (New Focus 8301), driven by a multi-axis controller (New Focus 8732).

Throughout the RX mirror's lifetime it suffered from hysteresis effects, as it had no closed loop feedback to verify its position upon a completed movement. The effects were large enough that APOLLO had to have Russet, our local instrument operator at APO, recalibrate its position on roughly a weekly time scale, which proved to be inconvenient. Later in its lifetime the power supply for one direction of one of the actuator axes failed. This meant that we no longer had full remote adjustability of the RX mirror; it required manual adjustment if the misalignment grew large enough.


With these issues in mind we desired a new setup that either had closed loop feedback, or a custom setup with effectively no hysteresis, fairly quick response, and most importantly just a fully working steering mirror. After some digging we opted to go for a custom setup as many of the ready-made closed loop feedback systems were fairly expensive. After carefully considering



**Figure 2.1:** Layout of main APOLLO optical bench components, that reside next to the range laser. The orientation is rotated 90° counterclockwise from what one would see standing at the back of the telescope, such that the left end is actually “up” in the real system. The outgoing laser beam begins as the thin green beam towards the lower left, travelling up, striking a reflective patch on a rotating optic called “T/R,” responsible for switching between transmit/receive modes. The receive beam initially comes back along the transmit path, but misses the reflective patch of the T/R, becoming the yellow beam on the left side, before entering the receiver housing as the right-going green beam. ~ 8% of the received light is sent to a CCD camera, indicated by the blue beam. The (original) RX mirror is seen at the very lower left with the three-axis actuated (red) mount.

a number of possible designs for the custom attachment to a mirror mount, a final design was settled upon. The design included a basic mirror mount which we would modify in order to bolt a custom-made mounting plate onto. The plate would then have highly accurate stepper motors attached to it, responsible for moving the turning screws of the mirror mount.

The design, seen in Figure 2.2 began with a basic mirror mount: Newport model 9852-k (“k” option to disclude knobs at the ends of the actuator screws), an updated version of the old 8852 mirror mount. Two Nema 17 stepper motors with custom shaft length and molex connectors were ordered from Lin Engineering, for 2-axis control (translation not needed). Motor hardware control is handled by a multi-axis motor controller from Trinamic Motion Control, model TCM-343, along with a 12 V AC/DC converter to power it. Serial communication to the motor controller is governed by a Raspberry Pi 2, model B, aptly named “PIRX”. User-commanded movements are sent through a telnet connection between APOLLO’s control computer (“HOUSTON”) and PIRX, and are recorded to a log file for later reference, when needed.




**Figure 2.2:** Design assembly of the new RX mirror setup. The mirror mount itself has been made transparent, to better visualize the custom components. Pre-made designs of the sprockets were not available, and were modelled to described specifications. The large sprockets on the mirror mount turning screws lack teeth because we just need information about relative proximity of components (the sprockets on motor shafts were designed prior to deciding to skip tooth modelling). The kevlar belt between the sprockets is not displayed. Slotted screw holes for mounting the stepper motors were used, with the pivot at one screw, to allow for adjustment of belt tightness. Custom screw holes were drilled into the side of the mirror mount, to allow for mounting of the custom components to its side.

Sprockets from WM Berg are used on the ends of the mirror mount screws (GP32A32-30), and shaft ends of the stepper motors (GC32A28-14) which are clamped down via hub-style clamps (CG1-12-A). Extended length mirror mount screws were needed to accommodate the sprocket and set screw hub, which were also obtained from Newport: model 9313-k. Timing chains connect the pairs of sprockets using a high-strength Kevlar core due to the small bend radius of the timing chain around the sprockets, also allowing for up to seven degrees of angular misalignment in the chain. I fabricated custom add-on mounting hardware, with carefully-designed positioning

of stepper motor mount points for proper belt tension and various clearance constraints, in the UCSD Physics Department Machine Shop.

Though the custom setup does not include closed-loop feedback to verify position, there was no hysteresis observed during the testing/characterization phase. With a pointing range of tens of arcseconds and resolution of  $< 0.02$  arcseconds, the completed assembly, seen in Figure 2.3, was installed in the summer of 2015, and has effectively eliminated the previous need for re-adjustment.



**Figure 2.3:** Top: completed RX assembly. Unseen: motor controller board, signal processing/routing board, PIRX. Bottom: after installation

# Chapter 3

## Proving APOLLO's Worth: The Absolute Calibration System (ACS)

### 3.1 ACS Concept

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 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  $\Delta t$  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  $\Delta t$  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. The ACS delivers select pulses out of the Cs-locked 80 MHz train at  $\sim 1.5 \frac{\text{photon(s)}}{\text{pulse}}$  directly to the APD (no lunar trip). Only those ACS pulses that coincide with expected fiducial and lunar gates are permitted to illuminate the APD.

Thus, we deliver an optical clock to the 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



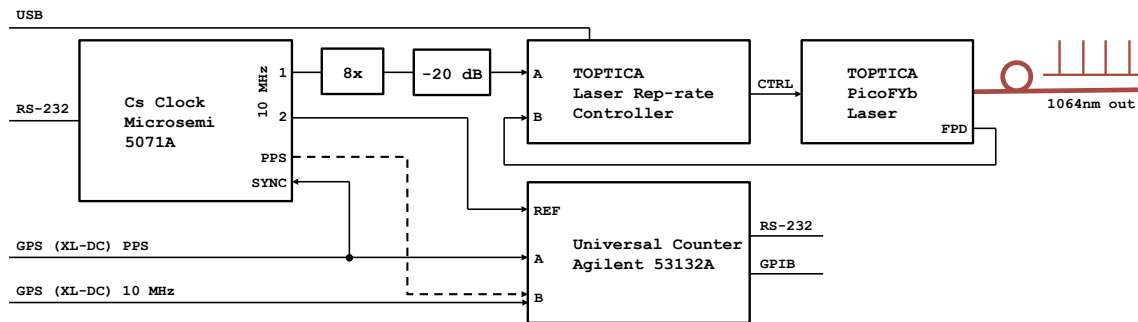
of systematic error. Naturally, it is important that we make every attempt to do so, but we can still realize a transition from precision to accuracy (in a more guaranteed manner, in fact) by employing this overlay technique—now standard practice in APOLLO data acquisition.

## 3.2 ACS Overview

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 APD delivery optics. Block diagrams for these three sub-systems are shown in Figures 3.1, 3.2, and 3.4, respectively. 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.

### 3.2.1 ACS laser and clock overview

The ACS Enclosure (Figure 3.1) houses the Cs frequency standard (Microsemi 5071A; stability can be seen in 4.1), 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.



**Figure 3.1:** 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.

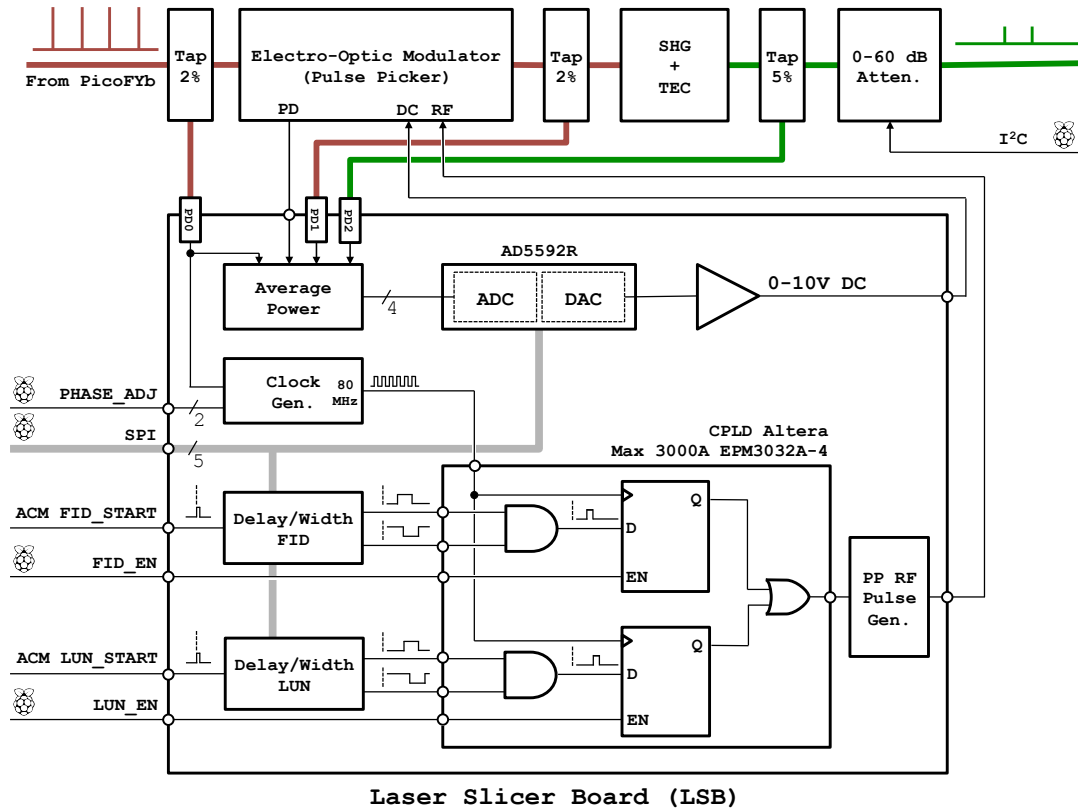
The ACS laser is a turn-key system that can be operated remotely. It is a customized TOPTICA PicoFYb that is diode-pumped and uses an all-fiber (Ytterbium-doped) configuration based on a ring oscillator with a semiconductor saturable absorption mirror (SESAM). The laser pulse width at the output of the 12 m optical fiber is  $\sim 10$  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 3.5.


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  $1 \mu\text{Hz}$  resolution in  $\sim 10$  s gates. In an alternative configuration (requiring a cable exchange, see dashed line in Figure 3.1), the time difference between the PPS signals from the two clocks is measured to a resolution of 0.1 ns. The results of this comparison are summarized in Section 4.1.

### **3.2.2 ACS Instrument Monitoring and Communications**

In addition to heading-up the effort on ACS-to-APD delivery optics mounting design, which is summarized in Section 3.4, I was placed in charge of software packages related to communications with the Cs clock, UC, and ACS enclosure's digital temperature controller (DTC). We monitor the Cs clock status on a ten minute cadence primarily to keep watch (and keep a log) for potential fatal errors, of which one had already occurred near the end of 2017 and since handled. We additionally have the ability to remotely command the Cs clock to do a variety of things (checking internal Cs logs remotely after the fatal error was quite useful), primarily to occasionally re-sync its PPS signal to that of the GPS-disciplined XL-DC clock for long-term stability of the Cs clocktrain phasing.

The DTC is polled on a two minute cadence to track temperature swings over time and correlate with other log data as needed to identify anomalies or features of interest; We also have the ability to remotely adjust parameters governing how the internal PID controller controls the temperature of the enclosure. Both the Cs clock and DTC are connected via RS-232 to a terminal server that we command and monitor via telnet. Additionally, the UC can be remotely commanded via a GPIB line to change measurement parameters, if the need ever arises.



**Figure 3.2:** 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  $\sim 1$  ACS  $\frac{\text{photon(s)}}{\text{pulse}}$  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<sup>2</sup>C or SPI).

### 3.3 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 3.2). 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  $\sim 1.5 \frac{\text{photon(s)}}{\text{pulse}}$  to the APD array. The EOM is controlled by the Laser Slicer Board (LSB), which monitors gate (one for fiducial returns, one for lunar returns) request signals from the ACM, 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 3.2). 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.

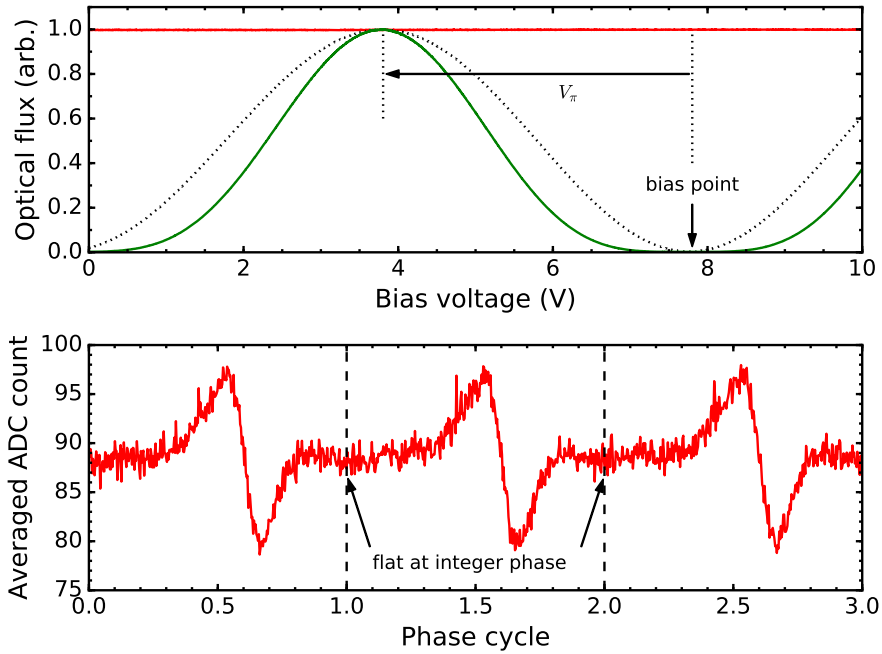
The EOM (Photline NIR-MX-LN-10-PD-P-P) is a fast (10 GHz) Mach-Zehnder interferometer in which the index of refraction of one arm, and therefore the relative phase of the light in the two interferometer arms, can be controlled by an external voltage. A change in bias voltage of  $V_\pi \sim 4 \text{ V}$  will toggle the EOM from destructive to constructive interference, with an

extinction ratio of 30 dB (on top of the  $\sim 5$  dB insertion loss). The quadratic behavior of the SHG results in  $\sim 60$  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 10 V and monitoring the average laser power after the EOM (PD1). Figure 3.3 (top) shows an example sweep. The value of  $V_\pi$ , 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  $V_\pi$  (manual potentiometer on the LSB) in order to flip the EOM state from blocked to open for one or more pulses.

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 3.2.



**Figure 3.3:** 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 8 V minimum and switch by amplitude  $V_{\pi}$ . 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.


Considering a single gate type for what follows, a pulse-request window is created by first stretching the trigger pulse to  $\sim 120$  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 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 Fig. 3.2). The direction of the phase adjustment is determined by the high/low state of one GPIO line, (LEADLAG) from PIRX, while a phase step is triggered by a rising TTL edge on the second GPIO line, DELAYCTL. After 1536 rising edges, the clock phase cycles by 360° (0.234° of phase per step). These lines are labeled PHASE\_ADJ in Fig. 3.2. 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 3.3 (bottom).

Similar to the Cs clock and universal counter, I took charge of communications using serial, SPI, and I<sup>2</sup>C protocols to the variable attenuator, adjustable delay chips, the clock multiplier/divider for phase adjustments and the ADC/DAC chip for power monitoring and pulse picker DC bias setting.






**Figure 3.4:** 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 3.2) establishes single-photon illumination of the APD array. The alternate injection site is discussed in the text.

### 3.4 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 3.4 shows the optical system that does this. First, the light exits the fiber and a lens creates a free-space collimated beam. Next, a  $45^\circ$  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  $45^\circ$  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.



**Figure 3.5:** Zoomed-in view of the receiver/optic mount design at the primary injection site. The receiver wall is forced to be transparent, to better view the optic-side hardware as well. In view is the custom optic and its mount, the fiber collimator and 45° dichroic mirror, and the neutral-density filter. The larger drawn circle whose normal is parallel to the receiver wall represents the primary mirror profile. Additionally, the drawn small silvered patch in the center of the custom optic can be seen.

I was placed primarily in charge of design, fabrication and modification of the various optical mounting hardware, custom optic design and receiver wall housing to accommodate the newly added components. Of critical importance was designing the mount to hold the ACS-coupling optic such that it could fit within the baffles of the receiver housing while not blocking any APOLLO signal photons. This limited the ability to expand the mount along the horizontal direction to easily handle the aforementioned constraints, adding to the challenge of placement for the spring plungers used to press the optic against one side of the internal portion of the mount. The design can be seen in Figure 3.5, while pictures of the completed/installed product can be

seen in Figure 3.6.




**Figure 3.6:** Left: The completed custom optic and mount, as seen from the APD’s perspective. Notice at this angle the small silvered patch looks exactly circular. Not seen are the 3 spring plungers (2 at the base, 1 in upper left corner from this view) on the opposite side to secure the optic while still avoiding obscuration of the primary mirror profile. Right: Overhead view of the installed custom optic on a magnetic kinematic mount. Bounding it on top/bottom are baffles within the receiver housing, creating the need for bevelled edges on those sides. To the right, the 45° dichroic mirror mount and the edge of the fiber collimator can be seen.

### 3.5 ACS Laser

The setup of the TOPTICA PicoFYb laser is depicted in Figure 3.7. 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  $\lambda = 1064.29 \text{ nm}$ ,  $\Delta\lambda \approx 0.191 \text{ nm}$  (FWHM) and a pulse width of  $\sim 10 \text{ ps}$ —as measured via an autocorrelator using a  $\text{sech}^2$  fit. An



**Figure 3.7:** Schematic diagram of the ring-shaped fiber oscillator. See text for details, and Ref. [1] for patent information.


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 3.1). 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 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 duration of 10 ps—longer than a more natural  $\sim 2$  ps for this type of laser—was chosen in order to minimize nonlinear effects in the 12 m PM980 fiber delivery following the laser. In particular, self-phase modulation (SPM) could have a detrimental effect on the efficiency of the second harmonic generation (SHG) stage that follows. After the fiber delivery, the pulse exhibits a slight temporal lengthening of about one picosecond to  $\sim 11$  ps, and a spectral broadening by a factor of 2 to  $\sim 0.4$  nm.



**Figure 3.8:** 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.

The pulse picker introduces an optical loss of about 5 dB. Therefore a fundamental power of  $\sim 11$  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 3.8). 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  $\frac{\text{photon(s)}}{\text{pulse}}$  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.

**Table 3.1:** 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

As implemented in the ACS system, we find approximate laser power and pulse energy levels as detailed in Table 3.1. The three taps shown in Figure 3.2 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 of ten of effective attenuation. Together with other losses and detection efficiencies, we end up capturing approximately one  $\frac{\text{photon(s)}}{\text{pulse}}$  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.

This chapter, in part, has been excerpted from the material as presented in Classical and Quantum Gravity, 2017. 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.; Zach, A. [19]

# Chapter 4

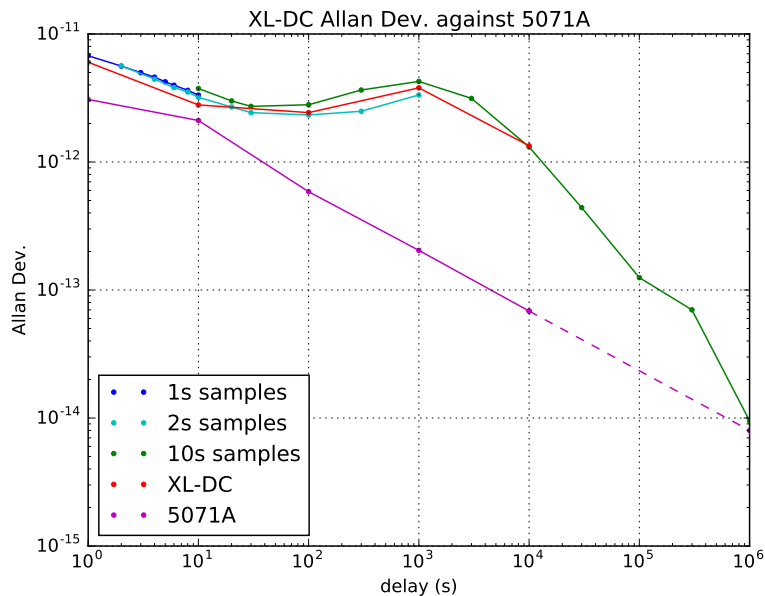
## Early ACS Results

### 4.1 Clock Comparison

The aforementioned 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). Allan deviation measurements of the Cs clock, and the XL-DC as compared against the Cs, can be seen in Figure 4.1.

We initially set up the cesium clock and clock frequency comparison monitoring in February of 2016, before fully implementing ACS capabilities in September of 2016 and then switching to using the cesium clock as APOLLO timebase at the start of 2017. For this short time period of ACS operation prior to switching the clocks, I put together a method of correcting range photon timing for discrete laser shot periods from the clock comparison data, using a rolling time-average over frequency differences during run times, incorporating information regarding XL-DC self-applied voltage corrections. This method (again, for this time period only) is used in

conjunction with ACS-derived corrections, described in Section 5.1. A separate effort, led by Yan Liang, developed a method for back-correcting a decade of APOLLO data based on these statistics. This effort, and a detailed comparison of the two clocks, is described elsewhere [26].



**Figure 4.1:** Clock Allan deviation measurements. Cs clock statistics (magenta; lower), provided by the manufacturer, were measured against a hydrogen maser. The dashed line projects the measurements to longer times. The XL-DC Allan deviation was measured with the UC at 1 s, 2 s, and 10 s averaging (blue, cyan, green, respectively), using the Cs clock as a reference. To estimate the Allan Deviation of the XL-DC clock alone, the red points show the quadrature subtraction of the Cs clock deviation (magenta) from the average sampled data. Assuming that the Cs clock projection is accurate, the Cs clock is more stable than the XL-DC for time intervals less than 2 weeks. Beyond that, the GPS-based disciplining of the XL-DC improves its stability relative to the free-running Cs standard. The degradation of the XL-DC stability at  $10^3$  s is likely due to atmospheric fluctuations that burden the GPS solution.

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  $3.0 \times 10^{-13}$  due to gravitational redshift. Our Cs clock was measured by the manufacturer to have a static frequency offset of  $-1.3 \times 10^{-13}$  relative to a hydrogen maser standard, for a net frequency offset from the



UTC second of  $1.7 \times 10^{-13}$ . The resulting phase accumulation should be roughly 15 ns per day. Indeed, we measure a  $\sim 14$  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.


## 4.2 Example ACS run and products

This section illustrates initial gains delivered by the ACS, while providing details on the techniques by which the results are realized. The following section provides a first-look into the longer-term behavior of the ACS calibration measures. An updated method of studying ACS gains and its results are presented in Chapter 5.

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. Of these, we get roughly 11,000 useful ACS photons for each gate type, while about 5,000 and 3,000 photons contribute to the fiducial and lunar return peaks, respectively. 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 4.2, 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. But we know *when*, in TDC-space, any lunar photon should appear. Those arriving at the


expected time are colored red in Figure 4.2.



**Figure 4.2:** 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. Black dots, laying within the yellow stripes, are events that are considered “ambiguous”—we cannot be certain if they are an actual lunar event, or an ACS event. Blue dots are the remainder—largely representing 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.

Note that the ACS photons (yellow) in Figure 4.2 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 3.3). 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 4.3). The number of stripes appearing in Figure 4.2 depends on the adjustable width of the pulse selection window as described in

Section 3.3. 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.



**Figure 4.3:** Relative timing between 50 MHz XL-DC clock and 80 MHz ACS laser pulses. Absent frequency shifts between the signals, 5 possible alignments of the 80 MHz laser pulse train are exhibited relative to any given 50 MHz rising edge (then the pattern repeats). The resulting comb of laser spikes has a 2.5 ns spacing.

The masking of ACS photons in the left panel of Figure 4.2 looks five times worse than it really is, because only one in five stripes is “live,” or possible, for a given shot. This is another important phenomenon to understand—tagging photons as ACS relies on a similar piling-up operation analogous to what is done for the lunar range photons. Photons are piled according to their phasing relative to a particular 50 MHz clock event, such that we visualize events arriving at 12.5 ns intervals easily, as seen in Figure 4.4. Fits are performed to these strong signals, and any photons within  $2.5\sigma$  of their centroids are tagged as “likely ACS.” Any photons within the same range of the lunar signal’s centroid are tagged as “likely lunar.” Any with both tags are the ambiguous photons in Figure 4.2, and the remainder are tagged as background. Performing the sorting in this manner, rather than relying on the ACS comb to tag ACS photons, allows us to minimize the ambiguity between ACS and lunar photon identification.

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 any lunar-gate ACS photons were captured 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 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 ( $\delta_{\text{shot}}$ ), 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 4.1. 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 4.5). The range error is the Gaussian fit centroid, expressed in millimeters. The estimated error of the centroid,  $\mu$ , is simply the Gaussian standard deviation  $\sigma$  divided by  $\sqrt{N_{\text{pairs}}}$ . As is apparent in

**Table 4.1:** APOLLO range errors are small, and can be determined precisely depending on how many channels and shots are considered.

Channel	$\delta_{\text{shot}}$	$N_{\text{pairs}}$	Range error (mm)
16	0	74	$-1.70 \pm 3.58$
16	5	730	$-0.13 \pm 1.09$
1–16	0	826	$-0.55 \pm 1.01$
1–16	5	8918	$0.30 \pm 0.32$
1–16	9	15234	$0.47 \pm 0.24$

Figure 4.5, the fit is rather good. Reduced chi-squared ( $\chi^2_\nu$ ) measures for all fits used in the present analysis are statistically probable. The spread seen in ACS time differences in Figure 4.5, 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., [27, 20]), 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  $\sim 1000$  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 vs. 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 4.1. 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 4.1 of 0.355 mm, the  $\chi^2$  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.

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 4.2. 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  $-4.3 \times 10^{-12}$  fractional frequency offset. Between shots 7000–9000, we see a reversal of 16 TDC bins, translating to a rate of  $4 \times 10^{-12}$ . Across the lunar round-trip time of 2.5 s, we might therefore expect a time measurement error of  $-4.3 \times 10^{-12} \cdot 2.5 \times 10^{12} \text{ ps} \cdot 0.15 \frac{\text{mm}}{\text{ps}} = -1.6 \text{ mm}$  and 1.5 mm, respectively (a 3.1 mm difference). Table 4.2 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.

**Table 4.2:** Clock drift is easily seen/recovered ( $\delta_{\text{shot}} = 5$ , all channels).

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

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 ( $2.5 \text{ s times } 20 \frac{\text{shots}}{\text{s}}$ ), 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 4.3 indicates statistically consistent “other” offsets with a weighted mean of  $1.44 \pm 0.34 \text{ mm}$ . 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 4.2 values to the weighted mean from Table 4.3, we verify that the 1–7000 shot range has a clock-induced offset of  $-1.75 \pm 0.51 \text{ mm}$  and the 7000–9000 shot range has an offset of  $1.25 \pm 0.81 \text{ mm}$ : again consistent with expectations from the independent measurement of clock frequency ( $-1.6$  and  $1.5 \text{ mm}$ ).

A final test shows that the ACS can nail down differential channel offsets in ten minutes,

**Table 4.3:** 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	$N_{\text{pairs}}$	Range error (mm)
1–7000	5992	$1.65 \pm 0.38$
7000–9000	1749	$0.68 \pm 0.72$

**Table 4.4:** Differential channel offsets (partial set; channel 15 is not used).

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

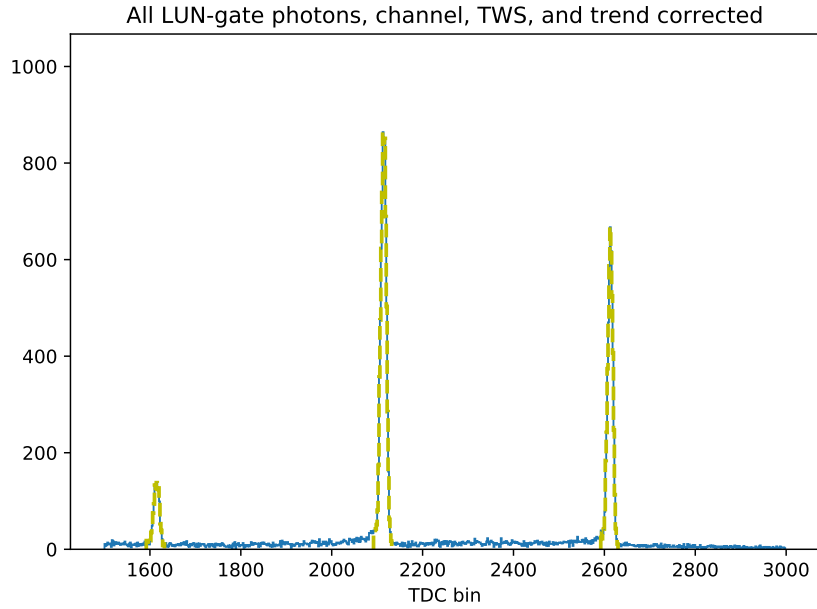
compared to months in the past. In a single ACS-LLR overlay run, we get an independent measure of two types of channel-specific timing offsets: large offsets (1 ns scale) common to both fiducial and lunar gates, due to essentially static delays that are different for each channel; and differential offset 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; More on this in section 5.2.

Table 4.4 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 now have a method to correct our now-characterized few-millimeter-scale errors, as discussed in Chapter 5.





**Figure 4.4:** The apparent comb in TDC space has features every 2.5 ns, or  $\sim 100$  TDC bins, while the *true* comb is spaced at 12.5 ns, or  $\sim 500$  TDC bins. This plot shows the reconstructed true comb via information from the “time within seconds” (TWS) counter (in the ACM). We can determine the phasing of each shot, using the knowledge that ACS 80 MHz pulses appear at discrete phases relative to the 50 MHz clock train. This, combined with the TDC number for each event, result in a histogram of photons spaced at 12.5 ns. All LUN-gate photons shown have had relative channel offsets, UC trend data, and TWS phasing applied to sharpen-up pulse alignment. Gaussian fits to the signals are displayed as red dashed lines. For the ACS delay width used in the pulse processing for this run ( $\sim 20$  ns wide), we are guaranteed to get either one or two ACS pulses per shot. For a given 50 MHz clock trigger (commanded by the ACM), sometimes the second pulse captured is early relative to the middle ACS pulse (captured every shot) and sometimes it is late—this is the reason for the appearance of *three* ACS pulses in this plot. This run tended to capture the later pulse more frequently than the earlier pulse, as evidenced by the relative strength of the signals.



**Figure 4.5:** Example histogram corresponding to the final row in Table 4.1. 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.


### 4.3 Range error statistics

Section 4.2 highlighted key analysis schemes available using the ACS. Chapter 5 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. Here, we apply the techniques from Section 4.2 to all APOLLO operations since 2016 September 12 (through 2017 April 3) to understand overall performance and stability.

Figure 4.6 summarizes the results. The top panel corresponds to the measure presented in Table 4.1, using  $\delta_{\text{shot}} = 5$  and all channels. The middle panel corresponds to the “simultaneous” measure of Table 4.3—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 4.6 appears to be much tighter.

Figure 4.6 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  $1.2 \times 10^{-11}$  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 4.1 and treated extensively in a separate paper [26]. The reduced  $\chi^2$  measure for the Cs clock-drift-only distribution indicates no spread



**Figure 4.6:** Channel-aggregated range offsets (0.5 mm bins); Blue: GPS-based; Red: cesium-based data (darkest is overlap). The top histogram compares ACS fiducial to lunar photons having the same shot number and is an “everything” 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.

beyond the uncertainties contributed by the ACS statistics (corresponding to errors in Tables 4.1, 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 4.6). 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 4.4 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 [26]. We also see static offsets at the few-millimeter scale, which is now routinely corrected for as best and accurately as we can, for each run contributing to APOLLO range data. This, and further study of similar APOLLO performance measures, are detailed in Chapter 5.

## 4.4 Initial 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  $\sim 3$  mm of random error—about double the previously-reasoned  $\sim 1.5$  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 [26]. 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.

This chapter, in part, was excerpted from the material as presented in *Classical and Quantum Gravity*, 2017. 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.; Zach, A. [19]

Figure 4.1 was excerpted from the material as presented in *Classical and Quantum Gravity*, 2017. Liang, Y.; Murphy, Jr., T. W.; Colmenares, N. R.; Battat, J. B. R.; [26]

# Chapter 5

## ACS Corrections to Range Data

Chapter 4 described a method of ACS study we now colloquially refer to as the “local  $\Delta t$ ” method, referring to its reliance on comparing nearby photons, and sensitivity to intra-run timing behavior (explicit study of clock drift). Although the local  $\Delta t$  method of extracting timing offsets is informative and powerful in its own right, we decided a different pursuit would end up being more descriptive of TDC-space dependent timing variations, and more sensible to use in the context of range data reductions.

In summary, the story of the ACS up to this point has been:

- We desire to calibrate our ranging data.
- To accomplish that, we want to have a system sending optical truth information to our detector.
- To find out how much bias APOLLO exhibits, we ask how much the timing system alters the truth.
- Results show our instrumentation is distorting the truth, but now we can measure the magnitude.

- How do we use this knowledge to recover truth timing, consistent with the processing of range results?

This chapter seeks to answer the last question in the list by illustrating a newly formulated method of utilizing ACS photons to correct range photon round trip times. One important note about nomenclature in the following: “offsets” refers to the timing errors that are actually *measured* by the ACS, whereas “corrections” are the adjustments to compensate for the offsets; One is the negative of the other.

## 5.1 ACS Corrections Concept

The ultimate goal in utilizing ACS corrections is to get an accurate measure of elapsed time between range photon signals (FID/LUN). We first present a summary regarding the timeline of APOLLO events and how range results are formed, to motivate consistent application of ACS corrections. Then the concept of how the correction scheme works will be explained, and why we can trust what it tells us.

### 5.1.1 APOLLO Range Photon Timing

Figure 5.1, a rough timeline of APOLLO events, shows an idealized scenario wherein a fiducial range photon and a corresponding (same shot) lunar-gate return photon are present in the same channel, with the ACS pulse-train overlaid. The laser fires, reaching the FPD (fast photodiode) a short time later before crossing the telescope’s axis intersection, the point at which a lunar prediction time for photon time-of-flight is based (used for gate scheduling and data processing). A time,  $T$  (the prediction time *if* our prediction were perfect), later, the laser pulse again crosses the same point after returning from the Moon. Shortly thereafter the light reaches the APD which produces an ECL START signal eventually seen by the TDC. Prior to receiving the returning lunar photons, the FPD signal is cable-delayed such that it

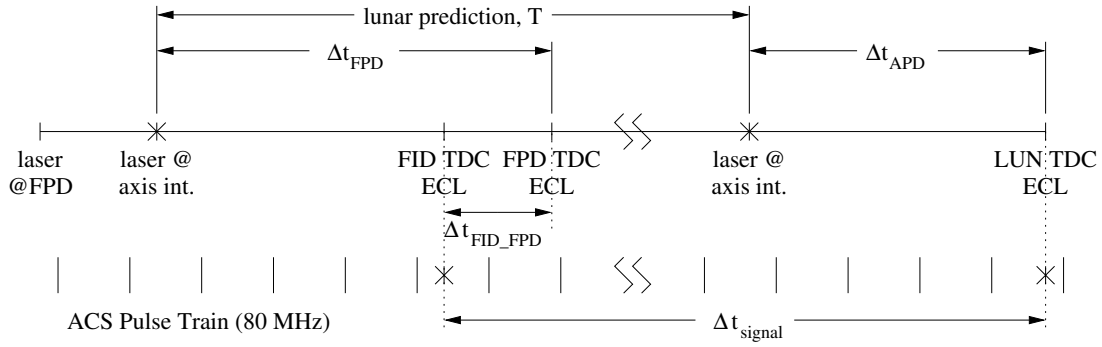


reaches the APD within the fiducial gate (the same gate that captures photons returning from the local corner cube). Recall ADC/TDC channel 15 is reserved specially for receipt of the FPD signal, every shot. The differential measurement, then, is referenced to the time when the range laser beam crosses the telescope axis intersection, which is insensitive to telescope motion, to avoid unnecessary complication in the equations of motion. This  $\Delta t_{\text{signal}}$  is analogous to  $\Delta t_{\text{APD}} + T - \Delta t_{\text{FPD}} + \Delta t_{\text{FID-FPD}}$  in the figure.

The FPD acts as a timing “anchor,” as it is guaranteed to exist for every laser shot, and has low jitter due to it timing a many-photon signal while employing a constant fraction discriminator. The predicted lunar time is based on the FPD, and as such, the gate positioning. However, in order to form a truly differential measurement, we need a reference optical signal that takes the same path through the telescope as the lunars and uses the same detector and timing electronics. This is where the fiducial signal comes into play, while also informing us of APD/TDC channel offsets, which the FPD cannot address.

What we do in practice to arrive at a measure of  $\Delta t_{\text{signal}}$  is not quite the same, as we are not guaranteed to get a returned FID or LUN photon for each shot. Even if there are both FID and LUN records in the same shot, a truly differential measurement would need matching channels also, decreasing the likelihood of a match. Therefore, in order to achieve the level of required statistical precision, some aggregation over shots and channels is necessary.

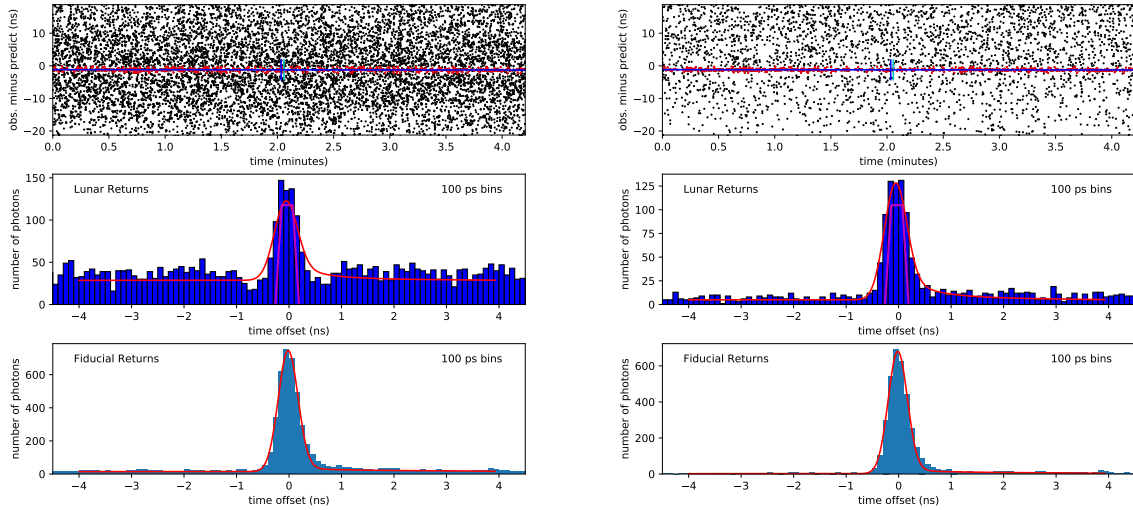
The difference in time between the FPD record (contained in *every* fiducial gate) and fiducial photons is a (different) value for each channel, but is found uniquely for each run by fitting a histogram of the difference in their TDC timestamps. This allows us to reference range photons against the FPD record reliably - the resulting product is a channel/shot-aggregated histogram of fiducial photon timing, the width of which is primarily due to the laser pulse width and jitter of timing electronics. A modified Gaussian functional fit, asymmetric due to detector crosstalk and photon diffusion into the conduction band, is performed on the histogram, and the parameters recorded for later use on the lunar photons.



**Figure 5.1:** The FPD signal, a proxy measure for laser fire time, tells us when to expect lunar return photons to hit the detector, when used in conjunction with a predicted round trip time ( $T$ ).  $\Delta t_{\text{FPD}}$  is defined as the time between the outgoing laser pulse reaching the axis intersection and the receipt of the cable-delayed FPD START signal at the TDC. Similarly,  $\Delta t_{\text{APD}}$  is defined as the time between the returning laser pulse at the axis intersection and receipt of the LUN\_START signal at the TDC. After the laser is fired, we record the TDC value of the FPD record,  $\text{TDC}_{\text{FPD}}$ , the digital time since the start of the fiducial gate. We call the absolute time of this  $t_{\text{FPD}}$ , and expect the lunar photons to register as LUN\_STARTS at time  $\Delta t_{\text{LUN}} = t_{\text{FPD}} + T - \Delta t_{\text{FPD}} + \Delta t_{\text{APD}}$ . Note that the TDC ECL signal is not the same as the recorded TDC time, as some processing time exists between the two that is measured by a subroutine called CALTDC, which is discussed later on. We know to a high degree of accuracy how long it takes light to travel from the telescope axis intersection to the fiducial cornercube which allows us to make a differential timing measurement between the outgoing laser pulse (at the axis intersection) and returning laser pulse (again at the axis intersection) — effectively *measuring* the value of  $T$ .

Round trip times for returning lunar photons are formed by referencing each lunar photon against the timestamp of that shot’s FPD record, the number of 50 MHz pulses between the lunar and fiducial gate, and the centroid of the FID-FPD histogram, including the appropriate channel offset correction (for that lunar photon). The fiducial functional fit is convolved with a trapezoid representing the reflector timing width and tilt at the time of ranging, given the total libration (lunar and Earth topocentric correction), into a functional fit to describe the histogram of lunar round trip times minus their predicted times [3]. A one-parameter fit is performed by sliding the functional fit over the lunar signal until best-overlaid, resulting in a fine-scale timing correction relative to the prediction, which can be seen in figure 5.2.

Recall though, that the ultimate goal is to use the ACS information to correct  $\Delta t_{\text{signal}}$



**Figure 5.2:** Left: plot of range photon returns for LLR+ACS run 190416k—a more recent run with strong lunar return signal, with a duration of 5000 shots ( $\sim 5$  minutes long). The bottom subplot is the histogram of fiducial photons fitted with the aforementioned asymmetric functional fit, representative of system performance. The upper subplot is a scatter plot of the lunar photons relative to their predicted return times. A line is fit to this, and lunar photon times relative to this fit appear in the middle subplot, fitted by the fiducial fit convolved with the trapezoidal reflector profile based on lunar libration, in magenta. Right: the same run, after ACS photons were identified and subtracted.


appropriately. How, then, do we stitch together the FID/LUN gates using the ACS pulse train (80 MHz)? This first relies on an assertion that has, for the most part up to now, remained unspoken: APOLLO timing measurements misrepresent the truth of ACS photons to us in the same manner as timing measurements of FID/LUN photons, since, within each gate type, the detector and electronics have no knowledge of which photons are which. We then envision a scheme wherein we correct timing distortions to guarantee ACS pulses are reported at integer multiples of 12.5 ns. In doing so, we could then take any individual FID or LUN photon timestamp and apply a correction to it based on where it lands in TDC-space, before the FID histogram is formed and round trip times are calculated. Applying corrections at the start of data processing ensures all subroutines used in the formation of a normal point are using the most accurate

information available.

Given the desire to parallel the main reduction process as much as possible, which involves some TDC-dependent corrections, we first attempted to adapt the local  $\Delta t$  method, described in the previous chapter, to focus on ACS photons close in temporal proximity to the actual FID/LUN signals. However, we were left starving for adequate statistics in order to achieve the desired level of precision for correction estimates. As such, we cannot continue to rely on the local  $\Delta t$  method. This created the need for developing a different method of describing TDC-dependent corrections.

### 5.1.2 Correction Scheme

In order to motivate the newer method, we return to a discussion of the comb of ACS photons first shown in scatter-plot form in Figure 4.2. In order to assign a correction to a photon timestamp from TDC-space, we of course would need our calibration photons to probe the appropriate region within TDC-space. As previously mentioned, photon return times are asynchronous relative to the 50 MHz clock such that their distribution in TDC-space is uniform. However, we can control the width of the distribution to 20 ns by adjusting gate positioning relative to the closest clock tick in reference to the predicted return times. The upper plot in Figure 5.3 shows a channel-aggregated example of the return distributions for both the FID and LUN gates, *without* adjusting for static differences between channels. The middle plot shows the same distributions, but with the FID and LUN ACS photon combs overlaid on top, mostly overlapping the range returns and thus probing the same region of TDC space. The bottom plot shows the same information, now accounting for relative channel offsets, resulting in a tightened-up histogram of comb teeth (and range returns, though barely discernible on the scale of ACS tooth photon counts). An iterative procedure to perform least squares fits of basic gaussian forms to each tooth is undertaken, and the fitted parameters are recorded for later reference and use in forming ACS-derived timing offsets.



**Figure 5.3:** Histograms of TDC-space range + ACS returns for run 190416k. The top and middle plots do *not* include relative channel offset corrections, while the bottom *does*. The programmable delay windows have been expanded to add 2 more comb teeth than earlier operation (e.g. as can be seen in the left panel of Figure 4.2), and the positioning slightly shifted, to ensure the range signals are completely encompassed by ACS data on a routine basis.

To the extent we trust that the ACS comb teeth represent calibration timing accurately, further motivated in Section 5.1.3, we can inform ourselves of TDC-dependent timing structure by using the measured tooth positions (against a reference tooth) as compared to expectation (exact  $n \cdot 2.5$  ns intervals). Gaussian fits to comb teeth are performed on a per-channel, per-gate-type (FID; LUN) basis, allowing us to describe timing offsets at each tooth position. Without ACS information between these points, the best we can do is to suggest that offsets within said space lie on an interpolated line between the adjacent offsets, resulting in an offset “mapping.” The idea, then, is to apply corrections at each tooth location in TDC-space to position the ACS comb teeth directly on top of the expected truth intervals; We then apply the same corrections (albeit

interpolated between) to the FID/LUN range photons, because we know they should be affected by timing offsets in the same manner as ACS photons. Offset estimates for regions outside the TDC values bounded by the ACS comb see an extrapolation argument based on a characteristic interpolated line slope, discussed in section 5.2.

**Table 5.1:** Example calculation of timing offsets from fitted teeth centroids.

Quantity description	Example values, ns				
Fitted tooth centroids	80.50	83.08	85.43	88.05	90.52
Fitted centroids, relative to reference	0.00	2.58	4.93	7.55	10.02
Expected centroids, relative to reference	0.00	2.50	5.00	7.50	10.00
Fitted - expected centroids (raw offsets)	0.00	0.08	-0.07	0.05	0.02
Mean subtracted offsets (final offsets)	-0.016	0.064	-0.086	0.034	0.004

So, how do we actually accomplish finding and utilizing ACS corrections? Table 5.1 shows a numerical example of the scheme, and Figure 5.4 graphically illustrates the concept. Timing offsets are formed at each tooth position, calculated as the deviation from  $n \cdot 2.5$  ns of the separation from a “reference” tooth in the comb, where  $n$  is an integer. Notice that this implicitly forms a reference “grid” of exact 2.5 ns spacings, with one grid marker lined up at the position of the reference tooth (the center tooth of the comb in Figure 5.4). If APOLLO were an ideal system, we would expect the comb teeth centroids to lie directly atop each reference grid marker. In this picture, tooth spacing deviations (i.e., ACS-derived timing offsets) are analogous to the difference between each tooth position and the position of the nearest reference grid marker, resulting in an offset at each tooth position. The *unweighted* mean of the offsets are subtracted (using a weighted mean makes the resulting offset values dependent on which tooth is chosen as the reference), finally resulting in a non-zero offset at the reference tooth; The negatives of these offsets are the resulting corrections. One can interpolate between discrete offsets to form an offset mapping for that region of TDC space; This is cleverly referred to as the “interpolated mapping” method of applying ACS corrections.

Now we have a way of adjusting any photon timestamp that lands in the region of the ACS

comb! Recall the combs are scheduled in such a manner that they mostly overlay the range returns, as shown earlier. Any photons landing outside the boundaries of the interpolated mapping see an extrapolated line (whose value starts at the offset magnitude for the corresponding endpoint) based on a distribution of interpolated offset lines; More on this in 5.2.

A major question remains, however: are we able to derive ACS corrections to meaningful levels of uncertainty, given the statistics involved? Typical runs are about 5000 shots, and we run the ACS light levels at about  $1.5 \frac{\text{photons}}{\text{gate}}$  (FID/LUN both), over all channels. There are 15 channels taking range data (recall channel 15 is reserved for the FPD) and about 10 comb teeth present in TDC space typically. This means 7500 total photons (in one gate type) across  $15 \cdot 10$  elements, so about 50 photons per tooth, per channel. Widths of comb teeth tend to be about 5.25 TDC bins on average (a recent detector improvement results in tighter timing widths; New data is being progressively built-up corresponding to this newer hardware period to study using the methods described herein). A statistical reduction of  $\sqrt{50} \sim 7$  results in knowing the tooth centroid to the  $\sim \frac{3}{4}$  TDC bin level. Frequently during operation the ACS is turned off during the first portion of a run to ensure proper lunar range signal acquisition, bringing the effective ACS shot duration down from 5000 to  $\sim 4000$ , such that tooth centroid uncertainty becomes  $\sim 0.85$  TDC bins. Corrections rely on deviations of paired tooth separations from truth, so that we would only know corrections to the  $\sqrt{2} \cdot 0.85 \sim 1.2$  TDC bin level, given the RSS of the uncertainties from both teeth ( $\sqrt{2}$  assumes the uncertainties are  $\sim$  the same). In practice, we see applied photon corrections at the 1-2 TDC bin level, indicating that applied corrections could be pulled around at a level commensurate with the correction magnitude, resulting in low signal-to-noise corrections. We need to do better, especially since typical normal point (NP) uncertainties tend to be at the  $\sim 10$  ps, or  $\sim 0.4$  TDC bin level.

Perhaps we try aggregating across time, channels and/or gates in some consistent manner to improve the correction uncertainties to better levels? We can improve statistics over time by successively including more runs, so naturally we primarily pursue aggregation over time. Tooth

*separations*, not tooth centroids, are chosen as the measure of merit to collect statistics on, since they govern departure from “truth” (hence the need for a reference tooth). We ensure comparison of the “same” paired tooth separations over long timespans via a sensible tooth labelling scheme based on TDC position relative to fixed label “boundaries” spaced at 100 TDC bins ( $\sim 2.5$  ns). Tracking tooth separations effectively decouples the eventual aggregate offset measures from any absolute position in TDC-space, removing any long-term timing offset dependence that absolute tooth TDC-space positions may be sensitive to. Though there may exist temporal structure in the positions of comb teeth across time, it may still be possible to aggregate paired tooth separations accurately, assuming larger-scale structure ( $\gtrsim$  few TDC bin applied correction level) is common to all comb teeth.

We imagine a scheme wherein we take a weighted mean—assuming each collection of tooth separations is well-represented as such—over time, to represent each paired tooth separation as a single point over the entire collection of separations. The resulting product is an aggregated form of the tooth separations in principle delineating where teeth *actually* manifest themselves in TDC-space, on a relative scale; The ensuing pattern is affectionately referred to as a “parent” comb, one for each channel/gate combination. I.e., a particular parent comb tooth label is chosen to act as a reference timing position of “0” and positions of parent comb teeth can be built-out from there using the aggregate tooth separations, while simultaneously building a reference grid of exact 2.5 ns spacings, relative to the reference timing position. Recall at this point we have the means to extract timing offsets at the location of each parent tooth, or an “offsets mapping”, just by differencing the relative reference grid timings from the relative parent comb timings (or “positions”). Then, we can try to re-capture differences between unique runs, channels and gate types by sliding the parent comb until best overlaid with each of the run/channel/gate ACS combs, or “daughter” combs, to inform us of where the daughter comb *actually* is, based on the parent comb structure. We accomplish this by first placing the aforementioned reference time of “0” at the location of the daughter comb tooth whose label matches the parent reference label. Finally, a



one-parameter least squares fit is performed to shift the parent comb until *best*-overlaid onto the daughter comb, incorporating information on the uncertainties of each.

Figure 5.5 provides a preview of the final outcome of this process, in a similar format to Figure 5.4 for a single channel. The full collection of offset mapping points for each gate (for channel 14) are shown (though the reader will recall only those parent teeth matching the daughter teeth labels participate in the sliding fit) due to capturing earlier/later comb teeth separations for certain runs when building the parent comb; Differing sets of delays used for ACS pulse selection over the long aggregation timespan involved are the main reason for this. Indeed, this is primarily why the end points of the mapping exhibit significantly larger uncertainties as compared to the center region: a lesser total number of “end teeth” are available to aggregate over tooth separations. Note that the scale of extrapolated slope uncertainty is smaller than the end point uncertainties, so a photon timestamp must travel further away before seeing a more appreciable increase in the uncertainty due to the contribution from the extrapolated slope. It is pleasing to see that the structure of the offsets is fairly smooth, such that for the most part, there are no strong discontinuities between offsets.

Further details on the aggregate version of the correction concept is deferred until Section 5.2 in favor of first more explicitly motivating confidence in the validity of the correction scheme. What conditions need be met to use this interpolated mapping properly? I.e., how do we know what we are doing is accurate? Are the necessary conditions being met? The following section explores answers to these questions.

### **5.1.3 Establishing Trust in the Interpolated Mappings**

We must make sure that four main concepts are addressed in proper utilization of corrections:

1. Timing offsets affect ACS photons in the same manner as FID/LUN range photons.

2. The ACS timing in raw TDC-space accurately reflects the 12.5 ns ACS pulse period.
3. Consistency of ACS timing between the two gate types, thus propagating the time base over the 2.5 s round-trip interval.
4. Consistency of time domain (vs TDC domain) representation for both ACS and range photons.

Certainty in these four concepts may seem simple after explanation, but it is important to explicitly be aware of these issues and how they are handled.

The first point relates primarily to the physical interaction between the photons and the detector. If the ACS photons *were* of a different wavelength than range photons (they are not), the APD would respond differently, resulting in bias. Additionally, as previously mentioned, there is no way for the measurement system to know which photons are which and somehow treat them in a special manner.

The second point is relevant because ACS photons do not distribute themselves in 12.5 ns separated pulses within raw TDC space, but form a comb of 2.5 ns separations, as shown in Figure 4.3. Due to the high degree of accuracy of the clock, and absent a frequency shift between the timing of ACS photons (the 80 MHz optical clock) and the timing of TDC space (50 MHz clock)—now both based on the same Cs clock—we can trust the separation of comb teeth in raw TDC space is an exact multiple of 2.5 ns. For those early ACS runs when the XL-DC clock was still used as the timebase, we are able to correct for the accumulated phase difference due to a frequency shift between the clocks using the clock comparison data obtained by the Universal Counter in Figure 3.1. As previously described, to the extent that the spacing of paired teeth differ from  $n \cdot 2.5$  ns indicates a timing offset at the location of the comb teeth.


The third point then asks whether the ACS corrections are derived in a consistent manner between gate types. After all, we are referencing comb teeth to each gate type independently and thus use differing corrections. Well, ACS pulses do stem from a continuous pulse train between

the gate types, and the 80 MHz frequency does not differ meaningfully over that timescale (or the run as a whole, for that matter), given the accuracy of the clock. This fact, combined with the confidence that our raw TDC-space tooth separations of  $n \cdot 2.5$  ns can be trusted to reflect the timing deviations (from 12.5 ns) between ACS pulses, assuage any concerns over inconsistency of ACS timing between gates.

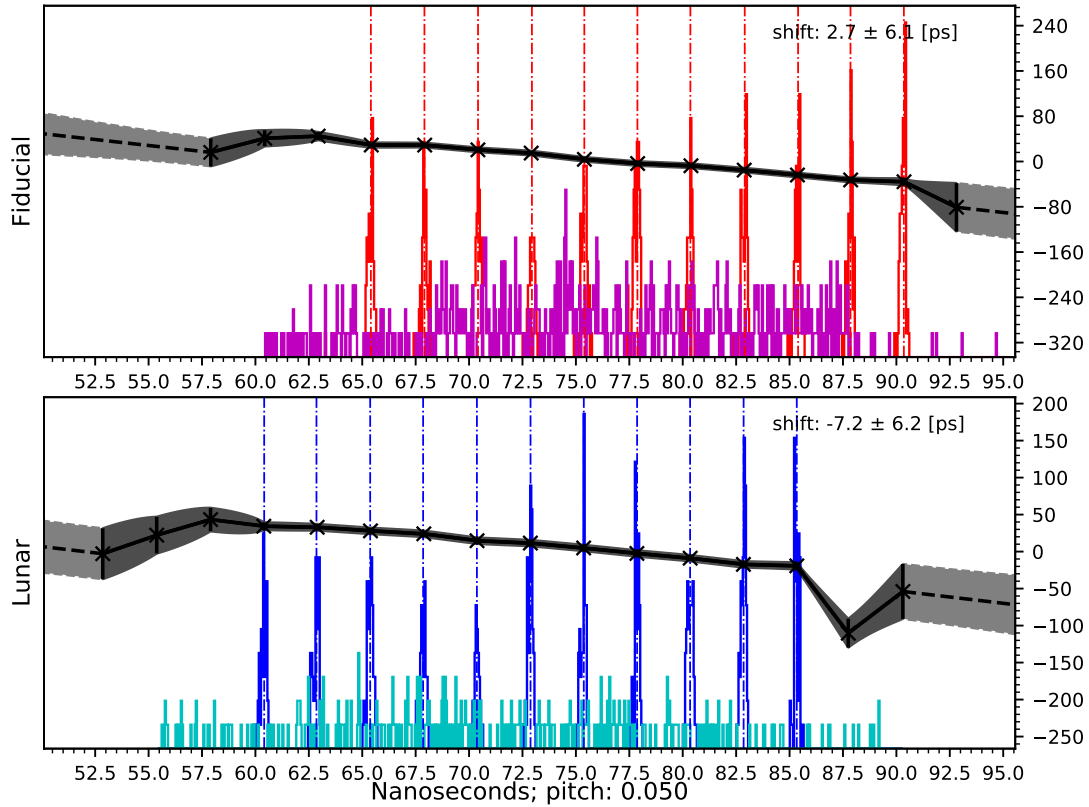
The fourth matter is crucial to ensure coherent timing between ACS and range photons. The concept to address it is simple: use the same conversion from TDC space to time space for the ACS photons as the range photons, and we can trust they are reflecting the same TDC timing. Once before and once after a run, a subroutine called “CALTDC” is called which aids us in understanding the rough-scale time structure intrinsically present in TDC space. This subroutine passes a single START signal to all 16 channels of the TDC, which also all see the same STOP signal, both derived from the 50 MHz clock. The result is a pair of pulses separated at integer multiples of 20.00 ns. One thousand events of five different pulse spacings (20, 40, 60, 80, 100) are sent to all channels, a quadratic functional fit is performed for each channel separately, and the parameters are recorded. This functional form is what is used to convert both the FID-gate and LUN-gate photon TDC numbers to time. The use of the CALTDC functional form takes some burden off of the ACS offset mappings by first removing rougher-scale TDC-only structure. The ACS offset mappings then reflect any structure from upstream timing electronics, the APD, and any residual TDC timing structure not removed by CALTDC. During application of the ACS offset mappings, we check to make sure the same CALTDC functional forms were used to form the mappings as what is being used to convert range photon TDC records to time.

Finally, we know we are defining and applying ACS corrections consistently by virtue of photon indistinguishability (to the detector), the definition of comb teeth to be spaced at  $n \cdot 2.5$  ns apart, the continuous pulse train and using the same CALTDC mapping between ACS photons and range photons. Without any one of these four pieces, the consistency would fail. If photons from differing sources were treated differently, bias would be introduced (though could be perhaps

studied if that *were* the case). Without the assertion/trust of  $n \cdot 2.5$  ns spacing in raw TDC space, we do not know how much ACS photons/comb teeth should be spaced by and thus could not derive corrections using this interpolated mapping. Without a continuous ACS pulse train, the current implementation for the ACS would not even work. It would likely have to be closer to the originally-envisioned ACS scheme of using a pulse-on-demand timing system. If we were missing the CALTDC mapping, we could convert to time using a straight  $25 \frac{\text{ps}}{\text{TDC bin}}$  conversion as long as we do that the same for ACS photons and range photons. Though of course doing so would be poor practice, given that we very likely can do better.



**Figure 5.4:** Cartoon diagram (not real data) of basic interpolation correction method concept. Reference grid markers are dotted vertical black lines, spaced at 2.5 ns exactly. Tooth centroids are vertical dash-dot lines of their respective gate color. Raw offsets are calculated as the narrow gray dash-dot line, defaulting to 0 at the center tooth (the chosen reference tooth) before the mean of the offsets is subtracted, shown as thick dashed gray lines. The solid black line, atop the dashed gray line for the valid span of range photon distributions, represents the final offsets that are utilized for corrections. A horizontal black dash-dot line is drawn to indicate 0-magnitude offsets. Note that the magnitudes of some offsets are *exaggerated* compared to real data, to better visualize the concept. The FID-gate mapping is drawn such that the reference grid perfectly overlays the fiducial range photon distribution. A deliberate shift between the LUN-gate reference grid and lunar range photon distribution, as well as a shift between the FID/LUN range distributions themselves, are introduced to indicate that the scheme can handle these imperfections.



**Figure 5.5:** Time-aggregated offset mapping for run 190416k, channel 14, with mapping offsets included. Interpolated offsets between each mapping point are shown as the solid black line. The darker fill-between region indicates the uncertainty throughout the interpolated portions of the offset mapping. Correspondingly, extrapolated portions of the offset mapping are indicated with dashed black lines, and a lighter-shaded fill-between area.

## **5.2 Aggregating the ACS Corrections to Meaningful Levels**

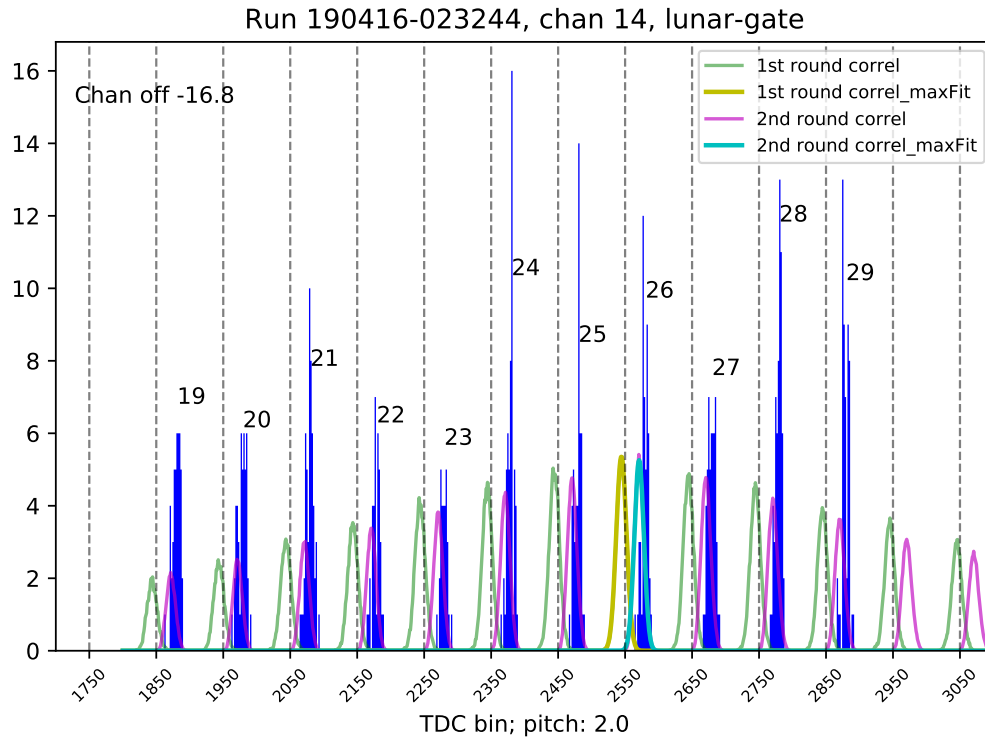
As the preview shown in Figure 5.5 suggests, we have found excellent stability in tooth separations over time, and know what a comb looks like for a given channel and gate type. Section 5.2.2 will clarify how this is accomplished and distilled into a usable result following the prescription discussed at the end of Section 5.1.2. Before diving straight into aggregating tooth separations however, there are a couple more concerns surrounding consistency in doing so—explored in the following section.

### **5.2.1 Channel Offsets and Tooth Labelling**

First we need to make sure we're comparing the "same" tooth pairings across time/channels/gates, accomplished by labelling the teeth appropriately and consistently. We assign tooth labels based on their TDC values: a tooth with a fitted centroid, adjusted by that channel's relative offset (to account for any potential hardware differences between runs), that falls between 2550 and 2649 would be labeled as tooth "26". A tooth between 2650 and 2749 is labeled as tooth "27", etc. i.e., the median value of each range (2600, 2700, etc) is divided by 100 to result in the appropriate label. There are label-checking subroutines in place to correct potentially incorrectly applied labels based on 1) the presence of more than one particular label for a particular channel and gate, and 2) what the labels are for the same tooth in the comb pattern on other channels; i.e., around the same channel-offset-adjusted TDC position.

The ability to accurately measure channel offsets is of great import for a variety of reasons in addition to consistent tooth labelling. As previously alluded to in Section 4.2, two types of channel offsets exist: one on a per-gate-type basis, mainly concerning timing alignment, and one on a differential-gate basis (comparing the same channel between gate types), which can result in differing range results between the channels; A very undesirable prospect. These "differential channel offsets", prior to the ACS, used to take months of collecting enough range data to yield

an acceptable measure of “truth” against which to compare range results on a per-channel (rather than aggregated) basis. The ACS can accomplish this in a single run, with flying colors.



**Figure 5.6:** Plot of ACS comb data and associated products for channel 14, run 190416k in the lunar gate, with its channel offset (printed at upper left) *already subtracted*. Tooth label boundaries (where the assigned label would change) are indicated by vertical dashed lines, with an arrow illustrating the assigned label to each tooth location. First-round correlation results against channel 1’s comb are plotted in green, with the maximum signal indicated by a Gaussian fit in yellow. After the “cleaned” reference histogram is formed, the correlation is performed again; The second-round results appear in magenta, with the maximum signal shown as the cyan Gaussian fit. Fluctuations in individual channel comb teeth can affect the determination of channel offsets from the first-round correlation such that when re-combined to form a new reference comb, the effective centroids of the new signal can vary slightly as compared with the first-round correlation results; This primarily contributes to differences in channel offset results between the correlation rounds. Regardless, using a “cleaned” second-round reference comb ensures a more well-behaved response in the correlation amongst all channels.

Channel offsets are found via a correlation of ACS comb data in each channel against a reference comb by sliding the histograms over each other until best-matched. In particular, this is a two-step process. One channel (usually channel 1) is chosen to act as a reference against which

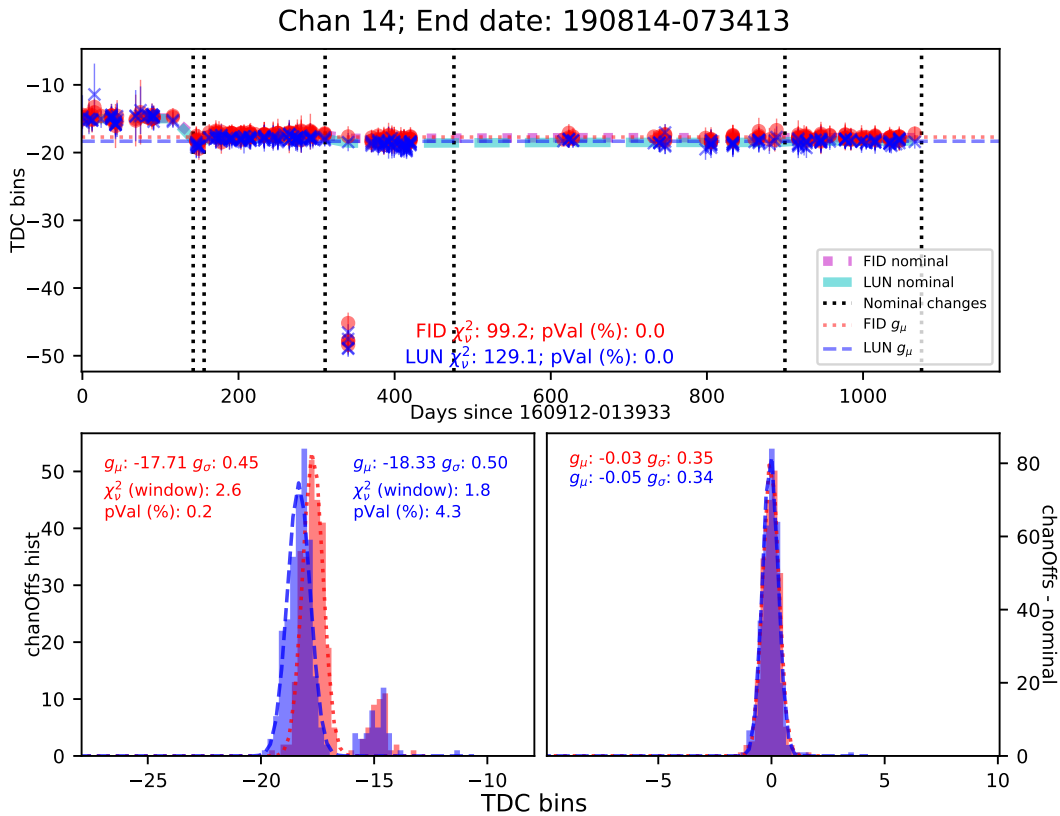


all other channel combs are slid against. The resulting maximum signal is fitted with a Gaussian, and offsets relative to the reference are recorded.

In hopes of reducing any bias dependent on variations in tooth strength or structure based on the original reference comb, a second-round correlation is performed. All channel combs are laid atop each other by applying the first-round offsets; Then for each x value, a median is taken of the corresponding y-values (frequencies), across channels, to form a “cleaned” comb histogram. The median operation has the effect of equalizing the strength and structure of comb teeth across channels, which we then use as the reference signal in the second-round correlation. The original reference channel also participates in this second-round correlation. Again, the resulting maximum is recorded and converted to a channel offset along with an estimate of uncertainty from the fit. Once all channel offsets have been found, a weighted mean (using the errors reported by the correlation maximum signal fit) over the offsets is calculated and removed from the raw channel offsets. A visual aid for the correlation and tooth labelling processes appears in Figure 5.6, for a single run/channel/gate.

The application of channel offsets to ACS combs during tooth labelling ensures the labels are insensitive as possible to any changes in the positioning of teeth within TDC space over time due to changes in the hardware (which usually results in differing channel offsets, and thus where teeth are manifested within TDC space). As such, deriving channel offsets accurately on a per-run basis for each gate type separately not only establishes consistency of tooth labelling over long timespans (and channels/gates), but also ensures we are doing our best to apply the most accurate ACS corrections that we possibly can to each run uniquely. Studying these channel offsets across time informs us of any potential outlying data that should not end up contributing to the aggregation of tooth separations towards a measure of global ACS offset mappings. Furthermore, portions of the ACS reduction and other analyses are sensitive to the establishment of nominal channel offset values for various time periods—another reason for careful study.

Figure 5.7 shows channel offsets for each gate type across time for channel 14, chosen



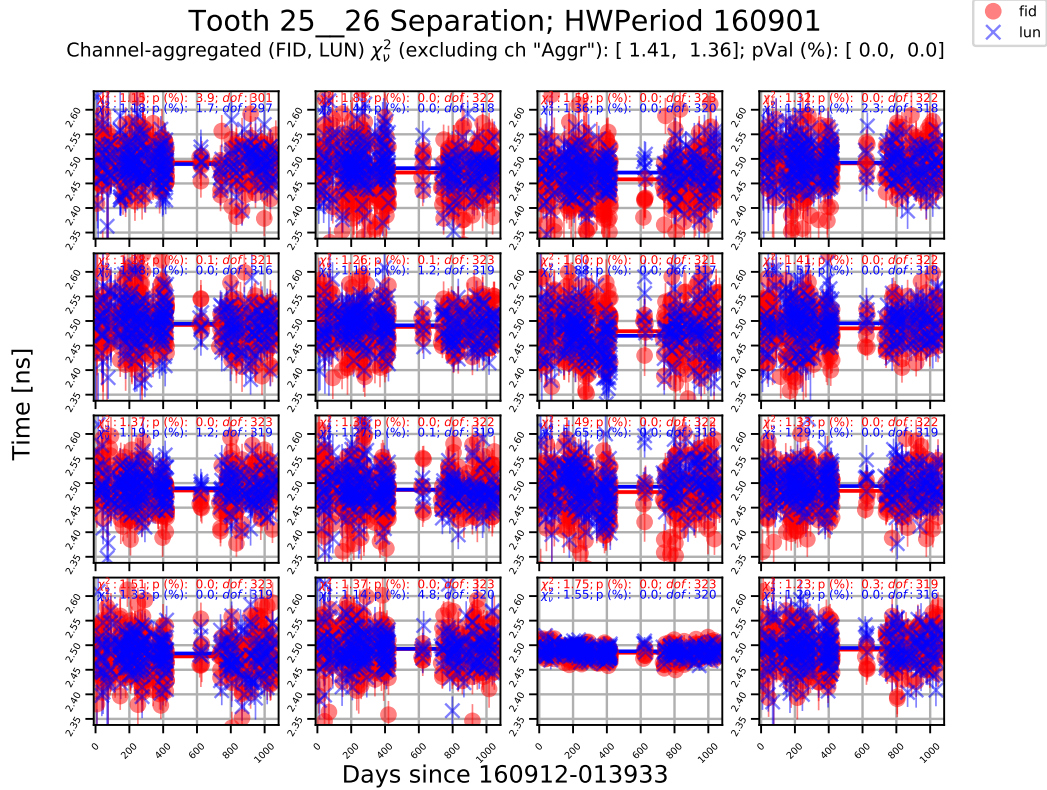
**Figure 5.7:** Channel offsets for channel 14 across time. Bottom left: histogram of offsets for the time range studied, with a windowed Gaussian fit establishing a measure for nominal for that time period. Goodness-of-fit measures are for the *windowed* fits. The fiducial gate is in red, the lunar in blue, and the darkest region is their overlap. Top: scatter plot of the associated histogram. Vertical black dashed lines represent where offsets were observed to be different enough to justify a new nominal value. Centroids of the windowed histogram fit appear as horizontal blue/red lines. Thick cyan and magenta lines, representing nominal values as a function of time, are overplotted. Note that the linear trend seen in these lines right before the first vertical dashed line is *not* an interpolated measure of channel offsets, but just an artifact of plotting. Runs within the night of August 18, 2017 experienced some cause-unknown temporary shift in channel offsets, primarily due to erratic behavior in channel 16. Notice the secondary peak in the bottom left subplot corresponds to the collection of offsets in the left-most region of the scatter plot, prior to experiencing a sudden shift in its relative offset (due primarily to another channel). Bottom right: residuals of channel offsets, as compared to nominal in each time period.

primarily to demonstrate a larger per-gate-type offset bias. Histograms are built of the offsets, a windowing routine is performed around the maximum signal to arrive at a reliable Gaussian fit, and a nominal offset measure is recorded. Residuals against nominal values for each period

appear in the bottom right; Since the periods have already been well-established at this point, these residuals appear neatly around 0. We can perform this routine over multiple time ranges individually to arrive at new well-defined nominal values any time large enough discrepancies are seen. Taking note that the differential channel offsets (the difference in fitted centroids on the lower-left subplot) for channel 14 tend to be  $\sim 0.6$  TDC bin, or  $15 \text{ ps} \sim 2.25 \text{ mm}$  one-way range, further demonstrates the importance of measuring these effects. Shortly after the Cs clock was installed, the electronics experienced a jump in offsets, primarily occurring in channel 11; Then again within weeks, though much less severe, for channel 16, both for unknown reasons. Other time periods were established due to features seen primarily in channels 1 and 16. The cause is currently unknown, though the primary suspect is old age. A follow-up study is planned in the future.

## 5.2.2 Aggregation of Tooth Separations

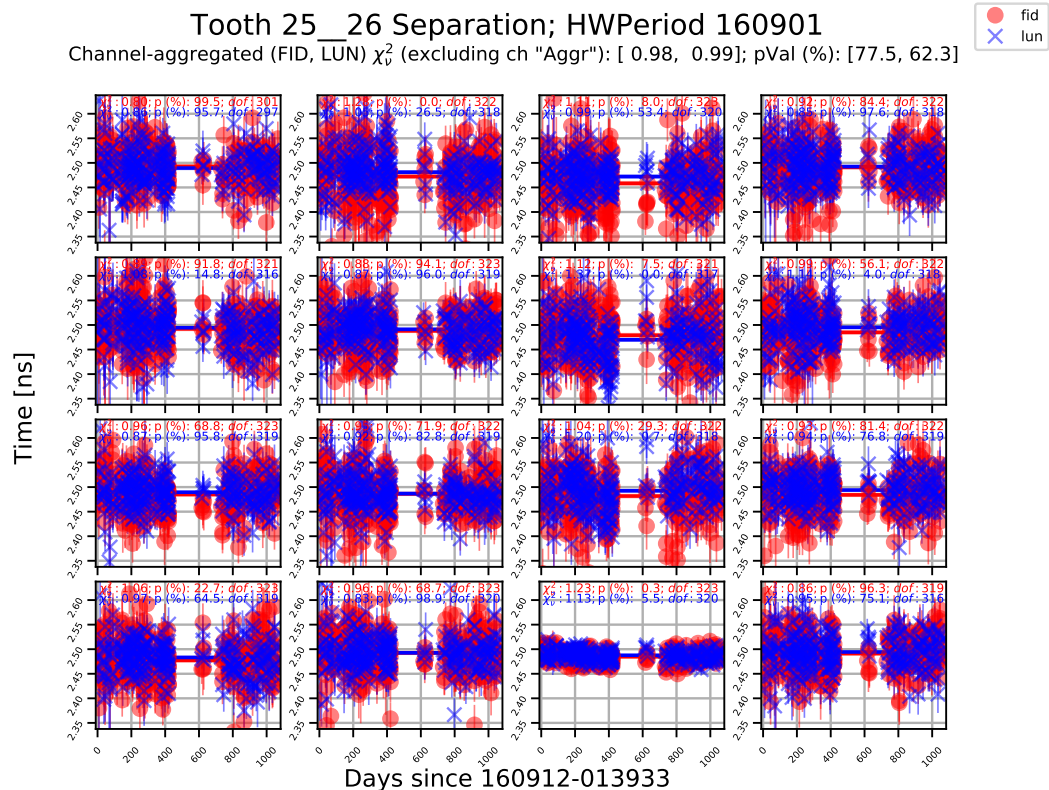
We utilize most of the ACS runs since the beginning of ACS over which to aggregate tooth separations; This encompasses the time periods between 160912 and 190814, as shown in Figure 5.8. In mid-August of 2019, the detector was replaced, necessitating the start of a new collection of ACS statistics for the new hardware. A lapse appears in ACS data collection between mid-November 2017 to the beginning of May 2018 due to the Cs clock experiencing a fatal error and having to deal with a lengthy shipping and repair process. In the time period shown, 332 runs contribute to the statistics. A slightly lesser number contributes to each collection of paired tooth separations (per gate type, per chan) as some tooth fits are rejected based on quality control reasons (thresholding conditions, large deviations between estimated vs fitted centroid, whether the fitted sigma falls within range of what we would expect for a Gaussian signal with a given number of contributing photons, etc.). Utilizing this large collection of run data, we result in a statistical reduction of  $\sim 18$ , allowing us to determine nominal tooth separations to the  $\frac{1.2}{18} \sim 0.07$  TDC bin level (recall the 1.2 value was arrived at in 5.1.2).



**Figure 5.8:** Example plot of tooth separations between teeth “25” and “26”, over time, for each channel and gate type separately. Channels run in increasing order left — right, top — bottom. The subplot for what would be channel 15 is actually the channel-aggregated tooth separations, on a per-run basis.  $\chi^2_v$ , corresponding p-values\*100 and the degrees of freedom (DOF) are shown for each gate, channel, relative to the weighted mean. An overall  $\chi^2_v$  measure,  $\frac{\sum_{ch=1}^{16} \chi^2_{ch}}{\sum_{ch=1}^{16} DOF_{ch}}$  for each gate type, excluding channel 15 (the per-run channel-aggregated tooth separations), is shown near the top of the plot;  $\sim 1.4$  for both gate types.

Despite the presence of a larger overall  $\chi^2_v \sim 1.4$ , there does not seem to be any evident temporal structure in the plots to suggest we could not use a weighted mean; Just that the tooth errors are underestimated. The tooth separation errors then are inflated by  $\sim \sqrt{1.4}$ , resulting in an overall  $\chi^2_v$  of  $\sim 1$ ; The same data as in Figure 5.8, after tooth separation uncertainty inflation, appears in Figure 5.9. This means we end up knowing the tooth separations to the  $\frac{1.2}{18} \cdot \sqrt{1.4} \sim 0.08$  TDC bin level, or  $\sim 0.3$  mm one-way distance.

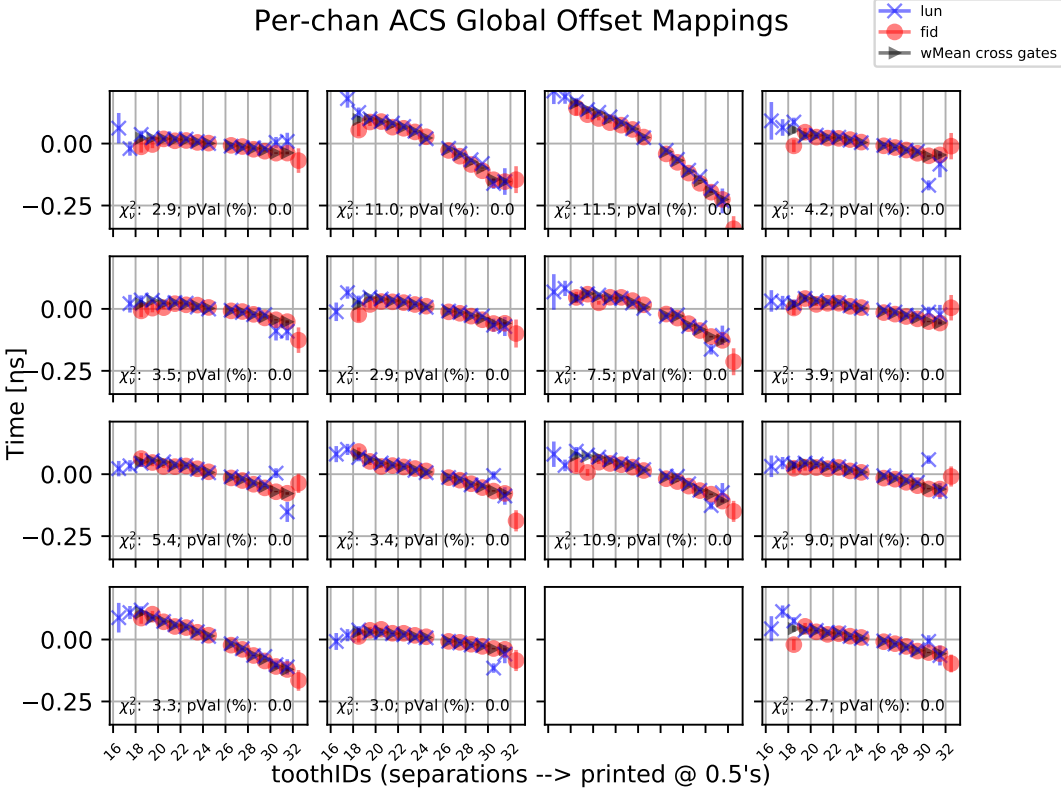
Given the seeming similarity in the variation of weighted means seen across channels/gate



**Figure 5.9:** Same data as Figure 5.8, after tooth separation uncertainty inflation has taken place. Note the  $\chi^2_v$  values near unity and p-value percentages above the threshold of 5%.

types, we later also tried aggregating the weighted means over channels but not gates, and tried aggregating over the two gate types, but not channels. Neither was a course of action resulting in confidence, even after inflating the tooth separation uncertainties. This can be seen in Figures 5.10 and 5.11. Note that the values shown are the weighted means of tooth separations *after conversion to offsets* (by subtracting-out the expected separation; more on this in section 5.2.3). Notice these plots show almost all tooth separations to be systematically lower than truth (to the left of the reference tooth, the corresponding tooth separations are negative-valued, so a positive offset indicates the separation is smaller than expected; to the right of the reference, the corresponding tooth separations are positive-valued, so *negative* offsets indicate the same thing), suggesting APOLLO is counting TDC time too slowly in some manner for each gate type; This may point to

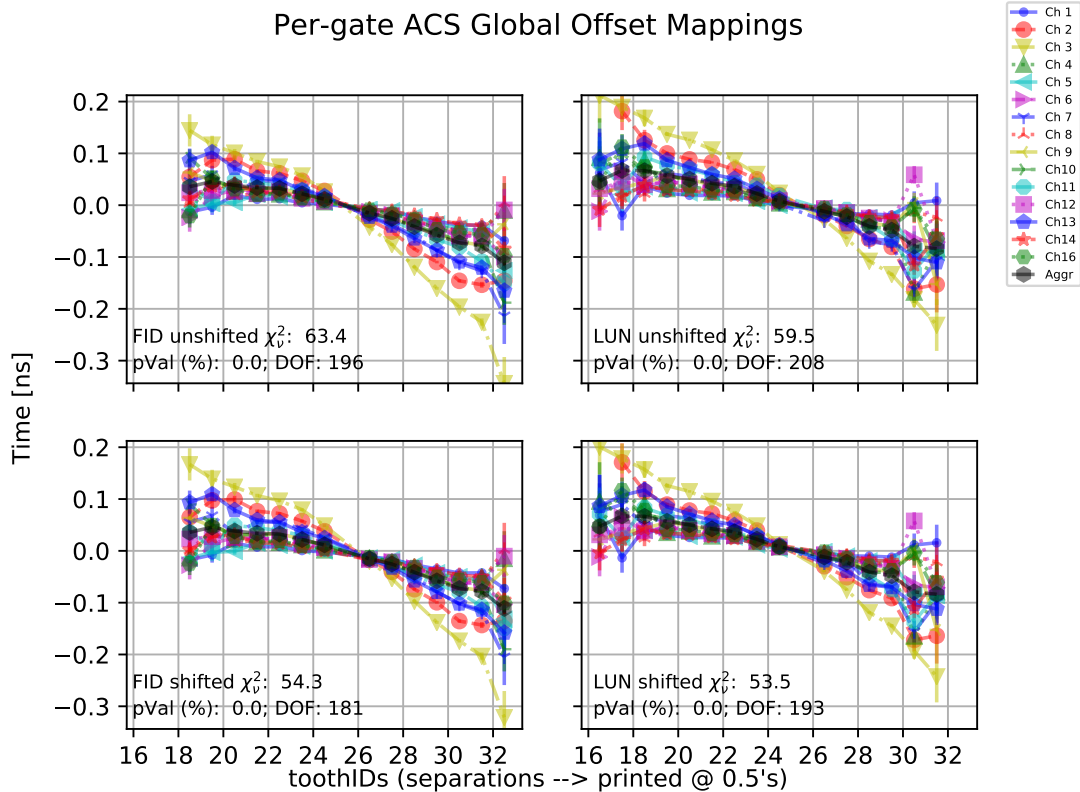
some failure in the CALTDC subroutine to describe a linear trend of timing which, in application of the CALTDC mapping, would not have been removed from the photon timestamps.



**Figure 5.10:** Attempt at aggregating tooth separations over the FID and LUN gates but not channels using a weighted mean for each tooth ID, represented by the black points. The x-axis represents the label each tooth has been assigned.  $\chi^2_v$ s and p-values are printed near the bottom of each subplot, suggesting a lack of confidence in the result.

### 5.2.3 Processing the Aggregate Tooth Separations

We have demonstrated that the tooth separations are consistent over long time spans, but each channel and gate type possesses unique characteristics, preventing further aggregation. From this point, we can form the aforementioned “parent” combs (one for each channel and gate combination) representing where ACS teeth truly are in TDC-space, and their associated offsets. We begin building-out the locations of the parent comb teeth relative to the chosen



**Figure 5.11:** Attempt at aggregating tooth separations across channels, for each gate type separately using a weighted mean for each tooth ID, represented by the black points. The bottom row differs from the top in that it first attempts to bring each channel’s mapping best in-line with the entire collection by shifting an appropriate amount. Again,  $\chi^2_v$ s and p-values are printed near the top of each subplot, suggesting a lack of confidence in the result.

reference tooth when calculating the separations, and correspondingly build-out the truth positions relative to the reference tooth. Uncertainties for the parent teeth locations are equivalent to the separation uncertainties (which already include the uncertainty of the reference tooth location). The uncertainty at the location of the reference is assigned the mean uncertainty of adjacent tooth separation measures.

Now that we have the parent comb formed, we can compare where its teeth landed relative to the simultaneously-built truth reference grid, analogous to what is seen in Figure 5.4 (though in Figure 5.4, the reference tooth was chosen to be at the left-most tooth, for ease of introducing the concept). Uncertainties on each offset value are assigned to be the same as the uncertainties in

the parent tooth locations, considering the mathematical operation is just the difference between parent tooth locations against reference grid markers (contain “no” uncertainty). This results in an aggregate form of the interpolated mappings for use in correcting each channel/gate TDC-space timing.

In an attempt to recover run-to-run variations, we return to each particular (“daughter”) ACS comb (unique run, channel, gate combination) and ask “where should this offset mapping (parent comb) be placed in TDC space to *best* match the tooth positions of these particular comb teeth”? A one-parameter least squares fit to systematically shift the parent comb one direction in TDC-space, is performed to answer as much, the flow/logic of which follows.

First, pick out *just* the parent comb teeth that correspond to the particular labels of the daughter comb we’re fitting to: this ensures that during the sliding fit, it does not try putting, for example, parent tooth “25” at the location of a daughter tooth with a label other than “25”. Then, pick a particular tooth in the daughter comb we are fitting to, and place the matching tooth of the parent comb at the same TDC (time) location. Recall the relative locations of the parent teeth tell us where the globally-derived offsets should be placed.

Finally, the one-parameter least squares fit is executed, trying to best-align the positions of the offsets against the positions of tooth centroids in this particular comb, accounting for the uncertainty in each measure while fitting. Uncertainty on the sliding fit parameter is combined with the uncertainty contributions of the offset mapping during application of corrections. To this end, the magnitude and structure of ACS corrections stay the same for each comb of a particular channel and gate type we deal with across runs, but the TDC values those corrections occur at will be different.

Figure 5.12 displays an ACS comb with the corresponding offset mapping locations overlaid on top, after the sliding fit has already occurred. The reader may note this example has a shift value close in magnitude to that of its uncertainty; Unfortunately the sliding fit is the primary contributor of uncertainty in the end, but we will see it still ends up well below a typical NP




uncertainty. Additionally, recall we showed a “preview” of this outcome in Figure 5.5, which acts as a complement plot to Figure 5.12 by displaying the actual offset mappings in a format similar to the conceptual plot in Figure 5.4. It is again worthwhile to note that the structure of the offset mappings reveals TDC space to suffer no significant discontinuous and/or nonlinear timing effects from tooth to tooth, showing that interpolation between offsets is a more-than-reasonable pursuit.

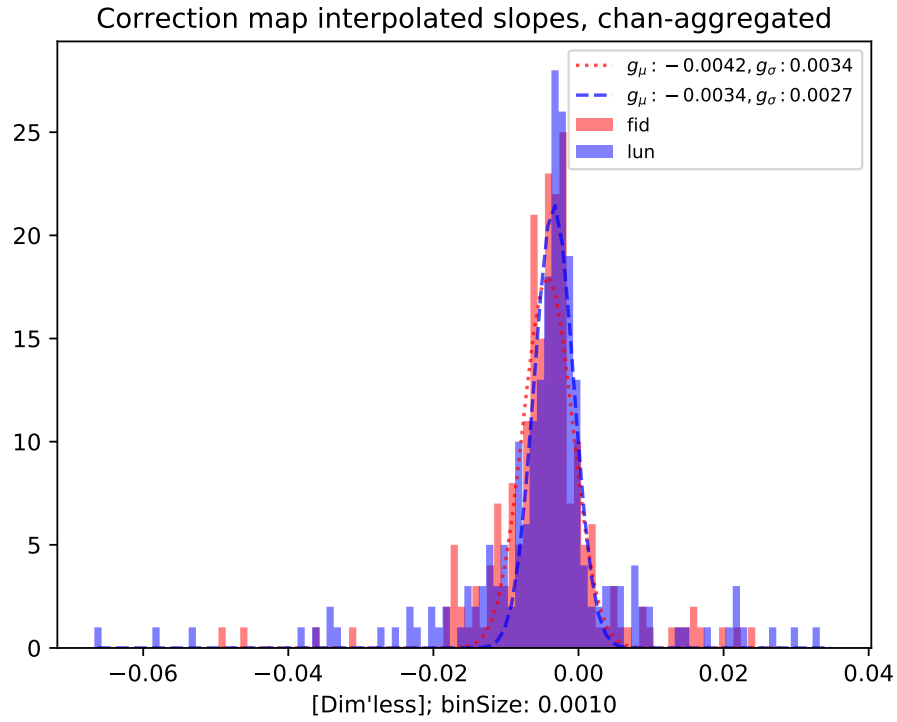
In order to fully complete the mapping, we return to the concern regarding corrections outside the bounds of offsets: how are photons that lie outside these bounds handled, since we cannot interpolate there? Fortunately since the mapping is generally well-overlaid with the range return signal, the only photons landing here will be background photons and do not play much part in the final answer. However, a distribution of interpolated line slopes is fit with a gaussian to see if a monotonic trend of offsets exists in one direction of TDC space. We aggregate the interpolated line slopes across channels and tooth pairings, for each gate type separately, as the statistics present are not sufficient to expect to see meaningful channel-distinct slope distributions; This can be seen in Figure 5.13. The distribution is largely composed of negative-valued slopes, indicating the presence of a monotonic trend in TDC space, which can be explicitly seen in Figure 5.5; In other words, most tooth spacings are systematically  $< 2.5$  ns, an example of which can be seen from the low-valued (compared to 12.5 ns) weighted means in Figures 5.8 and 5.9, suggesting APOLLO counts time too slowly. The fitted Gaussian centroid to this distribution is used as a characteristic slope for the extrapolated lines on either side of the offsets mapping; The Gaussian sigma is used to permit an extrapolated slope uncertainty that propagates into photon correction uncertainties sampling from the extrapolated region. We can see a manifestation of the negative-going slope in almost all interpolated lines in Figure 5.5; The centroid of each interpolated line distribution is represented by the dashed lines on the bounds of the offset mapping. The uncertainty of these distributions can be visualized especially towards the left edge of the upper subplot in Figure 5.5 where it begins to dominate over the uncertainty

in the last mapping point causing the uncertainty bounds to begin to curve.

The final product is an ACS offset mapping for each channel/gate, for each run, allowing us the ability to accurately capture the timing of photon events in-range of the mapping! We have shown the method for deriving ACS corrections and why we can trust the derivation, but how are the corrections actually applied to range photons? We must ensure that the application is not somehow biasing the final round trip time and are being used correctly.



**Figure 5.12:** Demonstration of how well daughter comb teeth line up to the parent comb for lunar gate channel 14 in run 190416k. Daughter comb data is displayed in blue, with red dash-dotted gaussian forms representing the parent comb teeth locations (all with the same amplitude and width): this is where we expect the daughter comb to *really* be. TDC positions have been converted to time relative to the STOP pulse that ends the TDC gate using the appropriate CALTDC-derived quadratic conversion. While it may seem like this ACS comb is closer to the end of the 100 ns TDC range, there exists processing time between the receipt of a START signal and beginning of the TDC gate—routinely on the scale of  $\sim 14$  ns for every channel (measured by CALTDC for each run uniquely). Given the small offset magnitudes involved, it is difficult to see deviations from nominal positions, though a close look does reveal differences in the locations of the parent and daughter tooth locations. The in-set plot shows the response of the sliding error ( $\chi^2$ ) function for a user-fed continuous set of timing shifts, to more clearly demonstrate the fit has well-defined error feedback.



**Figure 5.13:** Distribution of interpolated line slopes from time-aggregated correction mappings, aggregated across channels and tooth pairings. The fiducial gate is in red, the lunar in blue, and their overlap is the darkest region. Units are dimensionless, as the corrections and photon timestamps are both expressed in ns. The fitted centroids deviate from 0 by a little more than  $1\sigma$ , indicating that we can trust slightly more than distrust the presence of a monotonic trend away from 0 in TDC space.

## 5.3 Correction Scheme Application

### 5.3.1 ACS-enabled Run Corrections

Fortunately, insertion of ACS corrections into the main reduction pipeline is fairly straightforward, and occurs near the start of any data processing. Each run's ACS offset mappings (unique for each channel and gate type) are first formed by sliding the corresponding parent comb against the ACS combs to determine run-specific offsets, done for all ACS runs outside the main reduction pipeline to avoid further processing time for range results. Once running the main reduction, run data is grabbed and the photon timestamps are converted from TDC space to time using the CALTDC results. Recall the ACS offset mappings were derived from ACS comb teeth that had been converted from TDC space to time also using the CALTDC results. Thus, at this point, the ACS corrections and range photon timestamps are expressed on a consistent basis with each other. Now we can simply place each range photon's timestamp onto the offset mapping and ask what correction should be applied. This correction is applied to the range photon timestamp, after which the reduction pipeline resumes normal operation/processing to extract a differential time measurement.

In applying corrections, however, we need to estimate how much random uncertainty we have added to the nominal round trip time by introducing these corrections. We do not immediately have a functional form for a characteristic correction upon which to apply a formal uncertainty propagation procedure; In absence of this, we decided the use of a straight mean over the applied corrections is the simplest and most appropriate method for a characteristic correction result. So, we can think of the uncertainty on that characteristic correction as the formal error propagation of a straight mean:  $\sqrt{\sum_i \frac{\sigma_i^2}{N}}$ , where  $i$  is a dummy index and  $\sigma_i$  is the uncertainty on correction  $c_i$ . Recall this  $\sigma_i$  corresponds to the uncertainty on the weighted mean of paired tooth separations over the aggregation time period.

Of course, we do not want to include correction uncertainties for non-signal photons,

as they likely have large uncertainties due to being near or outside the boundaries of ACS mappings, where the uncertainty is much larger. So we window the applied corrections for use in a characteristic correction (and its corresponding uncertainty) around photons that are near the centroid of the functional fit in a  $2\text{--fitted--}\sigma$  window (for LUNs, part of the reflector timing profile is added to the representative  $\sigma$ ).

The estimation of random uncertainty added by the ACS corrections does not account for the uncertainty in the sliding fit parameter, however. Since each sliding fit represents some intrinsic offset between the parent and daughter combs, it is considered a systematic uncertainty and is not eligible to be statistically reduced by repeated sampling of range photons. However, any timing artifacts that would cause an offset between the daughter and parent combs are expected to be common to all channels (recall calibrations of relative channel offsets have already occurred). Therefore the uncertainty in the sliding fit parameter is allowed to be statistically reduced by combining the uncertainties across channels (*as if* we were averaging the sliding fit values, though we're not — they've already been added into the per-channel per-gate offset mappings):

$\sigma_{slide,tot} = \frac{\sqrt{\sum_{ch \neq 15} \sigma_{ch}^2}}{15}$ . This result is then RSS'd into the windowed correction uncertainty result.

An average magnitude of ACS correction uncertainty is around 6 ps for either gate type. The average number of (windowed, around  $\mu$  of functional fit) FID corrections for a 5000 shot run is  $\sim 3000$ , and  $\sim 200$  for the LUNs, corresponding to the number of *real* FID and LUN photons. The RSS of the two characteristic uncertainties would be dominated by the LUN-side, given the significantly lower number of records. If all correction uncertainties had the same magnitude (6 ps), we would expect  $\sqrt{\frac{N \cdot \sigma^2}{N^2}} = \frac{\sigma}{\sqrt{N}} = \frac{6}{\sqrt{200}} \sim 0.4$  ps as the characteristic *random* uncertainty on the collection of applied corrections.

A typical magnitude for the sliding fit uncertainty is around 7.5 ps. Combining the uncertainties across channels yields a characteristic uncertainty of  $\sim 2$  ps for the sliding fit for each gate type individually; Combining the gates gives  $\sqrt{2} * 2 \sim 2.8$  ps. Clearly this uncertainty will dominate over the random contribution, resulting in a final uncertainty of  $\sim 2.8$  ps—significantly

less than the  $\sim 10$  ps typical NP uncertainty. In practice, the mean over actual collections of estimated ACS uncertainties is closer to  $\sim 3.3$  ps.


To be clear, the final product in applying corrections to ACS runs containing calibration data is a single number and associated uncertainty estimating the whole-run differential offset. To quickly recap, corrections are applied on a per-photon basis (thus, per-gate, per-channel), and the collection of corrections for each gate type is windowed around the respective range signals. A mean over the windowed corrections is taken, and a characteristic uncertainty calculated as described above, for each gate type. The means are then converted back to offsets (multiplied by negative one), and the difference between them is taken to form a differential offset measure:  $O_{differential} = O_{lun} - O_{fid}$  (where  $O_{lun}$  and  $O_{fid}$  are the mean offset measures for the respective gates), with the final uncertainty the RSS of the two gate offset uncertainty measures.

### 5.3.2 Updated Range Error Statistics

Great! Enough description; How do ACS corrected results compare to non-corrected ones, for those runs containing calibration data? The same date range (160912-190814) of data as above was used in a batch comparison, separated by which clock was used, to check the ACS offsets and their effect on RTT results. Figure 5.14 shows two plots to this effect, with units kept in the time-domain as that is the actual measurement performed. Note that not all the runs previously shown (i.e. in Figure 5.8, for the same *reduced* run range) are included here, as some may have had minimal lunar range signal and thus unusable for published range results. We observe from the Cs distribution in the left-hand side of Figure 5.14 that our round trip times tend to be biased by  $\sim 12$  ps or  $\sim 2$  mm one-way range distance, according to the ACS differential offset estimates. This is roughly in-line with our preliminary assessment of the few mm in Section 4.3, in which we used the local  $\Delta t$  method of correction estimation.

The right-side plot of Figure 5.14 indicates a characteristic  $\chi^2$ -like measure (one for each run) comparing the corrected dataset of NPs against the uncorrected round trip times:

$\frac{(RTT_{\text{corrected}} - RTT_{\text{uncorrected}})^2}{\sigma_{\text{corrected}}^2 + \sigma_{\text{uncorrected}}^2}$ . A straight mean over these values indicates the overall  $\chi^2_{\text{v}}$ -like value for the entire collection (separated by time base). The Cs clock result of  $\sim 1.2$  then suggests a scale of  $\sim \sqrt{1.2} \cdot \sqrt{2} \sim 1.7$  times the typical uncertainty in a single normal point ( $\sqrt{2}$  is included given the denominator is the RSS of uncorrected with corrected errors, which individually are on roughly the same scale).



**Figure 5.14:** Data for the Cs, XL-DC clocks are plotted in red and blue, respectively. Left: Estimates of differential ACS offsets. Photon RTT offsets for each gate type are windowed the same as for the offset uncertainty estimate, a straight average is taken, and the difference (LUN - FID) between averages is recorded as each datapoint appearing here. Note that corrections are signed quantities, but the distribution of differences of means coincidentally has its left edge right at 0 ps. Right:  $\chi^2$  measures formed from the square of RTT differences over their combined uncertainties.  $\chi^2_{\text{v}}$  is found by taking the mean.

As for the XL-DC data presented in Figure 5.14, a fit was not attempted as it clearly does not have enough statistics to appear as a non-Gaussian distribution. Similar to the initial range error results shown in section 4.3, we can again marvel at the difference in widths of each clock's correction distribution, in part (as previously alluded to) due to the guarantee that there exists no clock drift relative to the ACS pulses when using the Cs clock, though also lower intrinsic jitter in the clock. It should be again noted by the weighted mean over XL-DC corrections, 25 ps or  $\sim 4$  mm one-way range, and spread of  $\sim 12$  ps or  $\sim 2$  mm, indicates the inaccuracy of past APOLLO data is still nowhere near inaccurate enough to explain the best model residuals, which



are substantially larger.

### 5.3.3 Corrections to ACS-disabled Runs

We have demonstrated the ability to utilize ACS corrections for those runs that have ACS enabled during ranging operation, but what about those runs lacking in ACS information? The addition of photon detections due to ACS information interferes with telescope pointing feedback from detector event rates which are indicated during ranging on a 4x4 grid corresponding to the APD elements. This effectively masks the actual lunar signal underneath a “floor” of ACS detections on runs/nights when observing conditions already make it difficult to ascertain the correct positioning to see the lunar return signal. Can we somehow extend the corrections method described above to sensibly correct runs when ACS is disabled due to low lunar signal? This section seeks to describe an avenue for realizing such a pursuit, and ultimately show results using the concept.

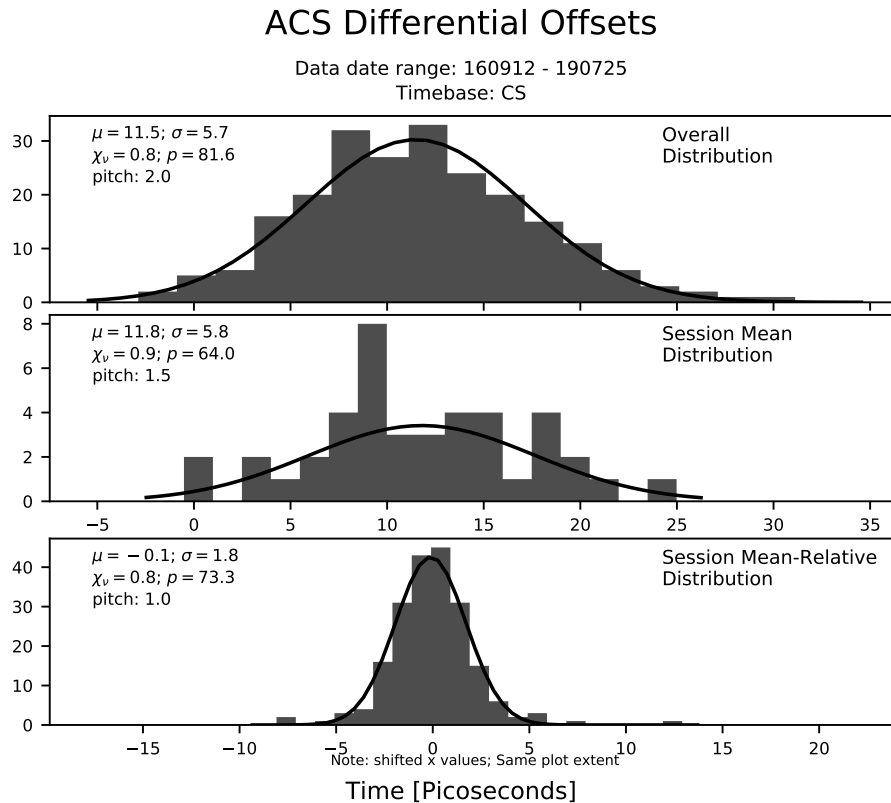
The main options we have available to us are the building of distributions of offsets over the entire course of ACS operation, and consideration of session-specific information (“sessions” here refers to the collection of runs performed in a single night, or session of observation, typically lasting  $< 1$  hour), where available to study. If we have enough ACS information within a session, we could for instance build a trendline to explore the possibility of monotonic trends in offset magnitudes over the course of the session, to inform where offsets for interleaved ACS-disabled runs might land. Alternatively, we can be informed what a likely offset *should* be for ACS-disabled runs based on “global” (across long timespans) ACS offset distributions and/or a likelihood argument using those distributions and a session weighted mean (if a session contains a non-zero amount of ACS-enabled runs).

Before moving on, it should be noted that while we gave the consideration of trendlines within sessions an honorable mention here, we ultimately did not have strong enough reason to utilize them. The amount of per-session information as compared to using a simple session

weighted mean was not significantly different. Beyond that, we would add unnecessary complication in the use of trendlines and how to calculate appropriate uncertainties based on how many contributing points exist in each session; Especially in teasing-out when it is most appropriate to use trendlines over simpler session-mean information. A more uniform and simple approach is desired.

The most obvious global distribution to build would be a distribution of just the differential whole-run ACS-enabled run offsets seen on the left in Figure 5.14, which we refer to as the “overall” distribution. This distribution reflects what the ACS delivers across *all* runs, and is therefore a representative offset and uncertainty. We can also calculate a session weighted mean (we will use the abbreviation  $w_{\text{Mean}}$  from here-on) offset value, for those sessions containing enough ACS-enabled runs to justify trust in a representative measure of a full-session characteristic offset; We say a session must have at least 3 ACS-enabled runs to contribute to this distribution which we refer to as the “session  $w_{\text{Mean}}$ ” distribution. Correspondingly, we can form what we call the “session  $w_{\text{Mean}}$  relative” distribution by comparing the individual ACS run offsets within a contributing session to *that session's* weighted mean (also subject to the 3+ runs in-session condition). These distributions, and Gaussian fits to them, can be seen in 5.15. The histograms are plotted on a uniform scale such that their relative spreads are true to the reader's eye, revealing the individual offsets, as related to their session weighted mean, to potentially contain some valuable information.

The most immediate and direct route of correcting ACS-disabled runs would be to calculate the most likely value of ACS-enabled offsets, i.e., the centroid of the *overall* distribution, to use as an offset value. Uncertainty for the offset would simply be the spread of this distribution. This is well and good for cases where 0 ACS-enabled runs are in the session, but we question whether we can do a better job for those sessions containing *some* ACS information, by using the information gathered regarding session  $w_{\text{Means}}$  and how individual offsets within each session compare to those.



**Figure 5.15:** Distributions, for the Cs clock, of ACS offsets for the run range specified. The overall distribution, previously seen in Figure 5.14, is presented at the top. The session wMean and session wMean relative distributions are shown in the middle and bottom plots, respectively. Note the *significantly* reduced spread of the session wMean relative distribution as compared to the two above, indicating the likelihood of capturing session-specific systematics.

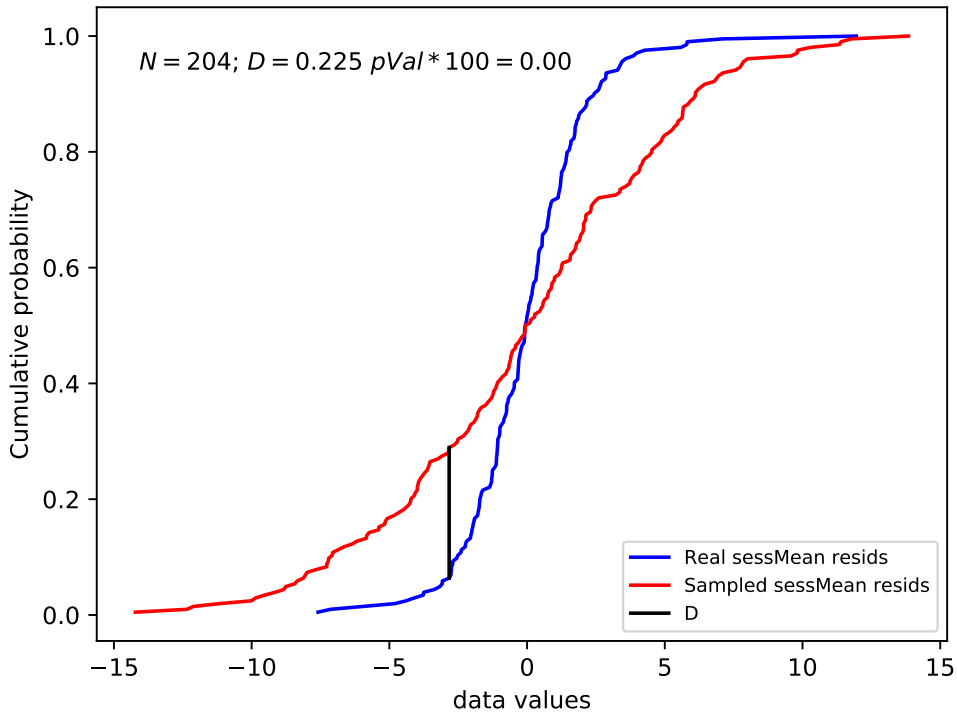
Prior to describing *how* we might utilize the per-session information, it would behoove us to demonstrate that said data is actually sensitive to systematic influences on a per-session basis. It could instead be the case that each ACS-enabled run offset in a session is simply a random sample taken from the overall distribution, and how they relate to the session wMean is not actually capturing any unique information. How might we show this, one way or the other?

To address this, we form another session wMean relative distribution, based on random samples taken from the functional fit to the overall distribution, in the exact same manner as the real session wMean relative distribution. That is, e.g., say one of our contributing sessions has 4 ACS-enabled run offsets in it. As previously described, we would calculate a

$w_{\text{Mean}}$  of these 4 offsets, and record the residuals of the 4 run offsets as compared to the session mean into the session  $w_{\text{Mean}}$  relative distribution. In similar fashion, we take 4 random samples from the overall distribution, calculate a mean from the random samples (unweighted mean in this case, as the random samples don't have individual uncertainties assigned to them), and record the residuals of the random samples from this mean into our random-sampled copy of the session  $w_{\text{Mean}}$  relative distribution. After considering all contributing sessions, we can ask whether or not these two distributions are derived from the same parent distribution by using a Kolmogorov-Smirnov test, which examines the maximum deviation of cumulative distribution functions (CDF).

The results of this can be seen in figure 5.16. Choosing a standard p-value threshold of 0.05, anything below which indicates  $> 95\%$  confidence the distributions come from *different* parents, we can confidently assert that some systematic per-session information being captured in how offsets on a given night relate to their  $w_{\text{Mean}}$ . It should be additionally noted that the structure of the two session  $w_{\text{Mean}}$  relative distributions being compared, as indicated by the instantaneous slope at any given point of their respective CDFs, suggests the randomly-sampled version to be significantly more spread-out than the real version. This suggests that on a given night, offsets tend to cluster more around one part of the overall distribution as opposed to being randomly distributed, resulting in a tighter relation to the session  $w_{\text{Mean}}$ .

Now that we have justified the use of session-specific information, how might we utilize it? We seek to come up with a fair estimate of what correction (and corresponding uncertainty) should be applied to ACS-disabled runs in a particular session. If the session has any ACS-enabled runs, we start by again calculating a  $w_{\text{Mean}}$ . We then could simply use the  $w_{\text{Mean}}$  as the offset for each ACS-disabled run in the session; However we also have information about individual run offsets on a global scale: the overall distribution. Consider the case that the calculated  $w_{\text{Mean}}$ , given the limited number of ACS-enabled runs in the session, lies right on top of the centroid of the overall distribution; Using this resulting value as-is seems to be a fair thing to

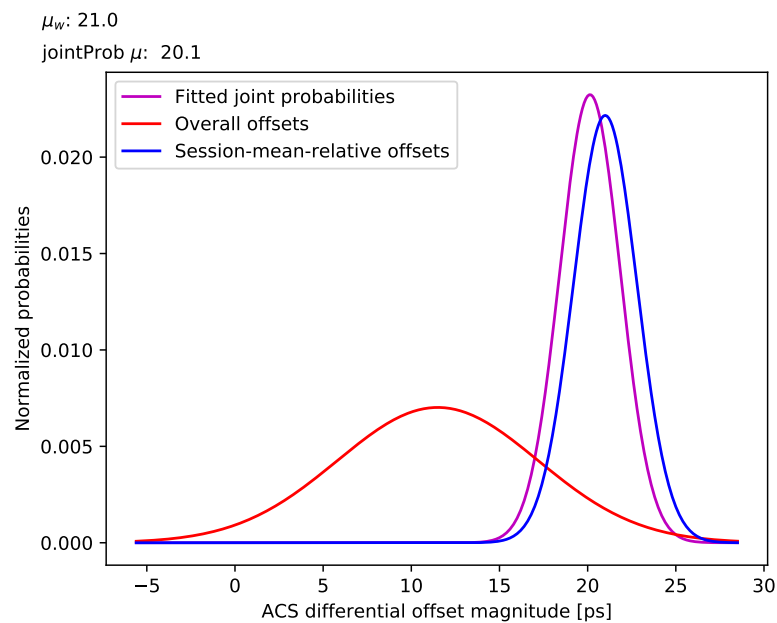


**Figure 5.16:** Cumulative distribution functions of the `session wMean` relative distribution and its randomly sampled (from the `overall` distribution) analog. Only data utilizing the `Cs` clock as the timebase, and for sessions containing 3+ ACS-enabled runs, contribute. The maximum deviation of their CDFs, referred to as the “KS statistic” and denoted as `D` is found, and used to judge the probability that the distributions are samples from the same parent. The x-axis represents the magnitude of offset residuals relative to their respective session `wMeans`, in picoseconds.

do. However now imagine the case that the calculated `wMean` is  $\sim 2 - \sigma$  discrepant relative to the `overall` distribution. We likely would not want to take that result at face-value for use on ACS-disabled runs in the session, as it would be fairly unlikely for *all* runs in the session to carry such an extreme offset, unless the `session wMean` relative distribution happened to be much more like a delta function. So, we would like to adjust the resulting value to make a little more sense/be a little more trustworthy.

This is where our use of the joint probability operation comes into play: it allows us to shift the raw estimated `wMean` for use in individual run offsets to a a bit more consistent with the

centroid of the overall distribution by some marginal amount. We offset the session  $w_{\text{Mean}}$  relative probability distribution by the estimated  $w_{\text{Mean}}$ , then take the joint probability of that with the overall probability distribution to result in an adjusted estimate of the session  $w_{\text{Mean}}$  offset to be used on ACS-disabled runs in that session. This is accomplished by a simple multiplication of the two distributions, as we assume the variables to be independent of each other; i.e., for a given ACS offset, we cannot know a-priori how it will compare to its  $w_{\text{Mean}}$ . Figure 5.17 illustrates the concept.



**Figure 5.17:** Example using a 2-ACS-run session on the night of 190427 with larger derived offsets contributing to the estimate of the session  $w_{\text{Mean}}$ , shown in the upper left. The overall (red) and session  $w_{\text{Mean}}$  relative (blue) distributions use the same histogram bins, are normalized and multiplied together after the centroid of the session  $w_{\text{Mean}}$  relative distribution is placed at the estimated session  $w_{\text{Mean}}$ . The product (magenta) is renormalized and fitted with a gaussian to determine the adjusted session  $w_{\text{Mean}}$ , also displayed in the upper left.

At best, the confidence we have in the estimate of a session  $w_{\text{Mean}}$  of a particular session is described by  $\frac{\sigma_{\text{session } w_{\text{Mean}} \text{ relative}}}{\sqrt{N}}$ <sup>1</sup>, where  $N$  is the number of ACS-enabled run corrections in that particular session. This corresponds to the uncertainty in placement of the session

<sup>1</sup>In practice we use the formal error propagation on the  $w_{\text{Mean}}$  so as to be sensitive to the estimated uncertainties within a session, but this expression is theoretically equivalent, and more easily illustrates the points to follow.

wMean relative distribution, relative to the overall. Assume for a moment that we know the session wMean to a very high degree, as if we somehow had, e.g., 1000 ACS-enabled runs within the session such that this positioning uncertainty is effectively 0. We would expect the distribution of those 1000 points relative to the session mean to spread out in a similar manner to the session wMean relative distribution, given that is what we see globally. Then we ask “for an individual run offset, by how much would we expect it to vary, relative to the wMean?” That would be  $\sigma_{\text{session wMean relative}}$ . So, the total effective uncertainty we assign to this ACS-disabled run correction ends up being  $\sqrt{\sigma_{\text{session wMean relative}}^2 + \frac{\sigma_{\text{session wMean relative}}^2}{N}}$ , which is added in quadrature to the NP uncertainty for that run. Of course for sessions without any ACS information, we must rely on just the centroid and spread of the overall distribution. Table 5.2 summarizes the uncertainty scaling, bounded by the spread of the overall distribution for 0 ACS runs in-session, and the spread of the session wMean relative distribution as N approaches infinity.

**Table 5.2:** Scaling of ACS-disabled run correction uncertainties, as a function of number of ACS runs in the corresponding session. The scaling sensibly decreases monotonically as N grows, asymptotically approaching the spread of the session wMean relative distribution.

N		Uncertainty (ps)
0		5.7
1		2.55
2		2.20
3		2.08
4		2.01
...		...
∞		1.8

The reader will notice the jump in estimated uncertainty between  $N = 0$  and  $N = 1$  elucidating the difference in having no ACS information in a session vs having *some*. The reader *may* also notice an unusual discrepancy between uncertainties for sessions with  $N > 1$  compared to the typical expected uncertainty mentioned in section 5.3.1. How can runs without ACS information be assigned a lower uncertainty than the typical uncertainty of a run *with* ACS

information in it? Something must be wrong either with the formation of the `session wMean` relative distribution, or how we assign uncertainties to runs with ACS data in them. Further study reveals the method of final uncertainty estimation results in *overestimations* in general despite the straightforward nature behind the idea (given the absence of a functional form for a characteristic correction). This is acceptable, though, due to its magnitude being significantly smaller than a typical NP uncertainty even with overestimation. However we still do *not* want to allow runs without ACS information to achieve a lesser uncertainty than the typical value. In order to compensate, we include a third term proportional to the typical ACS uncertainty in the expression under the square root above. The value is  $1 - \sigma$  (of the distribution of estimated ACS uncertainties) larger than the typical estimated ACS uncertainty, because even with many runs in a session, the confidence of an ACS-disabled run's offset should theoretically never be better than an ACS-enabled run. The table, after this adjustment, can be seen in Table 5.3.

**Table 5.3:** *Adjusted* scaling of ACS-disabled run correction uncertainties, as a function of number of ACS runs in the corresponding session.

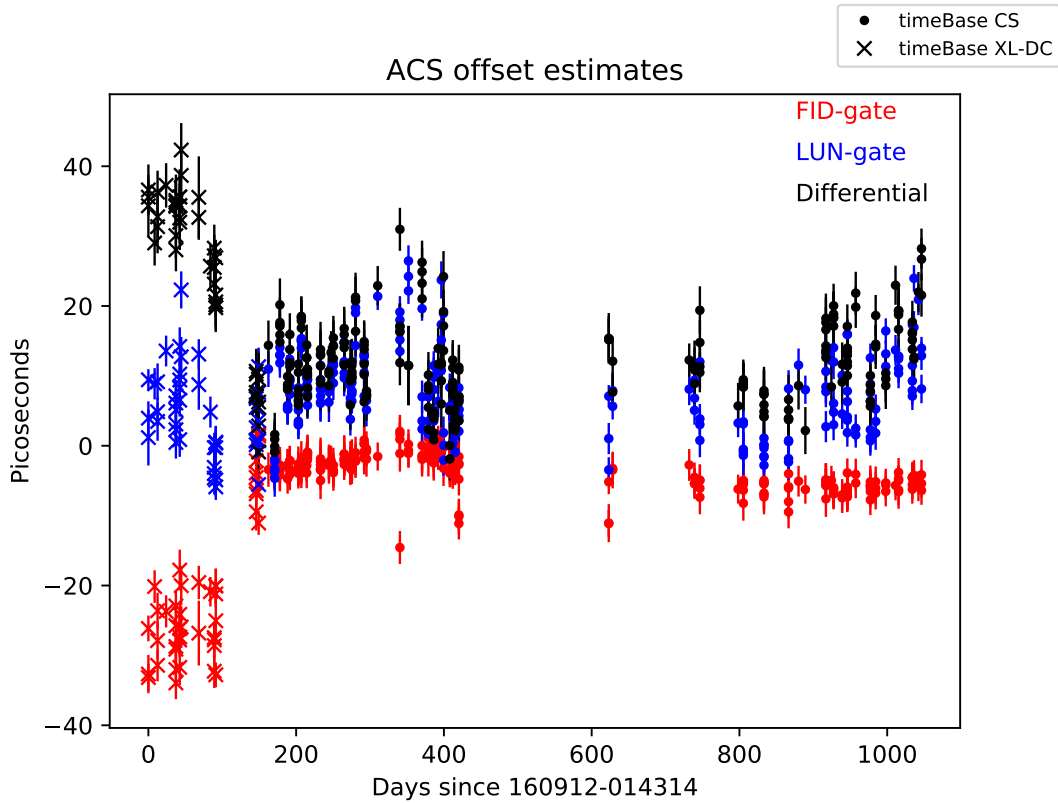
N		Uncertainty (ps)
0		5.7
1		4.33
2		4.14
3		4.07
4		4.04
...		...
$\infty$		3.94

### 5.3.4 ACS Offsets Over Time

As an interesting follow-up, we can look at the differential offsets (the values contributing to the `overall` distribution) and their constituents (per-gate-type, since the differential measure is the difference between them) as a function of time, shown in Figure 5.18. Clearly some time-variability is present in the ACS offsets, given the obvious sinusoid present, with a period



of  $\sim 1$  year. It is reassuring to note that the structure present is not wildly varying. Indeed, it further corroborates the story that APOLLO seems to be a fairly well-behaved system overall. Additionally, it provides even more justification for the *need* to be sensitive to session-specific information and visualizes why the spread of the `session wMean` relative distribution is so tight.



**Figure 5.18:** Scatter plot of same data appearing in the left panel of Figure 5.14, but also displaying the individual gate offsets for visualization. The data is split by timebase such that all “X” points represent the XL-DC and “.” points represent the Cs clock. Notice the comparatively more spread-out lunar data; This is because the offset estimates are dependent on how many registered (meaning, part of the main range signal) photons are in each gate, and lunar returns can be highly variable, as previously mentioned.

Further study of the time-varying structure is planned in the near future. We can compare the sinusoidal structure seen to temperature records to elucidate a connection, for example. We look forward to teasing-out other clever ways of applying the ACS.

## 5.4 Conclusions and Future Work

A robust method of estimating ACS corrections as a function of TDC-space, conveniently relying on long-timescale trends in TDC structure, has been developed for those runs containing ACS data. In summary:

- ACS data present in run: use that run's ACS correction and uncertainty.
- No ACS data in run, but  $> 0$  ACS runs in session: use `session wMean`, `session wMean relative` and `overall` distributions to form correction and uncertainty.
- No ACS data in run *or* session: use `overall` distribution for correction and uncertainty.

We are fortunately able to derive ACS corrections to a confidence level significantly less than the typical  $\sim 10$  ps uncertainty of a NP, due to the high quality of ACS information we are able to extract, and the ability to reliably aggregate long stretches of collected calibration data. As such, we are able to correct the  $\sim 2$  mm range bias of APOLLO without suffering appreciable additions to the uncertainty of the data, thus preserving APOLLO as a  $\sim$  few mm *precise* measuring device as well.

Moreover, the study of these offset estimates on a per-session basis have shown themselves to contain useful and statistically significant systematic information, reinforced by the tightness seen in the structure of the offsets over time. We are able to sensibly utilize these per-session estimates to additionally correct those APOLLO data runs that did not contain calibration data, due to lunar signal deficits.

Further study of the ACS has revealed it to be of even more value than originally envisioned, and continues to verify that APOLLO data does not contribute to the  $\sim 15$  mm-level uncertainty as indicated by the best model residuals available. We are not only able to track and visualize the variations in APOLLO inaccuracy over long timespans, but also to correct for these variations. Furthermore, we have demonstrated the *need* for such a calibration system, given the

non-static nature of range accuracy corrections. We continue to expect that outcomes of the ACS effort will drive model improvement and eventually realize and/or corroborate order-of-magnitude gains in understanding of fundamental gravity.

The ACS will continue to be operated as frequently as possible during observation sessions in order to facilitate the procedures explained here as well as explore others. A systematic study of the effects on ACS timing as a function of proximity in time to when the range laser fires to explore the potential effects of laser EMI on APOLLO is planned. Adjustments to the APD readout electronics have been recently studied and will soon be installed at the site; We hope for a potential  $\sim 20\%$  gain in system timing precision, improving the precision of the APOLLO system overall. Additionally, the ACS continues to allow for independent systematics testing in an offline-mode (not during ranging) for potential study individual effects of the various hardware components on timing results. It is exciting to have such a powerful tool to explore our apparatus with and we look forward to other future studies.

# **Appendix A**

## **Final notes**

# Bibliography

- [1] Haering R and Zach A “Fiber laser comprising a ring-shaped resonator” U.S. Patent 8 457 164 B2, June 4, 2013
- [2] Nave C R Coupling constants for the fundamental forces URL <http://hyperphysics.phy-astr.gsu.edu/hbase/Forces/couple.html>
- [3] Murphy T W 2013 *Reports on Progress in Physics* **76** 076901 (*Preprint* 1309.6294)
- [4] Nordtvedt K 1968 *Phys. Rev.* **170** 1186–1187
- [5] Bender P L, Currie D, Poultney S K, Dicke R H, Eckhardt D, Kaula W M, Mulholland J D, Plotkin H H, Silverberg E C and Faller J 1973 *Science* **182** 229
- [6] Samain E, Mangin J F, Veillet C, Torre J M, Fridelance P, Chabaudie J E, Féraudy D, Glentzlin M, Pham Van J, Furia M, Journet A and Vigouroux G 1998 *Astron. Astrophys. Suppl. Ser.* **130** 235
- [7] Shelus P J 1995 *IEEE Trans. Geosci. Remote Sensing* **GE-23** 385
- [8] Dickey J O, Bender P L, Faller J E, Newhall X X, Ricklefs R L, Ries J G, Shelus P J, Veillet C, Whipple A L, Wiatt J R, Williams J G and Yoder C F 1994 *Science* **265** 482–490 ISSN 0036-8075 (*Preprint* <https://science.sciencemag.org/content/265/5171/482.full.pdf>)
- [9] Williams J G, Turyshev S G and Boggs D 2012 *Class. Quant. Grav.* **29** 184004 (*Preprint* 1203.2150)
- [10] Müller J, Hofmann F and Biskupek L 2012 *Class. Quant. Grav.* **29** 184006
- [11] Müller J and Biskupek L 2007 *Class. Quant. Grav.* **24** 4533–4538 URL <https://doi.org/10.1088%2F0264-9381%2F24%2F17%2F017>
- [12] Williams J G, Turyshev S G and Boggs D H 2004 *Phys. Rev. Lett.* **93**(26) 261101 URL <https://link.aps.org/doi/10.1103/PhysRevLett.93.261101>
- [13] Murphy T W, Nordtvedt K and Turyshev S G 2007 *Phys. Rev. Lett.* **98**(7) 071102 URL <https://link.aps.org/doi/10.1103/PhysRevLett.98.071102>

- [14] Soffel M, Klioner S, Müller J and Biskupek L 2008 *Phys. Rev. D* **78**(2) 024033 URL <https://link.aps.org/doi/10.1103/PhysRevD.78.024033>
- [15] Müller J, W Murphy T, Schreiber K, J Shelus P, Torre J M, Williams J, Boggs D, Bouquillon S and G Francau a 2012 *Journal of Geodynamics*
- [16] Fienga A, Laskar J, Exertier P, Manche H and Gastineau M 2015 *Celestial Mechanics and Dynamical Astronomy* **123** 325–349
- [17] Pitjeva E V 2013 *Sol. Syst. Res.* **47** 386 (*Preprint* 1308.6416)
- [18] Reasenberg R D, Chandler J F, Colmenares N R, Johnson N H, Murphy T W and Shapiro I I 2017 Modeling and Analysis of the APOLLO Lunar Laser Ranging Data *Proceedings, 7th Meeting on CPT and Lorentz Symmetry (CPT 16): Bloomington, Indiana, United States, June 20-24, 2016* pp 137–140
- [19] Adelberger, E G and Battat, J B R and Birkmeier, K J and Colmenares, N R and Davis, R and Hoyle, C D and Ruixie, L H and McMillan, R J and Murphy, Jr, T W and Schlerman, E and Skrobol, C and Stubbs, C W and Zach, A 2017 *Class. Quant. Grav.* **34** 245008 (*Preprint* 1706.09550)
- [20] Murphy Jr T W, 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. Quant. Grav.* **29** 184005
- [21] Viswanathan V, Fienga A, Manche H, Courde C, Torre J, Exertier P and Laskar J 2016 Updates from INPOP ephemerides: Data reduction model and parameter estimation using IR LLR data from OCA *Proceedings, 20th International Workshop on Laser Ranging: Potsdam, Germany, October 9–14, 2016*
- [22] Hofmann F and Müller J 2016 Updates of the IfE LLR analysis model and new fit of relativistic parameters *Proceedings, 20th International Workshop on Laser Ranging: Potsdam, Germany, October 9–14, 2016*
- [23] Yagudina E 2016 Parameters of new version EPM Lunar ephemeris on the base of LLR observations 1970-2016 years *Proceedings, 20th International Workshop on Laser Ranging: Potsdam, Germany, October 9–14, 2016*
- [24] Hofmann F and Müller J 2018 *Classical and Quantum Gravity* **35** 035015 URL <https://doi.org/10.1088%2F1361-6382%2Faa8f7a>
- [25] Battat J B R, Murphy T W, Adelberger E G, Gillespie B, Hoyle C D, McMillan R J, Michelsen E L, Nordtvedt K, Orin A E, Stubbs C W and Swanson H E 2009 *Publ. Astron. Soc. Pac.* **121** 29
- [26] Liang Y, Murphy T W, Colmenares N R and Battat J B R 2017 *Class. Quant. Grav.* **34** 245009 (*Preprint* 1706.09421)

- [27] 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 (*Preprint* 0710.0890)