Implementing Real-Time Calibration in Advanced LIGO Control Software

Dane Stocks and Joseph Betzwieser
LIGO Livingston Observatory
LIGO SURF Program
DCC T1700318

August 4, 2017

Abstract

The digital error and control signals of Advanced LIGO's differential arm length control servo are used to reconstruct gravitational wave (GW) strain, $h(t)$. Currently, three different calibration pipelines produce $h(t)$ with varying errors and latencies. The real-time operating system in the front end computers runs CALCS, which performs infinite impulse response (IIR) filtering and control operations on 16384 Hz clock cycles. Current limitations of these filters yield systematic errors which a second pipeline, the GDS, corrects in low-latency using finite impulse response (FIR) filtering on computers distinct from the front end computers. The third pipeline, the DCS, implements FIR filtering to condition archived data, and is used to recalibrate entire data sets when dropouts occur in real-time. To prepare for O3 in late 2018, we construct a new, self-contained calibration pipeline in the front end computers which uses FIR filtering to calculate $h(t)$. This new front end model will replace the current calibration system in place, removing the redundancy of multiple, separate pipelines while reducing latency in control room access to calibrated GW strain.

1 Introduction

Advanced LIGO (Laser Interferometric Gravitational-Wave Observatory) is a system of enhanced Michelson interferometers designed to measure relative changes in the spacetime interval between two sets of test masses. The arms of a detector contain 4 km-long Fabry-Perot cavities, which are filled with light from an Nd:YAG 1064 nm laser (Figure 1a). Laser light continuously resonates in these

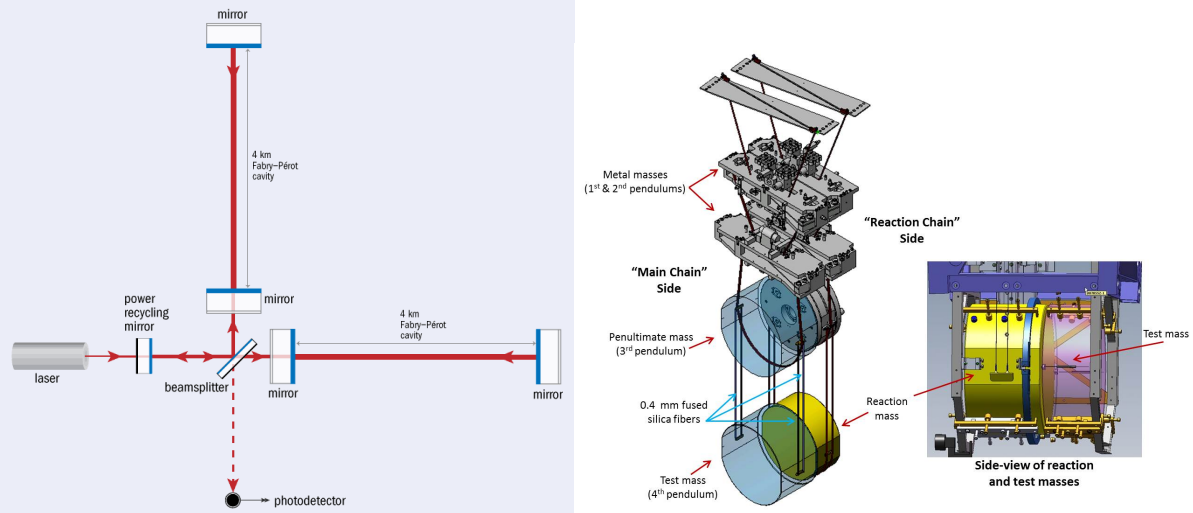

(a) Simplified diagram of an Advanced LIGO detector

(b) Multi-stage pendulum suspension, with both main and reaction chain shown

FIGURE 1: Basic detector infrastructure

cavities, forming extremely stable beams which destructively interfere at the anti-symmetric port of the detector. A passing gravitational wave (GW) produces a differential change in the X and Y arm lengths, ΔL_{ext}, which in turn causes a phase shift in the beams to be recorded at the photodetector. When the cavities are held on resonance, or on "lock," the laser power fluctuations measured at the photodetector are proportional to gravitational wave strain.

In practice, the strain measured by the interferometer is due to various sources of noise, and not just the astrophysical signals of interest. Multiple design measures have been taken to reduce noise and help reach the necessary detector sensitivity. The mirrors at the end of each Fabry-Perot cavity are the bottom stages of quadruple-pendulum systems [6], and the interferometer components rest on active seismic isolation platforms, which assist in keeping lock at low frequencies by reducing motion to the level of 10^{-11}–10^{-12} m Hz$^{-1/2}$ between 1 and 10 Hz [4].

Unwanted residual displacement remains, however, and additional methods are needed to keep the interferometer on resonance: in Advanced LIGO, this is accomplished through the use of five length sensing control loops. The differential arm (DARM) length control loop is the main focus of calibration efforts, as the final gravitational wave time-series $h(t)$ is produced from DARM length, defined as

$$\Delta L_{\text{ext}}(t) = L_{\text{x}}(t) - L_{\text{y}}(t) = h_{\text{ext}}(t)L \tag{1}$$

where $L_{\text{x}}(t)$ and $L_{\text{y}}(t)$ are the distances from the input test mass to the end test

mass of the X and Y arms at time t, $h_{\text{ext}}(t)$ is strain due to gravitational waves and external sources of noise, and L = 3994.5 m is the mean arm cavity length [1].

The outline of this paper is as follows: Section 2 provides an overview of the DARM loop and its components, Section 3 details the structure of the current calibration pipelines used to calculate strain, and Section 4 introduces a new pipeline, designed to run in the front end computers, and describes the progress made in constructing this model and results of its testing thus far.

2 The DARM loop

A real-time control system, the CDS (Control and Data System), operates in the front end computers of the interferometer. Sensor and actuator electronics are placed throughout the detector, and PCIe input/output chassis stream data from these electronics to the front end clusters. The computers in this cluster perform control algorithms and run the DARM length feedback loop in real-time. A digital filter bank, $D(f)$, and two transfer functions comprise this feedback system; the length sensing function, $C(f,t)$, and the actuation function, $A(f,t)$, are complex-valued functions of frequency subject to time-dependent corrections.

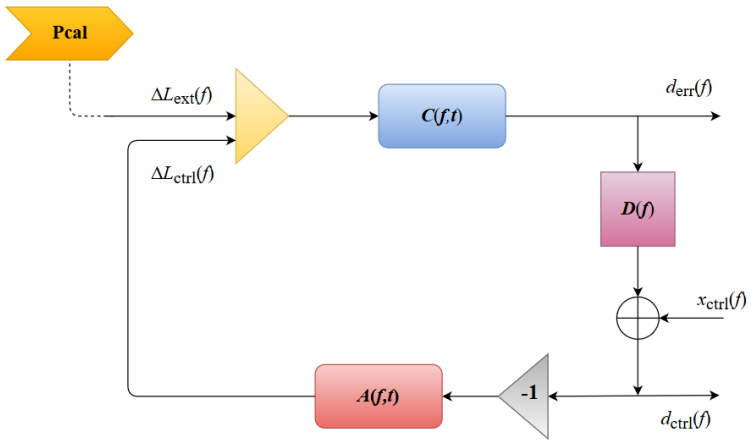

FIGURE 2: The DARM control loop keeps the detector's mirrors on resonance by suppressing external differential arm displacements. The transfer functions $C(f,t)$ and $A(f,t)$ are calibrated using different types of signal injections. Pcal, x_{ctrl}, and x_{T} (not shown in actuation filter block) are injection points used for determining the values of the time-dependent κ_{C}, κ_{T}, and κ_{PU} correction factors.

2.1 Sensing function

The sensing function, $C(f,t)$, describes the interferometer's optical response to $\Delta L_{\rm res}(f)$. That is, it transforms residual displacement in meters into a digtized error signal, $d_{\rm err}$, in arbitrary computer counts.

$$d_{\rm err}(f) = C(f,t)\Delta L_{\rm ext}(f) \tag{2}$$

$d_{\rm err}$ (DARM error) represents laser power fluctuations at the photodetector, and is sampled at a rate of 16384 Hz. The constituent parts of $C(f,t)$ include:

- $\kappa_{\rm C}(t)$ – optical gain, dependent upon the amount of light stored in the Fabry-Perot cavities. It describes how many digital counts are generated in $d_{\rm err}$ given a reference DARM displacement at time t.
- $(1 + if/f_{\rm c}(t))^{-1}$ – approximated coupled-cavity pole between the signal-recycling and arm cavities. Above the pole frequency, $f_{\rm c}$, the finite average storage time of photons in the resonant arm cavities severely attenuates the detector's response to $\Delta L_{\rm res}(f)$. During O1, the pole frequency at LLO was measured to be 388 Hz [1].
- $Q(f)$ – time-independent part of sensing function, including the laser power response of the readout port photodetector, the frequency response of the analog-to-digital converter electronics in the sensing chain, and the time correction $\tau_{\rm C}$ due to light travel in the cavities L/c as well as digitization and computational delays in the electronics.

A simplified representation of the sensing function is:

$$C(f,t) = \frac{\kappa_{\rm C}(t)}{1 + if/f_{\rm c}(t)}Q(f) \tag{3}$$

2.2 Digital filters

A set digital filters transforms the DARM error signal into a DARM control signal in double-point precision [7]:

$$d_{\rm ctrl}(f) = D(f)d_{\rm err}(f) \tag{4}$$

These filters are known exactly, and unlike the sensing and actuation functions, introduce no uncertainty into the overall frequency response function of the interferometer, $R(f)$. $d_{\rm ctrl}$ is passed into the actuation function to produce $\Delta L_{\rm ctrl}(f)$, a displacement which suppresses the sensed residual motion of the interferometer.

2.3 Actuation function

The main-chains of the suspended cascading pendula possess twin reaction-chains that are suspended in parallel (Figure 1b). The top mass of each chain is controlled independently using optical shadow sensors and electromagnetic actuators (OSEMs). For the bottom three suspended masses, actuation is provided separately to each stage of the main-chain via the corresponding mass of the reaction-chain. The DARM control signal, d_{ctrl}, is distributed to each actuator in the form of voltage, and a reaction force displaces each test mass at the ends of the detector: $\Delta L_{\mathrm{ctrl}}(f) = -A(f,t)d_{\mathrm{ctrl}}(f)$, where the minus sign is included by convention. On the upper intermediate (U) and penultimate (P) stages, digital-to-analog converters drive electromagnets to induce a magnetic force on the masses of the main-chain. On the test (T) stage, a digital-to-analog converter drives an electrostatic system which produces dipole-dipole interactions between the test mass and a series of electrodes attached to the corresponding reaction-chain stage [2]. This system is called the electrostatic drive (ESD).

The actuation response of the test mass stage is known to vary over time. Charge accrues on the test masses during observation runs and alters the strength of the ESDs [5]. While the strength of the actuators of the upper intermediate and penultimate stages are not expected to fluctuate with time, they are tracked in the event of unexpected equipment failure. Hence, the time-dependency of the actuation function is tracked via two variables: 1) $\kappa_{\mathrm{T}}(t)$, the test stage actuation scaling factor, and 2) $\kappa_{\mathrm{PU}}(t)$, the combined actuation scaling factor for the P and U stages.

Explicitly, d_{ctrl} is distributed to each actuator and multiplied by a series of digital filters, $F_i(f)$, where $i = \mathrm{T}, \mathrm{PU}$ to represent the test and combined penultimate and upper intermediate stages, respectively. If we define $A_{\mathrm{T}}(f)$ and $A_{\mathrm{PU}}(f)$ as the frequency dependence of each stage's actuator, then the total actuation is written as

$$A(f,t) = \Big[\kappa_{\mathrm{T}}(t)F_{\mathrm{T}}(f)A_{\mathrm{T}}(f) + \kappa_{\mathrm{PU}}(t)F_{\mathrm{PU}}(f)A_{\mathrm{PU}}(f)\Big]e^{-2\pi i f\tau_A} \tag{5}$$

where τ_A is the delay involved with digital-to-analog conversions between d_{ctrl} and the voltages across the actuators.

2.4 Response function

The DARM loop works to keep the differential arm length of the detector constant, and the digital error and control signals the loop outputs are used to characterize differential displacements sensed by the interferometer. Abstractly, we need to accurately model the detector's response, $R(f)$, to infinitesimal external displace-

ments ΔL_{ext} to be able to extract information about incident GW signals:

$$\Delta L_{\text{ext}}(f) = R(f) d_{\text{err}}(f) \tag{6}$$

Inspection of Figure 2 shows that

$$d_{\text{err}} = d_{\text{err}} D(f)(-1) A(f) C(f) + \Delta L_{\text{ext}}(f) C(f)$$

$$d_{\text{err}}\Big(1 + D(f)A(f)C(f)\Big) = \Delta L_{\text{ext}}(f) C(f)$$

$$d_{\text{err}} = \frac{\Delta L_{\text{ext}}(f) C(f)}{1 + G(f)}$$

where for simplicity we have approximated $C(f, t)$ as $C(f)$ and ignored its time dependent factors (same for $A(f, t)$). Additionally, we define the open loop transfer function $G(f) = D(f)A(f)C(f)$. From (6) it follows

$$\Delta L_{\text{ext}}(f) = R(f) \frac{\Delta L_{\text{ext}}(f) C(f)}{1 + G(f)}$$

$$R(f) = \frac{1 + G(f)}{C(f)}$$

We can re-write (6) using (1) to explicitly solve for h:

$$h(f) = \frac{1}{L}\left(\frac{1 + G(f)}{C(f)}\right) \cdot d_{\text{err}}$$

Notice from (4) that d_{err} can be distributed as

$$h(f) = \frac{1}{L}\left(C(f)^{-1} \cdot d_{\text{err}} + D(f)A(f) \cdot \frac{d_{\text{ctrl}}}{D(f)}\right)$$

$$h(f) = \frac{1}{L}\Big(C(f)^{-1} \cdot d_{\text{err}} + A(f) \cdot d_{\text{ctrl}}\Big) \tag{7}$$

$h(t)$ is calculated in Advanced LIGO's current calibration pipelines according to (7), with convolution replacing direct multiplication as the operation between the DARM signals and the transfer function models (as both are converted to the time domain to reduce latency).

2.5 Modeling loop through calibration

Minimizing the uncertainty present in $h(t)$ involves making precise measurements of the parameters in the actuation function A and sensing function C which are not known absolutely.

One method used to make such measurements involves a radiation pressure actuator called a *photon calibrator* (Pcal). A Pcal device placed near each end test mass (ETM) shoots 1047 nm light from a power-modulated Nd^{3+}:YLF laser onto the front surface of the mirror [2], where it reflects and is recorded by an auxiliary photodetector. Photon calibration can produce signals at specific frequencies, $x_T^{PC}(f)$, that are much larger than ΔL_{ext} at that frequency. These signals are known as calibration "lines" becuase they appear as large peaks in spectral density plots of d_{err} [3], and they facilitate high-precision measurements of DARM suppression.

Two additional types of calibration lines are x_{ctrl} and x_T, digital excitations meant to track time dependent paramters associated with the actuation function. While x_{ctrl} is injected directly into d_{ctrl}, x_T is injected only at the test mass actuator stage. A total of six calibration lines are used at LLO to track the time varying elements of the DARM loop, and they are listed in the following table.

Num.	Sym.	Freq. (Hz)	Type	Purpose
1	f_{tst}	35.9	ETM ESD	TST + PU actuation
2	f_{pcal}	34.7	Pcal	DARM control actuation
3	f_{ctrl}	33.7	x_{ctrl}	DARM control actuation
4	f_{pcal2}	331.3	Pcal	DARM sensing
5	f_{pcal3}	1083.1	Pcal	Calibration check
6	f_{pcal4}	3001.1	Pcal	Calibration check

To see how these calibration lines are used to derive the time-dependent correction factors (TDCFs) of the transfer function frequency models, see [8].

3 Current calibration process

Calibration seeks to reconstruct GW strain incident upon the detector using the digital signals d_{err} and d_{ctrl}. For its first and second observation runs (O1 and O2), Advanced LIGO reproduced $h(t)$ using a system of three separate calibration pipelines, which are listed in Table 1. This section details the function, purpose, and characteristics of each pipeline.

1. Front end calibration – CALCS

TABLE 1: An overview of the key differences between each calibration system used in Advanced LIGO, including their respective locations and structure.

Pipeline	**Latency**	**Location**	**Input**	**Filtering**
CALCS	real-time	front end computers	d_{err} and d_{ctrl}	IIR
GDS	O(10 s)	DMT	$\widetilde{\Delta L_{\mathrm{res}}}$ and $\widetilde{\Delta L_{\mathrm{ctrl}}}$	FIR
DCS	O(weeks)	LDAS	d_{err} and d_{ctrl}	FIR

The CALCS system operates in real-time within the front end computers of the interferometer. d_{err} and d_{ctrl} are picked-off from the DARM loop at the locations shown in Figure 2. Because it is located on the same computers as the DARM front end model, CALCS has direct access to the parameters of the infinite impulse response (IIR) filters for C and A which are used by the DARM loop, and these filters can be replicated exactly in the calibration process.

The IIR filters for C^{-1} are applied to d_{err} to produce $\widetilde{\Delta L_{\mathrm{res}}}$, which is output at 16384 Hz. Similarly, d_{ctrl} is branched and distributed to each of three IIR filter modules for the individual actuation functions: A_{T}, A_{P}, and A_{U}. These are the test, penultimate, and upper intermediate actuation functions, respectively. The convolution of d_{ctrl} with each IIR actuation filter produces a corresponding fraction of $\widetilde{\Delta L_{\mathrm{ctrl}}}$:

$$\widetilde{\Delta L_{\mathrm{res}}} = C(f)^{-1} d_{\mathrm{err}}$$

$$\widetilde{\Delta L_{\mathrm{ctrl}}} = \Delta L_{\mathrm{T}} + \Delta L_{\mathrm{P}} + \Delta L_{\mathrm{U}} = A(f) d_{\mathrm{ctrl}}$$

The constituent control displacements are downsampled to 4096 Hz and written out as partially calibrated data products. Simultaneously, $\widetilde{\Delta L_{\mathrm{ctrl}}}$ is sampled at 16384 Hz and delayed to be time synchronous with the calculated $\widetilde{\Delta L_{\mathrm{res}}}$. These displacements are summed to produce $\widetilde{\Delta L_{\mathrm{ext}}}$, the primary output of the CALCS pipeline.

Due to lack of sufficient time for testing before observing runs began, the current front end calibration scheme does not apply time-dependent correction factors (TDCFs). Additionally, features of the calibration such as super-Nyquist properties cannot be accurately modeled with IIR filters, and as such the data products of CALCS contain systematic errors which range from a few percent in the mid-frequency band (50–500 Hz) to nearly 10% at high frequencies (1–3 kHz). These shortcomings of the current front end calibration model necessitate the use of an additional online pipeline to make corrections.

2. Low-latency online calibration – GDS

The Global Diagnostic System (GDS) is an online, low-latency calibration pipeline written in `gstlal` codebase, a combination of LIGO Algorithm Library (LAL) tools wrapped in GStreamer [7]. The partially calibrated $\widetilde{\Delta L_{\mathrm{res}}}$, ΔL_{T}, ΔL_{P}, and ΔL_{U} computed by CALCS are broadcast to the Data Monitoring Tool (DMT) through the Data Acquisition System (DAQ) Data Concentrator, where they are written to frames. The GDS pipeline, located within the DMT, reads these data from shared memory.

Finite impulse response (FIR) filters are used to model the components of the inverse sensing and actuation functions not accounted for in CALCS, such as accurate time delays and the super-Nyquist features of both transfer functions. These filters are called "correction" filters, as they refine the estimates of the residual and control displacements made by CALCS, and are represented by C_{corr}^{-1} and A_{corr}. The process of generating FIR filters from the frequency domain models of the inverse sensing and actuation functions, as given in Section 2, is as follows:

1. *High-pass filter* — components of the frequency spectrum below 9 Hz are "rolled off" by multiplying them by half of a Hann window raised to the fourth power.

2. *Low-pass inverse sensing* — the inverse sensing function tends toward infinity at high frequencies (as the detector's response is attenuated for $f > f_{\mathrm{c}}$), so the frequency components of C above 6 kHz are multiplied by half of a Hann window.

3. *Artificial delay* — an artificial delay equal to half the number of taps in the produced FIR filter is added to center the filter in time. This delay is undone in the GDS pipeline by advancing the corresponding output by an equal number of computer cycles.

4. *Inverse Fourier transform* — after the Nyquist component of the frequency response is zeroed out, the inverse Fourier transform of each frequency domain model produces the desired time domain FIR filter.

5. *Tukey window* — the filter is passed through a Tukey window so that it falls off smoothly near its edges.

In the GDS pipeline, ΔL_{P} and ΔL_{U} are summed before filtering to reduce computational costs. The FIR filters in the actuation channels are longer in length than the inverse sensing filter (6 seconds vs. 1 second), stemming from the fact that significant noise attenuation is needed at lower frequencies ($<\sim$10 Hz). ΔL_{PU} =

$\Delta L_{\mathrm{P}} + \Delta L_{\mathrm{U}}$ and ΔL_{T} are downsampled to 2048 Hz from 4096 Hz to further improve computational efficiency. An applied sinc table minimizes the effects caused by aliasing.

A_{corr} and C_{corr}^{-1} are applied through direct convolution, producing the following displacements:

$$\Delta L_{\mathrm{res}}^{\mathrm{G}} = C_{\mathrm{corr}}^{-1} * \widetilde{\Delta L_{\mathrm{res}}}$$
$$\Delta L_{\mathrm{PU}}^{\mathrm{G}} = A_{\mathrm{corr}} * \Delta L_{\mathrm{PU}}$$
$$\Delta L_{\mathrm{T}}^{\mathrm{G}} = A_{\mathrm{corr}} * \Delta L_{\mathrm{T}}$$

where the superscript "G" reflects that the displacement is generated in the GDS pipeline. After $\Delta L_{\mathrm{PU}}^{\mathrm{G}}$ and $\Delta L_{\mathrm{T}}^{\mathrm{G}}$ are upsampled to the full 16384 Hz rate of the calibration model, the TDCFs discussed in Sections 2.1 and 2.3 are applied to each displacement. The residual displacement $\Delta L_{\mathrm{res}}^{\mathrm{G}}$ is divided by $\kappa_{\mathrm{C}}(t)$ to correct for the optical gain's fluctuations in time. Currently, a single factor $\kappa_{\mathrm{PU}}(t)$ represents the combined time-dependence of the penultimate and upper intermediate actuation strength, while $\kappa_{\mathrm{T}}(t)$ corrects for the varying strength of the ESD caused by charge build-up. Each kappa is sampled at 16 Hz because they are slowly-varying, so they must be upsampled to the rate of the model before being applied.

Following kappa corrections, the GDS pipeline outputs strain:

$$h(t) = \frac{\Delta L_{\mathrm{res}}^{\mathrm{G}}}{\kappa_{\mathrm{C}}(t)} + \Delta L_{\mathrm{PU}}^{\mathrm{G}} \kappa_{\mathrm{PU}}(t) + \Delta L_{\mathrm{T}}^{\mathrm{G}} \kappa_{\mathrm{T}}(t)$$

A calibration state vector, described in [7], monitors the fidelity of the computed $h(t)$ by ensuring that the TDCFs are within an accepted range of values, and that the FIR filters have had sufficient time to settle.

3. *High-latency offline calibration – DCS*

The third and final calibration system used in Advanced LIGO is the Data and Computing Systems (DCS) pipeline. The DCS pipeline is not connected to the online network which establishes communication between the CALCS and GDS systems: it functions offline by reproducing GW strain from archived data. The DCS system is used for posterity; frequent data dropouts occur in the front end computers because of detector maintenance and bugs in system hardware and software, and the need to recalibrate large sets of data after-the-fact is not uncommon.

The DCS pipeline reads raw d_{err} and d_{ctrl} from the DARM loop pick-off and applies FIR filters containing the full reference models of $C^{-1}(f)$ and $A(f)$, not just the correction components as used in the GDS. Additionally, instead of reading

TDCF values computed in the front end computers, a file containing the parameters used to calculate the TDCFs is hard-coded into the pipeline. This corrects potential systematic errors that arise out of dropouts in real-time computation of TDCFs. The high-latency production of GW strain typically occurs weeks to months after the raw data is archived, and is carried out by the LIGO Data Grid computing clusters. For a complete diagram of the Advanced LIGO calibration scheme, see LIGO-G1501518-v15.

4 Complete front end calibration pipeline

The goal of calibration is swift and accurate reporting of $h(t)$: this gives astronomers the maximum response period to perform electromagnetic follow-up of gravitational wave events. A complete search pipeline that runs in the front end computers of the interferometer would provide operators in the control room calibrated strain in extremely low latency. When implemented, this self-contained system will yield $h(t)$ as a raw data product, similar to the way d_{err} and d_{ctrl} are currently produced. In this section, we show that such a pipeline (referred to as the "front end model" after this point) can feasibly replace the current calibration system in use.

The front end model we have built closely replicates the calibration performed by the DCS pipeline. We generate FIR filters to use in the front end system from "L1DCS_1175961600.npz", a file containing measurements of the inverse sensing and actuation function parameters for that GPS time[1]. The inverse sensing filter is 16384 taps in length, while the actuation filter is 12288 taps long. Running at a rate of 16384 Hz, the inverse sensing filter is 1 second long, while the actuation filter is 6 seconds in length as both actuation paths run at the lower rate of 2048 Hz.

1. Inverse sensing chain

Whitened d_{err} is sent into the model from the DARM loop at 16384 Hz. All testing of the front end model to date has taken place on dedicated testing computers at LLO, not the main interferometer (L1). Due to hardware constraints of these computers, the inverse sensing FIR filter initially took too many computer cycles to complete. Consequently, the filter was split to run in parallel on two separate computer cores. In addition to the primary front end model, named `x2calcs1`, an auxiliary model (`x2calcs2`) was created to host the second half of the inverse sensing filter.

Intercommunication between model files takes a single computer cycle. Thus, Sending d_{err} to `x2calcs2` results in a one cycle delay. The delay of the total inverse

[1] This file can be located at the following repository.

sensing filter is defined to be half its length, or 8192 cycles. Sending the second half of the filtered signal back to `x2calcs1` to be summed produces another unit delay, and the inverse sensing chain has a total delay of 8194 cycles.

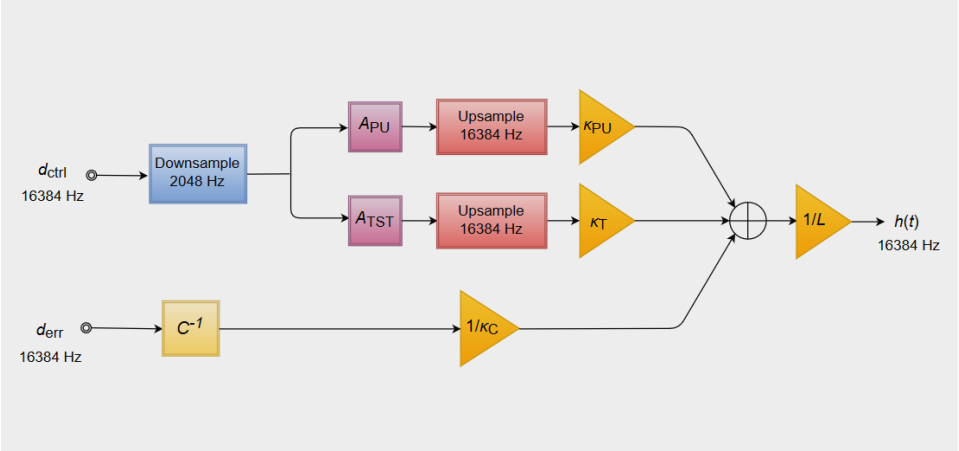

FIGURE 3: A simplified diagram of the complete front end calibration pipeline. Note that the injected kappa values are sampled at 16 Hz, and the cubic spline algorithm used to upsample ΔL^{F}_{PU} and ΔL^{F}_{T} ("F" referencing a displacement generated in the front end model) from 2048 Hz to 16384 Hz is used to resample the kappas to 16384 Hz.

2. Actuation chain

Whitened d_{ctrl} is sent into the model at 16384 Hz. A sinc table, shown in Figure 4a, downsamples d_{ctrl} to 2048 Hz. This sinc table is implemented as an FIR filter and produces a 791 cycle delay. The downsampled d_{ctrl} is then branched into both actuation channels and filtered by A_{PU} and A_{T}. Each filter is 12288 taps long, resulting in a 6144 cycle delay at the lower rate. At the full rate of 16384 Hz, this corresponds to a 49152 cycle delay.

Each filtered signal is then upsampled from 2048 Hz to 16384 Hz using cubic spline interpolation (Figure 4b).

Based on the structure of the code used in implementing the spline algorithm, a delay of 16 cycles is produced during the upsampling process. The actuation path accounts for a total delay of 49959 cycles. The inverse sensing and actuation paths have a relative difference of 41765 cycles, and a ring buffer delays the inverse sensing chain by this amount before summing all three channels to produce ΔL^{F}_{ext}.

3. Evaluating FIR filters

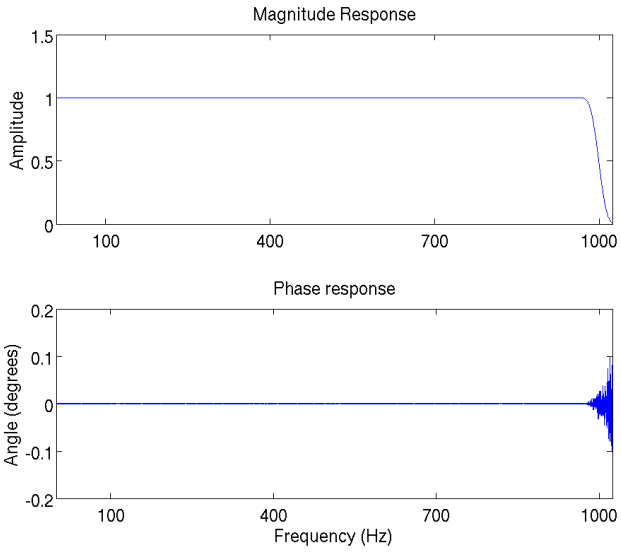

(a) SINC TABLE: the sinc table used to downsample d_{ctrl} from 16384 Hz to 2048 Hz is a sharp low pass filter.

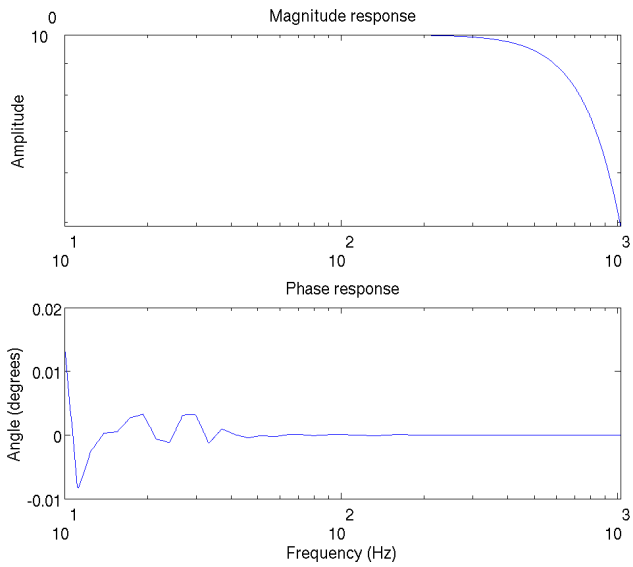

(b) CUBIC SPLINE: note the magnitude response of the cubic spline filter begins to roll off noticeably around $f = 200$ Hz. This introduces a roll off in the overall actuation transfer function.

FIGURE 4: The sinc table and cubic spline FIR filters used in the resampling processes of the front end model.

The frequency domain models of the inverse sensing and actuation functions from the DARM loop provide a benchmark against which the generated FIR filters are compared. This comparison allows the calibration group to quantize the error associated with approximating the transfer functions with FIR filters in the new front end model. Figure 5 displays a comparison of the inverse sensing filter, Figure 6 shows the penultimate/upper intermediate actuation filter responses, and Figure 7 shows the comparison of the test actuation filter.

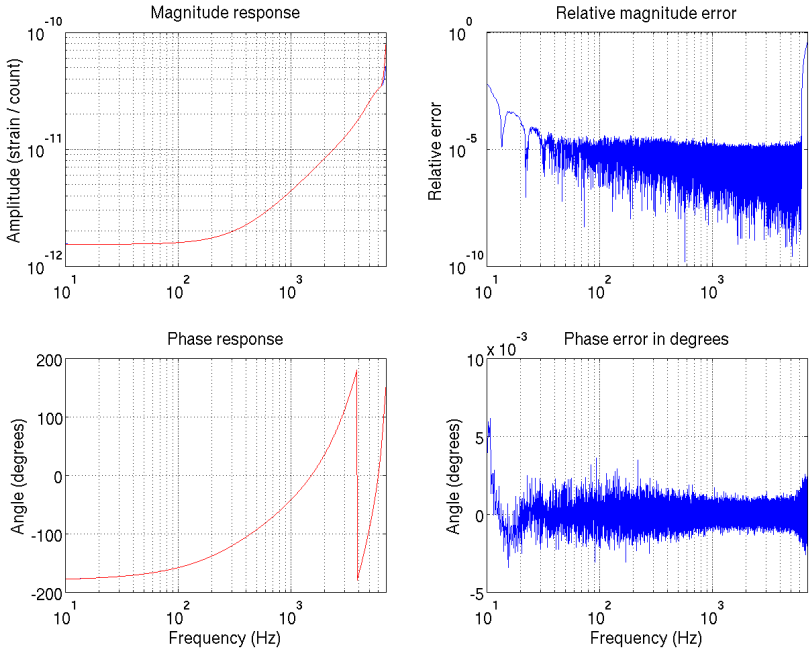

FIGURE 5: INVERSE SENSING FILTER COMPARISON: (top left, bottom left) — the magnitude and phase responses of the front end model FIR filter (blue) are plotted against the frequency response of the DARM loop calibration (red). (top right, bottom right) — error in magnitude response peaks around 10^{-2} (ignoring effects near Nyquist f_{Nq} = 8192 Hz), while error in phase is on the order of 10^{-3}.

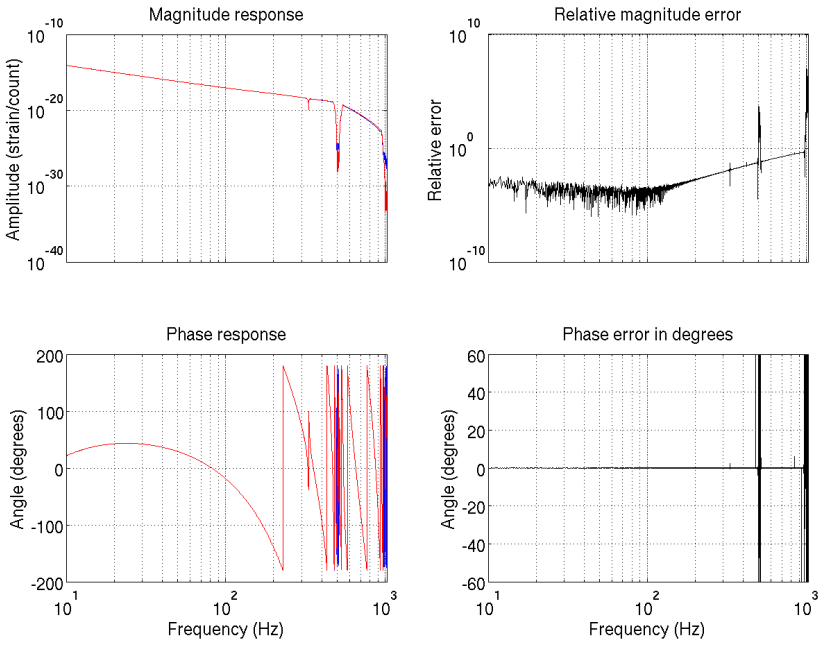

FIGURE 6: ACTUATION PU FILTER COMPARISON: (top left, bottom left) — the magnitude and phase responses of the front end model FIR filter (blue) are plotted against the frequency response of the DARM loop calibration (red). Note that the front end model varies from the DARM loop frequency response primarily at the notches located at $\approx$ 500 Hz and 1000 Hz. (top right, bottom right) — relative magnitude error is on the order of 10^{-3} until about 150 Hz, where an increasing roll off in magnitude response causes greater error.

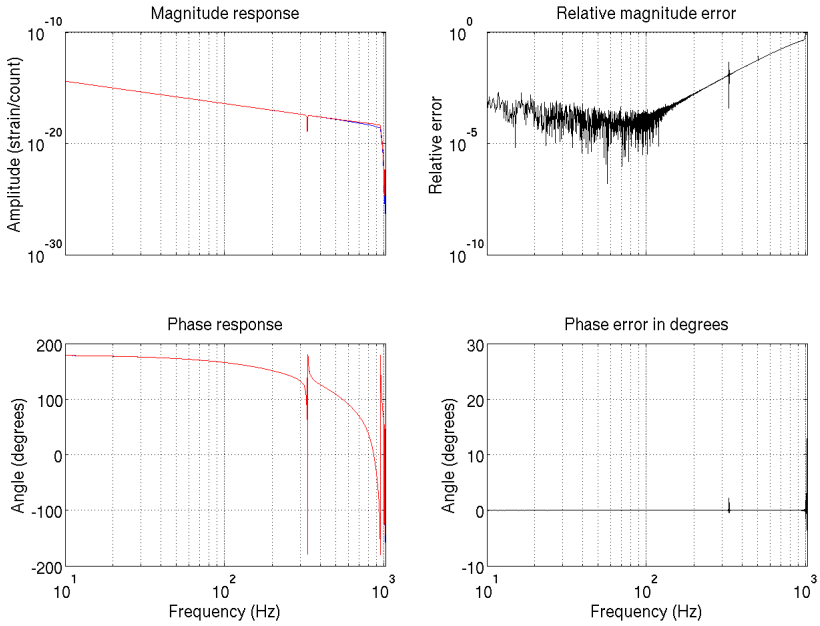

FIGURE 7: ACTUATION TST FILTER COMPARISON: (top left, bottom left) — the magnitude and phase responses of the front end model FIR filter (blue) are plotted against the frequency response of the DARM loop calibration (red). Note the TST filter has the same roll off in magnitude beginning around 150 Hz. (top right, bottom right) — error in magnitude is on the order of 10^{-3}–10^{-4} until the roll off, and phase error is virtually zero.

The roll off in magnitude occurring at high frequencies in the actuation FIR filters is well-understood. Figures 6 and 7 display the transfer function of the full actuation path, which contains the sinc table and cubic spline filters in addition to the actuation filter. As shown in Figure 4b, the cubic spline filter has a roll off beginning around 150 Hz. To verify that the actuation filters we produce in the front end model match the frequency domain model from the DARM loop, we divide out the transfer function of the sinc table and cubic spline. The resultant transfer functions of the corrected A_{PU} and A_{T} filters are shown in Figure 8 and Figure 9.

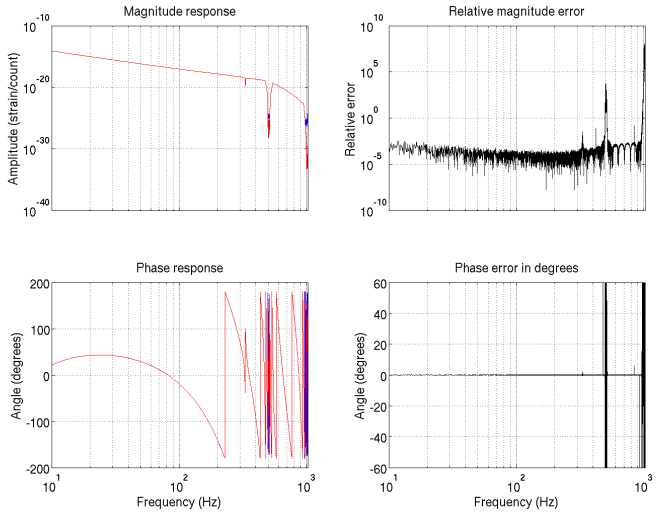

FIGURE 8: ACTUATION PU FILTER CORRECTED: (top left, bottom left) — front end model filter response (blue), DARM loop frequency response (red). Note the roll off in magnitude at high frequencies has been removed. (top right, bottom right) — relative magnitude error is reduced to $\approx 10^{-4}$ throughout the frequency band except at the notches.

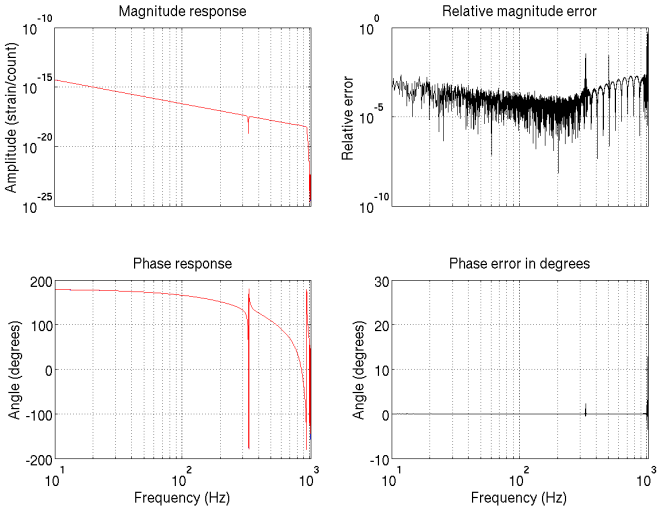

FIGURE 9: ACTUATION TST FILTER CORRECTED: (top left, bottom left) — front end model filter response (blue), DARM loop frequency response (red). Note the roll off in magnitude at high frequencies has been removed. (top right, bottom right) — relative magnitude error is reduced to $\approx 10^{-4}$ throughout the frequency band.

Because we know the roll off of the cubic spline leads to an error in the overall actuation path transfer functions, a systematic error in the strain of the front end model is introduced. For the time being, this error is acceptable because actuation contributes little to strain at high frequencies. For example, at 300 Hz, the A_{T} filter has an error of 1%, and A_{T} contributes only 20% to strain at this frequency (see aLIGO LLO Logbook 29622). Thus, errors in actuation add a 0.2% systematic error to the output strain at 300 Hz (error from A_{PU} filter is even smaller and can be safely ignored).

4. Comparison of calibrated strain

To test the viability of the front end model, we inject 600 seconds of data from the past into the pipeline and check if the output strain matches the strain that the DCS pipeline calculated during this interval of time. Figure 10 displays an amplitude spectral density (ASD) plot which shows how the calibrated strain measurements between the two systems compare across the entire frequency band of Advanced LIGO searches.

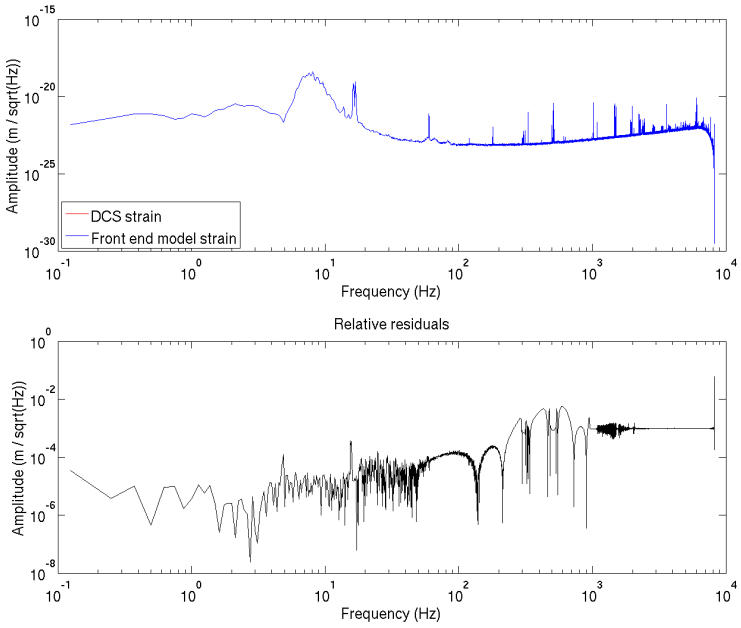

FIGURE 10: (top) — ASD comparison of the strain computed by the DCS calibration pipeline (red) vs. the new front end model (blue). (bottom) — the relative error between the two systems reaches a maximum of ≈ 0.5% near the violin mode frequencies.

The new front end model matches the DCS pipeline to better than 1% across all frequencies. This suggests that the front end model can be used successfully as the primary calibration system during O3, and the redundancy of the CALCS and GDS pipelines can be removed.

4.1 Moving forward

Next week, we will move to L1 and implement the front end model there to begin testing with live data. In the final draft of this paper, we will report the findings of the testing and indicate whether further modifications of the model are needed before Advanced LIGO's next observing run. If the move to L1 is successful, we will begin the process of generating the kappas in this new pipeline.

References

[1] B.P. Abbott *et al.* Calibration of the Advanced LIGO detectors for the discovery of the binary black-hole merger GW150914. Feb. 2016, revised Feb. 2017.

[2] Craig Cahillane *et al.* Calibration Uncertainty for Advanced LIGO's First and Second Observation Runs. Apr. 2016, revised July 2017.

[3] D Tuyenbayev *et al.* Improving LIGO calibration accuracy by tracking and compensating for slow temporal variations. *Classical and Quantum Gravity*, 34(1):015002, 2017.

[4] F. Matichard *et al.* Advanced LIGO two-stage twelve axis vibration isolation and positioning platform. Part 1: Design and production overview. *Precision Engineering*, 40:273–286, Apr. 2015.

[5] M. Hewitson *et al.* Charge measurement and mitigation for the main test masses of the GEO 600 gravitational wave observatory. *Classical and Quantum Gravity*, 24(24):6379, Nov. 2007.

[6] S. M. Aston *et al.* Update on quadruple suspension design for advanced ligo. *Classical and Quantum Gravity*, 29(23):235004, 2012.

[7] Xavier Siemens *et al.* Time Domain Calibration in Advanced LIGO. (LIGO-T1400256-v17), Apr. 2015.

[8] D. Tuyenbayev and the LSC Calibration Team. Tracking temporal variations in the DARM calibration parameters. July 2015, revised Sept. 2015.