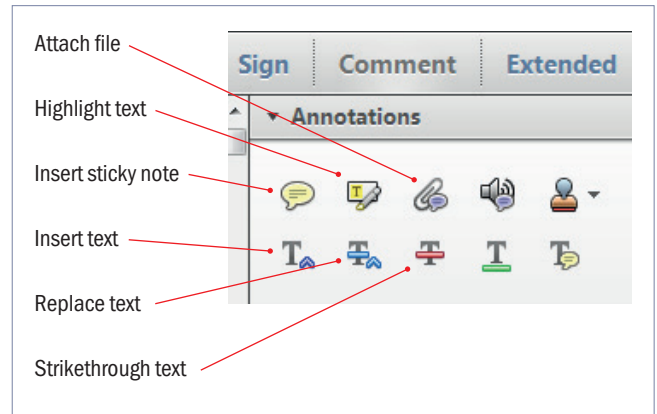


Making corrections to your proof

Please follow these instructions to mark changes or add notes to your proof. You can use Adobe Acrobat Reader (download the most recent version from <https://get.adobe.com>) or an open source PDF annotator.

For Adobe Reader, the tools you need to use are contained in **Annotations** in the **Comment** toolbar. You can also right-click on the text for several options. The most useful tools have been highlighted here. If you cannot make the desired change with the tools, please insert a sticky note describing the correction.

Please ensure all changes are visible via the 'Comments List' in the annotated PDF so that your corrections are not missed.




Do not attempt to directly edit the PDF file as changes will not be visible.




Replacing text

To replace text, highlight what you want to change then press the replace text icon, or right-click and press 'Add Note to Replace Text', then insert your text in the pop up box. Highlight the text and right click to style in bold, italic, superscript or subscript.




Inserting text

Place your cursor where you want to insert text, then press the insert text icon, or right-click and press 'Insert Text at Cursor', then insert your text in the pop up box. Highlight the text and right click to style in bold, italic, superscript or subscript.




Deleting text

To delete text, highlight what you want to remove then press the strikethrough icon, or right-click and press 'Strikethrough Text'.




Highlighting text

To highlight text, with the cursor highlight the selected text then press the highlight text icon, or right-click and press 'Highlight text'. If you double click on this highlighted text you can add a comment.



QUERIES

Page 1


AQ1 

Please check that the names of all authors as displayed in the proof are correct, and that all authors are linked to the correct affiliations. Please also confirm that the correct corresponding author has been indicated. Note that this is your last opportunity to review and amend this information before your article is published.

AQ2 

Please specify the corresponding author and provide his/her email address.

Page 2

AQ3 

Please be aware that the colour figures in this article will only appear in colour in the online version. If you require colour in the printed journal and have not previously arranged it, please contact the Production Editor now.

AQ4 

Please provide an affiliation for author "J Burguet-Castell"

Page 24

AQ5 

If an explicit acknowledgment of funding is required, please ensure that it is indicated in your article. If you already have an Acknowledgments section, please check that the information there is complete and correct.

Page 29

AQ6 

Please check that the funding information below is correct for inclusion in the article metadata. Division of Physics: PHY-1506360, PHY-1607178, PHY-1607585.

Page 30

AQ7 

We have been provided with ORCID iDs for the authors as below. Please confirm whether the numbers are correct. A D Viets 0000-0002-4241-1428, Duncan A Brown 0000-0002-9180-5765.

AQ8 


Only permanent or persistent web links should be used in reference lists. Examples of acceptable links include: Digital Object Identifier (DOI), PubMed identifier (PMID), PubMed Central reference number (PMCID), SAO/NASA Astrophysics Data System (ADS) Bibliographic Code, and arXiv e-print number. Please review your web links and provide updates as required.

AQ9 

Please check the details for any journal references that do not have a link as they may contain some incorrect information. If any journal references do not have a link, please update with correct details and supply a Crossref DOI if available.



AQ10 

Please update the year and publication details if appropriate in references [14, 33, 34, 35, 37, 38, 39, 40, 41, 42].

AQ11 

Please update the publication details if appropriate in references [28, 42].

Reconstructing the calibrated strain signal in the Advanced LIGO detectors

A D Viets^{1,2} , **M Wade**³, **A L Urban**⁴, **S Kandhasamy**⁵,
J Betzwieser⁵, **Duncan A Brown**⁶ , **J Burguet-Castell**,
C Cahillane⁴, **E Goetz**^{7,8}, **K Izumi**⁷, **S Karki**^{7,9}, **J S Kissel**⁷,
G Mendell⁷, **R L Savage**⁷, **X Siemens**¹, **D Tuyenbayev**^{7,10}
and **A J Weinstein**⁴

¹ University of Wisconsin-Milwaukee, Milwaukee, WI 53201, United States of America

² Concordia University Wisconsin, Mequon, WI 53097, United States of America

³ Kenyon College, Gambier, OH 43022, United States of America

⁴ California Institute of Technology, Pasadena, CA 91125, United States of America

⁵ LIGO Livingston Observatory, Livingston, LA 70803, United States of America

⁶ Syracuse University, Syracuse, NY 13244, United States of America

⁷ LIGO Hanford Observatory, Richland, WA 99352, United States of America

⁸ University of Michigan, Ann Arbor, MI 48109, United States of America

⁹ University of Oregon, Eugene, OR 97403, United States of America

¹⁰ University of Texas Rio Grande Valley, Brownsville, TX 78520, United States of America

AQ1

AQ2

E-mail: ~~xxxx~~

Received 27 October 2017, revised 12 March 2018

Accepted for publication 13 March 2018

Published



CrossMark

Abstract

Advanced LIGO's raw detector output needs to be calibrated to compute dimensionless strain $h(t)$. Calibrated strain data is produced in the time domain using both a low-latency, online procedure and a high-latency, offline procedure. The low-latency $h(t)$ data stream is produced in two stages, the first of which is performed on the same computers that operate the detector's feedback control system. This stage, referred to as the front-end calibration, uses infinite impulse response (IIR) filtering and performs all operations at a 16384 Hz digital sampling rate. Due to several limitations, this procedure currently introduces certain systematic errors in the calibrated strain data, motivating the second stage of the low-latency procedure, known as the low-latency `gstlal` calibration pipeline. The `gstlal` calibration pipeline uses finite impulse response (FIR) filtering to apply corrections to the output of the front-end calibration. It applies time-dependent correction factors to the sensing and actuation components of the calibrated strain to reduce systematic errors. The `gstlal` calibration pipeline is also used in high latency to

recalibrate the data, which is necessary due mainly to online dropouts in the calibrated data and identified improvements to the calibration models or filters.

Keywords: gravitational waves, LIGO, calibration, digital signal processing, interferometry, general relativity, astrophysics

AQ3

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

AQ4

1. Introduction

Gravitational waves (GWs) represent a new messenger for astronomy, carrying information about compact objects in the local universe such as neutron stars and black holes. To date, the laser interferometer gravitational-wave observatory (LIGO) and the Virgo detector have observed several transient gravitational-wave signals from merging stellar-mass black hole binaries [1–5] and a binary neutron star system [6], and have recently finished the second observing run (O2) of the Advanced LIGO era. There are two LIGO observatories in North America: one in Hanford, WA (H1) and another in Livingston, LA (L1). The Virgo detector is located in Cascina, Italy.

Each LIGO detector consists of two orthogonal arms, L_x and L_y , roughly 4 km in length (figure 1). External gravitational-wave signals are measured from changes in the differential arm (DARM) length,

$$\Delta L_{\text{free}}(t) = \Delta L_x(t) - \Delta L_y(t), \quad (1)$$

between pairs of test masses at opposite ends of L_x and L_y . Differential length variations are measured using interferometric techniques [8]: the test masses consist of highly reflective mirrors that form a pair of resonant Fabry–Pérot cavities. Input laser light passes through a beamsplitter, enters the resonant cavities, then recombines out of phase at an output photodiode. Any power fluctuations measured by the output photodiode will then correspond to differential arm length changes. To further improve sensitivity, a power recycling mirror at the input reflects laser light back into the arms of the interferometer to increase the laser power stored in the arms. A signal recycling mirror at the output enhances the interferometer’s sensitivity in the frequency band of interest [9–11].

However, we cannot measure $\Delta L_{\text{free}}(t)$ directly because the differential arm length degree of freedom is held in resonance by servos actuating the test mass positions in response to external stimuli, actively suppressing low-frequency fluctuations in $\Delta L_{\text{free}}(t)$ [8]. The servos are part of a feedback control system, known as the DARM control loop, whose residual signal, known as the error signal, is our primary observable, and from which any external length disturbance needs to be reconstructed. In practice, this is done by modeling various transfer functions in the DARM control loop, usually with reference measurements taken prior to the start of science-quality observing runs, and then applying them as filters in the time domain [7, 12].

Calibration was originally done in the frequency domain. The idea of producing a more convenient time-domain calibration came from GEO600 and has been employed since Initial LIGO’s second science run [12]. The time-domain calibration offers two important benefits over a frequency-domain calibration. First, time-domain calibration can be provided with negligible latency, which is what we are striving to achieve, while frequency-domain calibration contains inherent latency. Second, time-domain calibration is a more flexible data product that allows for any length Fourier transform that is needed or desired by data analysis efforts.




Figure 1. Simplified diagram of an Advanced LIGO detector. A set of four highly reflective mirrors (input test mass Y , end test mass Y , input test mass X , end test mass X) act as test masses, forming an orthogonal pair of Fabry–Pérot cavities. Laser light enters from the lower left, passes through a beamsplitter, and enters resonant cavities in each 4 km arm. The laser light then recombines out of phase at the gravitational-wave readout port, which records an error signal used to actuate the end test masses to hold each cavity in resonance. For clarity, only the optics of the lowest suspension stage are shown. *Inset:* an illustration of one of the dual-chain, quadruple pendulum suspension systems with actuators. This figure is reproduced with permission from [7].

Absolute displacement calibration is achieved using auxiliary laser systems, called photon calibrators [13]. They induce fiducial test mass displacements via radiation pressure with force coefficients derived from laser power measurements traceable to SI units. The powers of these laser beams that reflect from the test mass surfaces are measured using calibrated sensors, thus providing continuous absolute calibration when the interferometers are operating in their nominal configurations. The overall $1\text{-}\sigma$ uncertainty in the displacements induced by photon calibrators is 0.75% [13] with the long-term stability of the calibrated length variations verified during year-long observing runs [14]. They thus also provide stable references for continuously monitoring temporal variations in the responses of the interferometers as described in section 2.3

Changes in the responses of the detectors occur due to, for example, slight drifts in the alignment of the optics. These changes are small enough that they do not impact operation of the interferometers. However, if uncompensated, they would cause systematic biases in the

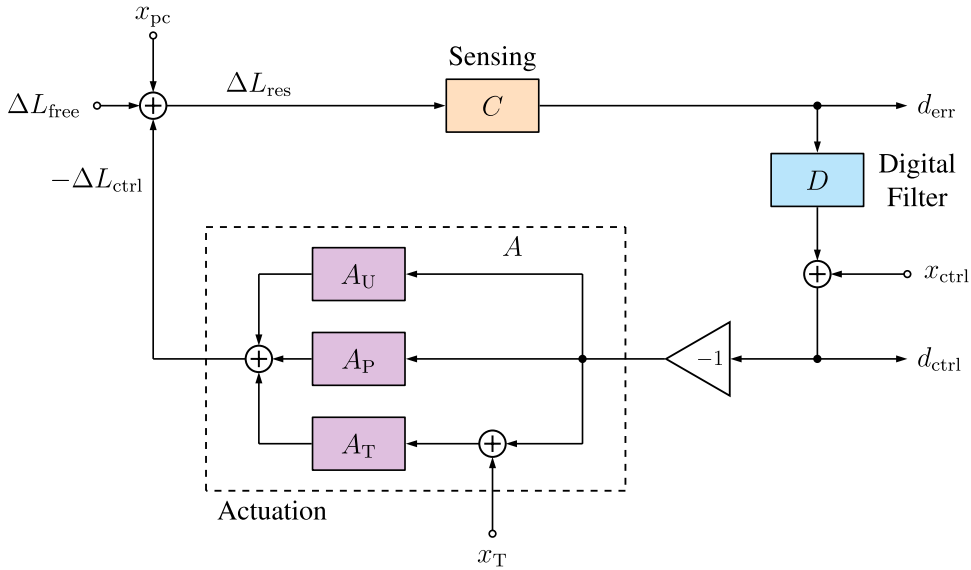


Figure 2. Block diagram of the Advanced LIGO differential arm (DARM) length feedback control loop. Here, the filter C is a sensing function representing the interferometer's response to changes in the measured differential arm length; D is a digital control filter; and the filters A_U , A_P and A_T represent actuation force-to-length responses for the upper-intermediate (U), penultimate (P) and test mass (T) stages of LIGO's quadruple pendulum suspension system. Length disturbances from external sources, including gravitational waves, enter the loop as ΔL_{free} . The controlled length, ΔL_{ctrl} , is subtracted from this to give a length change residual, ΔL_{res} , which passes through the sensing function resulting in a digital error signal, d_{err} . The error signal is then filtered through a set of digital filters D to create a digital control signal, d_{ctrl} . The digital filter D is a linear combination of a low-pass filter and notch filters that prevent excitation of resonances in the test mass suspensions, among other things. The digital control signal d_{ctrl} is used to actuate against the suspension pendula to hold the optical cavities in resonance. Calibration of Advanced LIGO strain data consists of reconstructing ΔL_{free} from d_{err} , d_{ctrl} and models of the filters C^{-1} and $A = A_U + A_P + A_T$. Lastly, to check accuracy and precision we inject known sinusoidal excitations (i.e. calibration lines) at three points in the loop, represented by x_{pc} (injected through radiation pressure by a photon calibrator), x_{ctrl} (injected into the control signal), and x_T (injected into the test mass actuation stage).

calibrated strain data. This would negatively impact estimation of parameters of astrophysical gravitational-wave sources [15]. Fortunately, variations known to have a significant impact, such as laser-power fluctuations in the Fabry–Pérot cavities, are small and slow compared to the rapid timescale (16384 Hz) on which we record and analyze data. To a good approximation, we can treat them as either time-dependent magnitudes of DARM loop transfer functions, or changes in the zero or pole frequencies of those transfer functions [7, 16].

A simplified model of the DARM control loop is shown in figure 2. Any length disturbances that arise from external sources, including gravitational waves and displacement noise, enter the loop as $\Delta L_{\text{free}}(t)$. A controlled differential length, $\Delta L_{\text{ctrl}}(t)$, compensates part of this external displacement, giving a residual displacement

$$\Delta L_{\text{res}}(t) = \Delta L_{\text{free}}(t) - \Delta L_{\text{ctrl}}(t). \quad (2)$$

The residual is directly sensed by the detector and read out as an error signal

$$d_{\text{err}}(t) = C * \Delta L_{\text{res}}(t) \quad (3)$$

where C is a sensing function representing the opto-mechanical response of the interferometer to changes in the DARM length. The operation

$$F * g(t) = \int_{-\infty}^{\infty} F(\tau) g(t - \tau) d\tau \quad (4)$$

denotes convolution of some time domain filter F with some signal $g(t)$. As indicated by equation (3), our model relies on the linearity of the detector's response, which is enforced by holding the differential degree of freedom slightly off anti-resonance. This is done by digitally adding a very small DC offset, and causes a small amount of light to exit the interferometer at the readout port [17].

The error signal is also used to create a control signal, $d_{\text{ctrl}}(t)$. The control signal feeds into the actuation system, which acts on the suspension pendula to hold the Fabry–Pérot cavities in resonance (figure 2). The controlled differential length is

$$\Delta L_{\text{ctrl}}(t) = A * d_{\text{ctrl}}(t) \quad (5)$$

where A is a filter representing the combined electromechanical actuation response from each pendulum stage. In terms of the observable $d_{\text{err}}(t)$ and its corollary $d_{\text{ctrl}}(t)$, the external differential length is therefore

$$\Delta L_{\text{free}}(t) = C^{-1} * d_{\text{err}}(t) + A * d_{\text{ctrl}}(t) \quad (6)$$

where C^{-1} is the filter inverse of C .

The external differential length can also be computed from only the error signal using the response function that relates the DARM length to the error signal: $\widetilde{\Delta L}_{\text{free}}(f) = \widetilde{R}(f) \widetilde{d}_{\text{err}}(f)$, where the tildes denote Fourier transformation into the frequency domain. The response function is given by

$$\widetilde{R}(f) = \frac{1 + \widetilde{A}(f) \widetilde{D}(f) \widetilde{C}(f)}{\widetilde{C}(f)}. \quad (7)$$

Historically, we have used both the error and control signals when reconstructing ΔL_{free} [12], as shown in equation (6), for two main reasons. First, this allows the computation of $\Delta L_{\text{free}}(t)$ to be insensitive to the digital filters D , which are often changed to improve the control loop design¹¹. Second, the predetermined control signal $x_{\text{ctrl}}(t)$, shown in figure 2, will not appear in the resulting $\Delta L_{\text{free}}(t)$ spectrum due to its location in the feedback loop between the $d_{\text{err}}(t)$ and $d_{\text{ctrl}}(t)$ pickoff points.

The dimensionless strain $h(t)$ used for detecting gravitational waves is a time series derived from $\Delta L_{\text{free}}(t)$,

$$h(t) = \frac{\Delta L_{\text{free}}(t)}{L}, \quad (8)$$

where $L = (L_x + L_y)/2$ is the average measured arm length. Since fluctuations in L are negligible, the process of calibration therefore amounts to reconstructing $\Delta L_{\text{free}}(t)$ from $d_{\text{err}}(t)$ and $d_{\text{ctrl}}(t)$. Hereafter we will use $\Delta L_{\text{free}}(t)$ or $h(t)$ interchangeably to refer to the final calibrated data product.

¹¹ Now that time dependent correction factors are included in the calibration process, this is actually no longer true. However, it remains a historical motivation for how methods were developed.

Equation (6) establishes a method for reconstructing calibrated strain data from the error and control signals. We construct accurate physical models for the filters C^{-1} and A (then divide by L , which does not change). In section 2, we outline the important physical properties of transfer functions associated with these filters and describe how we correct for small variations over time. A detailed discussion of the full interferometric response, including its effect on calibration uncertainty, can be found elsewhere [7, 15, 16]. In section 3, we describe the software pipelines used to perform the calibration in Advanced LIGO’s first and second observing runs. In section 4 we discuss plans for future development, and in section 5 we conclude.

2. Frequency domain models

2.1. Sensing function

After ΔL_{ctrl} is subtracted from ΔL_{free} , the sensing function C converts any residual test mass displacement ΔL_{res} into a digitized error signal representing the remaining fluctuations in laser power at the output photodiode. This signal is measured as d_{err} at the gravitational-wave readout port (figure 1) with a sample rate of 16384 Hz. The sensing function also accounts for responses of photodiodes and their analog electronics, light-travel time through the Fabry–Pérot cavities, and digital time delays and downsampling filters.

A standard suite of calibration measurements, as discussed in [7, 13], informs a static reference model for the sensing function. The static reference model is represented by

$$\begin{aligned} \tilde{C}^{\text{static}}(f) &= \left(\frac{H_C}{1 + if/f_{cc}} \right) \\ &\times \left(\frac{f^2}{f^2 + f_s^2 - iff_s/Q} \right) \\ &\times C_R(f) \exp[-2\pi if\tau_C]. \end{aligned} \quad (9)$$

The gain H_C gives the number of digital counts in d_{err} per unit differential length change, in meters. The pole f_{cc} is called the coupled cavity pole frequency; it represents the characteristic frequency beyond which the detector response to gravitational waves is significantly attenuated due to finite average photon storage time in the Fabry–Pérot cavities. Both LIGO interferometers are designed to have the same coupled cavity pole frequencies, but they differ due to differing losses in the optical cavities. During O2, the model reference values were $f_{cc} = 360.0$ Hz at H1 and $f_{cc} = 376.0$ Hz at L1. f_s is the resonant frequency of the optical spring of the signal recycling cavity (SRC), with a model reference value of 6.91 Hz at H1 and 0 Hz at L1 during O2. An optical spring exists in an opto-mechanical cavity if there is a linear relationship between the length of the cavity and the radiation pressure on the mirrors. Q is the quality factor of the SRC. The next parameter in equation (9), τ_C , combines computational delay in acquiring the digital signal, the light-travel time across the length of each arm, and a time advance of $\sim 11.7 \mu\text{s}$ necessary to approximate the majority of the detector response with a single pole. With this advance, the single-pole approximation is negligibly close to full analytical model across the relevant frequency band. The last factor, $C_R(f)$, encodes the remaining frequency dependence above ~ 1 kHz due to photodiode electronics and analog and digital signal processing filters. All parameters are measured before the start of data collection by actuating the test masses with a swept-sine signal, injected through a combination of the photon calibrator (x_{pc}), the control signal (x_{ctrl}), and the test mass actuation stage (x_{T}), shown in figure 2. These measurements are made roughly once every few months and form a set of reference parameters for the full calibration model [7].

Parameters of the sensing function—specifically, the gain H_C , the coupled cavity pole frequency f_{cc} , the spring frequency f_s , and the quality factor Q of the SRC—have been found to vary slowly with time. These factors are all found to change on a timescale of \sim minutes, which is slow compared to the timescale of recording data ($\sim 10^{-4}$ s). The full time-dependent sensing function model is

$$\begin{aligned} \tilde{C}(f; t) = & \kappa_C(t) \left(\frac{H_C}{1 + if/f_{cc}(t)} \right) \\ & \times \left(\frac{f^2}{f^2 + f_s(t)^2 - iff_s(t)/Q(t)} \right) \\ & \times C_R(f) \exp[-2\pi if\tau_C] . \end{aligned} \quad (10)$$

Variations in the gain H_C due to changes in laser power and alignment are captured by the factor $\kappa_C(t)$, which is set to a value of 1 in the reference model and fluctuates by less than 10% with time.

The reference model for C^{-1} , which converts power variations sensed at the anti-symmetric port of the interferometer to differential arm length variations and is represented in the frequency domain by $1/\tilde{C}^{\text{static}}(f)$, is plotted in figure 3.

2.2. Actuation function

The differential arm length is controlled using electromagnetic actuators on the top three pendulum stages and an electrostatic drive actuator on the bottom stage (figure 1). While the topmost stage is connected to a seismic isolation system, its actuation strength is relatively small in the Advanced LIGO sensitive frequency band, and the DARM feedback controls are not sent to this topmost stage. The lowest three stages—referred to as the upper intermediate (U), penultimate (P), and test (T) mass stages—are dominant above 10 Hz, and all these stages are displaced in concert [18–21]. The control signal (d_{ctrl}) is distributed to each stage in parallel, converted from digital counts to force, and used to actuate against the reaction mass at that stage. The net result of this is the controlled length differential, ΔL_{ctrl} (figure 2).

As with the sensing function, calibration measurements inform a static model for the actuation function. The static reference model for the actuation is represented by a counts-to-length transfer function,

$$\tilde{A}^{\text{static}}(f) = [\tilde{A}_U(f) + \tilde{A}_P(f) + \tilde{A}_T(f)] \exp[-2\pi if\tau_A] \quad (11)$$

where $\tilde{A}_i(f)$ represents the frequency response of the i th suspension stage actuator. Filtering that converts digital counts to actuation strength and splits the frequency content of d_{ctrl} across each stage of the actuation is folded into these functions for brevity. The low frequency content of d_{ctrl} is directed to the higher stages of the actuation system and the high frequency content is directed to the lower stages. The computational time delay τ_A accounts for digital-to-analog conversion. The modeled actuation transfer functions are plotted in figure 3.

The actuation function also has time-dependent gains for each stage. The fully time-dependent actuation function model is

$$\begin{aligned} \tilde{A}(f; t) = & [\kappa_U(t)\tilde{A}_U(f) + \kappa_P(t)\tilde{A}_P(f) + \kappa_T(t)\tilde{A}_T(f)] \\ & \times \exp[-2\pi if\tau_A] . \end{aligned} \quad (12)$$

The frequency-independent factors $\kappa_U(t)$, $\kappa_P(t)$, and $\kappa_T(t)$ represent slowly-varying scalar fluctuations over time. Changes are typically at the level of a few percent and vary on the




Figure 3. Frequency-domain models corresponding to the inverse sensing (C^{-1} , *left*) and actuation (*right*) transfer functions as measured at LIGO Hanford Observatory (H1). The inverse sensing function converts d_{err} counts to residual displacement, while the actuation function converts d_{ctrl} counts to controlled displacement. Note, the factor of $\sim 10^9$ difference between inverse sensing and total actuation around 10 Hz is due to a very large gain applied to d_{err} to create d_{ctrl} . At low frequency the inverse sensing is dominated by a low pass filter around 8 Hz, while above ~ 340 Hz it is affected by f_{cc} and above ~ 1 kHz by analog-to-digital conversion. In the actuation response, there are several notch filters used to avoid mechanical resonances. Above ~ 30 Hz, only the test mass stage (T) is actuated against. The corresponding curves at LIGO Livingston Observatory (L1) are qualitatively similar, but differ slightly in scale and frequency dependence; see e.g. [7].

timescale of \sim hours. κ_T represents the time-dependence of the electrostatic drive at the test mass stage of the actuation chain. Physically, this time variation is understood to be due to charge build-up on the electrostatic drive. The effect of κ_P and κ_U has so far only been tracked as a combined effect κ_{PU} , which we have observed to vary very little with time. Values of κ_{PU} that stray from unity are understood to originate from errors present in the model and occasional small changes in computational delays in the DARM feedback loop.

Figure 4 shows the relative strength of each component of the calibration across the aLIGO sensitive frequency range. The contribution of the actuation function to the final calibration diminishes above several hundred Hertz, after which the calibration is dominated by the sensing function. In the 30–300 Hz range, the actuation and sensing both contribute significantly to ΔL_{free} .

2.3. Tracking and compensating for small variations

The calibration models vary in time as the detectors fluctuate and settle throughout operation. While in Initial LIGO the only time-dependent correction applied to the calibration was the gain of the sensing function [12], compensation for temporal variations is more complex in Advanced LIGO. The time-dependent correction factors (TDCFs) for the current instruments include κ_C , f_{cc} , κ_T , κ_{PU} , f_s , and Q [16]. In the H1 detector, the optical spring frequency f_s and quality factor Q of the SRC are tracked, but at L1, $f_s \approx 3$ Hz, which is below the sensitive frequency range of the Advanced LIGO detectors, and so the time-dependence of SRC detuning parameters is not currently tracked at L1.




Figure 4. Amplitude spectral densities for $\Delta L_{\text{res}} = C^{-1} * d_{\text{err}}(t)$, $\Delta L_{\text{T}} = A_{\text{T}} * d_{\text{ctrl}}(t)$, $\Delta L_{\text{P}} = A_{\text{P}} * d_{\text{ctrl}}(t)$, $\Delta L_{\text{U}} = A_{\text{U}} * d_{\text{ctrl}}(t)$, and ΔL_{free} across the relevant frequency band at H1. These were obtained as intermediate products from the high-latency `gstlal` calibration pipeline (section 3.3). Note the visible calibration lines, indicated with arrows: $f_{\text{T}} = 35.9$ Hz is injected using the electrostatic drive, and $f_1^{\text{pc}} = 36.7$ Hz, $f_2^{\text{pc}} = 331.9$ Hz, and $f_3^{\text{pc}} = 1083.7$ Hz are injected using the photon calibrator.

The TDCFs are measured by injecting loud, known sinusoidal excitations (i.e. calibration lines) and tracking their amplitude and phase in d_{err} over time. While the calibration line placement varies in frequency at H1 compared to L1, calibration lines at both detectors serve similar purposes. Figure 2 shows the injection point of each calibration line. Length excitations due to all injected signals are suppressed by the DARM control loop, and their influence on d_{err} can be predicted with models for C , D , and A . Appendix A discusses how each TDCF is computed using the calibration lines.

Calibration accuracy can be quantified using the calibration lines injected with the photon calibrator, through comparison of their amplitude and phase as seen in the photon calibrator injection channel and in $h(t)$. Amplitude spectral densities (ASDs) of $h(t)$ near each photon calibrator line frequency are shown in figure 5 to demonstrate the effect of correcting for small detector variations. The data shown are taken from the L1 detector over a period of 72h surrounding binary black hole merger GW170104 [4]. Consecutive Fourier transforms were used to compute strain data in the frequency domain, using a relatively long integration time of 100s in order to average out any short-duration transient signals (which might otherwise contaminate the measured line height)¹² and to give a frequency resolution of 0.01 Hz. Without compensating for time variations, there can be up to a 15% systematic error in magnitude and 1–4° systematic error in phase across Advanced LIGO’s sensitive frequency range, which can impact parameter estimation [3, 22–31]. Improved calibration is obtained by compensating for these time-varying parameters. As with taking reference measurements of the sensing and actuation functions, accurate tracking of the TDCFs during observation relies on the accuracy of the photon calibrator.

¹² Transient gravitational wave signals could be events like GW170104, which swept from 30 Hz to 500 Hz in less than a second. These are signals that only last in our detector for a short period of time but can in general span a wide range of frequencies. The sinusoidal injections, by contrast, only exist at a single frequency but last throughout the observing run.

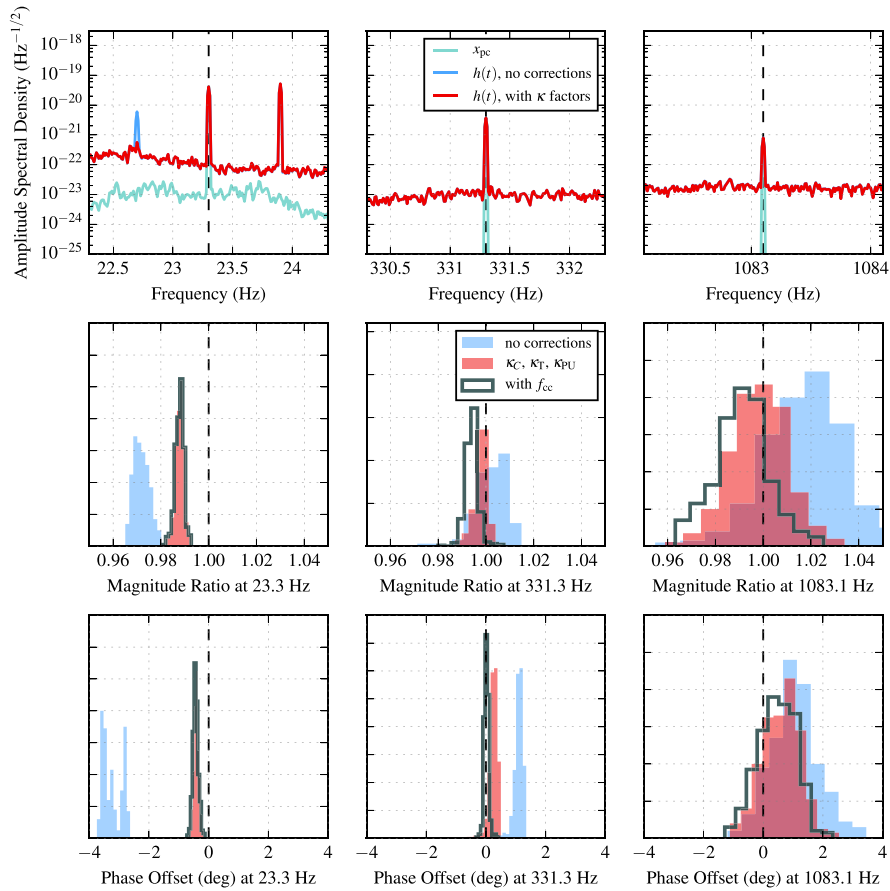


Figure 5. Impact on calibration accuracy from small deviations in the L1 detector over time. Data shown were collected between GPS seconds 1167 536 722 and 1167 695 629, on 4 and 5 January 2017. These deviations are quantified by scalar time-dependent factors $\kappa_C(t)$, $\kappa_T(t)$, and $\kappa_{PU}(t)$, which are applied to the time domain inverse sensing and actuation filters. In addition, the time-dependence of the coupled cavity pole frequency f_{cc} is tracked but not compensated for. In the H1 detector, the optical spring frequency f_s and quality factor Q of the SRC are also tracked but not compensated for. *Top:* amplitude spectral densities of $h(t)$ near the calibration line frequencies. Three of these lines—which live at low, mid and high frequencies, respectively—are injected through radiation pressure against the y -end test mass using a photon calibrator (x_{pc}). *Middle:* Histograms of the magnitude ratio $|\tilde{h}(f)|/|\tilde{x}_{pc}(f)|$, measured in the frequency domain at each x_{pc} line. A series of 100 s Fourier transforms of $h(t)$ and $x_{pc}(t)$ were computed covering 72 h of L1 data surrounding the binary black hole merger signal GW170104 [4]. Note that uncorrected strain contains 1–2% systematic errors in magnitude at all three x_{pc} lines, which is addressed by compensating for TDCFs. The scale of random errors in these measurements is set by the signal-to-noise ratio of the injected signal. However, correcting for f_{cc} would introduce an overall $\sim 1\%$ magnitude offset at mid and high frequencies due to systematics in the measurement of the sensing and actuation functions [15]. *Bottom:* Histograms of the phase offset between \tilde{x}_{pc} and \tilde{h} at each x_{pc} line over the same period. Note that correcting for TDCFs and f_{cc} completely removes the systematic error in phase at 331.3 Hz, which is in Advanced LIGO’s most sensitive band. Results from H1 are similar, although its calibration lines are at slightly different frequencies. During O2, there was also a larger variation in f_{cc} at H1 compared to L1.

Table 1. Summary of the significant differences and inherent purpose of the different calibration pipelines. The number in parentheses next to each pipeline name indicates the numerical shorthand assigned to that pipeline. For example A^0 indicates that this is the actuation function used by the front-end pipeline and ΔL_{free}^0 is the calibrated ΔL_{free} reconstructed by the front-end pipeline. The accuracy listed is the additional systematic error in magnitude when compared to the high-latency `gstlal` calibration pipeline, which produces the most accurate $h(t)$ we are capable of producing at the time of writing. This is not the total calibration uncertainty for the $h(t)$ produced by each pipeline. A full study of the total calibration uncertainty can be found in [15].

Pipeline	Latency	Accuracy	Input	Type of filtering	Purpose
<code>Front-end</code> (0)	Real-time	$\lesssim 14\%$	d_{err} and d_{ctrl}	IIR	Immediate commissioning feedback
Low-latency <code>gstlal</code> (1)	$\mathcal{O}(10\text{s})$	$\lesssim 2\%$	Output of <code>front-end</code> (ΔL_{res}^0 and ΔL_{ctrl}^0)	FIR	Low-latency searches
High-latency <code>gstlal</code> (2)	$\mathcal{O}(\text{weeks})$	n/a	d_{err} and d_{ctrl}	FIR	Final science results

Having sketched the computations required for generation of a precisely calibrated $h(t)$ data stream, we now discuss how this is accomplished in Advanced LIGO.

3. Calibration pipelines

Rapid analysis of data is essential in the era of gravitational-wave astronomy [32] and a key ingredient to this is providing high-precision calibration as quickly as is possible to the analysis pipelines. There are three different methods in place for producing calibrated $h(t)$ data during aLIGO's first and second observing runs, each with varying degrees of precision and latency. The method for producing the lowest latency calibrated data involves using the front-end computer system, which is the system that controls the instrument feedback loops. However, we do not yet have the software developed to provide a calibrated data product from the `front-end` that also meets the calibration accuracy requirements for gravitational-wave data analysis activities. Therefore, a more accurate low-latency calibration is produced using a combination of calibration software written for the front-end computers, known as the front-end calibration pipeline, and a `GStreamer`-based [33] pipeline, known as the `gstlal` calibration pipeline, that improves on the accuracy of the front-end calibration at the cost of adding a few seconds of latency.

Additionally, there is often a need to recalibrate the data at a later point in time. The most common reasons for recalibrating the data are dropouts of data somewhere in the low-latency system and improved models for the calibration developed over time. The high-latency, offline calibration is therefore the most accurate calibrated strain data produced for Advanced LIGO. The same `gstlal` calibration pipeline that is used in the low-latency procedure is used for the offline calibration, except it is run in a different mode with a few varying steps and different data and filter inputs.

Table 1 summarizes the differences between each calibration pipeline. Figure 6 shows a comparison of the amplitude spectral density of each calibration pipeline and figure 7 shows a comparison of the response function as derived from each calibration pipeline to the reference model response function. From these figures, it is clear that the improvement from the front-end calibration pipeline to the low-latency `gstlal` calibration pipeline is significant,

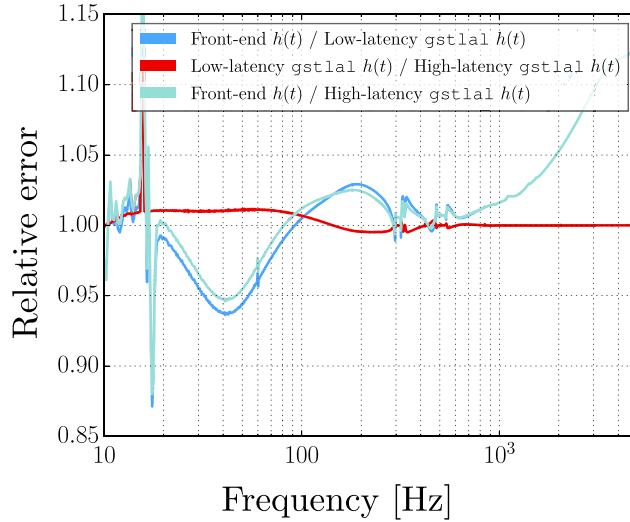


Figure 6. A comparison of the output from the front-end calibration pipeline, low-latency `gstlal` calibration pipeline, and high-latency `gstlal` calibration pipeline at L1. The plot shows the ratio of the ASDs of the front-end $h(t)$ to the low-latency `gstlal` $h(t)$ and the ratio of the ASDs of the low-latency `gstlal` $h(t)$ to the high-latency `gstlal` $h(t)$. The calibration undergoes about a 5-10% improvement between the front-end $h(t)$ and the low-latency `gstlal` $h(t)$ and an additional few percent improvement between the low-latency `gstlal` $h(t)$ and the high-latency `gstlal` $h(t)$. The biggest cause for the improvement from the front-end $h(t)$ to the low-latency `gstlal` $h(t)$ is the application of some of the TDCFs in the low-latency `gstlal` $h(t)$. The difference between the low- and high- latency `gstlal` calibration pipeline products is caused primarily by flaws in the front-end pipeline’s actuation filters that are uncompensated in the `gstlal` calibration pipeline. This is, however, not the case at H1, where the low- and high-latency `gstlal` calibration pipelines produce nearly identical outputs across the relevant frequency band.

while the improvement from the low-latency `gstlal` calibration pipeline to the high-latency `gstlal` calibration pipeline is more minor. While figure 7 does illustrate how the response function as derived from each calibration pipeline compares to the reference model response function, this figure does not portray the systematic and statistical uncertainty inherent to the reference model. For a detailed discussion of the full calibration uncertainty, see [15].

Notation: Throughout the following sections, quantities computed by the front-end calibration pipeline will be denoted with a superscript 0, quantities computed by the low-latency `gstlal` calibration pipeline will be denoted with a superscript 1, and quantities computed by the high-latency `gstlal` calibration pipeline will be denoted with a superscript 2. This notation also distinguishes the filters used in each pipeline. For example, C^0 is the sensing function IIR filter used in the front-end calibration pipeline and C^2 is the sensing function FIR filter used in the high-latency `gstlal` calibration pipeline.

3.1. The front-end calibration pipeline

The front-end calibration pipeline has the advantage of being directly hooked into all of the other front-end computer systems, thereby allowing seamless access to all of the appropriate instrument models and parameters. This enables the calibration model to remain up-to-date




Figure 7. A comparison of the response function derived from the output of the front-end calibration pipeline, low-latency `gstlal` calibration pipeline, and high-latency `gstlal` calibration pipeline at L1 to the reference model response function. The plot on the left shows the comparison across all frequencies in order to display the effects of the low- and/or high-pass filters used by each pipeline. The plot on the right is zoomed in to only display the frequency range (10 Hz, 5 kHz). This comparison captures the differences in the derived FIR and IIR filter representations of the calibration actuation and sensing functions. The effects of TDCFs is not included in this comparison.

and in-sync with the instrument. However, the front-end calibration infrastructure is limited by two features common to real-time computer systems. Originally only delays that were an integer number of digital sampling time intervals, where the digital sampling time interval for a 16384 Hz channel is 61 μ s, could be easily implemented using the front-end system. This hurdle has been overcome and will be discussed more below. Additionally, super-Nyquist poles cannot be easily modeled in the front-end calibration, which leads to a systematic error at high frequencies. The super-Nyquist poles are due to real features of the analog electronics and [cannot](#) be modeled in the digital front-end system because the poles are above the Nyquist frequency of the front-end system.

3.1.1. IIR filters The front-end calibration pipeline uses infinite impulse response (IIR) filtering techniques. Each component of the static reference models for the inverse sensing function and the actuation function [are](#) modeled using zero-pole-gain or second-order section IIR filters. Since it is hard to fit the transfer function measurements with just zero-pole-gain or second-order-section models, these filters introduce additional systematic errors in the calibrated data.

The IIR filters used for the actuation function are exact copies of the digital IIR filters used in the actuation path for interferometer controls whenever possible. The full suspension model is simplified, reducing the Q factor of some resonances while removing other high-frequency resonances, such as violin modes, before converting from a continuous representation to a discrete, IIR filter representation. The inverse sensing function is modeled using a single zero, a DC gain, and a high-frequency roll-off. In order to convert from the continuous representation of the actuation and inverse sensing functions to the discrete representation, we use the built-in MATLAB functions `c2d` with the bilinear (Tustin) method, which does the continuous to discrete conversion [34], and `minreal`, which removes very close zero/pole pairs [35]. High-frequency effects in the actuation and inverse sensing paths are not replicated in the front-end calibration pipeline.

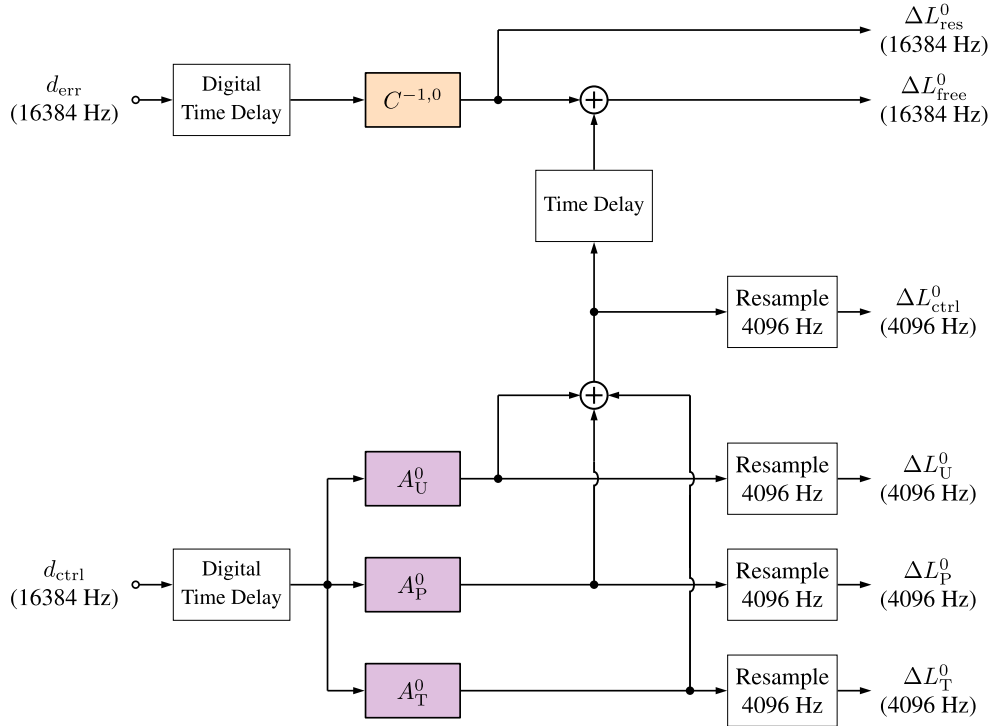


Figure 8. Simplified diagram of the front-end calibration pipeline. The digital error d_{err} and digital control d_{ctrl} signals are picked off from the DARM model in the front-end computers. This process adds a single digital-sampling-period delay in the d_{err} and d_{ctrl} signals, shown above as the ‘Digital Time Delay’ box. The error and control signals are then filtered with the relevant IIR filter representations for the inverse sensing function and the actuation function, respectively. The control signal path is delayed by an amount representing the net delay between the control and error signal paths before being combined with the error signal to form ΔL_{free}^0 . The filtered control and error signals are also outputted from the pipeline. Details of the calculation of the TDCFs are not shown.

3.1.2. Front-end calibration pipeline overview. The function of the front-end calibration pipeline is to pick-off the signals d_{err} and d_{ctrl} and use equation (6) to calculate the calibrated residual signal ΔL_{res}^0 and the calibrated control signals ΔL_{U}^0 , ΔL_{P}^0 , and ΔL_{T}^0 , where the combined calibrated control signal is $\Delta L_{\text{ctrl}}^0 = \Delta L_{\text{U}}^0 + \Delta L_{\text{P}}^0 + \Delta L_{\text{T}}^0$. The final calibrated output of the front-end calibration is ΔL_{free}^0 . Here we outline the basic steps involved in computing each of these outputs, and these steps are shown pictorially in figure 8.

Pick off error and control signals. Since the error and control digital signals d_{err} and d_{ctrl} are picked off from the DARM front-end model and passed to the front-end calibration model, there is a one digital-sampling-period delay.

Apply IIR filters. The IIR filters for the inverse sensing function are applied to the d_{err} signal to produce ΔL_{res}^0 , which is also saved as an output at a sample rate of 16 384 Hz. The d_{ctrl} signal is split into three equivalent paths. One path is filtered with the IIR filter models for the upper intermediate actuation function A_{U}^0 , another one is filtered with the IIR filter models for the

penultimate actuation function A_p^0 , and the third one is filtered with the IIR filter models for the test actuation function A_T^0 .

Resample $\Delta L_{p,U,T}^0$ outputs. After d_{ctrl} is filtered with the three stages of the actuation, the output of each of these filtering processes is resampled from 16384 Hz down to 4096 Hz and saved as the outputs ΔL_p^0 , ΔL_U^0 , and ΔL_T^0 . This downsampling is done to reduce the amount of storage memory required to save these data.

Combine and apply delay for ΔL_{ctrl}^0 . The filtered output of the actuation path (before it is resampled down to 4096 Hz) is combined together to form ΔL_{ctrl}^0 at a sample rate of 16384 Hz. The relative delay between the actuation and sensing paths is then applied to ΔL_{ctrl}^0 to ensure they combine with the correct phase. This delay requires sub-sample phase accuracy to more accurately represent the difference in the delay along each path and is implemented using a Thiran fractional delay filter [36].

Compute ΔL_{free}^0 . The output ΔL_{res}^0 and ΔL_{ctrl}^0 are added together to form ΔL_{free}^0 at a sample rate of 16384 Hz. This is the final output of the front-end calibration pipeline.

3.1.3. Time-dependent correction factors. The infrastructure to track the TDCFs is built into the front-end calibration pipeline. In order to compute the TDCFs, several data streams need to be demodulated at given calibration line frequencies. The demodulation technique used by the front-end calibration pipeline is to first mix a local oscillator at the given frequency with the incoming signal, then apply a low-pass filter with a corner frequency of 0.1 Hz to the mixed output. The demodulated lines are used to compute the TDCFs as described in equations (A.7), (A.8), (A.10), and (A.11). Finally, a 128s running average is performed using a low-pass filter with a corner frequency of ≈ 0.008 Hz, similar to the procedure used in the `gstlal` calibration pipeline, described in appendix B. The second low-pass filter is necessary to avoid sporadic transients that could impact the $h(t)$ data stream if corrections were applied. The TDCFs are then downsampled to 16 Hz and recorded for comparison with the TDCFs computed in the `gstlal` calibration pipeline.

The TDCFs computed in the front end are not used to constantly correct the sensing and actuation functions in any of the calibration pipelines. This is mainly due to the long response time of the IIR filters used in the front-end models that suffer from lasting effects of small transients.

3.1.4. Relative accuracy of the front-end calibration pipeline. A clear benefit of the front-end calibration process is that calibrated strain data can be produced with sub-second latency. Since the future of gravitational-wave astronomy relies on very low-latency detection candidate identification, we are working on moving the entire calibration process into the front-end pipeline. However, work to improve the accuracy of the front-end calibration pipeline to the desired level is still underway.

While the current level of accuracy of ΔL_{free}^0 is acceptable for commissioning purposes, it is relatively poor when compared to the other calibrated data products discussed in the next sections and does not yet meet the standards required for many astrophysical analyses. Systematic errors are introduced by the IIR filters, and it is challenging to build IIR filters that accurately model super-Nyquist features of the calibration, such as high-frequency poles in the sensing function due to a transimpedance amplifier and analog whitening filters. We have, however, had good success in modeling these features using FIR filters. Additionally, we can

not apply a phase advance in the real-time code used by the front-end calibration pipeline, and this results in initially recorded data with a fixed and known number of cycles of delay, relative to their actual occurrence in the interferometer. The sum total of these systematic errors is shown in figure 7, which shows a systematic error as large as 11% in magnitude at high frequencies and a very large phase error across the frequency band (10 Hz, 5 kHz).

The front-end calibration pipeline currently cannot apply the TDCFs to correct the sensing and actuation functions also due to limitations involved with using IIR filtering. The systematic error introduced by not applying the TDCFs is evident from figure 6 in the frequency range from 20 Hz–300 Hz as a $\sim 5\%$ error in magnitude.

3.2. The low-latency `gstlal` calibration pipeline

Some of the shortfalls of the front-end calibration procedure that limit the accuracy of the data are compensated for in a calibration pipeline called the low-latency `gstlal` calibration pipeline. It takes the output data streams from the front-end calibration pipeline, and after applying corrections, it produces the calibrated $h(t)$ time series with latency on the order of several seconds. During O2, the total latency of the `gstlal` calibration pipeline was improved from 10–14 s to 5–9 s. The software package used for this second low-latency calibration pipeline is `gstlal`, a package that wraps LIGO algorithm library (LAL) software [37] with the audio/video streaming software package `GStreamer` [33, 38]. The `gstlal` calibration pipeline applies high-frequency corrections that could not be applied in the front-end calibration, applies the appropriate time delay to each component of the calibration, compensates for the scalar TDCFs and computes a state vector that records the integrity of $h(t)$. An example release version of the `gstlal` calibration pipeline can be found in [39].

The `gstlal` calibration pipeline is also used for all high-latency recalibration that may be required. Both the low- and high-latency `gstlal` calibration pipelines use finite impulse response (FIR) filters. The reasons for using FIR filters instead of IIR filters, which are used by the front-end calibration pipeline, are the following: (1) When recalibrating the data offline, it is useful to recalibrate the data in short, parallelized stretches, which requires FIR filters as opposed to IIR filters; (2) The low-latency calibration system uses both a primary calibration pipeline setup and a redundant pipeline setup to safeguard against online system failures, and in order for the output of the primary and redundant pipelines to be identical, the pipelines cannot depend on their start time or long-term history; (3) All calibration products outside of the front-end need to be reproducible for arbitrarily chosen start times and durations.

3.2.1. FIR filters. The FIR filters used by the low-latency `gstlal` calibration pipeline model the super-Nyquist features of the inverse sensing function that are not captured by the front-end calibration pipeline, corrections for systematic errors introduced by the IIR filters used in the front-end calibration pipeline, and accurate time delays in both the sensing and actuation functions. We do not attempt to correct high-frequency effects in the actuation path, because the actuation function only contributes to the calibration at low frequencies. The process for creating discrete representations, FIR filters from established continuous representation models of the sensing and actuation is the following:

1. *High-pass filter* Seismic noise in the raw data channels is too high at low frequencies to measure any gravitational-wave signals, and the digital system has finite dynamic range. Low frequencies are therefore rolled off by multiplying the frequency components below 9 Hz by half of a Hann window raised to the fourth power.

Inverse Sensing FIR Filter Comparison

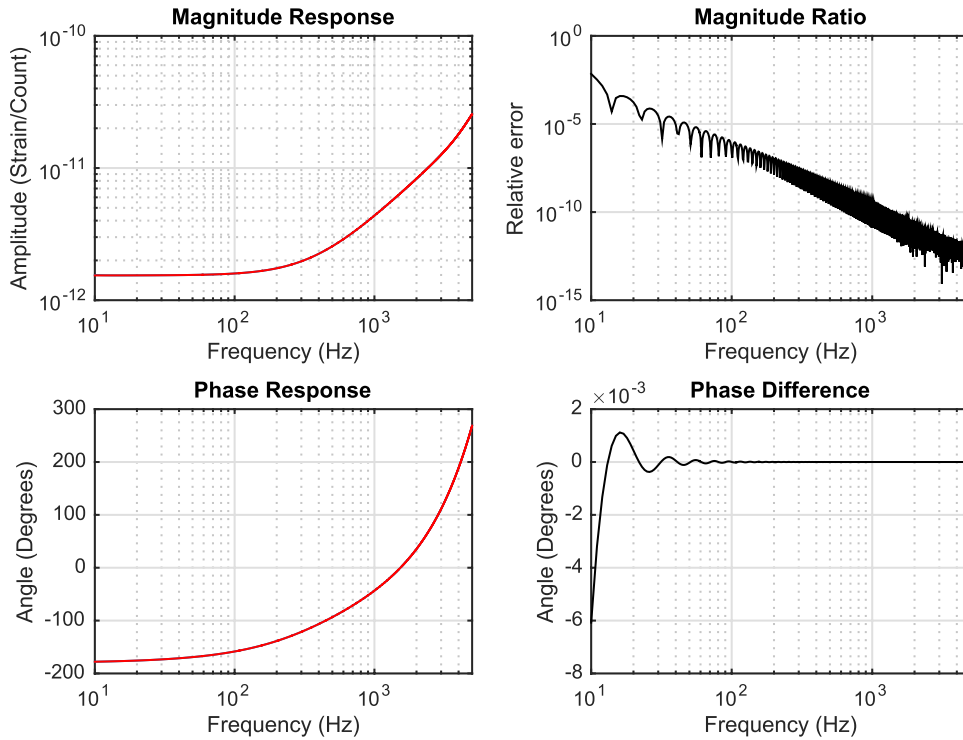


Figure 9. Comparison of the inverse sensing FIR filter to the inverse sensing model. The left panel shows a bode plot of the frequency response of an example inverse sensing FIR filter for L1 (red) and the original inverse sensing model (blue). The top right plot shows the relative error in magnitude between the original model and the frequency response of the FIR filter. The bottom right plot shows the phase difference in degrees between the original model and the frequency response of the FIR filter. As can be seen in the plots, the FIR filter is a very accurate representation of the derived model for the inverse sensing function above 10 Hz.

2. *Low-pass filter in the sensing path* Since the inverse sensing function tends toward infinity at high frequencies, a low-pass filter is applied to the inverse sensing function to roll off high frequencies ($f \gtrsim 6$ kHz) smoothly, in addition to the the high-frequency roll-off in the front end. This is done by multiplying the high frequency components by half of a Hann window.
3. *Artificial delay* An artificial delay is added to the FIR filter that is set to be half of the length of the filter. This delay is undone within the `gstlal` calibration pipeline by advancing the filter output by an equivalent number of samples. The reason for the delay is to center the FIR filter in time, avoiding edge effects while filtering and making the filter non-causal, with output depending on both past and future inputs.
4. *Inverse Fourier transform* The Nyquist component is zeroed out and then the inverse Fourier transform is computed to obtain the FIR time-domain filter.
5. *Tukey window* A Tukey window function is applied to the resulting time-domain FIR filter to ensure it falls off smoothly at the beginning and end of the time-domain filter response. All of the above steps are performed using standard `Matlab` software packages [40].

The fidelity of the FIR filters is checked by taking the frequency response of the final FIR filter with the artificial delay removed and comparing the resulting magnitude and phase to the original frequency-domain model. Example comparison plots are shown in figure 9. The differences between the frequency response of the FIR filter and the original frequency domain model are less than 0.1% in magnitude and 0.01° in phase from 10 Hz to 5 kHz. Additionally, the fidelity of these FIR filters can be seen by examining figure 7, which shows the combined effect of the FIR filters through the response function.

3.2.2. Low-latency `gstlal` calibration pipeline overview. Rather than reading in the final calibration product of the front-end, ΔL_{free}^0 , the `gstlal` calibration pipeline reads in the components of ΔL_{free}^0 separately. The contribution from the actuation is split into three terms that correspond to displacements from actuators at the three lowest stages of the suspension system: ΔL_{T}^0 , ΔL_{p}^0 , and ΔL_{U}^0 . The contribution from the inverse sensing function is ΔL_{res}^0 . Figure 10 shows a simplified diagram of the pipeline. The following procedure describes the workflow of the low-latency `gstlal` calibration pipeline chronologically:

Fill in missing data. Sometimes, chunks of data in one or more channels being read in by the `gstlal` calibration pipeline are missing or corrupted. These dropouts generally occur with a frequency of a few per week and can range from one second in length to several hours in some cases. Oftentimes, they occur in all the input channels at once, but this is not always the case. This missing or corrupted data is filled in at the beginning of the pipeline and time-stamped appropriately to produce a continuous stream. Most channels are filled in with zeros, except for the coherence uncertainty channels (discussed in appendix B), which are filled in with ones. This replacement also occurs when an input sample's magnitude is outside of the expected range of $[10^{-35}, 10^{35}]$, to prevent arithmetic underflows and overflows.

Add ΔL_{p}^0 and ΔL_{U}^0 . Since the same FIR filter and time-dependent correction are applied to both, ΔL_{p}^0 and ΔL_{U}^0 can be added before filtering, to save computational cost. The result is referred to as ΔL_{pU}^0 . ΔL_{T}^0 is not combined because it receives a separate time-dependent correction factor $\kappa_{\text{T}}(t)$.

Downsample the actuation channels. The three actuation channels are read into the `gstlal` calibration pipeline at a sample rate of 4096 Hz. To keep computational costs manageable, this is downsampled to 2048 Hz before applying the FIR correction filters. This is necessary for only the actuation channels, due to the longer filter length (discussed more below). The downsampling is done using the stock `GStreamer` element `audioresample` [41]. A sinc table is used to filter the input before downsampling, in order to minimize aliasing effects. The sinc table is an `acausal` filter, centered in time so as to avoid affecting the phase. It therefore adds a very small latency (≈ 0.05 s) to the pipeline.

Apply FIR filters. The FIR correction filters are applied to correct the front-end estimates of the components of DARM:

$$\Delta L_{\text{T}}^1(t) = [A_{\text{corr}} * \Delta L_{\text{T}}^0](t), \quad (13a)$$

$$\Delta L_{\text{pU}}^1(t) = [A_{\text{corr}} * \Delta L_{\text{pU}}^0](t), \quad (13b)$$

$$\Delta L_{\text{res}}^1(t) = [C_{\text{corr}}^{-1} * \Delta L_{\text{res}}^0](t), \quad (13c)$$




Figure 10. Simplified diagram of the low-latency `gstlal` calibration pipeline. The partially-calibrated outputs of the front-end calibration pipeline are used as inputs and filtered by the appropriate FIR filters for the actuation and inverse sensing paths. The filtered output is corrected by the TDCFs and added together. The result is divided by L to give the final results for $h(t)$. Details of the calculation of the TDCFs and the state vector calculation are not shown.

where C_{corr}^{-1} is the digital filter applied to correct the front-end estimate of the residual displacement, and A_{corr} corrects the front-end estimate of the controlled length differential. This correction is the same for all three stages of the actuation because the components that cannot be modeled in the front end, such as time delays and anti-imaging filters, are the same for each stage.

Since the inverse sensing path contains most of the information above 1 kHz (see figure 4), it is necessary to filter it at the full $h(t)$ sample rate of 16384 Hz. The inverse sensing correction filter C_{corr}^{-1} is 1 s in length and accounts for roughly half of the computational cost of running the pipeline. The actuation filters are 6 s in length, mainly because the actuation channels output from the front end have a lot of noise below 10 Hz, which requires a longer filter to attenuate sufficiently. Due to the length of these filters, it is necessary to filter at 2048 Hz instead of the original actuation sample rate of 4096 Hz provided by the front end. Since the actuation path contributes a small amount to the total strain above 1 kHz (see figure 4), this adds a small ($\lesssim 2\%$) systematic error to $h(t)$ from 1 kHz to 1.5 kHz.

The FIR filtering is performed in the time domain through direct convolution. The timestamps of the filtered result are then advanced by the appropriate amount to compensate for the artificial delay that was built into the FIR filters. To test timestamp accuracy, the calibration lines injected with the photon calibrator can be recovered from $h(t)$ and compared with the injected sinusoids to ensure the phase is consistent. This was done with three of the calibration lines at L1, and the measured timestamp differences were negligibly small compared to the sampling periods of the filters.

The artificial delay built into the FIR filters is half the length of the FIR filter. Therefore, filtering of the actuation path adds a latency of 3 s to the output of the pipeline, due to the 6 s

length of the actuation filters. This is a significant fraction of the 5–9 s latency of the low-latency `gstlal` calibration pipeline. Much of the remaining latency is due to the 4 s length of output $h(t)$ data files, which also accounts for the stated range in latency.

Upsample the actuation channels. After filtering at 2048 Hz, $\Delta L_T^1(t)$ and $\Delta L_{PU}^1(t)$ are upsampled to the full $h(t)$ sample rate of 16384 Hz. A sinc table is again used to filter the input, to attenuate frequencies close to the Nyquist frequency.

Apply time-dependent correction factors. The TDCFs are applied by multiplying the relevant quantities in a sample-by-sample manner. Correcting the filtered output with the TDCFs as a simple multiplication in the time domain is possible because their variation in time is slow compared to the lowest frequencies in the calibrated frequency band¹³.

Add inverse sensing and actuation paths. With correction FIR filters and time-dependent corrections applied, the residual displacement is added with the sum of the two controlled displacements to produce the measured free disturbance, ΔL_{free}^1 . This is divided by the average interferometer arm length L to produce the strain $h(t)$.

3.2.3. Time-dependent correction factors. For the TDCF computation, the `gstlal` calibration pipeline does not ingest any partially computed output from the front-end calibration pipeline but rather performs the computation from the same raw inputs as used by the front-end calibration pipeline. Demodulation in the `gstlal` calibration pipeline is done by multiplying each sample by $e^{-i\omega t}$, where t is the GPS time associated with that sample and ω is the angular frequency of a calibration line. We then downsample to 16 Hz to keep computational cost manageable and apply a low-pass filter H . The result is a complex value representing the amplitude and phase of the calibration line of interest in the signal being measured. For an oscillation with amplitude a and phase angle ϕ ,

$$\begin{aligned} & H * \{e^{-i\omega t} [a \cos(\omega t - \phi) + n(t)]\} \\ &= H * \left\{ \frac{a}{2} \left[e^{-i\phi} + e^{i(-2\omega t + \phi)} \right] + n(t)e^{-i\omega t} \right\} \approx \frac{a}{2} e^{-i\phi}, \end{aligned} \quad (14)$$

where $n(t)$ is noise, that is, anything other than the sinusoidal injection.

Downsampling a time series containing a ~ 300 Hz oscillation to 16 Hz would cause aliasing. Therefore, an anti-aliasing filter is included in the resampling process¹⁴. This effectively forms part of the low-pass filtering process, attenuating the signal to 1% of its original amplitude at the Nyquist frequency. The low-pass filter applied after downsampling is a 20 s Hann window applied in the time domain at 16 Hz. After this, the TDCFs are computed as described in appendix A.

Although the calculation of the TDCFs is running constantly, it produces an accurate result only when the interferometer is in a low-noise state, and even then, the resulting time series is quite noisy. This is largely due to imperfect coherence of the calibration lines in the error signal. We describe the procedure used to manage this issue in appendix B.

¹³ Computational cost could be reduced if the TDCFs were applied and the actuation channels summed before filtering, since the control correction filter would only be applied once. However, due to the excess noise below 10 Hz in the front-end output, multiplying by the TDCFs before high-pass filtering adds significant noise in the relevant frequency band due to noisy fluctuations in the TDCFs.

¹⁴ Due to the need to process complex data, the element used here is different from the one used for resampling the actuation path.

κ_T , κ_{PU} , and κ_C are used to correct $h(t)$, while the time-dependence of the coupled cavity pole f_{cc} and the optical spring frequency f_s and quality Q of the SRC are not corrected for, due to the difficulty of applying frequency-dependent corrections in [realtime](#). Computed values of f_{cc} can stray as far as 20 Hz from the reference model values and f_s varies from about 6 Hz to 9 Hz at H1.

3.2.4. State vector. In addition to computing $h(t)$, the `gstlal` calibration pipeline also computes a bitwise state vector that denotes the integrity of $h(t)$ at a sample rate of 16 Hz. The top level bits of the state vector are the summary bits used to determine whether or not $h(t)$ should be used for astrophysical analysis. For more details about the state vector definition during O2, see [appendix C](#).

3.2.5. Relative accuracy of the low-latency `gstlal` calibration pipeline. The low-latency `gstlal` calibration pipeline delivers a relatively accurate $h(t)$ data product with a latency of 5–9s. As we dive deeper into the era of gravitational-wave astronomy, work continues to improve both the accuracy and the latency of this pipeline.

The static systematic errors currently present in the low-latency `gstlal` calibration pipeline include errors introduced by downsampling the actuation path and any systematic bias from the FIR filters. [Figure 7](#) shows that the sum total of these systematic errors are typically less than 5% in magnitude and a few degrees in phase across the relevant frequency band (20 Hz, 5kHz), with narrow frequency ranges containing increased systematic error. Investigations are still underway to understand and improve the systematic error in these narrow frequency regions. At times, the systematic errors could be much higher than discussed here due to, for example, human errors in the low-latency system.

In addition to the above static systematic errors, the low-latency `gstlal` calibration pipeline does not correct for the time-changing f_{cc} , f_s , and Q factors. This systematic error is not assessed in this paper and will be addressed by a future publication.

3.3. High-latency `gstlal` calibration pipeline

The `gstlal` calibration pipeline is also used to produce a second calibration of the archived raw data in high latency. The high-latency calibration is generally produced weeks to months after acquisition of raw data products. The `gstlal` calibration pipeline run in high latency is quite similar to its low-latency counterpart and is depicted in [figure 11](#).

Instead of reading in the partially-calibrated outputs of the front-end, we calibrate in high latency using d_{err} and d_{ctrl} directly¹⁵. Therefore, the FIR filters applied by the high-latency pipeline contain the full static reference models for the inverse sensing and the actuation. Otherwise, the production of the filters is equivalent to the the description given in [section 3.2.1](#).

In the high-latency calibration, the model-based constants used in computing the TDCFs are recorded in a file which is ingested into the `gstlal` calibration pipeline. This is in contrast to these constants being read in directly from the front-end system. This allows us to correct large systematic errors in the TDCFs caused by mistakes made in low-latency in recording these factors in the front-end system.

The production of high-latency calibration is accomplished by running multiple jobs in parallel on the LIGO Data Grid computing clusters [\[42\]](#). Each job produces 4096s of calibrated

¹⁵ While it would certainly be possible to approach the high-latency calibration with the same methodology as the low-latency calibration, which is to correct the output of the front-end calibration pipeline, it is often impractical to develop the appropriate correction filters that model all the changes occurring in the front-end calibration pipeline during an observing run due to changes caused by human error that inevitably occur in realtime systems.

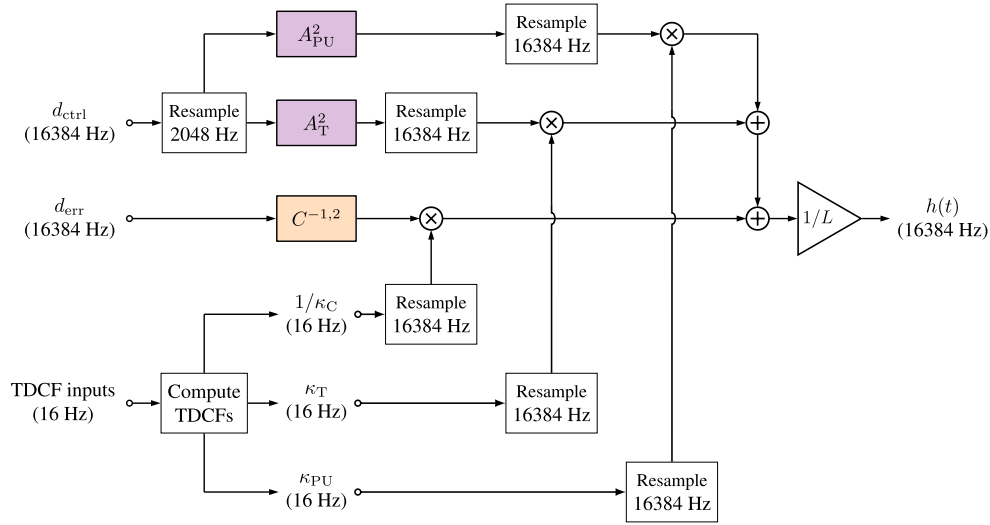


Figure 11. Diagram of the high-latency `gstlal` calibration pipeline. The digital error d_{err} and digital control d_{ctrl} signals are calibrated directly using the full inverse sensing and actuation models. Just as with the low-latency `gstlal` calibration pipeline, the error and control paths are corrected with the appropriate TDCF. A superscript 2 indicates association with the high-latency `gstlal` calibration pipeline. Details of the TDCF calculation and the state vector calculation are not shown.

data, using the corresponding stretch of d_{err} and d_{ctrl} with the addition of several minutes of input data at the beginning and end, to allow all filtering processes to settle.

In general, differences between the low-latency and high-latency versions of calibration can vary widely, depending on whether significant systematic errors exist in low-latency that can be corrected in high-latency. For the majority of O2, however, differences are under 2% in magnitude from 10 Hz to 5 kHz, as is seen in figure 6. The systematic errors that exist in the high-latency `gstlal` calibration are the same as those in the low-latency `gstlal` calibration pipeline, with the exception of any human errors that were present in the low-latency process. Some model improvements are also made between the low- and high-latency calibration iterations, which is the cause of the rough 2% and 1–3° improvement seen in figure 7.

4. Current and future development

During Advanced LIGO’s first two observing runs, the low-latency strain $h(t)$ was produced in the time domain using a combination of IIR filtering processes in the front-end computers and FIR filtering processes in the downstream `gstlal` calibration pipeline. The `gstlal` calibration pipeline also applies time-dependent correction factors derived from stable, calibrated reference displacements induced by photon calibrators to the actuation and sensing components of $h(t)$ to improve calibration accuracy in the relevant frequency band. The `gstlal` calibration pipeline is used in high latency to produce an additional, more accurate calibration that corrects for missing data and some known systematic errors in the low-latency $h(t)$ data. During O2, the systematic error and uncertainty achieved by the high-latency calibration (as well as much of the low-latency `gstlal` calibration) is less than 5% in magnitude and 3° in phase in the frequency band of 20 Hz to 5 kHz. For details of the frequency-dependent uncertainty budget for O2, see [15].

In the future, we hope to improve the latency and accuracy of the low-latency calibration data products. Work is in progress to implement a method for applying time- and frequency-dependent corrections, i.e. for the time-dependence of the cavity pole f_{cc} and the optical spring frequency f_s and quality factor Q of the SRC. As described in section 2.3 and appendix A, these are already computed in both the front-end and the `gstlal` calibration pipeline. In order to compensate for their time-dependence and improve $h(t)$ accuracy, an algorithm to compute a new inverse sensing filter and smoothly replace the previous filter is being developed. The time-dependence of the cavity pole is known to be a significant contributor to systematic error in $h(t)$ across the relevant frequency band (see, e.g. figure 5). The time-dependence of f_s and Q is significant at H1 in the lowest frequency band of concern, mainly below about 20 Hz. An additional change that should lead to a small improvement in the calibration is to compute κ_P and κ_U individually instead of tracking and applying their approximate combined effect in the `gstlal` calibration pipeline as κ_{PU} . This would better represent the true model of the actuation as expressed in equation (12), and it could have an impact at low frequencies as ΔL_P and ΔL_U differ in their frequency dependence (see figure 4).

The online calibration latency has improved during O2 by about 5 s. This was mainly due to the removal of the latency associated with applying an anti-aliasing filter during the resampling required for the calculation and application of the TDCFs. Further improvements can be made, primarily in two ways: (1) The length of the gravitational-wave frame files could be reduced from four seconds, which is the current standard, to one second; (2) The length of the actuation filter in the `gstlal` calibration pipeline can be reduced. Shortening the actuation filters is straightforward, but if not done with care, spectral leakage will corrupt the $h(t)$ spectrum. For this reason, we plan to incorporate more effective high-pass filtering into the FIR filters, as the current methodology involves simply rolling off low frequencies with half of a Hann window raised to the fourth power (first step in section 3.2.1).

Work is also underway to develop a complete front-end calibration model, as the eventual goal of the LIGO calibration team is to perform the low-latency calibration entirely in the front-end computers. There are several advantages to having the low-latency calibration entirely in the front end. The front end's direct access to interferometer models and parameters may reduce the occurrence of human errors in calibration. Calibrating in the front end also removes data transfer steps, reducing the likelihood of data dropouts.

The primary challenges that must be overcome to realize a complete front-end calibration are the ability to implement FIR filtering and to assign timestamps to data that are not real time. Due to the length of the FIR filters used in calibration as well as the overhead associated with each 16 kHz front-end computational cycle, the front-end code needs to be modified to distribute the calculations for a single FIR filter over multiple computer cores and spread a single lower rate FIR calculation over multiple 16 kHz cycles. A fundamental redesign of the interferometer's data acquisition system is needed to be able to advance the calibrated data's timestamp in real-time, as the system currently writes a single 64 s-long frame file at a time with all the data contained having the same timestamp, as well as broadcasting data as it comes in to low-latency pipelines in 1 s-long frame files.

5. Conclusion

Overall, the calibration process for Advanced LIGO has been expanded to include a very low-latency calibrated strain data product and an unprecedented level of accuracy in the

final, high-latency strain data when compared with Initial LIGO. The latency achieved during O2 run was 5–9 s, representing a 5 s improvement made since the beginning of the run. Typical calibration uncertainty levels in Initial LIGO ranged from 10% in magnitude and 10° in phase to 20% in magnitude and 20° in phase [43], while results for the calibration uncertainty around the time of GW170104 [4] are less than 5% in magnitude and 3° in phase [15]. Future Advanced LIGO observing runs promise further improvement to both the latency and accuracy of the calibration process with an end-goal of having sub-second-latency calibrated strain data with percent-level accuracy available for analysis. Such an achievement will help push the envelope forward for the new era of gravitational-wave astronomy and multi-messenger astronomy.

AQ5

Acknowledgments

The authors would like to thank Chris Wipf, Evan Hall, Jolien Creighton, Patrick Brockill, John Zweizig, Kipp Cannon, Bruce Allen, Chris Pankow, Chad Hanna, and Les Wade for helpful discussions and feedback while developing the methods discussed here. The authors were supported by National Science Foundation grants PHY-1607178, PHY-1607585, and PHY-1506360. LIGO was constructed by the California Institute of Technology and Massachusetts Institute of Technology with funding from the National Science Foundation and operates under cooperative agreement PHY-0757058. This paper carries LIGO Document Number LIGO-P1700236.

Appendix A. Calculation of time-dependent correction factors

The method used to measure and compensate for temporal variations in the interferometer response (see section 2.3) was described in [16]. Details of the method are reproduced here for completeness and to show how the additional factors for the detuning of the SRC are included.

Measurement of the factor κ_T is done through comparison of a photon calibrator line at frequency f_1^{pc} to another calibration line at frequency f_T , injected through x_T into only the test mass suspension stage. κ_{PU} is measured by comparison of the line at f_1^{pc} to another line at f_{ctrl} , which is added to the control signal and distributed to each of the lowest three stages in the quadruple pendulum (see figure 2). All three of these lines are placed between 10 and 40 Hz, where the stages of the actuation function are comparable in magnitude, to reduce systematic errors. The quantities κ_C and f_{cc} are measured using the f_2^{pc} line near 330 Hz, in the middle of Advanced LIGO's most sensitive frequency band. The quantities f_s and Q are measured by the lowest photon calibrator line f_4^{pc} , placed near 8 Hz, where their effect is significant. The remaining line at f_3^{pc} is used as a check on calibration uncertainty above 1 kHz. Calibration lines injected using the electrostatic drive or the photon calibrator can be seen in the $h(t)$ spectrum (see figure 4). The DARM line, x_{ctrl} , does not appear in $h(t)$ due to the location where it is injected, before d_{ctrl} is read out. Table A1 summarizes the purpose of each calibration line.

To solve for the TDCFs, we need to know their effect on the amplitude and phase of the calibration lines in d_{err} . From figure 2,

$$\tilde{d}_{\text{err}} = \kappa_C \frac{\tilde{C}_{\text{res}}}{1 + i f / f_{\text{cc}}} \left(\tilde{\Delta L}_{\text{free}} + \tilde{x}_{\text{pc}} - \tilde{\Delta L}_{\text{ctrl}} \right), \quad (\text{A.1})$$

where \tilde{C}_{res} is the static reference model sensing function with dependence on the cavity pole removed. Note that we have neglected the detuning of the SRC, as this has little impact at the

Table A1. Summary of the purpose of each calibration line.

Line	Purpose	Frequency
f_{ctrl}	Computation of κ_{PU}	10–40 Hz
f_{T}	Computation of κ_{T}	10–40 Hz
f_1^{pc}	Computation of κ_{T} and κ_{PU}	10–40 Hz
f_2^{pc}	Computation of κ_{C} and f_{cc}	~ 330 Hz
f_3^{pc}	Check on high-frequency calibration	~ 1 kHz
f_4^{pc}	Computation of f_{s} and Q	~ 8 Hz

calibration line frequencies used to compute κ_{T} , κ_{PU} , κ_{C} , and f_{cc} . Referring to figure 2, we see that

$$\widetilde{\Delta L}_{\text{ctrl}} = \kappa_{\text{T}} \widetilde{A}_{\text{T}} \left(\widetilde{D} \widetilde{d}_{\text{err}} + \widetilde{x}_{\text{ctrl}} - \widetilde{x}_{\text{T}} \right) + \kappa_{\text{PU}} \widetilde{A}_{\text{PU}} \left(\widetilde{D} \widetilde{d}_{\text{err}} + \widetilde{x}_{\text{ctrl}} \right). \quad (\text{A.2})$$

Solving for $\widetilde{d}_{\text{err}}$ in terms of the injected excitations, we obtain the expression

$$\widetilde{d}_{\text{err}} = \frac{\widetilde{\Delta L}_{\text{free}} + \widetilde{x}_{\text{pc}} - \kappa_{\text{T}} \widetilde{A}_{\text{T}} (\widetilde{x}_{\text{ctrl}} - \widetilde{x}_{\text{T}}) - \kappa_{\text{PU}} \widetilde{A}_{\text{PU}} \widetilde{x}_{\text{ctrl}}}{(S \widetilde{C}_{\text{res}})^{-1} + (\kappa_{\text{T}} \widetilde{A}_{\text{T}} + \kappa_{\text{PU}} \widetilde{A}_{\text{PU}}) \widetilde{D}}, \quad (\text{A.3})$$

where we have defined

$$S \equiv \frac{\kappa_{\text{C}}}{1 + \text{i}f/f_{\text{cc}}}. \quad (\text{A.4})$$

To solve for κ_{T} , we demodulate d_{err} at the electrostatic drive calibration line f_{T} and the first photon calibrator line f_1^{pc} , which are separated by only ~ 1 Hz. This yields two simple equations:

$$\widetilde{d}_{\text{err}}(f_{\text{T}}) = \frac{\kappa_{\text{T}} \widetilde{A}_{\text{T}} \widetilde{x}_{\text{T}}}{(S \widetilde{C}_{\text{res}})^{-1} + (\kappa_{\text{T}} \widetilde{A}_{\text{T}} + \kappa_{\text{PU}} \widetilde{A}_{\text{PU}}) \widetilde{D}} \Big|_{f_{\text{T}}} \quad (\text{A.5})$$

and

$$\widetilde{d}_{\text{err}}(f_1^{\text{pc}}) = \frac{\widetilde{x}_{\text{pc}}}{(S \widetilde{C}_{\text{res}})^{-1} + (\kappa_{\text{T}} \widetilde{A}_{\text{T}} + \kappa_{\text{PU}} \widetilde{A}_{\text{PU}}) \widetilde{D}} \Big|_{f_1^{\text{pc}}}. \quad (\text{A.6})$$

Then, we take the following ratios of $\widetilde{d}_{\text{err}}$, $\widetilde{x}_{\text{pc}}$, and \widetilde{x}_{T} at the two line frequencies to estimate κ_{T} :

$$\kappa_{\text{T}} \approx \frac{1}{\widetilde{A}_{\text{T}}(f_{\text{T}})} \cdot \frac{\widetilde{d}_{\text{err}}(f_{\text{T}})}{\widetilde{x}_{\text{T}}(f_{\text{T}})} \left(\frac{\widetilde{d}_{\text{err}}(f_1^{\text{pc}})}{\widetilde{x}_{\text{pc}}(f_1^{\text{pc}})} \right)^{-1} \cdot \frac{\widetilde{R}^{\text{static}}(f_{\text{T}})}{\widetilde{R}^{\text{static}}(f_1^{\text{pc}})}. \quad (\text{A.7})$$

In equation (A.7), we have treated the ratio of the denominators of equations (A.5) and (A.6) as being constant in time, an approximation that depends on the line frequencies f_1^{pc} and f_{T} being close together. $\widetilde{d}_{\text{err}}(f_{\text{T}})$, $\widetilde{x}_{\text{T}}(f_{\text{T}})$, $\widetilde{d}_{\text{err}}(f_1^{\text{pc}})$, and $\widetilde{x}_{\text{pc}}(f_1^{\text{pc}})$ are constantly measured by the calibration pipelines, while the complex constant $\widetilde{R}^{\text{static}}(f_{\text{T}}) / (\widetilde{A}_{\text{T}}(f_{\text{T}}) \widetilde{R}^{\text{static}}(f_1^{\text{pc}}))$ depends only on static functions at the line frequencies obtained from measurements, and is therefore read into the calibration pipelines, along with several other constants used in computing the TDCFs.

A similar method is used to compute κ_{PU} with the DARM line f_{ctrl} and the same photon calibrator line f_1^{pc} , yielding the result

$$\kappa_{\text{PU}} \approx -\frac{1}{\tilde{A}_{\text{PU}}(f_{\text{ctrl}})} \left[\frac{\tilde{d}_{\text{err}}(f_{\text{ctrl}})}{\tilde{x}_{\text{ctrl}}(f_{\text{ctrl}})} \left(\frac{\tilde{d}_{\text{err}}(f_1^{\text{pc}})}{\tilde{x}_{\text{pc}}(f_1^{\text{pc}})} \right)^{-1} \cdot \frac{\tilde{R}^{\text{static}}(f_{\text{ctrl}})}{\tilde{R}^{\text{static}}(f_1^{\text{pc}})} + \kappa_{\text{T}} \tilde{A}_{\text{T}}(f_{\text{ctrl}}) \right]. \quad (\text{A.8})$$

As seen in equation (A.8), κ_{PU} depends on κ_{T} , which is computed first. The optical gain correction factor κ_{C} and the cavity pole f_{cc} depend on both κ_{T} and κ_{PU} under this formalism. To measure them, we make use of a higher calibration line injected with the photon calibrator, $f_2^{\text{pc}} \sim 330$ Hz, where the first term in the denominator of equation A.3 dominates. Solving for $S(f_2^{\text{pc}})$, we have

$$S(f_2^{\text{pc}}) = \frac{1}{\tilde{C}_{\text{res}}} \left(\frac{\tilde{x}_{\text{pc}}}{\tilde{d}_{\text{err}}} - \tilde{D} [\kappa_{\text{T}} \tilde{A}_{\text{T}} + \kappa_{\text{PU}} \tilde{A}_{\text{PU}}] \right)^{-1} \Big|_{f_2^{\text{pc}}}. \quad (\text{A.9})$$

Then the expressions for κ_{C} and f_{cc} are simply

$$\kappa_{\text{C}} = \frac{|S(f_2^{\text{pc}})|^2}{\Re[S(f_2^{\text{pc}})]}, \quad (\text{A.10})$$

$$f_{\text{cc}} = -\frac{\Re[S(f_2^{\text{pc}})]}{\Im[S(f_2^{\text{pc}})]} f_2^{\text{pc}}. \quad (\text{A.11})$$

In order to compute the optical spring frequency f_{s} and quality factor Q of the SRC, we need to consider the full model of the sensing function described in equation (10). In a similar fashion to the derivation above, we express d_{err} in terms of the injected excitations and the TDCFs, this time including the full sensing function model:

$$\begin{aligned} \tilde{d}_{\text{err}} &= \frac{f^2}{f^2 + f_{\text{s}}^2 - \text{i}ff_{\text{s}}/Q} \frac{\kappa_{\text{C}} \tilde{C}_{\text{res}}}{1 + \text{i}f/f_{\text{cc}}} \left(\tilde{\Delta}L_{\text{free}} + \tilde{x}_{\text{pc}} - \tilde{\Delta}L_{\text{ctrl}} \right) \\ &= \frac{\tilde{\Delta}L_{\text{free}} + \tilde{x}_{\text{pc}} - \kappa_{\text{T}} \tilde{A}_{\text{T}} (\tilde{x}_{\text{ctrl}} - \tilde{x}_{\text{T}}) - \kappa_{\text{PU}} \tilde{A}_{\text{PU}} \tilde{x}_{\text{ctrl}}}{(1 + \xi)(S \tilde{C}_{\text{res}})^{-1} + (\kappa_{\text{T}} \tilde{A}_{\text{T}} + \kappa_{\text{PU}} \tilde{A}_{\text{PU}}) \tilde{D}}, \end{aligned} \quad (\text{A.12})$$

where we have defined

$$\xi \equiv \frac{f_{\text{s}}^2 - \text{i}ff_{\text{s}}/Q}{f^2} \quad (\text{A.13})$$

and again used equation (A.2) to substitute for $\tilde{\Delta}L_{\text{ctrl}}$. Demodulating d_{err} at the lowest photon calibrator line ($f_4^{\text{pc}} = 7.93$ Hz at LHO) yields the expression

$$\tilde{d}_{\text{err}}(f_4^{\text{pc}}) = \frac{\tilde{x}_{\text{pc}}}{(1 + \xi)(S \tilde{C}_{\text{res}})^{-1} + (\kappa_{\text{T}} \tilde{A}_{\text{T}} + \kappa_{\text{PU}} \tilde{A}_{\text{PU}}) \tilde{D}} \Big|_{f_4^{\text{pc}}}. \quad (\text{A.14})$$

Therefore, $\xi(f_4^{\text{pc}})$ can be computed in the `gstlal` calibration pipeline using

$$\xi(f_4^{\text{pc}}) = -1 + \left[\mathcal{S} \tilde{C}_{\text{res}} \left(\frac{\tilde{x}_{\text{pc}}}{\tilde{d}_{\text{err}}} - (\kappa_{\text{T}} \tilde{A}_{\text{T}} + \kappa_{\text{PU}} \tilde{A}_{\text{PU}}) \tilde{D} \right) \right]_{f_4^{\text{pc}}}, \quad (\text{A.15})$$

and f_s and Q take the simple form

$$f_s = f_4^{\text{pc}} \sqrt{\Re[\xi(f_4^{\text{pc}})]}, \quad (\text{A.16a})$$

$$Q = -\frac{\sqrt{\Re[\xi(f_4^{\text{pc}})]}}{\Im[\xi(f_4^{\text{pc}})]}. \quad (\text{A.16b})$$

Note that the calculation of f_s and Q depends on the other TDCFs, which are computed first. Here, we have assumed that the effect of SRC detuning is very small at the higher frequencies used to compute κ_{T} , κ_{PU} , κ_{C} , and f_{cc} . This is valid as long as f_s is small compared to those frequencies. We have not neglected the effect of the previously computed TDCFs on f_s and Q , as the optical gain and actuation strength clearly have significant contributions here.

Appendix B. Smoothing of time-dependent correction factors

As noted in section 3.2.3, the computed TDCFs cannot be applied to $h(t)$ directly due to both excessive noise and the fact that the measurements of the calibration lines are inaccurate when detector noise increases. Here, we describe the method used to attenuate noise in the TDCFs and to discriminate between times of acceptable and unacceptable measurements.

The first step is to accept or reject computed TDCFs based on the coherence between the injection channels x_i and error signal d_{err} at the calibration line frequencies. To compute the coherence between two signals $x(t)$ and $y(t)$ at a frequency f , the signals are demodulated at the chosen frequency using a local oscillator. Then, the coherence is computed as

$$\gamma_{xy}^2(f) = \frac{|\langle \tilde{x}^*(f) \tilde{y}(f) \rangle|^2}{\langle |\tilde{x}(f)|^2 \rangle \langle |\tilde{y}(f)|^2 \rangle}, \quad (\text{B.1})$$

where the angled brackets denote averages, computed using a low-pass filter, and the superscript asterisk denotes complex conjugation. Coherences are computed using 10s chunks of input data, and $n_{\text{d}} = 13$ independent, consecutive values are then averaged, so that each averaged coherence is based on 130s of input data. The average coherence is used to compute the normalized random error in the magnitude of the transfer function \hat{H}_{xy} at each calibration line frequency:

$$\epsilon \left[\left| \hat{H}_{xy} \right| \right] \approx \sqrt{\frac{1 - \gamma_{xy}^2}{2n_{\text{d}}\gamma_{xy}^2}} \quad (\text{B.2})$$

as derived in [44]. This uncertainty is computed in the front end and read into the `gstlal` calibration pipeline. If the estimated uncertainty is below the chosen threshold (currently 0.004 at LHO and 0.02 at LLO), a computed TDCF is accepted and passed downstream; otherwise, it is flagged as a gap. Each TDCF is accepted or rejected based on the estimated uncertainty at the frequency of every calibration line used in its computation. Very roughly speaking, the coherence of the calibration lines is measured to be acceptable about half of the time during an observing run, although this varies widely with time.

Table B1. A summary of the meaning of each bit in the calibration state vector.

Bit	Short descriptor	Long descriptor
0	HOFT_OK	$h(t)$ was successfully computed (logical AND of bits 2–4, 11, 13, 17, and 25)
1	OBS_INTENT	Interferometer is in ‘observation intent’ mode
2	OBS_READY	Interferometer is in ‘observation ready’ mode
3	HOFT_PROD	$h(t)$ was produced
4	FILTERS_OK	Calibration filters settled in
5	NO_STOCH_HW_INJ	No stochastic hardware injections present
6	NO_CBC_HW_INJ	No compact binary coalescence hardware injections present
7	NO_BURST_HW_INJ	No burst hardware injections present
8	NO_DETCHAR_HW_INJ	No detector characterization hardware injections present
9	NO_GAP	The input data was present (not a gap)
10	KAPPA_SMOOTHING_OK	TDCFs smoothing algorithm is settled in
11	KAPPA_TST_SMOOTH_OK	Smoothed κ_T output is in expected range
12	KAPPA_TST_MEDIAN_OK	Median array used for κ_T smoothing not dominated by bad coherence time
13	KAPPA_PU_SMOOTH_OK	Smoothed κ_{PU} output is in expected range
14	KAPPA_PU_MEDIAN_OK	Median array used for κ_{PU} smoothing not dominated by bad coherence time
15	Not in use	Not in use
16	Not in use	Not in use
17	KAPPA_C_SMOOTH_OK	Smoothed κ_C output is in expected range
18	KAPPA_C_MEDIAN_OK	Median array used for κ_C smoothing not dominated by bad coherence time
19	F_CC_SMOOTH_OK	Smoothed f_{cc} output is in expected range
20	F_CC_MEDIAN_OK	Median array used for f_{cc} smoothing not dominated by bad coherence time
21	SUS_COH_OK	Coherence uncertainty for SUS calibration line is acceptable
22	DARM_COH_OK	Coherence for DARM calibration line is acceptable
23	PCALY_LINE1_COH_OK	Coherence for first photon calibrator line is acceptable
24	PCALY_LINE2_COH_OK	Coherence for second photon calibrator line is acceptable
25	NO_UNDERFLOW_INPUT	Magnitude of input for all channels is between 1×10^{-35} and 1×10^{35}
26	F_S_SMOOTH_OK	Smoothed f_s output is in expected range
27	F_S_MEDIAN_OK	Median array used for f_s smoothing is not dominated by bad coherence time
28	Q_SMOOTH_OK	Smoothed Q of SRC output is in expected range
29	Q_MEDIAN_OK	Median array used for Q smoothing is not dominated by bad coherence time

The accepted values are then passed to a 128s running median, where each new value replaces the oldest value currently in the median array. If computed values are being rejected due to unacceptable coherence, the algorithm will recognize the gap flag associated with those values and instead replace the oldest value in the array with the previously computed median. Thus, when detector noise increases, the reported TDCFs stabilize at the most recent median value computed during the low-noise state. The choice to use a running median instead of, e.g. a running average, is intended to make the result insensitive to occasional outliers that occur in

the computed TDCFs, resulting in a smoother time series. The drawback is that computing a median over 128 s of data can be computationally expensive. To offset this extra cost, we have taken advantage of the fact that, in a running median, the previous median is known. This can be used to reduce the number of operations from $\mathcal{O}(\ell^2)$ to $\mathcal{O}(\ell)$, where ℓ is the length of the median array.

A short, 10 s running average is used after the running median in order to remove ‘kinks’ left after the running median. This is computationally cheap and reduces the high-frequency content of the smooth TDCF time series.

All of the gating and smoothing operations are carried out at 16 Hz, before the TDCFs are upsampled to 16384 Hz and used to correct $h(t)$.

Appendix C. Calibration state vector

The `gstlal` calibration pipeline computes a bitwise state vector that summarizes the integrity of the calibration at a sample rate of 16 Hz. The definition of each bit in the calibration state vector during O2 is shown in table B1. A value of zero for a bit means the bit is False and a value of 1 means the bit is True.


The zeroth bit of the state vector indicates that $h(t)$ is considered valid and usable for astrophysical analyses. Bit 1 in the state vector is determined based on state information provided by the operator and the front-end system. If bit 1 is set to True, then the operator has determined that no commissioning activities are ongoing and science-quality data should be available. Similarly, bit 2 is determined based on state information provided by the front-end system and indicates that the interferometer has reached a nominal low-noise configuration. Bit 3 indicates whether or not strain was computed by the pipeline, but this bit does not necessarily indicate the validity of the strain data at that time.

Bit 4 indicates whether the various FIR filters have settled in since the interferometer reached a nominal low-noise state as determined by bit 2. This bit will be turned off for N seconds after the observation-ready state is reached and will also be turned off for N seconds before the observation-ready state is lost, where N seconds is the total filter settle time for all FIR filters in the pipeline. This is due to the non-causal nature of the FIR filters, as discussed in section 3.2.1.

Bits 5–8 indicate whether any known hardware injections are present in the data. A hardware injection is achieved by actuating the end test masses of the interferometers to mimic an astrophysical signal. Hardware injections are performed for both transient and continuous gravitational-wave emission [45]. Bit 9 indicates whether there are times that no input data was received by the calibration pipeline. Bits 10–24 and 26–29 are all associated with the TDCFs calculation and its smoothing process (see appendix A). In particular, bits 11, 13, and 17 indicate whether or not the computed TDCFs that are used to correct $h(t)$ (κ_C , κ_{PU} , and κ_T) are within some expected range of values. If the TDCFs stray outside of the expected range, this is a red flag that something about the current reference calibration model may now be drastically inaccurate. Bit 25 indicates whether the input could have caused arithmetic underflows or overflows.

Data analysts are advised to use the logical AND of bits 0 and 1 to determine whether to analyze data for astrophysical signals. Bit 0 indicates the overall integrity of the $h(t)$ calculation while bit 1 indicates that the operator has determined the interferometer is in trustworthy, science-quality data-taking state.

AQ6

AQ7 **ORCID iDs**A D Viets  <https://orcid.org/0000-0002-4241-1428>Duncan A Brown  <https://orcid.org/0000-0002-9180-5765>AQ8 **References**

- AQ9 [1] Abbott B P *et al* (LIGO Scientific Collaboration and Virgo Collaboration) 2016 *Phys. Rev. Lett.* **116** 061102
- [2] Abbott B P *et al* (LIGO Scientific Collaboration and Virgo Collaboration) 2016 *Phys. Rev. Lett.* **116** 241103
- [3] Abbott B P *et al* (LIGO Scientific Collaboration and Virgo Collaboration) 2016 *Phys. Rev. X* **6** 041015
- [4] Abbott B P *et al* (LIGO Scientific Collaboration and Virgo Collaboration) 2017 *Phys. Rev. Lett.* **118** 221101
- [5] Abbott B P *et al* (LIGO Scientific Collaboration and Virgo Collaboration) 2017 *Phys. Rev. Lett.* **119** 141101
- [6] Abbott B P *et al* (LIGO Scientific Collaboration and Virgo Collaboration) 2017 *Phys. Rev. Lett.* **119** 161101
- [7] Abbott B P *et al* (LIGO Scientific Collaboration) 2017 *Phys. Rev. D* **95** 062003
- [8] Izumi K and Sigg D 2016 *Class. Quantum Grav.* **34** 015001
- [9] Abbott B P *et al* (LIGO Scientific Collaboration and Virgo Collaboration) 2016 *Phys. Rev. Lett.* **116** 131103
- [10] Aasi J *et al* 2015 *Class. Quantum Grav.* **32** 074001
- [11] Buonanno A and Chen Y 2002 *Phys. Rev. D* **65** 042001
- [12] Siemens X *et al* 2004 *Class. Quantum Grav.* **21** S1723
- [13] Karki S *et al* 2016 *Rev. Sci. Instrum.* **87** 114503
- AQ10 [14] Characterization of the Photon Calibrator data from LIGO's sixth Science Run <https://dcc.ligo.org/LIGO-P1100013/public>
- [15] Cahillane C *et al* 2017 *Phys. Rev. D* **96** 102001
- [16] Tuyenbayev D *et al* 2016 *Class. Quantum Grav.* **34** 015002
- [17] Martynov D V *et al* 2016 *Phys. Rev. D* **112** 2004
- [18] Robertson N A *et al* 2002 *Class. Quantum Grav.* **19** 4043
- [19] Aston S M *et al* 2012 *Class. Quantum Grav.* **29** 235004
- [20] Carbone L *et al* 2012 *Class. Quantum Grav.* **29** 115005
- [21] Matichard F *et al* 2015 *Precis. Eng.* **40** 273
- [22] Abbott B P *et al* 2016 *Astrophys. J. Lett.* **833** L1
- [23] Abbott B P *et al* (LIGO Scientific Collaboration and Virgo Collaboration) 2017 *Phys. Rev. Lett.* **118** 121101
- [24] Abbott B P *et al* (LIGO Scientific Collaboration and Virgo Collaboration) 2017 *Phys. Rev. Lett.* **118** 121102
- [25] Abbott B P *et al* 2017 *Astrophys. J.* **839** 12
- [26] Abbott B P *et al* 2017 in prep
- [27] Yunes N, Yagi K and Pretorius F 2016 *Phys. Rev. D* **94** 084002
- [28] Chamberlain K and Yunes N 2017 [arXiv:1704.08268](https://arxiv.org/abs/1704.08268)
- AQ11 [29] Abbott B P *et al* (LIGO Scientific Collaboration and Virgo Collaboration) 2016 *Phys. Rev. Lett.* **116** 221101
- [30] Chernoff D F and Finn L S 1993 *Astrophys. J.* **411** L5
- [31] Abbott B P *et al* (LIGO Scientific Collaboration and Virgo Collaboration) 2016 *Phys. Rev. Lett.* **116** 241102
- [32] Singer L P *et al* 2014 *Astrophys. J.* **795** 105
- [33] GStreamer <https://gstreamer.freedesktop.org>
- [34] MATLAB c2d <https://mathworks.com/help/control/ref/c2d.html>
- [35] MATLAB minreal <https://mathworks.com/help/control/ref/minreal.html>
- [36] Laakso T and Valimaki V 1996 *IEEE Signal Process.* **13** 30
- [37] LAL <https://wiki.ligo.org/DASWG/LALSuite>

- [38] gstlal <https://wiki.ligo.org/DASWG/GstLAL>
- [39] gstlal-calibration <https://versions.ligo.org/git/gstlal/tree/gstlal-calibration?h=gstlal-calibration-1.1.7-v1>
- [40] MATLAB <https://mathworks.com/help/matlab/>
- [41] GStreamer audioresampler <https://gstreamer.freedesktop.org/data/doc/gstreamer/head/gst-plugins-base-plugins/html/gst-plugins-base-plugins-audioresample.html>
- [42] LIGO Data Grid <https://wiki.ligo.org/LDG/>
- [43] Abadie J *et al* (VIRGO, LIGO Scientific) 2012 arXiv:1203.2674 [gr-qc]
- [44] Bendat J S and Piersol A G 2011 *Random Data: Analysis and Measurement Procedures* vol 729 (New York: Wiley)
- [45] Biwer C *et al* 2017 *Phys. Rev. D* **95** 062002