

Validating gravitational-wave detections: The Advanced LIGO hardware injection system

C. Biwer,¹ D. Barker,² J. C. Batch,² J. Betzwieser,³ R. P. Fisher,¹ E. Goetz,² S. Kandhasamy,³ S. Karki,⁴ J. S. Kissel,² A. P. Lundgren,⁵ D. M. Macleod,³ A. Mullavey,³ K. Riles,⁶ J. G. Rollins,⁷ K. A. Thorne,³ E. Thrane,^{8,27} T. D. Abbott,⁹ B. Allen,^{5,10,11} D. A. Brown,¹ P. Charlton,¹² S. G. Crowder,¹³ P. Fritschel,¹⁴ J. B. Kanner,⁷ M. Landry,² C. Lazzaro,^{15,16} M. Millhouse,¹⁷ M. Pitkin,¹⁸ R. L. Savage,² P. Shawhan,¹⁹ D. H. Shoemaker,¹⁴ J. R. Smith,²⁰ L. Sun,²¹ J. Veitch,²² S. Vitale,¹⁴ A. J. Weinstein,⁷ N. Cornish,¹⁷ R. C. Essick,¹⁴ M. Fays,²³ E. Katsavounidis,¹⁴ J. Lange,²⁴ T. B. Littenberg,²⁵ R. Lynch,¹⁴ P. M. Meyers,²⁶ F. Pannarale,²³ R. Prix,⁵ R. O’Shaughnessy,²⁴ and D. Sigg²

¹*Syracuse University, Syracuse, New York 13244, USA*

²*LIGO Hanford Observatory, Richland, Washington 99354, USA*

³*LIGO Livingston Observatory, Livingston, Louisiana 70754, USA*

⁴*University of Oregon, Eugene, Oregon 97403, USA*

⁵*Albert-Einstein-Institut, Max-Planck-Institut für Gravitationsphysik, D-30167 Hannover, Germany*

⁶*University of Michigan, Ann Arbor, Michigan 48109, USA*

⁷*LIGO, California Institute of Technology, Pasadena, California 91125, USA*

⁸*Monash Centre for Astrophysics, School of Physics and Astronomy,
Monash University, Victoria 3800, Australia*

⁹*Louisiana State University, Baton Rouge, Louisiana 70803, USA*

¹⁰*University of Wisconsin-Milwaukee, Milwaukee, Wisconsin 53201, USA*

¹¹*Leibniz Universität Hannover, D-30167 Hannover, Germany*

¹²*Charles Sturt University, Wagga Wagga, New South Wales 2678, Australia*

¹³*Bellevue College, Bellevue, Washington 98008, USA*

¹⁴*LIGO, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, USA*

¹⁵*Center for Relativistic Astrophysics and School of Physics, Georgia Institute of Technology,
Atlanta, Georgia 30332, USA*

¹⁶*INFN, Sezione di Padova, I-35131 Padova, Italy*

¹⁷*Montana State University, Bozeman, Montana 59717, USA*

¹⁸*SUPA, University of Glasgow, Glasgow G12 8QQ, United Kingdom*

¹⁹*University of Maryland, College Park, Maryland 20742, USA*

²⁰*California State University Fullerton, Fullerton, California 92831, USA*

²¹*The University of Melbourne, Parkville, Victoria 3010, Australia*

²²*University of Birmingham, Birmingham B15 2TT, United Kingdom*

²³*Cardiff University, Cardiff CF24 3AA, United Kingdom*

²⁴*Rochester Institute of Technology, Rochester, New York 14623, USA*

²⁵*University of Alabama in Huntsville, Huntsville, Alabama 35899, USA*

²⁶*University of Minnesota, Minneapolis, Minnesota 55455, USA*

²⁷*OzGrav: The ARC Centre of Excellence for Gravitational-wave Astronomy,
Hawthorn, Victoria 3122, Australia*

(Received 23 December 2016; published 27 March 2017)

Hardware injections are simulated gravitational-wave signals added to the Laser Interferometer Gravitational-wave Observatory (LIGO). The detectors’ test masses are physically displaced by an actuator in order to simulate the effects of a gravitational wave. The simulated signal initiates a control-system response which mimics that of a true gravitational wave. This provides an end-to-end test of LIGO’s ability to observe gravitational waves. The gravitational-wave analyses used to detect and characterize signals are exercised with hardware injections. By looking for discrepancies between the injected and recovered signals, we are able to characterize the performance of analyses and the coupling of instrumental subsystems to the detectors’ output channels. This paper describes the hardware injection system and the recovery of injected signals representing binary black hole mergers, a stochastic gravitational wave background, spinning neutron stars, and sine-Gaussians.

DOI: [10.1103/PhysRevD.95.062002](https://doi.org/10.1103/PhysRevD.95.062002)

I. INTRODUCTION

The Advanced Laser Interferometer Gravitational-wave Observatory (Advanced LIGO) is a network of two interferometric gravitational-wave detectors located in Hanford, WA, and Livingston, LA [1]. The Advanced LIGO detectors are part of a global network of current and planned detectors including Virgo [2], GEO600 [3], KAGRA [4], and LIGO India [5]. The first direct observations of gravitational waves, both from binary black hole mergers, were made in Advanced LIGO's first observing run [6,7].

In order to make confident statements about gravitational-wave events, Advanced LIGO employs studies to understand both transient and persistent noise artifacts [8], and the calibration of the detectors [9,10]. In addition, analysis pipelines for detection and parameter estimation of gravitational-wave signals employ a wide range of different techniques to mitigate the impact of non-Gaussian and nonstationary noise in the detectors' data. Testing these analyses and characterizing the detectors involves carrying out "hardware injections" in which we simulate the detectors' response to a gravitational-wave signal. Hardware injections are used to validate "software injections," where simulated signals are added to the data as part of an analysis pipeline without any physical actuation occurring; software injections are used for high-statistics evaluation of the performance of analyses.

Hardware injections have several other uses. Following a detection candidate, we study similar simulated gravitational-wave signals through the use of repeated injections. These hardware injections provide an end-to-end check for the search and parameter estimation analyses to recover signals in the detectors' data. The recovery of hardware injections provides an additional check of the sign of the calibration between the Advanced LIGO detectors using astrophysical waveforms and the recovery measures the time delay of the signal in the controls system; the calibration of the detectors is checked by other means as well [9,10]. In addition, we can check for instrumental and environmental channels that respond to changes in differential arm length variations from gravitational-wave signals.

Another use for hardware injection in Initial LIGO were "blind injections" which were hardware injections known only to a small team [11,12]. Blind injections simulate the detection and characterization of a real astrophysical signal. No blind injections were carried out during Advanced LIGO's first observing run. There are no plans to perform blind injections in future observing runs.

To create a hardware injection we physically displace the detectors' test masses. The mirrors in the arms act as "freely falling" test masses [13]. Advanced LIGO measures the differential displacement along the two arms $\Delta L = L_x - L_y$, and the output channel to analyses is gravitational-wave strain $h = \Delta L/L$ where $L = (L_x + L_y)/2$ [13]. Differential displacement of the test masses mimics the detectors' response to a gravitational-wave signal.

The detectors' response to a true gravitational-wave is not exactly the same as the detectors' response to physically displacing the test masses [14,15]. However, the difference is well understood, and it is only relevant at high frequencies [14,15]. In addition, the actuators apply a force to the test masses in their suspensions whereas a true gravitational-wave does not.

Advanced LIGO uses different actuators to perform hardware injections than Initial LIGO. In Initial LIGO, the test masses were displaced using magnets mounted on the optic itself; however, these actuators are no longer used to move the test masses due to displacement noise [16,17]. In Advanced LIGO's first observing run, hardware injections were realized with two different actuation methods: electrostatic drive systems [18] and photon radiation pressure actuators referred to as "photon calibrators" [19]. Starting in December 2015 the photon calibrators have been the only actuator used to perform hardware injections since their actuation range available for hardware injections is larger.

During Advanced LIGO's first observing run a wide variety of waveforms were injected. Advanced LIGO is sensitive to astrophysical sources of gravitational waves including binary black hole and/or neutron star mergers [20,21], the stochastic gravitational-wave background [22], and spinning neutron stars [23]. Hardware injections for each of these astrophysical sources were performed. In addition, detector characterization studies injected series of sine-Gaussians across the Advanced LIGO frequency range.

This paper describes how we inject signals into the Advanced LIGO detectors with the photon calibrators in Sec. II. Section III describes the results from analyses that used hardware injections in Advanced LIGO's first observing run. This includes the recovery of binary black hole merger signals in Secs. III A and III B, the stochastic gravitational-wave hardware injection in Sec. III C, and a population of spinning neutron stars in Sec. III D. Although binary neutron star hardware injections were performed we do not discuss their recoveries in this paper. A description of the detector characterization analysis to check the response of instrumental and environmental changes to differential displacement of the test masses is described in Sec. III E. Finally, Sec. IV summarizes the hardware injections from Advanced LIGO's first observing run.

II. HARDWARE INJECTION PROCEDURE

Each different type of astrophysical source has different signal characteristics and properties, and hence different technical requirements for the hardware injection system. In particular, the difference in the time duration of the sources.

The time in Advanced LIGO's frequency range for compact-object binary mergers can last a fraction of a second to minutes depending on the component masses. The signal enters Advanced LIGO's frequency range from low frequency, and as the two component masses inspiral closer together they sweep upward in frequency [6]. The

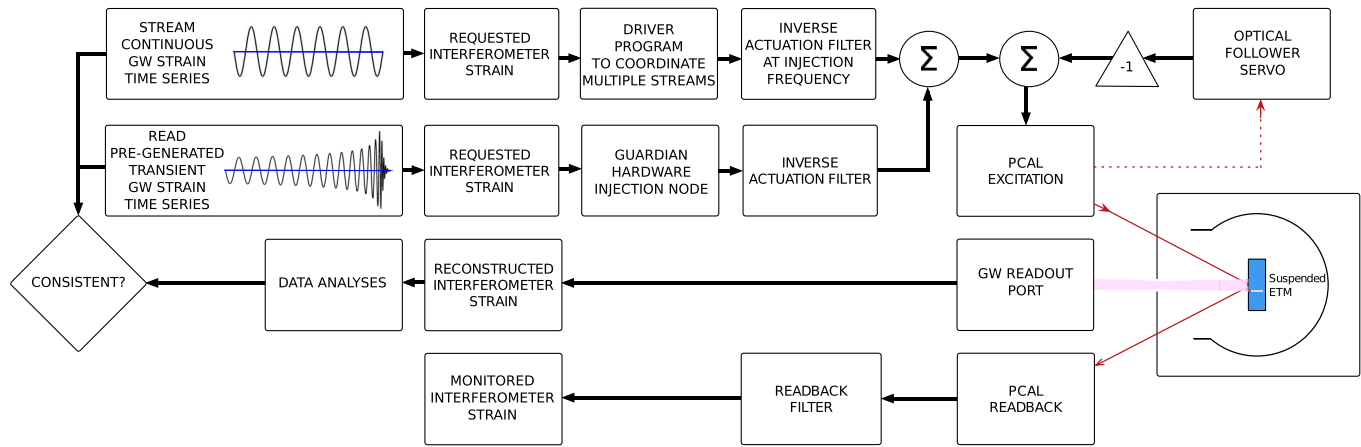


FIG. 1. Block overview of the Advanced LIGO hardware injection system. Time series for transient and continuous-wave injections are generated and sent to the photon calibrator (PCAL). The signal modulates the laser power of the photon calibrator to displace the end test mass (ETM) in a way that mimics a gravitational wave (GW) passing through the detector. The optical follower servo has its own pick-off of the light that is sent towards the ETM indicated by a dashed line. The gravitational-wave strain of the detector is analyzed and checked for consistency by the analysis’ developers. A photodetector that receives the light reflected from the test mass is used to monitor and verify the injected signal.

merger’s termination frequency and waveform length is determined by the masses of the two objects. For example, GW150914 terminates at 250 Hz after about 0.2 s above 35 Hz [6], whereas the inspiral-only portion of a binary neutron star waveform with both component masses equal to $1.4 M_{\odot}$ terminates at 1527 Hz after about 36 s above 35 Hz.

Gravitational-waves from a spinning neutron star will be present in the data for the full duration of an observing run. Spinning neutron stars emit continuous gravitational waves at an almost constant frequency which is Doppler modulated by Earth’s motion [23]. The gravitational-wave frequency slowly evolves as the pulsar spins down [23].

The stochastic gravitational-wave background will persist in the data throughout the observing run. The stochastic background is the superposition of many events that combine to create a low-level broadband nondeterministic signal [22].

Nonastrophysically motivated injections for detector characterization studies use a succession of short duration (<1 s) sine-Gaussians across Advanced LIGO’s frequency range.

We categorize hardware injections into two classes: “transient injections” that are localized in time, and “continuous-wave injections” that are active throughout the duration of the observing run. Examples of transient injections include simulated binary black hole and/or neutron star mergers, sine-Gaussians, and stochastic background signals. These signals have a finite duration. The simulated stochastic background is included as a transient injection since we increase the amplitude of the waveform in order to limit it to a short segment of data. Continuous-wave injections simulate a synthetic population of rapidly spinning neutron stars (which we designate in shorthand as pulsars, although such a source need not emit electromagnetic pulsations detectable at the Earth).

Separate automation processes control transient and continuous-wave injections. Figure 1 shows a schematic of the two pathways that generate and transmit gravitational-wave strain time series to the photon calibrator. In this section we work through Fig. 1, beginning at the top left and working clockwise, in order to describe the processes that control the transient and continuous-wave injections.

We generate the simulated gravitational waveforms for transient injection signals prior to injection. The system for managing the automated processes of the Advanced LIGO detector subsystems is Guardian [24]. Guardian manages the transient hardware injections, it reads the next scheduled injection’s time series and transmits the data to the digital control system of the detector at the scheduled time.

Continuous-wave injections are generated in real time. A streaming time series of simulated gravitational waves from a synthetic population of spinning neutron stars, described by astrophysical parameters, including the strain amplitude, sky location, and initial frequency, is transmitted to the digital controls system of the detector. A driver program called `psinject` (“pulsar injection”) coordinates the simultaneous generation and buffering of multiple streams of signals representing pulsars [25].

The transient and continuous-wave signals in the digital controls system of the detector are sent to an actuator that displaces the test masses to simulate the detector’s response to a gravitational wave signal. In Advanced LIGO’s first observing run, we used the electrostatic drive systems [18] and photon calibrators [19] as actuators for hardware injections. Each actuator has its own actuation pathway in the controls system; however, in Fig. 1 we show only the photon calibrators’ pathway.

Hardware injections are carried out by actuating one of the end test masses (ETM) of the interferometer and thus inducing differential interferometer strain variations that simulate the response to an incident gravitational wave. We only need to apply a force on one ETM to induce differential strain variations in the interferometer. The common arm length degree of freedom of the interferometer, $(L_x + L_y)/2$, is controlled by its own servo. If an actuator lengthens the x -arm by applying a force on the ETM, then the common arm length servo will promptly shorten the y -arm length to suppress the change in the common arm length degree of freedom. This creates differential interferometer strain variations that are partially suppressed by the differential arm length feedback servo.

The differential arm length degree of freedom of the interferometer is controlled by a feedback servo loop that actuates the longitudinal position of one of the ETMs [10]. The differential arm length feedback control loop suppresses apparent ETM displacements resulting from noise sources, signal injections, and gravitational waves. Because this servo suppresses the injected waveform, reconstructing the unsuppressed injected strain requires correcting for the action of this servo. The correction for the response of the common and differential arm length servos used in reconstructing the gravitational-wave strain was described in Ref. [10].

The actuators for the servo that controls the differential arm length degree of freedom are electrostatic drive systems. These actuators apply forces via fringing field gradients from electrodes patterned onto a reaction mass separated by a few millimeters from the back surface of the ETM [18].

The electrostatic drive systems were used at the beginning of Advanced LIGO's first observing run for injecting simulated signals. They successfully injected the waveforms for the GW150914 and stochastic background hardware injection analyses. However, the actuation range available for hardware injections is restricted because they are part of the differential arm length servo which consumes a significant fraction of its total actuation range in maintaining stable servo operation.

In order to inject a larger parameter space of waveforms, for example binary black hole and/or neutron star mergers at closer distances, we transitioned to photon calibrators for hardware injections. Since December 2015, we use a photon calibration system to displace the ETM in a way that simulates the effect of a gravitational wave signal. This is depicted on the right of Fig. 1.

A photon calibrator system uses an auxiliary, power-modulated laser with two beams impinging on the ETM located at the end of the x -arm of the interferometer. The photon calibrator on the other arm, the y -arm, is used for calibrating the detector output [19]. The two beams are diametrically opposed on the surface of the ETM, adjusted to have equal powers, and positioned to minimize unintended torques and deformations of the surface which could cause errors in the expected displacement.

The Advanced LIGO photon calibrators employ a feedback control system referred to as the “optical follower servo” [19,26]. This servo, with a bandwidth of ~ 100 kHz, facilitates simulated signal injection via ETM actuation. This ensures that the laser output power modulation closely follows the analog voltage waveform injected at the servo input.

Digital infinite impulse response (IIR) compensation filters, called the “inverse actuation filters,” convert the requested interferometer strain signal (a digital signal) into an estimate of the photon calibrator optical follower servo input signal (an analog signal) required to achieve the desired length actuation. There is an analogous set of filters for the electrostatic drive system; however, we focus on the photon calibrators here. These filters are designed to compensate for several factors. There is compensation for (i) the force-to-length transfer function of the suspended ETM, (ii) the signal conditioning electronics that includes a digital anti-imaging filter, the digital-to-analog converter gain, and an analog anti-imaging filter, and (iii) the optical follower servo transfer function. Phase delays of the anti-imaging filters and physical time delays of the digital control system cannot be compensated by the inverse actuation filters because the digital IIR filters allowed by the Advanced LIGO control system must be causal. These delays, on the order of $240 \mu\text{s}$, are taken into account during injection recovery.

The digital signals from the transient and continuous-wave injection pathways are passed through the inverse actuation filters, summed, and sent to the photon calibrator; see Fig. 1. Sporadic, unintended interruptions occurred in the Hanford injection system during Advanced LIGO's first observing run, in which the buffering failed to keep up with real-time injection. The cause was not tracked down because the interruptions occurred at apparently random times, but the dropouts may be related to periods of high traffic on the controls system computer network. The sudden termination introduces a step function to the inverse actuation filters that has a large response at high frequencies. The effect of these dropouts, should they recur, will be mitigated by the use of point-by-point, Fourier-domain inverse actuation functions, using a separate, constant coefficient for each of the injected spinning neutron stars, all of which are extremely narrowband. This is shown in the continuous-wave injection pathway in Fig. 1. Transient injections were not affected. Guardian sets the gain after the inverse actuation filters to zero while there is no active transient injection so unintended signals do not propagate into the detector data.

The strain actually injected into the interferometer is determined using the photon calibrator read-back signal generated by a power sensor that monitors the laser light reflected from the ETM, as shown at the bottom of Fig. 1. The output of this sensor is converted to injected interferometer strain using the read-back filter that compensates for the force-to-length transfer function as well as digital and analog filters in the signal read-back pathway. In the case of




FIG. 2. The maximum displacement of an ETM using the photon calibrator (blue). For the sinusoidal force induced by sinusoidally power modulated laser beams, $F = mA$ implies that the induced displacement is given by $x = -F/(m\omega)^2$. The dashed blue curve indicates that the fidelity of the induced displacements degrades above 1 kHz due to the need to roll off the inverse actuation filters to maintain stability near the Nyquist frequency. The maximum displacement of the ETM required for two optimally oriented compact binary waveforms that contain an inspiral, merger, and ringdown are shown for reference. A 30–30 M_\odot binary at 100 Mpc (yellow) and 1.4–1.4 M_\odot binary at 100 Mpc (red) were generated using the SEOBNRv2 approximant [28]. Note that the required displacement for the 1.4–1.4 M_\odot binary exceeds the maximal photon calibrator displacement at high frequencies.

hardware injections, however, the excitation channel is calibrated by taking a transfer function measurement between the excitation channel and the read-back photodetectors. This transfer function is then incorporated within the inverse actuation filters. This provides a calibration accuracy on the order of a few percent, sufficient for the hardware injection analysis. For better calibration, however, we can compare the recovered signal and the injected signal as measured by the read-back photodetector.

There are some limitations to the photon calibrator system. First, the photon calibrator has a limited actuation strength. Figure 2 shows the maximal displacement of the ETM using the photon calibrator system. The photon calibrator can provide up to ~ 1 W of peak power, but the force-to-length response of the ETM transfer function scales as the inverse-square of frequency [27]. Thus, the photon calibrator is limited in the amount of induced ETM displacement, especially at higher frequencies. Second, signal fidelity above 1 kHz is limited due to the shape of the anti-imaging filters and the desire to roll off the compensation filters close to the Nyquist frequency such that the compensation filters remain stable. Nonetheless, the photon calibrator is able to provide precise, calibrated displacements of the ETM in response to many astrophysical waveforms.

III. RESULTS

In this section we describe results from hardware injection analyses in Advanced LIGO’s first observing run. Hardware injections were used as (i) an end-to-end test of searches and parameter estimation analyses, (ii) an additional check of the calibration, and (iii) a method to check for instrumental and environmental channels that are coupled to the detectors’ output channels. Binary black hole waveforms with parameters similar to GW150914 and GW151225 were used to test modeled and unmodeled analyses described in Secs. III A and III B. A simulated stochastic gravitational-wave signal was recovered with the analysis described in Sec. III C. The population of simulated spinning neutron stars analysis in Sec. III D also provided an additional check of the calibration; they were used to verify the overall sign of the detectors’ calibration and to measure the time delay between the hardware injection excitation channel and the detectors’ output channels. It is possible to use other hardware injections to perform these checks of the calibration as well. These astrophysical signals were injected coherently into the Hanford (H1) and Livingston (L1) detectors. Section III E describes the study that injects a series of sine-Gaussians across the Advanced LIGO frequency range to check for instrumental and environmental channels that respond to differential arm length variations.

A. Compact binary coalescence gravitational-wave hardware injections

Advanced LIGO observed two binary black hole mergers (GW150914 and GW151226) and a third detection candidate (LVT151012) during its first observing run [6,7]. After each detection was made, hardware injections were used to simulate gravitational-wave sources with similar parameters to each event in the detector. Verifying that these hardware injections were recovered by the search and parameter estimation analyses was part of the validation of each detection. Compact binary coalescence searches use matched filtering to correlate Advanced LIGO data with a bank of gravitational-wave templates [29]. Here we consider hardware injections analyzed by the PyCBC search for gravitational waves [30,31] described in Refs. [21,32]. Parameter estimation analyses were used to analyze the hardware injections and check for consistency with GW150914 and GW151226. We ran the same code used to characterize the detected events [33,34]. We show the recovery of hardware injections with parameters taken from posterior distributions of parameter estimation results for GW150914 [34] and GW151226 [7,21].

For GW150914 and GW151226, we injected ten waveforms coherently into the two detectors after collecting enough data to confidently establish a detection. The GW150914 hardware injections were generated with the SEOBNRv2 waveform approximant and included systems with component spins aligned with the angular momentum




FIG. 3. A comparison of the signal-to-noise ratio ρ from software injections and the recovered signal-to-noise ratio of the hardware injection. Parameters for the hardware injections were drawn from the posterior distributions for GW150914 (circles) and GW151226 (triangles). The software injection ρ is the mean and 1σ error from the recovery of 50 software injections filtered with the injected waveform near the time of the injection. The threshold on ρ is indicated by the gray region. The arrows indicate the coherent injection affected by a nearby noise transient.

of the binary [28]. The GW150914 waveforms had a total mass from $[68 M_{\odot}, 79 M_{\odot}]$ in the source frame, mass ratios from 1 to 1.8, and distances from $[250 \text{ Mpc}, 530 \text{ Mpc}]$. Mass ratio is defined as m_1/m_2 where $m_1 > m_2$. These signals were injected October 2 to October 6, 2015. The GW151226 hardware injections were generated with the precessing waveform approximant IMRPhenomPv2 [35,36] and injected on January 11, 2016. The GW151226 waveforms had a total mass from $[25 M_{\odot}, 30 M_{\odot}]$ in the source frame, mass ratios from 1 to 4.3, and distances from $[240 \text{ Mpc}, 580 \text{ Mpc}]$. For both the GW150914 and GW151226 waveforms the sky positions were selected to be on the same triangulation ring as the corresponding astrophysical event.

Figure 3 shows the reported PyCBC matched-filter signal-to-noise ratio ρ versus the expected ρ calculated using software injections. Figure 3 includes 19 of the 20 hardware injections performed for GW150914 and GW151226. The expected ρ was calculated using software injections, in which signals are added to the data without any physical actuation. Detector data within hours of the hardware injections was selected for adding software injections since the sensitivity of the detectors does not significantly vary on these time scales [8]. The normalization of ρ implies that the ρ measured for a population of identical signals in different realizations of the detector noise

will be $\int df |\tilde{h}(f)|^2 / S_h(f)$ [29]. The recovered software injections were found to be consistent with the expectation.

All of the hardware injections are coherent but an astrophysical signal can have a different ρ in each detector. Hanford and Livingston have their own angular sensitivity and noise spectra that affects ρ for an event [13]. One of the GW150914 hardware injections reported $\rho < 5.5$ in Livingston. In order to manage computational considerations, the analysis requires a single-detector signal-to-noise ratio of at least 5.5. Thus, this injection was not “detected.” A signal-to-noise ratio < 5.5 for this injection, with an expected signal-to-noise ratio of 6, is consistent with the variation of the matched-filter output in Gaussian noise [29].

In Fig. 3 there is one GW150914 hardware injection that was recovered with a signal-to-noise ratio of 16.1 and 10.9 in Hanford and Livingston respectively; however, the injection had an expected signal-to-noise ratio of 22.1 and 13.4. This injection was recovered with a lower signal-to-noise ratio because a loud transient noise artifact was present in the Livingston data shortly after the hardware injection.

There are a variety of transient noise artifacts that adversely affect the search [8,9]. The search includes a signal consistency test to mitigate their effect. PyCBC reports a detection statistic $\hat{\rho}$ [37] that combines information about the matched-filter signal-to-noise ratio ρ , where consistency is determined with a reduced χ^2 statistic χ_r^2 [38]. The χ_r^2 statistic [38] downweights the significance of the noise transients.

While hardware injections are an important end-to-end test, software injections are useful because a large number can be performed without disturbing the detector or significantly reducing the duty cycle of the detectors. Figure 3 shows the software injections to be consistent with the recovery of signals that propagate through the detectors; therefore we can generate large populations of software injections that are used in other studies to evaluate the search efficiency [30], detections [32], and binary merger rates [20,21].

Figure 4 shows the χ_r^2 statistic [38] versus ρ for hardware injections, a large population of software injections, and noise transients. Astrophysical events are indicated with stars. Hardware injections are indicated with squares. Software injections are denoted by pluses. These software simulations repeat the analysis many times to test the search across a large parameter space. The software injections in Fig. 4 were generated from a population of aligned-spin binaries with source-frame component masses between 2 to 98 M_{\odot} using the SEOBNRv2 waveform approximant [28]. The population of software injections is randomly distributed in sky location, orientation, distance, and time. The injection times are within the 39 day period around GW150914 reported in Ref. [32].

In Fig. 4 a highly significant astrophysical signal should be clearly separated from the background distributions. We see a separation of the software injections with

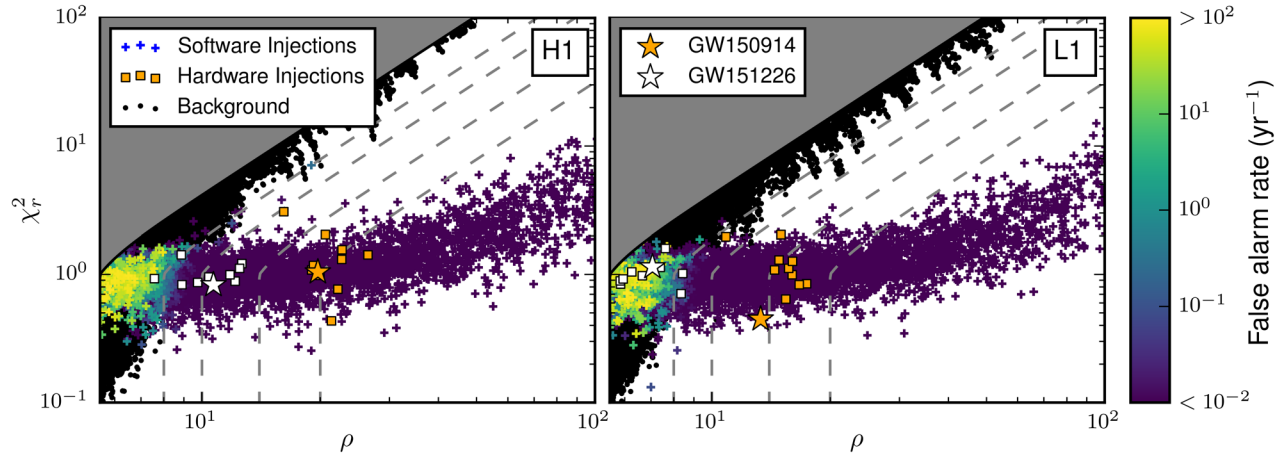


FIG. 4. PyCBC χ_r^2 statistic versus matched-filter signal-to-noise ratio ρ for each detector. Software injections are represented as pluses that are colored by false-alarm rate. The false-alarm rate is calculated using the time-slide algorithm described in Ref. [30]. The gravitational-wave events GW150914 and GW151226 are shown as stars. Hardware injections for GW150914 and GW151226 are represented as boxes. These are coherent software and hardware injections, and therefore the H1 and L1 plots are dependent on each other. Single-detector background distributions (black dots) are plotted; there was a threshold applied indicated by the gray region. Lines of constant detection statistic $\hat{\rho}$ are shown (gray dashed lines); plotted are $\hat{\rho} = \{8, 10, 14, 20\}$.

high significance (false-alarm rate $< 1/100 \text{ yr}^{-1}$) and background distributions. All ten GW150914 hardware injections are recovered with high significance. Although the GW151226 Livingston hardware injections are not visibly distinguishable from the background distribution in Fig. 4, seven hardware injections have a highly significant false-alarm rate ($< 1/100 \text{ yr}^{-1}$) since we combine data from both detectors. Two hardware injections were recovered with $1/10 \text{ yr}^{-1} > \text{false-alarm rate} > 1/100 \text{ yr}^{-1}$, a significance comparable to the gravitational-wave candidate LVT151012 ($1/2 \text{ yr}^{-1}$) reported in Advanced LIGO's first observing run [6]. The software and hardware injections with similar parameters to GW150914 and GW151226 found with high significance validates the search's ability to detect similar systems.

If a detection candidate is a true gravitational wave, we should be able to reproduce the morphology of the posterior distributions using the hardware injections as well as with software injection. Conversely any significant differences have the potential to highlight discrepancies between the observation and our waveform models, or errors in our data analysis. Here we focus on two parameters: chirp mass and sky location.

The chirp mass \mathcal{M} is defined as

$$\mathcal{M} = \frac{(m_1 m_2)^{3/5}}{(m_1 + m_2)^{1/5}}. \quad (1)$$

Here, m_1 and m_2 are the binary's component masses. The chirp mass is typically the best estimated parameter of a compact binary coalescence signal, since it dominates the phase evolution during inspiral. In Fig. 6 we show for all the GW151226 hardware injections the posterior distributions of the chirp mass minus the respective injected values, using the precessing waveform approximant IMRPhenomPv2 [35,36].

Most posteriors have comparable width. Hardware injections with low signal-to-noise ratio have broader distributions and in one case shows bimodality. The width of the 90% credible interval for the detector-frame chirp mass for GW151226 is $\sim 0.12 M_\odot$ [21], which is comparable to that found with the hardware injections. Verifying that the width and shape of the posterior distribution for the chirp mass of the candidate events is similar to those of the hardware injection analyses has been part of validating the parameter estimation results for each detection.

Sky maps from the parameter estimation analysis of GW150914 and GW151226 were shared with electromagnetic observatories [39,40] and were shown in Refs. [21,41]. In Fig. 5, we show a reconstructed Earth-bound coordinate sky map for GW151226 along with sky maps for two hardware injections. One of the two hardware injections (at GPS time 1136588346) has low signal-to-noise ratio and thus spans a larger sky area, although still near to the same triangulation ring. The other injection (at GPS time 136592747) is instead representative of the typical map: all other maps look similar to this and are not shown to avoid overcrowding.

A previous study used the parameter estimation method described above to validate another strategy used to interpret GW150914, by directly comparing data to simulations of Einstein's equations [42]. In that study the parameter estimates for GW150914 derived from IMRPhenomPv2 and numerical relativity agreed [42]. This study was repeated comparing hardware injections to numerical relativity, and found posterior distributions in mass and spin that were consistent with the IMRPhenomPv2 analysis. This comparison provided a timely, independent validation of this new parameter estimation strategy using real data and in this region of parameter space.

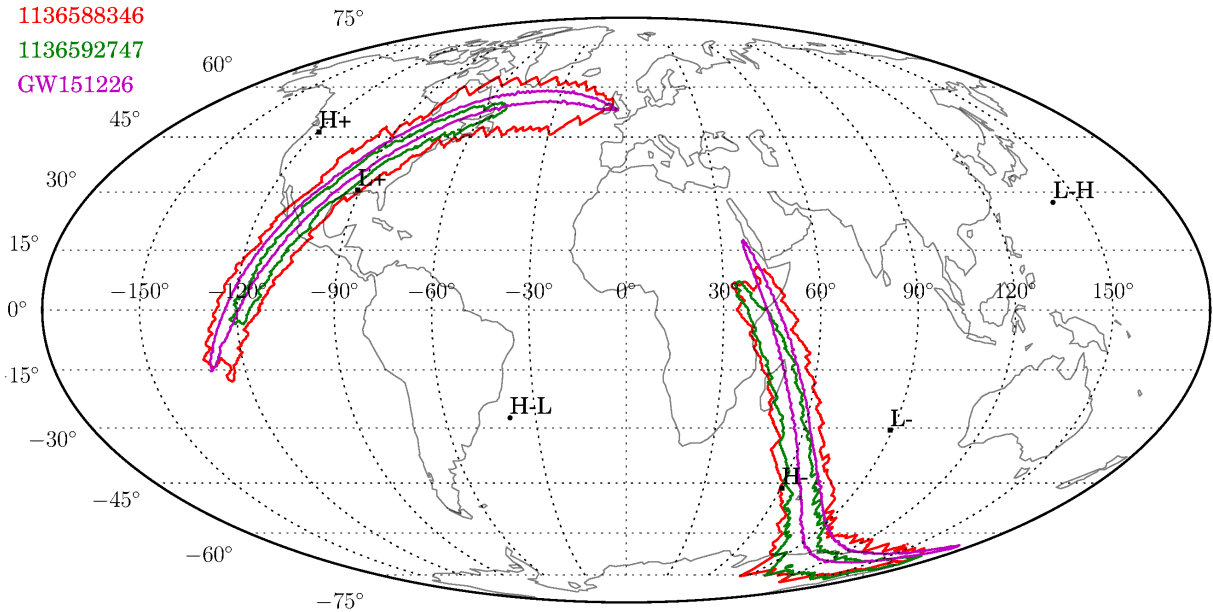


FIG. 5. The 90% confidence interval sky maps for two hardware injections (red and green) and GW151226 (magenta). The sky maps are shown in Earth-bound coordinates. H+ and L+ mark the Hanford and Livingston sites, and H- and L- indicate antipodal points; H-L and L-H mark the poles of the line connecting the two detectors (the points of maximal time delay). The two hardware injections are chosen to be representative of an average event (green) and a subthreshold event (red). We notice how all sky maps have support near the same ring of equal time delay between the two Advanced LIGO detectors.

B. Burst gravitational-wave hardware injections

There are astrophysical sources of gravitational waves that have poorly modeled or unknown waveforms, such as core-collapse supernovae [43]. In order to search for a wide range of unmodeled astrophysical sources, analyses search the Advanced LIGO strain data for short-duration, transient gravitational-wave events referred to as “bursts” [44]. Here, we look at injection recoveries using Coherent WaveBurst [45], BayesWave [46], and LALInferenceBurst [33,47]. These analyses produce reconstructed waveforms with




FIG. 6. Posterior probability density functions (PDF) for the chirp mass inferred from GW151226 hardware injections. The true value has been removed to center all distributions around zero. Hardware injections with very low signal-to-noise ratio show large width and in one case bimodality. The bimodal distribution comes from the injection at GPS time 1136588346 which is also shown in Fig. 5.

minimal assumptions about the waveform morphology. We compare these reconstructions to the injected waveforms of hardware injections.

Binary black hole waveforms were used to test the burst analyses. In addition to the ten GW150914 hardware injections described in Sec. III A, there were 24 waveforms injected with physical parameters similar to GW150914. Eight were nonspinning waveforms with equal source frame component masses and a total mass of $76 M_{\odot}$; 16 were aligned-spin with total masses from $[70 M_{\odot}, 80 M_{\odot}]$ in the source frame and mass ratios from 1 to 5. The waveforms were generated with the SEOBNRv2 approximant [28].

Since burst searches do not use gravitational-wave templates, they are less sensitive to compact merger signals than modeled searches [48]. The burst searches did not detect GW151226, and recovered only a single hardware injection mentioned in Sec. III A to validate that detection. The recovery of GW151226 hardware injections using the burst analyses are not discussed in this paper.

Coherent WaveBurst identifies coherent events in spectrographic data from the Advanced LIGO detectors constructed using a wavelet representation. It then reconstructs the gravitational waveform signal using a constrained likelihood method [45]. For signals consistent with compact binary coalescences, it also estimates the system’s chirp mass based on the time-frequency evolution of the signal [49].

The low-latency Coherent WaveBurst search recovered 28 of the 34 total injections. In Fig. 7 we show the recovered versus injected excess power signal-to-noise ratio and chirp mass. Note that the excess power signal-to-noise ratio [45] is




FIG. 7. The 28 recovered hardware injections by the low-latency Coherent WaveBurst search. The recovered excess power signal-to-noise ratio and chirp mass are consistent with expectations [49]. Left: Recovered excess power signal-to-noise ratio (Reconstructed SNR) versus injected excess power signal-to-noise ratio (Injected SNR). Right: Recovered chirp mass (Reconstructed) versus injected chirp mass (Injected).

distinct from the matched-filter signal-to-noise ratio ρ . Figure 7 matches the expectations for the recovered chirp mass from fitting the time-frequency evolution of the data [45,49]. These estimates help guide the initial response to detections. Subsequently, accurate estimates are obtained with template-based, fully coherent Bayesian parameter estimation [33].

BayesWave uses a sum of sine-Gaussian wavelets to model the gravitational-wave signal [46]. The reconstruction assumes an elliptically polarized gravitational wave, but no other constraints are imposed [46]. BayesWave investigated the 28 GW150914 hardware injections recovered by Coherent WaveBurst. Previous studies with software injections show that the recovered waveforms produced by BayesWave accurately match injected signals [44]. To measure the overlap between injected and recovered waveforms, we use the network match

$$\text{Match} = \frac{(h_{\text{inj}}|h_{\text{rec}})}{\sqrt{(h_{\text{inj}}|h_{\text{inj}})(h_{\text{rec}}|h_{\text{rec}})}} \quad (2)$$

where h_{inj} is the injected waveform, h_{rec} is the recovered waveform, and $(a|b)$ is the noise-weighted inner product summed over all interferometers [46]. The average network match between the injected and reconstructed waveforms is 94%. The 94% match is consistent with the average match found using software injections [44]. An example of a reconstructed waveform is shown in Fig. 8.

Coherent WaveBurst, BayesWave, and LALInferenceBurst provide sky localization estimates. Figure 9 demonstrates sky maps for one of the GW150914 hardware injections and GW150914 itself [34]. We see similar support in Earth-bound coordinates, and nearly identical structures around the triangulation rings. The right-hand panels of Fig. 9 highlight this with the posterior distributions for the time delay between the two detectors.

C. Stochastic gravitational-wave hardware injection

A stochastic gravitational-wave background is expected to arise from the superposition of many events, each of which are too weak to resolve, but which combine to create a low-level signal [22]. By cross-correlating data from two or more detectors, it is possible to detect low-level correlations hidden beneath the detectors' noise [50]. The stochastic background from unresolved binary black holes is a particularly promising source, potentially within reach of advanced detectors [51]. For every binary black hole observed by Advanced LIGO, there are many more, which contribute to the stochastic background. On October 23,




FIG. 8. BayesWave median reconstruction and 90% credible interval (blue) and the injected waveform (red); time is shown on the x axis and whitened strain on the y axis. The data has been whitened using the estimated noise curve from the time of the injection. The network match for this waveform is 98%.

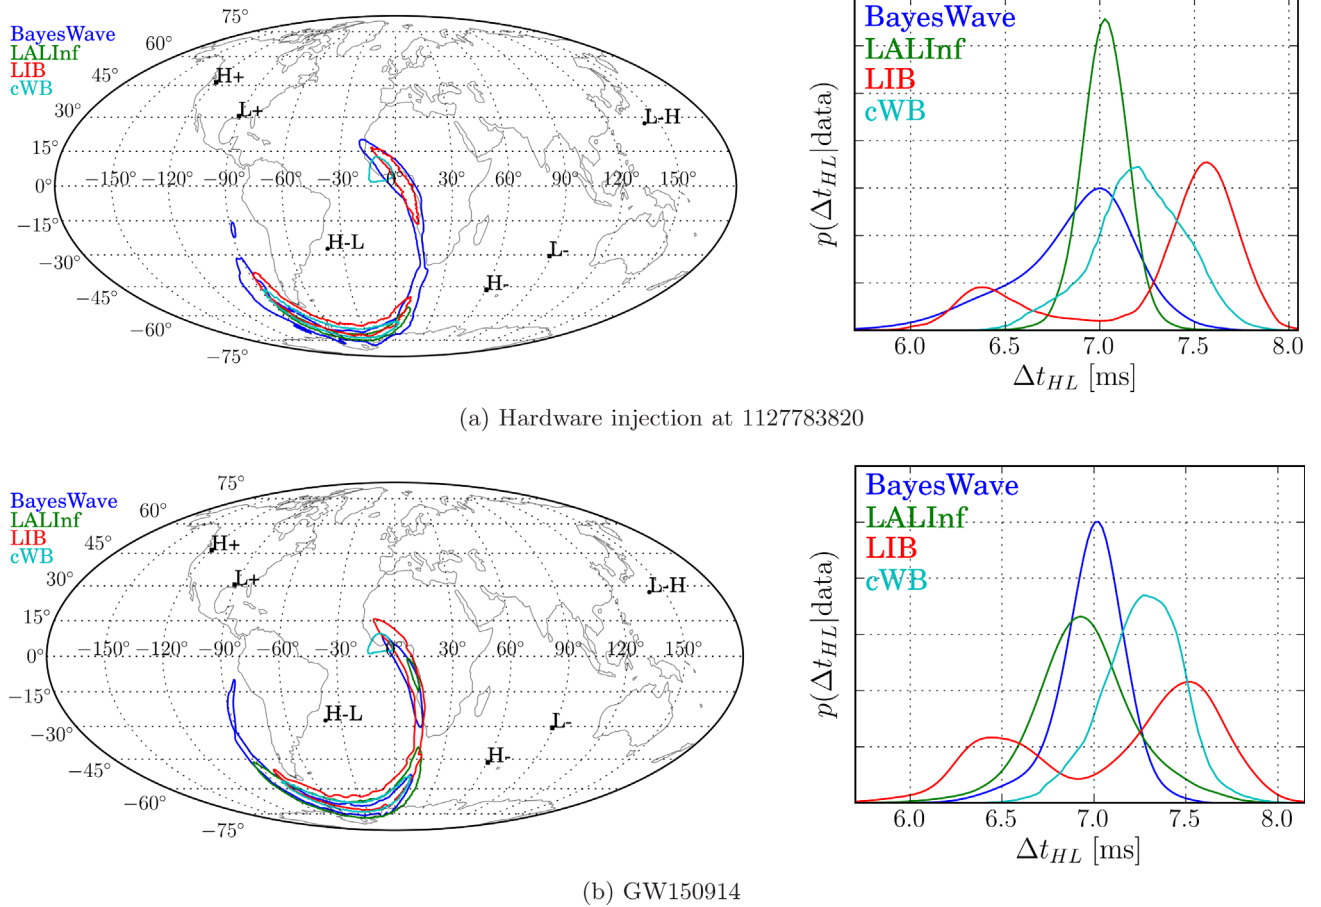


FIG. 9. Sky localization estimates for a GW150914 hardware injection and GW150914 itself from burst analyses: BayesWave, Coherent WaveBurst (cWB), and LALInferenceBurst (LIB). We include the parameter estimation analysis from Sec. III A (LALInf) for comparison [34]. Left: We show the localization maps in Earth-bound coordinates. Right: To highlight the similar positions relative to the detectors the marginal distributions for the time delay between the two detectors is shown. Note that the horizontal scale is much smaller than the 10 ms light travel time between the two detectors.

2015, a simulated stochastic gravitational-wave background signal was simultaneously injected into both detectors. The 600 s-long signal corresponded to one specific realization of an isotropic Gaussian background. The background from binary black holes is actually non-Gaussian, but the standard stochastic search makes no distinction between Gaussian and non-Gaussian signals.

The strength of a stochastic gravitational-wave signal is parametrized by the fractional contribution of the energy density in gravitational waves to close the Universe [50]:

$$\Omega_{\text{GW}}(f) = \frac{1}{\rho_c} \frac{d\rho_{\text{GW}}}{d \ln f}. \quad (3)$$

Here, ρ_c is the critical energy density of the Universe, f is frequency, and $d\rho_{\text{GW}}$ is the energy density between f and $f + df$. The injected signals were chosen such that $\Omega_{\text{GW}}(f) = 8.7 \times 10^{-5}$. This corresponds to a strain power spectral density of

$$\begin{aligned} S_h(f) &= \frac{3H_0^2}{10\pi^2} \frac{\Omega_{\text{GW}}(f)}{f^3} \\ &= (2.9 \times 10^{-23} \text{ Hz}^{-1/2})^2 \left(\frac{25 \text{ Hz}}{f}\right)^3 \end{aligned} \quad (4)$$

where H_0 is the Hubble constant. Figure 11 shows the spectrum of the injected signal and the amplitude spectral density of the noise at Hanford and Livingston near the time of the injection. The signal was low-pass filtered below 500 Hz prior to injection.

We carried out a cross-correlation search following the standard procedure [52]. The data was split into 50%-overlapping, 60 s intervals, and utilizing coarse-grained 0.25 Hz-wide frequency bins, we recovered an $\Omega_{\text{GW}}(f)$ of $(8.8 \pm 0.6) \times 10^{-5}$, consistent with the injected value. Our injection recovery assumes that $\Omega(f)$ is constant. The recovered signal [an unbiased estimator for $\Omega(f)$ in Eq. (3)] is shown in Fig. 10. The y axis shows the recovered signal as a function of the time lag between the detectors.




FIG. 10. Recovered stochastic signal as a function of lag time between detectors (blue). The detection statistic is an unbiased estimator for $\Omega(f)$ defined in Eq. (3). The injected signal is represented with a red circle. The peak of the recovered signal is at zero lag, indicating a successful recovery.

A peak at zero, and the absence of structure at other times, shows that the signal is recovered as expected.

D. Continuous gravitational-wave hardware injections

The recovery of hardware injections is used by continuous-wave searches as an end-to-end validation of the analyses in the presence of instrumental artifacts and imperfect instrument calibration. Coherent searches for gravitational waves from known pulsars are sensitive to deviations from the injected signal since a small bandwidth around the gravitational-wave frequency is integrated for months or




FIG. 11. Amplitude spectral density of the injected stochastic signal compared to the instrumental noise at Hanford and Livingston.

years [23]. These searches have the capability of monitoring the self-consistency of the interferometer calibration and, in particular, the long-term stability of absolute phase recovery. Continuous-wave searches can be implemented using a variety of methods, ranging from highly targeted searches based on ephemerides inferred from observed electromagnetic pulsations [23], to systematic templated searches for sources with previously unknown frequencies and sky locations [53–55], to searches for excess radiation flux in narrow frequency bands [56,57]. Here, we consider a coherent search based on Bayesian recovery of signal parameters [23] to validate the fidelity of hardware injections. This analysis can be used to cross-check elements of the instrument calibration, including proper coherence of data from interferometers separated by thousands of kilometers, which are sensitive to timing errors.

Figure 12 shows the posterior probability density function (PDF) for strain amplitude h_0 for simulated continuous sinusoidal gravitational waves emitted from a spinning neutron star, which has a signal frequency near 108.9 Hz and a nearly linear polarization. The signal is recovered with an amplitude consistent with the intended strength, within calibration and actuation uncertainties. Similarly, Fig. 12 also shows the recovered phase constant for this injection. In the analysis there is no compensation for the time delay between the hardware injection excitation channel and the detector output channel. The observed phase constant is consistent with the expected phase constant given the known (but uncompensated) time delay.

During Advanced LIGO’s first observing run, the hardware injections included 15 simulated continuous gravitational wave as would be emitted by spinning neutron stars. The signals were streamed in real time with frequencies ranging from 12–1991 Hz. Figures 13 and 14 show a summary of the agreement between recovered and intended amplitude and phase for the 14 injections with sufficient signal-to-noise ratio to permit recovery. Instrumental noise at the lowest-frequency injection (12 Hz) proved too large to permit recovery of simulated signal P13. There is evidence of a constant uncompensated time delay of about $150 \mu\text{s}$ in the time-domain actuation, which manifests as a phase delay increasing linearly with injection frequency in Fig. 14. In the future, a compensating timing advance will be included in the inverse actuation filter, and the simulated continuous-wave sources’ amplitudes will be increased for more precise and rapid validation of the hardware injection signal.

The overall sign of the calibration is important in order to detect and estimate the parameters of astrophysical signals correctly. An incorrect sign on the calibration would invert the signal in one detector and the parameter estimates would be incorrect. The continuous-wave injections were used as an additional check on the sign of the calibration between the Advanced LIGO detectors, since an incorrect sign would lead to a relative phase offset between the two




FIG. 12. P03 posterior probability density functions for the recovered strain amplitude and phase constant for the injected spinning neutron star signal at 108.86 Hz (referred to as P03). Note that the horizontal scales have suppressed zeroes. The dashed vertical lines indicate the intended injection amplitude and phase in radians. The red and green curves indicate the separately recovered amplitudes and phases for the Hanford and Livingston interferometers, respectively. The small discrepancies in amplitude (10%) and phase (0.1 radian = 5 degrees) fall within the uncertainties of the actuation system used for the injections.




FIG. 13. Comparisons of recovered signal amplitudes for the Hanford and Livingston signals for the 14 recovered continuous-wave injections (P00-P14). First panel: Fractional amplitude difference [(recovered minus injected)/injected]. Second panel: Amplitude values (recovered and injected).

detectors in Fig. 14. We found the sign of the calibration to be correct.

E. Loud hardware injections for detector characterization

Noise artifacts in Advanced LIGO data adversely affect the output of gravitational-wave search analyses [8,9]. In

searches for transient gravitational waves, some periods of time are excluded from the analysis to remove periods of poor data quality and known transient noise. These are known as “data quality vetoes” [8,9]. Removing periods of time with excess noise improves the performance of gravitational-wave searches [8,9]. Some of these data quality vetoes are derived from information recorded in




FIG. 14. Comparisons of recovered signal phases for the Hanford and Livingston signals for the 14 recovered continuous-wave injections (P00-P14). First panel: Phase difference (degrees) [recovered minus true]. Second panel: Conversion of phase difference magnitude to time difference. Third panel: Same as first panel but with a linear horizontal scale, to indicate the roughly linear dependence of residual phase difference, consistent with a constant known (but uncompensated for) time delay.

auxiliary channels. Auxiliary channels include instrumental channels that record degrees of freedom of the interferometer and its isolation systems as well as channels that monitor the environmental conditions around the instrument [58]. The environmental monitoring system includes seismic, acoustic, and electromagnetic data.

To avoid discarding true gravitational-wave signals, any auxiliary channels used for vetoes are first checked to ensure that they do not respond to gravitational-wave-like signals; i.e., changes in differential arm length. This process is referred to as a “safety check,” since a channel that has no sensitivity to gravitational waves is considered “safe” for use when constructing a veto. To test whether auxiliary channels respond to differential arm length changes, three sets of 12 loud (matched-filter signal-to-noise ratios > 100) transient hardware injections were performed at both detectors, and the auxiliary channel data were examined both qualitatively and quantitatively for signs of coupling.

Spectrograms were manually inspected at the time of hardware injections. These signals were very strong and

clear, with high signal-to-noise ratio, in channels that were expected to record differential displacement, e.g. interferometer differential sensing and actuation, and closely related degrees of freedom. No signs of coupling were found in thousands of other auxiliary channels, indicating that they may be used to construct vetoes. Hundreds of time-frequency representations of auxiliary channels were also inspected at the times of GW150914 and GW151226 with the same outcome [8].

Loud hardware injections were used to statistically assess the coupling. An algorithm based on a transformation using sine-Gaussians [59] was used to identify and parametrize noise transients by their time, frequency, and signal-to-noise ratio. The time of noise transient is compared with the times of the loud hardware injections.

For each channel, the number of noise transients that occurred within 100 ms of loud injections are counted and compared to the number that would be expected based on chance [60]. For any channel exhibiting a higher number of overlaps than expected by chance, the time-frequency behavior of the raw data is further investigated to see if

there is a plausible connection. We find that only obviously related channels, such as those in the sensing and actuation chain for the differential length control loop, were sensitive to the loud hardware injections.

IV. CONCLUSIONS

This paper presents the Advanced LIGO hardware injection system infrastructure for injecting simulated gravitational-wave signals into the detectors by displacing the test masses, and results from Advanced LIGO's first observing run. Hardware injections were used for validating analyses after a gravitational-wave detection, as an additional check of the calibration, and characterizing the detectors' response to differential arm length variations.

After the detection of GW150914 and GW151226, sets of binary black hole merger waveforms with similar parameters were injected to validate the search and parameter estimation analyses. The recovered signals were checked for consistency with the parameters of the injected waveforms, including signal-to-noise ratio, chirp mass, and sky position. Similarly, the stochastic background and continuous-wave searches used simulated waveforms as an end-to-end test.

In order to detect and estimate the parameters of astrophysical signals the calibration must be correct, and the continuous-wave injections provided an additional check of the calibration sign. They were also used to measure the time delay of the hardware injection pathway and checked that it was consistent with the predicted value from the calibration model.

Data quality vetoes are used to increase the performance of search analyses, and detector characterization hardware injections were used to identify output channels in the control system that can be used to construct data quality vetoes. After each gravitational-wave detection, we carried out a study to check for cross-couplings with the detectors'

output gravitational-wave strain channel. Channels that contained a trace of the injected signal were considered unsafe and excluded from data quality veto studies.

In the future, we plan to exclusively use the photon calibrators to inject simulated gravitational waves. Future work on the hardware injection system includes using point-by-point, Fourier-domain inverse actuation functions for each of the injected spinning neutron stars to mitigate the effect of data dropouts.

ACKNOWLEDGMENTS

LIGO was constructed by the California Institute of Technology and Massachusetts Institute of Technology with funding from the National Science Foundation (NSF), and operates under cooperative agreement PHY-0757058. Advanced LIGO was built under Grant No. PHY-0823459. Computations were carried out on the Syracuse University HTC Campus Grid which is supported by NSF Grant No. ACI-1341006. Fellowship support from the LIGO Laboratory for S. K. is gratefully acknowledged. C. B. and D. A. B. acknowledge support from NSF Grant No. PHY-1404395. K. R. acknowledges support from NSF Grant No. PHY-1505932. E. T. acknowledges support from the Australian Research Council Grant No. FT150100281 and CE170100004. P. S. acknowledges support from NSF Grant No. PHY-1404121. J. R. S. acknowledges support from NSF Grant No. PHY-1255650. J. V. acknowledges support from the Science and Technology Facilities Council Grant No. ST/K005014/1. J. L. and R. O. acknowledge support from NSF Grant No. PHY 1505629. C. B. would like to thank Laura Nuttall for providing useful suggestions and Collin Capano for the software injection data in Sec. III A. This paper carries the LIGO Document Number LIGO-P1600285.

-
- [1] J. Aasi *et al.* (LIGO Scientific Collaboration), *Classical Quantum Gravity* **32**, 074001 (2015).
 - [2] F. Acernese *et al.* (VIRGO Collaboration), *Classical Quantum Gravity* **32**, 024001 (2015).
 - [3] H. Grote (LIGO Scientific Collaboration), *Classical Quantum Gravity* **27**, 084003 (2010).
 - [4] Y. Aso, Y. Michimura, K. Somiya, M. Ando, O. Miyakawa, T. Sekiguchi, D. Tatsumi, and H. Yamamoto (KAGRA Collaboration), *Phys. Rev. D* **88**, 043007 (2013).
 - [5] B. Iyer *et al.*, LIGO-India Technical Report No. LIGO-M1100296-v2 (2011), <https://dcc.ligo.org/LIGO-M1100296/public>.
 - [6] B. P. Abbott *et al.* (Virgo and LIGO Scientific Collaborations), *Phys. Rev. Lett.* **116**, 061102 (2016).
 - [7] B. P. Abbott *et al.* (Virgo and LIGO Scientific Collaborations), *Phys. Rev. Lett.* **116**, 241103 (2016).
 - [8] B. P. Abbott *et al.* (Virgo and LIGO Scientific Collaborations), *Classical Quantum Gravity* **33**, 134001 (2016).
 - [9] J. Aasi *et al.* (Virgo and LIGO Scientific Collaborations), *Classical Quantum Gravity* **32**, 115012 (2015).
 - [10] B. P. Abbott *et al.* (LIGO Scientific Collaboration), [arXiv:1602.03845](https://arxiv.org/abs/1602.03845).
 - [11] J. Abadie *et al.* (Virgo and LIGO Scientific Collaborations), *Phys. Rev. D* **85**, 082002 (2012).
 - [12] J. Abadie *et al.* (Virgo and LIGO Scientific Collaborations), *Phys. Rev. D* **81**, 102001 (2010).
 - [13] B. P. Abbott *et al.* (Virgo and LIGO Scientific Collaborations), *Phys. Rev. Lett.* **116**, 131103 (2016).

- [14] K. Kawabe, Ph.D. thesis, University of Tokyo, 1998.
- [15] M. Rakhmanov, J. D. Romano, and J. T. Whelan, *Classical Quantum Gravity* **25**, 184017 (2008).
- [16] R. Weiss, Various reports of experiments conducted on the Barkhausen noise research (2009), <https://dcc.ligo.org/LIGO-T080355/public>.
- [17] J. R. Smith (LIGO Scientific Collaboration), *Classical Quantum Gravity* **26**, 114013 (2009).
- [18] L. Carbone *et al.*, *Classical Quantum Gravity* **29**, 115005 (2012).
- [19] S. Karki *et al.*, *Rev. Sci. Instrum.* **87**, 114503 (2016).
- [20] B. P. Abbott *et al.* (Virgo and LIGO Scientific Collaborations), *Astrophys. J.* **832**, L21 (2016).
- [21] B. P. Abbott *et al.* (Virgo and LIGO Scientific Collaborations), *Phys. Rev. X* **6**, 041015 (2016).
- [22] M. Maggiore, *Phys. Rep.* **331**, 283 (2000).
- [23] J. Aasi *et al.* (LIGO Scientific Collaboration), *Astrophys. J.* **785**, 119 (2014).
- [24] J. Rollins, Advanced LIGO Guardian documentation (2015), <https://dcc.ligo.org/LIGO-T1500292/public>.
- [25] LSC Data Analysis Software Working Group, LALSuite, <https://www.lsc-group.phys.uwm.edu/daswg/projects/lalsuite.html>.
- [26] L. Canete, Implementation of an optical follower servo for the aLIGO Pcal (2014), <https://dcc.ligo.org/LIGO-T1300442/public>.
- [27] E. Goetz *et al.*, *Classical Quantum Gravity* **26**, 245011 (2009).
- [28] A. Taracchini, Y. Pan, A. Buonanno, E. Barausse, M. Boyle, T. Chu, G. Lovelace, H. P. Pfeiffer, and M. A. Scheel, *Phys. Rev. D* **86**, 024011 (2012).
- [29] B. Allen, W. G. Anderson, P. R. Brady, D. A. Brown, and J. D. E. Creighton, *Phys. Rev. D* **85**, 122006 (2012).
- [30] S. A. Usman *et al.*, *Classical Quantum Gravity* **33**, 215004 (2016).
- [31] A. Nitz *et al.*, PyCBC software (2016), <https://doi.org/10.5281/zenodo.197080>.
- [32] B. P. Abbott *et al.* (Virgo and LIGO Scientific Collaborations), *Phys. Rev. D* **93**, 122003 (2016).
- [33] J. Veitch *et al.*, *Phys. Rev. D* **91**, 042003 (2015).
- [34] B. P. Abbott *et al.* (Virgo and LIGO Scientific Collaborations), *Phys. Rev. Lett.* **116**, 241102 (2016).
- [35] P. Schmidt, F. Ohme, and M. Hannam, *Phys. Rev. D* **91**, 024043 (2015).
- [36] M. Hannam, P. Schmidt, A. Bohè, L. Haegel, S. Husa, F. Ohme, G. Pratten, and M. Pürrer, *Phys. Rev. Lett.* **113**, 151101 (2014).
- [37] S. Babak *et al.*, *Phys. Rev. D* **87**, 024033 (2013).
- [38] B. Allen, *Phys. Rev. D* **71**, 062001 (2005).
- [39] L. Singer *et al.* (2015), <http://gcn.gsfc.nasa.gov/other/GW150914.gcn3>.
- [40] L. Singer *et al.* (2015), <http://gcn.gsfc.nasa.gov/other/GW151226.gcn3>.
- [41] LIGO and Virgo Scientific Collaborations, 2016, <https://losc.ligo.org/events/GW151226/>.
- [42] B. P. Abbott *et al.* (Virgo and LIGO Scientific Collaborations), *Phys. Rev. D* **94**, 064035 (2016).
- [43] C. L. Fryer and B. K. C. New, *Living Rev. Relativ.* **14**, 1 (2011).
- [44] B. P. Abbott *et al.* (Virgo and LIGO Scientific Collaborations), *Phys. Rev. D* **93**, 122004 (2016); **94**, 069903(E) (2016).
- [45] S. Klimenko *et al.*, *Phys. Rev. D* **93**, 042004 (2016).
- [46] N. J. Cornish and T. B. Littenberg, *Classical Quantum Gravity* **32**, 135012 (2015).
- [47] J. Skillings, *Bayesian Anal.* **1**, 833 (2006).
- [48] J. Abadie *et al.* (Virgo and LIGO Scientific Collaborations), *Phys. Rev. D* **81**, 102001 (2010).
- [49] V. Tiwari, S. Klimenko, V. Necula, and G. Mitselmakher, *Classical Quantum Gravity* **33**, 01LT01 (2016).
- [50] B. Allen and J. D. Romano, *Phys. Rev. D* **59**, 102001 (1999).
- [51] V. Mandic, S. Bird, and I. Cholis, *Phys. Rev. Lett.* **117**, 201102 (2016).
- [52] J. Aasi *et al.* (Virgo and LIGO Scientific Collaborations), *Phys. Rev. Lett.* **113**, 231101 (2014).
- [53] B. Abbott *et al.*, *Phys. Rev. D* **77**, 022001 (2008).
- [54] B. P. Abbott *et al.*, *Phys. Rev. D* **94**, 042002 (2016).
- [55] B. P. Abbott *et al.* Virgo and LIGO Scientific Collaborations, *Phys. Rev. D* **94**, 102002 (2016).
- [56] S. W. Ballmer, *Classical Quantum Gravity* **23**, S179 (2006).
- [57] B. P. Abbott *et al.* (Virgo and LIGO Scientific Collaborations), *Phys. Rev. Lett.* **107**, 271102 (2011).
- [58] A. Effler, R. M. S. Schofield, V. V. Frolov, G. González, K. Kawabe, J. R. Smith, J. Birch, and R. McCarthy, *Classical Quantum Gravity* **32**, 035017 (2015).
- [59] S. Chatterji, L. Blackburn, G. Martin, and E. Katsavounidis, *Classical Quantum Gravity* **21**, S1809 (2004).
- [60] J. R. Smith, T. Abbott, E. Hirose, N. Leroy, D. Macleod, J. McIver, P. Saulson, and P. Shawhan, *Classical Quantum Gravity* **28**, 235005 (2011).