Characterizing Quasi-Periodic Low-Frequency Noise During O4 at LIGO Livingston

Lanie DeMarcay[1], Jane Glanzer[2]

[1]Davidson College, Davidson, NC 28035, USA
[2]LIGO Laboratory, California Institute of Technology, Pasadena, CA 91125, USA

Abstract.

During LIGO's fourth observing run, O4, the Livingston detector's binary neutron star (BNS) range exhibited recurring oscillations, appearing as noise in the 30-50 Hz band of the gravitational-wave strain data. This project uses frequency-domain analysis to identify auxiliary channels that may be responsible for this noise with the goal of mitigating it before the next observing run. It is important to reduce this excess noise as it is located in the region where the interferometer is particularly sensitive to the mergers of higher mass binary black holes. We investigated whether auxiliary channels were coupled with the strain oscillations by using coherence and bicoherence analyses. The survey comparing ordinary coherence between the strain and thousands of auxiliary channels found no channel with a consistent linear relationship, suggesting the coupling may instead be nonlinear. We have since shifted to a segmented bicoherence survey, sensitive to quadratic phase coupling between frequency components, which has revealed structured, localized patterns for several channels within the oscillation's frequency band. Continued analysis across additional detector subsystems and oscillation days is needed to determine whether these patterns indicate a genuine physical driver of the noise.

1. Introduction

Gravitational waves, often thought of as "ripples in spacetime", are generated by asymmetrically accelerating massive objects such as black holes and neutron stars. First predicted by Albert Einstein in his general theory of relativity [1], these waves contain critical information on the behavior of some of the most energetic events in the universe. The first direct detection of gravitational waves occurred nearly a century after they were first theorized on September 14, 2015, when the Laser Interferometer Gravitational-Wave Observatory (LIGO) observed the spiraling and merger of a binary black hole system [2]. This marked the growth of an entirely new field of gravitational-wave astrophysics, and since that first detection over a decade ago, the LIGO-Virgo-KAGRA Collaboration (LVK) has cataloged 390 gravitational-wave events, dramatically expanding our understanding of the universe [3].

LIGO operates two observatories: one in Livingston, Louisiana (LLO) and one in Hanford, Washington (LHO) [2]. The facilities use modified Michelson interferometers

to measure small changes in optical path lengths due to space time fluctuations, measuring gravitational waves (GW) as a time series. The event produces a change in the differential arm (DARM) lengths such that, $\Delta L_{\text{free}} = \Delta L_x(t) - \Delta L_y(t)$, which is related to the gravitational wave strain, $h(t) = \Delta L_{\text{free}}(t)/L$. Here, L is the proper length of one of the interferometer's arms, which is approximately 4 km for the LIGO detectors. The ratio of the differential arm-length change to the proper arm length gives the dimensionless gravitational-wave strain, $h(t)$, which represents the change in the interferometer's arm length caused by the gravitational wave. The calibrated gravitational wave strain, $h(t)$, is not directly measured; it is reconstructed from the interferometer's digitized error signal, which encodes differential arm (DARM) displacement as fluctuations in the laser power at the gravitational wave readout port [4]. The GW strain being measured by these interferometers is incredibly small, on the scale of 1/1000th of the diameter of the nucleus of an atom, and therefore requires extraordinarily precise instrumentation [5]. Due to this extreme sensitivity, LIGO data is susceptible to noise from a variety of environmental and instrumentation factors such as passing trains [6] or scattered light [7, 8].

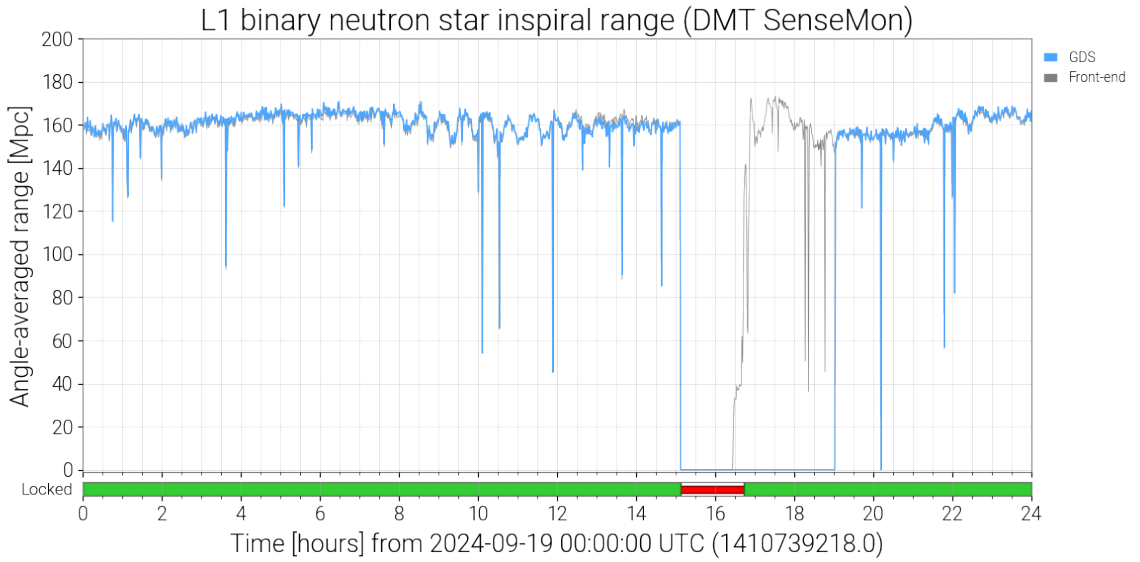
Understanding and mitigating these noise issues is referred to as Detector Characterization (DetChar). The DetChar team monitors the status of the detectors, investigates data quality, examines sources of instrumental noise, and ensures that noise artifacts are not misidentified as an astrophysical signal [9]. They work with numerous auxiliary channels, such as accelerometers and seismometers, to monitor various parts of the detector. One important metric for quantifying LIGO's sensitivity is the binary neutron star (BNS) range, defined as the distance to which a pair of 1.4 solar mass neutron stars can be detected with a signal-to-noise ratio (SNR) of 8, and is given by

$$R = \frac{1}{\rho_0}\left(\frac{5\pi}{24c^3}\right)^{1/2}(GM)^{5/6}\pi^{-7/6}\frac{1}{2.26}\left[4\int_0^\infty \frac{f^{-7/3}}{S_n(f)}\,df\right]^{1/2} \tag{1}$$

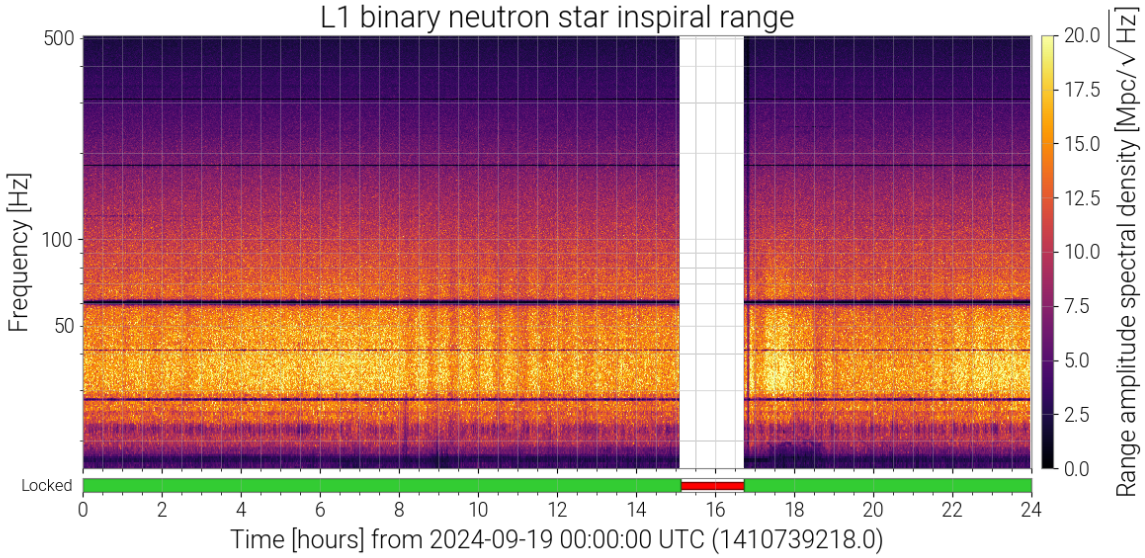
[10]. Because of the weighting of $f^{-7/3}$, low-frequency noise is particularly damaging to the BNS range and can significantly impact LIGO's detection rate of high mass binary black hole mergers.

During the fourth LIGO observing run (O4), the BNS range at the Livingston Observatory (LLO) exhibited approximately 30-minute periodic oscillations [11], which manifested as excess noise in the 30–50 Hz band of the main gravitational-wave strain channel, as seen in Figure 1. The correlated behavior observed across many auxiliary channels prompted a systematic investigation into the origin of these oscillations.

Identifying the origin of such oscillations is important for two reasons. First, if the cause can be understood, the oscillations may be mitigated or eliminated, improving the quality of GW strain data in the low frequency range. Second, understanding this noise source contributes to the broader knowledge base of how environmental and instrumental effects impact LIGO data. This project begins the process of searching for the the origin of these oscillations by looking for auxiliary channels whose relationship to the gravitational-wave strain is specifically tied to the oscillation, first through coherence



(a) Time series of the binary neutron star (BNS) inspiral range for the LIGO Livingston detector on September 19, 2024. The detector exhibits approximately 30-minute periodic oscillations from 08:00 UTC to 13:00 UTC. The gaps in the figures correspond to periods when the detector lost lock, meaning that the interferometer was no longer operating stably for strain measurements, which can be caused by environmental disturbances, such as earthquakes.



(b) Time–frequency spectrogram of the binary neutron star (BNS) inspiral range for the LIGO Livingston detector on September 19, 2024. The plot displays the 30-minute periodic oscillations in discolored striated bands in the 30–50 Hz range from 08:00 UTC to 13:00 UTC.

Figure 1

and, since that proved insufficient, through bicoherence.

2. Objectives

The goal of this project is to identify the source, or at least narrow down the location, of the quasi-periodic BNS range oscillations observed at LLO during the O4 period, building on the previous summer's time-lagged cross-correlation work by Gabriel Grant and Jane Glanzer (LIGO SURF 2025) [12]. To achieve this objective, rather than searching for channels that are simply correlated with the strain at all times for unrelated reasons, this project segments the strain data into oscillation "peak" and "trough" epochs and searches for auxiliary channels whose behavior is specifically associated with these oscillation epochs. We used two complementary statistical lenses to determine if a given auxiliary channel was linked with the oscillation, and not just generically coupled to the strain: ordinary (linear) coherence, and bicoherence, which is sensitive to a nonlinear, quadratic form of phase coupling. Identifying such a channel, or narrowing the search down to a particular interferometer subsystem, would mark real progress toward mitigating this noise source and would add to the broader understanding of how environmental and instrumental noise impacts LIGO's strain data.

3. Methods

This project builds on last summer's work by Gabriel Grant and Jane Glanzer through LIGO SURF 2025, whose goal was to better understand these oscillations as they appear in the BNS inspiral range at LLO. Over the course of this summer, the analysis developed into a four-stage pipeline, and each stage was built out into its own standalone, reusable codebase rather than a collection of one-off scripts, so that the search could be pointed at any new day or channel list without hand-editing the underlying analysis logic. The four stages are: (0) cataloging every observed instance of the oscillation across O4 and picking a representative subset of days to fully analyze, (1) segmenting each chosen day's strain into peak and trough epochs, (2) a segmented linear coherence survey between the strain and auxiliary channels, and (3) a segmented bicoherence survey, undertaken after the linear survey failed to identify a channel with a consistent peak-versus-trough coherence. Each of these four stages is described in detail below; what each survey actually found is reported separately in Section 4.

3.1. Stage 0: Finding and cataloging every oscillation instance

Before any strain data was analyzed, every instance of the oscillations had to be identified and logged. This turned out to be a genuinely time-consuming task, as I manually examined the DetChar summary pages [9] across the 909-day calendar span of O4 [13] and cataloged each instance in a working Overleaf document (DCC LIGO-T2600326-v1). The process for a candidate day was as follows: I began by examining the time series and time–frequency spectrogram on the LIGO DetChar summary pages

to visually check for an instance of the oscillation (Figure 1). These are automatically generated daily plots produced by the detector-characterization team. A day was flagged as a candidate whenever it displayed the characteristic roughly 30-minute quasi-periodic oscillations concentrated in the 30–50 Hz band of the strain channel, consistent with the behavior first reported by Glanzer [11].

Every flagged day was then logged by hand in a running Overleaf catalog, recording four pieces of information for each event: the calendar date, UTC start time, UTC end time, and a qualitative clarity rating. The clarity rating uses five levels to describe how clearly the oscillation stood out from the background BNS-range noise on a given day. The resulting catalog now covers the entire O4 observing run and represents a substantial portion of the early summer's work. Building it required manually reviewing the DetChar summary pages for every day on which LIGO Livingston collected data during O4, before any subsequent strain-data analysis could begin for correlation with auxiliary channels.

The data from the catalog of candidate days were then converted into a plain CSV file to be used in an exploratory-analysis script. This code was developed to first check for any obvious patterns between the instances of the oscillations. It checked for basic distributions to identify any commonalities in start time, end time, and duration of every logged oscillation, and it also examined patterns for a number of other factors, such as time of day, day of the week, and time of year. A single summary plot of all of these results is shown in Figure 2. As can be seen, there does not appear to be a clear pattern in the oscillations, allowing us to rule out an obvious time-based cause and move into the coherence analysis framework.

Using the results from this section of the analysis, we selected nine representative days for the full coherence analysis: June 7, 2023; November 5, 2023; December 14, 2023; March 30, 2024; September 5, 2024; September 18, 2024; September 19, 2024; October 21, 2025; and November 17, 2025. These were deliberately chosen to be spread across different months and points in the observing run, favoring high-clarity events with different oscillation morphologies, so that any coupling ultimately found would be more likely to reflect something genuinely tied to the oscillations.

3.2. Stage 1: Epoch segmentation

For each of the nine selected days, the calibrated strain channel (`L1:GDS-CALIB_STRAIN`) is bandpassed to the 30–50 Hz band of interest and notch-filtered at known frequencies, such as 41 Hz to remove a known suspension resonance and 60 Hz to remove the known power grid frequency, to remove loud spikes that would otherwise impact the epoch segmentation code. A band-limited RMS (BLRMS) is then computed: the bandpassed strain is broken into short, fixed-length windows called a "stride," typically 60–90 seconds long, and the root-mean-square amplitude within each window becomes one sample of the BLRMS trace. The result is a slowly varying shape that rises when the 30–50 Hz oscillation is loud and falls when it is quiet. Once the loud noise spikes are

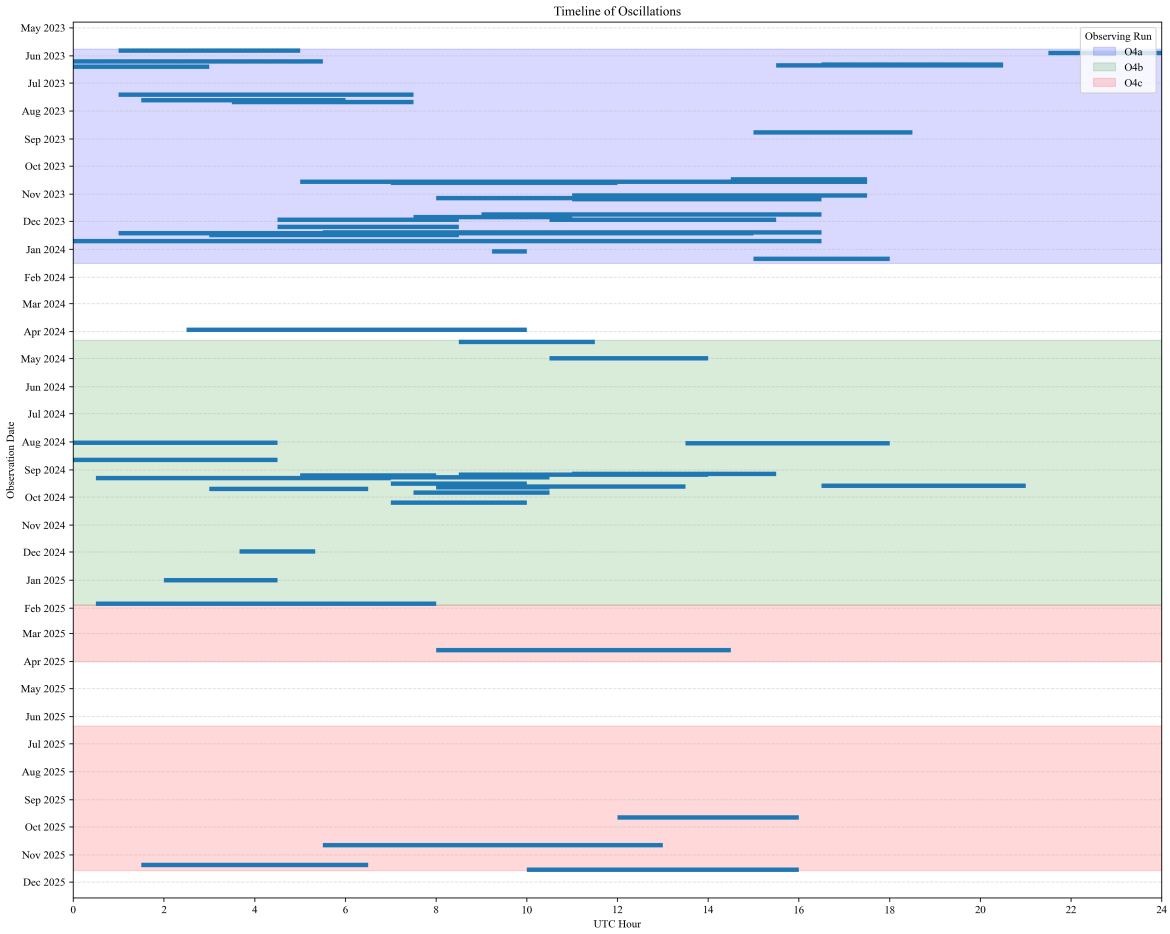


Figure 2: Timeline of the observed oscillations throughout the fourth observing run (O4), with each horizontal line representing an individual oscillation observation. The horizontal axis gives the UTC hour at which the oscillation was observed, while the vertical axis shows the observation date. The shaded regions indicate the O4a, O4b, and O4c observing periods. The gap within O4c from April 1 to June 11, 2025, corresponds to the commissioning period between the two O4c observing periods.

clipped, the BLRMS trace is normalized onto the interval $[0, 1]$.

The normalized BLRMS is lightly smoothed and then thresholded at two percentiles: samples above the upper percentile, typically around 55%, are labeled part of a peak epoch, and the sufficiently long gaps between peak epochs are labeled trough epochs. This peak/trough segmentation is the most important idea underscoring the entire search code structure. It is what allows both the coherence and bicoherence analyses to isolate behavior specific to the oscillation itself, by comparing what a channel looks like during the oscillation against what that same channel looks like when the oscillation is quiet.

A central practical challenge in this stage is that the oscillation does not have a single, uniform shape across days: some days show sharp, narrow, "fin-like" peaks,

while others show broad, flat-topped, "mesa-like" peaks, and the overall amplitude of the oscillation also varies significantly from day to day. A single fixed set of segmentation parameters (bandpass band, BLRMS stride, peak/trough percentile thresholds, minimum epoch duration, smoothing width, clipping aggressiveness) that performed well on one day worked poorly on other days. This was addressed with an interactive, per-day tuning notebook: for a given day, the notebook fetches and bandpasses the strain, computes the BLRMS, applies the epoch-labeling logic above, and overlays the resulting peak (blue) and trough (orange) shaded regions directly on the BLRMS trace, as shown in Figure 3. The user then compares and tweaks the parameters for a "Version A" and "Version B," until whichever version's shading visibly matches the true oscillation pattern is selected and saved.

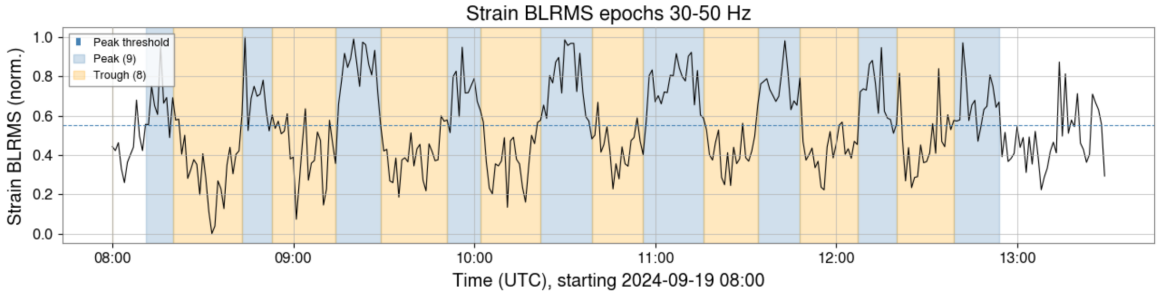


Figure 3: Normalized 30–50 Hz strain BLRMS for the September 19, 2024 oscillation (08:00–13:30 UTC) using the optimized segmentation parameters. The detected peak (blue) and trough (orange) epochs closely follow the observed oscillation, producing long, contiguous intervals that accurately distinguish periods of high and low strain amplitude.

3.3. Stage 2: Linear coherence search

3.3.1. Defining coherence Coherence is the first tool to use when examining whether a given signal is related to another signal in the frequency domain. It measures, at each frequency f, how much of the power in one signal can be linearly predicted from the other. Formally, coherence between the strain channel and an auxiliary channel is

$$C_{xy}(f) = \frac{|S_{xy}(f)|^2}{S_{xx}(f)\, S_{yy}(f)}, \text{[14]} \tag{2}$$

where $S_{xy}(f)$ is the cross-spectral density between the strain and the auxiliary channel, and $S_{xx}(f)$, $S_{yy}(f)$ are their respective power spectral densities. Coherence is bounded between 0, completely unrelated at that frequency, and 1, perfectly linearly related at that frequency. It is important to note that it only captures a linear relationship: if channel A at frequency f predicts channel B at that *same* frequency f, coherence will find it, but if the relationship only shows up as some combination of frequencies (as described in Section 3.4), coherence cannot detect it. Coherence was chosen as

the starting point for this project precisely because it is the simplest and most well-understood tool available, so it was a reasonable first hypothesis to rule in or out.

3.3.2. Running the segmented coherence survey Rather than computing a single coherence estimate over an entire day's observing window, which would fail to differentiate between the oscillation's "peak" and "trough" behavior, coherence is computed separately for every individual epoch identified in Stage 1. The results from this segmented coherence survey are each averaged into a mean peak-coherence spectrum and a mean trough-coherence spectrum per channel. Channels are then ranked by a Δ score, defined as the difference between the summed coherence during peak epochs and during trough epochs, integrated over the 30–50 Hz band of interest. A large positive Δ score indicates a channel whose linear relationship with the strain gets stronger during the oscillation, rather than a channel that happens to be coherent with the strain at all times. To further filter these results, we added a 95% statistical significance threshold, so that coherence values below this level are not statistically significant beyond pure noise, regardless of how large a channel's raw Δ score is. A final consistency metric was also computed: the fraction of individual peak or trough segments whose coherence independently clears this significance threshold. This metric protects against the scenario where a channel's promising average score is actually being driven by a single unusually loud segment, rather than by a signal that is genuinely and consistently present across many separate peak epochs.

The survey itself is organized by interferometer subsystem, since running all approximately two thousand auxiliary channels at once is both computationally intensive and hard to interpret. Grouping by physical subsystem allowed us to better interpret and draw trends across our results. We used ten subsystems: arm-length stabilization (ALS), angular sensing and control (ASC) divided into loops and sensors, input mode cleaner and pre-stabilized laser (IMC/PSL), length sensing and control and calibration (LSC/CAL), output mode cleaner and thermal compensation (OMC/TCS), squeezed-light (SQZ), and suspension (SUS) channels at the corner, end-X, and end-Y stations.

We ran the coherence analysis across each subsystem and day. For every subsystem and day, the script produces an HTML report containing the epoch-overview plot (the same tuned BLRMS trace shown in Figure 3), a ranked channel table sorted by Δ score, and, for every channel, a four-panel summary plot. The top-left panel overlays the averaged peak- and trough-epoch coherence spectra with the 95% significance threshold. The top-right panel shows the peak-minus-trough difference as a function of frequency, highlighting which frequency bands are most coherent during the oscillation. The bottom-left panel compares the strain's amplitude spectral density (ASD) to the noise this auxiliary channel could plausibly account for during peak and trough epochs. The bottom-right panel summarizes the numerical scores and consistency metrics as plain text. Examples of the ranked table and a representative four-panel plot are shown in Section 4.1, alongside the results they were used to produce.

Once we had completed running the entire subsystem analysis, we developed a pipeline that combines all ten subsystems' channel data into one day-wide ranked list, allowing us to compare across each full day for any particularly interesting results. We then took the results from each of the nine days and developed a cross-day script that pools every day's ranked list across the full nine-day trial and examines whether a channel's or a subsystem's high score recurs across multiple days, rather than appearing on only one. Recurrence across days would provide much stronger evidence of a real, physical coupling than a single day's result, since a one-off high score could easily be driven by an unrelated coincidence that happened to occur on that particular day and had nothing to do with the oscillation mechanism itself. What these two comparison pipelines found, and why it motivated the shift to bicoherence, is described in Section 4.1.

3.4. Stage 3: The segmented bicoherence search

3.4.1. Bicoherence as the next step Once the linear coherence survey had been run thoroughly enough to rule out an easy answer, we determined that the coupling mechanism itself was likely something ordinary coherence could not detect, meaning it was a nonlinear relationship. Coherence, by its definition (Equation 2), can only detect a relationship that remains at the same frequency in both signals: if channel A's signal at 30 Hz predicts channel B's signal at 30 Hz, coherence is able to detect it, regardless of what else might be happening at other frequencies. A coherence search, by construction, is blind to any nonlinear mechanism entirely, even if the coupling is perfectly real and strong. Bicoherence exists to cover this type of behavior, which is what made it the obvious next step once coherence had been thoroughly exhausted.

3.4.2. Defining bicoherence Bicoherence is a frequency-domain method for detecting a specific kind of nonlinear relationship, quadratic coupling between signals [15]. This framework examines whether the power at frequency f_1 combined with the power at frequency f_2 in one signal predicts power at the sum frequency $f_1 + f_2$ in the other. Two frequency components at f_1 and f_2 are said to be phase-coupled if they combine to generate a third component at the sum frequency $f_3 = f_1 + f_2$ with a fixed, non-random phase relationship. In our work, the bicoherence algorithm takes two frequency components from the same auxiliary channel and tests whether their interaction is coupled to the gravitational-wave strain at the sum frequency, $f_1 + f_2$. This behavior is formally captured by the bispectrum,

$$B(f_1, f_2) = \langle X(f_1)\, X(f_2)\, X^*(f_1 + f_2) \rangle\,, \tag{3}$$

where $X(f)$ is the Fourier transform of the signal being examined and $\langle \cdot \rangle$ denotes averaging over many independent time segments; X^* is the complex conjugate.

If f_1 and f_2 really are phase-coupled through some real, physical, quadratic mechanism, the phase of the product $X(f_1)X(f_2)X^*(f_1 + f_2)$ will be nearly the same across segments. If f_1 and f_2 are not quadratically phase coupled, the phase of that same

product will be random from segment to segment, and averaging drives the bispectrum to zero. Bicoherence is simply the bispectrum normalized onto the interval $[0, 1]$, so that its value can be interpreted the same way across different channels and different noise levels. Using the normalization introduced by Hagihira *et al.* [16], we get

$$b^2(f_1, f_2) = \frac{|B(f_1, f_2)|^2}{\langle |X(f_1)X(f_2)X^*(f_1 + f_2)| \rangle^2}. \tag{4}$$

Bicoherence approaches 1 if f_1 and f_2 are genuinely and consistently phase-coupled, and nears 0 if they are not. Poloskei *et al.* [15] further note that a genuine physical coupling tends to produce a broad, connected region of elevated bicoherence across nearby (f_1, f_2) pairs, while a false positive driven by chance tends to appear instead as a single, scattered, isolated point. This distinction is essential for interpreting the results discussed in Section 4.2.

3.4.3. The bicoherence analysis The bicoherence script was adapted from code originally written by a collaborator of Dr. Jane Glanzer's, based on a general-purpose fast Fourier transform cross-bicoherence method. The code takes three input signals, a, b, and y, computing the cross-bicoherence between a and b (the auxiliary signals) against y (the gravitational-wave strain data). In this project's current single-channel search, a and b are both set to the same auxiliary channel's signal. The script therefore examines whether a given single auxiliary channel's own content at f_1 and f_2 combines to produce excess coupling in the GW strain at $f_1 + f_2$. Section 5 discusses the natural extension of setting a and b to two different auxiliary channels as a step for future work.

The same underlying peak/trough logic used for coherence applies here, but the epoch definitions from Stage 1 are now used as the individual time windows over which bicoherence itself is computed. Bicoherence is estimated separately for every individual peak epoch and every individual trough epoch, for every channel. The in-band search region used for scoring is restricted to (f_1, f_2) pairs whose sum, $f_1 + f_2$, falls inside the studied frequency band of 30–50 Hz where the oscillations reside. Once again, we notch out a list of known instrumental line frequencies, such as the suspension violin modes and the 60 Hz power line, so that an already-known spectral line is not mistaken for a genuine, oscillation-specific coupling. The bicoherence script catalogs a number of detailed statistics for every channel and epoch, so that results can eventually be compared consistently across subsystems and days rather than examined one channel plot at a time.

For each channel and epoch, `z_excess` measures how strongly the number of bicoherence bins exceeding the significance threshold differs from the number expected under the Gaussian noise background. These epoch-level `z_excess` values are averaged across all of a channel's peak epochs for a given day and then ranked against that day's own noise floor, the median `z_excess` across every channel surveyed in that report, producing the `z_over_floor` score used to rank channels. A further `z_score` standardizes `z_over_floor` for comparison across channels and days by dividing by the

median absolute deviation (MAD) of `z_excess` across the subsystem. This gives a value indicating how many robust deviations a given channel sits above the median, allowing us to compare across subsystems and days using one consistent metric.

For every channel, the script also builds a panel containing that channel's peak and trough epoch heatmaps in chronological order, using the same blue-for-peak, orange-for-trough color convention used throughout the rest of the pipeline; an example of this panel is shown in Figure 6. These per-channel panels are collected into one subsystem-wide report, ranked by the aforementioned `z_over_floor` metric, so that channels with coupling clearly above the subsystem's own noise floor rise to the top, making it easier to scan for interesting patterns. Every channel is also scored on three further measures within its own subsystem, so channels remain comparable across subsystems: recurrence, how often a channel lands in that day's own top ranks; cluster area and power, how broad and connected its excess coupling region is; and stability, how little its peak-coupling frequency pair moves around from day to day. These three values are averaged into one `composite_score`, used to flag the most promising channels for closer inspection.

4. Results

The methods described above were used to run a full, subsystem-by-subsystem coherence survey and bicoherence survey across the nine oscillation days selected in Stage 0. This section reports what each of those two surveys actually found, and, in the case of coherence, why the search ultimately did not succeed.

4.1. Segmented coherence results

The segmented coherence survey was run across 2,373 auxiliary channels examined for each day. Figure 4 shows an example of the ranked-channel table this survey produces for a single subsystem and day. On any single day, a handful of channels reliably showed elevated Δ scores worth a closer look, and taken in isolation those scores looked encouraging. But when we ran the cross-day recurrence check described in Section 3.3, we found that no single channel or subsystem emerged as a clear and consistent linear signal. A channel with a strikingly high Δ score on one day would, on the next day examined, show a score with no statistical significance. None of the promising channels survived being checked against a second, independent day, motivating the switch to bicoherence.

4.2. Segmented bicoherence results

Switching to the bicoherence analysis, which searches for quadratic phase coupling, we again found several channels with clear, structured patterns, such as the one shown in Figure 6. The plots showed stripes, spots, and other connected features concentrated in the 30–50 Hz in-band region. One interesting single-day candidate identified so far is `SUS-ETMX_L3_OPLEV_YAW_OUT_DQ`: an optical lever on the end-X test mass suspension,

#	Channel	Δ score (peak−trough)	Score (peak)	Score (trough)	Max coh (peak)	Max coh (trough)	Consistency peak	Consistency trough	Plot
1	PEM-EX_MAG_EBAY_SUSRACK_X_DQ	12.76	77.00	64.24	0.738	0.699	1.00 (9/9)	1.00 (8/8)	view
2	PEM-EX_MAG_EBAY_SEIRACK_Z_DQ	10.09	75.51	65.42	0.730	0.694	1.00 (9/9)	1.00 (8/8)	view
3	PEM-EX_MAG_EBAY_SEIRACK_X_DQ	8.50	77.22	68.72	0.736	0.693	1.00 (9/9)	1.00 (8/8)	view
4	PEM-EX_MAG_EBAY_SUSRACK_Z_DQ	8.50	77.12	68.63	0.734	0.692	1.00 (9/9)	1.00 (8/8)	view
5	PEM-EX_ADC_4_11_OUT_DQ	5.88	77.35	71.47	0.742	0.704	1.00 (9/9)	1.00 (8/8)	view
6	PEM-CS_MAG_LVEA_VERTEX_Y_DQ	5.59	78.08	72.49	0.729	0.685	1.00 (9/9)	1.00 (8/8)	view
7	PEM-MX_MIC_VEA_MIDX_DQ	5.47	67.12	61.65	0.241	0.212	0.00 (0/9)	0.25 (2/8)	view
8	PEM-CS_MAG_LVEA_OUTPUTOPTICS_Z_DQ	5.25	77.21	71.95	0.722	0.676	1.00 (9/9)	1.00 (8/8)	view
9	PEM-CS_MAG_LVEA_OUTPUTOPTICS_X_DQ	5.02	78.84	73.82	0.721	0.691	1.00 (9/9)	1.00 (8/8)	view
10	PEM-EX_ADC_4_13_OUT_DQ	4.50	79.66	75.16	0.733	0.698	1.00 (9/9)	1.00 (8/8)	view
11	PEM-CS_MAG_LVEA_INPUTOPTICS_Y_DQ	4.08	77.85	73.77	0.723	0.677	1.00 (9/9)	1.00 (8/8)	view
12	PEM-CS_MAG_LVEA_VERTEX_X_DQ	3.41	77.79	74.38	0.738	0.692	1.00 (9/9)	1.00 (8/8)	view
13	PEM-MX_ACC_BEAMTUBE_VEA_Z_DQ	2.97	69.77	66.80	0.518	0.485	0.11 (1/9)	0.25 (2/8)	view
14	PEM-EX_MAG_EBAY_SEIRACK_Y_DQ	2.79	76.79	74.00	0.733	0.695	1.00 (9/9)	1.00 (8/8)	view
15	PEM-CS_MAG_LVEA_INPUTOPTICS_X_DQ	2.72	77.81	75.09	0.726	0.689	1.00 (9/9)	1.00 (8/8)	view
16	PEM-EY_AMON_BEAMTUBE_4000Y_DQ	2.50	80.35	77.85	0.756	0.712	1.00 (9/9)	1.00 (8/8)	view
17	PEM-MX_ACC_BEAMTUBE_VEA_X_DQ	2.32	69.97	67.65	0.539	0.504	0.00 (0/9)	0.25 (2/8)	view
18	PEM-CS_MAG_LVEA_OUTPUTOPTICS_Y_DQ	2.27	79.15	76.87	0.709	0.678	1.00 (9/9)	1.00 (8/8)	view
19	PEM-CS_ADC_4_23_OUT_DQ	2.26	75.38	73.12	0.321	0.293	0.78 (7/9)	0.75 (6/8)	view
20	PEM-CS_AMON_BEAMTUBE_40X_DQ	1.83	77.95	76.12	0.733	0.685	0.89 (8/9)	1.00 (8/8)	view
21	PEM-MX_ACC_BEAMTUBE_1900X_X_DQ	1.77	70.59	68.82	0.345	0.306	0.11 (1/9)	1.00 (8/8)	view
22	PEM-EX_MAINSMON_EBAY_1_DQ	1.76	77.66	75.90	0.732	0.689	1.00 (9/9)	1.00 (8/8)	view
23	PEM-CS_ACC_LVEAFLOOR_YCRYO_Z_DQ	1.69	80.50	78.81	0.733	0.691	1.00 (9/9)	1.00 (8/8)	view
24	PEM-EY_MAINSMON_EBAY_3_DQ	1.66	77.72	76.06	0.732	0.682	1.00 (9/9)	1.00 (8/8)	view
25	PEM-EX_MAINSMON_EBAY_2_DQ	1.61	76.96	75.35	0.732	0.697	1.00 (9/9)	1.00 (8/8)	view
26	PEM-CS_MAINSMON_EBAY_3_DQ	1.57	74.69	73.12	0.733	0.684	0.89 (8/9)	1.00 (8/8)	view
27	PEM-CS_AMON_BEAMTUBE_40Y_DQ	1.33	76.44	75.11	0.607	0.374	0.89 (8/9)	1.00 (8/8)	view
28	PEM-MX_AMON_BEAMTUBE_2000X_DQ	1.24	77.28	76.04	0.580	0.485	0.67 (6/9)	1.00 (8/8)	view
29	PEM-EX_MAINSMON_EBAY_3_DQ	1.22	77.07	75.85	0.732	0.689	1.00 (9/9)	1.00 (8/8)	view
30	PEM-CS_MAG_LVEA_INPUTOPTICS_Z_DQ	1.21	78.84	77.63	0.730	0.690	1.00 (9/9)	1.00 (8/8)	view

Figure 4: Example ranked-channel table from a segmented coherence HTML report (PEM subsystem, September 19, 2024). Channels are sorted and color-coded by Δ score, with the peak score, trough score, maximum peak/trough coherence, and peak/trough consistency ratios shown alongside. Clicking a channel name in the live report opens that channel's own four-panel summary plot.

measuring yaw, side-to-side rotation about the vertical axis. On September 19, 2024, this channel showed a connected in-band bicoherence excess that appeared across its individual peak-epoch segments.

Following up on this candidate channel using the same recurrence check described in Section 3.4, we checked the same channel against two additional oscillation days spread elsewhere across the observing run: November 5, 2023 and October 21, 2025. Unlike the September 19, 2024 result, neither of these two additional days reproduced a comparably clear or consistently structured in-band excess for this channel (Figure 7). While the September 19 result showed an interesting structure, it did not pass the recurrence check: if this channel were truly responsible for the coupling, a similar pattern should appear across multiple days on which the oscillation is present. As with the coherence results, this suggests that bicoherence has not yet detected the coupling mechanism that drives the oscillations, at least not for this channel.

Across the days and subsystems processed so far, in the two surveys covering nine oscillation days and 2,373 auxiliary channels, no single channel or subsystem has emerged as a clear, dominant, and consistently recurring linear or nonlinear coupling. This motivates the future work described below, as we hope to discover the cause of this noise before the next Intermediate Run 1, IR1.

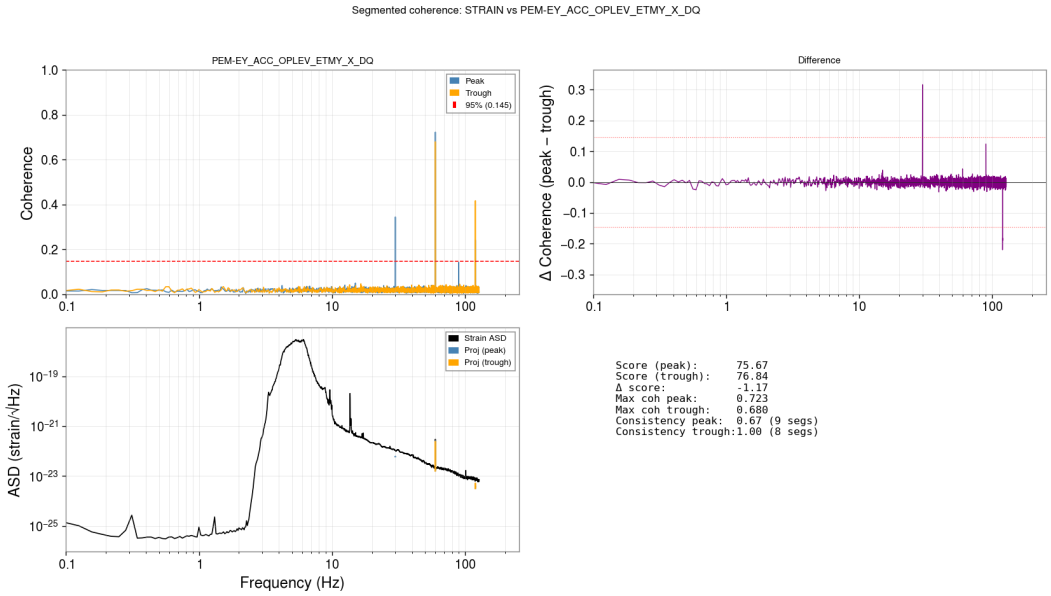


Figure 5: Four-panel segmented coherence summary for the auxiliary channel `PEM-EY_ACC_OPLEV_ETMY_X_DQ` on September 19, 2024. This channel is in the Physical Environment Monitoring (PEM) subsystem: an accelerometer on the optical lever tracking beam alignment in the end-Y arm station. These results are representative of the survey as a whole: no consistent peak-versus-trough coherence signature, motivating the shift to the bicoherence code.

5. Next Steps: Cross Bicoherence

The next step for this project is an extension of the bicoherence method itself. Rather than testing whether a single auxiliary channel's own content at f_1 and f_2 combines to produce excess power in the strain at f_1+f_2, the plan is to test whether the content of two different auxiliary channels, at f_1 and f_2 respectively, jointly produces coupling into the strain at their sum frequency. This is motivated directly by the fact that single-channel cross-bicoherence results have, so far, been inconsistent across days, as described in Section 4.2. For example, if the true physical mechanism driving the oscillation involves an interaction between two separate parts of the interferometer, such as motion in one subsystem modulating a different signal picked up by an entirely different subsystem, rather than either signal being coupled into the strain on its own, then no single-channel search, whether linear or quadratic, could ever find it. This motivates the shift to a cross-bicoherence framework.

Fortunately, this extension does not require writing a new script from scratch: the underlying function used throughout this project already accepts three separate input signals, a, b, and y, and was only ever restricted to the single-channel case in this project's current usage because a and b were both set to the same channel's own signal. The real challenge in this extension is not the code, but rather the scale: testing every possible pair among roughly two thousand auxiliary channels is far larger than testing

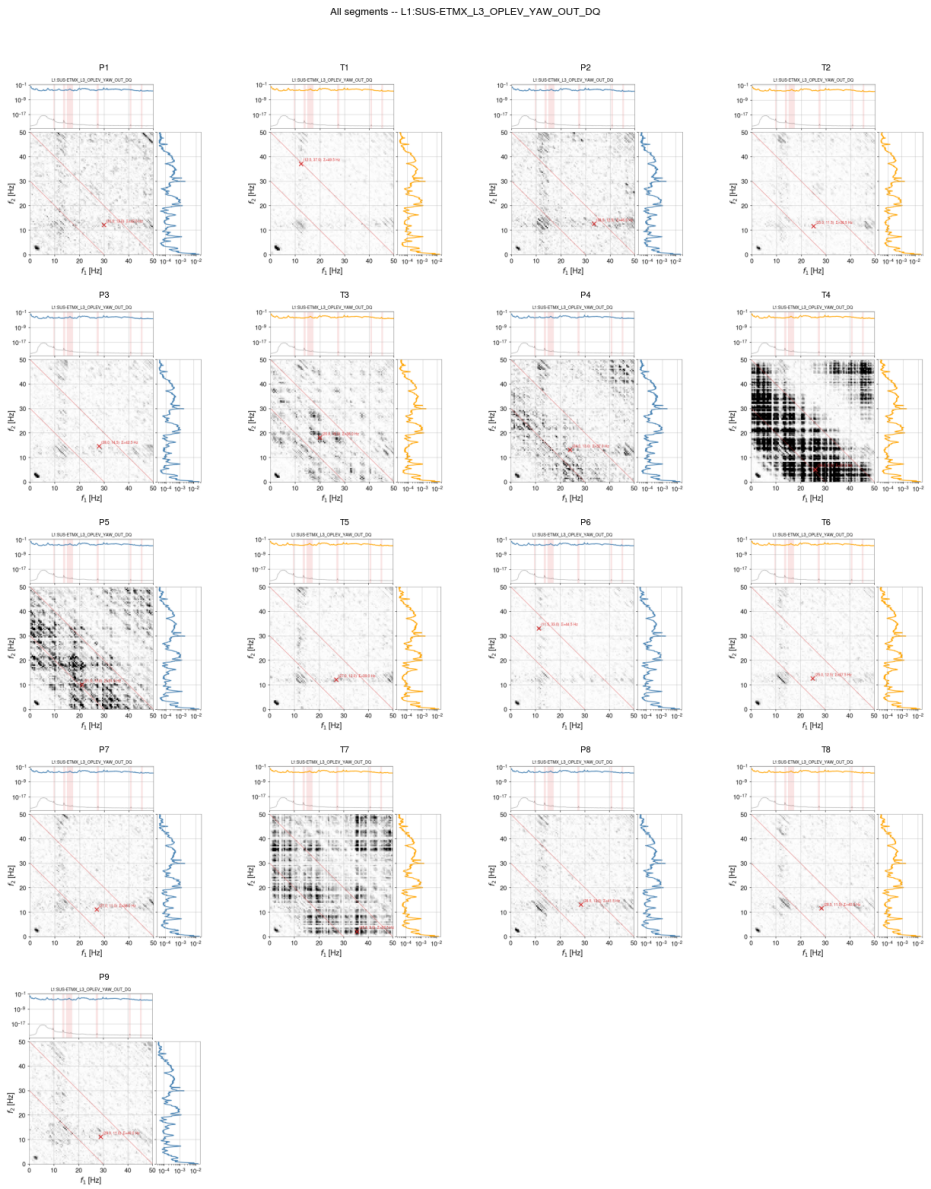


Figure 6: Peak and trough bicoherence heatmaps for the auxiliary channel `SUS-ETMX_L3_OPLEV_YAW_OUT_DQ` on September 19, 2024. This channel is an optical lever on the end-X test mass suspension (SUS), measuring yaw, side-to-side rotation about the vertical axis. A clear pattern appears within the 30–50 Hz in-band region of interest and is consistent across the individual epoch segments that make up the peak-epoch average. However, the pattern does not exhibit the expected "on-off" behavior: it does not consistently appear during the peak epochs and disappear during the trough epochs of the oscillation, suggesting that the observed coupling is likely not associated with the oscillation.

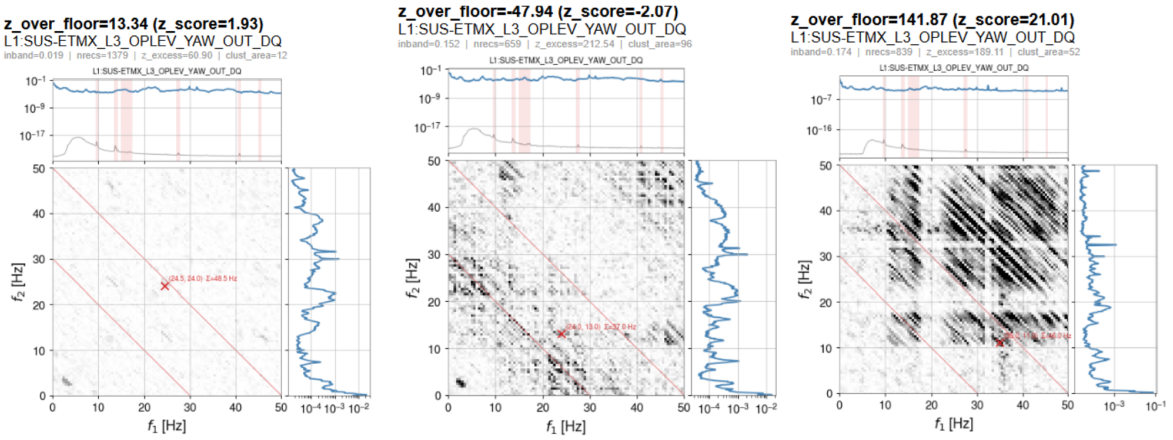


Figure 7: Whole-period bicoherence result for `SUS-ETMX_L3_OPLEV_YAW_OUT_DQ` on three separate oscillation days spread across O4: September 19, 2024; October 21, 2025; and November 5, 2023. The clear structure seen on September 19 does not recur with comparable strength on the other two days.

every channel individually, on the order of two million distinct pairs per day. The planned approach is instead to use a smaller channel list of channels more likely to be causing this behavior, rather than attempting an exhaustive sweep.

Acknowledgements

Thank you to Jane Glanzer for all of her help throughout this project and for being an amazing mentor. Thank you to the NSF and LIGO Lab for funding this summer, and to Caltech for hosting this program. Thank you to everyone who helped run and support this program, and to all of my fellow SURFs for making this such a great summer.

References

[1] Einstein A 1916 *Annalen Phys.* **49** 769–822

[2] Abbott B P *et al.* (LIGO Scientific Collaboration and Virgo Collaboration) 2016 *Phys. Rev. Lett.* **116**(6) 061102 URL https://link.aps.org/doi/10.1103/PhysRevLett.116.061102

[3] The LIGO Scientific Collaboration, the Virgo Collaboration and the KAGRA Collaboration 2026 *arXiv e-prints* arXiv:2605.27225 (*Preprint* 2605.27225)

[4] Sun L *et al.* 2020 *Classical and Quantum Gravity* **37** 225008 (*Preprint* 2005.02531)

[5] LIGO Lab Caltech: MIT 2016 LIGO: The First Observation of Gravitational Waves Caltech LIGO Video Gallery educational video on the LIGO video gallery page URL https://www.ligo.caltech.edu/video/ligo20160215v1

[6] Glanzer J, Soni S, Spoon J, Effler A and González G 2023 *Classical and Quantum Gravity* **40** 195015 (*Preprint* 2304.07477)

[7] Soni S *et al.* 2021 *Classical and Quantum Gravity* **38** 025016 (*Preprint* 2007.14876)

[8] Soni S, Glanzer J, Effler A, Frolov V, González G, Pele A and Schofield R 2024 *Classical and Quantum Gravity* **41** 135015 (*Preprint* 2311.05730)

[9] Soni S *et al.* (LIGO) 2025 *Class. Quant. Grav.* **42** 085016 (*Preprint* 2409.02831)

[10] Davis D and Capote E 2025 **111** 122005 (*Preprint* 2502.07253)

[11] Glanzer J 2024 aLIGO LLO Logbook 70426

[12] Glanzer J 2024 *LSU Doctoral Dissertations* URL https://repository.lsu.edu/gradschool_dissertations/6455

[13] LIGO-Virgo-KAGRA Collaboration 2025 LIGO-Virgo-KAGRA Collaboration successfully wraps up its fourth Observing Run LIGO Scientific Collaboration o4 began May 24, 2023 and ended November 18, 2025 URL https://ligo.org/ligo-virgo-kagra-collaboration-successfully-wraps-up-its-fourth-observing-run/

[14] Carter G C, Knapp C H and Nuttall A H 1973 *IEEE Transactions on Audio and Electroacoustics* **AU-21** 337–344

[15] Poloskei P Z, Papp G, Por G, Horvath L and Pokol G I 2018 *arXiv e-prints* arXiv:1811.02973 (*Preprint* 1811.02973)

[16] Hagihira S, Takashina M, Mori T, Mashimo T and Yoshiya I 2001 *Anesthesia & Analgesia* **93** 966–970