Mitigation of the instrumental noise transient in gravitational-wave data surrounding GW170817

Chris Pankow, Katerina Chatziioannou, Eve A. Chase, Tyson B. Littenberg, Matthew Evans,
Jessica McIver, Neil J. Cornish, Carl-Johan Haster, Jonah Kanner, Vivien Raymond,
Salvatore Vitale, and Aaron Zimmerman

Mitigation of the instrumental noise transient in gravitational-wave data surrounding GW170817

Chris Pankow,[1] Katerina Chatziioannou,[2] Eve A. Chase,[1] Tyson B. Littenberg,[3] Matthew Evans,[4] Jessica McIver,[5] Neil J. Cornish,[6] Carl-Johan Haster,[2] Jonah Kanner,[5] Vivien Raymond,[7] Salvatore Vitale,[8] and Aaron Zimmerman[2]

[1]*Center for Interdisciplinary Research and Exploration in Astrophysics (CIERA) and Department of Physics and Astronomy, Northwestern University, 2145 Sheridan Road, Evanston, IL 60208, USA*
[2]*Canadian Institute for Theoretical Astrophysics, 60 St. George Street, Toronto, Ontario, M5S 3H8, Canada*
[3]*NASA Marshall Space Flight Center, Huntsville, Alabama 35811, USA*
[4]*LIGO, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, USA*
[5]*LIGO, California Institute of Technology, Pasadena, CA 91125, USA*
[6]*eXtreme Gravity Institute, Department of Physics, Montana State University, Bozeman, Montana 59717, USA*
[7]*Cardiff University, Cardiff CF24 3AA, United Kingdom*
[8]*LIGO Laboratory and Kavli Institute for Astrophysics and Space Research, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, USA*

In the coming years gravitational-wave detectors will undergo a series of improvements, with an increase in their detection rate by about an order of magnitude. Routine detections of gravitational-wave signals promote novel astrophysical and fundamental theory studies, while simultaneously leading to an increase in the number of detections temporally overlapping with instrumentally- or environmentally-induced transients in the detectors (glitches), often of unknown origin. Indeed, this was the case for the very first detection by the LIGO and Virgo detectors of a gravitational-wave signal consistent with a binary neutron star coalescence, GW170817. A loud glitch in the LIGO-Livingston detector, about one second before the merger, hampered coincident detection (which was initially achieved solely with LIGO-Hanford data). Moreover, accurate source characterization depends on specific assumptions about the behavior of the detector noise that are rendered invalid by the presence of glitches. In this paper, we present the various techniques employed for the initial mitigation of the glitch to perform source characterization of GW170817 and study advantages and disadvantages of each mitigation method. We show that, despite the presence of instrumental noise transients louder than the one affecting GW170817, we are still able to produce unbiased measurements of the intrinsic parameters from simulated injections with properties similar to GW170817.

I. INTRODUCTION

The discovery of GW170817 [1] was a watershed moment in astrophysics. It was the first detection of a binary neutron star (BNS) through gravitational waves (GWs), and its association with GRB170817A [2] spurred a world-wide effort which revealed emission throughout the electromagnetic spectrum [3]. The signal is the loudest GW event detected so far, exhibiting a signal-to-noise ratio (SNR) of 18.8 in the LIGO-Hanford, 26.4 in the LIGO-Livingston [4], and 2 in the Virgo [5] detectors.

Initially identified by low-latency compact binary searches [6, 7], coincidence and rapid localization [8] were hindered because of a non-astrophysical transient present in the LIGO-Livingston detector about one second before the estimated merger time of the signal [1]. This transient signal (colloquially referred to as a *glitch* [9, 10]) was a high-amplitude, short-duration excursion of unknown origin in the control loop used to sense the differential motion of the interferometer arms, which is calibrated to produce GW strain data [4]. This excursion caused an overflow in the digital-to-analog converter of the optic drive signal and prevented the searches from registering an event trigger for GW170817 in the LIGO-Livingston detector in low latency. The cause of the excursion is un-

known, but such *overflow* glitches occur independently at a rate of roughly once every few hours in both LIGO detectors [11].

After the signal was visually identified in both LIGO detectors, it was clear that the noise transient in LIGO-Livingston would need to be mitigated in order to provide accurate sky localization maps and estimate the parameters of GW170817.

Three types of glitch removal were employed. The first method excised 0.2 seconds of time series data corrupted by the noise transient using the PyCBC *gating* algorithm [12, 13]. The second method is a variation on the first which allows for better fidelity in the process. The gating methods were developed in the context of detection, since glitches can corrupt the output of matched filtering searches [6, 12]. They are applied in an automated processes which identify statistically significant excursions in the data stream and excise them. Cognizant that any electromagnetic counterpart was likely to fade on the order of hours to days [14], the first method of data excision was used to produce and promptly disseminate an accurate sky map for GW170817 [15]. This resulted in the successful discovery of a counterpart electromagnetic signal [16]. However, while such data excision techniques are speedy to apply, portions of the underlying signal are also removed in the process. A third, more so-

phisticated, technique used the `BayesWave` algorithm [17] to model the glitch and subtract it from the data, preserving more of the signal. The glitch-subtracted data, produced days after the initial detection, was used for parameter estimation [1] and more detailed studies of the signal properties [18, 19].

The focus of this work is to evaluate the efficacy and fidelity of these techniques — *gating* and `BayesWave` — in the context of characterization of the intrinsic properties of the signal. A detailed study of the effect of gating and `BayesWave` cleaning on sky position reconstruction is deferred to a later study. In our analysis we use the `LALInference` package [20], a standard parameter estimation analysis for GW signals from compact binaries. This analysis models the signal with existing theoretical waveform templates and assumes that the noise is *stationary*, namely that its properties do not change appreciably on the timescale of the signal, and *Gaussian*, namely that the noise follows a colored Gaussian distribution. The presence of the noise transient in the data breaks both assumptions about noise behavior.

For a true GW source, such as GW170817, no control case is available for a glitch mitigation study; the true waveform and intrinsic source parameters are not known *a priori*. We circumvent this problem with a set of synthetic event injection studies, using data stretches which exhibit instrumental transients similar to the one afflicting GW170817. We inject synthetic GW170817-like signals concurrently with those noise transients and then apply the transient removal methods. We then recover the signal parameters, and compare the known properties of the injected synthetic signal to the posterior distributions obtained by analyzing the data with and without the subtraction procedure applied. As a reference, we also perform the same parameter estimation analysis with a zero-noise realization. This provides an ideal case for comparison with the real-noise analyses.

We find that modeling and subtracting the glitch using the `BayesWave` algorithm leads to robust and unbiased estimation of the parameters of the signal, including the masses, spins, and tidal deformabilities of the binary components. The latter is particularly relevant in the case of GW170817. The tidal deformability quantifies the deformation of a body in the presence of an external field [21] and depends sensitively on the equation of state of matter at supranuclear densities [22, 23]. Since it affects the binary evolution mostly during the latest stages [24] of its evolution, its main effect on the waveform temporally coincides with the glitch in the data (which, as mentioned above, happened roughly 1 second before the estimated merger time). Despite this, we show that the measurement of the tidal deformability presented in Refs. [1, 18, 19] are only minimally affected by the glitch once it has been successfully modeled and removed.

II. MITIGATION TECHNIQUES

In the presence of a non-astrophysical transient and a GW signal, the measured strain time series can be expressed [25, 26] as

$$s(t) = n(t) + h(t, \vec{\lambda}) + g(t, \Delta t) \,. \qquad (1)$$

In the above equation $n(t)$ is a stream of stationary and Gaussian noise, $h(t)$ is the GW signal that depends on the binary parameters $\vec{\lambda}$ [27], and $g(t)$ is the non-Gaussian, instrumentally-induced transient of duration Δt. The spectral profile of the noise is characterized by the one-sided power spectral density (PSD) denoted $S(f)$. The detection and parameter estimation algorithms make explicit assumptions about the nature of the data — namely that the Fourier transform of $n(t)$, $\tilde{n}(f)$ is related to the PSD by $2\langle \tilde{n}(f)\tilde{n}^*(f') \rangle = S(f)\delta(f - f')$. In other words, the noise is stationary and Gaussian distributed, given its ensemble average in the Fourier domain is proportional to the spectral density [28].

In the context of GW parameter estimation both of these assumptions determine the form of the log-likelihood function for the data, given a model for the GW signal $h(\vec{\lambda}')$ [20], defined as:

$$\log L(s|\vec{\lambda}') \propto \langle \tilde{s} - \tilde{h}(\vec{\lambda}') | \tilde{s} - \tilde{h}(\vec{\lambda}') \rangle. \qquad (2)$$

where the quantities with a tilde are the Fourier-domain transformed quantities from the time-domain equivalents in Eq. (1). The brackets denote a noise-weighted inner-product defined as $\langle a|b \rangle \equiv 4\Re \int_0^\infty \tilde{a}^* \tilde{b} / S(f) df$.

Eq. (2) is calculated from the residual after subtracting the expected signal $\tilde{h}(\vec{\lambda}')$ from the data $\tilde{s}$. In the presence of a glitch, the residual consists of both Gaussian noise and the instrumentally-induced transient. Moreover, the parameters of the signal can be severely biased as the analysis attempts to return residuals that are consistent with our assumption of Gaussian noise. The intrinsic amplitude of the glitch relative to the Gaussian noise sets the scale of the mismatched residual (e.g. with which h can g be matched), and if the glitch is comparable in SNR relative to the signal, it can overwhelm the likelihood calculated for the signal alone [29]. For a broad overview of the impact on the estimation of the signals properties of compact object coalescences, including the effects of longer-duration glitches, see [30]. We next describe strategies to mitigate glitch-induced biases by removing the glitch from the data, $s(t)$, with minimal impact of the recovered signal.

A. Frequency-independent subtraction

The initial solution to mitigating the noise transient near GW170817 was to smoothly taper the time series data to zero while the glitch was present [12, 13], see Fig. 2 of Ref. [1],

$$W(t) = [1 - w(t, t_0, t_w, t_{\text{taper}})] \,. \qquad (3)$$

In the above a windowing function w — typically a Tukey window — is applied around the center of the noise transient time t_0. The shape of the window is set by t_w, the half-duration of the data to zero out, and t_{taper}, the duration of the Tukey tapering on each side of the gated data. For GW170817, $t_0 = 1187008881.389$, $t_w = 0.1s$, and $t_{\text{taper}} = 0.5s$. The duration of the data affected by the gating window $\Delta t = 2t_w + 2t_{\text{taper}}$ [1]. Clearly, this also removes the signal $h(t)$ during Δt from the data series. Therefore, while this method is effective in removing the high amplitude power in the glitch, it might induce a different bias because the signal waveform template would also need to be windowed in the same way to match optimally. Section III A shows that, for a zeroed time window of $2t_w = 0.25$ seconds or less, this bias is within the tolerances set by the uncertainties in the posterior distributions for a preliminary analysis, and so the trade-off is acceptable.

B. Time-frequency area subtraction

A similar, but *frequency-dependent* solution was developed to further reduce damaging the fidelity of the noise-transient-subtracted data $s'(t)$

$$s'(t) = s(t) - g_{\text{TFA}}(t, t_0, \Delta t, \Delta f), \qquad (4)$$

where $g_{\text{TFA}}(t, t_0, \Delta t, \Delta f)$ is the glitch estimate produced by the Time-Frequency Area (TFA) algorithm.

The TFA algorithm starts by band-passing the data around the noise transient in the frequency band effected by the glitch (e.g., 50 to 800 Hz). This is done in the time-domain after de-trending and windowing the time-series.

The Fast Fourier Transform (FFT) of the band-passed data is then computed, and the amplitudes of frequencies corresponding to known strong features (e.g., calibration lines) are set to zero. If desired, additional band-passing can also be done in the frequency domain by setting the amplitude of out of band frequencies to zero as well. For the results presented here, frequencies near 300 Hz and 500 Hz were set to zero to remove strong narrow lines due to injected calibration lines and mechanical resonances [31], and frequencies above 700 Hz were zeroed to improve the band-pass filtering.

Finally, the band-passed and line-removed frequency-domain data are converted back to the time domain with an inverse FFT operation, and the glitch estimate $g_{\text{TFA}}(t, t_0, \Delta t, \Delta f)$ is produced by windowing this time-series around the time of the glitch. The Tukey window function, which is equal to 1 in the central region, is used throughout the TFA algorithm to ensure that the amplitude of the de-glitched time series is not modified by the windowing operation.

C. Noise-transient fitting and removal

Apart from generic considerations about the duration and bandwidth of the noise transient, the methods described in Secs. II A and II B are mostly agnostic to the morphology of the transient itself. They are expedient to apply, and significantly reduce biases in the recovered signal parameters. However, the signal is still modified in some way because of the excision applied. A more sophisticated treatment can be applied if one considers the coherent and incoherent decompositions of the data from the two instruments. While a true GW signal is *coherent* across the detector network, any instrumentally-induced transient is *incoherent*, assuming that the noise in the various detectors is uncorrelated. The morphology-independent algorithm BayesWave [17] exploits this distinction to reconstruct any coherent and incoherent power in the detectors through a continuous wavelet basis. BayesWave uses Morlet-Gabor wavelets, with morphologies

$$\Psi(t; A, f_0, \tau, t_0, \phi_0) = Ae^{-(t-t_0)^2/\tau^2} \cos\left[2\pi f_0(t-t_0)+\phi_0\right] \quad (5)$$

to simultaneously model both the coherent signal and the incoherent noise transient in Eq. (1). Each wavelet depends on five parameters: its amplitude A, quality factor $Q \equiv 2\pi f_0 \tau$, central frequency f_0, central time t_0, and phase ϕ_0.

BayesWave harnesses the inherent efficiency of Bayesian inference [17]. When a signal is coherent across the detector network it can be fully described in all detectors by the same set of Morlet-Gabor wavelets plus four extrinsic parameters characterizing its sky location, polarization, and the ratio of the two independent GW polarization amplitudes (ellipticity). An incoherent signal, on the other hand, can only be reconstructed with independent sets of wavelets in each detector. Since the incoherent model (termed the 'glitch model' in BayesWave literature) has more parameters ($5N_dN$, where N_d is the number of detectors and N the number of wavelets) than the coherent model ($5N + 4$, usually referred to as the 'signal model'), the latter will be preferred assuming both models fit the data equally well [32]. Any coherent signal can be separated from the incoherent power, as it can be modeled with fewer parameters [33].

In practice and for the signal and glitch near GW170817, this process is aided by the fact that a BNS inspiral and an overflow glitch are reconstructed by wavelets of largely different quality factors, Q. A BNS signal is characterized by a long duration (of order seconds to minutes), and is therefore best modeled by wavelets with a large quality factor. On the contrary, the glitch has a short duration and a more compact time-frequency signature, which is best recovered with wavelets with a small quality factor.

Once the coherent and the incoherent part of the signal have been simultaneously reconstructed, the incoherent portion of the data is subtracted from the individual in-

strument data streams. The instrumental transient — present only in the LIGO-Livingston detector — is removed while the coherent GW signal is only minimally affected, giving

$$s'(t) = s(t) - g_I(t), \qquad (6)$$

where $g_I(t)$ is BayesWave's reconstruction of the incoherent part of the data. In effect, BayesWave fits for the shape of $g(t)$ while leaving $h(t)$ and $n(t)$ unaffected in Eq. (1). As a consequence, the assumptions of stationary and Gaussian residuals once the GW signal has been subtracted are explicitly restored and normal parameter estimation techniques are applicable. We show in Sec. III B that this procedure leads to negligible biases in the parameters of the signal.

III. SIGNAL CHARACTERIZATION

The *overflow* glitch described in Sec. I preceded the merger time of GW170817 in the LIGO-Livingston instrument by about 1.1 seconds. As a result, the signal was originally only identified in the LIGO-Hanford detector. In addition to automated monitors [34, 35], manual checks and visual inspections of the data revealed both the glitch in LIGO-Livingston and the coincident nature of the signal. The glitch intersected the time-frequency track of the merger signal, immediately revealing the need to mitigate the impact of the glitch. In order to produce a timely alert of the event to partner astronomers, the frequency-independent data excision method of Sec. II A was used to generate a rapid sky position posterior. On longer time scales, the methods in II B and II C were applied, and ultimately, the final results presented in [1, 18, 19] used the data set produced with the method in II C.

In this section, we compare the performance and fidelity of methods II A, II B, and II C by injecting synthetic signals with parameters comparable to GW170817 in data from the two LIGO detectors. In Sec. III A we compare the effect of glitch-mitigation methods of Secs. II A, II B, and II C on a simulated signal in the absence of a glitch. Since the data are completely excised, the absence of a glitch is irrelevant to the investigation of the two gating methods. The use of BayesWave glitch-removal method on data where there is no glitch present should meanwhile change the results very little, which we confirm. We then explore the performance of BayesWave glitch removal in the analysis of signals injected on top of actual instrumental glitches in III B. For all the analyses of this section the PSD is calculated from on-source data with the technique presented in [36]. We also marginalize over the calibration uncertainty of the detector as explained in [37]. For simplicity, we assume priors on the amplitude (phase) marginalization spline points of 5% (3°).

A. Comparison of mitigation methods

In this subsection, we study the effect of the three glitch-mitigation methods described in Sec. II on source characterization. We select a time less than a half-hour after GW170817 in a contiguous set of data and hence preserve the state of the instruments as closely as possible and add a simulated GW170817-like signal to the data, using the IMRPhenomPv2 waveform family, see Sec. III B. For this specific time, no glitch was present within the duration of the simulated signal. The methods from II A, II B, and II C were applied to the simulated signal in exactly the same manner and temporal offset relative to GW170817.

In particular, understanding the consequences of the two glitch excision methods (Secs. II A and II B) is important since the methods are straightforward and quick to apply, and will likely continue to be used for rapid response to interesting GW events, especially as the detection rate increases and the probability of signals overlapping with glitches rises. The initial data excision is expected to impact the extracted parameters differently depending on the intersection of the excision and time-frequency track of the signal. In the case of GW170817, the intersection point of the track with the excision is relatively close to the merger time; the fixed time interval in between the glitch peak time and the signal implies the intersection point in the time-frequency plane is near ~ 100 Hz. As a result, the sweep of the signal through the most sensitive part of the instrument bandwidth — where many physical effects are most well measured — could be affected. While the mass parameters are typically well measured from lower frequencies, the effects of spin and tidal parameters on the waveform phasing begin to become measurable in this region. Relative to the length-in-band of GW170817 (~ 100 s from 25 Hz to merger), the glitch duration is small (the overflow itself is less than 5 ms), as is the excision duration (~ 0.5 s). Ideally, handling this excision would require a modification of the signal template to similarly excise the portion of the signal which is removed with the data. While not conceptually complex, adding this capability to the existing low latency GW searches would increase computational costs.

For this example, no significant glitching is present near the merger time of the injection. The methods in Section II A and II B each remove a fixed time-frequency volume from the data, and so the impact on the SNR of the signal should be the same with or without a glitch present. For the BayesWave method in Section IIC, the absence of a glitch means that only a small amount of power is expected to be removed in this case. The comparison between the different glitch-mitigation methods is displayed in Fig. 1. We present posterior distributions for various parameters of a signal injected, and the various mitigation techniques presented in Sec. II applied. The SNR computed from data s and template $h(\vec{\lambda}')$ is proportional to $\langle \tilde{s}|\tilde{h}(\vec{\lambda}')\rangle$, and as such, it is clear that re-
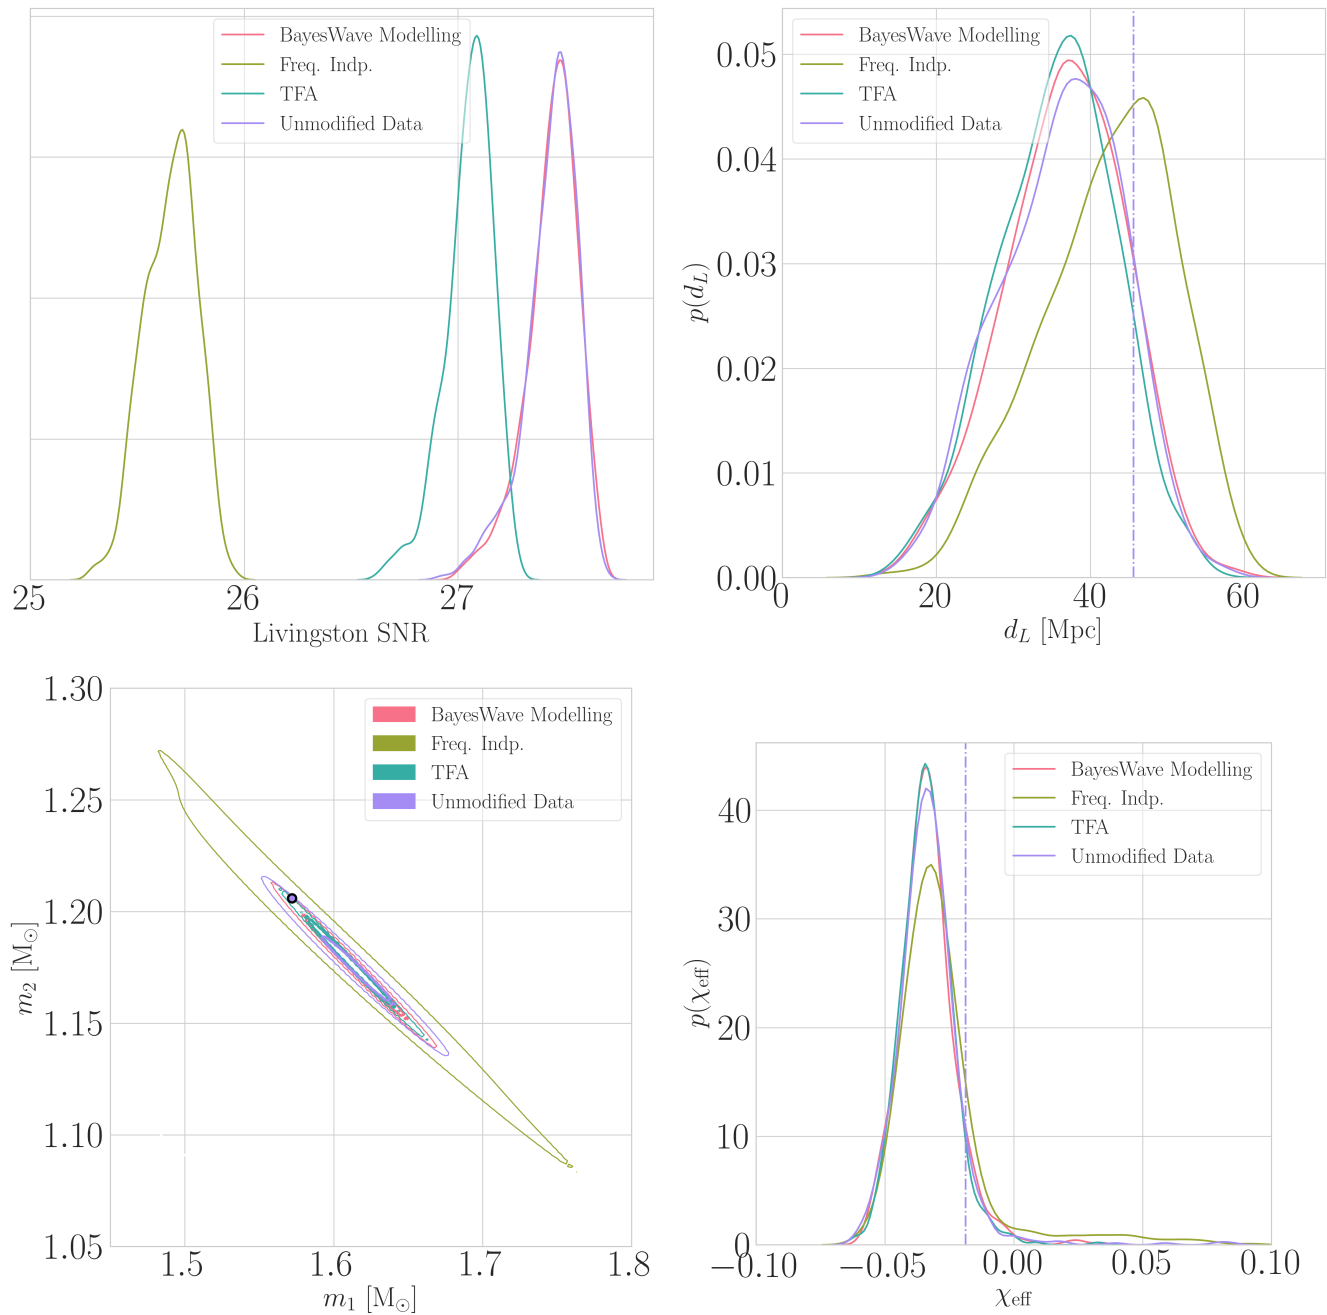
FIG. 1: Comparison of mitigation methods. We show posterior distributions for the SNR in LIGO-Livingston (top left), the signal distance (top right), the component masses (bottom left), and the effective spin (bottom right) for a simulated signal injected into data adjacent in time to GW170817, but without a strong glitch near the end time. The posteriors are computed after applying the frequency-independent excision method of Sec. II A (yellow, "Freq. Indp."), the frequency-dependent excision method of Sec. II B (purple, "TFA"), after modeling and subtracting the glitch as described in Sec. II C (pink, "BayesWave Modelling"), and when analysing the signal without any additional modification (green, "No Gating"). The two excision methods lead to small biases in the estimation of the signal parameter. Moreover, the corresponding uncertainties are increased due to signal power being removed.

moving some part of s will incur a systematic reduction in the SNR distribution over the parameter set. This effect is apparent as the SNR distribution in the LIGO-Livingston data is reduced by about two units ($\sim 7\%$) in the case of frequency-independent data excision and less than one unit ($\sim 2\%$) for the frequency-dependent subtraction method (top left panel). The distance distribution is similarly modified as it is inversely proportional to the SNR, and the event appears to be slightly further away than in reality (top right panel). The distance / inclination degeneracy then leads to a marginal concurrent effect on the joint posterior for these two parameters. Conversely, the SNR and distance distributions after the application of the BayesWave glitch subtraction are very similar to the ones obtained when the signal is injected but otherwise unmodified. This is the expected behavior, since there is no glitch to remove.

Also in Fig. 1, we show the effect of excision methods on the recovered detector-frame component masses (bottom left panel), the effective spin parameter[1] χ_{eff} (bottom right panel). The two dimensional marginalized posterior for the component masses shows that mass posteriors derived with different mitigation methods are broadly consistent with each other and give similar credible intervals along the small axis (chirp mass) of the posterior ellipse. However, the data excision methods lead to mass posteriors that encompass a broadened set of component masses and mass ratios. The widened posterior is indicative not only of the loss of SNR due to the flat portion of the window which entirely excises the data. The tapered portions of the window also induce spurious features in the Fourier domain which can serve to change the likelihood itself, altering the matching phase and amplitude of the data to a putative waveform.

This difference in the mass ratio estimate could affect the measurement of the spin parameters, due to the well-known spin-mass correlation [39]. Indeed, the posteriors for the effective spin parameter, shown in the bottom right plot of Fig. 1, exhibit a minor tail in the posteriors derived from excised data, reflecting the larger mass ratio uncertainty exhibited by the frequency independent and TFA excision data.

Overall, we find that both data excision methods recover a smaller SNR and, in the case of the frequency-independent method, can introduce small parameter biases, which are largely absent when the glitch is modeled and removed. These simple excision methods are an appealing option for quick glitch mitigation, due to their applicability to a large range of glitch morphologies, especially when rapid sky-localization is necessary. However, for follow-up studies the glitch-fitting method seems preferable, since it never introduce biases, and can deal

with cases where the glitch is of comparable length to the signal. In the next section we examine the performances of the glitch-fitting method using simulated signals.

B. Efficacy and veracity of transient removal

In the previous subsection we demonstrated that simple data excision techniques can lead to rapid source characterization with a tolerable degree of parameter biases. However, the recovered signal SNR is decreased, which implies that useful information about the signal has been lost. In this section we study in more detail the technique presented in Sec. II C, which attempts to fit the glitch and remove it while preserving the underlying signal and the Gaussian detector noise.

We produce GW170817-like signals and add them to the data of the LIGO-Livingston detector around four overflow glitches. To produce the signals we use the IMRPhenomPv2 [40–43] and the TaylorF2 [44] (including tidal corrections [45]) waveform models[2] and we place the signals with end times 1.2 seconds after the time assigned to the Livingston overflow glitch we are studying. Table I provides details for these simulated signals. We follow the same procedure from Ref. [1] in applying BayesWave to model and subtract the glitch from the Livingston data following the method described in Sec. II C. Figure 2 shows the time-frequency representation of each glitch together with the simulated signal before and after glitch subtraction. The simulated signal is clearly visible as a characteristic 'chirp' in both cases.

We then perform a coherent Bayesian analysis to estimate the posterior probability distribution of the physical parameters of the simulated signals using the parameter estimation code LALInference [20]. We obtain results by analyzing three instances of the synthetic event: (i) real data including the glitch, (ii) real data where the glitch has been subtracted, and (iii) simulated data with a zero-noise realization [3]. In all cases, the power spectral density used to compute the likelihood in Eq. (2) is fit to the data surrounding and including the event and glitch time [36] to ensure an even comparison between different runs. Of particular astrophysical interest are the physical aspects of the binary measurable through the gravitational-wave signal: its component masses (m_1, m_2) via total mass and mass

---

[1] The effective spin parameter is the mass-weighted sum of the binary components' spins along the orbital angular momentum [38] and one of the best measured spin combinations with GW data.

[2] When our study was started, no readily available waveform family which allows for both spin-precessional and finite-size tidal effects existed, though such a model has recently become available [46].

[3] A zero-noise realization run assumes that the noise is exactly zero at all frequencies, while the power spectral density stays finite. This is a tool often used in the gravitational-wave literature, for example [47–49], to assess parameter estimation in an controlled environment. Statistically, results obtained with a zero-noise realization are equivalent to the average results that would be obtained with a large sample of runs on Gaussian noise [50, 51].
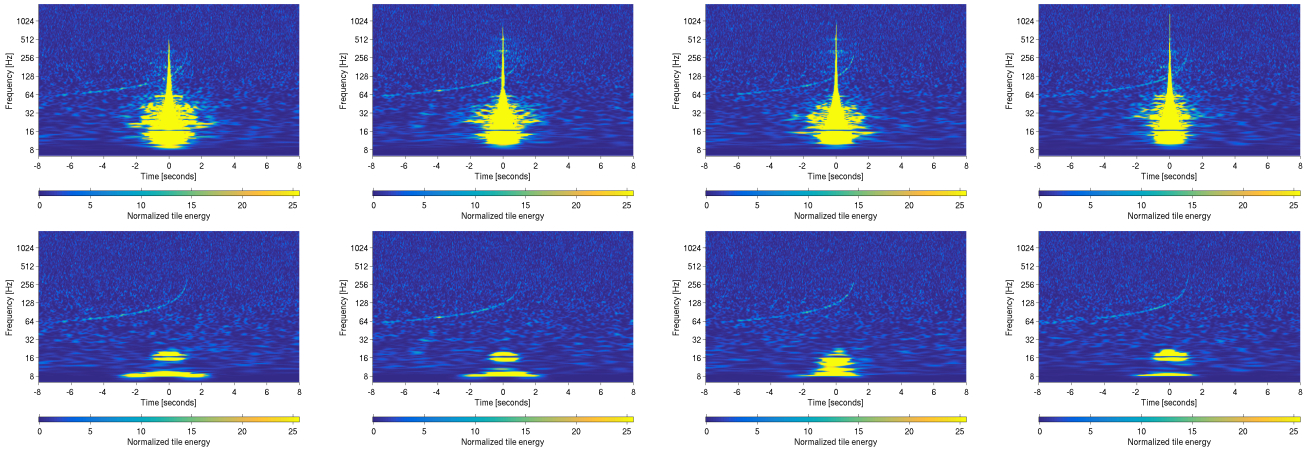
FIG. 2: Normalized signal amplitude (color axis) in the Livingston instrument for four synthetic events added to real interferometric data, as examined in a time-frequency representation. The top row is before the glitch-mitigation is applied, and the bottom row is after we have modeled and removed the glitch with `BayesWave`. The simulated chirping signal is clearly visible in the background of all. The glitch extends in frequency below the lowest frequency used in either the glitch fitting and removal or the parameter estimation procedures, and as such was not completely removed. However, since those frequencies are excluded from the analyses, the result is unaffected by the remaining glitch power.

ratio, mass-weighted spins ($\chi_{\rm eff}$), and tidal deformabilities (Λ_1, Λ_2). In order to check the efficacy of recovery for either, we used the `IMRPhenomPv2` family to synthesize and insert three examples of a precessing signal, and the `TaylorF2` to insert three tidally influenced, but spin-aligned, waveforms into the data. In Fig. 3 we present a selection of two-dimensional posteriors over combinations of the parameters for a precessing signal simulated with `IMRPhenomPv2` (top row) and three signals simulated with `TaylorF2` (second to fourth row). In all cases, we use the same waveform family to simulate the signal added into the data and as template in the parameter estimation algorithm.

In all four events shown in Fig. 3, we verify the recovery of the detector-frame mass parameters, since they dominate the phase evolution of the waveform. The left column of Fig. 3 displays the total mass – chirp mass posterior for all four simulated signals, with the zero-noise-realization and glitch-removed cases shown. The glitch-present recovery is often badly biased, enough so that we do not include it in Fig. 3 and instead refer to Table I to indicate their credible regions. The posteriors for the glitch removed and zero-noise cases are qualitatively very similar. From Fig. 3 it is seen that the mass recovery is consistent with the injected value in all cases, and the zero-noise and glitch-removed recovery encompasses similar values. This validates the premise that the glitch removal method does not bias the lower PN order parameters which influence the waveform.

The second column of Fig. 3 shows the two-dimensional posteriors for the mass ratio and the effective spin. We again find no bias in the parameter estimates due to the glitch removal. The posteriors for the glitch removal and the zero-noise analyses are minimally shifted with respect

to each other, which is consistent with the expected effect of noise realization on signal recovery and not evidence for a bias. Indeed, the specific noise realization of the data induces an additional shift on the posterior estimates of the order of the posterior variance. In addition, the top-right panel shows the posterior for the effective spin $\chi_{\rm eff}$ and the spin parameter $\chi_{\rm p}$ [41]. The latter parameter is an estimate of spin-precession in the waveform. As expected we again see no biases due to glitch removal.

Finally, the tidal parameter estimation also seems unaffected by the glitch removal procedure. The two dimensional marginalized posteriors for the component tidal parameters are presented in the right column, second through fourth row of Fig. 3. The posteriors obtained with our two analyses are both consistent with each other and capture the known value well within their credible intervals. Moreover, the recovered posterior structure is very similar to the actual posterior measured for GW170817 in [1], exhibiting similar boundaries and degeneracies.

Credible intervals for the parameters, as well as their injected values, are quoted in Table I for all three analyses (with glitch mitigation, without glitch mitigation, and with a zero-noise realization). The 90% credible intervals computed from the glitch-mitigated data are consistent with the ones from the zero-noise analysis. This is in stark contrast to the parameter estimates computed if the glitch is included in the data analyzed. The values recovered are well outside of the posterior in the glitch-free examples, and nowhere near the known values. In all four cases, the mass ratio is pushed to extremely high values. The same extreme displacement occurs for the tidal parameters, producing tidal deformability values in the thousands, strongly peaked for Λ_1 and nearly unmea-
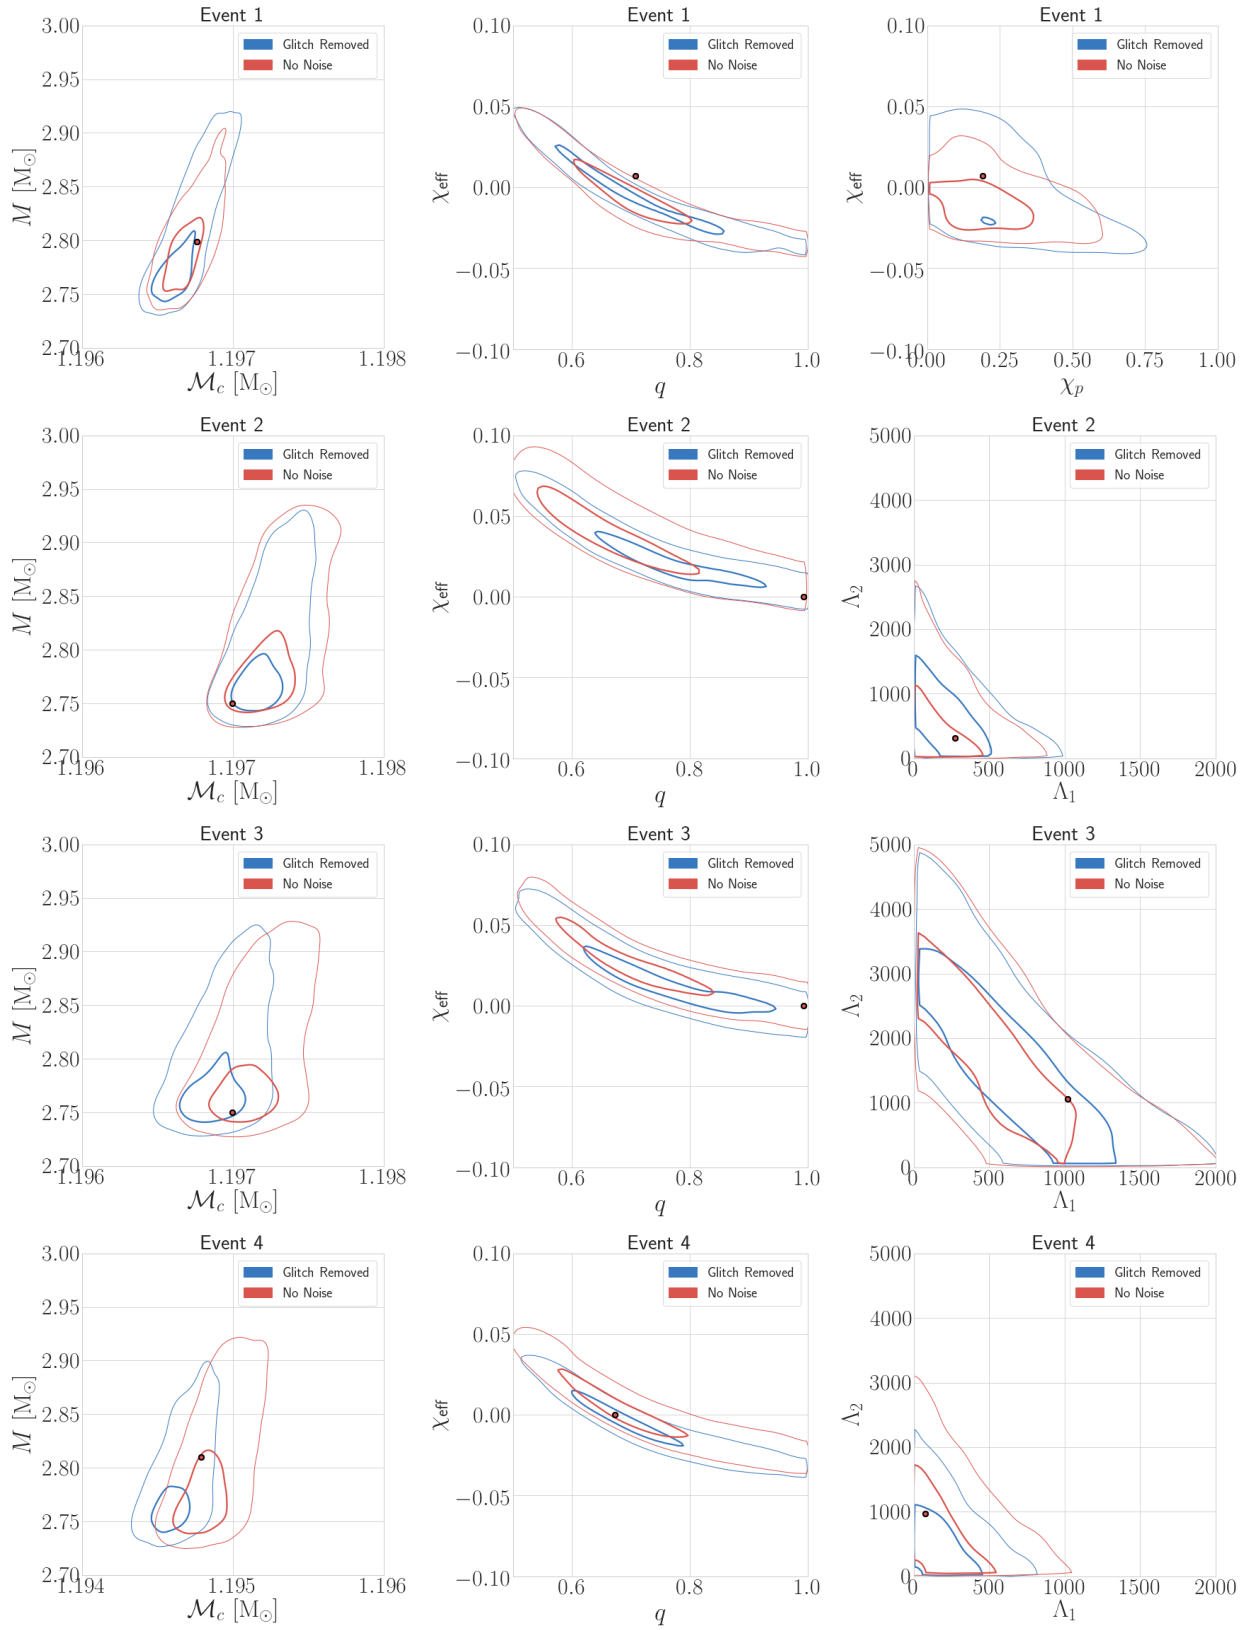
FIG. 3: Posterior densities for the masses, spins, and tidal deformabilities for four overflow glitches (top to bottom). The left column shows the posteriors in total mass and chirp mass (in the detector frame); the middle column shows mass ratio versus χ_{eff}; the last column shows χ_{eff} versus χ_{p} (first event, top row) or component tidal parameter posteriors Λ_1 and Λ_2 (remaining three events). Thick (thin) lines show the 50% (90%) credible regions. The true values of the parameters are represented by the black edged marker. In all cases the parameter estimates after removing the overflow glitch are consistent with estimates from injections in zero noise.

Glitch GPS time		True Value	Recovered Values		
			Zero-Noise	No Mitigation	With Mitigation
Event 1 (1186164816)	m_1 [$M_\odot$]	1.64	$1.64^{+0.24}_{-0.23}$	$24.44^{+2.62}_{-0.92}$	$1.64^{+0.25}_{-0.23}$
	m_2 [$M_\odot$]	1.16	$1.16^{+0.18}_{-0.13}$	$2.79^{+0.20}_{-0.92}$	$1.16^{+0.17}_{-0.14}$
	q	0.71	$0.70^{+0.24}_{-0.16}$	$0.11^{+0.02}_{-0.03}$	$0.70^{+0.23}_{-0.17}$
	χ_{eff}	0.007	$-0.01^{+0.04}_{-0.02}$	$0.38^{+0.02}_{-0.69}$	$-0.01^{+0.05}_{-0.03}$
	Λ_1	0	-	-	-
	Λ_2	0	-	-	-
Event 2 (1186691156)	m_1 [$M_\odot$]	1.38	$1.64^{+0.26}_{-0.24}$	$61.19^{+0.17}_{-0.96}$	$1.60^{+0.29}_{-0.20}$
	m_2 [$M_\odot$]	1.37	$1.16^{+0.19}_{-0.14}$	$8.19^{+0.07}_{-0.13}$	$1.19^{+0.16}_{-0.17}$
	q	0.99	$0.71^{+0.26}_{-0.17}$	$0.13^{+0.00}_{-0.00}$	$0.75^{+0.22}_{-0.20}$
	χ_{eff}	0	$0.04^{+0.05}_{-0.03}$	$0.88^{+0.00}_{-0.01}$	$0.02^{+0.05}_{-0.02}$
	Λ_1	275	193^{+529}_{-180}	1429^{+45}_{-14}	248^{+568}_{-227}
	Λ_2	309	503^{+1516}_{-467}	2621^{+2051}_{-2222}	668^{+1395}_{-599}
Event 3 (1186885739)	m_1 [$M_\odot$]	1.38	$1.60^{+0.28}_{-0.21}$	$40.18^{+0.44}_{-0.14}$	$1.59^{+0.28}_{-0.19}$
	m_2 [$M_\odot$]	1.37	$1.18^{+0.17}_{-0.16}$	$6.17^{+0.04}_{-0.07}$	$1.20^{+0.16}_{-0.17}$
	q	0.99	$0.74^{+0.23}_{-0.19}$	$0.15^{+0.001}_{-0.001}$	$0.75^{+0.21}_{-0.20}$
	χ_{eff}	0	$0.02^{+0.05}_{-0.03}$	$0.33^{+0.02}_{-0.02}$	$0.01^{+0.05}_{-0.02}$
	Λ_1	1018	588^{+1084}_{-525}	4958^{+40}_{-144}	645^{+1070}_{-574}
	Λ_2	1062	1566^{+2461}_{-1411}	$5.00^{+1.01}_{-1.40}$	1545^{+2420}_{-1382}
Event 4 (1186300855)	m_1 [$M_\odot$]	1.68	$1.63^{+0.25}_{-0.23}$	$38.71^{+0.08}_{-0.05}$	$1.62^{+0.26}_{-0.22}$
	m_2 [$M_\odot$]	1.13	$1.16^{+0.18}_{-0.14}$	$5.80^{+0.02}_{-0.02}$	$1.17^{+0.17}_{-0.15}$
	q	0.67	$0.71^{+0.25}_{-0.17}$	$0.15^{+0.001}_{-0.001}$	$0.72^{+0.23}_{-0.18}$
	χ_{eff}	0	$0^{+0.04}_{-0.03}$	$0.33^{+0.01}_{-0.01}$	$-0.01^{+0.04}_{-0.02}$
	Λ_1	77	240^{+605}_{-220}	4991^{+8}_{-27}	211^{+576}_{-193}
	Λ_2	973	728^{+1686}_{-668}	$0.14^{+0.02}_{-0.03}$	591^{+1502}_{-539}

TABLE I: Properties of the four simulated signals analyzed. The first column gives the GPS time of the overflow glitch around which we make the BNS injections; the second column list the binary parameters we study; the third column gives they injected (true values). The remaining three columns give the median recovered value for each parameter, as well as the 90% credible intervals in the case of a zero-noise-realization injection (fourth column), a parameter estimation analysis on data that include the glitch (fifth column) and on data obtained after glitch mitigation with `BayesWave` (sixth column).

surable for Λ_2. The uncertainties from the posteriors are wider than their glitch-free counterparts. Despite this, the credible intervals miss the known value by a wide margin, several times their own width. This implies that without mitigation, not only is the most basic parameter estimation biased, but simply incorrect.

IV. CONCLUSIONS

The focus of this work is to examine potential biases induced by various noise transient subtraction methods. There is a trade off between speed and fidelity in terms of their efficacy. Simple but fast methods which ignore the morphology of the noise transient or the data spectra are likely to introduce noticeable and predictable biases in signal amplitude recovery. Also worth noting is that the bias and widening of uncertainty induced by these methods may be more severe for weaker (lower SNR) signals — GW170817 was the highest SNR GW signal detected to date. Amelioration of these biases by introduc-

ing the same gating procedure to the template waveform remains impractical. In a low latency environment where sky maps are required without waiting for a full pass of a glitch extraction method like `BayesWave`, the time-frequency data excision in Section II B seems best suited and of appropriate latency. While parameter biases are reduced for more careful treatments of data excision, they are not a substitute for full glitch subtraction — particularly when high data fidelity is required to retrieve an accurate estimation of physical parameters. Indeed we find that the data produced after glitch-mitigation by `BayesWave` are consistent with our zero-noise study to within shifts that can be attributed to the presence of noise.

The consistent recovery of the spin and tidal parameters are important tests: the frequencies affected are not only in the most sensitive bandwidth of the GW interferometers, but also where subdominant effects on the phase of the waveform, such as those from spin and tidal deformabilities, become measurable. If the glitch removal procedure damaged the structure or coherence of the sig-

nal, then it would be expected that those parameters would exhibit significant bias. A study examining the effect of glitches on source property estimation by time-shifting glitches along the time-frequency track of signals is in preparation [30]. Another study which examines the biases in physical parameter recovery for binary black hole signals where the merger time overlaps a glitch can be found in [29].

The overflow glitch studied here is, in some sense, a best case scenario for glitch subtraction with BayesWave, because the signal and glitch have sufficiently different morphology. In particular, while the signal and glitch overlap in time-frequency (T-F) space, they are fit by wavelets having very different quality factor Q (which governs the number of cycles in the wavelet)–high Q for the BNS and low Q for the glitch. Generally, when glitches and signals do not overlap in T-F-Q space, a straightforward application of BayesWave will provide clean residuals for the PE analysis. Fortunately, that was the scenario presented by GW170817. Higher mass GW signals, such as BBH mergers, use lower Q wavelets and would therefore be more difficult to separate from a similar glitch with the current analysis tools. To mitigate against such a scenario in the future, we are developing glitch subtraction algorithms which include information about the network coherence of the signal, and lack thereof for any glitches.

Since multi-messenger astronomy involving NSs in compact binaries is expected to be highly featured in the upcoming Advanced LIGO/Virgo/KAGRA observing runs [52], confidence in the measured properties of GW signals is paramount. In the presence of short-duration noise transients near a merger signal, gating and data subtraction procedures are acceptable stopgap measures to obtain data products needed for rapid iden-

tification of electromagnetic counterparts. However, to obtain unbiased results without inflating uncertainties, we find that data-tailored glitch-removal techniques such as BayesWave are required.

[1] B. P. Abbott et al. (Virgo, LIGO Scientific), Phys. Rev. Lett. **119**, 161101 (2017), 1710.05832.

[2] B. P. Abbott et al. (Virgo, Fermi-GBM, INTEGRAL, LIGO Scientific), Astrophys. J. **848**, L13 (2017), 1710.05834.

[3] B. P. Abbott et al. (GROND, SALT Group, Oz-Grav, DFN, INTEGRAL, Virgo, Insight-Hxmt, MAXI Team, Fermi-LAT, J-GEM, RATIR, IceCube, CAAS-TRO, LWA, ePESSTO, GRAWITA, RIMAS, SKA South Africa/MeerKAT, H.E.S.S., 1M2H Team, IKI-GW Follow-up, Fermi GBM, Pi of Sky, DWF (Deeper Wider Faster Program), Dark Energy Survey, MASTER, AstroSat Cadmium Zinc Telluride Imager Team, Swift, Pierre Auger, ASKAP, VINROUGE, JAGWAR, Chandra Team at McGill University, TTU-NRAO, GROWTH, AGILE Team, MWA, ATCA, AST3, TOROS, Pan-STARRS, NuSTAR, ATLAS Telescopes, BOOTES, CaltechNRAO, LIGO Scientific, High Time Resolution Universe Survey, Nordic Optical Telescope, Las Cumbres Observatory Group, TZAC Consortium, LOFAR, IPN, DLT40, Texas Tech University, HAWC, ANTARES, KU, Dark Energy Camera GW-EM, CALET, Euro VLBI Team, ALMA), Astrophys. J. **848**, L12 (2017), 1710.05833.

[4] J. Aasi et al. (LIGO Scientific), Class. Quant. Grav. **32**, 074001 (2015), 1411.4547.

[5] F. Acernese et al. (VIRGO), Class. Quant. Grav. **32**, 024001 (2015), 1408.3978.

[6] C. Messick, K. Blackburn, P. Brady, P. Brockill, K. Cannon, R. Cariou, S. Caudill, S. J. Chamberlin, J. D. E. Creighton, R. Everett, et al., Phys. Rev. D **95**, 042001 (2017), URL https://link.aps.org/doi/10.1103/PhysRevD.95.042001.

[7] A. H. Nitz, T. Dal Canton, D. Davis, and S. Reyes, ArXiv e-prints arXiv:1805.11174 (2018), 1805.11174.

[8] L. P. Singer and L. R. Price, Phys. Rev. D **93**, 024013 (2016), URL https://link.aps.org/doi/10.1103/PhysRevD.93.024013.

[9] L. K. Nuttall, T. J. Massinger, J. Areeda, J. Betzwieser, S. Dwyer, A. Effler, R. P. Fisher, P. Fritschel, J. S. Kissel, A. P. Lundgren, et al., Classical and Quantum Gravity **32** (2015).

[10] B. P. Abbott, R. Abbott, T. D. Abbott, M. R. Abernathy, F. Acernese, K. Ackley, M. Adamo, C. Adams, T. Adams, P. Addesso, et al., Classical and Quantum Gravity **33** (2016).

[11] B. P. Abbott, R. Abbott, T. D. Abbott, M. R. Abernathy, F. Acernese, K. Ackley, C. Adams, T. Adams, P. Addesso, R. X. Adhikari, et al., Classical and Quantum Gravity **35** (2018).

[12] S. A. Usman, A. H. Nitz, I. W. Harry, C. M. Biwer, D. A. Brown, M. Cabero, C. D. Capano, T. Dal Canton, T. Dent, S. Fairhurst, et al., Classical and Quantum Gravity **33** (2016).

[13] A. Nitz, I. Harry, D. Brown, C. M. Biwer, J. Willis, T. D. Canton, et al., *ligo-cbc/pycbc: O2 production release 19* (2017), URL https://zenodo.org/record/852372#.Wz-QcVMvwxM.

[14] B. D. Metzger and E. Berger, Astrophys. J. **746** (2012).

[15] The LIGO Scientific Collaboration, The Virgo Collaboration, et al., GCN **21513**, 1 (2017).

[16] D. A. Coulter et al., Science **358**, 1556 (2017), 1710.05452.

[17] N. J. Cornish and T. B. Littenberg, Classical and Quantum Gravity **32**, 135012 (2015), URL http://stacks.iop.org/0264-9381/32/i=13/a=135012.

[18] B. P. Abbott et al. (Virgo, LIGO Scientific) (2018), 1805.11579.

[19] B. P. Abbott et al. (Virgo, LIGO Scientific) (2018), 1805.11581.

[20] J. Veitch et al., Phys. Rev. **D91**, 042003 (2015), 1409.7215.

[21] Éanna É.. Flanagan and T. Hinderer, Phys. Rev. D **77**, 021502 (pages 5) (2008), URL http://link.aps.org/abstract/PRD/v77/e021502.

[22] T. Hinderer, Astrophys. J. **677** (2008).

[23] T. Hinderer, B. D. Lackey, R. N. Lang, and J. S. Read, Phys. Rev. D **81** (2010).

[24] J. S. Read, C. Markakis, M. Shibata, K. Uryū, J. D. E. Creighton, and J. L. Friedman, Phys. Rev. D **79**, 124033 (2009), arXiv:0901.3258.

[25] L. S. Finn and D. F. Chernoff, Phys. Rev. D **47**, 2198 (1993).

[26] É. É. Flanagan and S. A. Hughes, Phys. Rev. D **57**, 4535 (1998).

[27] B. Allen, W. G. Anderson, P. R. Brady, D. A. Brown, and J. D. E. Creighton, Phys. Rev. D **85** (2012).

[28] B. S. Sathyaprakash and B. F. Schutz, Living Rev. Rel. **12**, 2 (2009), 0903.0338.

[29] J. Powell, Class. Quant. Grav. **35**, 155017 (2018), 1803.11346.

[30] J. McIver, T. J. Massinger, D. Davis, L. K. Nuttall, and R. Smith (2018).

[31] B. P. Abbott et al., CQG **33** (2016), 1602.03844.

[32] J. Veitch and A. Vecchio, Phys. Rev. **D81**, 062003 (2010), 0911.3820.

[33] J. Kanner, T. Littenberg, N. Cornish, M. Millhouse, E. Xhakaj, F. Salemi, M. Drago, G. Vedovato, and S. Klimenko, Phys. Rev. D **93**, 022002 (2015).

[34] R. Biswas, L. Blackburn, J. Cao, R. Essick, K. A. Hodge, E. Katsavounidis, K. Kim, Y.-M. Kim, E.-O. Le Bigot, C.-H. Lee, et al., Phys. Rev. D **88**, 062003 (2013), 1303.6984.

[35] R. Essick, L. Blackburn, and E. Katsavounidis, Classical and Quantum Gravity **30**, 155010 (2013), 1303.7159.

[36] T. B. Littenberg and N. J. Cornish, Phys. Rev. **D91**, 084034 (2015), 1410.3852.

[37] W. M. Farr, B. Farr, and T. Littenberg, Tech. Rep. LIGO-T1400682, LIGO Project (2015), URL https://dcc.ligo.org/LIGO-T1400682/public.

[38] E. Racine, Phys. Rev. **D78**, 044021 (2008), 0803.1820.

[39] C. Cutler and É. E. Flanagan, Phys. Rev. D **49**, 2658 (1994).

[40] M. Hannam, P. Schmidt, A. Bohé, L. Haegel, S. Husa, F. Ohme, G. Pratten, and M. Pürrer, Phys. Rev. Lett. **113**, 151101 (2014), 1308.3271.

[41] P. Schmidt, F. Ohme, and M. Hannam, Phys. Rev. D **91**, 024043 (2015), 1408.1810.

[42] S. Husa, S. Khan, M. Hannam, M. Pürrer, F. Ohme, X. J. Forteza, and A. Bohé, Phys. Rev. D **93**, 044006 (2016), 1508.07250.

[43] S. Khan, S. Husa, M. Hannam, F. Ohme, M. Pürrer, X. J. Forteza, and A. Bohé, Phys. Rev. D **93**, 044007 (2016), 1508.07253.

[44] T. Damour, B. R. Iyer, and B. S. Sathyaprakash, Phys. Rev. D **66**, 027502 (2002).

[45] J. Vines, É. É. Flanagan, and T. Hinderer, Phys. Rev. D **83**, 084051 (2011).

[46] T. Dietrich et al. (2018), 1804.02235.

[47] T. B. Littenberg, J. G. Baker, A. Buonanno, and B. J. Kelly, Phys. Rev. **D87**, 104003 (2013), 1210.0893.

[48] C. L. Rodriguez, B. Farr, V. Raymond, W. M. Farr, T. B. Littenberg, D. Fazi, and V. Kalogera, Astrophys. J. **784**, 119 (2014), 1309.3273.

[49] S. Vitale and H.-Y. Chen, Phys. Rev. Lett. **121**, 021303 (2018), URL https://link.aps.org/doi/10.1103/PhysRevLett.121.021303.

[50] S. Nissanke, D. E. Holz, S. A. Hughes, N. Dalal, and J. L. Sievers, Astrophys. J. **725**, 496 (2010), 0904.1017.

[51] M. Vallisneri, Phys. Rev. Lett. **107**, 191104 (2011), 1108.1158.

[52] B. P. Abbott, R. Abbott, T. D. Abbott, M. R. Abernathy, F. Acernese, K. Ackley, C. Adams, T. Adams, P. Addesso, R. X. Adhikari, et al., Living Reviews in Relativity **19**, 1 (2016).