



This is the accepted manuscript made available via CHORUS. The article has been published as:

Gravitational wave sources as timing references for LISA data

Tyson B. Littenberg

Phys. Rev. D **98**, 043008 — Published 9 August 2018

DOI: [10.1103/PhysRevD.98.043008](https://doi.org/10.1103/PhysRevD.98.043008)

Gravitational Wave Sources as Timing References for LISA Data

Tyson B. Littenberg¹

¹*NASA Marshall Space Flight Center, Huntsville, AL 35812, USA*

(Dated: July 25, 2018)

In the mHz gravitational-wave band, galactic ultra-compact binaries (UCBs) are continuous sources emitting at near-constant frequency. The signals from many of these galactic binaries will be sufficiently strong to be detectable by the *Laser Interferometer Space Antenna* (LISA) after $\sim\mathcal{O}(1)$ week of observing. In addition to their astrophysical value, these UCBs can be used to monitor the data quality of the observatory. This paper demonstrates the capabilities of galactic UCBs to be used as calibration sources for LISA by demanding signal coherence between adjacent week-long data segments separated by a gap in time of *a priori* unknown duration. A parameter for the gap duration is added to the UCB waveform model and used in a Markov-chain Monte Carlo algorithm simultaneously fitting for the astrophysical source parameters. Results from measurements of several UCBs are combined to produce a joint posterior on the gap duration. The measurement accuracy's dependence on how much is known about the UCBs through prior observing, and seasonal variations due to the LISA orbital motion, is quantified. The duration of data gaps in a two-week segment of data can be constrained to within ~ 0.2 s using $\mathcal{O}(10)$ UCBs after one month of observing. The timing accuracy from UCBs improves to $\lesssim 0.1$ s after 1 year of mission operations. These results are robust to within a factor of ~ 2 when taking into account seasonal variations.

I. INTRODUCTION

The *Laser Interferometer Space Antenna* (LISA) will be a unique astronomical observatory, unveiling the gravitational wave sky in the mHz band expected to be richly populated with sources ranging from stellar mass compact binaries in the galaxy, extreme mass ratio in-spirals in galactic centers, to supermassive black hole mergers at high redshifts [1].

Most expected LISA sources will be long-duration, while the data stream from the satellites will be interrupted due to periodic maintenance and random disturbances. Many signal classes are identified in the data stream using phase-coherent template waveform models, requiring data segments to be combined at accurate time intervals to prevent biases introduced by artificial phase shifts in the gravitational wave (GW) signals introduced by gap timing errors. While spacecraft hardware solutions for accurately timing the duration of gaps are understood, persistent and strong gravitational-wave sources can be used as calibration sources to provide an independent measure of start times for data segments. At the very least, studying the effectiveness of GW sources at measuring gap durations will quantify the error budget for timing accuracy required for unbiased astrophysical parameter estimation. If the instrumental timing accuracy is superior to what can be achieved using GW sources, the uncertainties in gap duration can be justifiably ignored by the science analysis. On the other hand, if hardware solutions provide comparable accuracy to the GW measurement, uncertainties in the gap duration should be incorporated into the science analysis in order to marginalize over the effects on parameter estimation. This paper will explore how GWs from ultra-compact binaries (UCBs) in the galaxy can be exploited as phase standards to measure the duration of data gaps.

Ultra-compact binaries (UCBs), mostly double white

dwarf systems, are prolific gravitational wave sources emitting at near-constant frequency in the LISA band [2]. LISA will detect $\mathcal{O}(10^5)$ UCBs after $\mathcal{O}(1)$ year of observations [3]. The strongest UCBs will be detectable after ~ 1 week of observations. As the observing time of the LISA mission increases, the individual UCBs will be increasingly-well measured, improving the accuracy with which the sources can be used to monitor data quality as new data are acquired.

Ref. [4] studied the impact of data disturbances, including instantaneous drop outs, in the recovery of UCB-like monochromatic signals in a LISA-like data stream and demonstrated success in removing disturbances, although frequent (\sim daily) disruptions were problematic. The effect of data gaps on parameter estimation of UCBs was studied in Ref. [5] where it was confirmed that minimizing the number of gaps, as opposed to the duration of each gap, results in the least degradation to the measurement of the astrophysical parameters of UCBs. The study assumed perfect knowledge of the gap start and stop times.

This paper will quantitatively demonstrate how well the brightest UCBs serve as timing standards by demanding coherence across gaps between two short-duration (one week) data segments after one, two, six, and twelve months of prior observing. The results show that, after one month of observations, the 25 brightest binaries provide constraints on the duration of gaps in data to better than ~ 0.2 s, improving to $\lesssim 0.1$ s after one year of observing.

Incorporating data characterization parameters into the model for astrophysical sources as a way of characterizing gravitational wave data is not a new idea. Using gravitational wave sources to validate ground-based gravitational wave data calibration has been proposed by Ref. [6] in the case where an ensemble of joint gamma-ray burst and binary neutron star observations are avail-

able. For a more immediate application, calibration uncertainty has been identified as an additional source of measurement error [7]. The method described in Ref [8] is routinely used to marginalize over calibration uncertainty in the analysis of compact binary mergers observed by ground based detectors (e.g. [9]). This paper applies a similar conceptual approach to space-based gravitational wave data.

The paper is organized as follows: Section II describes the model for the LISA observatory and the UCB signals, and modifications made to include the gap parameters. Section III describes the analysis, including details about the data and source simulations in Section III A and the numerical results in Section III B. Section IV summarizes the key results, addresses potential weaknesses in this study, argues why they are permissible, and outlines future directions of research based on these findings.

II. DATA MODEL

LISA will be a constellation of three free-flying spacecraft, each on independent heliocentric orbits approximately one AU from the sun. Each is equipped with an optical metrology system to precisely monitor the distances between free-falling test masses housed within the spacecraft. From the inter-spacecraft ranging measurements, two Michelson-like interferometry signals are synthesized digitally to simultaneously measure the two GW polarizations. The simulated detector data in this study is consistent with the instrument noise levels quoted in the LISA L3 Mission Proposal [1]. This analysis takes advantage of the noise-orthogonal A and E data streams constructed from linear combinations of the Michelson-like time delay interferometry channels [10]. The A and E data streams assume identical noise in each of the six interferometer links of the spacecraft, and so are an idealized case that will not be realized during mission operations, but are suitable for this proof-of-concept study.

Because the UCBs observable by LISA are widely separated, the orbital evolution due to GW emission is small over the mission lifetime, and thus the sources appear in the data stream at near-constant frequency. Thus, UCBs do not produce the dramatic “chirp” waveforms that are customarily thought of in the context of GW observations of compact binaries. The simple frequency evolution of UCBs is a double-edged sword. On one hand it dramatically simplifies the waveform modeling, as beyond leading order post-Newtonian terms are negligible. On the other hand, it is the higher-order post-Newtonian terms in the in-spiral (and the full-relativistic effects in the merger, etc.) that encode the information about the constituents of the binary (e.g. masses and spins). Furthermore, a subset of UCBs will be interacting with one another, and their orbital evolution will be dominated by mass transfer (e.g. AM CVn systems), causing the orbital period to *increase* during LISA observations [11]. As a consequence, masses of (and the distance to) typical

UCBs are not directly observable, but rather combined into phenomenological parameters for the GW amplitude \mathcal{A} , frequency f_0 (defined at some fiducial time t_0), and linear frequency evolution \dot{f} which may, in principle, contain relativistic effects *as well as* contributions from mass transfer.

The remaining observables are two angles defining the sources position on the sky α, δ and three angles describing the source orientation with respect to the observer ι, ψ, ϕ_0 . For this work the location and orientation angles are nuisance parameters which are marginalized in this analysis, though they are of paramount importance for multimessenger observations [12–15]. Together, these quantities form a parameter vector θ which describes a single frequency-domain UCB template. This study uses the fast-slow decomposition originally described in Ref. [16] to forward-model the response of the LISA detector $\tilde{h}(f; \theta)$ to the gravitational wave source.

Using the UCBs as time standards requires modifying the waveform $\tilde{h}(f; \theta)$ with a “calibration parameter” δt which determines the duration of the gap between two segments of data. Without loss of generality, it is assumed that the start and stop time for the first segment of data, t_0 and t_f , respectively, are perfectly known. The δt then parameterizes the absolute start time of the second segment as $t_f + \delta t$ (see Eq. 1). A uniform prior $p(\delta t)$ for the gap duration is used with support between ± 20 s.

$$\begin{aligned} \tilde{h}(f; \theta) &\rightarrow \begin{cases} \tilde{h}(f, t_0; \theta) & \text{for } t_0 < t < t_f \\ \tilde{h}(f, t_f + \delta t; \theta) & \text{for } t > t_f \end{cases} \\ p(\delta t) &= U(\pm 20 \text{ s}) \end{aligned} \quad (1)$$

The choice of a 40 second window for the start time of the second segment is completely arbitrary. At the time of this writing, the expected duration of data gaps in the LISA data stream has not yet been determined. However, the total gap duration will not effect the results reported here. For short segments of data UCBs are effectively radiating at constant frequency, requiring months to years of observing to detect any orbital evolution. The frequency drift that will occur over the interval of a data gap is undetectable. Consequently, using UCBs to measure gap durations with two weeks of data basically comes down to phase-matching two sinusoids of identical frequency. For a single binary, there is a degeneracy in the total gap duration and the integer number of phase windings of the signal, however that degeneracy is broken, with reasonable priors on the gap duration, by using multiple binaries as a comb of phase standards at different frequencies. The effect of shifting the prior window on the segment start time by hours, or even days, is negligible.

III. ANALYSIS

The goal of this work is to establish an understanding of the capability for astrophysical sources to be used as timing standards. This capability is demonstrated by analyzing simulated data containing high S/N UCBs using a parallel tempered Markov Chain Monte Carlo (MCMC) pipeline adapted from Ref [17], which was specifically designed for galactic binary searches in LISA data.

A. Simulations and Procedure

The simulated galactic population of UCBs used for this study is the same as in the first generation of LISA Data Challenges [25]. The galaxy simulation combines a population of detached white dwarf binaries from [18] and interacting binaries (e.g. AMCVn-type systems) from [19]. The binaries are distributed randomly in a model of the galaxy and assigned random orientations with respect to the observer. From ~ 1 to ~ 4 mHz the LISA band will be dominated by the millions of UCBs which blend together to form an irresolvable foreground. Above ~ 4 mHz the remaining UCBs are isolated from one another in frequency and have high S/N making them easily detectable. From the galaxy simulation, binaries are ranked by S/N and the 25 loudest systems with $f_0 > 4$ mHz are selected for study. There are high S/N binaries at lower frequencies, however they will be contending with overlapping signals and the galactic confusion noise which make for a more complicated analysis owing to the computational cost of determining the confusion noise level, contending with its time-varying power due to variations in the orientation of the LISA detector plane induced by the satellites' orbital motion, and potential covariances with sources nearby in frequency to the binary in question. In general, exploiting UCBs for data quality purposes will be most straightforward using isolated binaries at high frequency.

This study does not include binaries already discovered through electromagnetic observations known to be LISA sources [15, 20, 21]. While often discussed as potential calibration sources for LISA—they are commonly referred to as “verification binaries”—population synthesis simulations predict the existence of significantly stronger GW sources than the verification binaries. The known LISA binaries are at lower frequency where the noise characteristics are more complicated, due to the dense population of UCBs at those frequencies, meaning longer integration times will be required for these binaries to be detectable. The high frequency binaries observable by LISA are easier to detect and characterize, and therefore will serve as better phase standards over shorter integration times.

Figure 1 shows the strain amplitude of the simulated sources as a function of their GW frequency, compared to the LISA **noise amplitude spectral density** curve and an estimate of the confusion noise level after one year of observations, as described in [22]. The sources are




FIG. 1: Circles show locations of the UCBs in frequency-amplitude space colored by their S/N after two weeks of observing. The black line is the LISA **noise amplitude spectral density** curve, including the contribution from the unresolved foreground of low-frequency UCBs in the galaxy (dashed line). The foreground is the residual signal from UCBs after all resolvable signals have been identified and subtracted. This example is from the expected residual after one year of LISA operations. For this study, only binaries with frequencies above the confusion noise were selected to avoid complexities due to seasonal variations of the foreground and source overlap.

colored by their S/N which depends on the location and orientation and so can not be simply read off from the height above the **noise** curve.

The parameter δt introduced to measure gap duration is correlated with the frequency parameter f_0 of each UCB signal. To improve the measurement of δt , results from several binaries are combined, and the knowledge about each source that builds up over time is incorporated into the analysis. To put it differently, when analyzing short segments of data, priors on the source parameters constructed from earlier observing times are enforced. While the degeneracies between δt and source parameters are not broken, constraints on the astrophysical parameters which build up over time will propagate through the analysis to improve the determination of δt in newly acquired data. If this procedure were to be adopted during mission operations, the analysis would need to be diligent about which data were used to fit gap durations and which were used to build priors to avoid “double-counting” the data i.e., using the same segments for analysis as were used to build priors.

For this study, separate priors are constructed after one month, three months, six months, and one year of prior observations for each high S/N UCB. From the posteriors inferred from these “pilot” MCMC runs, uniform priors are constructed for the eight source parameters spanning the 90% credible intervals of the marginalized posteriors. Distilling the full posteriors to a set of 1-dimensional uniform distributions is overly conservative, and it ignores

valuable information about the source parameters which, when included, would make the priors more representative of what is known about each binary. However, for this proof-of-principal exploratory study, this short cut is sufficient considering that at this time the actual scenario of the instrumental uncertainties in gap frequency and duration are yet to be determined. Developing efficient ways of building priors from previous observations, and updating posteriors as more data are acquired, will be an important development for LISA data analysis given the long-duration signals which will be observed.

To demonstrate the δt measurement capabilities, two weeks of data are simulated as two separate one-week segments. The data are analyzed demanding signal coherence in the two segments, fitting for the start time of the second segment relative to the end of the first along with the source parameters. This is done independently for each of the 25 brightest UCBs to determine the marginalized posterior distribution of δt for each injection i

$$p(\delta t|d_i) = \int p(\delta t, \theta|d_i) d\theta \quad (2)$$

Because δt is common to each data segment, analysis of each binary is an independent measure of the gap duration. Thus the joint posterior on δt is approximately the product of the marginalized distributions $p(\delta t|\mathbf{d}) = \prod_i p(\delta t|d_i)$ (ignoring covariances between UCB parameters introduced by each sharing the δt parameter). The marginalized posteriors for δt for the high S/N binaries are, to a good approximation, Gaussians. For simplicity, each $p(\delta t|d_i)$ is fit to a normal distribution so that the product can be computed analytically. More accurate approaches would be to use a hierarchical analysis to characterize the posterior on δt (e.g., see Ref. [23] for a GW-centric example) or, better yet, to perform a global fit to the 25 UCBs with δt being common for each source, but such complexity is not needed for this exploratory study.

For one-week data segments the strength of the signal, and therefore its utility as a calibration source, will depend on the source's position relative to the plane of the LISA constellation, causing the measurement accuracy of the source, and therefore the calibration parameter, to vary over the course of each year. To account for the dependence on the location of LISA in its orbit, the analysis is repeated once during each quarter of the year to bracket the range of outcomes for the galaxy realization and mission configuration used in this study.

B. Results

First covariances in a single source between the gap duration parameter δt and f_0 are investigated, to motivate the importance of building up priors from the LISA observing time in order to exploit UCBs as calibration sources in short segments of data. Correlations with




FIG. 2: Two-dimensional marginalized posterior distribution function for f_0 and gap duration δt (shifted so that 0 is the true value) from analyzing two weeks of data without assuming any prior knowledge of the source, demonstrating the covariance between source frequency f_0 and gap duration. Horizontal lines enclose the 90% credible region after one month (light blue, solid line) and twelve months (dark green, dot-dashed line) of prior observing. The black dotted line is the true value of f_0 for the simulated source. Priors constructed from pilot runs on UCBs will propagate to improved constraints on the gap duration.

other phase parameters of the UCB were not found to be important, as \dot{f} and ϕ_0 are poorly constrained when analyzing short data segments. Figure 2 shows the marginalized two-dimensional posterior distribution functions between δt and f_0 . The black dashed line is the true value for the simulated source, while the colored horizontal lines show the 90% credible intervals of the source parameters after one (light blue, solid lines) and twelve (dark green, dot-dashed lines) months of prior observation. As is evident in the correlations between the source and gap duration parameters, constraints on δt will rely on information about the sources which will build up over the mission lifetime. Note that for prior observation times of greater than one year, the improvement in measuring δt saturates i.e., further restricting the prior range on f_0 does not notably restrict $p(\delta t|d)$, as it is fundamentally limited by the short-duration segments being analyzed.

After adopting uniform priors over the 90% credible intervals from the pilot runs, each of the 25 brightest binaries are independently analyzed and the marginalized posteriors on δt are fit to a Gaussian distribution. Using uniform priors over the 90% credible intervals of the pilot runs, and multiplying Gaussian fits to the δt posteriors, is a simplified approach to what would be deployed for a production-quality version of this analysis, but is sufficient for exploring the potential for UCBs to function as phase standards in LISA data. Figure 3 shows a representative marginalized posterior on δt from the MCMC along with the Gaussian fit, demonstrating the consistency between the sampled distribution and the Gaussian




FIG. 3: Example Gaussian fit (black, dashed line) to the marginalized posterior distribution function for δt (green, solid histogram). The posterior is shifted in time so that the true value is at 0. Fitting the posterior to a Gaussian is a satisfactory approximation for this study. This demonstration used a ~ 7 mHz source after two months of prior observing.

approximation.

From the analytic fits to each $p(\delta t|d_i)$, the joint distribution on the gap duration from the ensemble of high-frequency, high S/N binaries is computed as the product of the Gaussian fits to the individual posteriors. The left panel of Figure 4 shows the 2σ envelopes for the joint posterior as a function of how many binaries are used in the fit. The four bands correspond to the joint posteriors using the four different prior observing scenarios in this analysis: One month (light blue), two months (dark blue), six months (light green), and twelve months (dark green). From this figure it is apparent that the posterior is converging towards the true value at $\delta t - \delta t_0 = 0$. The right panel, using the same color scheme, quantifies how the standard deviation (σ) of the joint distribution depends on the number of binaries in the fit, reaching a value between ~ 0.2 and ~ 0.1 when 20 binaries are employed in the joint fit, depending on the prior observing time. Note that, should the LISA science data be sampled at or below 1 Hz, the UCB measurement provides sub-sample accuracy on its measurement of the gap duration.

Because of the short duration data segments used to fit for the gap duration, constraints derived from the analysis of UCBs will exhibit seasonal variation as the orientation of the LISA constellation changes with respect to the source during the year. To test the impact of these seasonal variations, the analysis is repeated simulating two week segments of data from each quarter of the year and the joint posteriors on δt are compared after combining the top 25 binaries in the catalog. Figure 5 shows the joint posterior $p(\delta t|\mathbf{d})$ for the different quarters of the year. While the variation is apparent in the results, the fidelity of the δt constraints are robust to within a factor

of ~ 2 throughout the LISA orbit.

IV. DISCUSSION

This paper investigated how bright UCBs could be used as phase standards in LISA data to constrain the duration of data gaps. This work demonstrates that UCBs can be used as calibration sources on short segments of data. The duration of a data gap in a two-week segment of data can be constrained to within ~ 0.2 s using $\mathcal{O}(10)$ UCBs after one month of observing. The timing accuracy from UCBs improves to $\lesssim 0.1$ s after 1 year of observing. This result is robust to seasonal variations as the orientation of the LISA constellation changes with respect to the GW sources.

The results of this study are a useful benchmark for timing accuracy of gap durations for LISA, as it concerns UCBs. If gap durations are constrained through spacecraft “house-keeping” data to much better precision than found here, the effect will be negligible in the analysis of UCBs. Alternatively, should there be some ambiguity in the duration of a data gap to larger than the amounts demonstrated here, UCBs can be used to better-constrain the absolute start times of data segments.

Future investigations aiming to measure the impact of data gaps on LISA observations should consider how the uncertainties in gap duration as constrained by UCBs influence the characterization of other astrophysical sources, such as supermassive binary black hole mergers, to determine the conditions under which the level of timing accuracy demonstrated here is sufficient for other sources.

The analysis reported here is a proof-of-concept demonstration using methods similar to production-level algorithms for GW data analysis [8, 17, 24]. That being said, there are some places where the numerical experiment performed here would be vulnerable due to simplifications made in this study.

The procedure of constructing uniform priors over the 90% credible intervals from marginalized posteriors of the pilot runs is overly simplistic. Information obtained from the pilot runs, including the shape of the distributions and covariances between parameters, is lost by this approach. More egregiously, there is the possibility that the true source parameters will be excluded by the overly-simplified priors. It is not expected that addressing these issues would lead to qualitatively different conclusions about the UCBs’ utility in constraining the duration of data gaps, as the analyses on the week-long segments of data basically saturate their priors. To put it differently, small changes to the priors will lead to similarly small differences in the posteriors because the likelihood on the short segments of data is nearly constant over the priors from the pilot runs. In a production-quality analysis these details will need to be addressed even if their effect is small. To that end, investigations into strategies to build joint priors on the astrophysical

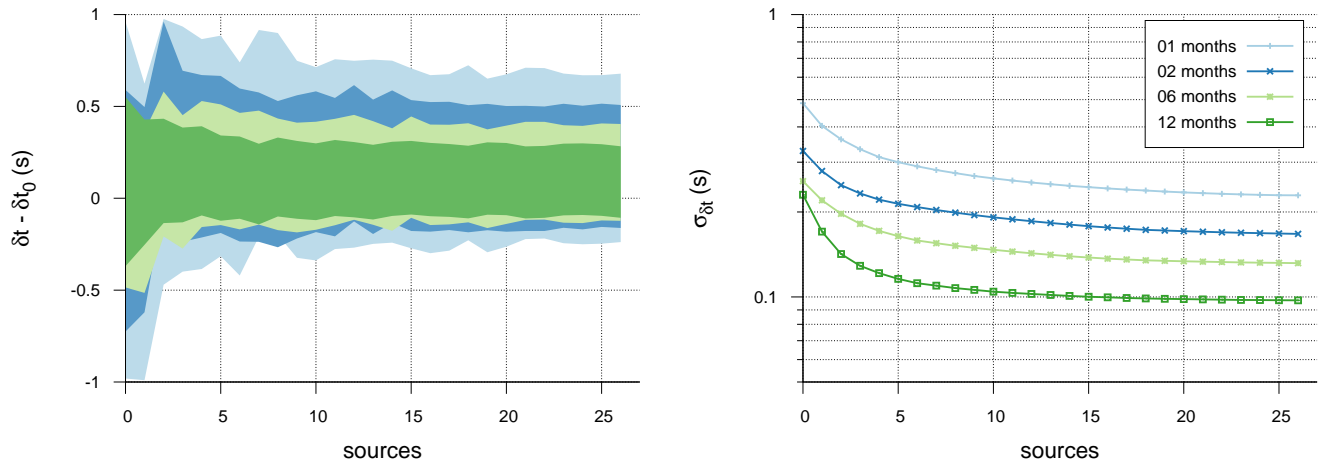


FIG. 4: Left panel: 2σ intervals of the joint posterior on the gap duration δt separating two 1-week segments of data as a function of the number of sources in the fit. The joint posterior $p(\delta t|\mathbf{d})$ converges towards the true value of $\delta t - \delta t_0 = 0$ as the number of sources increases. Different color bands correspond to priors built from one month (light blue) two months (dark blue), six months (light green) and twelve months (dark green) of observing. Right panel: The same results now showing the standard deviation $\sigma_{\delta t}$ of the joint posterior as a function of the number of sources, reaching 0.2 s (0.1 s) for 1 month (12 months) of prior observing.

parameters from posteriors inferred by pilot runs, which are more representative of the full N -dimensional distributions, while still being efficiently implementable in stochastic sampling algorithms, have begun.

Another weakness of this work is its dependence on a single realization of the UCB galactic population. The top 25 binaries out of a population of $\sim 10^4$ resolvable sources are in the tail of the source distribution, and characteristics of this sub-population will have a large variance. However, high frequency binaries in the galaxy are generically high S/N sources, and while this study used 25 binaries, the constraints on δt begin saturating after ~ 5 to ~ 15 binaries, depending on the prior observation time. A more thorough investigation into the variability of the LISA UCB sources, including into the population of high frequency detached binaries which are of scientific interest beyond their utility as calibration sources, as a result of different population simulations, is an area of ongoing study.

Finally, the dependence on the prior observations of binaries to confine source parameters when analyzing the short data segments raises the question of how those priors are obtained. The pilot runs to establish the priors did not consider that those data will also contain gaps. Therefore, the procedure demonstrated here assumes a reliable gap measurement from monitoring the house-keeping data during the first several months of the mission operations. In its absence, assuming uncertainties in the gap durations from the instrumental measurements are dominated by statistical error as opposed to a systematic offset, the performance demonstrated here would still be achievable, but the convergence time, i.e. the amount of prior observing needed for establishing similarly precise priors, would increase. A study incorporating such a

level of detail should be undertaken as the LISA mission design reaches maturity, leading to more realistic scenarios for the acquisition of data and the uncertainty budget in the data gaps.

This work demonstrates that UCBs can be used to measure the duration of data gaps and, as a corollary, places an upper limit on the uncertainty of the instrumentally-measured gap durations for it to have a negligible effect on UCB measurement. This paper also provides guidance for how analysis of LISA data will interact with the uncertainty budget of the observatory, and how UCBs can provide, at a minimum, a reserve means for maintaining phase-coherence in the LISA data stream.

V. ACKNOWLEDGMENTS

The author would like to acknowledge contributions to the analysis software used here by Neil Cornish and Travis Robson. This work was supported by NASA grant NNH15ZDA001N-APRA. The author would also like to acknowledge the support of the NASA LISA Study Office.

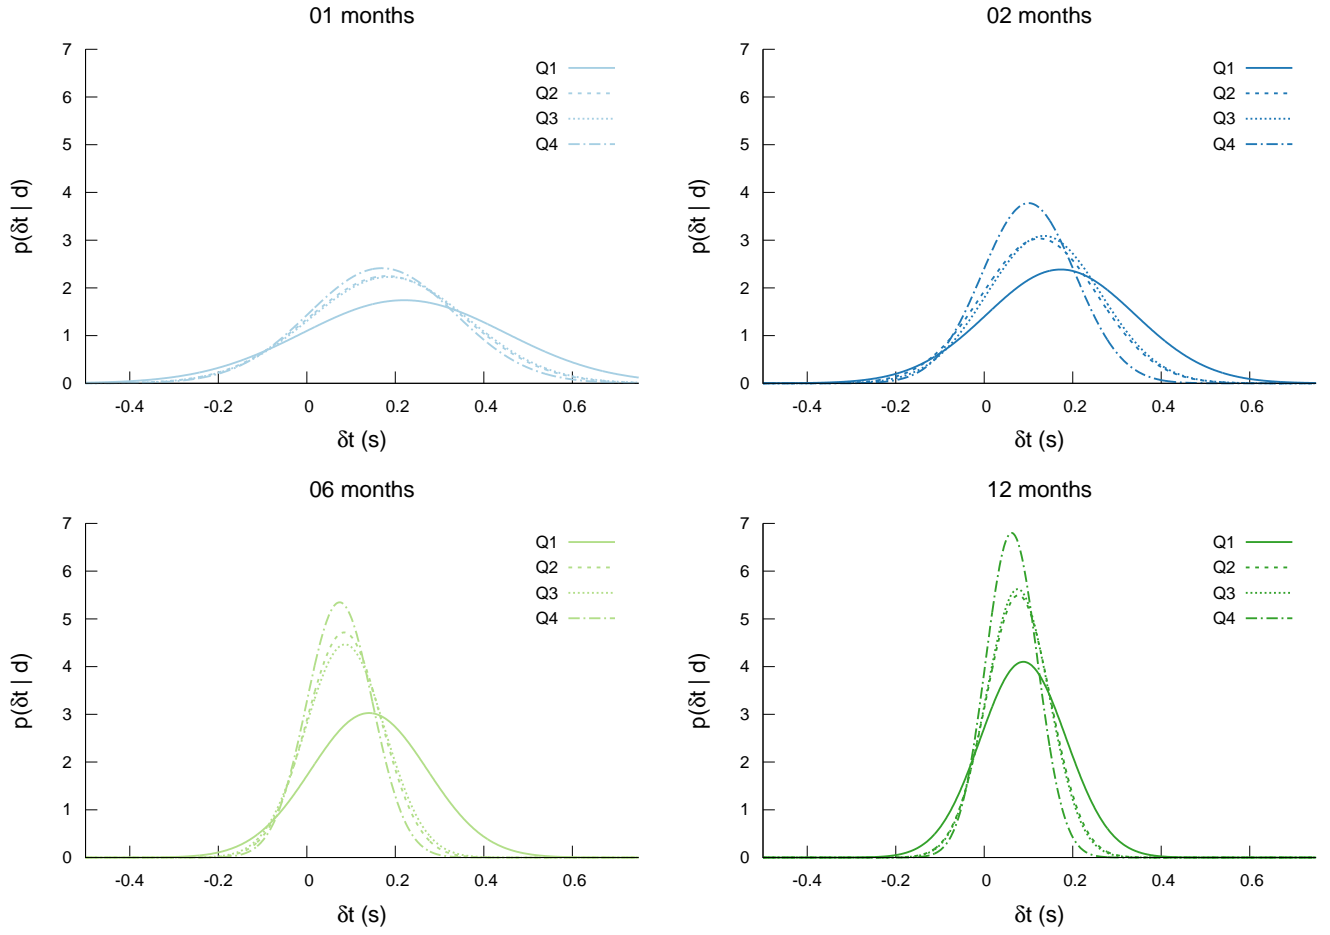


FIG. 5: Demonstration of the seasonal variation in the joint posterior of δt using the 25 brightest UCBs after one (top left), two (top right), six (bottom left), and twelve (bottom right) months of prior observing. The different distributions are the results from simulating data during each quarter of the LISA constellation's yearly orbit. While the seasonal variation is evident, the constraints on δt are of similar scale.

-
- [1] P. Amaro-Seoane, H. Audley, S. Babak, J. Baker, E. Barausse, P. Bender, E. Berti, P. Binetruy, M. Born, D. Bor-toluzzi, et al., ArXiv e-prints (2017), 1702.00786.
 - [2] D. Hils, P. L. Bender, and R. F. Webbink, *Astrophys. J.* **360**, 75 (1990).
 - [3] N. Cornish and T. Robson, in *Journal of Physics Confer-ence Series* (2017), vol. 840 of *Journal of Physics Con-ference Series*, p. 012024, 1703.09858.
 - [4] S. E. Pollack, *Classical and Quantum Gravity* **21**, 3419 (2004), URL <http://stacks.iop.org/0264-9381/21/i=14/a=007>.
 - [5] J. Carré and E. K. Porter, ArXiv e-prints (2010), 1010.1641.
 - [6] M. Pitkin, C. Messenger, and L. Wright, *Phys. Rev. D* **93**, 062002 (2016), 1511.02758.
 - [7] S. Vitale, W. Del Pozzo, T. G. F. Li, C. Van Den Broeck, I. Mandel, B. Aylott, and J. Veitch, *Phys. Rev. D* **85**, 064034 (2012), 1111.3044.
 - [8] W. M. Farr, B. F. Farr, and T. B. Littenberg, Tech. Rep., LIGO Project, <https://dcc.ligo.org/LIGO-T1400682/public/main> (2015).
 - [9] B. P. Abbott, R. Abbott, T. D. Abbott, M. R. Aber-nathy, F. Acernese, K. Ackley, C. Adams, T. Adams, P. Addesso, R. X. Adhikari, et al., *Physical Review Let-ters* **116**, 241102 (2016), 1602.03840.
 - [10] M. Tinto, F. B. Estabrook, and J. W. Armstrong, *Phys. Rev. D* **65**, 082003 (2002).
 - [11] G. Nelemans, L. R. Yungelson, and S. F. Portegies Zwart, *Mon. Not. Roy. Astron. Soc.* **349**, 181 (2004), astro-ph/0312193.
 - [12] A. Cooray, A. J. Farmer, and N. Seto, *Astrophys. J.* **601**, L47 (2004), arXiv:astro-ph/0310889.
 - [13] T. B. Littenberg, S. L. Larson, G. Nelemans, and N. J. Cornish, *Mon. Not. Roy. Astron. Soc.* (2012), first published online December 24, 2012 doi: 10.1093/mn-ras/sts507, arXiv:1207.4848 [astro-ph.IM].
 - [14] S. Shah and G. Nelemans, *Astrophys. J.* **790**, 161 (2014), 1406.3599.
 - [15] V. Korol, E. M. Rossi, P. J. Groot, G. Nelemans, S. Toon-en, and A. G. A. Brown, *Mon. Not. Roy. Astron. Soc.* **470**, 1894 (2017), 1703.02555.
 - [16] N. J. Cornish and T. B. Littenberg, *Phys. Rev. D* **76**, 083006 (2007), arXiv:0704.1808.
 - [17] T. B. Littenberg, *Phys.Rev.* **D84**, 063009 (2011), 1106.6355.
 - [18] S. Toonen, G. Nelemans, and S. Portegies Zwart, *Astron-omy and Astrophysics* **546**, A70 (2012), 1208.6446.
 - [19] G. Nelemans and C. A. Tout, *Mon. Not. Roy. Astron. Soc.* **356**, 753 (2005), astro-ph/0410301.
 - [20] A. Stroeer and A. Vecchio, *Class. Quantum Grav.* **23**, S809 (2006), arXiv:astro-ph/0605227.
 - [21] T. Kupfer, V. Korol, S. Shah, G. Nelemans, T. R. Marsh, G. Ramsay, P. J. Groot, D. T. H. Steeghs, and E. M. Rossi, ArXiv e-prints (2018), 1805.00482.
 - [22] N. Cornish and T. Robson, ArXiv e-prints (2018), 1803.01944.
 - [23] I. Mandel, *Phys. Rev. D* **81**, 084029 (2010), URL <http://link.aps.org/doi/10.1103/PhysRevD.81.084029>.
 - [24] J. Veitch, V. Raymond, B. Farr, W. Farr, P. Graff, et al., *Phys.Rev.* **D91**, 042003 (2015), 1409.7215.
 - [25] <https://gitlab.in2p3.fr/stas/MLDC>