

# Bayesian Model Selection for LISA Pathfinder

Nikolaos Karnesis,<sup>1,\*</sup> Miquel Nofrarias,<sup>1</sup> Carlos F. Sopena,<sup>1</sup> Ferran Gibert,<sup>1</sup> Michele Armano,<sup>2</sup>  
Heather Audley,<sup>3</sup> Giuseppe Congedo,<sup>4</sup> Ingo Diepholz,<sup>3</sup> Luigi Ferraioli,<sup>5</sup> Martin Hewitson,<sup>3</sup>  
Mauro Hueller,<sup>6</sup> Natalia Korsakova,<sup>3</sup> P.W. McNamara,<sup>7</sup> Eric Plagnol,<sup>8</sup> and Stefano Vitale<sup>6</sup>

<sup>1</sup>*Institut de Ciències de l'Espai, (CSIC-IEEC), Campus UAB, Facultat de Ciències, Torre C-5, 08193 Bellaterra, Spain*

<sup>2</sup>*ESAC, European Space Agency, Camino bajo del Castillo s/n, Urbanización  
Villafranca del Castillo, Villanueva de la Cañada, 28692 Madrid, Spain*

<sup>3</sup>*Albert-Einstein-Institut, Max-Planck-Institut für Gravitationsphysik und  
Universität Hannover, Callinstrasse 38, 30167 Hannover, Germany*

<sup>4</sup>*Department of Physics, University of Oxford, Keble Road, Oxford OX1 3RH, UK*

<sup>5</sup>*ETH Zürich, Institut für Geophysik, Sonneggstrasse 5, 8092 Zürich*

<sup>6</sup>*Dipartimento di Fisica, Università di Trento and INFN, Gruppo Collegato di Trento, 38123 Povo, Trento, Italy*

<sup>7</sup>*European Space Technology Centre, European Space Agency, Keplerlaan 1, 2200 AG Noordwijk, The Netherlands*

<sup>8</sup>*APC, Université Paris Diderot, CNRS/IN2P3, CEA/Ifnu, Observatoire de Paris,  
Sorbonne Paris Cité, 10 Rue A. Domet, 75205 Paris Cedex 13, France*

(Dated: June 6, 2014)

The main goal of the LISA Pathfinder (LPF) mission is to fully characterize the acceleration noise models and to test key technologies for future space-based gravitational-wave observatories similar to the eLISA concept. The data analysis team has developed complex three-dimensional models of the LISA Technology Package (LTP) experiment on-board LPF. These models are used for simulations, but more importantly, they will be used for parameter estimation purposes during flight operations. One of the tasks of the data analysis team is to identify the physical effects that contribute significantly to the properties of the instrument noise. A way of approaching this problem is to recover the essential parameters of a LTP model fitting the data. Thus, we want to define the simplest model that efficiently explains the observations. To do so, adopting a Bayesian framework, one has to estimate the so-called Bayes Factor between two competing models. In our analysis, we use three main different methods to estimate it: the Reversible Jump Markov Chain Monte Carlo method, the Schwarz criterion, and the Laplace approximation. They are applied to simulated LPF experiments where the most probable LTP model that explains the observations is recovered. The same type of analysis presented in this paper is expected to be followed during flight operations. Moreover, the correlation of the output of the aforementioned methods with the design of the experiment is explored.

## I. INTRODUCTION

LISA Pathfinder (LPF) [1, 2] is an European Space Agency mission that will serve as a technology demonstrator for a future space-based gravitational-wave observatory like eLISA [3]. The LPF mission will prove geodesic motion by monitoring the relative acceleration of two test masses in nominally free-fall in the frequency band of 1 to 30 mHz. The main instrument on-board LPF is the LISA Technology Package (LTP), which is a suite of experiments with the aim of measuring and characterizing the different contributions to the differential acceleration noise between the two test masses. This characterization is the main task of the data analysis team. To that end, dedicated experiments are going to be performed in order to estimate the unknown parameters of the system. And for that purpose, a number of parameter estimation methods and models of the LTP have been implemented [4–6]. The main question that arises is about the suitability of the different models implemented, or in simpler terms, which model can describe

better the observations of the experiment.

The motivation to implement an algorithm that will help us classify our models is due to the nature of the LTP system. There has been a lot of work trying to understand the instrument, and the models implemented are based on theoretical and experimental measurements from test campaigns [7, 8]. Selecting the most probable model is crucial for the analysis for two main reasons: first, the most suitable LTP model that describes the observed physical effects is chosen, and secondly, overfitting situations are avoided, together with biased estimation of the parameters of the system.

One could use several criteria and algorithms that classify competing models, but in the end, working in a Bayesian framework, the main aim is to calculate the *Bayes Factor*, which is a comparison between the *evidences* of the models [9, 10]. The evidence is defined as the probability of the data  $\vec{y}$  given the model, that is, the probability distribution on  $\vec{y}$  that quantifies the predictive capabilities of the given model.

Most of the approaches are based upon the likelihood evaluated at the maximum and a penalty for the number of parameters in the model consisting in multiplying by the so-called *Occam Factor*. The *Occam's razor* is the principle that states that the simplest hypothesis that

---

\*Electronic address: karnesis@ieec.uab.es

explains the observations is the most favorable. In our case we assume that we have to compare two different LTP models; model  $M_1$  and a simpler one,  $M_2$ , over a data set  $\vec{y}$ . If  $M_1$  is a more complex model, it presents more predictive capabilities than  $M_2$ , which translates to more disperse evidence over the data set  $\vec{y}$ . Thus, in the case where the data is compatible with both models,  $M_2$  will turn out to be more probable than  $M_1$ , without having to express any subjective dislike for complex models [10]. In other words, it is almost certain that the more parameters in a model the better the fit to the data. But taking it to the extreme, we can imagine a model with as many parameters as the data to fit. In that case we have over-parameterized the model while the aim is to include only the parameters which *substantially* improve it.

In this paper we investigate several methods to compare competing LTP models giving emphasis to Reversible Jump Markov Chain Monte Carlo techniques. The plan of the paper is as follows: In Sec. II, we make a brief introduction to Bayesian methods and in Sec. III, the LTP experiments and models are thoroughly explained and we investigate several applications of our Bayesian techniques to LTP experiments. In Sec. IV, we compare the available methods and discuss their output in connection with the experiment design. We end with a summary and conclusions in Sec. V. The codes for the different algorithms are integrated into the LISA Technology Package Data Analysis (LTPDA) Matlab toolbox [11, 12], that comes together with proper documentation and help for the user.

## II. BACKGROUND ON BAYESIAN STATISTICS

An algorithm that automatically penalizes higher-dimensional models is the Reverse Jump Markov Chain Monte Carlo (RJCMCMC) algorithm. The RJCMCMC method [13–17] is widely used when dealing with nested models, meaning that we need to compare a set of models, where simpler models are a subset of a more complicated one. In fact, the RJCMCMC algorithm is the generalized case of Markov Chain Monte Carlo (MCMC) methods that is capable of sampling the parameter space and at the same time jumping between models with different dimensionality.

At this point it would be convenient to describe the philosophy of the Bayesian framework and the Metropolis algorithm and then move to the general case of the RJCMCMC algorithm. The Bayes rule can be summarized by the following equation:

$$\pi(\vec{\theta}|\vec{y}) = \frac{\pi(\vec{y}|\vec{\theta})p(\vec{\theta})}{\pi(\vec{y})}, \quad (1)$$

where  $\vec{\theta}$  is a model parameter vector,  $\pi(\vec{y}|\vec{\theta})$  is the *likelihood* of the parameters  $\vec{\theta}$  over the data-set  $\vec{y}$ , and  $\pi(\vec{\theta}|\vec{y})$  and  $p(\vec{\theta})$  are the *posterior* and the *prior* distributions of the parameters respectively. Note that the *evidence*  $\pi(\vec{y})$ ,

or *marginal likelihood*, is often neglected in parameter estimation algorithms, as it serves only as a normalisation constant:

$$\pi(\vec{\theta}|\vec{y}) \propto \pi(\vec{y}|\vec{\theta})p(\vec{\theta}). \quad (2)$$

The Metropolis algorithm is one of the MCMC-type methods available that is widely used for parameter estimation purposes. It is based on sampling the parameter space by proposing new samples  $\vec{\theta}_n$  and evaluating the likelihood at each step. By sampling the posterior distribution, probability distributions to the parameters to be estimated are assigned. It is certain that given a large amount of steps in the parameter space, the Metropolis algorithm will converge to the set of parameters  $\vec{\theta}_{\text{MAP}}$  that maximise the likelihood.

### A. Calculating the Bayes Factor

The evidence of a hypothesis X given the data-set  $\vec{y}$ ,  $\pi_X(\vec{y})$ , states the support for this hypothesis, or in other words, how much the data favors a given model. In our case, the hypothesis is the model implemented to describe the LTP system. Consequently, given two different models X and Y, the Bayes Factor  $B_{XY}$  is a comparison between the evidences of model X and model Y given by their ratio

$$B_{XY} = \frac{\pi_X(\vec{y})}{\pi_Y(\vec{y})}, \quad (3)$$

where

$$\pi_K(\vec{y}) = \int \pi(\vec{\theta}_K, \vec{y}) d\vec{\theta}, \quad K = X, Y. \quad (4)$$

This integral is extremely costly to evaluate, specially when the model becomes complicated with higher dimensionality. In the next sections, a selection of estimators of the Bayes Factor are described. Most of them evaluate directly the Bayes factor, while other techniques are used to estimate the evidence for each model. All of them are used to investigate our LPF models and mission experiments. In the end, we can draw conclusions about the competing models given the estimated value of the Bayes Factor. If  $B_{XY} < 1$ , the evidence is negative and the observations support model Y. If  $B_{XY} > 1$ , the evidence is positive and model X is more favourable than model Y.

### B. The Reversible Jump Markov Chain Monte Carlo Algorithm

The RJCMCMC algorithm is a robust and efficient tool to estimate the Bayes Factor. It can be shown [13, 15] that after a large number of iterations it will converge to the true value of the Bayes Factor. The only drawback is the computational cost of the algorithm. When more than three models are being compared, meaning that

many transdimensional moves have to be performed, a considerable amount of time is required for convergence. The algorithm implemented in this work is a special case of the *Metropolized Carlin and Chib* method [14]. More specifically, let us suppose that we have a total number  $K$  of models to compare given a data set  $\vec{y}$ . Then, the recipe for our RJMCMC method can be summarised in the following steps:

1. Initialization: Choose an initial model  $k$  and the corresponding parameters  $\vec{\theta}_k$ .
2. Apply the Metropolis algorithm for model  $k$ . This step is also called the “in model step”.
3. Generate new  $\vec{\theta}_{k'}$  from a multivariate Gaussian PDF and a random number  $\rho \in [0, 1]$  from a uniform distribution. This is the step where we propose a new model  $k'$ .
4. Calculate the acceptance ratio  $\alpha'$ :

$$\alpha' = \min \left[ \frac{\pi(\vec{y}|\vec{\theta}_{k'})p(\vec{\theta}_{k'})g(u_{k'})}{\pi(\vec{y}|\vec{\theta}_k)p(\vec{\theta}_k)g(u_k)} |\mathbf{J}|, 1 \right], \quad (5)$$

where  $g(u)$  is the proposal distribution from where the “dimension matching” parameters  $u$  are drawn [18], and  $|\mathbf{J}|$  is the Jacobian:

$$|\mathbf{J}| = \left| \frac{\partial(\vec{\theta}_{k'}, u_{k'})}{\partial(\vec{\theta}_k, u_k)} \right|. \quad (6)$$

5. If  $\rho < \alpha'$  we accept the new model  $k'$  with parameters  $\vec{\theta}_{k'}$  and set  $\vec{\theta}_k = \vec{\theta}_{k'}$ .
6. Iterate from step 2 until convergence.

The old set of parameters is connected to the new one by a well defined function  $\vec{\theta}_k = q(\vec{\theta}_{k'}, u)$  (and of course  $\vec{\theta}_{k'} = q'(\vec{\theta}_k, u')$ ). We use *independent* proposals, so  $\vec{\theta}_k = q(\vec{\theta}_{k'}, u) = u$  and  $\vec{\theta}_{k'} = q'(\vec{\theta}_k, u') = u'$ , thus, the Jacobian term in equation (5) is unity. The algorithm spends most of the time iterating “inside” the model that best describes the data. The RJMCMC method auto-penalizes high dimension models, also by taking into account the priors  $p(\vec{\theta}_k)$  of each model  $k$ . They serve as an Occam Factor integrated within the algorithm. Convergence is achieved if two main conditions are satisfied. First, the condition of *reversibility*, which is stated in a simple way: The proposal function must be invertible, meaning that we can jump from the proposed parameters back to the current parameters. And second, we must satisfy the *dimension matching* condition which in our case is always true since we use independent proposals in the acceptance ratio. After convergence has been

achieved, a good approximation to the Bayes Factor is given by [15, 19–21]

$$B_{XY} = \frac{\# \text{ of iterations in model X}}{\# \text{ of iterations in model Y}}. \quad (7)$$

### C. Other approximations

While the RJMCMC method directly estimates the Bayes Factor, the other methods implemented in our work make an approximation to the evidence of each model. We implemented them as a cross-check for the RJMCMC method, but they also provide certain freedom of choice, depending on the nature of the problem, as well as the data available.

The first approximations we consider are the Laplace approximations. The Laplace approximations perform a comparison between the volume of the models in the parameter space and the volume of the uncertainty ellipsoid of the parameters [9]. This comparison is feasible if we assume that we work within a high Signal-to-Noise Ratio (SNR) and therefore the posterior PDF is Gaussian near the maximum likelihood parameters,  $\vec{\theta}_{\text{MAP}}$ . Another requirement is that a sufficient large number of samples must have been collected. Then, the evidence in Eq. (4) becomes:

$$\pi_X(\vec{y}) \simeq (2\pi)^{D_X/2} |\mathbf{H}|^{1/2} \pi(\vec{\theta}_{\text{MAP},X}, |\vec{y}), \quad (8)$$

where  $D_X$  is the number of dimensions of model X and  $\mathbf{H}$  is the Hessian matrix of the posterior. There are two main variations of the Laplace approximation. In the first one we make use of the Fisher Information Matrix  $\mathbf{F}$ , calculated at  $\vec{\theta}_{\text{MAP}}$ , as an approximation to the expected covariance matrix [21, 22]. Then, the evidence of the model becomes:

$$\pi_X(\vec{y}) \simeq (2\pi)^{D_X/2} |\mathbf{F}|^{-1/2} \pi(\vec{\theta}_{\text{MAP},X}, |\vec{y}) \quad (9)$$

Of course, the main limitations of this method are associated with the confidence we have on the calculation of the Fisher Matrix. Furthermore, as expected, the results appear to be poorer in comparison with the other methods as we move towards lower SNR areas. We can follow the notation of [21] and call this particular approximation the Laplace-Fisher (LF) approximation. Another well-known variation is the Laplace-Metropolis (LM) estimator of the marginal likelihood [9]. In this case, we use all necessary components for the calculation of the evidence from previous MCMC estimates. The parameters  $\vec{\theta}_{\text{MAP}}$  are extracted from the chains of a MCMC parameter estimation run for the particular model, while we use the weighted covariance matrix of the chains  $\Sigma$ , using a Minimum Volume Ellipsoid (MVE) or a Minimum Covariance Determinant estimator (MCD) [23]. The MVE method has been implemented and integrated in the LT-PDA toolbox. In this case, the evidence of model X can be written as

$$\pi_X(\vec{y}) \simeq (2\pi)^{D_X/2} |\Sigma|^{1/2} \pi(\vec{\theta}_{\text{MAP},X}, |\vec{y}). \quad (10)$$

The LM method is considered to be a very reliable tool for the computation of the evidence of a model [9]. The third method we have used is the Schwarz-Bayes Information Criterion (SBIC) and is based on the following: After the assumption that the priors for each model follow a multivariate Gaussian PDF, the Schwarz criterion is defined as:

$$S \simeq \ln(\pi(\vec{\theta}_{\text{MAP},X}, |\vec{y}|)) - \ln(\pi(\vec{\theta}_{\text{MAP},Y}, |\vec{y}|)) - 1/2(D_X - D_Y)\ln(n), \quad (11)$$

where  $D_X$  and  $D_Y$  are the dimensions of each model and  $n$  is the number of samples in the data. It can be proven that if  $n \rightarrow \infty$ , then  $S \rightarrow B_{XY}$  and the Schwarz criterion can be a good approximation to the Bayes factor. In fact,  $n$  must be chosen carefully so that  $n = N_{\text{eff}}$ , where  $N_{\text{eff}}$  is the number of effective samples in the data that represent the growth of the Hessian matrix of the log-likelihood [9].

#### D. Implementation

If the system's transfer function is  $\mathbf{H}_s(\vec{\theta})$  and we consider known injection signals,  $\vec{x}$ , then the output of the system,  $\vec{y}$ , can be expressed as:

$$\vec{y} = \mathbf{H}_s(\vec{\theta})\vec{x} + \vec{n}, \quad (12)$$

where  $\vec{n}$  is the overall instrument noise and  $h(\vec{\theta}) = \mathbf{H}_s(\vec{\theta})\vec{x}$  is the response or template of the system. The noise in the LTP experiment can be approximated with Gaussian noise since we work in a high SNR regime and we can safely consider that any dependence on the system parameters is negligible. Consequently, the likelihood for a model with parameters  $\vec{\theta}$  can be written as

$$\pi(\vec{y}|\vec{\theta}) = C \exp\left[-\frac{1}{2}\langle \vec{y} - h(\vec{\theta}) | \vec{y} - h(\vec{\theta}) \rangle\right], \quad (13)$$

where the angular brackets denote the noise weighted inner product [24]

$$\langle \vec{a} | \vec{b} \rangle = 2 \int_0^\infty \left[ \tilde{a}^*(f)\tilde{b}(f) + \tilde{a}(f)\tilde{b}^*(f) \right] / S_n(f), \quad (14)$$

and where  $S_n(f)$  is the power spectral density of the noise. For the proposal distribution appearing in the acceptance ratio of the Metropolis-Hastings algorithm a multivariate Gaussian distribution is used. Although the choice of the proposal distribution should not affect the final estimated parameter PDFs, it greatly affects the convergence speed of the chains to the target posterior distribution. The covariance of the multivariate Gaussian is calculated beforehand or during the search phase using the Fisher Information Matrix (FIM) [22]. In the case of the RJMCMC method, the FIM is calculated once for each model before the execution of the algorithm. It is computed numerically or analytically depending on the

model. The LTP models are coded either in State-Space format [25–27], or in pure analytical format in the acceleration domain [28].

The first step in the analysis is to apply a Fourier transformation to the data and estimate the power spectral density of the noise  $\vec{n}$ , i.e.  $S_n(f)$ . As described in [4], the nature of the LTP experiment allows us to combine all the available information from different investigations. This means that we can combine the posterior distribution for each data set for each model. Then, we can apply MCMC analysis strategies to sample the posteriors and make use of any of the criteria described in section II C or sample the joint parameter spaces with the RJMCMC method that was described in section II B. Technically, since the RJMCMC algorithm can be seen as a generalised MCMC method, it performs a parameter search by simultaneously sampling the likelihood and consequently it maps the joint posterior distribution and assigns PDFs to the parameters. The extra point is that we can estimate directly the ratio of the evidences of the models and test the hypothesis made for the data.

### III. MODELING THE LTP EXPERIMENT

As a technology demonstrator of a space-based gravitational-wave observatory, the LPF mission will place two 2 kg Test Masses (TM<sub>1</sub> and TM<sub>2</sub>) in free fall. The goal is to estimate the residual differential acceleration between TM<sub>1</sub> and TM<sub>2</sub> [6]. To minimize all external forces acting on TM<sub>1</sub> along the x-axis, a drag-free control loop [30] has been designed. The coordinates of TM<sub>2</sub> are controlled via electrostatic actuators and the Spacecraft (SC) is controlled by micro-Newton thrusters. More specifically, the main components of the LPF mission are:

- The Gravitational Reference Sensor (GRS) [31]. It consists of the test-masses and the vacuum chamber around them. It is mounted with two identical electrostatic actuators to control the 6 degrees of freedom of each test-mass, and also to apply forces and torques to keep the test-masses in free fall.
- The Optical Metrology System [32] consists of the optical bench, its subsystems and the processing computer. It performs the sensitive optical measurements of the positions of the test-masses along the x-axis. For the simplified one-dimensional model version, we consider two interferometer inputs  $\vec{i}$  and two outputs  $\vec{o}$  of the system. The measured displacements are  $x_1$  and  $x_{12}$ , the distance of TM<sub>1</sub> to the SC and the distance between TM<sub>1</sub> and TM<sub>2</sub> respectively.
- The current design of the propulsion system for LPF is based on Cold Gas [2] thrusters. The main function of the thrusters is to maintain the reference test-masses in free fall conditions in the measured bandwidth.

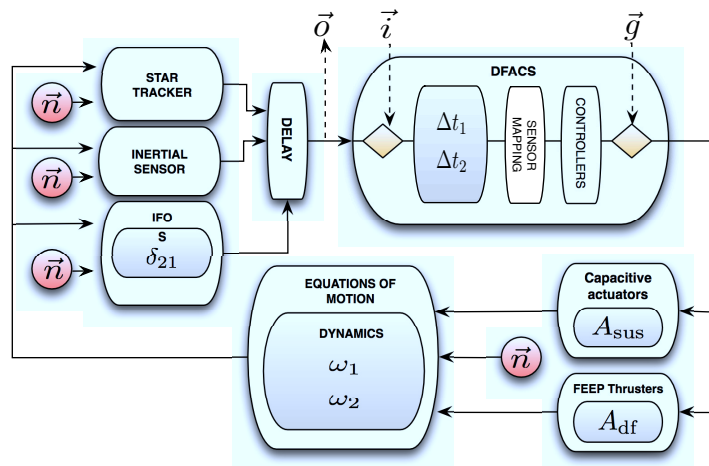


FIG. 1: Schematics of a simplified LTP State-Space Model (SSM). It is composed by smaller SSMs of the various subsystems, each one a standalone SSM, here represented with white boxes. The rhombi represent the main injection inputs to the system. We use mainly the first injection port, which represents the interferometric inputs to the system, while the second rhombus stands for the capacitance actuators injection ports. The symbol  $\vec{n}$  represents noise contributions from various sources and finally, the parameters to fit, for the sake of convenience, are located in the gray boxes inside of the respective SSMs. The interferometer signals  $\vec{i}$  are injected to the controllers (DFACS), where the commanded forces are generated and applied through the capacitance actuators and the thrusters of the space-craft. In the last two of experiments of [29] “out of loop” forces  $\vec{g}$  are applied to the three bodies (TMs and SC) of the system. The resulting movement of the bodies is monitored by the Inertial Sensor, the Star Tracker and the interferometer. Here, the interferometer output is denoted as  $\vec{o}$ .

- Finally, the Drag-Free and Attitude Control System (DFACS) calculates all forces and torques acting on the SC and test-masses and computes the commanded forces/response of the system in order to maintain the nominal free fall of  $TM_1$ .

The LTP is a complicated instrument composed by a plethora of subsystems and coupled control loops. During the mission, there will be dedicated experiments with the aim of characterizing the different noise contributions and couplings. In the following section, the LTP models implemented for system identification experiments are described. We present simulations of the planned system identification experiments and use all techniques described to perform model selection.

### A. The LTP system dynamics

For the purposes of this work we consider a simplified version of the LTP operating in the so-called main science mode and we confine the system dynamics to the one-dimensional case. There are two approaches in modeling the dynamics of the three-body system: The first one, shown in Fig. 1, is to fully describe the dynamics together with the controllers in a state-space format (SSM). In this case, it is very convenient to represent the model as a closed loop system, due to closed-loop dynamics and controllers of the instrument. The white boxes represent standalone SSMs, all together *assembled* to form the final LTP experiment [25]. Each submodule has injection  $u(t)$

and output  $y(t)$  ports and can be represented as:

$$\begin{aligned} \dot{x}(t) &= A \times x(t) + B \times u(t), \\ y(t) &= C \times x(t) + D \times u(t), \end{aligned} \quad (15)$$

where  $x(t)$  are the *states* of the system,  $A$  is the state matrix,  $B$  is the input matrix,  $C$  is the output matrix, and  $D$  is the feed-through matrix [25]. This form of implementation has numerous advantages, like modularity and flexibility in modeling the LPF mission. For example, the user is able to combine any given SSM module and use a custom noise model for each particular subsystem.

The second approach is to represent the dynamics in an analytical way in the acceleration domain [28] which is described in section III C. A more in-depth investigation on the analysis in acceleration domain and its practical advantages is to appear soon.

For the purposes of this paper, we simulate the experiments as explained in [4, 5, 25, 29, 33]. Then, the complete system is controlled by optical readouts that measure the positions of the Test Masses,  $x_1$  and  $x_{12}$ , where the  $x$ -axis is defined by the line joining them. In order to perform parameter estimation we inject sinusoidal signals to each interferometric channel alternately, or direct forces to the three bodies of the system (test-masses and SC), and measure the response of the system. For this type of experiments, we can define the following time-series arranged as vectors:

$$\vec{o} = \begin{pmatrix} o_{x_1} \\ o_{x_{12}} \end{pmatrix}, \quad \vec{i} = \begin{pmatrix} i_{x_1} \\ i_{x_{12}} \end{pmatrix}, \quad \vec{n} = \begin{pmatrix} n_{x_1} \\ n_{x_{12}} \end{pmatrix}, \quad (16)$$

$$\vec{g} = \begin{pmatrix} g_1 - g_{SC} \\ g_2 - g_1 \end{pmatrix}, \quad (17)$$

where  $\vec{o}$  is the measured signals vector,  $\vec{i}$  is the interferometer injection signals vector,  $\vec{g}$  is the commanded forces signals vector, and  $\vec{n}$  is the overall noise for both channels. Here  $g_1$  and  $g_2$  are the commanded forces applied to  $TM_1$  and  $TM_2$  respectively while  $g_{SC}$  is the force applied to the SC. The noise contributions of each LTP subsystem are denoted in Fig. 1 as  $\vec{n}$ . It can be instrumental or read-out noise, or it may originate from external sources (like solar radiation).

For the sake of simplicity, we use the one-dimensional version of the models. The set of parameters that we recover from the system is  $\{\omega_1, \omega_2, \Delta t_1, \Delta t_2, A_{df}, A_{sus}, \delta_{21}\}$ :  $\omega_1$  and  $\omega_2$  stand for the electrostatic stiffness of the test-masses to the surroundings,  $\Delta t_1$  and  $\Delta t_2$  are time delays,  $A_{df}$  and  $A_{sus}$  the capacitance and thruster gains respectively, and  $\delta_{21}$  denotes the cross-coupling between the two interferometric channels.

As described in [29, 33], during the planned mission experiments, sinusoidal signals of different frequencies are going to be injected to the control loop, simulating displacements of the test masses. Besides the LTP dynamics system, the data analysis team has modeled noise sources for the various subsystems of the LTP mission based on theoretical predictions or characterization of each element from on-ground test campaigns. The main noise contributions come from the thrusters, the interferometric readouts, and the capacitance actuators [34]. We can simulate synthetic noise for any particular experiment and perform parameter estimation exercises in order to test the system identification algorithm's readiness. A MCMC search of the parameter space for such a simulated run returns satisfactory results as shown in Table I.

TABLE I: MCMC parameter estimation results from simulated experiments. The experiments being analyzed here, are the two first described in [4] and [29].

Parameter	Real value	Estimated $\pm \sigma$
$\omega_1$	$1.3 \times 10^{-6}$	$(-1.2999 \pm 0.0002) \times 10^{-6}$
$\omega_2$	$1.9 \times 10^{-6}$	$(-1.8999 \pm 0.0002) \times 10^{-6}$
$A_{df}$	0.82	$0.8201 \pm 0.00025$
$A_{sus}$	1.08	$1.080004 \pm 3 \times 10^{-6}$
$\Delta t_1$	0.2	$0.20020 \pm 5 \times 10^{-5}$
$\Delta t_2$	0.2	$0.2003 \pm 2 \times 10^{-4}$
$\delta_{21}$	0.0004	$0.0004001 \pm 1 \times 10^{-7}$

By using the LTPDA machinery, synthetic noise can be generated and the response of a known system can be simulated. For the following study cases, the LTP default parameters were always set to the same values given in Table I. It should be noted that there is no reason to have any strong prior belief about the true values of the parameters. Therefore, it is reasonable to assume uniform priors for all the tests.

## B. Application to a simplified LTP model

In order to demonstrate a first application to the LTP, we can investigate a model selection case that was first encountered during a data analysis exercise [33]. This type of exercises are organized by the DA team with the aim of testing the developed tools in more realistic scenarios. In this case, the ‘‘true’’ values of the parameters of the system to identify, were totally unknown. In that situation, the data analysis team was able to perform parameter estimation, but it was noted that the fit was improved when two extra parameters were introduced. These parameters are the guidance delays,  $\Delta t_1$  and  $\Delta t_2$ . They are simply time delays to the application of the guidance signals (see Fig. 1), due to operations of the Data Management Unit (DMU), which is the on-board computer controlling the LTP experiment. The final robustness of the fit indicates that the new parameters substantially improve the model.

This problem can be better studied by reducing it to a model selection problem and can be tested with synthetic data-sets. The simulated data source is a LTP system with the default characteristics given in Table I with the exception of the delays, which were set to:  $\Delta t_1 = \Delta t_2 = 0$ . With this configuration, we assume a true system where the application signals are applied instantaneously. For our first try-out we injected ‘‘fake’’ interferometric displacements, while for the second investigation we used the same structure of injections but with much lower Signal to Noise Ratio (SNR  $\sim 5$ ).

In order to determine the importance of the two extra parameters, we have to verify which model describes better the particular data-set: The seven parameter model X with parameters  $\vec{\theta}_X = \{\omega_1, \omega_2, \Delta t_1, \Delta t_2, A_{df}, A_{sus}, \delta_{21}\}$ , or the five parameter model Y with parameters  $\vec{\theta}_Y = \{\omega_1, \omega_2, A_{df}, A_{sus}, \delta_{21}\}$ . While both models, X and Y, are capable of explaining the observations, we expect the simpler model Y to be more favorable from a RJMCMC output, since the extra parameters  $\Delta t_1$  and  $\Delta t_2$  are not significant for the data. The evolution of the Bayes factor for such an investigation for two different levels of SNR is shown in Fig. 2.

TABLE II: Results for the Guidance delay investigation with the low SNR experiment. See text for details.

Method	$B_{XY}$
RJMCMC	0.309
LF	0.124
LM	0.078
SBIC	0.768

The results we obtain for all the approximations verify that the simpler model Y is much more probable than the more complicated model X. For the particular experiments proposed in [29] where the SNR is high, the Bayes factor computed tends to zero. In fact, for the case of the




FIG. 2: (Color online) First  $3 \times 10^4$  iteration of the RJMCMC output when comparing a seven and a five-dimensional LTP model (models X and Y respectively). Since the models are not competitive when the SNR = 60, the blue line tends asymptotically to zero.

RJMCMC method, there is no single iteration “inside” the more complex model. This changes dramatically depending on the nature of the problem and, of course, as we show in section IV A, on the SNR. In Table II we present only the low-SNR experiment, for the sake of comparison between the methods. Each method seems to favor the simpler model but they are not in total agreement between them. This is to be expected, as the SNR of this investigation is very low and the approximations of the evidence become more sensitive. For this particular case of the injections, the models are not competitive and therefore, the resulting estimated Bayes Factor is extremely small. For a direct comparison, the RJMCMC algorithm requires more than  $10^8$ - $10^9$  iterations.

### C. More realistic applications

During the last years a series of data analysis exercises have been performed. These so called Operational Exercises are organized by the LPF Data Analysis team with the aim of training the scientists and engineers for the upcoming flight operations. Data generated by the ESA simulator are used and real-life shifts and duties are rehearsed. The ESA simulator for the LPF is an industry developed software that is considered the most accurate simulator to date. The scientific and organizational results are valuable since they reproduce situations and data analysis challenges that are expected to occur during flight operations.

An outcome of these data challenges was the realiza-

tion of the necessity to change or manipulate the model of the dynamics of the system shortly after receiving telemetry from the satellite, thus facing once again model selection problems. Expressing the dynamics in the acceleration domain [28], one could write a simplified model for the differential acceleration, like the one below:

$$a_{12} = \left[ \frac{d^2}{dt^2} + \omega_2^2 \right] x_{12} + (\omega_2^2 - \omega_1^2) x_1 - AF_2 + AF_1 \quad (18)$$

where  $a_{12}$  is the differential acceleration and  $F_1$  and  $F_2$  denote the applied forces on the first and second test-masses respectively. The real motion  $\vec{x}$  of the test-masses can be approximated by the delayed interferometer read-outs:

$$\vec{x} = \begin{pmatrix} x_1 \\ x_{12} \end{pmatrix} = \begin{pmatrix} o_1(t - \tau_{ifo}) \\ o_{12}(t - \tau_{ifo}) \end{pmatrix} \quad (19)$$

The parameters appearing in Eq. (18) and (19) are the stiffnesses of the two test-masses  $\omega_1$  and  $\omega_2$ , the interferometer read-out delay  $\tau_{ifo}$ , and the gains of the capacitance actuators  $A_{sus}$  (here represented as  $A$  for simplicity). Note that for this first approximation we have assumed identical actuators  $A$  for both test-masses. The commanded forces  $g_1$  and  $g_2$  are available as telemetry, but the *real* applied forces on the three bodies are to be determined by the measurements and the analysis itself. For example, in the real data-stream there might be additional delays or even filtering of the applied forces coming from the controllers, so in reality a gain  $A$  might be proven to be frequency dependent:  $A(f)$ . This situation will appear in the received telemetry and we need the means to disentangle those two physical effects in a quantitative way.

For the particular simulated data-set, the model was not able to remove all the injected signals and even for the simple case of Eq. (18), this led to a biased estimate for the parameters. Apart from a simple time delay on the commanded signals, there might be another process that causes a difference between the commanded  $g_1$  and  $g_2$  and the actual applied forces  $F_1$  and  $F_2$ . For a first approximation of such a process, we assume a single real pole filter, filtering the time-series of the applied force on the second test-mass. For this investigation we can propose two models where  $AF_i(t) = Ag_i(t - \tau_C)$  and  $AF_i = A(\frac{f_o}{f - jf_o})g_i$  respectively. Here  $\tau_C$  denotes the actuators time delay. In the end we can apply the model selection methods to these two models: the first one, X, where the applied forces are time delayed, and the second one, Y, where the forces are frequency dependent (filtered by a single real pole filter). The calculated Bayes Factor between those two models is:

$$BF_{XY} = \frac{\pi_X(\vec{y})}{\pi_Y(\vec{y})} = 9.9478 \times 10^{-11}, \quad (20)$$

clearly indicating that the most probable process on the forces is the one described by model Y. The real pole was




FIG. 3: (Color online) Power spectra of simulated differential acceleration between the two test-masses. The gray curve represents the reference noise measurement, while the light blue curve is the differential interferometer readout. The value of the computed Bayes Factor can be confirmed with the comparison of the equivalent estimated residual acceleration for each model. The differential interferometer read-out is also plotted for comparison.

estimated to be  $f_o = 1.963 \pm 0.001$  Hz and was confirmed for the complete set of system identification experiments. The estimated acceleration residuals for both models can be seen in Fig. 3.

The same analysis strategy applies of course to more complicated versions of the analytical models, where the number of parameter increases with the more terms of the equation are added, as happens in the following more realistic case. As already described in section III, the two test masses are controlled with identical actuators. In reality, a small misbalance between the electrostatic actuators surrounding each test mass might be present. If we introduce this “asymmetry” to the system, immediately for the simple case of Eq. (18), we can increase the dimensionality of the model by four parameters:

$$\begin{aligned}
 a_{12} = & \left[ \frac{d^2}{dt^2} + \omega_2^2 \right] x_{12} + (\omega_2^2 - \omega_1^2) x_1 \\
 & - A_2 \left( \frac{f_{o2}}{f - j f_{o2}} \right) g_2(t - \tau_{C2}) \\
 & + A_1 \left( \frac{f_{o1}}{f - j f_{o1}} \right) g_1(t - \tau_{C1}). \quad (21)
 \end{aligned}$$

These parameters are gains, delays and filters that are different for the actuator of each test mass. Theoretically the highest in dimensions model of Eq.(21) can describe the observations, but the problem to solve appears to be over-parametrized. A solution is to generate a set of nested models under the highest in dimensions of Eq. (21) and apply the RJMCMC algorithm. The result of such a run is shown in Fig. 4 and it reveals the most favorable model and consequently the underlying procedure that describe best the physical system. For the par-




FIG. 4: A RJMCMC run on a set of nested LTP models. There is a clear preference for the five-dimensional model for the given data-set. The data were produced with a “perfect” model where the two respective actuators were identical.

ticular simulation we can verify that the five-dimensional model is the best, concluding that no “asymmetry” in the hardware of the LTP is present. This changed in the following Data Challenge, where the Bayes factor between the simple model of Eq. (18) and a six-parameter asymmetric one, is greater than  $10^6$ , clearly supporting the correct higher dimension model where  $A_1 \neq A_2$ .

#### IV. USING THE BAYES FACTOR FOR EXPERIMENT OPTIMIZATION

In this last section we explore the capability of the implemented model selection framework, not only as a set of tools for data analysis purposes, but also as a way to evaluate the efficiency of the experiments we are planning to run in the satellite. As we are going to show, by comparing different models under different input signal conditions, we can safely determine the best range of parameters that define our experiments, or verify the injection frequencies that maximize the information extracted from the system.

##### A. The Bayes Factor as function of the SNR

It has been shown [21, 35] that there is a dependence of the Bayes Factor output on the SNR regime of the investigations. This, of course, holds true in the case of the LTP as it can be seen in Fig. 5. This figure was created by simulating LTP experiments for each value of the SNR, while the injection signals were single frequency






FIG. 5: (Color online) The Bayes Factor as a function of SNR computed using different methods.

( $f = 0.01\text{Hz}$ ) sinusoidal inputs into the system. The LTP models under comparison were quite similar with the exception of a different realization of the response model of the thrusters. It is clear that above the critical value of the  $\text{SNR} = 21$  the results obtained with the different techniques are consistent and in good agreement. Below that value of the SNR we cannot make clear decisions about the competing models, as the wrong model is showing preference, or we poorly approximate the Bayes Factor.

Although this SNR limit varies, as expected, depending on the type of investigation and model, the current result is already providing an estimation of the required power of the injection signals that we need to consider in the LTP experiment. This information and the method used here will be of particular interest for the design of in-flight experiments for the LPF mission.

### B. The Bayes Factor as function of the injection frequencies

Furthermore, for system identification experiments, as in the case of the LTP, the computed evidence of a model depends on the design of the experiment itself. The information obtained from the system differs depending on the injection frequencies. An interesting study is to explore in detail this relation. A four- (X) and a five-dimensional (Y) models are examined, given different injection fre-




FIG. 6: (Color online) The Bayes Factor as a function of the injection frequencies in the first interferometric channel. The computed evidence of model X is stronger when the sinusoidal injection signals have a frequencies around  $f_{i_{x_1}} \simeq 0.0006\text{ Hz}$  and  $f_{i_{x_1}} \simeq 0.05\text{ Hz}$  respectively.

quencies. More precisely, since the difference between the models is the cross-coupling  $\delta_{21}$  as shown in Fig. 1, which describes the signal leakage from the first to the differential interferometric channel, we examine the Bayes Factor given different injection frequencies to the first channel, while keeping constant the injection to the differential channel ( $f_{i_{x_{12}}} = 0.2\text{ Hz}$ ). The SNR of this experiment is kept at the “low” value of 28.

The data generation model is mounted with a “perfect” interferometer ( $\delta_{21} = 0$ ) and model X is the same as the one used to produce the data, while model Y is the one with the extra parameter  $\delta_{21}$ . The expected outcome of this exploration is that if the system is more sensitive to the  $\delta_{21}$  parameter at some particular frequencies, we must detect an increase in the Bayes Factor which underlines a more clear decision towards the correct model.

In Fig. 6 we can see the corresponding Bayes Factor versus the injected frequencies to the first channel. Given the low SNR of the investigation, while model X should be more favorable, a preference for the more complex model Y is shown for a certain set of frequencies. This result is a clear indicator of the set of preferred frequencies that can be injected to the system for its characterization given the current configuration and SNR. Indeed, injections around  $f \simeq 0.0006\text{ Hz}$  and  $f \simeq 0.05\text{ Hz}$  promote the identification of the correct model, while injections at both the high and low frequency limit, together with  $f \simeq 0.01\text{ Hz}$ , may induce the analysis into an error. This frequency dependence must be associated to the sensitiv-

ity of the experiment to a given parameter, the parameter  $\delta_{21}$  in this particular case. This observed dependency, when considering a more realistic model, will be of particular interest in the selection of injection signals for the experiments to be run in-flight.

## V. CONCLUSIONS

We have implemented three different methods to compare competing models of the LTP experiment on-board the LPF mission: The RJMCMC algorithm, the Laplace approximations, and the Bayes Information Criterion. The results from each method seem to be in agreement, but the output strongly depends on the expected SNR and of course on the models under investigation. Considering the LPF mission planned experiments, the SNR is high enough to safely use any of all the available techniques, but probably the most computationally demanding methods will be used for off-line analysis to confirm our first computations.

The RJMCMC algorithm (together with the Laplace methods and the Bayes Information Criterion) employed in this work has been integrated in the LTPDA toolbox as part of the LPF data analysis software. The RJMCMC algorithm is by far the most computationally costly, but at the same time it is the more suitable one when we compare more than two nested models or we work with inputs with low SNR. The Laplace-Metropolis and the Laplace-Fisher methods are reliable when we work in the high SNR regime, but they also require significant computing time, specially when one has to use outlier detection methods to estimate the weighted covariance matrix. On the other hand, the Laplace-Fisher approximation is limited by the use of the Fisher Information Matrix, which for the case of LTP state space models is computed numerically.

Moreover, an attempt to associate the output of the

mentioned methods with the actual system identification experiment has been made. We have used different experiment setups to demonstrate that the Bayes Factor depends not only on the SNR, but also on the injection frequencies to the system.

The developed algorithms were successfully applied to model selection problems for the LPF data analysis for the first time. Two different cases of LTP model selection problems have been investigated over data-sets that were produced by both the LTPDA and the ESA simulator. For the first case, we have considered an easy case of five- and seven-dimensional state-space models, where the importance of the extra two parameters was examined. These two extra parameters are time delays caused by the LPF hardware and they can be characterized as essential parameters of the model. For the second case, we explored the most suitable dimensionality of analytic models. There, the simplest model that described efficiently the observations was recovered, excluding the more complicated ones that caused over-fitting issues. This type of analysis is expected to be performed during operations due to the broad spectrum of possible applications, like identifying external disturbances that result into forces applied to the three-body system.

## Acknowledgments

We acknowledge support from contracts AYA-2010-15709 (MICINN) and 2009-SGR-935 (AGAUR). NK acknowledges support from the FI PhD program of the AGAUR (Generalitat de Catalunya). MN acknowledges a JAE-doc grant from CSIC and support from the EU Marie Curie CIG 322288. GC acknowledges support from the Beecroft Institute for Particle Astrophysics and Cosmology.

- 
- [1] F. Antonucci, M. Armano, H. Audley, G. Auger, M. Benedetti, et al., *Class.Quant.Grav.* **29**, 124014 (2012).
  - [2] P. McNamara, F. Antonucci, M. Armano, H. Audley, G. Auger, M. Benedetti, P. Binétruy, J. Bogenstahl, D. Bortoluzzi, N. Brandt, et al., in *Astronomical Society of the Pacific Conference Series*, edited by G. Auger, P. Binétruy, and E. Plagnol (2013), vol. 467 of *Astronomical Society of the Pacific Conference Series*, p. 5.
  - [3] P. A. Seoane et al. (eLISA Collaboration) (2013), 1305.5720.
  - [4] M. Nofrarias, C. Röver, M. Hewitson, a. Monsky, G. Heinzel, K. Danzmann, L. Ferraioli, M. Hueller, and S. Vitale, *Physical Review D* **82**, 1 (2010), ISSN 1550-7998, URL <http://link.aps.org/doi/10.1103/PhysRevD.82.122002>.
  - [5] G. Congedo, L. Ferraioli, M. Hueller, F. De Marchi, S. Vitale, et al., *Phys.Rev.* **D85**, 122004 (2012), 1108.0862.
  - [6] F. Antonucci, M. Armano, H. Audley, G. Auger, M. Benedetti, P. Binétruy, C. Boatella, J. Bogenstahl, D. Bortoluzzi, P. Bosetti, et al., *Classical and Quantum Gravity* **28**, 094006 (2011), ISSN 0264-9381, URL <http://stacks.iop.org/0264-9381/28/i=9/a=094006?key=crossref.41fc94b94ce025596175f5fe463e5a86>.
  - [7] F. G. Cervantes, R. Flatscher, D. Gerardi, J. Burkhardt, R. Gerndt, M. Nofrarias, J. Reiche, G. Heinzel, K. Danzmann, L. G. Boté, et al., in *Astronomical Society of the Pacific Conference Series*, edited by G. Auger, P. Binétruy, and E. Plagnol (2013), vol. 467 of *Astronomical Society of the Pacific Conference Series*, p. 141.
  - [8] H. Audley, K. Danzmann, A. Garcia Marin, G. Heinzel, A. Monsky, et al., *Class.Quant.Grav.* **28**, 094003 (2011).
  - [9] R. Kass and A. Raftery, *Journal of American Statistical Association* **90** (1995).
  - [10] D. J. C. MacKay, *Information Theory, Inference, and Learning Algorithms* (Cambridge University Press,

- 2003).
- [11] M. Hewitson, M. Armano, M. Benedetti, J. Bogenstahl, D. Bortoluzzi, et al., *Class.Quant.Grav.* **26**, 094003 (2009).
- [12] *LTPDA Matlab toolbox* (2005), URL <http://www.lisa.aei-hannover.de/ltlda/>.
- [13] P. J. Green, *Biometrika* pp. 711–32 (1995).
- [14] H. F. Lopes and M. West, *Statistica Sinica* **14**, 41 (2004).
- [15] P. Dellaportas, J. J. Forster, and I. Ntzoufras, *Statistics and Computing* pp. 27–36 (2002).
- [16] A. Stroeer and J. Veitch, *Phys. Rev. D* **80**, 064032 (2009), URL <http://link.aps.org/doi/10.1103/PhysRevD.80.064032>.
- [17] R. Umstätter, N. Christensen, M. Hendry, R. Meyer, V. Simha, J. Veitch, S. Vigeland, and G. Woan, *Phys. Rev. D* **72**, 022001 (2005), URL <http://link.aps.org/doi/10.1103/PhysRevD.72.022001>.
- [18] S. J. Godsill, *JOURNAL OF COMPUTATIONAL AND GRAPHICAL STATISTICS* **10**, 230 (2001).
- [19] T. B. Littenberg and N. J. Cornish, *Phys. Rev. D* **80**, 063007 (2009), URL <http://link.aps.org/doi/10.1103/PhysRevD.80.063007>.
- [20] F. Bartolucci, L. Scaccia, and A. Mira, *Biometrika* **93**, 41 (2006), ISSN 0006-3444.
- [21] N. Cornish and T. Littenberg, *Physical Review D* **76**, 1 (2007), ISSN 1550-7998, URL <http://link.aps.org/doi/10.1103/PhysRevD.76.083006>.
- [22] M. Vallisneri, *Physical Review D* **77**, 1 (2008), ISSN 1550-7998, URL <http://link.aps.org/doi/10.1103/PhysRevD.77.042001>.
- [23] P. J. Rousseeuw and M. Leroy Annick, *Robust Regression and outlier detection* (John Wiley and Sons, 1987).
- [24] L. Barack and C. Cutler, *Phys.Rev.* **D69**, 082005 (2004), [gr-qc/0310125](http://arxiv.org/abs/gr-qc/0310125).
- [25] M. Nofrarias, F. Antonucci, M. Armano, H. Audley, G. Auger, M. Benedetti, P. Binétruy, J. Bogenstahl, D. Bortoluzzi, N. Brandt, et al., in *Astronomical Society of the Pacific Conference Series*, edited by G. Auger, P. Binétruy, and E. Plagnol (2013), vol. 467 of *Astronomical Society of the Pacific Conference Series*, p. 161.
- [26] M. Hewitson, M. Diaz-Aguilo, and A. Grynagier, *Tech. Rep. S2-AEI-TN-3069* (2011).
- [27] A. Grenadier and M. Weyrich, *Ph.D. thesis* (2008).
- [28] G. Congedo, R. Dolesi, M. Hueller, S. Vitale, and W. J. Weber, *Phys. Rev. D* **88**, 082003 (2013), URL <http://link.aps.org/doi/10.1103/PhysRevD.88.082003>.
- [29] S. Vitale, *Tech. Rep. S2-UTN-TN-3045* (2007).
- [30] W. Fichter, P. Gath, S. Vitale, and D. Bortoluzzi, *Classical and Quantum Gravity* **22**, S139 (2005), ISSN 0264-9381, URL <http://stacks.iop.org/0264-9381/22/i=10/a=002?key=crossref.87240f1d344f89a51d67135cf12a4c33>.
- [31] R. Dolesi, D. Bortoluzzi, P. Bosetti, L. Carbone, A. Cavalleri, et al., *Class.Quant.Grav.* **20**, S99 (2003).
- [32] G. Heinzl, C. Braxmaier, M. Caldwell, K. Danzmann, F. Draaisma, et al., *Class.Quant.Grav.* **22**, S149 (2005).
- [33] M. Nofrarias, L. Ferraioli, G. Congedo, M. Hueller, M. Armano, et al., *J.Phys.Conf.Ser.* **363**, 012053 (2012), 1111.4916.
- [34] F. Antonucci, M. Armano, H. Audley, G. Auger, M. Benedetti, P. Binétruy, J. Bogenstahl, D. Bortoluzzi, P. Bosetti, N. Brandt, et al., *Classical and Quantum Gravity* **29**, 124014 (2012), URL <http://stacks.iop.org/0264-9381/29/i=12/a=124014>.
- [35] J. Veitch and A. Vecchio, pp. 1–10 (2008), [arXiv:0801.4313v3](http://arxiv.org/abs/0801.4313v3).