

# Free-flight experiments in LISA Pathfinder

JI Thorpe<sup>1</sup>, CJ Cutler<sup>2</sup>, M Hewitson<sup>3</sup>, O Jennrich<sup>4</sup>, P Maghami<sup>1</sup>,  
S Paczkowski<sup>3</sup>, G Russano<sup>5</sup>, S Vitale<sup>5</sup> and WJ Weber<sup>5</sup> for the LPF  
Team

<sup>1</sup>Gravitational Astrophysics Lab, NASA Goddard Space Flight Center, 8800 Greenbelt Road,  
Greenbelt, MD USA 20771

<sup>2</sup>NASA Jet Propulsion Laboratory, 4800 Oak Grove Drive, Pasadena, CA 91109

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

<sup>4</sup>European Space Technology Centre, European Space Agency, Keplerlaan 1, 2200 AG  
Noordwijk, The Netherlands

<sup>5</sup>Dipartimento di Fisica, Università di Trento and Trento Institute for Fundamental Physics  
and Application / INFN, 38123 Povo, Trento, Italy

E-mail: james.i.thorpe@nasa.gov

**Abstract.** The LISA Pathfinder mission will demonstrate the technology of drag-free test masses for use as inertial references in future space-based gravitational wave detectors. To accomplish this, the Pathfinder spacecraft will perform drag-free flight about a test mass while measuring the acceleration of this primary test mass relative to a second reference test mass. Because the reference test mass is contained within the same spacecraft, it is necessary to apply forces on it to maintain its position and attitude relative to the spacecraft. These forces are a potential source of acceleration noise in the LISA Pathfinder system that are not present in the full LISA configuration. While LISA Pathfinder has been designed to meet its primary mission requirements in the presence of this noise, recent estimates suggest that the on-orbit performance may be limited by this ‘suspension noise’. The drift-mode or free-flight experiments provide an opportunity to mitigate this noise source and further characterize the underlying disturbances that are of interest to the designers of LISA-like instruments. This article provides a high-level overview of these experiments and the methods under development to analyze the resulting data.

## 1. Introduction

The basic operating principle of an interferometric gravitational-wave detector is the measurement of fluctuations in space-time curvature via the exchange of photons between pairs of geodesic-tracking references separated by large baselines [1]. A key challenge for implementing such a detector is the development of an object whose worldline approximates a geodesic – an *inertial* particle. The LISA Pathfinder (LPF) mission [2] will validate the technology of drag-free test masses for use as inertial references in a future space-based gravitational wave detector such as the Laser Interferometer Space Antenna (LISA).

The technique of drag-free flight for disturbance reduction [3, 4] can be briefly summarized as follows. A reference or ‘test’ mass is placed inside a hollow housing within the host spacecraft (SC). On orbit, the test mass is allowed to float freely inside the housing while a sensor system monitors the position and attitude of the test mass relative to the SC. This information is used

by a control system which commands the SC to follow the orbit of the test mass. In doing so, the SC isolates the test mass from external disturbances. By actuating the SC rather than the test mass (as is done in a traditional inertial guidance system) to maintain the relative position and attitude, the test mass is isolated from noise associated with the actuation itself. The remaining residual acceleration noise of the test mass results from forces local to the SC as well as external forces that are not absorbed by the SC (e.g. magnetic).

In the LPF implementation [5], two 46 mm cubic Au-Pt test masses are contained in a single SC. One of the test masses is designated as the reference test mass (RTM) and the SC performs drag free flight along one linear degree of freedom (DoF) while controlling the remaining five DoFs (three angular and two linear) with an electrostatic suspension system. The second or non-reference test mass (NTM) is located  $\sim 38$  cm away from the primary test mass along the drag-free axis and is used as a witness to assess the residual acceleration of the RTM. An interferometric metrology system [6], monitors relative displacements between the RTM and NTM with a precision of  $\sim 10$  pm/ $\sqrt{\text{Hz}}$ .

A key difference between the LPF configuration and the LISA configuration is that the NTM must be electrostatically suspended along the sensitive DoF because the SC cannot simultaneously follow the trajectories of both the RTM and NTM along the same DoF. This suspension force has a noise component which represents a disturbance in the LPF measurement that is not present in LISA-like configurations. With all noise sources at the design requirement levels for LPF, this residual suspension noise contribution is not a significant contribution to the overall measured differential acceleration noise in the LISA measurement band of  $0.1 \text{ mHz} < f < 100 \text{ mHz}$  (see Figure 1a). The dominating term over this band are various types of spurious forces on the test masses, precisely the phenomena of interest to the designers of future gravitational wave instruments.

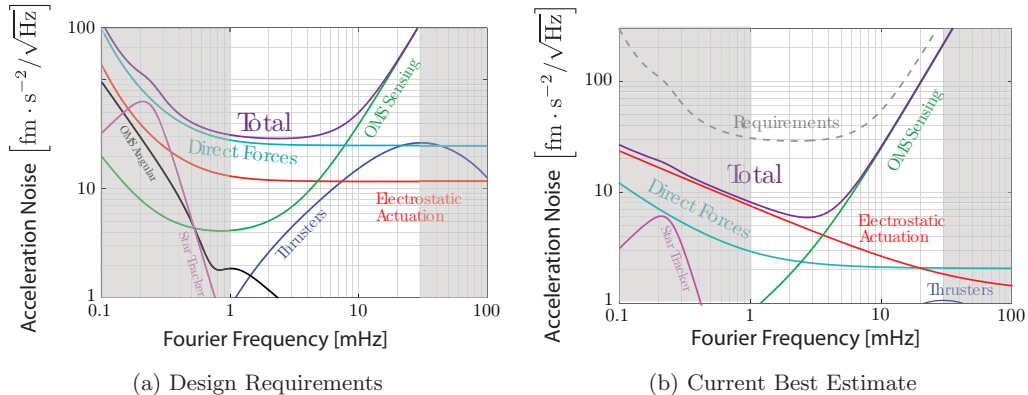


Figure 1: Breakdown of noise sources for measurement of acceleration of the reference test mass relative to the non-reference test mass along the sensitive DoF. The left panel shows the design requirement levels whereas the right panel shows the current best estimates based on ground test campaigns of flight hardware and system modeling.

Figure 1b shows the current best estimate (CBE) for the differential acceleration noise in LPF. As would be expected from a conservative set of design requirements, the CBE levels for all terms are lower than the corresponding ones for the requirements. However, the dominating term in the measurement band is now expected to be suspension noise. This expectation is supported by stringent limits placed on the magnitudes of unknown or unmodeled forces on the LPF inertial sensor using torsion pendulums [7, 8].

To be clear, LPF performance at either the requirement or CBE level would accomplish the

goal of validating drag-free flight as a technique for realizing inertial reference sensors for LISA-like observatories. Nevertheless, a direct measurement of the forces acting on the test masses would provide additional valuable information to the designers of such observatories.

## 2. Free-Flight Experiments

The electrostatic suspension of the NTM in LPF is needed to counteract forces in the NTM-SC system that differ from those in the RTM-SC system, which are suppressed by the drag-free control loop. These include disturbance forces that are local to the test masses, such as residual gas disturbances, thermal disturbances, electrostatic forces, magnetic forces, etc. In addition, the static gravitational gradient differs at the RTM and NTM location, leading to a constant bias in the force that must be applied to the NTM along the sensitive axis. Although this bias will be minimized through the use of compensation masses designed based on pre-flight gravitational models, it is expected that residual gravitational accelerations along the sensitive DoF will be on the order of  $10^{-10}$  m/s<sup>2</sup>. This sets the amplitude, and consequently the noise, of the required suspension force.

In the *drift mode* or *free-flight* experiment, the compensation of the static field experienced by the NTM is performed with a series of discrete ‘kicks’ rather than with a continuously-applied force. Between the kicks, the electrostatic actuation of the NTM along the sensitive DoF is turned off, allowing the NTM to drift under the influence of the static forces.

### 2.1. The LTP Drift Mode Experiment

The drift mode controller designed for the LISA Test Package (LTP)[9] fixes both the length of the drift period ( $\sim 350$  s) and the duration of the kicks ( $\sim 1$  s). A Kalman-filter based observer tracks the motion of the NTM during the free flight and estimates the impulse required to maintain the NTM position. The amplitude of the subsequent kick is adjusted to deliver this impulse and the process is repeated. Figure 2a shows the displacement of the RTM relative to the NTM for a segment of simulated data from the LTP drift mode. It consists of a series of repeated quasi-parabolic flights (the trajectory of the NTM is not a true parabola due to the influence of other force terms such as a linear spring term that couples the NTM to the SC) with a duration of 350 s. Figure 2b shows the corresponding acceleration of the RTM relative to the NTM, which shows a series of discrete kicks separated by nearly constant acceleration during the free flight segments.




Figure 2: Example data for a LTP drift-mode simulation. The left panel shows the displacement of the RTM relative to the NTM whereas the right panel shows the acceleration.

## 2.2. A proposed ST7-DRS Free-Flight Experiment

An alternative drift-mode control design has recently been developed for possible use with the NASA-provided payload known as the Disturbance Reduction System (DRS)<sup>1</sup>. The ST7 design is based on a modified dead-band controller, where the dynamical state of the system determines the control mode. The two-dimensional phase space of NTM position and NTM velocity is used to describe the system state and a region centered around the nominal position and zero velocity is designated as the dead band. So long as the NTM position and velocity remain in this dead band region, no suspension control is applied along the sensitive DoF. When the NTM state drifts out of this deadband region, a suspension controller is engaged and remains on until the NTM state re-enters the dead-band region. The resulting trajectory of the NTM is more complex than that of the LTP drift mode experiment. In general, the duration of the free-flight segments vary from segment to segment. A 300 ks simulation yielded an approximately normal distribution of free-flights with a mean duration of  $\sim 357$  s and a standard deviation of  $\sim 14$  s. While at first glance, this irregularity in the free flight durations might seem to be disadvantageous, it has a potential advantage for data analysis because the few long segments can be used to estimate spectral information at lower frequencies. Additionally, unlike the LTP design, the DRS control design automatically adjusts to changes in the static gravity gradient and is robust to fairly significant changes. This means that the controller would not have to be re-optimized in-flight as the residual gravity gradient changes, for example due to fuel consumption.

## 2.3. Laboratory Experiments

Torsion pendulums with a LPF-like TM suspended as a torsion member have been used to measure the small forces relevant to the free-fall purity in eLISA and LPF. External torques on the suspended test mass provide the ability to mimic the LPF gravity gradients that must be compensated in orbit, making the torsion pendulum a test bed for the LPF actuation and the free-fall mode. To simulate a large DC acceleration, the pendulum can be rotated by an angle  $\Delta\Phi$  with respect to the inertial sensor, such that a DC torque  $N_{DC} = -\Gamma \cdot \Delta\Phi$  is required to keep it centered. This is analogous to the differential actuation force,  $m \cdot \Delta g$ , needed to compensate the difference in residual gravity gradient between the RTM and NTM in LISA Pathfinder. The differential forces applied by a pair of electrodes to produce this torque can be made similar to the force levels required in flight, but can also be varied at will on ground by choosing the pendulum rotation, something that is not possible in orbit.

The pendulum torque sensitivity is around  $1 \text{ fN} \cdot \text{m}/\sqrt{\text{Hz}}$  at 1 mHz, corresponding to an equivalent acceleration of  $50 \text{ fm}/\text{s}^2\sqrt{\text{Hz}}$ , and thus near the LPF noise specification. The pendulum can thus allow a quantitatively significant test of the free-fall mode and our analysis algorithms, with real data exhibiting gaps as well as the large dynamic range needed in free-fall mode, and thus readout and dynamic system linearity challenges. The on-ground experiment will also allow more flexibility to explore different control strategies, by varying flight and impulse time or control points, and different dynamic configurations, thanks to the possibility to have a variable stiffness.

The torsion pendulum consists of a hollow replica of the LISA Pathfinder test mass, suspended by a thin silica torsion fiber, and hangs inside a Gravitational Reference Sensor (GRS) prototype [8]. The sensitive degree of freedom is  $\Phi$ , the rotation around the z-axis. A breadboard version of the front-end electronic chain provides torque authority of  $\sim 200 \text{ fN} \cdot \text{m}$  by applying desired actuation voltages a diagonal pair of  $x$  electrodes.

The equation of motion for the torsional degree of freedom for the test-mass is

$$I\ddot{\Phi} = -I\omega_0^2\Phi - \frac{I}{\tau}\dot{\Phi} + N(t) \quad (1)$$

<sup>1</sup> the DRS is part of NASA's Space Technology 7 (ST7) mission

where  $I$  is the moment of inertia,  $\omega_0 = \sqrt{\frac{\Gamma}{I}} = \frac{2\pi}{T_0}$  is the pendulum resonance angular frequency (and  $T_0$  the period),  $\Gamma$  is the pendulum rotational elastic constant and  $\tau$  is the energy decay time. It is possible to soften electrostatically the pendulum by applying DC constant voltages to lengthen the pendulum period from roughly 465 s, without applied fields, to as much as  $T_0 \approx 830$  s, to allow flight times comparable to those foreseen for LPF,  $T_{fly} = 250$  s.

Our free-fall test consists of three measurements. The first is the pendulum background torque noise level in absence of any applied force, converting the measured angular displacement into torque  $N_m = I\ddot{\Phi} + \Gamma\Phi + \frac{I}{\tau}\dot{\Phi}$  and then estimating the noise power spectral density  $S_{N_m} = S_N$ . This is performed with the pendulum rotated such that the TM is centered without applied torques.

Then, to simulate a large DC acceleration, the pendulum is rotated by a large angle with respect to the GRS. In the current configuration, a rotation angle  $\Phi_{EQ} \approx 2$  mrad requires a DC torque of roughly  $13 \text{ pN} \cdot \text{m}$  to keep the test mass centered, a differential force of roughly  $1.3 \text{ nN}$  (with electrostatic softening of the pendulum, this gives  $\Phi_{EQ} \approx 5.3$  mrad, see Figure 3b). The second measurement is then the torque noise  $S_{N_m} = S_{N'} + S_{N_{act}}$  with the test mass continuously actuated, introducing excess actuations noise.

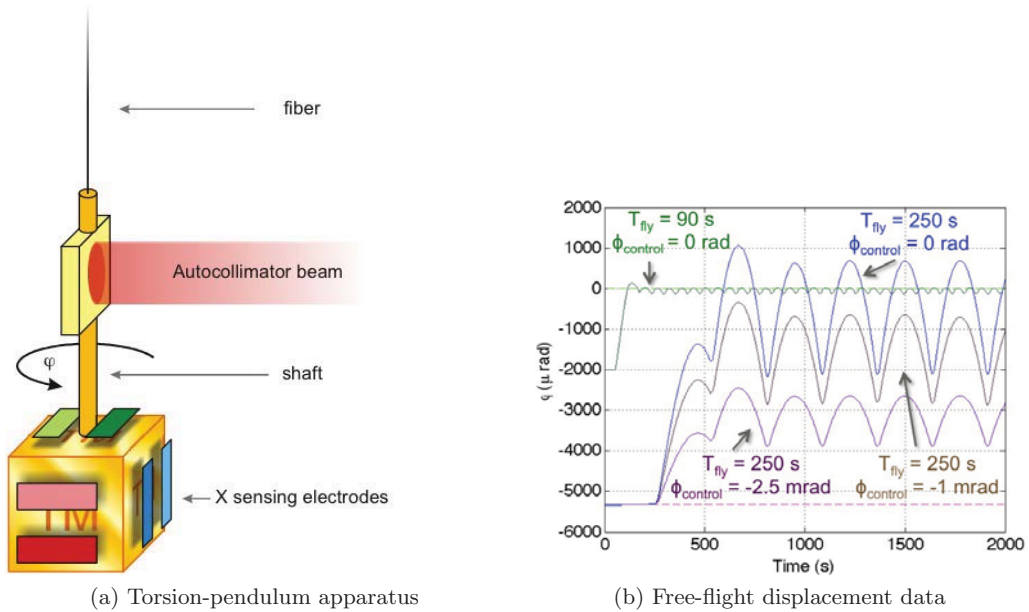


Figure 3: Example data from a torsion pendulum facility used to simulate LPF free-flight experiments. The displacement timeseries is for an experiment performed with  $T_{fly} = 90$  s and 250 s,  $\Phi_{control} = 0, -1, -2.5$  mrad.

The third measurement is then the torque noise when the free-fall control scheme is actuated to hold the TM centered on average, with the same pendulum rotation as in the second phase. Torque impulses are applied periodically with a duty cycle  $\chi$ , with average amplitude  $N_{ACT} \approx -\Gamma\Phi_{EQ}/\chi$ . The actuation-free torque noise  $S_{N_m} = S_{N'}$  is measured by estimating the power spectral density during the free motion of the pendulum. Ideally, in the absence of any error introduced by the data analysis or the large dynamic range, the actuation-free noise of the first measurement is recovered.

The pendulum dynamics in between two impulses is a free oscillation around the equilibrium point. The torsional spring is small and positive, in contrast to the small negative spring

expected in orbit. Because the stiffnesses are small and the flight time relatively short, 250 s flight compared to the 830 s free-oscillation period, the motion is similar to a parabolic flight in both cases.

The motion is periodically forced, by the impulses applied, to come back to a single initial position with a chosen velocity, to allow a periodic flight. To do that, a control scheme is implemented, where an observer estimates the pendulum position  $\Phi(t_n)$  and velocity  $\dot{\Phi}(t_n)$  before each impulse with least squares fitting of the pendulum rotation data. Then a controller estimates the impulse intensity  $N_{imp(n)}$  needed to reach the ideal initial point for the next cycle, using the pendulum dynamic constants and the flight and impulse times,  $T_{fly}$  and  $T_{imp}$ . We can also vary the set point angle, or average position of the mass during the free flight, to change the actuation level.

Example preliminary data are shown in Figure 3b, with flight times of 90 and 250 s, using pendulum periods of 482 and 830 s, respectively, employing also different controller set points. The controller has been successfully employed with a variety of control configurations. The next step is that of debugging and understanding the torque noise in our on-ground free-fall mode and beginning to test our flight data analysis algorithms on the experimental data.

### 3. Data Analysis Challenge

At their core, the goals of the LPF drift mode experiments are the same as those conducted in the standard science mode: measurement of the spectrum of the acceleration of the RTM relative to the NTM and estimation of system parameters such as the gravity gradients, stiffness terms, etc. The general analysis strategy for LPF experiments begins with models of the expected acceleration that are parametrized by relevant system parameters. Fitting these models to the data provides estimates of the system parameters and the fit residuals can be used to estimate the relative test mass acceleration.

The drift mode data poses two unique challenges to this approach. The first is the presence of the kicks, which represent a high-noise configuration of the NTM and can't be used for spectral estimation. Consequently, they must be excised in some way. The second is the size of the free-flight signals relative to the noise levels of interest. In displacement, the free flights have an amplitude of  $\sim 10 \mu\text{m}$ , compared with a displacement sensitivity of  $7 \text{ pm}/\sqrt{\text{Hz}}$ . In acceleration, the constant gravity gradient is  $10^{-10} \text{ m/s}^2$ , compared with expected noise levels of  $\sim 10^{-14} \text{ m/s}^2/\sqrt{\text{Hz}}$ . Residuals in fitting the free-flight terms caused by small parameter or model errors can significantly impact the resulting estimate of the relative acceleration between the NTM and RTM. The effect of such residuals is exacerbated by the presence of the data gaps around the excised kicks, to which we now turn our attention.

After the ‘deterministic’ portion of the free-flight data has been removed, the next step is to estimate the spectrum of the residuals. The data from individual free-flights is well suited to this purpose and standard spectral estimation techniques can be applied. Unfortunately, the minimum Fourier frequency at which the spectrum can be estimated is determined by the length of the free flight segments. For the  $\sim 350 \text{ s}$  flights for the LTP experiment, this limits the estimation to  $f > 3 \text{ mHz}$ . To move to the lower end of the LTP band (1 mHz) and to the full LISA band (0.1 mHz), requires combining data from successive flights. Since the portions of data during the kicks cannot be used, this amounts to the not-uncommon problem of estimating the underlying spectrum in the presence of gaps.

The general effect of data gaps is to introduce systematic biases into the estimated spectrum of the underlying continuous process. The precise nature of this bias depends both on the characteristics of the gaps (their duration, number, and grouping) as well as on the spectrum of the signal. The primary effect on LPF free-flight data is to ‘fill in the bucket’ in the LPF sensitivity. This is due to power from the upper end of the LPF band aliasing into lower frequencies and is exacerbated by the large range in noise power between the minimum noise of



$\sim 3 \times 10^{-14} \text{ m/s}^2/\sqrt{\text{Hz}}$  at  $f \sim 3 \text{ mHz}$  and the noise of  $\sim 5 \times 10^{-12} \text{ m/s}^2/\sqrt{\text{Hz}}$  at  $f \sim 100 \text{ mHz}$ . To illustrate this effect, we generated a series of ‘noise-only’ mock LTP drift mode data using a linear state-space simulator of the LTP experiment. The simulator was configured to run in nominal science mode (no free-flights) with the capacitive actuation noise of the NTM along the sensitive axis artificially turned off to mimic the noise environment in the drift mode experiment. The red trace in Figure 4a shows the average residual RTM-NTM acceleration noise for an ensemble of 100 runs of the simulator with different noise seeds, representing the expected level in a drift mode experiment. Gaps of 4 s duration separated by 350 s continuous segments were then artificially placed in the data to mimic the portion of the data that would be removed to avoid the influence of the kicks during drift mode. The increase in gap size over the 1 s kick duration allows the suppression of transients generated during downsampling of the data from 10 Hz to 1 Hz and an estimate of the acceleration from the measured displacement timeseries using finite differencing. As a zeroth-order estimate for the data in the gaps, a linear ‘patch’ is placed in each gap connecting the beginning point with the end point of the gap. The spectral estimate from this linearly patched data is the green trace in Figure 4a. It is clear that the linearly patched data gives a severely overestimated noise level over the entire LPF band. For reference, Figure 4c shows the fractional error of the spectral estimates over the lower portion of the LPF/LISA frequency band and demonstrates that the spectrum recovered using linear patches overestimates by factors of  $30 \sim 10^3$ .

#### 4. Approaches to Data Analysis

In this section, we briefly outline some of these techniques developed as part of this ongoing effort, with the details left to future publications.

##### 4.1. Windowing of data gaps

One strategy for mitigating the adverse effects of data gaps on the estimation of spectra from timeseries is to employ windowing. When estimating spectra from timeseries without gaps, it is typical to first multiply the time series by a window function that smoothly tapers to zero at the beginning and end of the timeseries. This helps to suppress artifacts caused by mismatch between the beginning and the end of the timeseries. This same strategy can be applied to data with gaps by applying a window to each individual segment of continuous data. The cyan line in Figure 4a shows the result of the spectral estimate made using a  $\sin^2$  window in each free-flight segment. The estimate of the spectra made from this windowed data is a far better approximation of the original data than that made using the linear patches. However, the windowed estimate shows a large excess relative to the original spectrum near the inverse gap separation of  $1/354 \text{ s} \approx 2.8 \text{ mHz}^2$  and harmonics. This is emphasized in the zoom-in around 2.8 mHz in Figure 4b. The cyan trace in Figure 4c shows the fractional error in the spectral estimate with the windowing method relative to the original data. With the exception of the inverse gap separation frequencies, the error is less than or near unity.

A more rigorous test of the windowing method is the Kolmogorov-Smirnov (KS) test, a statistical technique to quantify the probability that two samples are drawn from different underlying distributions. In our case, we can treat each frequency bin as a separate random process and compare the distributions of the power in each bin across the ensemble of 100 simulations for the spectra estimated from original signal with no gaps to the spectra estimated using the windowed method. The cyan line in Figure 4d shows the value of the KS statistic as a function of frequency for the windowed method. The horizontal red lines show values of the KS statistic corresponding to probabilities of 90% and 99% that the two data sets are drawn from *different* distributions. The fact that the cyan trace rises at Fourier frequencies below

<sup>2</sup> To clarify, 4 s gaps separated by 350 s free-flight segments gives a center-to-center gap spacing of 354 s.




Figure 4: Comparison of methods for spectral estimation of RTM acceleration relative to NTM with data gaps corresponding to LTP drift mode kicks of 1 s duration separated by 350s free-flight segments. A linear state-space simulator of LTP was used to generate an ensemble of 100 noise realizations of LTP science mode data and data gaps were artificially inserted allowing for the comparison of the estimated spectra with and without gaps. The upper plots show the spectra of the original signal (red), the signal with gaps (green), the spectra estimated with the window method (cyan), and the spectra estimated with the constrained-gaussian patching method (blue). The upper left plot shows the lower portion of the LPF band (where addressing the gaps is important) and the upper right plot zooms in around the inverse free-flight duration. The lower left plot shows the fractional error between the spectrum estimated with gaps and that without gaps for the three approaches. The lower right plot shows the Kolmogorov-Smirnov (KS) statistic comparing the ensemble distribution of power for the original data, the windowed data, and the constrained-gaussian filled data as a function of frequency bin. The horizontal dashed lines show critical values of the KS statistic corresponding to probabilities that the two samples arise from *different* distributions. Both the window and patching method yield significant bias at the inverse free-flight duration frequency and harmonics, but the bias of the patching method is more sharply concentrated near this frequency and is less significant at lower frequencies.

$\sim 1$  mHz indicates that the distributions of the spectra made from the un-gapped and windowed data differ in a statistically-significant way. At higher Fourier frequencies, the value of the KS statistic decreases, but there is no straightforward way to translate this into a confidence that the two samples are drawn from the *same* distribution.



#### 4.2. Constrained-Gaussian Gap Patching

One approach to dealing with data gaps and the difficulties they cause with estimating spectra is to fill the gaps with fabricated data. While the fabricated data in these patches will necessarily bias the spectral estimate, if the gaps are not too numerous and the patch data are chosen carefully, the bias can be far less than what arises from the gaps themselves. To minimize the bias resulting from the patches, the data in the patches should have the same spectral content as the missing data it is replacing. The first requirement is that the data points in the patches have proper correlations with *one another*, this is relatively straightforward problem of drawing random samples of (unevenly-sampled) data with a corresponding spectrum. The second requirement is that the data points in the patches have proper correlation with the *existing data*. In Appendix A, we demonstrate that enforcing the correlation with the existing data can be accomplished by biasing the mean values of the random data used to make the patches, a technique we call *Constrained-Gaussian Gap Patching* or CG patching. The values for these biases can be entirely computed from the existing data and a model of the underlying spectrum.

Once patched, standard approaches can be used to compute the spectrum of the entire data set as if it were continuous, without suffering from the bias introduced by the gaps. In principal this is a catch-22, one must know the spectrum in order to pick the patches that will allow the spectrum to be measured. In practice, an approximate *a priori* model of the spectrum will suffice. If necessary, an iterative procedure can be used in which an initial guess for  $S_y$  is used to generate the patches, a new value for  $S_y$  is estimated from the patched data series, and that spectrum is used to determine new patches.

The CG-patching approach was applied to the mock ‘noise-only’ drift-mode data in Figure 4a. The ensemble-average spectra is shown in the blue trace in Figure 4a and matches the average spectra from the original data (red) quite closely for all frequencies except the inverse gap separation frequency and its harmonics. The zoom-in around the inverse gap separation frequency in Figure 4b shows that, compared to the similar artifact in the windowed data, the artifact in the CG-patched data is not as large in amplitude or bandwidth. This might be of limited utility if there is a feature of interest in the underlying spectrum that is particularly close to the inverse gap separation frequency. The performance of the CG-patching is further illustrated by the fractional error (Figure 4c) and KS statistic (Figure 4d), both of which illustrate the CG-patching method’s improvement in performance relative to the windowed approach at low Fourier frequencies.

#### 4.3. Other methods

While the window and CG-patching methods have received the most attention for potential application to data from free-flight experiments, other methods are also being studied. One approach is to recognize that if an approximate model for the underlying spectra is known (as is also required for the CG-patching method) and the location of the gaps are known, the spectral bias induced by the gaps can be analytically removed. This is a classic example of a deconvolution problem with the associated computational difficulties. However, if the data is aggressively low-pass filtered, the problem becomes tractable.

An altogether unique approach is to consider only the data in the free-flights themselves and perform time-domain fits on the free-flight trajectories. Each of these fits will have parameters such as residual gravity gradient that are the same as those for the first stage, explained at the beginning of Section 3, of the methods outlined above. However, there will also be some variation in these parameters between successive flights that is caused by low-frequency (Fourier frequencies lower than the inverse free-flight time) acceleration noise in the system. By analyzing the variation in parameters recovered from fits to a series of free-flight trajectories, an estimate of acceleration noise at Fourier frequencies below the inverse free-flight frequency can be made

without the need to deal with data gaps.

## 5. Conclusions

The free flight experiments planned for LPF will provide an opportunity to gather data that is even more representative of a full-scale LISA like instrument than the standard LPF science mode. This improvement will allow for better characterization of the small forces that ultimately limit the performance of LPF, LISA, and other future instruments requiring low-disturbance environments such as advanced geodesy or fundamental physics missions. The challenges in extracting this information from the free-flight data are significant, but progress in overcoming them is being made with a number of independent techniques. The challenge of estimating the relative RTM/NTM acceleration noise in the presence of data gaps has been successfully addressed by a number of techniques with the current limitation being the ability to combine this with accurate removal of the deterministic free-flight signal. It is fully expected that the remaining challenges will be resolved in time for launch and operations.

## Appendix A: Generating Constrained-Gaussian Patches

In the time domain, spectral information is encoded in the two-point correlation function,  $C^{jk}$ , which measures the correlation between the samples  $y^j$  and  $y^k$ . For stationary processes, the two-point function depends only on the separation between the samples,  $C^{jk} = \mathcal{C}(j - k)$ , and is directly related to the spectral density,

$$C^{jk} = \int_{-\infty}^{+\infty} \frac{1}{2} S_y(f) e^{-2\pi f T(j-k)}. \quad (2)$$

Under the further assumption that  $y^j$  measures a Gaussian random process, the two-point function can be used to express the probability distribution function (PDF) for  $y^j$ ,

$$P(\vec{y}) \propto \exp \left[ -\frac{1}{2} \sum_{j,k} C_{jk} y^j y^k \right], \quad (3)$$

where  $C_{jk}$  is the inverse of  $C^{jk}$  such that  $C_{jl} C^{lk} = \delta_k^j$ .

For the case of data with gaps, it is useful to separate the data  $y_j$  into portions inside the gaps,  $j \in \mathcal{G}$ , and data between the gaps  $j \in \bar{\mathcal{G}}$ . The sum in (3) can then be divided into three parts,

$$\ln P(\vec{y}) = \text{const} - \frac{1}{2} \left[ \sum_{\bar{\mathcal{G}}, \bar{\mathcal{G}}} C_{jk} y^j y^k + 2 \sum_{\bar{\mathcal{G}}, \mathcal{G}} C_{jk} y^j y^k + \sum_{\mathcal{G}, \mathcal{G}} C_{jk} y^j y^k \right]. \quad (4)$$

The first term in (4) is a fixed constant depending on measured data that can be absorbed into the overall normalization constant. The final term depends entirely on the data within the gaps, and just describes the PDF for (unevenly sampled) Gaussian noise with a particular spectrum. The middle term encodes the relationship between data inside the gaps and data outside the gaps. Faithfully reproducing that relationship is the key to generating good patches for filling gaps in measured data. Since we are interested in generating the data within the gaps, we can factor the summation of the middle term in (4) and define  $\lambda_j \equiv \sum_j^{\bar{\mathcal{G}}} C_{jk} y^j$  such that that term becomes  $2 \sum_j^{\mathcal{G}} \lambda_j y^j$ . The PDF can then be re-written as

$$\ln P(\vec{y}) = \text{const} - \frac{1}{2} \sum_{\mathcal{G}, \mathcal{G}'} C_{jk} (y^j - \Delta^j)(y^k - \Delta^k), \quad (5)$$

where  $\Delta^j \equiv -\sum_k^{\bar{G}} C^{jk} \lambda_k$ . Comparing (5) with the standard expression for multivariate Gaussian distributions, one finds that the data in the gap is described by a multivariate Gaussian with covariance  $C_{jk}$  and non-zero means  $\Delta^k$ . Note that the values for  $\Delta^k$  are entirely determined by the data outside the gap.

The prescription for *Constrained Gaussian Gap Filling* is thus as follows: make a guess for spectral density  $S_y(f)$  describing the data; compute an estimated two-point correlation function,  $C_{jk}$ ; compute the  $\Delta^k$  for each point on the gap based on the data outside the gap and the estimated  $C_{jk}$ , draw data from a multivariate Gaussian distribution with covariance  $C_{jk}$  and mean  $\Delta^k$ .

## References

- [1] Bondi H, Pirani F A E and Robinson I 1959 *Royal Society of London Proceedings Series A* **251** 519–533
- [2] Armano M and et al 2009 *Class. Quant. Grav.* **26**
- [3] Lange B 1964 *AIAA Journal* **2** 1590–1606 URL <http://dx.doi.org/10.2514/3.55086>
- [4] deBra D B 1997 *Classical and Quantum Gravity* **14** 1549 URL <http://stacks.iop.org/0264-9381/14/i=6/a=026>
- [5] Antonucci F, Armano M, Audley H, Auger G, Benedetti M, Binetruy P, Boatella C, Bogenstahl J, Bortoluzzi D, Bosetti P, Brandt N, Caleno M, Cavalleri A, Cesa M, Chmeissani M, Ciani G, Conchillo A, Congedo G, Cristofolini I, Cruise M, Danzmann K, Marchi F D, Diaz-Aguilo M, Diepholz I, Dixon G, Dolesi R, Dunbar N, Fauste J, Ferraioli L, Fertin D, Fichter W, Fitzsimons E, Freschi M, Marin A G, Marirrodriaga C G, Gerndt R, Gesa L, Giardini D, Gibert F, Grimani C, Grynagier A, Guillaume B, Cervantes F G, Harrison I, Heinzl G, Hewitson M, Hollington D, Hough J, Hoyland D, Hueller M, Huesler J, Jeannin O, Jennrich O, Jetzer P, Johlander B, Killow C, Llamas X, Lloro I, Lobo A, Maarschalkerweerd R, Madden S, Mance D, Mateos I, McNamara P W, Mendes J, Mitchell E, Monsky A, Nicolini D, Nicolodi D, Nofrarias M, Pedersen F, Perreur-Lloyd M, Perreca A, Plagnol E, Prat P, Racca G D, Rais B, Ramos-Castro J, Reiche J, Perez J A R, Robertson D, Rozemeijer H, Sanjuan J, Schleicher A, Schulte M, Shaul D, Stagnaro L, Strandmoe S, Steier F, Sumner T J, Taylor A, Texier D, Trenkel C, Tombolato D, Vitale S, Wanner G, Ward H, Waschke S, Wass P, Weber W J and Zweifel P 2011 *Classical and Quantum Gravity* **28** 094002 URL <http://stacks.iop.org/0264-9381/28/i=9/a=094002>
- [6] Audley H, Danzmann K, Marin A G, Heinzl G, Monsky A, Nofrarias M, Steier F, Gerardi D, Gerndt R, Hechenblaikner G, Johann U, Luetzow-Wentzky P, Wand V, Antonucci F, Armano M, Auger G, Benedetti M, Binetruy P, Boatella C, Bogenstahl J, Bortoluzzi D, Bosetti P, Caleno M, Cavalleri A, Cesa M, Chmeissani M, Ciani G, Conchillo A, Congedo G, Cristofolini I, Cruise M, Marchi F D, Diaz-Aguilo M, Diepholz I, Dixon G, Dolesi R, Fauste J, Ferraioli L, Fertin D, Fichter W, Fitzsimons E, Freschi M, Marirrodriaga C G, Gesa L, Gibert F, Giardini D, Grimani C, Grynagier A, Guillaume B, Cervantes F G, Harrison I, Hewitson M, Hollington D, Hough J, Hoyland D, Hueller M, Huesler J, Jeannin O, Jennrich O, Jetzer P, Johlander B, Killow C, Llamas X, Lloro I, Lobo A, Maarschalkerweerd R, Madden S, Mance D, Mateos I, McNamara P W, Mendes J, Mitchell E, Nicolini D, Nicolodi D, Pedersen F, Perreur-Lloyd M, Perreca A, Plagnol E, Prat P, Racca G D, Rais B, Ramos-Castro J, Reiche J, Perez J A R, Robertson D, Rozemeijer H, Sanjuan J, Schulte M, Shaul D, Stagnaro L, Strandmoe S, Sumner T J, Taylor A, Texier D, Trenkel C, Tombolato D, Vitale S, Wanner G, Ward H, Waschke S, Wass P, Weber W J and Zweifel P 2011 *Classical and Quantum Gravity* **28** 094003 URL <http://stacks.iop.org/0264-9381/28/i=9/a=094003>
- [7] Carbone L, Ciani G, Dolesi R, Hueller M, Tombolato D, Vitale S, Weber W J and Cavalleri A 2007 *Phys. Rev. D* **75**(4) 042001 URL <http://link.aps.org/doi/10.1103/PhysRevD.75.042001>
- [8] Cavalleri A, Ciani G, Dolesi R, Hueller M, Nicolodi D, Tombolato D, Wass P J, Weber W J, Vitale S and Carbone L 2009 *Classical and Quantum Gravity* **26** 094012 URL <http://stacks.iop.org/0264-9381/26/i=9/a=094012>
- [9] Grynagier A 2010 The drift mode for LISA Pathfinder, control description Tech. Rep. S2-iFR-TN-3005 Universitat Stuttgärt