



Publication Year	2022
Acceptance in OA	2024-05-06T13:52:42Z
Title	Transient acceleration events in LISA Pathfinder data: Properties and possible physical origin
Authors	Armano, M., Audley, H., Baird, J., Binetruy, P., Born, M., Bortoluzzi, D., Castelli, E., Cavalleri, A., Cesarini, A., Chiavegato, V., Cruise, A. M., Dal Bosco, D., Danzmann, K., De Deus Silva, M., Diepholz, I., Dixon, G., Dolesi, R., Ferraioli, L., Ferroni, V., Fitzsimons, E. D., Freschi, M., Gesa, L., Giardini, D., Gibert, F., Giusteri, R., Grimaldi, C., Grzymisch, J., Harrison, I., Hartig, M. S., Heinzl, G., Hewitson, M., Hollington, D., Hoyland, D., Hueller, M., Inchauspé, H., Jennrich, O., Jetzer, P., Johlander, B., Karnesis, N., Kaune, B., Korsakova, N., Killow, C. J., Lobo, J. A., López-Zaragoza, J. P., Maarschalkerweerd, R., Mance, D., Martín, V., Martin-Polo, L., Martin-Porqueras, F., Martino, J., McNamara, P. W., Mendes, J., Mendes, L., Meshksar, N., Nofrarias, M., Paczkowski, S., Perreur-Lloyd, M., Petiteau, A., Plagnol, E., Ramos-Castro, J., Reiche, J., Rivas, F., Robertson, D. I., RUSSANO, Giuliana, Sala, L., Sarra, Paolo, Slutsky, J., Sopena, C. F., Sumner, T., Texier, D., Thorpe, J. I., Vetrugno, D., Vitale, S., Wanner, G., Ward, H., Wass, P., Weber, W. J., Wissel, L., Wittchen, A., Zannoni, Carlo, Zweifel, P., LISA Pathfinder Collaboration
Publisher's version (DOI)	10.1103/PhysRevD.106.062001
Handle	http://hdl.handle.net/20.500.12386/35078
Journal	PHYSICAL REVIEW D
Volume	106

Transient acceleration events in LISA Pathfinder: properties and possible physical origin

M Armano,¹ H Audley,² J Baird,³ P Binetruy,^{4,†} M Born,² D Bortoluzzi,⁵ E Castelli,^{6, a} A Cavalleri,⁷ A Cesarini,⁸ V Chiavegato,⁶ A M Cruise,⁹ D Dal Bosco,⁶ K Danzmann,² M De Deus Silva,¹ I Diepholz,² G Dixon,⁹ R Dolesi,⁶ L Ferraioli,¹⁰ V Ferroni,⁶ E D Fitzsimons,¹¹ M Freschi,¹ L Gesa,^{12,†} D Giardini,¹⁰ F Gibert,^{6, b} R Giusteri,² C Grimaldi,⁸ J Grzysch,¹³ I Harrison,¹⁴ M S Hartig,² G Heinzel,² M Hewitson,² D Hollington,¹⁵ D Hoyland,⁹ M Hueller,⁶ H Inchauspé,^{3, c} O Jennrich,¹³ P Jetzer,¹⁶ B Johlander,¹ N Karnesis,¹⁷ B Kaune,² N Korsakova,³ C J Killow,¹⁸ J A Lobo,^{12,†} J P López-Zaragoza,¹² R Maarschalkerweerd,¹⁴ D Mance,¹⁰ V Martín,¹² L Martin-Polo,¹ F Martin-Portueras,¹ J Martino,³ P W McNamara,¹³ J Mendes,¹⁴ L Mendes,¹ N Meshksar,¹⁰ M Nofrarias,¹² S Paczkowski,² M Perreux-Lloyd,¹⁸ A Petiteau,^{4,3} E Plagnol,³ J Ramos-Castro,¹⁹ J Reiche,² F Rivas,^{6, d} D I Robertson,¹⁸ G Russano,^{6, e} L Sala,⁶ P Sarra,²⁰ J Slutsky,²¹ C F Sopuerta,¹² T Sumner,¹⁵ D Texier,¹ J I Thorpe,²¹ D Vetrugno,⁶ S Vitale,⁶ G Wanner,² H Ward,¹⁸ P Wass,^{15,22} W J Weber,⁶ L Wissel,² A Wittchen,² C Zanoni,⁶ and P Zweifel¹⁰

(LISA Pathfinder Collaboration)**

¹European Space Astronomy Centre, European Space Agency, Villanueva de la Cañada, 28692 Madrid, Spain

²Albert-Einstein-Institut, Max-Planck-Institut für Gravitationsphysik und
Leibniz Universität Hannover, Callinstrasse 38, 30167 Hannover, Germany

³Université Paris Cité, CNRS, Astroparticule et Cosmologie, F-75013 Paris, France

⁴IRFU, CEA, Université Paris-Saclay, F-91191 Gif-sur-Yvette, France

⁵Department of Industrial Engineering, University of Trento, via Sommarive 9, 38123
Trento, and Trento Institute for Fundamental Physics and Application / INFN

⁶Dipartimento di Fisica, Università di Trento and Trento Institute for
Fundamental Physics and Application / INFN, 38123 Povo, Trento, Italy

⁷Istituto di Fotonica e Nanotecnologie, CNR-Fondazione Bruno Kessler, I-38123 Povo, Trento, Italy

⁸DISPEA, Università di Urbino Carlo Bo, Via S. Chiara, 27 61029 Urbino/INFN, Italy

⁹The School of Physics and Astronomy, University of Birmingham, B15 2TT Birmingham, United Kingdom

¹⁰Institut für Geophysik, ETH Zürich, Sonneggstrasse 5, CH-8092 Zürich, Switzerland

¹¹The UK Astronomy Technology Centre, Royal Observatory, Edinburgh, Blackford Hill, Edinburgh EH9 3HJ, United Kingdom

¹²Institut de Ciències de l'Espai (CSIC-IEEC), Campus UAB, Carrer de Can Magrans s/n, 08193 Cerdanyola del Vallès, Spain

¹³European Space Technology Centre, European Space Agency, Keplerlaan 1, 2200 AG Noordwijk, The Netherlands

¹⁴European Space Operations Centre, European Space Agency, 64293 Darmstadt, Germany

¹⁵Physics Department, Blackett Laboratory, High Energy Physics Group, Imperial
College London, Prince Consort Road, London SW7 2BW, United Kingdom

¹⁶Physik Institut, Universität Zürich, Winterthurerstrasse 190, CH-8057 Zürich, Switzerland

¹⁷Department of Physics, Aristotle University of Thessaloniki, Thessaloniki 54124, Greece

¹⁸SUPA, Institute for Gravitational Research, School of Physics and
Astronomy, University of Glasgow, Glasgow G12 8QQ, United Kingdom

¹⁹Department d'Enginyeria Electrònica, Universitat Politècnica de Catalunya, 08034 Barcelona, Spain

²⁰OHB Italia S.p.A, Via Gallarate, 150 - 20151 Milano, Italy

²¹Gravitational Astrophysics Lab, NASA Goddard Space Flight Center, 8800 Greenbelt Road, Greenbelt, MD 20771 USA

²²Department of Mechanical and Aerospace Engineering, MAE-A, P.O.
Box 116250, University of Florida, Gainesville, Florida 32611, USA

(Dated: May 25, 2022)

We present an in depth analysis of the transient events, or glitches, detected at a rate of about one per day in the differential acceleration data of LISA Pathfinder. We show that these glitches fall in two rather distinct categories: fast transients in the interferometric motion readout on one side, and true force transient events on the other. The former are fast and rare in ordinary conditions. The second may last from seconds to hours and constitute the majority of the glitches. We present an analysis of the physical and statistical properties of both categories, including a cross-analysis with other time series like magnetic fields, temperature, and other dynamical variables. Based on these analyses we discuss the possible sources of the force glitches and identify the most likely, among which the outgassing environment surrounding the test-masses stands out. We discuss the impact of these findings on the LISA design and operation, and some risk mitigation measures, including

experimental studies that may be conducted on the ground, aimed at clarifying some of the questions left open by our analysis.

PACS numbers: 04.80.Nn, 07.05.Kf, 95.55.-n

[†] Deceased.

^a Current address: Gravitational Astrophysics Lab, NASA Goddard Space Flight Center, 8800 Greenbelt Road, Greenbelt, MD 20771 USA

^b Current address: isardSAT SL, Marie Curie 8-14, 08042 Barcelona, Catalonia, Spain

^c Current address: Institut für Theoretische Physik, Universität Heidelberg, Philosophenweg 16, 69120 Heidelberg, Germany

^d Current address: Universidad Loyola Andalucía, Department of Quantitative Methods, Avenida de las Universidades s/n, 41704, Dos Hermanas, Sevilla, Spain

^e Current address: INAF Osservatorio Astronomico di Capodimonte, I-80131 Napoli, Italy

^{**} Corresponding authors:
 eleonora.castelli@nasa.gov
 lorenzo.sala@unitn.it
 stefano.vitale@unitn.it

I. INTRODUCTION

The European Space Agency (ESA) launched and operated the LISA Pathfinder (LPF) mission [1, 2] between December 2015 and July 2017. The scientific goal of LPF was to demonstrate that parasitic forces on a test mass (TM), to be used as a geodesic reference in the future LISA gravitational wave observatory [3], may be suppressed below the required level.

To that aim, the mission carried a miniature version of one of the LISA interferometric arms, that is, two kilogram-size free-orbiting TMs, separated by a few tens of centimeters, and an interferometric readout measuring their relative acceleration along the line joining their respective centers of mass.

The mission surpassed its goals and found that the acceleration due to parasitic forces had a power spectral density (PSD) better than LISA requirements across the entire observational frequency band [20 μ Hz, 1 Hz] [1].

Acceleration noise was found to be stationary over the week-long measurement runs, allowing a consistent PSD estimation. In addition, the measured PSD was found to be rather stable over the more than one year duration of the mission science operations, except for the decrease of the Brownian noise following the corresponding decrease of the pressure surrounding the TMs [1].

Besides this quasi-stationary noise, the acceleration data series also contained isolated events with different shapes and amplitudes which we call *glitches*. These events are *de facto* signals and can be modeled and subtracted from the data, a procedure which is essential, at least for the most energetic ones, to get a consistent estimate of the PSD of the underlying quasi-stationary noise over the entire data series.

The purpose of this paper is to present: a comprehensive statistical description of glitches (Sec. III, IV); an analysis of their physical properties, including the results of a coincidence analysis with other auxiliary data series (Sec. V); a discussion of the possible physical sources that are excluded by the two preceding analyses and the sources which instead might closely match the experimental evidence (Sec. VI).

II. SUMMARY DESCRIPTION OF THE EXPERIMENT

A. The LISA Technology Package

The instrument flown on LPF, the LISA Technology Package (LTP), has been described in detail in [4]. Here we summarize its essential features.

The LTP, depicted in Fig. 1, carried two cubic Au-Pt




FIG. 1. Rendering of the LISA Technology Package. The rendering shows the two test-masses hosted inside their respective electrode-housings (some of the electrodes are not represented), and the vacuum chambers enclosing both test-masses and electrode housings. The picture also shows the high stability optical bench hosting all interferometric read-outs, and many other features of the instrument, launch lock, UV-light based test-mass neutralizer, etc. that are not relevant here. [Credits: ESA/ATG medialab]

TMs each with a mass of 1.928 kg and size 46 mm. During operation these TMs had no mechanical contact with their surroundings, and were *free falling*, each one at the center of a housing leaving 3 to 4 mm clearance gaps to the faces of the TM. Each of these *electrode housings* (EH) carried a series of electrodes facing all faces of the respective TM. These electrodes were used for two purposes. First they were part of a capacitive sensor of the motion of the TM relative to its housing, for all degrees of freedom (DOF). Second, they were used to apply feedback electrostatic forces and torques to the TM, whenever needed.

The main sensor for the relative motion of the TMs was a heterodyne laser interferometric system, called the Optical Metrology System (OMS) [5]. For the purpose of this paper it is important to recall that the interferometer measured six different DOF: the relative displacement $\Delta x(t)$ between the the TMs along the sensitive x -axis, joining their respective centers of mass; the relative displacement $x_1(t)$ along the x -axis of one of the TM, called TM1, relative to the interferometer optical bench and, as a consequence, also relative to the spacecraft; the two angles of rotation $\eta(t)$ and $\phi(t)$ for both TMs, around the y -axis and z -axis respectively. These six TM DOF, and also the remaining six, were also measured at all times by the capacitive sensors [6]. However the interferometric readout was approximately three orders of magnitude more sensitive than the capacitive one for all DOF for which they were both available. All the measurements above have a sampling rate of 10 Hz.

Each EH, with its respective TM, was hosted inside a vacuum chamber that was vented to space via a dedicated valve. The chamber was needed to handle the vacuum on ground, and because the outgassing within the spacecraft once in orbit was too large to achieve the desired vacuum level around the TMs. Thus the chamber was evacuated and sealed on ground and then, once on orbit, vented to space via a dedicated duct. In what follows we call the Gravitational Reference Sensor (GRS) the system of the TM, its EH and vacuum chamber, and all related accessories.

Besides the measurement of the TM motion, other physical quantities have been measured throughout the mission. In particular we measured: the magnetic field vector at various locations, via a dedicated set of magnetometers [7]; the temperature at various critical locations, via a dedicated set of thermistors [8]; the cosmic ray flux with a radiation monitor [9, 10]; finally, two additional interferometric readouts, one monitoring frequency fluctuations (F interferometer) and the other one monitoring common mode noise sources as a reference (R interferometer) [5]. The R interferometer time series has therefore been subtracted from the Δx , x_1 and F interferometer time series.

B. Dynamical controls and data series formation

LPF was a controlled dynamical system consisting of the spacecraft and the two TMs. More specifically, the spacecraft was forced to follow one of the TMs (TM1) along x via an active control loop, using the spacecraft cold gas micro-thrusters as actuators [11], known as drag-free control.

The TM rotation along ϕ and η was kept fixed relative to the spacecraft by an active loop using electrostatic torques. These torques were applied via the above mentioned electrodes.

No force was applied along x on TM1, while a control loop (electrostatic suspension) kept the distance between the two TMs nominally fixed, by applying a suitable electrostatic force along x on the other TM (TM2). All other degrees of freedom were also controlled, but the details are not relevant here.

As the distance between the TMs was actively controlled, the out-of-the-loop differential disturbance force per unit mass acting on the TMs, Δg_e , could not be identified with the in-loop relative acceleration measurement $\Delta \ddot{x}$ [2]. This acceleration, which was estimated numerically [12], on the contrary had to be corrected for the known applied feedback forces per unit mass $g_c(t)$ [2].

In addition, acceleration data series were also corrected for the following effects.

- Measured inertial forces per unit mass due to spacecraft rotation $g_i(t)$, which include the specific centrifugal force and the apparent angular acceleration [1]. These effects will not be relevant for LISA.
- The forces per unit mass generated by the motion of the TMs through static force gradients in the spacecraft, as LISA data can also be corrected for those. Such force acting on TM i is well approximated by the linear model $-\omega_i^2 x_i$, as described in [1].
- The spurious pick up $g_{CT}(t)$ of spacecraft motion along different DOF, due to cross-talk [13, 14], which also includes the pick-up of the common mode motion of the TMs, described by a term $\delta_{x_1} \ddot{x}_1$.

Thus the out-of-the-loop, differential force per unit mass data series used in the following analyses, can be written as:

$$\begin{aligned} \Delta g(t) = & \Delta \ddot{x}_{OMS}(t) + \omega_2^2 \Delta x_{OMS}(t) + (\omega_2^2 - \omega_1^2) x_{1,OMS}(t) \\ & - g_c(t) - g_i(t) - g_{CT}(t). \end{aligned} \quad (1)$$

Note that, in Eq. 1, we have attached the suffix OMS to all coordinates, to indicate that these have been measured by the relevant interferometers, and not by the capacitive sensors. For these we will use the GRS suffix. Note also that all the ω 's above, as well as δ_{x_1} , have been measured in dedicated calibration experiments [13]. In particular $\omega_2^2 \approx -4.5 \times 10^{-7} \text{ s}^{-2}$ is negative, while the differential stiffness $(\omega_2^2 - \omega_1^2)$ is roughly 20 times smaller and thus neglected in this discussion here.

$\Delta g(t)$ in Eq. 1 is our best estimate for $\Delta g_e(t)$. However the series is corrupted by any disturbance $n_{OMS}(t)$, random noise or spurious signal, affecting the differential interferometer readout Δx_{OMS} . Such disturbance enters into $\Delta g(t)$ in Eq. 1, both through $\Delta \ddot{x}_{OMS}(t)$, and through $\omega_2^2 \Delta x_{OMS}(t)$. Thus, the residual noise in Δg can be evaluated:

$$\Delta g = \Delta g_e(t) + \ddot{n}_{OMS}(t) + \omega_2^2 n_{OMS}(t). \quad (2)$$

For the sake of the following discussion, it is important to note the following.

- The suspension control loop has significant gain only at low frequency, with a 3 dB cut-off at about 4 mHz. Thus $-g_c(t)$ becomes a good representation of $\Delta g_e(t) + \ddot{n}_{OMS}(t) + \omega_2^2 n_{OMS}(t)$ for signals at mHz and below, while in the opposite limit, it is $\Delta \ddot{x}_{OMS}$ that mostly contributes to $\Delta g(t)$.
- Δg measures the difference between the forces acting on the two TMs along x . Thus there is no way of discriminating the contributions of the individual

TM forces; for instance the differential acceleration could be entirely caused by force on one TM or the other.

- $\Delta g > 0$ corresponds to a force pushing the TMs one toward the other.

C. Data runs

The data we consider here are made of uninterrupted *runs* during which the TMs and the satellite were in steady control conditions, with no purposely applied stimulus of any nature¹. We have analyzed data from the three different kind of runs listed below.

1. The week-long runs, with physical conditions adjusted to reach the lowest noise level, which we used to estimate the quasi stationary noise PSD [2]. During the majority of these runs, the temperature of the LTP was kept at about 22 °C. However in two of these runs the temperature was lowered to about 11 °C, to decrease the outgassing rate and hence pressure and Brownian noise. One of these is the best noise performance run of February 2017 [1].
2. Shorter runs, with operating conditions slightly different from those used to achieve the best noise performance, but still with low noise and quasi-stationary behaviour. These also include runs during which the spacecraft control used an alternative set of micro-thrusters, based on colloidal propellants. In these runs the noise performance was slightly worse than in the low noise runs [15]. Runs in 1 and 2 include a total measurement time of 138.4 d.
3. A set of runs of lower stability, the origin of which we describe as follows. In May 2017, in an attempt to further decrease pressure, we lowered the temperature to about 0 °C, a value outside the nominal operating range of the instrument. The instrument entered into a rather unstable state, with a rate of glitches so high as to make the estimate of the background noise meaningless below ~ 1 mHz. When the temperature was raised again to 11 °C, the instrument went back to its ordinary behaviour. These runs lasted a total of 11.9 d.

In the rest of the paper we call the runs in 1 and 2 *ordinary runs* (OR), and those in 3, at about 0 °C, *cold runs* (CR).

¹ Only in one these, a two day long run, a small sinusoidal signal of less than 100 fN amplitude was injected in the suspension control loop to calibrate its force authority. Its effect on Δg was negligible.

III. GLITCH DETECTION AND PARAMETER ESTIMATION

Glitches were detected as localized signals in the $\Delta g(t)$ data series, emerging from noise after some data pre-processing. They fall in two broad categories: impulse carrying glitches, and high frequency, low-impulse glitches. The detection method and the parameters we estimate are different for the two categories. We describe both in the following.

A. Impulse-carrying glitches

The first category includes 432 signals (98 in OR, 334 in CR) that carry significant total impulse per unit mass $\Delta v = \int_0^\infty \Delta g(t) dt$.

Note that Δv is also the zero-frequency limit of Fourier transform of the Δg signal. Thus these signals have significant energy within the low frequency band of the control loop, and hence show up also as features within the feedback force data series $g_c(t)$.

We manually detected glitches of this kind by first filtering the data with a low-pass, finite impulse response filter, consisting of a 100 s-long Blackman-Harris window with an effective roll-off frequency at 10 mHz.

The filter suppressed the intense high frequency noise coming from the double time derivative of the interferometer readout noise, and made the glitch fairly visible above the remaining noise within some continuous stretch of the $\Delta g(t)$ data series, that we call the glitch stretch.

Subsequently, for each glitch, a polynomial was fitted to the data in two 1000 s long stretches, one immediately preceding the glitch stretch, and the other immediately following it. The polynomial was of first order for glitch stretches shorter than 1000 s (see later for the definition of duration), and of second order for longer glitches. The best-fit polynomial was then subtracted from the data. The result of this procedure, for one of the longest glitches, is shown in Fig. 2, left panel.

Such background subtraction was necessary to get rid of the long term drift that affected all data, mostly due to the gravitational signal from propellant tank depletion and long term temperature variations. The use of a second order polynomial for the longest glitches was able to accommodate some drift rate variation over many hours duration.

To these pre-processed data, we fitted a simple signal template (described further down) in the time domain, properly low-passed with the same filter used for the data (see Fig. 2). For OR, the main purpose of such fitting was to remove the glitch from the data. This allowed us to estimate the PSD of the underlying noise on the




FIG. 2. Left: example of an impulse carrying glitch. The picture shows: (green) the native data after low-pass filtering and background subtraction; (orange) the fitting template; (gray) the residual after subtraction of the template. Right: example of a fast, low impulse glitch. The figure shows data after the 0.5 Hz low-pass filtering.

entire data series, thus reaching the lowest attainable frequencies. To this aim, we subtracted the unfiltered version of the best fit signal $h(t)$ from the native $\Delta g(t)$ time series.

For ordinary runs, the procedure was indeed quite effective [1], leaving a residual time series with the same PSD, within errors, as that of the series from which the glitch stretch was simply removed. This comparison is only possible down to the lowest frequency that could be attained with both time series (see Fig. 3).




FIG. 3. Amplitude spectral density (ASD) $S_{\Delta g}^{1/2}(f)$ of quasi-stationary noise, vs the frequency f , of the best noise run of February 2017. Green data points: ASD of glitch-free stretches only. Orange data points: ASD of residuals after glitch removal. Errors represent 68% confidence intervals.

For CR runs the data series were rather complex, with very short time intervals between glitches and quite a number of overlapping ones. Such complexity reduced

the quality of the fitting and did not allow for glitch subtraction to any level useful for PSD estimation below 1 mHz. Nevertheless the fit allowed estimating the glitch key parameters (see below).

Of 432 events, 48 were fitted to the heuristic template [1]:

$$h_2(t) = \frac{\Delta v}{\tau_1 - \tau_2} \left(e^{-t'/\tau_1} - e^{-t'/\tau_2} \right) \Theta(t'), \quad t' = t - t_0, \quad (3)$$

Here: t_0 is the glitch occurrence time; $\Theta(t)$ is the Heaviside step function; the decay time is set by the larger of the time constants τ_1 and τ_2 , while the rise time by the shorter of the two; finally, Δv is the total transferred impulse per unit mass.

For 384 of the 432 glitches, the template in Eq. (3) simplified to its limit² for $\tau_1 \rightarrow \tau_2 \rightarrow \tau$:

$$h_1(t) = \frac{\Delta v}{\tau^2} t' e^{-t'/\tau} \Theta(t'), \quad t' = t - t_0 \quad (4)$$

Two exceptional glitches, the longest ones and both in

² Note that this template corresponds to the first-order shapelet considered in [16].

OR, required a third exponential function

$$h_3(t) = \Delta v \left(\frac{\tau_1 e^{-t'/\tau_1}}{(\tau_1 - \tau_2)(\tau_1 - \tau_3)} + \frac{\tau_2 e^{-t'/\tau_2}}{(\tau_2 - \tau_3)(\tau_2 - \tau_1)} + \frac{\tau_3 e^{-t'/\tau_3}}{(\tau_3 - \tau_1)(\tau_3 - \tau_2)} \right) \Theta(t'), \quad t' = t - t_0, \quad (5)$$

to account for some fine structure of the glitch onset, and to allow for good quality glitch subtraction.

For all templates above, and for the glitches fitted to them, $h(t)$ never crosses the level $h(t) = 0$. We found 3 glitches that still carry a significant impulse, but, contrary to the others, do cross $h(t) = 0$ during their evolution, even though only once. These glitches are fitted to

$$h_c(t) = h_2(t) + \tau_d \dot{h}_2(t), \quad (6)$$

with τ_d a constant that can have any sign.

Note that all templates above and their related glitches leave no net step in the Δg data series after glitch subtraction, as the glitch signal goes to 0 at both its ends.

Given the non-optimal nature of a time domain fit in colored noise, we have not performed any dedicated goodness-of-fit test, besides the check on the PSD of residuals. In particular we have not made any systematic comparison among different templates. Actually in some cases, for the shortest glitches of smaller amplitude, a fit with a filtered version of a Dirac delta or with a simple exponential gave comparable results.

In addition to allowing for glitch removal, the fitting also allowed us to estimate the values of all template parameters and their errors, as described in App. A. From these parameter values, for the sake of further analysis, we also calculated an effective duration Δ for the templates in Eqs. (3) to (5) as the time interval following t_0 in which there is 99% of its signal energy, defined as $\int_0^t h^2(t') dt'$. In the simplest case of Eq. (4), this corresponds to $\Delta \sim 4.20 \tau$. For the template in Eq. (6) we define the duration as that of $h_2(t)$ in that same equation.

B. Fast, low-impulse glitches

A second population of glitches, which included 4 glitches in ordinary runs and 152 in cold runs, is characterized by short duration and minimal total impulse. Specifically,

- The overall duration of the glitch is compatible, within errors, with that of the convolution of the impulse response of the filter used to estimate the second time derivative, with that of the low-pass

filter.

- They carry no significant impulse (per unit mass) $\Delta v_{\text{glitch}} = \int_0^\infty \Delta g(t) dt$ (see Sec. IV B). As a consequence, they show no detectable counterpart in the feedback force time series $g_c(t)$.
- They are detected in the data after low-pass filtering with a filter (Blackman-Harris, 2 s long) with a roll-off frequency of $\simeq 0.5$ Hz, *i.e.*, significantly higher than that used for the other category of glitches. Actually, the majority shows up with no filtering at all.

Of the 4 glitches of this kind observed in OR, 2 were well subtracted from the data by fitting them to the filtered version of the second time derivative of a step in the Δx time series. As expected, given the high frequency nature and the lack of impulse, the subtraction had no detectable effect on the PSD. Given such a negligible impact on the data, the 2 remaining glitches that were discovered in a second search were neither fitted nor subtracted.

This kind of glitch was not subtracted from CR data either, as the fitting turned out to be unfeasible in most of the cases. Thus, for the sake of the following analysis, for all glitches in this category we will only consider the impulse Δv_{glitch} and the time of occurrence, defined as the time when $\Delta g(t)$ reaches the maximum of its absolute value.

C. Other spurious signals in the data

In addition to the glitches described above, we found a few signals in Δg caused by the impact of the spacecraft with micrometeoroids. The impact caused a well identified acceleration of the spacecraft [17]. As mentioned in Sec. II B, $\Delta g(t)$ has been corrected for the acceleration of the spacecraft. However, for the most energetic events a residual signal was still found in the data due to calibration errors. Glitches of this kind are well understood and could have been suppressed by a better calibration. Thus we are not going to discuss them here any further.

We will also not discuss spikes in the data, some of which were periodic, observed upon operation of particular devices and that could be reproducibly suppressed by turning their source off.

IV. GLITCH PARAMETER STATISTICS

We have performed a statistical analysis of the observed glitch parameters. In the following we give the results of such analysis, separately for the different glitch categories.

A. Impulse-carrying glitches

In this section we present the main statistical features of the glitches that are fitted to the models of Eqs. (3) to (5). We only briefly discuss at the end the properties of the few glitches fitted to Eq. (6).

1. Occurrence rate and waiting time

Fig. 4 shows the histograms of the waiting time ΔT between impulse-carrying glitches for both ordinary and cold runs. ΔT is defined as either the time between two subsequent glitches, the vast majority of the samples, or the time between the starting moment of the run and the first glitch, relevant for very short runs with few glitches.




FIG. 4. Histogram of the waiting time ΔT for OR (orange) and CR (blue). Straight lines, and associated shadowed areas, represent, respectively, the Bayesian fit to an exponential distribution, and the corresponding 68% confidence interval. For the CR, the fit is limited to data with $\Delta T \leq 0.2$ d. The rates from the Bayesian estimation for OR, and CR with $\Delta T \leq 0.2$ d, are, respectively, $\lambda = 0.96^{+0.11}_{-0.09} \text{ d}^{-1}$ and $\lambda = 32^{+2}_{-2} \text{ d}^{-1}$.

For both OR and CR, the result of a Lilliefors test [18] is compatible with ΔT being exponentially distributed. Data for OR are well fitted to an exponential distribution with average rate $\lambda = 0.96^{+0.11}_{-0.09} \text{ d}^{-1}$ (3). Bayesian rate estimation is described in App. B.

Data for the cold runs, and for $\Delta T \leq 0.2$ d, are also compatible with an exponential distribution with $\lambda = 32^{+2}_{-2} \text{ d}^{-1}$. However the distribution shows a clear excess tail for longer times, amounting to some excess counts at longer times, see Fig. 4, originating from data taken at different temperatures (see below).

³ This rate of $\lambda = 0.96^{+0.11}_{-0.09} \text{ d}^{-1}$ is apparently slightly higher than that reported in our preliminary search [1] of $\lambda = (0.78 \pm 0.02) \text{ d}^{-1}$. We have traced back this apparent discrepancy to the smaller subset of runs used in [1], and to a mistake in reporting the error. The event rate estimated with the current Bayesian analysis, at 68% confidence level on the same subset of runs, gives $\lambda = 0.75^{+0.13}_{-0.09} \text{ d}^{-1}$, which is compatible with the current estimation at $\simeq 1\sigma$.

Fig. 5 shows the time evolution of the rate for OR over the course of the mission, which is consistent with a time independent value.

Fig. 6 shows the evolution of the rate λ and of the system temperature T_{LTP} during CR. The reported temperature is that of the bay that contained the LTP, as the readout electronics for all thermometers on the LTP itself saturated at about 8 °C.

Note that the rate variations of Fig. 6 fully explain the few excess counts at longer waiting times in the CR histogram of Fig. 4.




FIG. 5. Glitch occurrence rate λ , during OR, vs time t from launch, throughout the entire mission. Points are calculated by grouping glitches observed during runs the start times of which differ by less than a month. Vertical error bars are Bayesian estimates assuming exponential distribution, and corresponds to 68.3% (1σ) likelihood (see App. B). Horizontal error bars correspond to the total duration of the considered epoch. The dashed line, and the associated gray shaded area, represent, respectively, the mean rate and its error from the Bayesian estimate. The blue shaded area indicates the epoch of CR.




FIG. 6. Glitch occurrence rate λ (left scale), and LTP bay temperature (right scale), during CR, as a function of time since the beginning of cooldown. Small temperature changes before day 12 were the result of adjustments of heaters settings aimed at stabilizing the system behaviour. Re-heating to ordinary conditions started at day 12.

2. Impulse and duration

Fig. 7 shows the impulse and the duration of all glitches that are fitted to the templates in any of Eqs. (3) to (5). For reference the figure reports, for any given duration, also the amplitude of a glitch, well fitted to the template

of Eq. (4) that would have $|\Delta v|/\sigma_{\Delta v} \equiv \text{SNR} = 3$. Here $\sigma_{\Delta v}$ is the error on Δv , and SNR stands for signal-to-noise ratio. For OR, the line refers to the sensitivity of the February 2017 run, and is then a lower limit for the other, less sensitive runs.




FIG. 7. Absolute impulse per unit mass $|\Delta v|$, and duration Δ , of impulse carrying glitches fitted to templates in Eqs. (3) to (5). Upper left panel: the 81 positive impulse glitches observed during OR. Upper right panel: the 17 negative impulse glitches observed during OR. Lower left panel: the 306 positive impulse glitches observed during CR. Lower right panel: the 28 negative impulse glitches observed during CR. For reference, the gray dashed line represents, for any given duration, the amplitude of a glitch of the kind in Eq. (4) that would have $\text{SNR} = 3$. In the two upper panels, SNR is calculated for the lowest noise OR of February 2017. In the lower panels, the line is calculated for the sensitivity of the CR. For a detailed analysis of the glitch SNR, see Fig. 9.

In addition, Fig. 8 gives separately the time evolution of $|\Delta v|$ and Δ for the glitches in OR.




FIG. 8. Absolute impulse per unit mass $|\Delta v|$ (top), and duration Δ (bottom) as a function of time from launch for the glitches of Fig. 7. Note that the apparent clustering corresponds to the different measurement runs. The blue shaded area corresponds to the epoch of CR.

Fig. 9 shows the histogram of their SNR, with its evident lower bound at $\text{SNR} \sim 3$.




FIG. 9. Histogram of the SNR for all glitches of Fig. 7. The lowest observed value is $\text{SNR} \sim 3$. The probability density refers to the logarithm of SNR.

For the sake of further discussion, we also report in Fig. 10 the histogram of the absolute peak value $|\Delta g_{\text{max}}|$ for the glitches in Fig. 7.

3. Zero-crossing glitches

The preceding figures do not include the 3 glitches that are fitted to the template in Eq. (6). We summarize their properties in Table I.




FIG. 10. Histogram of peak amplitude $|\Delta g_{\text{max}}|$ for all glitches of Fig. 7. The probability density is for the logarithm of the amplitude.

TABLE I. Observed glitches corresponding to the template in Eq. (6). Values above the double horizontal line refer to OR. Values below that same line, to CR.

Δv [pm/s]	Δ [s]	τ_d [s]
0.10 ± 0.01	35 ± 8	-37 ± 9
0.81 ± 0.03^a	111 ± 3	72 ± 4
0.42 ± 0.02	143 ± 50	-29 ± 22

^a Corresponds to an event in $\ddot{x}_{1,\text{OMS}}(t)$, see Sec. V C.

B. Fast, low-impulse glitches

As glitches of this category were rare in OR, we limit our analysis to CR. Fig. 11 in the upper panel shows the histogram of the waiting time.




FIG. 11. (upper) Histogram of the waiting time ΔT for fast, low impulse glitches in CR. (lower) Histogram of $(\Delta v^2_{\text{glitch}} / \Delta v^2_{\text{rms}})$ for low impulse glitches. The orange line is the properly normalized distribution for a chi-square with one degree of freedom.

A Lilliefors test [18] for the exponential distribution

on such histogram fails, indicating some departure from Poisson statistics. Some subsets of the glitches show indeed a marked clustering, especially near the beginning of the re-heating procedure.

We must note though that the mean rate was different at different times during the CR, while the test assumes a single global distribution. A test on subsets of the data show that some subset is still compatible with an exponential distribution.

Fig. 11 in the lower panel reports the histogram of the impulse Δv_{glitch} , estimated by numerically integrating the short data stretch containing the glitch, normalized to its estimated error Δv_{rms} . We derive Δv_{rms} , by performing the same numerical integration used to estimate Δv_{glitch} , on random selected, glitch-free stretches of data, of same length and within the same run.

The histogram is quantitatively compatible ($p = 0.23$) with a chi-square distribution with one degree of freedom, that is, with the hypothesis that Δv_{glitch} is normally distributed with zero mean and standard deviation equal to Δv_{rms} .

V. JOINT ANALYSIS WITH OTHER TIME SERIES

In an attempt to understand the nature of glitches we have analysed some other data series that have been measured synchronously with Δg throughout the mission.

A. Discriminating between force and readout effect using the GRS Δx sensor.

The motion of both TM along x and relative to the spacecraft has been measured at all times by the capacitive sensor of the GRS. From their measured coordinates $x_{1,\text{GRS}}(t)$ and $x_{2,\text{GRS}}(t)$, we have formed a measurement of their relative displacement, independent of $\Delta x_{\text{OMS}}(t)$:

$$\Delta x_{\text{GRS}}(t) = x_{2,\text{GRS}}(t) - x_{1,\text{GRS}}(t) + n_{\text{GRS}}(t) \quad (7)$$

with $n_{\text{GRS}}(t)$ the measurement noise.

The difference between these two measurements only contains the difference between the noise terms:

$$\Delta x_{\text{OMS}}(t) - \Delta x_{\text{GRS}}(t) = n_{\text{OMS}}(t) - n_{\text{GRS}}(t), \quad (8)$$

and would immediately reveal a spurious signal within $n_{\text{OMS}}(t)$, if such signal were large enough to be detected against the relatively noisy GRS data.

We have used the $\Delta x_{\text{OMS}}(t) - \Delta x_{\text{GRS}}(t)$ data series to discriminate glitches that may have been caused by such

signals within $n_{\text{OMS}}(t)$, from those due to the true force $\Delta g_e(t)$.

Specifically, we have used the information that the glitch signal $\ddot{n}_{\text{OMS}}(t) + \omega_2^2 n_{\text{OMS}}(t)$, would follow one of the templates in Eqs. (3) to (5). If we define $h(s)$ to be the Laplace transform of such a template in Δg , then the Laplace transform of the relevant associated glitch in $n_{\text{OMS}}(t)$ would be given by $h(s)/(s^2 + \omega_2^2)$.

As $h(s)$ is in all cases a rational function of s , then $n_{\text{OMS}}(t)$ would carry a diverging term $\propto e^{+\sqrt{-\omega_2^2}t} \simeq e^{+t/1.5\text{ks}}$, which after a few thousand seconds would dominate the data.

We illustrate the concept, for one of the glitches, in Fig. 12. The glitch is clearly visible in both $\Delta x_{\text{GRS}}(t)$ and $\Delta x_{\text{OMS}}(t)$, while it disappears in their difference. In the figure we also show the inverse Laplace transform of $h(s)/(s^2 + \omega_2^2)$, *i.e.* the signal one would observe in $\Delta x_{\text{OMS}}(t) - \Delta x_{\text{GRS}}(t)$, if the glitch were due to a spurious signal in $n_{\text{OMS}}(t)$. For the glitch in question, the picture clearly shows that this source for the glitch is ruled out.




FIG. 12. $-\Delta x_{\text{OMS}}(t)$ (black), $-\Delta x_{\text{GRS}}(t)$ (orange), and their difference (green) for an impulse-carrying glitch ($\Delta v = 22.1 \text{ pm/s}$, $\Delta = 7.82 \text{ s}$). The negative signs on the first two data series have been used to show that the residuals in the difference are dominated by the large noise in $\Delta x_{\text{GRS}}(t)$. Also shown (dashed line) is the signal one would observe in $\Delta x_{\text{OMS}}(t) - \Delta x_{\text{GRS}}(t)$, if the glitch were due to a spurious signal in $n_{\text{OMS}}(t)$.

We have calculated such hypothetical signal in $n_{\text{OMS}}(t)$, that is, the inverse Laplace transform of $h(s)/(s^2 + \omega_2^2)$ for all glitches, both in the OR and CR. We have found that such signals would have been clearly visible in the data if a sufficiently long observation time were available after the glitch occurrence time (up to 10^4 s for the weakest glitches). On the contrary, we have found none in the actual data. Only for a few events in the CR we could not reach any conclusion due to lack of sufficient observational data after the glitch.

Thus impulse-carrying glitches are due to true forces

acting on the TMs, and are not artifact due to motion readout.

Low momentum, high frequency glitches consist of a feature in Δx_{OMS} , and are not visible within the feedback force time series. However, they could still be either a spike in true acceleration, or just a feature in the interferometer output.

As already mentioned, glitches of this kind consist of a step, or of a few points outlier in Δx , or some variation of those. In the case of just a feature in the interferometer, $n_{\text{OMS}}(t) - n_{\text{GRS}}(t)$ would then contain a similar feature (for these fast signals, we are neglecting the term $\omega_2^2 n_{\text{OMS}}$). Unfortunately the largest steps are a few tens of pm high, while the resolution on step detection in the $\Delta x_{\text{OMS}}(t) - \Delta x_{\text{GRS}}(t)$ is not better than ~ 1 nm. Similar limitations hold for the detection of outliers. Thus we were not able to discriminate between true force and interferometer readout for this category of glitches.

B. Associated differential torque

In close analogy with what we did with $\Delta g(t)$, we measured the differential out-of-the-loop torque per unit moment of inertia on the TMs, both around y and around z . We have built these quantities by subtracting the control torques per unit moment of inertia from the measured differential angular accelerations of the TMs. For instance such differential torque for the z -axis, $\Delta \gamma_\phi$, is defined as

$$\Delta \gamma_\phi = \ddot{\phi}_2 - \ddot{\phi}_1 + (N_{\phi_1} - N_{\phi_2}) / I_{zz}, \quad (9)$$

where N are the commanded torques, I_{zz} is the common moment of inertia around z , and the subscripts indicate the TM. The subscript η indicates the rotation about the y -axis.

As for the case of Δg , angular accelerations are corrected for some small gradient effect. More important, as the rotational motion of the spacecraft is rather intense, and is common mode for both TMs, torques and angular rotations have been re-calibrated to maximize the rejection of such a large common mode disturbance.

For each of the impulse carrying glitches within the $\Delta g(t)$ time series, we have fitted both the $\Delta \gamma_\phi(t)$ and the $\Delta \gamma_\eta(t)$ time series to exactly the same template in Eq. (3) and (4), with same time of occurrence and same time parameters, leaving the amplitude $\Delta \Omega_{\phi,(\eta)}$, an effective increase in angular velocity, as its sole fitting parameter. In addition, we have also included in the fitting model a parabolic background.

We have estimated the uncertainty on $\Delta \Omega_{\phi,(\eta)}$ by repeating the above fitting procedure over the sliding stretch $\Delta \gamma_{\phi,(\eta)}(t_1 + \delta t, t_2 + \delta t)$, with t_1 and t_2 the time bounds

of the actual stretch that contains the glitch, and δt a sliding delay. We have thus generated a series $\Delta \Omega_{\phi,(\eta)}(\delta t)$, with $|\delta t| \leq 20(t_2 - t_1)$, which shows, at $\delta t = 0$, a marked peak above the background jitter, whenever the torque is significant (see Fig. 13).

As the $\Delta \Omega_{\phi,(\eta)}(\delta t)$ time series has an intrinsic auto-correlation over a scale $\delta t \simeq t_2 - t_1$, it contains in practice only $\simeq 40$ independent data points. Thus its root mean square cannot be used, as is, for error estimation. We then use, more cautiously, as an estimate for the error, the peak absolute value $\delta \Omega_{\phi,(\eta)}$ of the series, calculated on the data outside the central stretch $|\delta t| \leq t_2 - t_1$. More specifically, if the central peak does not exceed, in absolute value, $\delta \Omega_{\phi,(\eta)}$, we take $|\Delta \Omega_{\phi,(\eta)}| \leq \delta \Omega_{\phi,(\eta)}$.

If, on the contrary, the peak exceeds in absolute value $\delta \Omega_{\phi,(\eta)}$, we take this as the error. It must be noted though that, for Gaussian statistics, the maximum absolute value among 40 independent samples falls in the interval $(2.4_{-0.4}^{+0.5}) \sigma$. Thus the confidence interval associated with such an error is greater than 95%.




FIG. 13. An example of the $\Delta \Omega_\phi(\delta t)$ series for one of the very few glitches with significant associated torque. The value in $\delta t = 0$ is the result of the fit for native data. The arrows indicate the interval excluded for the sake of error evaluation.

A practical way to report the results of this analysis, is to introduce an effective *lever arm* defined as

$$r_{\phi,(\eta)} = \frac{I_{zz,(yy)}}{m} \frac{\Delta \gamma_{\phi,(\eta)}}{\Delta g} = \frac{I_{zz,(yy)}}{m} \frac{\Delta \Omega_{\phi,(\eta)}}{\Delta v} \quad (10)$$

For a single force, applied normal to one of the x or y faces of either TM, r_ϕ would be the distance between the force application point and the center of the face. A similar interpretation holds for r_η . Note that for a real point-like force, given the size of the TM, the maximum value for both $|r_\phi|$ and $|r_\eta|$ would be 23 mm, while there is no upper limit for a distribution of forces.

A particular interesting case, further discussed in Sec. VI, is that of a force resulting from a voltage difference between the TM and one of the electrodes facing its x -face. These electrodes were used for the control loop on

TM2, but also for the angular control of both TMs. Such a voltage difference would have produced a force with a lever arm $|r_\phi| \sim 11$ mm [19].

The results of the analysis for r_ϕ in OR are shown in Fig. 14. A list of the glitches with any of the lever arms significantly different from zero is reported in Table II.

Fig. 14 shows that, as expected from Eq. (10), the error on r_ϕ decays approximately as $|\Delta v|^{-1}$, decreasing below about 1 cm for $|\Delta v| \gtrsim 2$ pm/s. For only 3 of the 98 OR impulse-carrying glitches we detect an effective armlength that is significantly different from zero. In a slight majority (62%) of the cases, where we find an armlength compatible with zero, we have sufficient resolution to exclude the hypothesis of an electrode origin to the glitch. Specifically, 62% of the OR glitches are both compatible with $r_\phi = 0$ and have “upper limit” uncertainty of less than 7 mm (or 2/3 of the characteristic “electrode arm” of 11 mm). If any of these relatively high SNR glitches had had an 11 mm effective armlength, they would have been detected with peaks more than 50% above our background in Fig. 13.

This fraction decreases to 50% for glitches with $\Delta \leq 1$ min, while it increases to 91% for glitches with duration $\Delta \geq 1.6$ ks, a duration threshold that will be discussed in Sec. VI.

$\Delta\gamma_\eta(t)$ is significantly more noisy than $\Delta\gamma_\phi(t)$. This gives larger errors on r_η than on r_ϕ , and makes less likely to find lever arms significantly different from zero. For instance, just for only about 32% of the glitches the error upper bound is smaller than (2/3) 11 mm. We give this figure here just for the sake of comparison with the case of r_ϕ , but $r_\eta = 11$ mm has no special significance.

The complexity of data during CR applies also to the $\Delta\gamma_{\phi,(\eta)}(t)$ series that contain multiple fast events and other non stationary features. This makes the results of our search procedure more noisy, reducing, for instance, the fraction of glitches, for which the upper error bound to r_ϕ is less than (2/3) 11 mm, to about 45%, and to 39% if $\Delta \leq 1$ min.

Moreover, we note that, of the 3 zero-crossing glitches transferring impulse in Table I, none of them is associated to a measurable torque, measured with the same analysis as in this section. The error on r_ϕ associated to the stronger ones (one in OR and one in CR), is lower than (2/3) 11 mm.

As a final note, in addition to the glitches listed in Table II, our procedure also finds 3 glitches in the CR for which the peak in $\Delta\Omega_{\phi,(\eta)}(t)$ is significantly displaced from the true time of occurrence. We believe that these events are due the accidental coincidence between a force glitch and an unrelated feature in the torque data. Thus we have not included them in the table.

TABLE II. Glitches with any of the lever arms significantly different from zero. Figures above the horizontal line refer to OR, figures below the same line to CR. Errors correspond to confidence $> 95\%$ (see text).

Δv [pm/s]	Δ [s]	r_ϕ [mm]	r_η [mm]
-8.23 ± 0.03	101.1 ± 0.4	-2.3 ± 0.2	–
1.6 ± 0.2	1560 ± 140	-9.5 ± 3.5	–
22.22 ± 0.01	7.82 ± 0.01	0.5 ± 0.2	-1.2 ± 0.4
-25.68 ± 0.02	2100 ± 5	-1.2 ± 0.7	–

C. Other interferometer channels

1. Low-impulse glitches

We find a significant number of coincidences between the low impulse glitches and the various interferometer channels, $x_{1,\text{OMS}}$, F, and R (see Sec. II).

a. R/F channels. Out of 152 low impulse glitches detected in CR, some correspond to events in the R or F channels, respectively 55 and 4. In the latest part of CR, tens of glitches of this kind appear with non-Poissonian arrival time after the heaters were turned on to bring the temperature back to standard operating conditions. Those events showed a rapid deviation in the R channel, with a characteristic shape lasting some tens of seconds due to the internal control loops of the optical interferometer system [20]. The shape in either R or F is then different from that of Δg .

b. x_1 channel. Some of the low impulse glitches belonging to CR were found to coincide with fast events in the $x_{1,\text{OMS}}$ interferometer, detected as fast features in $\ddot{x}_{1,\text{OMS}}(t)$. Among the 152 two-sided glitches detected in CR, 28 showed counterpart in $\ddot{x}_{1,\text{OMS}}(t)$.

The total number of glitches showing a coincidence in the interferometric channels is 81/152, given that some of them show multiple coincidences. However due to the complexity of the data of all time series, some coincidences may have gone undetected.

Out of the 4 low impulse events detected in OR, 2 showed coincidence to events in the interferometer time series, one in R and the other one in $x_{1,\text{OMS}}$.

2. Impulse-carrying glitches

We find only one coincidence between an impulse-carrying glitch and the other interferometer channels. This is the coincidence between the zero-crossing glitch on the second line of Table I and a spike in the $\ddot{x}_{1,\text{OMS}}(t)$ time series. The spike had no detectable counterpart in the time series of the force commanded onto the space-




FIG. 14. Absolute lever arm $|r_\phi|$, as defined in Eq. (10), as a function of absolute impulse $|\Delta v|$ (left) and duration Δ (right), for impulse-carrying glitches of OR. The black data points represent the glitches for which the lever arm has been found to be significantly different from zero. The green horizontal segments represent the upper bound for $|r_\phi|$ for glitches for which $|\Delta\Omega_\phi| \leq \delta\Omega_\phi$. For 9 of the 98 impulse-carrying glitches of the OR, we were not able to perform the analysis on the $\Delta\gamma_\phi(t)$ data series. The gray horizontal line refers to the 11 mm “electrical” reference, as described in text.

craft by the drag-free control loop, as expected for such a fast feature.

In analogy with the case of the differential measurement, a spike like this might be either due to a spike in $g_1 - G$, with g_1 the force per unit mass on TM1 along x , and G that on the spacecraft, or to a feature in the $x_{1,\text{OMS}}$ interferometer readout. When interpreted as a force, its total impulse, $\Delta v \simeq 0.8 \text{ nm/s}$, is much larger than the corresponding impulse, Δg ($\Delta v \sim 0.8 \text{ pm/s}$, in Table I). This rules out that the spike is in g_1 , and leaves only the options that it might have been due to a force impulse on the spacecraft, or to a feature in the $x_{1,\text{OMS}}$ interferometer readout.

D. Other time series

We also analysed the time series of magnetometers and of the various thermometers of the LTP, without finding any coincidence.

We also inspected the time series of the inertial forces, calculated as explained in [1], and found no coincidence. Note that, as stated in Sec. II B, these series had been subtracted from the observed acceleration to form Δg . Thus a true force glitch in any of these series, should have emerged from the original acceleration series, and should have disappeared in the subtraction. Thus our check was against non-force data artifacts in the inertial force series that may have been transferred to Δg upon subtraction, or against miscalibration in the subtraction, that would have left a residual glitch in Δg .

We also inspected the time series of the forces on the TMs along directions orthogonal to x , but the force sensitivity along these axes is too low, with respect to that

along x , to give any significant information. This reduction stems from both the lack of an interferometric readout on y and z , and from these degrees of freedom being used to control the motion of the spacecraft.

Note that the effects of the acceleration of the spacecraft has already been subtracted from data as illustrated in Sec. II B. In particular, the $\delta_{x_1}\ddot{x}_{1,\text{OMS}}$ term in g_{SC} in Eq. (1) subtracts the effect of spacecraft acceleration being picked up by the finite common mode rejection of the differential interferometer, and the term $(\omega_2^2 - \omega_1^2)x_{1,\text{OMS}}(t)$ subtracts the effect of the “harmonic” coupling to the spacecraft acceleration. Thus our inspection of the $x_{1,\text{OMS}}$ series is a search either for possible disturbances affecting simultaneously both Δg and $x_{1,\text{OMS}}$, or for incomplete subtraction due to calibration errors.

VI. DISCUSSION

We discuss here the implications of the observations described so far for the identification of the possible sources of the observed glitches.

Before we start the discussion, it is useful to give some clarifications on the CR, during which the instrument was operated well outside its nominal working range.

Besides the increased glitch rate, we observed two major effects of such non standard operation mode. First, all interferometric channels of the OMS were characterized by increased noise and spurious signals. Second, the system was subject to a significant mechanical distortion.

We could detect such distortion by monitoring the time series $\Delta x_{\text{GRS}}(t) - \Delta x_{\text{OMS}}(t)$ that measures, to first order, the difference between the relative displacement of the two EHs to that of the two TMs.

During CR this difference kept changing over time, with a total variation of up to a few μm , likely due to the thermal distortion of the mounting structure of the two GRS. This distortion must have put under severe stress the interface between the GRS and the glass structure of the OMS, which, on the contrary, is virtually undistorted by temperature because of its very low temperature coefficient.

A. Fast, low-impulse glitches

We discuss next the fast, low-impulse glitches. The fast time scale (Fig. 2 right), the absence of any feedback force signal, the absence of any net impulse (Fig. 11, lower panel), and, most importantly, the coincidence with events in other interferometer channels for the majority of them, are all features that point, for these glitches, to an explanation as interferometer anomalies.

Note that glitches of this kind were very rare in ordinary conditions (less than one per month), while their production has been boosted by the relatively unstable situation of the CR, in some occasion in the form of clusters that violate the Poisson condition for random occurrence times (Fig. 11, upper panel).

We were not able to trace back the true generating mechanism behind these anomalies. It is however worth mentioning that similar events also showed up in a similar interferometer flown on the GRACE Follow-on mission [21]. The interferometer readout exhibited phase jumps that would translate for us into steps in Δx , which were traced back to mechanical disturbances generated by thrusters activation. Thus mechanical stress may be the root cause of these interferometer anomalies on LPF.

Since the mechanism behind these transients is not fully understood, we cannot predict whether they will occur in LISA in a similar fashion. For instance, while LPF interferometry reached a sensitivity $32 \text{ fm}/\sqrt{\text{Hz}}$ [5], LISA is expected to be limited at $10 \text{ pm}/\sqrt{\text{Hz}}$.

Nevertheless it is worth mentioning that the science degradation resulting from this kind of glitch, if there were to be any in the LISA data, is not expected to be severe, given their very fast time scale affecting just a few data points, and the lack of a low frequency component due to the corresponding lack of impulse.

In addition we note that the largest of these events in the interferometer outputs might be detectable on ground. Though much noisier than in flight, LPF interferometry could demonstrate on ground $\simeq \text{pm}/\sqrt{\text{Hz}}$ sensitivities [22]. This should allow detection and study of, for instance, the largest, tens of pm high, steps in the interferometer output. The effect of environmental conditions on such glitch production could then be investigated.

B. Impulse-carrying glitches

Contrary to fast low-impulse glitches, we were able to demonstrate that a large fraction of the impulse-carrying ones consists of true force events acting on one or both TMs. Though we could not extend the demonstration to the smallest ones, the shape similarity with the larger ones, and the continuity of the parameter distribution, make it very likely that all glitches of this kind, or at least the vast majority of them, consist of true force events.

In the following, after discussing their main statistical features, we discuss the possible sources for such force events, with the aim of ruling out the unlikely ones, and identifying the most likely ones.

1. Glitch taxonomy and some implications for their origins

To gain more insight into the nature of the force glitches, let us now discuss their statistical properties.

First, their time of occurrence has been following Poisson statistics during both OR and CR (Fig. 4). We observed neither any significant clustering nor any repeated pattern. Thus these glitches appear to be due to independent sources and to occur at random times.

In particular, during OR in stable conditions and within the specified operating conditions for the LTP, we observed a constant mean rate of occurrence (Fig. 5) throughout the $\simeq 1.2 \text{ yr}$ of the mission science operations. During this same period of time, the pressure around the TM had been decreasing by almost an order of magnitude [1], and many changes of operational settings had been taking place in between any of the different noise runs we have been discussing here, both to maintain the orbit, and to perform dedicated investigations [13].

Cooling of the system to near 0°C increased this rate by more than one order of magnitude. However such a rate increase was not uniform across the glitch parameter space.

More specifically, Fig. 7 shows two basic features:

- a population of glitches with $\Delta \gtrsim 1.6 \text{ ks}$, which accounts for 37% of all positive glitches in OR, for just 0.7 % in CR, and is absent, for negative impulse glitches, in all runs;
- in both OR and CR, positive impulse glitches constitute the vast majority of all glitches. For convenience we remind here that a positive impulse pushes the TMs one toward the other.

In Table III we list the number of glitches in these three categories: ($\Delta v > 0$, $\Delta < 1.6 \text{ ks}$), ($\Delta v > 0$, $\Delta > 1.6 \text{ ks}$), and ($\Delta v < 0$, $\Delta < 1.6 \text{ ks}$), for both OR and CR. The

table also contains the projected counts one would have observed during the 11.9 d long CR, had the statistics not been affected by the cooldown (see App. C). The result of this projection shows that:

- Counts for positive glitches of duration less than 1.6 ks are more than twenty times larger than the largest projected value. Thus cooldown has strongly increased the rate of these glitches.
- Counts for positive glitches of duration larger than 1.6 ks are compatible with the projected values, and thus the rate of these glitches has not increased upon cooldown.
- Counts for negative glitches, all of which are short, are more than four times larger than the value expected from the OR statistics. Thus cooldown has affected the rate of these glitches too, though to a lesser extent than that of positive shorter glitches.

We have also compared the OR glitch parameter distributions with that for CR, for the categories of glitches that are significantly populated in both type of runs.

Statistical tests on the equality of two multivariate distributions are still subject to a debate that goes well beyond the scope of this paper. To get nevertheless a sense of the similarity between the two distributions, we have performed marginal Kolmogorov-Smirnov tests for Δ , Δv , and $\Delta g_{\max}$. This last parameter is $\propto |\Delta v|/\Delta$, and thus mixes somewhat the populations of the two other parameters. The resulting p -values are reported in Table IV.

Though these considerations are not at all a proof, they are nevertheless suggestive that in OR, the glitches belong to three different families.

- The first is a family of positive impulse glitches with $\Delta \geq 1.6$ ks.

For this family Δv increases rapidly with Δ . For reference, we calculate that the line Δv [pm/s] = $\Delta^2/(0.65 \text{ ks})^2$ is the lowest power-law upper bound to all glitches in this family.

TABLE III. Observed and expected number of glitches listed by duration and impulse sign.

Runs	$\Delta v > 0$	$\Delta v > 0$	$\Delta v < 0$
	$\Delta < 1.6 \text{ ks}$	$\Delta > 1.6 \text{ ks}$	
OR	51	30	17
CR	304	2	28
OR→CR ^a	1-13	0-9	0-6

^a Counts for CR, projected from the observed counts and rate in OR. Intervals correspond to 90% confidence.

Though the shape of distribution might be affected by the detection threshold from below, such quadratic upper bound does not appear to be due to any obvious selection effect.

Remarkably, the rate of these glitches has not been affected by the thermo-mechanical stress conditions of the cooldown.

As the threshold $\Delta \geq 1.6$ ks has no particular physical meaning, this family may also include some of the shorter duration glitches. Actually if the quadratic upper bound above were used to define the family, the family would include 3 more glitches in OR and 2 more in CR. However the conclusions about the count projections would not be modified by these adjustments.

- The second family is composed of positive Δv glitches with $\Delta < 1.6$ ks. This coincides basically with the glitch population of this kind during the CR.

Indeed, given that the rate of the glitches of the long positive impulse family was unaffected by cooldown, even if that family extended to shorter duration, very few samples would contaminate the CR distribution with $\Delta < 1.6$ ks.

About 70% of the glitches within this family have duration less than one minute. However the family also includes a sparse sample of glitches with duration that can take values up to just below the $\Delta = 1.6$ ks threshold.

The rate of these glitches has been greatly affected by the thermo-mechanical conditions of cooldown. Actually, as in CR negative impulse glitches are only about 8% of total, the time/temperature evolution of the rate in Fig. 6 refers basically to the glitches in this family.

- Finally, the third is the relatively small family of negative impulse glitches, again with $\Delta \leq 1.6$ ks.

About 53% of these have duration of less than a minute, and these too include samples with longer duration, approaching $\Delta = 1.6$ ks.

The rate of this family has also been affected by cooldown, though to a lesser extent, but there are

TABLE IV. Resulting p -values from the marginal Kolmogorov-Smirnov test, for parameter distributions in OR and CR.

Category	p_{Δ}	$p_{\Delta v}$	$p_{\Delta g_{\max}}$
$\Delta v > 0, \Delta \leq 1.6 \text{ ks}$	0.26	0.33	0.31
$\Delta v < 0$	0.35	0.28	0.28

not enough samples to assess if the rate has been affected by temperature.

2. Possible physical sources of force glitches

In the light of all the evidence above, we now discuss the possible physical sources of these force glitches.

a. Platform acceleration and inertial forces. An obvious cause of transient “events”, in space-borne differential accelerometers, may be some corresponding events in the acceleration of the spacecraft, which would be picked up because of the finite common mode rejection of the instrument. An obvious example of this is the already mentioned case of a micrometeoroid hit.

In our case we can rule out this source, as we have been correcting the data, as described in Sec. II B, for the coupling to spacecraft motion. In addition we have inspected the $x_{1,OMS}(t)$ data series, to check for any residual coincidence, possibly due to some residual inaccuracy in the data correction, finding none. Similar data correction and inspection also rule out, as sources of glitches, inertial forces due to spacecraft rotation.

b. Thermal effects. Also the lack of coincidence with thermometer readings allows us to rule out some possible explanations.

The correlation of temperature and temperature gradient variations with Δg has been investigated by a series of dedicated experiments during the operation of LPF [23]. We have measured the dependence of Δg on the average temperature $\bar{T}$ of all thermometers on each EH. Upon heating the EH we found a complex behaviour with a relatively prompt response and coefficient $\partial\Delta g/\partial\bar{T}|_p$, plus a delayed response with coefficient $\partial\Delta g/\partial\bar{T}|_d$, likely due to the delayed heating of distant sources. The former was found to be pretty constant over time, at about $0.5 \text{ pm s}^{-2} \text{ K}^{-1}$. The coefficient $\partial\Delta g/\partial\bar{T}|_d$ was instead found to decrease by about a factor 5 over the course of the mission and to level off at $\simeq 0.5 \text{ pm s}^{-2} \text{ K}^{-1}$, paralleling the decrease in pressure.

A similar pressure-dependent behaviour was also, as expected, found for $\partial\Delta g/\partial\Delta T_{EH_i}$, where ΔT_{EH_i} is the difference of temperature across the i -th EH. For both EHs, these coefficients leveled off at $\simeq 10 \text{ pm s}^{-2} \text{ K}^{-1}$.

With such sensitivities, to explain the smallest of the observed Δg glitches with a glitch in $\bar{T}$, one would need amplitudes of order mK at the beginning of mission and of many tens of them at the end. We would have detected such glitches, as our resolution is of order of tens μK in $\bar{T}$ for a 100 s glitch following the template in Eq. (4).

The same applies to a glitch in the differential temperature ΔT_{EH} , for which glitch amplitudes would need again to be of order of mK, and where our sensitivity is in the

μK range thanks to a dedicated low noise temperature differential readout [8].

Therefore such hypothetical temperature glitches would have been detected within their relative time series. We believe that this rules out the hypothesis that glitches may be due to thermal transients in the system.

c. Gravitational signals. The lack of any significant permanent change in the force, upon the occurrence of any of these force events, rules out the possibility that they may consist of the gravitational signals from some amount of mass permanently changing position or leaving the system. This would be the case, for instance, for a large outgassing event from some point of the spacecraft or for some sudden leak of a large amount of propellant.

A different origin of a gravitational signal would be a body reversibly moving around its equilibrium position. An example of that is the 2 kg tungsten mass located within each GRS, at a few centimeters from the TM center to suppress the static gravitational field on the TM [24] (see Fig. 15).




FIG. 15. Schematic representation of the GRS, horizontal section. The yellow square is the Au-Pt TM. The orange hollow square represents the EH, whose mid section carries four symmetric holes, one of which is the input port for the laser beam (in red). The green part is the section of a titanium structure that holds the various elements together, while in gray is the specially shaped tungsten gravitational balance mass. The brown circle is the section of the vacuum chamber, while the cyan rectangle represent the optical window to transmit the laser beam.

The gravitational force gradient per unit mass acting on the TM due the balance mass is $-\omega_{1,2}^2 \simeq 5 \times 10^{-7} \text{ s}^{-2}$ (see Sec. II B) [25]. Thus, to produce our smallest glitches with a peak amplitude about fm s^{-2} , one would need a displacement of the balance mass out of its equilibrium position peaking at $\sim 2 \text{ nm}$ and then getting back to rest.

The balance mass is the largest source of gravitational gradient, the gradient from other sources being signifi-

cantly smaller, as the gradient from the rest of the GRS drops to about $4 \times 10^{-7} \text{ s}^{-2}$, due to compensation [25], and as the contribution of farther apart components decays as the cube of the distance. For instance, the optical bench contributes a gradient of about $6 \times 10^{-8} \text{ s}^{-2}$, the other GRS $3 \times 10^{-8} \text{ s}^{-2}$, and the rest of spacecraft $1 \times 10^{-8} \text{ s}^{-2}$ [26].

Characteristic frequencies of mechanical parts surrounding the TM are in the kHz range. Thus, even the overdamped mechanical motion of those parts takes place on time scales much shorter than that of the vast majority of the observed glitches. In order to produce a glitch by way of this physical mechanism, one would need parts to move because of some transient thermo-mechanical distortion, and not by some mechanically excited free motion.

We foresee two major damaging patterns for such a thermo-mechanical distortion. The first is an expansion or contraction of the GRS around the TM due to a temperature fluctuation. The major effect of such distortion would be to move the balance mass relative to the TM. Given the construction details, this effect adds a term of order $\simeq 0.3 \text{ pm s}^{-2} \text{ K}^{-1}$ to the temperature coefficient $\partial \Delta g / \partial \bar{T}|_p$ discussed above. The possibility of glitches originating this way has been already ruled out.

The second distortion pattern that may give origin to a force transient, is a rigid displacement along x of either of the two GRS relative to its own TM. This is the case, for instance, if the struts that attach the GRS to the LTP bay thermally expand or contract, moving both GRS in opposite directions, and changing their distance by some amount δl .

By moving the main sources of gradient, such distortion would cause both a signal $\Delta g = -\omega_{2,\text{GRS}}^2 \delta l$, with $-\omega_{2,\text{GRS}}^2 \sim -\omega_2^2 \sim 4 \times 10^{-7} \text{ s}^{-2}$, and a signal $x_{\text{OMS}} - x_{\text{GRS}} = \delta l$. As already mentioned, we have observed both such signals upon the large distortion caused by cooldown.

Thus, a glitch originating from such a distortion pattern, or from any pattern that would move any of the GRS relative to its TM, should also show up in the $x_{\text{OMS}} - x_{\text{GRS}}$ series, peaking at $\Delta g_{\text{max}} / (-\omega_{2,\text{GRS}}^2)$, a number that ranges from a few nanometers to micrometers. As discussed in Sec. V A, we have inspected the series $x_{\text{OMS}} - x_{\text{GRS}}$ and found no corresponding glitch.

A distortion pattern moving the spacecraft, or one of its large components, relative to the entire LTP and non-detectable in $x_{\text{OMS}} - x_{\text{GRS}}$ would require much larger amplitudes. The difference of gradient on the two TMs due to the entire spacecraft, which is the quantity that would matter in this case, is of order 10^{-9} s^{-2} . Thus, to produce the observed glitches, the spacecraft would have to have moved along x by an amount ranging from micrometers to millimeters, requiring temperature changes

ranging from a fraction of a K to hundreds of K, without causing any detectable distortion within the LTP. We consider such a scenario to be quite unlikely.

In conclusion, we believe that a gravitational origin is an unlikely explanation for the vast majority of the observed glitches.

d. Magnetic force. It also appears unlikely that these force glitches are explained by some slow transient in the magnetic field. With *slow* here we mean that we are not considering eddy current effects, which we will discuss later.

The magnetic susceptibility of the TMs has been measured to be $\chi \sim 3 \times 10^{-5}$ [27] and its permanent magnetic moment $|\boldsymbol{\mu}| < 5 \text{ nAm}^2$ [28], though this is just an upper limit. The static magnetic field on board LPF was found to be $|\mathbf{B}| \sim 1 \text{ }\mu\text{T}$ [7], and, finally, from the lack of correlation between the magnetic field and force noise [29], we estimate the magnetic gradient to be less than $10 \text{ }\mu\text{T/m}$.

The force peak amplitudes of all observed glitches have $\Delta g_{\text{max}} \geq 10^{-15} \text{ m s}^{-2}$ (see Fig. 10). To reach such a force level, a glitch in the magnetic field gradient, on either of the TM, considering only the coupling to the TM induced moment, should peak at about $1 \text{ }\mu\text{T/m}$. Unless we take into account an unrealistically close source, a gradient like this would have produced a detectable signal on some of the magnetometers, at least for glitch duration larger than about 10 s. As glitch amplitudes were even 1000 times larger than this minimal one, many of the glitches would have produced large magnetometer signals.

Similar conclusions are obtained assuming that the glitch is in the magnetic field. To produce our force glitch with a glitch in the magnetic field, one would require a peak at least $0.1 \text{ }\mu\text{T}$, well above the detection threshold. Thus the lack of observed magnetic glitches rules out this explanation.

A further mechanism for force transient of magnetic origin, is that of currents induced via Seebeck effect by thermal gradients inside the TM. The TM material has a finite amount of Au and Pt precipitate that may create effective thermocouples at grain boundaries. The effect was noticed during magnetic characterization of the TMs on ground, during which temperature differences of order of a few K across the TM, due to manipulation by human hands, induced a magnetic moment peaking at some pA m^2 . With the value for the static gradient quoted above, this effect would just give a correction of $\simeq 0.1 \text{ pm s}^{-2} \text{ K}^{-1}$ to the coefficient $\partial \Delta g / \partial \Delta T_{\text{EH}_i}$ mentioned in the previous section, and its role in producing glitches is ruled out together with the rest of the thermal effects.

In addition to these low frequency effects, magnetic fields at high frequency may induce eddy currents within the TMs and then exert Lorentz forces on them [27].

The effect is thus quadratic and would convert the low frequency amplitude fluctuations of a high frequency magnetic spectral line into a corresponding low frequency force.

To give a scale of the effect, a recent finite-element electromagnetic calculation by the LISA project [30], has shown that the effect of a dipole of 1 mA m^2 located at a distance $d = 20 \text{ cm}$ from the TM and oscillating at the frequency of 100 Hz , would cause a force of $\Delta g \simeq 4 \text{ fm/s}^2$. The effect reaches its peak at 100 Hz , while at lower frequency the induced current decreases and above that the screening effect of the metallic EH attenuates the oscillating field.

The effect of a dipole source decreases with d^{-7} , so that at the closest distances of about 0.4 m between the TM and any active device on the LPF spacecraft the effect might be ~ 100 times smaller.

The spacecraft prime contractor performed a test campaign on ground against audio frequency magnetic lines [31] during LPF development. A few lines have been identified with peak amplitudes $< 1 \text{ nT}$ at the position of the TMs. In the point dipole model at a distance of $d \simeq 0.4 \text{ m}$, each line would be generated by a dipole of $\simeq 0.3 \text{ mA m}^2$ and would exert a static force $\Delta g \simeq 4 \times 10^{-3} \text{ fm s}^{-2}$. Even a glitch consisting of 100% amplitude modulation of any these lines - a behaviour not observed during test - would have then an amplitude orders of magnitude less than those in Fig. 10. However, we had no magnetometer on board sensitive to the audio band and thus we cannot exclude that additional, more intense, amplitude modulated lines had been generated once on orbit, as the operating conditions may have been significantly different from those during testing.

Also, the shape and timescale of the glitches are not easily reconciled with such an AC magnetic origin. Electromagnetic emission from electronics is usually modulated by noise, by switching among different operational settings, and by thermal variations, and we do not see how these may easily follow the reversible exponential behaviour lasting minutes to hours that would be required to generate the observed glitches.

One way though of producing a smooth time evolution is that of two lines of constant amplitude, the frequencies of which would slowly drift over time. If during some time interval these lines had a substantial overlap in the frequency domain, they would indeed generate a force on the TM.

Lines observed during testing were stable in frequency, but, again, we cannot exclude that other lines were present in flight. It seems, however, highly unrealistic that on orbit enough lines have been generated, with different enough drift rates and shapes, to explain the hundreds of glitches of Fig. 7, with parameter values that span a few

orders of magnitudes.

In conclusion, we believe that while the possibility that some of the observed glitches are due to eddy currents cannot be discarded, it is highly unlikely that this source may explain the majority of the observed glitches. Nevertheless we certainly recommend that in LISA a thorough testing is performed on ground, and that on-board diagnostic magnetometers with sensitivity up to 1 kHz are considered.

e. Electrostatic forces from GRS electronics Each LPF TM is electrically isolated from its surrounding with no detectable discharging path. Thus an event of charging because of cosmic rays [10, 32] or any other source of particles, would show up as a step in Δg , quite incompatible with the observed finite-impulse glitches.

In addition, all surfaces facing the TM are conducting and grounded, and would not accumulate free charge.

Still, glitches may be produced by spurious voltage transients in the electronics we use to control the TMs. More specifically, the mentioned electrodes facing the x -faces of the TM, are all driven by separate amplifiers. A voltage glitch in one of these amplifiers would certainly produce a force on the TM.

However, a single-electrode event like this would also generate a torque around z , with a lever arm $|r_\phi| = 11 \text{ mm}$, which is the main reason why we have searched for a torque component to the detected force glitches as explained in Sec. VB.

To span the observed range of glitch peak amplitudes, from 1 fm/s^2 to 1 pm/s^2 , with the electrode geometry, there would need to be a transient change in the mean square voltage at the actuation amplifier output between roughly 10 mV^2 and $10\,000 \text{ mV}^2$. This could occur in different forms:

- A transient “quasi-DC” voltage in the $100 \text{ }\mu\text{V}$ to 100 mV range, mixing with stray DC potential differences of order 100 mV , due to TM charge and/or stray “patch” voltages.
- A transient change in the roughly 1 V actuation audio-frequency carriers [2] by roughly $10 \text{ }\mu\text{V}$ to 10 mV .
- A transient electrode oscillation coherent with the 100 kHz sensing “injection” frequency [6] and mixing with the 0.6 V amplitude TM bias, in roughly this same amplitude range.
- A spontaneous AC oscillation, at some random frequency not associated with the actuation or injection, with amplitude in the 5 mV to 200 mV range.

While we cannot directly exclude any of these – though the 100 kHz excitation would have likely given some capacitive sensing error – they were not detected in dedicated

pre-flight tests, albeit relatively short (less than day per electrode), which could have detected such anomalies. Certainly such features can and should be investigated more thoroughly on ground for the LISA electronics.

A strong indicator that the glitches do not originate in the actuation electronics comes from the analysis of the possible torque component to the observed glitches. The findings in Sec. VB show that:

- There are in total 56 glitches within OR, spread over the entire parameter space, for which we would have been able to detect a lever arm of 11 mm (see Fig. 14). Only one of these is both incompatible with $r_\phi = 0$ and compatible with $r_\phi = 11$ mm (see Table II). Though there is no proof that this glitch is indeed of electrical origin, one might nevertheless take $1/56$ as a rough bound to the fraction of glitches that may be due to this source.
- For the CR glitch data, none of the 147 CR glitches for which there is sufficient resolution to resolve $r_\phi = 11$ mm have such an effective arm. The probability of such an event, using binomial statistics, and assuming the distribution is the same as during OR, is $p = 0.08$, a figure that does not allow us to reject the equal distribution hypothesis. Using both observations, OR and CR, the probability of such an occurrence becomes $p \leq 0.023$ with 95% confidence.
- For the kind of standard, audio-frequency electronics we are discussing here, minute to hour long transients which are not induced by some corresponding thermal transients, are quite unexpected. Of the 121 glitches with $\Delta \leq 1$ min and detectable 11 mm lever arm, none is found to have such a lever arm, which gives $p \leq 0.024$ with 95% confidence.

While this effective arm test is inconclusive for the smallest and fastest glitches, for which our sensitivity to a lever arm is reduced, most of our glitches are incompatible with a single-electrode electrical origin.

In addition to this, even for the smaller, faster glitches, the observed increase in rate upon cooling the spacecraft is not easily reconciled with an electrical origin.

More complex voltage events, simultaneously affecting more than one electrode – such as two adjacent electrodes which combine to give force without torque – are even less likely given the design of the electronics [33].

Testing for the different types of transient voltages that could produce glitches at the levels and rates observed in LPF would require dedicated detection circuitry and long measurements. The voltage levels are however accessible, and such testing in preparation for LISA is recommended.

f. Outgassing environment One candidate source of force is the exchange of momentum between the TM and the gas molecules surrounding it. This exchange, in the form of Brownian noise, dominated the noise budget at frequencies above about 1 mHz [2].

Gas pressure around the TM, that we deduced from the Brownian noise, decayed over the course of the mission, as the vacuum chambers were vented to space via venting ducts. Pressure went from about 10 μ Pa at the beginning of the mission, to about 1 μ Pa toward the end, following a power law function of time, strongly indicative of water outgassing [29].

The vacuum environment of the TM is rather unusual, as the vacuum chamber is densely packed with components: TM, EH, TM launch-lock mechanism, various cable bundles, etc. Thus, the outgassing surface to volume ratio is unusually high for a vacuum system, and the distribution of outgassing surfaces rather non symmetric around the TM.

In such an environment, one possible source of glitches may be an event of release of some metastably trapped gas from pores. Similar events are often observed in vacuum systems, due to so-called virtual leaks - cavities with a high impedance connection to the outside - that may trap gas and release it in bursts [34]. The phenomenon is also known to be triggered by mechanical stress and friction.

We have observed events similar to glitches in the pressure gauge time series during vacuum preparation of the GRS on the ground (see Fig. 16). For that test the GRS had been inserted in a wider vacuum chamber with its venting valve open. The chamber was pumped down and its temperature was raised to and maintained at $\simeq 115^\circ\text{C}$ for about 24 h to get rid of most of the adsorbed water, a standard procedure for vacuum systems known as bakeout.

The stretch containing the glitches in Fig. 16 was observed during final cooldown, during which the system was subject to a significant amount of thermo-mechanical stress due to the relatively rapid contraction. A similar behaviour with many spikes was also observed during the preparation of the other GRS, while no spikes were detected during a test with the empty chamber. However, we have no way of assigning with certainty the source of these gas emission spikes to the GRS interior, and we only show them here as an example of the phenomenon in vacuum systems.

Molecular simulations show that a molecule released from a surface nearby the TM can transfer a net momentum to it before leaving the system via the venting duct. For instance, using the same simulation method and simplified geometry described in [35, 36], we calculate that a water molecule with a Maxwell-distributed momentum, entering the laser port in the center of the




FIG. 16. Blue line and left scale: pressure in one of the two GRS during cooldown, after about 24 h pumping at about 115 °C for vacuum preparation (bakeout) on ground. Red line and right scale: temperature in the test facility. Vacuum preparations were performed by inserting the GRS, with its venting valve open, inside a wider vacuum chamber. The pressure and temperature shown in the figure are those of this wider chamber. The time origin is set at the end of the bakeout phase. (Data courtesy OHB Italia)

x -face of the EH (Fig. 1, 15) and hitting the TM, would exchange with the TM an average momentum per unit mass along x of $\Delta v \simeq -2.4 \times 10^{-22}$ m/s, before leaving the EH through any of its holes. This is about a factor 40 larger than the mean momentum of the distribution $\langle \Delta v \rangle \simeq 5.7 \times 10^{-24}$ m/s, an enhancement due to multiple collisions between molecule and TM, caused by the constrained geometry [35].

Using the figure above for the impulse per molecule: the glitch with largest TM impulse $\Delta v \simeq 1$ nm/s would correspond to roughly 4×10^{12} water molecules (0.13 ng) hitting the TM; the sum of all glitches in OR and during the entire mission, assuming the observed constant rate, would amount to $\simeq 1.5$ ng; those observed during CR would amount to $\simeq 0.15$ ng. These figures constitute a comparatively small amount of molecules, when compared with the total outgassing rate from each GRS, of the order of 100 $\mu\text{g/d}$ at the beginning of the mission and 10 $\mu\text{g/d}$ at the end.

We have also performed some molecular dynamics simulations with *Molflow* [37], on a more realistic model of the GRS. This model includes the real geometry of the EH, the tungsten balance mass, the cables, and many other details. The tool cannot calculate the total momentum transferred to the TM, however it calculates the differential pressure between its opposite x -faces. This is a reasonable proxy for the momentum transferred per unit time, though certainly underestimated, as it neglects the momentum along x transferred to the y and z faces.

We have simulated, as an example, an instantaneous emission of molecules from a point on the side of the

tungsten mass that faces the aperture in the center of the x -face of the EH, identical and opposite to the mentioned laser port (see Fig. 15). This emission indeed creates a glitch in the differential pressure with shape similar to the observed ones. In particular the profile never crosses the line $\Delta g = 0$. The timescale and the details of the profile however depend on the assumed sojourn time of the molecules on the various surfaces, and on other assumptions in the model.

What is independent of these details is the total transferred momentum Δv , that we calculate by integrating the force time profile. With m_m the total mass of the emitted molecules, we find $\Delta v/m_m \sim 0.5 \text{ nm s}^{-1} \text{ ng}^{-1}$. Note that the simulation shows that only 20% of the molecules emitted from the balance mass indeed enters the EH, while the others follow different paths. Thus, given also the intrinsic momentum underestimate of *Molflow*, this result is not inconsistent with that of the simplified simulation above.

The tungsten balance masses are a natural candidate for such gas-burst events. First, a microscope analysis has shown that its sintered material is porous, with micrometer size pores [38]. Second, the sign of the transferred impulse for this source would be positive, as for the great majority of the observed glitches.

Many other components may also trap gas, beginning with the various bundles of cables that connect the electrodes, and the various motors of the launch lock, to their respective electronics. Some of these sources have the proper position to create also negative impulse glitches that indeed we are able to reproduce with *Molflow*.

Note also that simulations show that the lever arm of a gas inflow from the main inlets to the EH, like that coming from the balance mass, is negligible. Thus, for the few observed glitches with non-zero lever arm, one should assume that the gas has been emitted by some source localized inside the EH.

However, no other source creates the same sort of cavity around one of the EH apertures as that created by the special shape of the tungsten balance mass. This is reflected in the fact that, for all other sources, the ratio between the number of molecules hitting the TM and that of those following different paths is always significantly smaller than that for the tungsten balance mass.

One more argument that may support the gas-release interpretation of glitches is the sensitivity of gas emission to thermo-mechanical stress, as is well illustrated by Fig. 16. As said, thermo-mechanical stress accompanied the CR and their highly increased glitch rate.

In particular, Fig. 6 shows that the occurrence rate switched almost reversibly, from the $\lambda \sim 1 \text{ d}^{-1}$ of ordinary runs, to the many tens per day of the CR, when crossing a comparatively narrow temperature range of a

few degrees. In addition the data are also suggestive of some slow transient relaxation at the lowest temperatures. This may indicate that the rate may be following some complex, non linear stress pattern due to the significant differential thermal contraction of the high numbers of equally complex contact interfaces within the GRS.

Though the likelihood of this source looks the highest among those discussed so far, a few aspects remain to be clarified. An important one is the time profile of the glitches and the associated wide distribution of Δ . Most of the spikes in the data of Fig. 16 have an almost instantaneous onset, followed by the decay pattern one would expect in a standard vacuum system: large gas releases saturate the adsorption speed of the chamber solid surfaces and are quickly pumped down, while smaller ones follow a slower decay in quasi-equilibrium with surface adsorption. This is the response one would expect for a fast, virtually instantaneous, release of gas from some pocket.

Our simulations show that the time profile of $\Delta g(t)$, in the case of an instantaneous emission of a group of molecules from a specific source, also consists of a rise followed by a decay. The rise time is due to the distribution of diffusion times from the source to the EH inlet. The decay is due to the diffusion of molecules, in the space between EH and TM, from the inlet to the final exit aperture. The time constant of such decay is substantially fixed and independent of the rise time, even for molecules emitted from inside the EH.

The time scales of both these branches depend on the assumed sojourn time of the molecules on the various surfaces they encounter along their path. For instance the decaying branch is well fitted by an exponential with a time constant ~ 45 times the mean sojourn time of the inside of the EH.

Sojourn times of molecules on metal surfaces depend exponentially on their binding energy. Thus they may vary by orders of magnitude, from 10^{-12} s to more than seconds, depending on the nature and state of the surface, the nature and amount of adsorbed species etc. It is not easy then to find reliable estimates for a specific situation. As the range of possible values is rather wide, by properly selecting the sojourn times in the model, we have been able to reasonably match, in our simulations, the observed shapes of Eq. (3) or (4).

However, once the sojourn time for the EH has been selected, the resulting duration for such particular match is fixed. Indeed, the duration of the decay branch depends only on that choice, and for $\sim 90\%$ of the observed glitches the template is that in Eq. (4), which only contains one time constant, so that the duration of the decay branch fixes the overall duration. The model of an instantaneous release of gas from one source would then not reproduce

the observed large variability of glitch duration.

The consequence of this is that, if glitches are due to gas release, for a large fraction of them their time profile must be dictated by the intrinsic time evolution of the gas release, while it can be limited by diffusion across the GRS only for some of the shortest ones.

We note that standard Fickian diffusion in simple geometries, like from the bottom of a very narrow pit or from the center of a spherical piece of material, does produce a time evolution of the gas outflow very close to that of our observed glitches, and that the time scale for diffusion may indeed be very long, depending on the gas species and the material.

However, we were not able to find in the literature any reference to events of slow gas release from pores or other imperfections. Though there are many qualitative reference to condensation of gas in such kind of defects, to the possibility that it gets released during pump-down, and to slow, diffusion limited gas motion in porous media, we were not able to find a specific measurement on single events showing such slow time evolution. The only hint we have of the possibility of some non instantaneous gas evolution is in very few of the peaks in Fig. 16, showing indeed some minute-long rise times.

It must be noted though that the scale of the events in Fig. 16 is orders of magnitude different in amplitude from the kind of release that would explain our glitches. For instance, the above mentioned *Molflow* simulation shows that the $\simeq 2$ ng of total water molecules emitted from the balance mass, needed to generate the glitch with the largest Δv , would generate a peak pressure of a few nPa at the venting valve, with some reasonable assumption for the sojourn times of molecules, well below the measurement resolution of the figure.

Additionally, we were not able to find any sound explanation for the near-quadratic dependence of the impulse, and hence the number of molecules that have hit the TM, on the duration of the event. Fickian diffusion time scales with the square of the length of the diffusion path, with no explicit dependence on the size of the fluid volume that is diffusing.

As a final difficulty for this interpretation, the 3 glitches in Table I that cross zero, and have non negligible impulse, are hard to fit to this picture, as none of our molecular simulations could reproduce a glitch for which Δg crosses zero.

The two with smallest impulse may still be compatible with some interferometer artifact, as we don't have the resolution in Δx_{GRS} to discriminate against such a case.

The smallest one might also be of electrical origin, as we do not have the resolution to evaluate the proper value of the lever arm.

The largest impulse one, however, is incompatible with

both those options, and would hence remain unexplained by any of the mechanisms we have considered in this section. The observation of a simultaneous event in the x_1 interferometer, never observed for all other impulse carrying glitches, may indeed indicate a different phenomenon, for which we do not have any reasonable model so far. It is worth reminding that leakage of the spacecraft acceleration into Δg has been corrected for, and, in this specific case, the correction was anyway negligible.

The possibility that the glitches have their origin in outgassing, exacerbated by mechanical stress, suggests careful avoidance of thermo-mechanical stress, including operation of the instrument near its integration temperature. Additionally, while the possible outgassing origin requires specific increased attention to the relevant procedures and testing of LISA GRS hardware, it also suggests caution, with additional testing and analysis, in considering any possible design changes.

Even more desirable is a dedicated experimental campaign to study if these gas release events exist for the kind of surfaces and elements that compose the GRS interior. Though the amount of gas released is of order of nanograms, its detection as a pressure transient in a properly designed vacuum system, with a high sensitivity pressure gauge or mass spectrometer, does not seem out of reach. In addition a dedicated experiment is also possible with the torsion pendulums used to test small forces on the TM of LPF and LISA [35]. Such experiments are currently under study.

VII. CONCLUSIONS

In summary:

- We have reasonable confidence that low-impulse glitches are due to rare transients in the interferometer readout and can be kept well under control in ordinary operating conditions.
- The fastest force glitches with less-than-a-minute duration may have different explanations, including electronics events, eddy current transients and outgassing events. For all these possibilities, proper ground testing is possible to both consolidate the understanding and reducing the risk that their rate impacts on LISA data quality.
- For the long, minute-to-hour force glitches, the only credible explanation that seems to match most of the observational evidence is outgassing. The kind of outgassing events that would explain the observation are somewhat different from the outgassing spikes observed in the vacuum system under transient conditions (Fig. 16). Thus, to consolidate this

hypothesis, a dedicated study is needed, and will be performed, also including appropriate experiments.

In conclusion, we are confident that, supported by the indicated tests and studies, the rate of glitches in LISA will be kept to a manageable level. As with other gravitational wave detectors, identification and mitigation of the remaining glitches will be part of the analysis pipelines that yield LISA's astrophysical data products. Initial efforts are already underway to develop and validate glitch mitigation techniques, informed by experience with LISA Pathfinder [16].

ACKNOWLEDGEMENTS

This work has been made possible by the LISA Pathfinder mission, which is part of the space-science program of the European Space Agency.

We thank Paolo Chiggiato and the vacuum, surfaces and coatings group from Conseil Européen pour la Recherche Nucléaire (CERN), for very helpful discussions about the LPF outgassing environment.

The Italian contribution has been supported by Istituto Nazionale di Fisica Nucleare (INFN) and Agenzia Spaziale Italiana (ASI), Project No. 2017-29-H.1-2020 "Attività per la fase A della missione LISA". The UK groups wish to acknowledge support from the United Kingdom Space Agency (UKSA), the Scottish Universities Physics Alliance (SUPA), the University of Glasgow, the University of Birmingham, and Imperial College London. The Swiss contribution acknowledges the support of the Swiss Space Office via the PRODEX Programme of ESA, the support of the ETH Research Grant No. ETH-05 16-2 and the support of the Swiss National Science Foundation (Projects No. 162449 and 185051). The Albert Einstein Institute acknowledges the support of the German Space Agency, DLR. The work is supported by the Federal Ministry for Economic Affairs and Energy based on a resolution of the German Bundestag (FKZ 500Q0501, FKZ 500Q1601, and FKZ 500Q1801). J. I. T. and J. S. acknowledge the support of the US National Aeronautics and Space Administration (NASA). The Spanish contribution has been supported by Contracts No. AYA2010-15709, No. ESP2013-47637-P, No. ESP2015-67234-P, No. ESP2017-90084-P and No. PID2019-106515GB-I00. Support from AGAUR (Generalitat de Catalunya) Contract No. 2017-SGR-1469 is also acknowledged. M. N. acknowledges support from Fundacion General CSIC (Programa ComFuturo). F. R. acknowledges an FPI contract from MINECO. The French contribution has been supported by the CNES (Accord Specific de projet CNES 1316634/CNRS 103747), the CNRS, the Observatoire de Paris and the University Paris-Diderot. E. P. and H. I.

would also like to acknowledge the financial support of the UnivEarthS Labex program at Sorbonne Paris Cité (ANR-10-LABX-0023 and ANR-11-IDEX-0005-02). N. K. would like to thank for the support from the CNES Fellowship.

Appendix A: Estimate of parameter errors

As the fitting procedure in Sec. III is non optimal and non linear, we estimated the fitting parameter covariance from the Cramér-Rao bound [39], assuming the noise is Gaussian, and using the measured PSD of the residuals. For glitches in OR, we have checked that an optimal filter procedure, for which the Cramér-Rao bound becomes an exact estimate, returns parameter values that are in agreement, within the estimated uncertainty, with those found with our non optimal method.

To this aim, once a glitch had been identified and fitted, we have: expanded the fitting function around the best fit parameter values, up to linear terms in the fitting parameters; applied the optimal linear filter method for multi-component signals [40]; calculated the fitting amplitudes; finally propagated these results back to that of the original fitting parameters. The results of both procedures agree within the uncertainty estimated from the Cramér-Rao bound, except for a few outliers.

We were not able to apply the optimal procedure also to the CR data, because of their above mentioned complexity. Thus, for consistency, in the following we use, for both OR and CR, the parameter values resulting from the non optimal, time domain procedure.

Our error estimates represent certainly a lower bound. In particular, given the length of the low pass filter we use for glitch identification, the uncertainty on Δ for the shortest glitches, $\Delta \lesssim 30$ s, may be significantly underestimated. However we stress that none of the results depends critically on the accuracy of such an uncertainty estimate, as parameter fluctuations within the glitch population

are significantly larger than their uncertainties.

Appendix B: Estimation of glitch rate

The rate of impulse-carrying glitches in both OR and subsets of CR is computed with a Bayesian analysis, as follows. According to the Lilliefors test, the distribution of the waiting times t_i is compatible with an exponential distribution. We can then provide an estimate of the rate λ . Under the assumption that each event is an independent random extraction from an exponential distribution with rate λ , we apply Bayes' theorem to the joint probability of n events, with a uniform prior on λ , to get its posterior distribution:

$$\mathcal{L}(\lambda) = \frac{S^{n+1}}{\Gamma(n+1)} \lambda^n e^{-S\lambda}, \quad S = \sum_i t_i \quad (\text{B1})$$

where Γ is the Gamma function. We estimate λ as the probability-maximizing value, and the asymmetric error bounds at 1σ confidence level as the quantiles of the probability distribution containing 68% of it. The result is compatible with the fit of the waiting time distributions, in Fig. 4.

Appendix C: Projection of glitch count rate

The projection of the glitch count rate from OR to CR has been calculated as follows. From the observed counts in OR, we made a Bayesian estimate of the posterior distribution for the probabilities of the three different glitch categories. In addition we assumed a Poisson distribution for the number of glitches during CR, with the posterior distribution for the rate derived from the OR as explained in Sec. IV. We then integrated the probabilities for the counts in the three categories, as calculated from the proper multinomial distribution, over all the posteriors above, to obtain the total probability for those same counts.

-
- [1] M. Armano et al. Beyond the Required LISA Free-Fall Performance: New LISA Pathfinder Results down to 20 μHz . *Phys. Rev. Lett.*, 120:061101, Feb 2018.
 - [2] M. Armano et al. Sub-Femto- g Free Fall for Space-Based Gravitational Wave Observatories: LISA Pathfinder Results. *Phys. Rev. Lett.*, 116:231101, 2016.
 - [3] P. Amaro-Seoane et al. Laser Interferometer Space Antenna. *arXiv:1702.00786 [astro-ph]*, February 2017. arXiv: 1702.00786.
 - [4] P. McNamara, S. Vitale, K. Danzmann and LISA Pathfinder Sci Working Team. LISA Pathfinder. *Classical and Quantum Gravity*, 25(11), 2008.
 - [5] M. Armano et al. Sensor Noise in LISA Pathfinder: In-Flight Performance of the Optical Test Mass Readout. *Phys. Rev. Lett.*, 126:131103, Apr 2021.
 - [6] LISA Pathfinder Collaboration et al. Capacitive sensing of test mass motion with nanometer precision over millimeter-wide sensing gaps for space-borne gravitational reference sensors. *Physical Review D*, 96(6):062004,

September 2017.

- [7] M. Armano et al. Spacecraft and interplanetary contributions to the magnetic environment on-board LISA Pathfinder. *Monthly Notices of the Royal Astronomical Society*, 494(2):3014–3027, 04 2020.
- [8] M. Armano et al. Temperature stability in the sub-milliHertz band with LISA Pathfinder. *Monthly Notices of the Royal Astronomical Society*, 486(3):3368–3379, 04 2019.
- [9] M. Armano et al. Measuring the Galactic Cosmic Ray flux with the LISA Pathfinder radiation monitor. *Astroparticle Physics*, 98:28 – 37, 2018.
- [10] M. Armano et al. Characteristics and Energy Dependence of Recurrent Galactic Cosmic-Ray Flux Depressions and of a Forbush Decrease with LISA Pathfinder. *The Astrophysical Journal*, 854(2):113, February 2018.
- [11] M. Armano et al. LISA Pathfinder platform stability and drag-free performance. *Phys. Rev. D*, 99:082001, Apr 2019.
- [12] L. Ferraioli, M. Hueller and S. Vitale. Discrete derivative estimation in LISA Pathfinder data reduction. *Classical and Quantum Gravity*, 26(9), 2009.
- [13] M. Armano et al. Calibrating the system dynamics of LISA Pathfinder. *Phys. Rev. D*, 97:122002, Jun 2018.
- [14] M.S. Hartig, S. Schuster and G. Wanner. Geometric tilt-to-length coupling in precision interferometry: mechanisms and analytical descriptions. *Journal of Optics*, 24(6):065601, June 2022.
- [15] G. Anderson et al. Experimental results from the ST7 mission on LISA Pathfinder. *Phys. Rev. D*, 98:102005, Nov 2018.
- [16] Q. Baghi et al. Detection and characterization of instrumental transients in LISA Pathfinder and their projection to LISA. *Phys. Rev. D*, 105:042002, Feb 2022.
- [17] J.I. Thorpe et al. Micrometeoroid events in LISA pathfinder. *The Astrophysical Journal*, 883(1):53, sep 2019.
- [18] H.W. Lilliefors. On the Kolmogorov-Smirnov Test for the Exponential Distribution with Mean Unknown. *Journal of the American Statistical Association*, 64(325):387–389, 1969.
- [19] L. Carbone et al. Upper limits to surface-force disturbances on LISA proof masses and the possibility of observing galactic binaries. *Physical Review D*, 75(4), 2007.
- [20] G. Hechenblaikner et al. Digital Laser Frequency Control and Phase-Stabilization Loops in a High Precision Space-Borne Metrology System. *IEEE Journal of Quantum Electronics*, 47(5):651–660, May 2011.
- [21] K. Abich et al. In-Orbit Performance of the GRACE Follow-on Laser Ranging Interferometer. *Phys. Rev. Lett.*, 123:031101, Jul 2019.
- [22] H. Audley et al. The lisa pathfinder interferometry-hardware and system testing. *Classical and Quantum Gravity*, 28(9), 2011.
- [23] M. Armano et al. Temperature-induced forces on geodesic reference masses for LISA: results from LISA Pathfinder. *In preparation*, 2022.
- [24] M. Armano et al. Constraints on lisa pathfinder’s self-gravity: design requirements, estimates and testing procedures. *Classical and Quantum Gravity*, 33(23):235015, 2016.
- [25] P. Sarra. LISA Technology Package Core Assembly Gravitational Analysis Report. Technical Report S2-CGS-RP-3001. Iss09, OHB-CGS, 2015.
- [26] H.D. Wealthy and W. C. Spacecraft-level gravity model results. Technical Report S2.ASU.TN.2270.ISS02, EADS-Astrium, 2008.
- [27] F. Antonucci et al. From laboratory experiments to LISA Pathfinder: achieving LISA geodesic motion. *Classical and Quantum Gravity*, 28(9):094002, 2011.
- [28] P. Sarra. TM FM3 Magnetic Moment Measurement Report. Technical Report S2-CGS-TR-3030, OHB-CGS, 2012.
- [29] M. Armano et al. In depth analysis of LISA Pathfinder performance results: time evolution, noise projection, physical models and implications for LISA. *In preparation*, 2022.
- [30] V. Ferroni and C. Zanoni. personal communication.
- [31] H.D. Wealthy and T. C. Prediction of spacecraft magnetic field based on recent testing. Technical Report S2.ASU.TN.2523.ISS03, EADS-Astrium, 2011.
- [32] M. Armano et al. Charge-induced force noise on free-falling test masses: Results from lisa pathfinder. *Phys. Rev. Lett.*, 118:171101, Apr 2017.
- [33] M. Armano et al. Analysis of the accuracy of actuation electronics in the Laser Interferometer Space Antenna Pathfinder. *Review of Scientific Instruments*, 91(4):045003, 2020.
- [34] D.M. Hoffman, B. Singh and J.H. Thomas, editors. *Handbook of vacuum science and technology*. Academic Press, San Diego, CA, 1998.
- [35] A. Cavalleri et al. Increased brownian force noise from molecular impacts in a constrained volume. *Phys. Rev. Lett.*, 103:140601, Sep 2009.
- [36] L. Carbone et al. Thermal gradient-induced forces on geodesic reference masses for lisa. *Physical Review D*, 76(10), 2007.
- [37] R. Kersevan and M. Ady. Recent Developments of Monte-Carlo Codes Molflow+ and Synrad+. In *Proc. 10th International Particle Accelerator Conference (IPAC’19)*, Melbourne, Australia, 19-24 May 2019, number 10 in International Particle Accelerator Conference, pages 1327–1330, Geneva, Switzerland, Jun. 2019. JACoW Publishing. <https://doi.org/10.18429/JACoW-IPAC2019-TUPMP037>.
- [38] E. Castelli. *LISA Pathfinder noise performance results: disturbances in the sub-mHz frequency band and projection to LISA*. PhD thesis, University of Trento, 2020.
- [39] S.M. Kay. *Fundamentals of statistical signal processing: Practical algorithm development*, volume 3. Pearson Education, 2013.
- [40] A. Ortolan, G. Vedovato, M. Cerdonio and S. Vitale. Optimal reconstruction of the input signal in resonant gravitational-wave detectors - data-processing algorithm and physical limitations. *Physical Review D*, 50(8):4737–

4743, 1994.