Precision charge control for isolated free-falling test masses: LISA pathfinder results

M. Armano,[1] H. Audley,[2] J. Baird,[3,4] P. Binetruy,[4,*] M. Born,[2] D. Bortoluzzi,[5] E. Castelli,[6] A. Cavalleri,[7] A. Cesarini,[8] A. M. Cruise,[9] K. Danzmann,[2] M. de Deus Silva,[10] I. Diepholz,[2] G. Dixon,[9] R. Dolesi,[6] L. Ferraioli,[11] V. Ferroni,[6] E. D. Fitzsimons,[12] M. Freschi,[10] L. Gesa,[13] D. Giardini,[11] F. Gibert,[6] R. Giusteri,[6] C. Grimani,[8] J. Grzymisch,[1] I. Harrison,[14] G. Heinzel,[2] M. Hewitson,[2] D. Hollington,[3,†] D. Hoyland,[9] M. Hueller,[6] H. Inchauspé,[4,15] O. Jennrich,[1] P. Jetzer,[16] N. Karnesis,[4] B. Kaune,[2] N. Korsakova,[17] C. J. Killow,[17] L. Liu,[6] I. Lloro,[13] J. A. Lobo,[13,*] J. P. López-Zaragoza,[13] R. Maarschalkerweerd,[14] F. Mailland,[18] D. Mance,[11] V. Martín,[13] L. Martin-Polo,[10] F. Martin-Porqueras,[10] J. Martino,[4] I. Mateos,[13] P. W. McNamara,[1] J. Mendes,[14] L. Mendes,[10] N. Meshksar,[11] M. Nofrarias,[13] S. Paczkowski,[2] M. Perreur-Lloyd,[17] A. Petiteau,[4] M. Pfeil,[3,19] P. Pivato,[6] E. Plagnol,[4] J. Ramos-Castro,[20] J. Reiche,[2] D. I. Robertson,[17] F. Rivas,[13] G. Russano,[6] G. Santoruvo,[18] P. Sarra,[18] D. Shaul,[3] J. Slutsky,[21] C. F. Sopuerta,[13] T. Sumner,[3] D. Texier,[10] J. I. Thorpe,[21] C. Trenkel,[3,22] D. Vetrugno,[6] S. Vitale,[6] G. Wanner,[2] H. Ward,[17] S. Waschke,[3] P. J. Wass,[3,15] W. J. Weber,[6] L. Wissel,[2] A. Wittchen,[2] and P. Zweifel[11]

(LISA Pathfinder Collaboration)

[1]*European Space Technology Centre, European Space Agency, Keplerlaan 1, 2200 AG Noordwijk, The Netherlands*
[2]*Albert-Einstein-Institut, Max-Planck-Institut für Gravitationsphysik und Leibniz Universität Hannover, Callinstraße 38, 30167 Hannover, Germany*
[3]*High Energy Physics Group, Physics Department, Imperial College London, Blackett Laboratory, Prince Consort Road, London SW7 2BW, United Kingdom*
[4]*APC, Univ Paris Diderot, CNRS/IN2P3, CEA/lrfu, Obs de Paris, Sorbonne Paris Cité, France*
[5]*Department of Industrial Engineering, University of Trento, via Sommarive 9, 38123 Trento, Italy and Trento Institute for Fundamental Physics and Application/INFN, 38123 Povo, Trento, Italy*
[6]*Dipartimento di Fisica, Università di Trento and Trento Institute for Fundamental Physics and Application/INFN, 38123 Povo, Trento, Italy*
[7]*Istituto di Fotonica e Nanotecnologie, CNR-Fondazione Bruno Kessler, I-38123 Povo, Trento, Italy*
[8]*DISPEA, Università di Urbino "Carlo Bo", Via S. Chiara, 27 61029 Urbino/INFN, Italy*
[9]*The School of Physics and Astronomy, University of Birmingham, Birmingham, United Kingdom*
[10]*European Space Astronomy Centre, European Space Agency, Villanueva de la Cañada, 28692 Madrid, Spain*
[11]*Institut für Geophysik, ETH Zürich, Sonneggstrasse 5, CH-8092 Zürich, Switzerland*
[12]*The UK Astronomy Technology Centre, Royal Observatory, Edinburgh, Blackford Hill, Edinburgh EH9 3HJ, United Kingdom*
[13]*Institut de Ciències de l'Espai (CSIC-IEEC), Campus UAB, Carrer de Can Magrans s/n, 08193 Cerdanyola del Vallès, Spain*
[14]*European Space Operations Centre, European Space Agency, 64293 Darmstadt, Germany*
[15]*Department of Mechanical and Aerospace Engineering, MAE-A, P.O. Box 116250, University of Florida, Gainesville, Florida 32611, USA*
[16]*Physik Institut, Universität Zürich, Winterthurerstrasse 190, CH-8057 Zürich, Switzerland*
[17]*SUPA, Institute for Gravitational Research, School of Physics and Astronomy, University of Glasgow, Glasgow G12 8QQ, United Kingdom*
[18]*OHB Italia SpA, Via Gallarate 150, 20151 Milano, Italy*
[19]*Fakultät Elektrotechnik und Informatik, Hochschule Ravensburg-Weingarten, Doggenriedstraße, 88250 Weingarten, Germany*
[20]*Department d'Enginyeria Electrònica, Universitat Politècnica de Catalunya, 08034 Barcelona, Spain*
[21]*Gravitational Astrophysics Lab, NASA Goddard Space Flight Center, 8800 Greenbelt Road, Greenbelt, Maryland 20771 USA*
[22]*Airbus Defence and Space, Gunnels Wood Road, Stevenage, Hertfordshire SG1 2AS, United Kingdom*

*Deceased.
†Corresponding author.
d.hollington07@imperial.ac.uk

The LISA Pathfinder charge management device was responsible for neutralizing the cosmic-ray-induced electric charge that inevitably accumulated on the free-falling test masses at the heart of the experiment. We present measurements made on ground and in flight that quantify the performance of this contactless discharge system which was based on photoemission under UV illumination. In addition, a two-part simulation is described that was developed alongside the hardware. Modeling of the absorbed UV light within the Pathfinder sensor was carried out with the GEANT4 software toolkit and a separate MATLAB charge transfer model calculated the net photocurrent between the test masses and surrounding housing in the presence of AC and DC electric fields. We confront the results of these models with observations and draw conclusions for the design of discharge systems for future experiments like LISA that will also employ free-falling test masses.

I. INTRODUCTION

The space-based gravitational wave observatory, the Laser Interferometer Space Antenna (LISA), will open up a completely new observational window on the Universe. By escaping noise sources associated with the terrestrial environment that limit ground-based detectors, low-frequency astronomical sources such as merging super-massive black holes, extreme mass-ratio inspirals and thousands of extended compact binary systems will be observed [1]. As a technology demonstrator for LISA, LISA Pathfinder was a European Space Agency mission that operated between March 2016 and July 2017. It successfully demonstrated that the residual acceleration between two free-falling test masses could be reduced to $\mathrm{fm\,s^{-2}}/\sqrt{\mathrm{Hz}}$ levels in a frequency range of 0.02–30 mHz, essential for a million-km scale gravitational-wave observatory [2,3].

A subset of forces contributing to this residual acceleration are related to charge accumulating on the test masses as a consequence of the high-energy charged particles present in the space environment [4,5]. Though the structures surrounding the test masses offer shielding from particles with < 100 MeV/nucleon [6,7], cosmic rays with higher energy inevitably deposit charge. A charged test mass then interacts with the conducting surfaces that surround it via the Coulomb force, or, at a much smaller level, with any residual magnetic fields via the Lorentz force [8]. Charging is caused by particle radiation from two main sources: galactic cosmic rays (GCRs) and solar energetic particles (SEPs) [9,10]. Simulations predicted that the ever present but variable flux of GCRs would charge the test masses positively at a rate of $+10$ to $+100$ e s^{-1}, while transient SEP events occurring at most a few times per year could enhance the charging rate by several orders of magnitude [6]. In addition, a significant net charge could also be transferred to the test masses during their initial release as well as after any subsequent regrabbing maneuvers. Due to work function differences between surfaces in last direct mechanical contact as well as triboelectric processes, this was predicted to be of order 10^8 e.

To mitigate the effects of test-mass charging, Pathfinder included a Charge Management System (CMS) as part of the LISA technology package (LTP) instrument. While other sensitive space-borne accelerometers, such as CHAMP [11], GRACE [12], GOCE [13] and MICROSCOPE [14] have relied on a physical electrical connection to control the charge of the test mass, in order to reach the desired force noise goals for LISA, a contactless method is essential. The Pathfinder CMS exploited the same principles of photoemission under UV illumination as successfully demonstrated in the Gravity Probe B (GP-B) mission [15]. In GP-B, bidirectional charge control was achieved by applying a strong local electric field at the point of illumination. However, in order to allow the possibility for full-precision acceleration measurements to continue during discharging, the Pathfinder system was designed to achieve charge control by differential illumination of the test mass and surrounding housing, avoiding large DC electrostatic fields.

In simple terms, such a system produces positive charging when the test mass is illuminated and photoelectrons are removed, while negative charging occurs when the test mass surroundings are illuminated and photoelectrons migrate to the test mass. However, a real system involves a number of complicating factors. The gold sensor surfaces are highly reflective at the UV wavelengths required for photoemission and no matter the illumination this results in UV light being distributed on both test mass and housing surfaces. Furthermore, the photoelectron yield can vary significantly due to unavoidable surface contamination during instrument assembly, even between surfaces that were prepared using an identical method. Finally the energy of the liberated electrons are of order 1 eV so that the volt-scale AC and DC voltages used for test mass capacitive sensing and actuation can strongly influence the discharging behavior.

The structure of the paper is as follows. After first describing the Pathfinder CMS hardware and the details of the discharging problem, we present a computer simulation of the system. This model is then used to interpret measurements made on ground and in flight with a comparison of the underlying surface properties being made. We conclude by discussing the lessons learned that will aid the development of a CMS for LISA.

II. THE LISA PATHFINDER CHARGE MANAGEMENT SYSTEM

Delivered by Imperial College London, the Pathfinder CMS consisted of three major hardware components, shown in Fig. 1. The UV Lamp Unit (ULU) housed six low-pressure mercury lamps (UVP, model 11SC-2) along with their associated electronics. Such lamps produce a spectrum of emission lines but only the one at 253.7 nm (4.89 eV) was used for discharging. A bandpass filter blocked deep UV light to protect optical elements from solarization while also preventing excess visible photons from entering the sensor. A system of lenses were used to collect and focus the light from an aperture into the Fiber Optic Harness (FOH). The FOH was a series of custom-made, UV-transparent optical fiber assemblies which routed the light from the ULU in the main body of the spacecraft to the two vacuum enclosures containing the gravitational reference sensors (GRS). The routing path for the fibers required two FOHs, joined at the interface of the thermal enclosure of the LTP. At the join, an attenuating spacer was used to adjust the level of light reaching the sensor with the length of each spacer being specific to each lamp. Each FOH assembly consisted of a bundle of 19 individual 200 μm diameter fibers, giving an effective diameter of over 1 mm but allowing a minimum bend radius of 25 mm. Finally the Inertial Sensor UV Kit (ISUK) was a custom-made titanium ultra-high vacuum feed-through with a 1 mm diameter, 75 mm long optical fiber that delivered the light into the sensor. Each GRS had three ISUKs that entered at the corners of the lower z-face of the sensor housing. Two ISUKs were directed toward the housing and one toward the test mass.

Each lamp housing contained a thermistor, an Ohmic heater and a UV-sensitive silicon carbide photodiode (ifw optronics, model JEC 0.1 s) for monitoring the UV power output from the lamp. During on-ground calibration of the instrument, the relationship between each lamp's photodiode reading and the UV power emitted from its ISUK was

measured. The photodiode reading was then used in flight to estimate the photon flux entering the sensor. The UV power entering the sensor was commandable in a range of order 1–100 nW with 256 output settings.

The Pathfinder test masses were 1.93 kg, 46 mm cubes made from 73:27 gold-platinum alloy. Each mass was surrounded by a cuboidal housing made from molybdenum accommodating 12 electrodes for sensing and actuation in six degrees of freedom. The gaps between electrode and test mass ranged from 2.9–4.0 mm depending on the axis. Test mass actuation was achieved with audio-frequency sinusoidal voltages on the sensing electrodes with maximum amplitudes of up to 12 V. The y and z housing faces also had an additional six "injection" electrodes, with 4.88 V applied at 100 kHz to capacitively bias the test mass at 0.6 V for position sensing [16]. The total test mass capacitance was 34.2 pF and this value will be used throughout the rest of the paper to convert the test mass charge to a test mass potential with respect to the electrically grounded housing.

Almost all of the test mass, housing and electrode surfaces were sputter-coated with gold except for features related to the caging system that held the test masses during launch [17]. Eight iridium caging fingers entered each sensor via holes at the corners of both z-faces. A concave hemisphere at the tip of each finger was designed to fit over an adjacent dome on each test mass corner and prevent slippage when the caging system was engaged. Both the fingers and the area where they made mechanical contact with the test mass remained uncoated in order to avoid the risk of the gold pealing during release. Each sensor also had two gold-platinum plungers at the center of each z-face that formed the Grabbing Positioning and Release Mechanism (GPRM) used to release the test masses into free fall. Once in space the caging fingers were permanently retracted 2.25 mm below the level of the housing z-face with a one-shot mechanism. Upon initial test mass release the GPRM plungers were retracted 6.5 mm below the housing z-face with the plungers returning to the same position following any regrabbing maneuvers. A flight test mass and electrode housing prior to integration can be seen in Fig. 2.

To measure the test mass charge, sinusoidally varying voltages at millihertz frequencies were applied to combinations of sensing electrodes. This induced a force or torque on the test mass with an amplitude that was proportional to the test mass charge. Typically this was done on the x or ϕ axes using the interferometer to measure the differential test mass motion or rotation with a higher precision than could be achieved with the capacitive sensor [18]. To limit the coupling to fluctuations in both stray actuation voltages as well as from surface patch effects [5,19], the test mass charge needed to be kept below approximately 2×10^7 elementary charges (equivalent to about ± 100 mV). During the mission this was primarily achieved by employing a so-called fast discharge scheme.

FIG. 1. Left: The flight ULU with external heaters attached. The ends of the six optics barrels are seen temporarily sealed with gold dust caps. Top right: An end-on view of an FOH with the 19 individual cores visible. Bottom right: A flight ISUK.
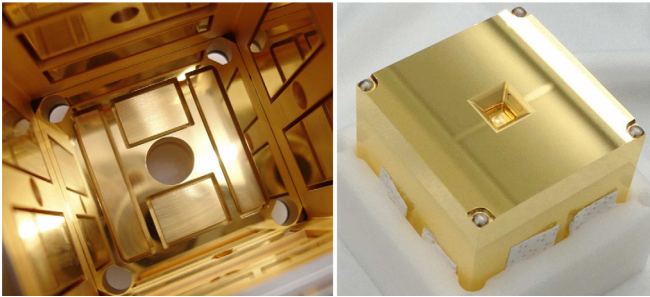
FIG. 2.　Left: A view up into the sensor housing, centered on the upper *z*-face. The test mass and lower *z*-face have yet to be integrated and the ISUK entry points are just out of shot. The two *x*-faces are on the left and right of the image while the *y*-faces can be seen at the top and bottom. The empty recesses where four of the caging fingers enter can be seen in the corners and the plunger entry hole can be seen in the center. The electrodes set within the walls are clearly visible. Right: One of the flight test masses with the lower half temporarily held in a protective case. Note the caging corner domes with the underlying gold-platinum surface exposed. Credit: OHB Italia SpA.

The test masses were left to charge gradually for predetermined intervals (typically a couple of weeks) before being discharged relatively quickly (typically tens of minutes). The discharge was performed either by a preconfigured command sequence uploaded from ground or by autonomous onboard software using an iterative sequence of measurement and illumination. An alternative approach, referred to as continuous discharge, was tested on several occasions during the mission. This scheme involved keeping the test mass charge permanently close to zero by continuously illuminating the sensor such that the charging-discharging rates were balanced.

A. Discharge properties

The performance of the Pathfinder CMS depended on the photoelectric properties of the predominantly gold coated GRS surfaces. Deposited in vacuum and measured *in situ*, the work function of pure gold is 5.2 eV [20]. However, upon exposure to air, stable adsorbates generally reduce the work function to around 4.2 eV [21]. It is therefore only the presence of surface contamination that makes photoemission possible with the 4.89 eV photons produced by the mercury lamps. The adsorbates are bound to the surface of the gold sufficiently well to persist in high vacuum and with moderate heating [22]. With 4.89 eV photons, differences in the surface contamination can produce values for the gold quantum yield (the number of photoelectrons emitted per absorbed photon) that vary between 10^{-6} and 10^{-4} [23]. Of the remaining exposed surfaces within the sensor, iridium has a work function ranging from 5.4–5.8 eV [24] and should therefore not have contributed to discharging currents in the Pathfinder system. Limited photoelectric measurements of 73∶27

gold-platinum suggest properties similar to those of gold [25].

Due to reflections within the sensor, when primary surfaces were illuminated some UV light was inevitably absorbed by opposing surfaces. This resulted in two competing photocurrents, one in the desired direction and one acting against it. For instance, if the test mass was primarily illuminated a photocurrent from the test mass toward the electrode housing was generated but there was also a counter-photocurrent from the housing surfaces toward the test mass. The magnitude of these two photocurrents depended on the amount of light absorbed and the quantum yield of the surfaces. As the photoelectrons had energies of around 1 eV they were also strongly affected by the volt-scale AC and DC voltages applied to the electrodes as well as any potential difference between the test mass and electrically grounded housing.

To assess the discharging behavior in flight only the apparent yield could be obtained, which is defined as the net change in test mass charge per photon injected into the sensor. We define a positive apparent yield as that which produced a positive current to the test mass—a net flow of photoelectrons from the test mass to the surrounding housing surfaces. A negative apparent yield is then due to a net flow of photoelectrons from the housing to the test mass. Note that the apparent yield differs from the quantum yield (the number of photoelectrons emitted per absorbed photon, for an individual surface), and is completely dependent on the state of the system at the time of illumination. The apparent yield varied with the lamp used, applied DC biases, the AC actuation scheme and with test mass charge. With the system in a particular actuation state an entire discharge curve for an individual lamp could be measured, essentially showing how the apparent yield varied with test mass potential (charge).

Figure 3 shows idealized discharge curves for both test mass and housing illuminations. For each illumination two saturation levels are observed. At large negative potentials all the photoelectrons originating from the test mass are able to flow away from it, while none from the housing have enough energy to overcome the potential difference and reach the test mass. At large positive potential differences the converse is true. The shape of the transition between the two saturation levels is dependent on the photoelectron energy distributions of both contributions. As a net photocurrent flows the test mass charge changes and the potential moves toward equilibrium. At this point the two opposing photocurrents balance and the apparent yield is therefore zero. We refer to this as the equilibrium potential. Finally, note that at a test mass potential of 0 V illuminating the test mass produces a positive apparent yield (net electrons away from test mass) while illuminating the housing produces a negative apparent yield (net electrons toward test mass). Continuous discharge in science mode required a negative
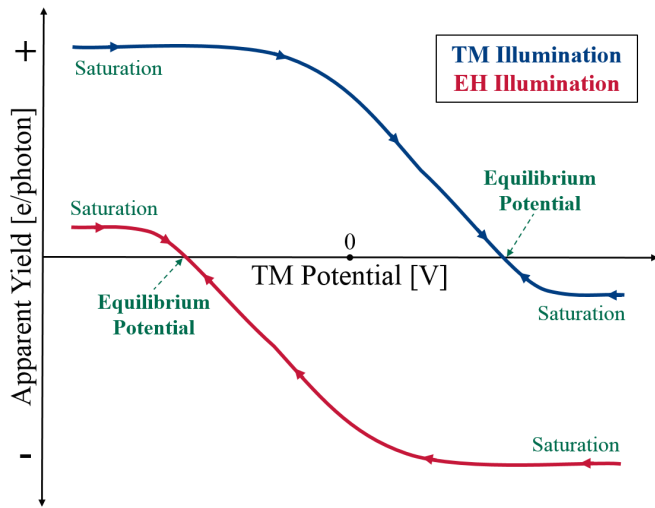
FIG. 3.   A cartoon of idealised discharge curves for both test mass (TM) and electrode housing (EH) illuminations.

apparent yield (and ideally bipolar capability) in order to keep the test mass potential around a set point of 0 V while countering the positive environmental charging.

During fast discharges the curves can effectively be shifted by applying DC voltages to the sensing electrodes which bias the test mass. By illuminating only when biases are temporarily applied, such a technique can aid the desired direction of discharge. It should be noted however that care needs to be taken when applying DC voltages to particular electrodes as if a significant number of photo-electrons are generated in their local region the shape of the curve can be unintentionally changed in addition to causing a shift.

In these idealized examples presented in Fig. 3, for a particular illumination the magnitude of the saturation level is considerably higher in the intended direction of discharge than the reverse direction. In reality this is not necessarily the case. Both the geometry and reflective properties can conspire such that opposing surfaces absorb similar amounts of light leading to similar photocurrent magnitudes. Alternatively, even if the amount of light in the desired direction dominates this can be undermined if the quantum yield of the opposing surfaces are significantly higher. Such a scenario can prevent bipolar discharge at 0 V without the aid of DC biases.

Measurements in 2010 using the torsion pendulum facility at the University of Trento [26], with a prototype GRS electrode housing and a gold-coated test mass, demonstrated the difficulty of negatively charging a neutral test mass without applying additional voltages [27]. The results were initially surprising in that bipolar discharge had been demonstrated successfully in several previous tests with similar gold-coated prototype GRS hardware [28,29]. The asymmetry of the discharge process was attributed to an imbalance in the quantum yields of the specific test mass and electrode housing surfaces under test [25].

Prior to flight no system had been observed where fast discharge was impossible: DC biases applied to the sensing electrodes had always managed to shift the test mass potential to sufficiently suppress the unwanted photocurrent. However, application of DC biases was not desirable in science mode and this would have made a continuous discharge mode unusable. Efforts to create more reproducible properties by baking surfaces under vacuum led to mixed results [22].

There were obvious concerns how such unpredictable variations in the quantum yield could affect the flight discharge system and a dedicated study looked at various ways of mitigating the risk. Possibilities included doping patches of the gold surface to produce higher but more predictable quantum yields or adding mirrors to the ISUK tips in order to redirect the light. However, the additional risks to the primary mission with both options were deemed too high and a more pragmatic solution was agreed upon. The handling and storage of sensor surfaces were strictly controlled and at final integration on ground the discharge properties were measured before and after the system underwent a low temperature bake-out under vacuum. The system then remained sealed but not actively pumped until it was vented to space when in orbit.

III. DISCHARGE SIMULATION

The complexity of the problem and the desire to assess differences between ground and in-flight measurements motivated the need to simulate the system. The first full simulation was created at Imperial College London and was described in Ref. [25] with another full model of the system developed independently by Airbus Defence and Space [30]. The updated Imperial College simulation described here consists of two parts. First, the propagation and absorption of the UV light within the sensor is modeled with a ray-tracing algorithm. The ray-trace output can then be fed into the second part of the simulation which approximates the sensor surfaces as a series of parallel plates. It calculates how many photoelectrons can flow given their energy distribution and the instantaneous electric fields present, before stepping in time and reevaluating.

A. Ray trace

The ray trace is written within the GEANT4 framework which is more commonly used to simulate the passage of high-energy particles through matter using Monte Carlo techniques [31–33]. It also includes a parametrizable microfacet-based reflection model and the ability to simulate arbitrary geometries and light sources. Using GEANT4 we model the propagation of the UV light within the sensor and calculate the percentage of total light absorbed by each distinct surface. As we shall see in the next section, we are concerned with which regions experience different instantaneous electric fields due to their effect on the emitted
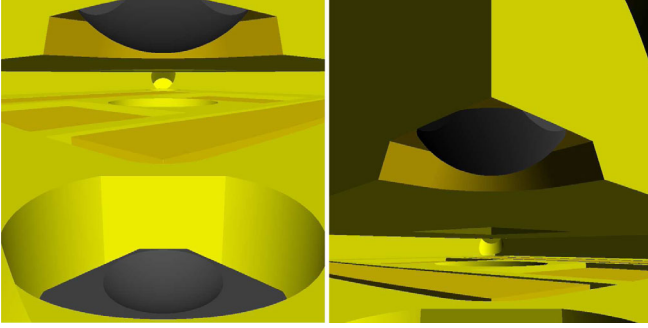
FIG. 4. Left: A view of the sensor as seen from an ISUK aimed at the housing. A retracted caging finger can be seen in the foreground with a concave hemisphere visible at its tip. With the finger engaged, this feature fit over the dome structure that can be seen just above on the corner of the test mass. Right: A view of the sensor as seen from an ISUK aimed at the test mass. The uncoated corner dome is clearly visible as well as part of the *x*-, *y*- and *z*-faces of the test mass.

photoelectrons. As such we end up defining relatively large areas where we calculate the percentage of absorbed light, e.g., each sensing electrode, where a table of such totals forms the ray-trace output.

The model includes the full flight geometry with the option of having the central plungers either engaged or retracted. Any of the six ISUKs can be used as a light source with the emitted light cone from each having been experimentally measured. It was found that the emission cone angle was strongly dependent on the length of the spacer chosen for each optical chain causing the cone half angles to vary from $7.4°$ to $19.2°$. The ISUKs themselves entered the sensors in the corners of the lower *z*-face and were angled $\pm20°$ to the *z*-plane. All results presented here used 10^7 randomly generated initial rays with this number producing results repeatable to two decimal places for each ISUK illumination. Figure 4 shows a view of the simulated sensor, as seen by either an ISUK aimed at the electrode housing or test mass.

The sensors contained three types of surface material: gold (the vast majority of surfaces by area), iridium (retracted corner caging fingers), and gold-platinum alloy (uncoated test mass corner domes and retracted central plungers). The refractive indices (which determine the specular reflection properties) of each material take literature values [34,35]. In the case of gold, a Pathfinder sample surface had also been measured and found to be in good agreement with literature values [25]. The surfaces within the sensor had a range of roughnesses that determine how light was scattered after undergoing reflection. These surface properties are summarized Table I.

An approximate criterion for the appearance of nonspecular behavior is when the scale of surface features are comparable or greater than the wavelength of illuminating light [36]. Given the scale of the surface roughness and that the UV light had a wavelength of 254 nm, reflection from

TABLE I. The different surface types within the sensor. The parameter R_a describes the surface roughness and is defined as the arithmetical mean deviation about the surface's mean height. While this parameter was measured for some flight surfaces only estimates were available for others.

Surface	Material	R_a	Comments
Test mass sides	Au	0.6 to 0.8 nm	Measured
Test mass domes	$Au_{0.7}Pt_{0.3}$	<250 nm	Estimated
Sensing electrodes	Au	30 to 40 nm	Measured
Injection electrodes	Au	50 to 80 nm	Measured
Electrode housing	Au	20 nm	Measured
Caging fingers	Ir	<200 nm	Estimated
Central plungers	$Au_{0.8}Pt_{0.2}$	<200 nm	Estimated

the test mass sides, electrodes and housing surfaces were expected to be specular. Depending on the final roughness, the test mass corner domes, caging fingers and central plungers may have had a diffuse component but were still expected to be dominated by specular reflection. The ray trace has the capability of modeling a nonspecular lobe component using a microfacet model described with a Gaussian angular distribution and parametrized by σ_α. However, without a direct measurement it is difficult to predict an appropriate σ_α and in any case only estimated R_a values were available. To assess the impact of these unknowns the ray trace was run with a range of reflection properties for the rougher surfaces to determine how the absorption percentages would be affected. Even when the surfaces were described as completely diffuse reflectors (an unrealistic extreme) the final results only differed from the completely specular case at the level of a few percent. For simplicity we present here the model results that consider all surfaces as completely specular, but this approximation should have little effect on the final results. Likewise, the effect of uncertainties in the refractive indices, sensor geometry, and ISUK output distributions are difficult to accurately quantify although adjusting each by feasible amounts has shown that none change the results significantly.

The output of the ray trace using the flight geometry and all six ISUK illuminations are shown in the Appendix with a color-map representation of the data for lamps 00 and 02 (which pointed at the test mass and electrode housing, respectively) shown in Fig. 5. The table shows that the majority of the light is absorbed by structures on the negative *z*-faces of the test mass and housing and there is effectively no absorption by positive *z*-electrodes for any illumination. For both test mass illuminations there is also a significant amount of light (24% and 36%) absorbed on the *x* and *y*-faces. Notice also the difference in results for ISUKs pointing in the same direction which is predominantly caused by each ISUK having a different light emission cone. What is most striking is that upon direct illumination of the housing around 60% of the total UV light is absorbed within gaps between the electrode and
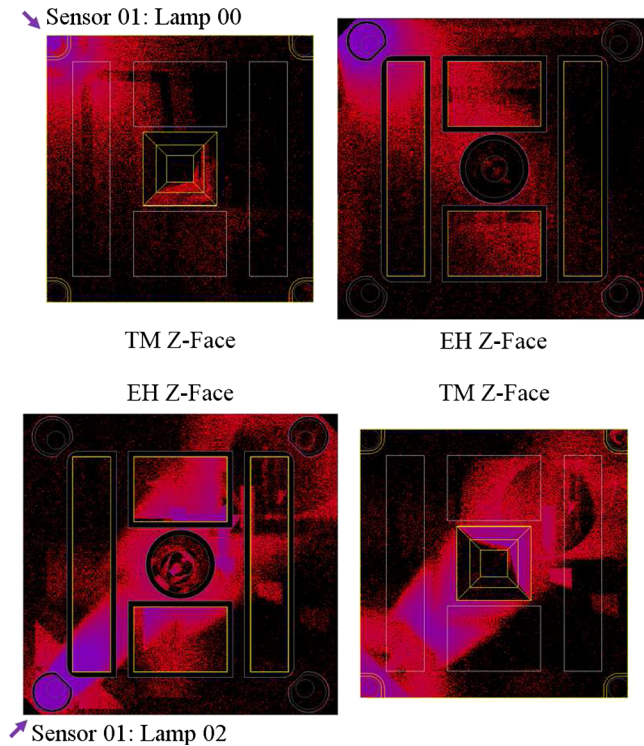
FIG. 5.   The distribution of absorbed UV light on the lower *z*-faces of the flight sensor. Top: Illumination provided by lamp 00 which is aimed at the test mass. The primary test mass surfaces are shown on the left while light reflected onto the housing surfaces can be seen on the right. Note that the primary light falls directly on the test mass corner with little reflected light reaching the electrode regions. Bottom: Illumination provided by lamp 02 which is aimed at the housing. The primary housing surfaces are shown on the left while light reflected onto the test mass surfaces can be seen on the right. Note that the primary light is mainly absorbed within the recess of the caging finger but some light also reaches the rest of the *z*-faces.

surrounding guard ring surfaces, most of which is within the recess that the caging finger closest to illumination is retracted into. A further 11–12% is also absorbed by the iridium caging finger itself. As will be discussed in the next section, this absorbed UV light is effectively wasted when it comes to discharging. Such suboptimal illumination is due to the original design of the caging mechanism (which did not include the corner caging fingers) being reviewed after the ISUK entry points had been finalized. This has also had consequences for the test mass illumination where the corner domes absorb a significant amount of light (around 43% and 23%). These general findings can be seen visually within the color maps.

B. Photoelectron flow model

Written in MATLAB, the photoelectron flow model uses the output of the ray trace to estimate the apparent yield of the system, given the photoelectric properties of individual surfaces as well as any applied AC or DC electric fields.

The model simplifies the problem by splitting the sensor up into pairs of adjacent surfaces on the test mass and electrode housing and treating them like parallel plates with a uniform electric field between them. The model calculates the net photoelectron flow between each surface pair while taking into account any instantaneous electrical potential difference as well as the energy distribution of the emitted photoelectrons. The system is then stepped in time, applied AC voltages are recalculated, and the test mass potential is updated due to any charge transfer in the previous step. The net flow is then evaluated again with this process repeated for a user-defined number of steps.

Though photoelectrons are emitted over a range of angles (described by a cosine distribution), the model only considers the component of an electron's kinetic energy that is normal to the emitting surface. As shown in Ref. [22], this is well described by a triangular distribution with an E_{max} equal to the difference between the UV photon energy (4.89 eV) and the surface work function, with the distribution peaking at around $E_{max}/5$. As the model approximates the surfaces as a series of parallel plates it is this energy component that will determine if a photoelectron can overcome any opposing electric potential. Such a triangular distribution has also been confirmed by unpublished measurements made at Imperial College London.

The model includes the 100 kHz, 4.88 V voltage applied to the six injection electrodes as well as the translational and rotational actuation voltages applied to the 12 sensing electrodes at frequencies of 60, 90, 120, 180, 240, and 270 Hz [16]. In flight the amplitude of these voltages varied depending on the axis and particular actuation mode but were typically around 2 V during the charge measurements that will be presented. All actuation biases were applied to pairs of electrodes with opposite polarity so as not to bias the test mass whereas the injection bias was applied with the same polarity to all electrodes causing the test mass potential to oscillate by ±0.6 V at 100 kHz. Given typical photoelectron energies of order 0.1 eV (equivalent velocity $\sim 2 \times 10^5$ ms^{-1}) and the largest gap being 4 mm, crossing times should be of order 20 ns. The highest-frequency AC voltage was 100 kHz, with a period that is therefore equivalent to 500 crossing times such that the AC fields appear quasistatic to the photoelectrons. Some charge measurements included additional DC voltages applied to particular sensing electrodes to intentionally shift the test mass potential which the model also includes.

As shown by the ray-trace results, a significant fraction of light is absorbed within housing gaps and recesses, particularly for lamps aimed at the electrode housing. Any photoelectrons generated from these surfaces have no adjacent region on the test mass and are instead opposite either another housing surface or the iridium caging finger. For two reasons it is very unlikely that these photoelectrons will reach the test mass. First, they are emitted randomly with a cosine distribution meaning that at most only half

will be directed toward the test mass. From the photo-electron's perspective, the deeper within a gap it originates the smaller the solid angle the test mass subtends, making it even less likely. The second and more significant reason results from work function differences between the iridium caging fingers and the surrounding gold-coated recess and to a lesser extent the small differences in work function across the gold surfaces within the gaps (so-called patch potentials [37]). These differences lead to contact potentials (also known as Volta potentials) of possibly a few tenths of a volt between gold surfaces and possibly up to a volt between the gold and iridium [38]. Given photoelectron energies of less than 1 eV from gold, electric fields within the gaps resulting from these contact potentials would have a strong influence on behavior. The upshot of these effects is that it is very unlikely that light absorbed within housing gaps can produce photoelectrons that can reach the test mass. Therefore the model simply ignores any UV light absorbed within the gaps as these cannot affect the net flow of charge to and from the test mass.

As mentioned, the model simplifies the electric field geometries within the sensor by only considering parallel field lines between adjacent surfaces. While uniform fields are an excellent approximation at the center of the electrode and test mass faces, toward the edges the field geometries become increasingly complicated. This is particularly true around the caging features and where the ray trace has already shown the majority of the light is absorbed. Fortunately, it is a photoelectron's final destination that is important rather than its exact trajectory and given their relatively low energy compared to the typical field strengths this tends to be inevitable. Treating the fields in this way significantly reduces the simulation run time, which is vital when fitting in a large parameter space. Similarly, a small amount of light (less than a few of percent) is absorbed by surfaces at the edges and corner of the housing that have no directly opposite region on the test mass. In these cases the first electric field the emitted photoelectrons encounter will be that between the test mass and grounded housing. Again, given their relatively low energies most of the time the photoelectron's destination will be determined simply by the direction of the field. As such, UV absorption in these regions is treated as if it were simply facing the test mass.

IV. PREFLIGHT MEASUREMENTS ON GROUND

Measurements of the flight sensor's discharge properties were carried out on ground between September and October 2014 by OHB Italia SpA (formally Compagnia Generale per lo Spazio) in Italy. The sensors were mounted in their flight vacuum enclosures but had yet to be integrated onto the spacecraft. To allow measurements to be made the central plungers were engaged and used to hold the test masses but the eight caging fingers were retracted so as to affect illumination as little as possible.

An electrometer (Keithley, model 6517B) was connected to the otherwise isolated plungers and used to directly measure the drain current from the test mass during illumination, unlike in flight where changes in the test mass potential with UV illumination were the only measurement possible.

To perform the measurements each of the two sensor vacuum enclosures were in turn mounted in a test-rig thermal-vacuum chamber. Illumination was provided by a single mercury lamp and an optics barrel identical to that used in the ULU, but two off-the-shelf fibers and a vacuum feedthrough were used to route the light to each ISUK. A mechanical chopper was placed in the optical chain to modulate the light signal and allow for a lock-in measurement that improved sensitivity. Unlike in flight, the light cone emitted from each ISUK was nominally identical and a series of calibration measurements with a UV power meter were used to determine the throughput of the optical chain. During the test all electrodes were electrically shorted to the housing with the isolated test mass in electrical contact only with the plungers. For each ISUK illumination a measurement was made with the test mass biased at both -5 and $+5$ V with respect to the housing. With photoelectron energies of around 1 eV, these biases are sufficient to completely turn off one of the two opposing photocurrents, allowing extraction of the saturation levels illustrated in Fig. 3. These measurements were repeated before and after a 24 hour, 115°C bake-out under vacuum (with a further 24 hours either side for heating and cooling of the system). The results are shown in Table II.

After first distinguishing between before and after bake-out, the measurements can be gathered into four groups: the saturated apparent yield of the test mass surfaces with direct illumination (one per sensor), the saturated apparent yield of the test mass surfaces with indirect illumination (two per sensor), the saturated apparent yield of the electrode housing surfaces with direct illumination (two per sensor), and the saturated apparent yield of the electrode housing surfaces with indirect illumination (one per sensor). Within each of these groups the distribution of absorbed UV light was nominally identical, and therefore any variation within a particular group should only be due to differences in the photoelectric properties of each illuminated area. Comparing these groups before and after the bake-out shows that there was no clear systematic increase in the apparent yields with some rising and some falling. However, what did occur was a reduced variation in the apparent yields within each group; not only between different corners of a single sensor but even between the two separate sensors. This is confirmed if the coefficient of variation (CoV) is considered for each group (the ratio of standard deviation and mean value, provided here as a percentage), with three out of the four groups showing a significant reduction after bake-out. The one exception is for the electrode housing when directly

TABLE II. The flight sensor apparent yields measured before and after bake-out, with statistical uncertainties. Each of the six lamps were directed at either the test mass (TM) or electrode housing (EH) but inevitably generated a photocurrent in both directions. By biasing the test mass by either ± 5 V the aim was to completely suppress one of these, saturating the net photocurrent in the opposite direction. A positive apparent yield represents a saturated photocurrent from the test mass while a negative apparent yield represents a saturated photocurrent from the housing. Note that the apparent yields for direct test mass illumination are an order of magnitude higher than the other measurements. The coefficient of variation (σ/μ) was also calculated (ignoring statistical uncertainties) for appropriate groups of measurement to quantify how the surface properties varied for a given illumination.

	Apparent yield (net electrons/injected photon)			
	Before bake-out		After bake-out	
	−5 V TM saturation	+5 V EH saturation	−5 V TM saturation	+5 V EH saturation
Sensor 01: Lamp 00 (TM)	$(+1.67 \pm 0.03) \times 10^{-5}$	$(-1.03 \pm 0.05) \times 10^{-6}$	$(+2.49 \pm 0.08) \times 10^{-5}$	$(-1.91 \pm 0.06) \times 10^{-6}$
Sensor 02: Lamp 01 (TM)	$(+4.65 \pm 0.54) \times 10^{-5}$	$(-1.35 \pm 0.32) \times 10^{-6}$	$(+2.18 \pm 0.03) \times 10^{-5}$	$(-1.99 \pm 0.05) \times 10^{-6}$
Sensor 01: Lamp 02 (EH)	$(+2.52 \pm 0.03) \times 10^{-6}$	$(-3.34 \pm 0.12) \times 10^{-6}$	$(+2.47 \pm 0.03) \times 10^{-6}$	$(-4.19 \pm 0.03) \times 10^{-6}$
Sensor 01: Lamp 04 (EH)	$(+2.35 \pm 0.05) \times 10^{-6}$	$(-1.87 \pm 0.03) \times 10^{-6}$	$(+2.84 \pm 0.05) \times 10^{-6}$	$(-3.67 \pm 0.05) \times 10^{-6}$
Sensor 02: Lamp 05 (EH)	$(+6.06 \pm 0.96) \times 10^{-6}$	$(-2.48 \pm 0.54) \times 10^{-6}$	$(+2.16 \pm 0.03) \times 10^{-6}$	$(-2.33 \pm 0.04) \times 10^{-6}$
Sensor 02: Lamp 06 (EH)	$(+7.90 \pm 1.24) \times 10^{-6}$	$(-4.28 \pm 0.75) \times 10^{-6}$	$(+2.82 \pm 0.11) \times 10^{-6}$	$(-5.12 \pm 0.20) \times 10^{-6}$
CoV (TM)	66.7%	19.0%	9.4%	2.7%
CoV (EH)	58.0%	35.1%	12.6%	30.4%

illuminated, which showed a comparable level of variation before (35.1%) and after bake-out (30.4%). The general increase in uniformity was presumably due to a change in the surface state: either the heating led to a more uniform redistribution of surface contaminants or it removed weakly bound adsorbates with the remaining tightly bound layers common to all surfaces. Finally, considering just the results after bake-out, the apparent yield for the test mass surfaces with direct illumination were around a factor of 6 greater than when the electrode housing surfaces were directly illuminated. As we shall see in Sec. V C, to explain these differences in apparent yield the ray trace is required to disentangle the affect of the total absorbed light from the intrinsic quantum yields.

It should be noted that there were several systematic differences between the measurements made on ground and the situation in flight. First, the optical path between the lamp and the ISUKs differed from that in flight, producing a light emission cone with a half angle of 12.3°, common to all of the ISUKs. This change meant slightly different regions within the sensor were inevitably illuminated, with potentially different photoelectric properties. The second difference was that during flight the central plunger was retracted and was electrically connected to the housing rather than the test mass. This had three consequences; it changed the light distribution within the sensor slightly, altered the electric field distribution in its vicinity, and any photoelectrons generated from it acted to charge the test mass negatively. The final (and arguably most significant) difference was the absence of the various time-varying sensing and actuation voltages present within the flight sensors. Experience from the Trento tests on GRS proto-type hardware showed that they have a significant influence on the overall discharging behavior [25]. By using the

models described in Sec. III, which can account for these differences, the ground and flight tests can however be analyzed together and their intrinsic surface quantum yield values compared. The results of this comparison will be presented in Sec. V C.

V. IN-FLIGHT MEASUREMENTS

During ground testing the entire system was pumped down to $<10^{-5}$ mbar and the design of the flight caging and venting mechanism allowed each vacuum enclosure to be closed while the outer chamber of the test rig was still under vacuum. The sensors were then left sealed (but with no active pumping) throughout spacecraft integration with the internal pressure estimated to have drifted up to around a millibar during this time. Following a successful launch on December 3, 2015, Pathfinder reached its operational orbit around the L1 Lagrange point in February 2016. The vacuum chambers housing the sensors were then vented to space on February 3, over 15 months after the discharge measurements were made on ground.

A series of initial calibration measurements showed that despite the systematic differences compared to the ground test, the discharge properties were qualitatively similar to those measured preflight. As before, lamps 00 and 01, which were aimed at the test masses, had significantly higher apparent yields than the lamps aimed at the electrode housing. As will be described in Sec. V C, dedicated measurements made later in the mission quantified these differences and allowed comparisons to be made with simulation. During these initial measurements it was also discovered that during spacecraft integration the fiber chains for lamps 01 and 05 had been inadvertently switched. Thus, lamp 05 was unintentionally attenuated

by a factor of ~63, while lamp 01 had no attenuation. Nevertheless, both were still capable of discharging although lamp 05 at a significantly reduced rate.

Following calibration, the CMS was commissioned for the start of the mission proper on the March 1, 2016 and operated successfully until the mission's end on the July 17, 2017. The six lamps were turned on a combined total of 418 times in flight (without fail) and were operated for a total of 421 hours. No significant signs of aging, such as output instability, were observed for any of the lamps. However, by the end of the mission usage levels were only just approaching those where such effects were predicted to start.

After initial release from the caging mechanism, residual test mass charges equivalent to -421 and -281 mV were measured on the test masses. The level of residual charge is related to the difference in work function of the last two surfaces in electrical contact as well as triboelectric processes. For reasons that remain unclear, subsequent grabs and releases throughout the rest of the mission produced residual test mass potentials systematically lower and ranged from -98 to -52 mV. During typical science mode operations the charging rate for both test masses was similar and drifted upward from around $+21 \, \mathrm{es}^{-1}$ at the start of the mission to around $+36 \, \mathrm{es}^{-1}$ by the end of the mission. This increase was predominately due to the declining solar cycle resulting in a gradually increasing GCR flux. Such rates were in good general agreement with preflight predictions [7] but dedicated experiments revealed additional charging rate dependencies on actuation scheme and test mass potential, with rates as high as $+57 \, \mathrm{es}^{-1}$ measured at an extreme test mass potential of -1 V. Understanding these dependencies will form the basis of a future paper, but suffice it to say that charging typically caused the test mass potential to increase by around $+12$ mV per day.

For the vast majority of the mission a fast discharge scheme was employed. During a weekly spacecraft station keeping period the CMS was used to discharge the test masses to a set level, e.g., -80 mV, taking around 30 minutes. Incident cosmic rays then gradually charged the test masses positively until the next station keeping, at which point the CMS reset the test masses' charge to the desired level. At the start of the mission this was performed weekly but was later reduced to every two or three weeks, allowing longer uninterrupted science runs. Initially these fast discharges were done open-loop by bringing the test mass to an equilibrium potential that had been shifted by applying appropriate DC voltages during the discharge. While this method was both simple and reliable it was also fairly inefficient when it came to lamp usage. Once confidence in the system's performance had increased these routine fast discharges were instead handled by the automated closed-loop CMS software. Both methods typically got within the equivalent of a few millivolts of the desired

test mass charge level, similar in scale to the systematic uncertainty in the measurement due to stray DC voltages on the sensing electrodes.

Two experiments were also performed to test the continuous discharge concept, successfully holding both test mass potentials within 10 mV of zero for several days. A detailed analysis of these tests will be presented in a future article.

A. Long-term stability

As explained previously, the apparent yield for a particular lamp is dependent on the state of the system at the time of the measurement. To make a clean assessment of the stability of the photoelectric properties over the course of the mission, a subset of measurements need to be chosen that were made with the system in a reproducible state. The automated discharges carried out during the regular station keeping periods offer an ideal data set for this. Functionally, the automated CMS software performed a discharge in the same way each time. After first measuring the test mass charge by applying a dither voltage, it then calculated the appropriate lamp setting and duration to reach the requested level of final test mass charge. Following illumination, it would remeasure the test mass charge and determine if it was within a predefined range of success. If the system had over-/undershot, the lamp settings would be recalculated and the sensor illuminated for a second time. This would continue until the desired discharge level had been reached or the time allocated for discharging had been exhausted. To aid the discharge, the CMS applied relatively large DC voltages to the sensing electrodes (± 5 V on x and y, ± 1 V on z) that shifted the test mass potential by approximately ± 1.27 V.

By June 2016 the automated CMS was in routine use, typically discharging from a few positive tens of millivolts to either -80 or -120 mV. Given the systematic nature of the discharges the CMS would typically use a particular lamp aimed at the electrode housing—lamp 02 for sensor 01 and lamp 06 for sensor 02—and also a similar lamp output setting each time. Combined with on-ground calibration measurements, the telemetered photodiode readings within the lamp blocks allow the number of photons entering the sensor during a discharge to be estimated. This can then be used with the change in test mass charge measured before and after illumination to calculate the apparent yield. In addition, a couple of early in-flight test measurements were made during commissioning at the end of February 2016 which allow the baseline of long-term stability to be extended. While these measurements were similar to the automated CMS discharges, as a precaution the lamps were commanded manually and only on for a relatively short amount of time at a low setting. They were also made when the absolute test mass charges were lower than the other measurements considered at levels equivalent to -275
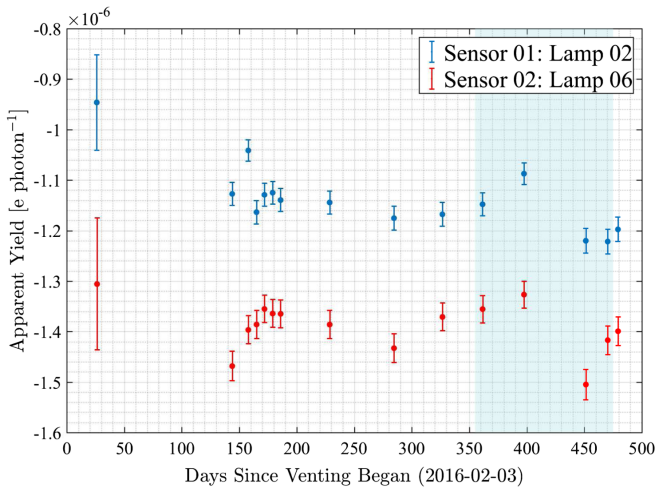
FIG. 6. The apparent yields observed during CMS fast discharges for lamp 02 coupled to sensor 01 and lamp 06 coupled to sensor 02. The first two measurements were made during commissioning 26 days after venting began. They contain a greater level of uncertainty but allow the time baseline to be extended significantly. The shaded blue regions show a four month period where the sensor temperatures were lowered to $12 \pm 1°C$ due to an unrelated experiment. For the majority of the mission the sensors were at $23 \pm 1°C$.

and -110 mV, respectively. These two measurements therefore have a larger statistical error but also the possibility of a systematic error due to the lamp setting used and the test mass potential at the time of illumination. The apparent yields obtained during CMS fast discharges for lamps 2 and 6 are shown in Fig. 6.

The apparent yield for both lamps were stable at about the 10% level approximately 140 days after the system was vented to space and remained so for the duration of the mission. In both cases the apparent yields were consistent with the measurements made during commissioning, although the aforementioned caveats should be kept in mind. The small drift observed in the measurements was correlated with the test mass potential at the start of the discharge which was not unexpected. Given the systematic nature of the discharges and the fact that the test masses tended to charge up in tandem, this also explains the ostensible correlation between the apparent yields for both lamps. With the DC biases applied during the fast discharges ($+5$ V on x and y, $+1$ V on z), the measured apparent yields should have been approaching the negative saturation levels for each lamp. As this level is dominated by the quantum yield of the electrode housing, we can say that it was stable with time under vacuum, cosmic-ray irradiation, and a temperature reduction of $11°C$ as well as with the UV irradiation from the discharging itself. It also shows that the optical chains between each lamp and the sensors did not degrade as this would have been observable in these measurements. While these measurements alone do not necessarily prove that the photoelectric properties

were unchanged compared to those made on ground, they do demonstrate that they reached a stable level once in space.

B. Discharge curves

Although the CMS was able to successfully discharge both test masses from day one of the mission, it is vital that it is understood why the system behaved the way it did such that its success can be replicated for LISA. This was particularly important given the variability in photoelectric properties that had been measured during Pathfinder development and the fact that the sensor's geometrically complicated corner regions absorbed significant amounts of light. Several dedicated experiments were performed to characterize each lamp's full discharge curve, such as that shown in Fig. 3, which shows the change in apparent yield with test mass potential. Again, a lamp's discharge curve is dependent on the area of the sensor illuminated, the quantum yield and work function of surfaces where UV light is absorbed, and the voltages present in regions where photoemission occurs. By measuring each lamp's full discharge curve a deeper understanding of the overall system is gained: it aids modeling and allows comparison with measurements made on ground.

The measurements described here were made on the February 1 and 2, 2017 and the June 23 and 24, 2017. They consist of spot measurements where the test mass charge was measured before and after relatively short illuminations as well as continuous measurements made during longer illuminations. For a particular test mass, lamps were used in sequence, with one used to charge the test mass positively while the next would charge negatively. To allow higher test mass potentials to be reached, during some illuminations DC biases of ± 4.8 V were applied to all x and y sensing electrodes that shifted the test mass potential by 1.18 V. Strictly speaking these measurements describe a separate discharge curve to those made with no applied DC biases as they change the local electric fields around the electrodes. However, they are presented together here as shifts in the curves relate to the fraction of photoelectrons emitted in the x and y sensing electrode regions and allow a more stringent test of our simulations. All of the measurements presented here were made with the AC actuation voltages as low as possible, with amplitudes typically less than 2 V. Note that the long-term measurements discussed in the previous section were made with significantly larger actuation voltages present, leading to an expected difference in the observed apparent yields. Finally, to limit systematic differences related to a lamp's photodiode calibration all measurements made with a particular lamp were made at the same output setting. The measured discharge curves for all six lamps are shown in Fig. 7.

What is immediately striking is the similarity between the curves for lamps 00 and 01, which were aimed at test mass 01 and 02, respectively. The positive parts of the
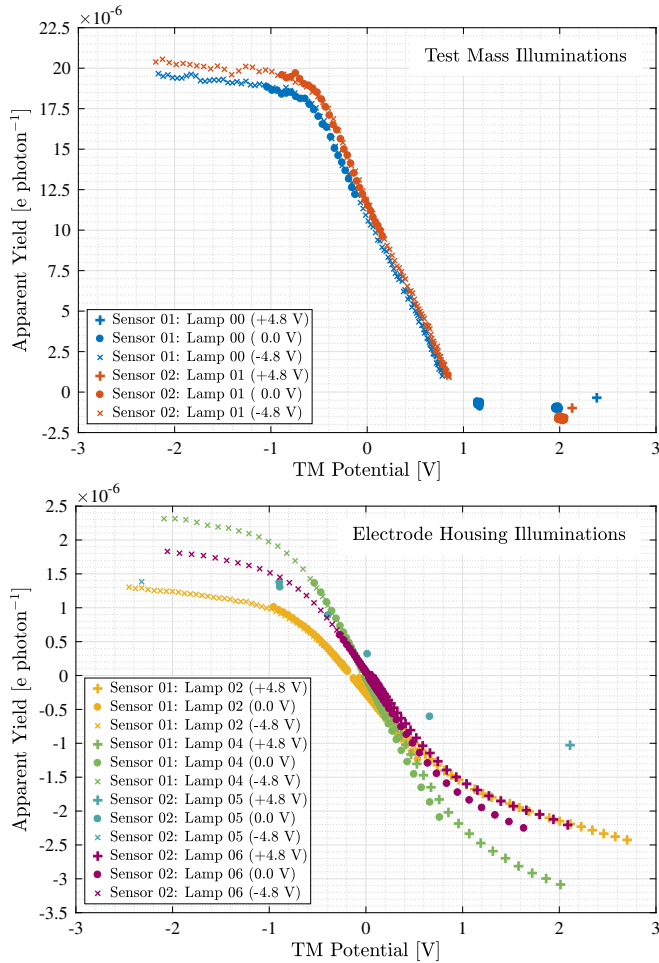
FIG. 7. Top: Measured discharge curves for lamps 00 and 01, which primarily illuminated the test masses in sensors 01 and 02. Bottom: Measured discharge curves for the four lamps which were aimed at the electrode housing of either sensor 01 or 02. Filled circles indicate measurements made with no applied DC biases, + symbols where +4.8 V was applied to all x and y sensing electrodes, and × symbols where −4.8 V was applied instead. Combined with its relatively low yield, the inadvertent attenuation of lamp 05's output made only spot measurements practical due to time constraints. Note that while the random error on individual points is small (error bars removed to improve clarity), there is a systematic uncertainty specific to each curve at the 10–20% level. This would scale an entire curve and is due to systematic uncertainties in the calibration of the optical chains.

curves dominated by emission from test mass surfaces differ by only a few percent, although there are subtle but significant differences if the negative apparent yield from housing surfaces is considered. Nonetheless, such quantitative similarities are slightly surprising given that the optical chain for lamp 00 featured heavy attenuation by design, whereas the chain for lamp 01 inadvertently had none. On its own this should not affect the discharge curves as the net change in the test mass charge should scale with the number of injected photons. However, as mentioned

previously the length of the attenuating spacer in each lamp's optical chain changed the angular distribution of light injected into the sensor: the greater the level of attenuation, the narrower the beam emitted. Indeed, this is seen in the UV ray-trace results (detailed in Sec. III A and tabulated in the Appendix) where the total absorbed by each test mass is similar (66.07% and 64.50%, respectively), but the amount absorbed by the gold-platinum corner domes is significantly different (42.58% and 23.17%, respectively) due to the difference in ISUK light distributions. Despite such a significant difference in where the light was absorbed, the similarity in the apparent yields suggests that the uncoated gold-platinum corner domes had a similar quantum yield to that of the gold. For both lamps the equilibrium potential was just below +1 V. The negative saturation yield for lamp 01 was about twice that of lamp 00 and spot measurements for both showed the level was roughly halved when +4.8 V was applied compared to those with no DC bias on the x and y sensing electrodes. The measurements also show that the positive yield had not completely saturated by −2.2 V and still showed a small negative slope.

The curves for the four lamps aimed at the electrode housing are qualitatively similar to each other but significantly different from those aimed at the test mass. For the four housing lamps, the amplitudes of the negative saturation levels were comparable to those of the positive saturation levels. Partly in consequence of this, the equilibrium potential of lamps 04 and 06 were close to zero and only lamp 02 had a significantly negative equilibrium potential (−150 mV) with no applied DC bias, making it the only lamp suitable for the automated CMS continuous discharge. The heavily attenuated lamp 05 had a light output distribution significantly different to the other three lamps and the most positive equilibrium potential (+250 mV). All four lamps had a fairly wide range in their saturation levels, although with the exception of lamp 05's positive level, none appeared to have fully saturated by ±2.5 V. Lamps 02, 04, and 06 showed clear shifts in negative apparent yields when +4.8 V was applied compared to those with no DC bias on the x and y sensing electrodes.

Comparing the discharge curves to the idealized cases shown in Fig. 3, both lamps primarily illuminating the test mass produced curves in qualitative agreement with expectation. The positive saturation level had an apparent yield over 10× higher than that at the negative level leading to a positive apparent yield at a test mass potential of 0 V. However, none of the four lamps aimed at the electrode housing matched the ideal case. While the positive saturation levels were comparable in amplitude to the negative levels observed with test mass illumination, the negative levels were around an order of magnitude lower than the equivalent obtained with test mass illumination. While this

behavior was not unexpected based on the ground tests it was nonoptimal and to fully understand it the simulated system described previously needs to be invoked.

C. Estimated photoelectric properties

To assess the measured discharge curves we first take the output of the ray trace (detailed in Sec. III A and tabulated in the Appendix) as the input to the photoelectron flow model (described in Sec. III B). The work function and quantum yield parameters of both the test mass and electrode housing surfaces are then varied independently for each curve while checking for consistency with the data. Although the model allows the photoelectric properties to be specified for individual surfaces (each electrode, test mass face, corner dome, etc.) we only considered average properties for either the test mass or housing to avoid arbitrary tuning. Despite this simplification, the model is able to capture many of the features observed in the measured curves with the results shown in Fig. 8.

For each discharge curve there were four free parameters: the average test mass quantum yield, the average housing quantum yield, the average test mass work function, and the average housing work function. It was found that the average quantum yields were heavily constrained by the data with the values for the test mass and housing also being independent of each other. Essentially, the quantum yields have to be compatible with the saturation levels at either extreme which typically leads to uncertainties of less than 10%. Prior measurements on ground had always found the work function of gold (previously exposed to air) to lie between 3.5–4.7 eV [25], and therefore only values in this range were considered. The interplay between the two opposing work functions (that define the photoelectron energy distributions) determine the shape and position of the transition between the two saturation levels, and it was found that there was some degeneracy in pairs of work functions that gave good results. Nevertheless, our primary concern was the quantum yields and the exact work function values had little influence on these.

The simulated discharge curves generally capture the shape of the measurements very well, particularly at positive yields for lamps aimed at the housing (02, 04, 05, 06). However, the model struggles to accurately reproduce the shifts in the negative yields for lamps 02 and 06 when +4.8 V was applied. One possible explanation is that in these cases the yield from the electrode regions were higher than the average value, causing a more prominent shift when the DC bias was applied. Tests have shown that the model is capable of reproducing the observed behavior by increasing the yield in the electrode regions and proportionally decreasing it elsewhere, thus keeping the same average value. However, there is no way of distinguishing between a single electrode having a much higher yield or several being slightly higher. As such we have stuck to a single average value, though it produces a poorer fit in these regions. It should be emphasized here that the smaller shift observed for lamp 04 and larger shifts for lamps 00 and 01 are well captured by a single average quantum yield. This observation increases confidence that the amount of light predicted to be absorbed in these regions by the ray trace is correct as it needs to be compatible with the average quantum yield. The model also has trouble reproducing the gradually increasing positive saturation levels for the test mass facing lamps, particularly lamp 00. While there is a small gradient in the simulated saturation level it is not as prominent as that observed. Again, tests suggest that by varying the yield in different regions but keeping the average similar the model can reproduce the behavior but there is no way of distinguishing between several arbitrary ways of tuning.

We now have estimates for the average quantum yields of each surface measured in flight. To allow comparison, the apparent yields measured on ground also need to be converted to average quantum yields. For the ground measurements only the illumination ratios obtained from the ray trace are required as only the discharge curve saturation levels were measured and different bias voltages were not applied to individual electrodes. However, as explained in Sec. IV the measurements made on ground had several subtle differences compared to those made in flight. As such the ray trace was rerun with appropriate ISUK light distributions (unlike flight, these were identical for the six lamps) and with the central caging plungers engaged. For lamps 00 and 01 which primarily illuminated the test mass, 66.5% of the total light injected was absorbed by test mass surfaces while 14.8% was absorbed by electrode housing surfaces. For the four lamps aimed at the housing, 12.9% of the total light injected was absorbed by test mass surfaces while 18.0% was absorbed by electrode housing surfaces. As discussed in Sec. III, the remainder of the light in both cases was absorbed by surfaces unlikely to influence discharging and was therefore ignored. Given the fraction of total light absorbed by each surface type, the apparent yields measured on ground (Table II) can be converted to average quantum yields. A comparison between the average quantum yields of each surface measured on ground and in flight is shown in Fig. 9.

The estimated quantum yields range from about 1 to 4×10^{-5} e per absorbed photon, consistent with typical historical measurements made on individual gold surfaces during hardware development [25]. The estimates suggest that for each lamp any changes in the quantum yield between ground and flight were relatively small, differing by a few tens of percent. For the test mass quantum yields the average change between ground and flight was 22%, while for the electrode housing surfaces it was 31%. In addition to the random uncertainty on the in-flight quantum yield estimates there are significant systematic uncertainties, approximately at the 10–20% level, related to the on-ground calibration of the UV power emitted.
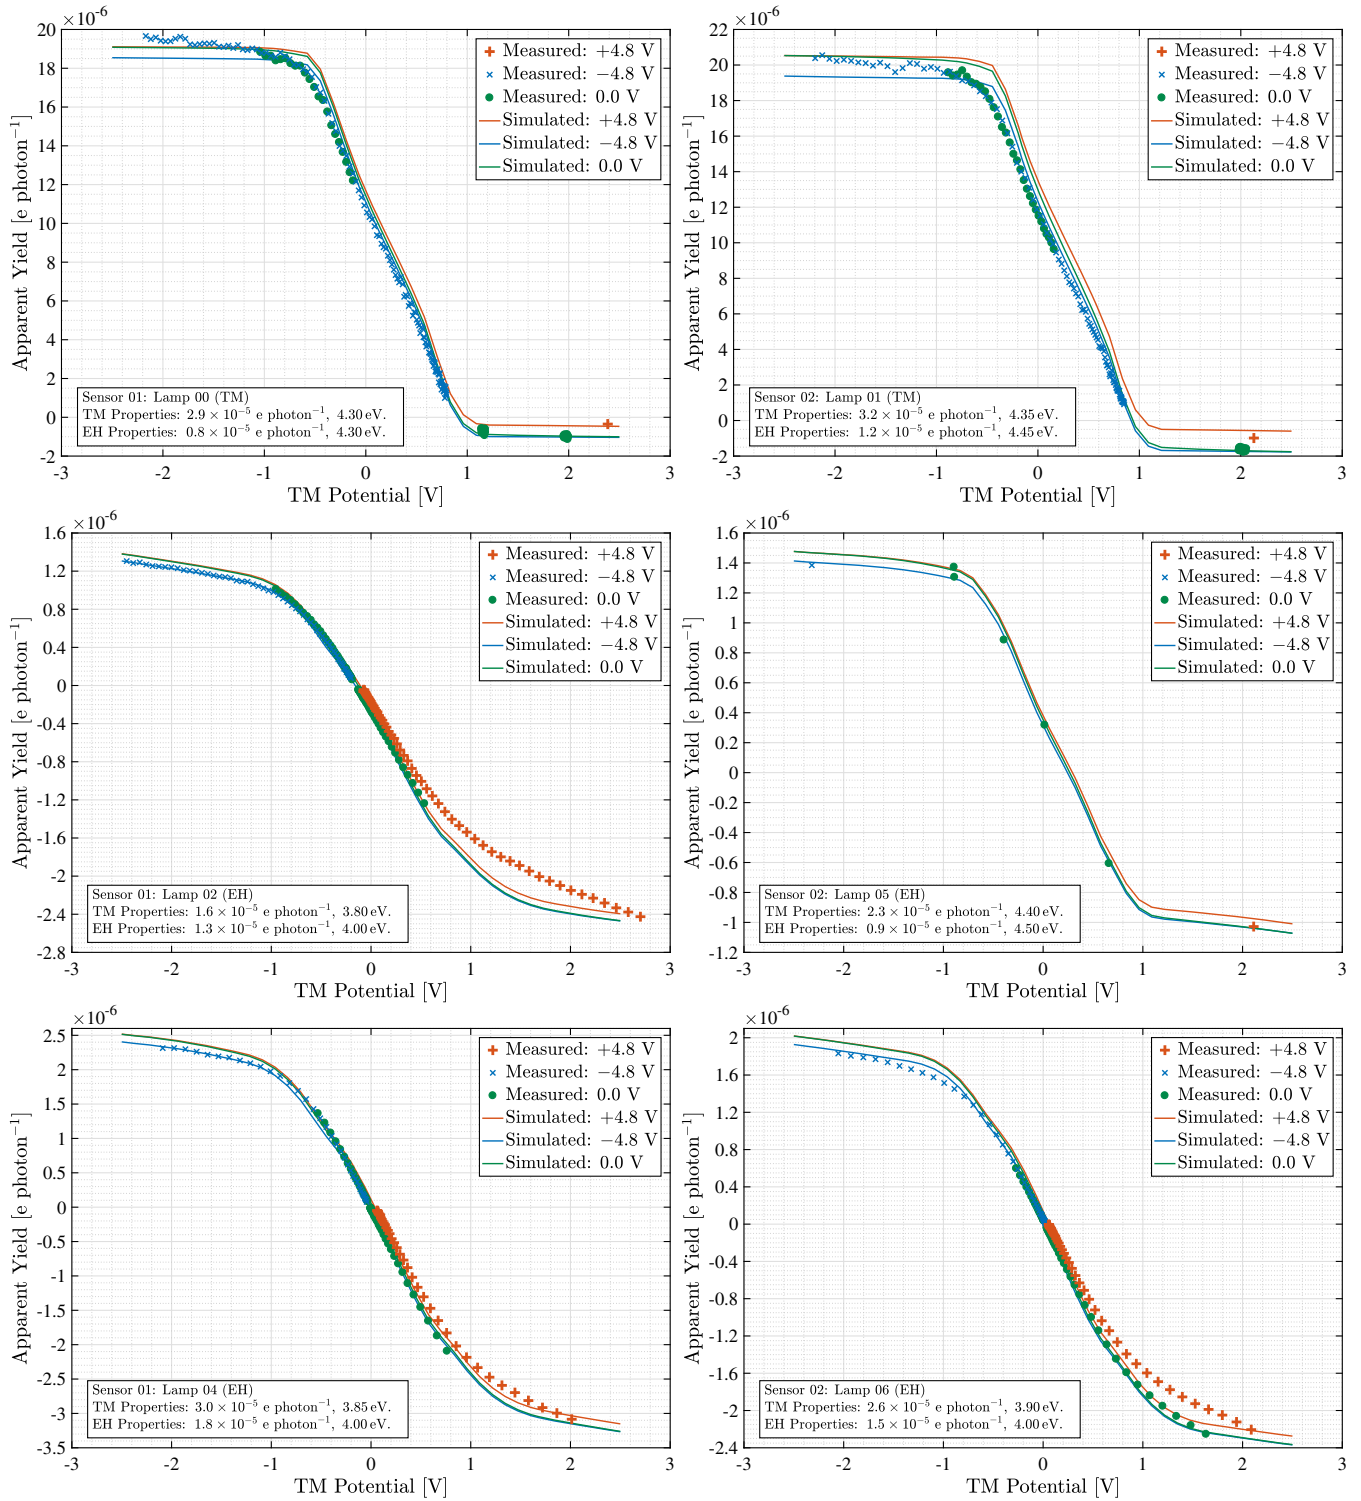
FIG. 8.   The six discharge curves measured in flight, described in Sec. V B. Those on the left are for sensor 01 while those on the right are for sensor 02. In each case simulated discharge curves have been overlaid. For a single set of photoelectric parameters three discharge curves were simulated for each lamp: one with no DC biases applied, one with +4.8 V applied to all x and y sensing electrodes, and one with −4.8 V applied to all x and y sensing electrodes.

For a particular lamp this systematic uncertainty in the number of photons reaching the sensor would act to scale the test mass and electrode housing quantum yields in the

same way, either increasing or decreasing them together. Looking at Fig. 9, this alone cannot explain the changes as several pairs of lamp measurements did not change in the
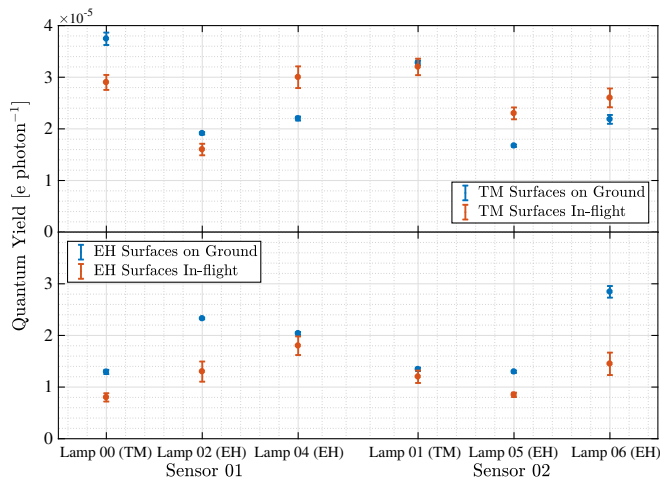
FIG. 9.   The average quantum yields from the test mass surfaces (TM, top panel) and the electrode housing surfaces (EH, bottom panel) for all six lamp illuminations. For each lamp the primary surface under illumination is specified in brackets. Quantum yield estimates are shown based on measurements made after bake-out on ground (Table II) as well as those made in flight (Fig. 8). Note that there are additional systematic uncertainties at the 10–20% level on the in-flight estimates due to the UV photodiode calibration.

same direction or by the same amount. Another aspect to note is that due to a slightly different illumination and a change in central plunger position, the exact same areas were not illuminated on ground and in flight. Therefore, even if the system had remained completely stable the measured quantum yields would not be expected to be identical as slightly different areas were being probed in each case.

While it was previously shown that the photoelectric properties were stable once in space, they do not appear to have become more uniform as the level of scatter on ground and in flight is comparable. The average level of the quantum yields for each surface type was also little changed. The average quantum yield (and standard deviation) for the test mass surfaces went from $(2.5 \pm 0.3) \times 10^{-5}$ e per absorbed photon on ground to $(2.6 \pm 0.3) \times 10^{-5}$ e per absorbed photon in flight, while the electrode housing surfaces went from $(1.9 \pm 0.3) \times 10^{-5}$ e per absorbed photon to $(1.2 \pm 0.2) \times 10^{-5}$ e per absorbed photon. The average values also highlight that the quantum yield from the test mass surfaces appeared to be systematically higher than those from the housing, both on ground and in flight. The possibility that nominally identical gold-coated surfaces should differ in this way will be discussed in the context of preparing for the LISA mission.

VI. TOWARD LISA

Despite its overall success, measurements made during development, testing, and in flight for the Pathfinder

discharge system have identified three key issues that face the LISA system. The first relates to the significant variability in the photoelectric properties of air-exposed gold surfaces, when illuminated with 253.7 nm UV light. Independent of any modeling, the apparent yields measured on ground with the flight sensor had coefficients of variation up to 30% for groups of measurements with nominally identical illuminations. Modeling the system also allowed the intrinsic quantum yields to be extracted and found that for nominally identical gold surfaces, stored and handled in controlled conditions and undergoing a 24 hour bake-out at 115℃, the quantum yield varied from 1 to 4×10^{-5} e per absorbed photon. Similar levels of variability have been observed previously and have been interpreted as an extreme sensitivity to the state of the surface [22]. Such variability makes it impossible to accurately predict the final performance of the discharge system and while the ability to apply DC biases makes it very unlikely that fast discharge would be prevented, continuous discharging (relying on bipolar transfer with no DC bias) could be impractical, as was the case on Pathfinder with three of the four electrode housing lamps. While an improved understanding of how surface contamination affects the quantum yield would certainly be beneficial, in real terms there is limited scope for what could be done differently for LISA. The surfaces will inevitably be exposed to air during integration and strict storage and handling procedures were already in place for Pathfinder. One area of investigation that could prove fruitful is the relationship between bake-out temperature and duration and how they relate to the quantum yield variability.

Alongside the variability, the quantum yield of the test mass surfaces as extracted from our model appear to be systematically higher than that of the housing by around a factor of 2. While an overlooked systematic error in the modeling is always possible it is difficult to envisage where such a large discrepancy could originate. If instead the asymmetry has a physical origin, seeking an explanation before moving toward LISA will be important in order to consolidate confidence in the design, production, and integration of the sensor. Possible overlooked differences between the test mass and housing gold surfaces include the underlying substrate, the surface roughness, and the relative temperatures when cooling after bake-out which could lead to contaminates being redistributed preferentially on one of the surfaces.

The affect of both variability and asymmetry in the quantum yields was exacerbated on Pathfinder by a significant amount of light being lost within recesses when the housing surfaces were illuminated. Combined with contact potentials preventing transport of the low-energy photo-electrons, this effectively meant 70% of the light was wasted when the housing was illuminated. This left the ratio of useful absorbed light between the housing and test mass at just 2.1 for direct housing illumination in flight. This was compared to a ratio of useful absorbed light between the test mass

and housing of 4.6 for direct test mass illumination in flight. Broadly speaking, for test mass illumination the apparent yield from the test mass would still dominate even if the quantum yield was around a factor of 4 higher for the housing than for the test mass. For housing illumination, only around a factor of 2 difference can be accommodated, which was coincidentally similar to the systematic difference in the observed test mass and electrode housing quantum yields. Though it could of course bring unforeseen risks, there is the possibility that the illumination scheme for LISA could be changed. For example, initial tests suggest that had the lamp 02 ISUK been angled at 8° rather than 20°, the absorption percentages in flight would have been 47.7% for the electrode housing (as opposed to 20.5% for Pathfinder), 15.2% for the test mass (9.9% for Pathfinder), with just 37.0% wasted, increasing the ratio of useful light to 3.2. However, a detailed investigation would be required to assess the many illumination alternatives and whether such changes were feasible from an engineering perspective. More importantly, there is the possibility that any changes could have unintended consequences elsewhere and there is therefore a strong argument to keep things the same, particularly given the success of the Pathfinder system.

Due to mass, power consumption, and lifetime advantages, the LISA CMS will almost certainly be based on UVLEDs rather than mercury lamps [39–41]. While changing the light source will not automatically address the issues associated with variable and asymmetric quantum yields, it could offer ways of mitigating against them. The use of UVLEDs opens up the possibility of driving them in a pulsed mode while synchronizing their output to the AC voltages present within the sensor [30,39,42]. Conceptually, the sensor would only be illuminated while the AC voltages were at levels conducive to the desired direction of discharge. Such a scheme should offer an increased dynamic range and robustness against variable and asymmetric quantum yields but would come at the cost of increased complexity and potentially a new set of unknown challenges.

VII. CONCLUSIONS

The Pathfinder CMS operated successfully over the entire mission from March 2016 to July 2017. It had to counter average cosmic-ray charging rates from $+21$ to $+36$ es^{-1} as well as post-release residual charge that produced test mass potentials as high -421 mV. During the mission over 50 fast discharges, two continuous discharging runs and several dedicated discharging tests were performed. The six mercury lamps were turned on a combined total of 418 times without fail and were operated for over 421 hours. The routine fast discharges carried out during station keeping kept the test mass charge within 100 mV of zero during all primary test mass acceleration runs and continuous discharge tests also demonstrated that the test mass charge could be maintained within 10 mV of zero for days, and likely weeks.

Alongside the production of the physical hardware, a two-part software model was developed at Imperial College London to simulate the overall discharging behavior. By modeling the entire sensor geometry and reflective properties of individual surfaces, the ray trace determined how much UV light was absorbed and where. This information could then be fed into the photoelectron flow model which approximated the system as a collection of parallel plates, greatly simplifying the electric field geometries. Taking into account the many AC and DC voltages within the sensor, as well as the photoelectron energy distributions, the surface quantum yields and work functions were then estimated from the original apparent yield measurements. It was found that this relatively simple model was able to reproduce most of the observed behavior with shortfalls due to a desire to limit the number of free parameters rather than an intrinsic limitation in the simulation. While a more rigorous treatment of the electric fields would have increased the complexity and therefore runtime of the simulation, it is doubtful a full implementation would have improved results. Instead, it is a lack of knowledge for how the photoelectric properties vary over particular surfaces that ultimately limits the simulation. There is no practical way of directly measuring the photo-electric properties of individual flight surfaces *in situ* and due to their propensity to vary upon reexposure to air, measurements would be of little use prior to integration. However, the simulation could be improved with direct roughness and reflectivity measurements, particularly for the iridium and gold-platinum surfaces.

The ability to measure each lamp's apparent yield on ground with the flight sensors proved very useful and increased confidence that the discharge system would work successfully in space. There are several ways in which similar measurements for LISA could be improved. First, the optical chains (including any attenuators) need to be identical to those used in flight as they affect the distribution of UV light emitted from the ISUKs and therefore the area illuminated within the sensor. Second, a lamp's entire discharge curve should be measured rather than just the saturation levels, ideally with AC voltages applied to electrodes that are representative of those experienced in flight. Less extrapolation would then be needed to predict the final in-flight performance and also allow a better test of simulation. Finally, and if time permitted, different DC biases could be applied to individual electrodes to assess how different regions contribute to the total apparent yield, either through absorbing a lot of light or having a high quantum yield. Not only would this provide a more stringent test of simulated predictions but also allow fast discharge voltage schemes to be adjusted to give optimal performance in flight.

Measurements made in flight showed behavior for each lamp that was qualitatively similar to that observed on ground and further analysis using the full simulation showed reasonable quantitative agreement. Despite the

estimated internal sensor pressure being around a millibar for over 15 months prior to the launch, the test mass quantum yields changed by an average of only 22% while for the electrode housing surfaces it was 31%. A separate analysis using measurements made during routine fast discharges showed that the apparent yields for lamps 02 and 06 were stable within the measurement uncertainty 140 days after the system was vented to space and consistent with another measurement made during commissioning just 25 days after venting. The measured apparent yields were stable with time under vacuum, UV irradiation from discharging, cosmic-ray irradiation, and a temperature reduction of about 11℃. This analysis also demonstrated that the optical chain transmission between each lamp and the sensors did not degrade over the course of the mission.

Moving toward development of the LISA discharge system, the issues of suboptimal illumination of the housing, variability in the quantum yield of air-exposed gold surfaces, and the possible asymmetry in test mass and housing yields will need to be addressed. While all the lamps were capable of fast discharge, these issues conspired to make all but one of the four housing illuminations less than ideal for automated continuous discharging as they lacked a significant negative apparent yield at a test mass potential of 0 V. Given the lessons learned from the Pathfinder CMS there are essentially two options for the LISA system if continuous discharge capability is deemed desirable. The first would be to understand and ultimately reduce the quantum yield variability and asymmetry allowing for a simple system similar to Pathfinder, though likely replacing the mercury lamps with UVLEDs. Alternatively, it could be accepted that these factors cannot be well controlled in a practical way and instead design a more complicated modulated CMS that can mitigate against their effect. Given the early stage of LISA development, it is likely that both strategies will continue to be explored.

APPENDIX: UV LIGHT ABSORPTION PERCENTAGES

TABLE III. Percentage of total UV light injected that is absorbed by each region, belonging to either the electrode housing (EH) or test mass (TM). X, Y, Z, and $+/-$ denote the different faces of the housing and test mass.

	Sensor 01						Sensor 02					
	Lamp 00 (TM)		Lamp 02 (EH)		Lamp 04 (EH)		Lamp 01 (TM)		Lamp 05 (EH)		Lamp 06 (EH)	
	EH	TM	EH	TM	EH	TM	EH	TM	EH	TM	EH	TM
$-$X Housing	0.17	0.01	1.27	0.24	0.83	0.05	2.30	12.09	0.30	0.02	1.21	0.22
$-$X Sensing electrode 01	0.00	0.00	0.21	0.09	0.00	0.00	0.14	0.13	0.00	0.00	0.18	0.09
$-$X Sensing electrode 02	0.00	0.00	0.02	0.02	0.00	0.00	1.72	1.17	0.00	0.00	0.02	0.02
$+$X Housing	2.85	7.44	0.88	0.06	1.26	0.24	0.52	0.03	1.15	0.21	0.59	0.04
$+$X Sensing electrode 01	0.78	0.94	0.00	0.00	0.02	0.02	0.00	0.00	0.02	0.02	0.00	0.00

(Table continued)

TABLE III. *(Continued)*

	Sensor 01						Sensor 02					
	Lamp 00 (TM)		Lamp 02 (EH)		Lamp 04 (EH)		Lamp 01 (TM)		Lamp 05 (EH)		Lamp 06 (EH)	
	EH	TM	EH	TM	EH	TM	EH	TM	EH	TM	EH	TM
+X Sensing electrode 02	0.07	0.07	0.00	0.00	0.20	0.09	0.00	0.00	0.18	0.09	0.00	0.00
−Y Housing	2.48	7.33	1.20	0.19	0.78	0.02	0.58	0.03	0.28	0.01	1.12	0.19
−Y Sensing electrode 01	0.66	0.63	0.16	0.07	0.00	0.00	0.00	0.00	0.00	0.00	0.13	0.06
−Y Sensing electrode 02	0.02	0.01	0.01	0.01	0.00	0.00	0.00	0.00	0.00	0.00	0.01	0.01
−Y Injection electrode 00	0.09	0.07	0.05	0.04	0.00	0.00	0.00	0.00	0.00	0.00	0.05	0.04
+Y Housing	0.21	0.01	0.83	0.02	1.18	0.19	1.78	12.01	1.06	0.19	0.50	0.02
+Y Sensing electrode 01	0.00	0.00	0.00	0.00	0.15	0.06	1.86	0.99	0.11	0.05	0.00	0.00
+Y Sensing electrode 02	0.00	0.00	0.00	0.00	0.01	0.01	0.01	0.00	0.01	0.01	0.00	0.00
+Y Injection electrode 00	0.00	0.00	0.00	0.00	0.05	0.04	0.14	0.12	0.06	0.04	0.00	0.00
−Z Housing	4.79	1.04	8.14	3.17	8.07	3.10	5.10	2.12	7.24	2.85	7.79	3.03
−Z Sensing electrode 01	0.01	0.00	5.19	0.91	0.09	0.13	0.43	0.97	0.03	0.05	4.24	0.94
−Z Sensing electrode 02	0.44	0.39	0.09	0.14	4.82	0.92	0.06	0.01	2.21	1.09	0.06	0.11
−Z Injection electrode 01	0.25	0.04	0.85	0.77	1.36	0.23	0.33	0.09	0.47	0.06	0.56	0.60
−Z Injection electrode 02	0.10	0.03	1.49	0.26	0.81	0.72	0.48	0.05	0.54	0.31	1.07	0.15
+Z Housing	0.48	0.01	0.11	0.01	0.12	0.01	0.18	0.01	0.21	0.01	0.14	0.01
+Z Sensing electrode 01	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
+Z Sensing electrode 02	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
+Z Injection electrode 01	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
+Z Injection electrode 02	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
All above surfaces	13.41	18.03	20.50	5.99	19.73	5.83	15.63	29.85	13.86	5.02	17.69	5.51
Test mass corner domes	⋯	42.58	⋯	0.77	⋯	0.74	⋯	23.17	⋯	0.69	⋯	0.69
Test mass corner regions	⋯	5.40	⋯	0.47	⋯	0.42	⋯	11.40	⋯	0.22	⋯	0.34
Test mass central recess	⋯	0.06	⋯	2.62	⋯	2.37	⋯	0.08	⋯	0.91	⋯	2.06
Central plungers (Wasted)	0.01	⋯	0.11	⋯	0.11	⋯	0.01	⋯	0.17	⋯	0.13	⋯
Caging fingers (Wasted)	5.64	⋯	11.80	⋯	11.61	⋯	2.97	⋯	11.08	⋯	11.26	⋯
Housing recesses (Wasted)	14.88	⋯	57.75	⋯	59.19	⋯	16.89	⋯	68.06	⋯	62.33	⋯
Total useful light	13.41	66.07	20.50	9.85	19.73	9.37	15.63	64.50	13.86	6.84	17.69	8.60

[1] P. Amaro-Seoane *et al.*, Laser interferometer space antenna, arXiv:1702.00786.

[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] M. Armano *et al.*, Beyond the required LISA free-fall performance: New LISA Pathfinder results down to 20 μHz, Phys. Rev. Lett. **120**, 061101 (2018).

[4] D. N. A. Shaul, H. M. Araújo, G. K. Rochester, T. J. Sumner, and P. J. Wass, Evaluation of disturbances due to test mass charging for LISA, Classical Quantum Gravity **22**, S297 (2005).

[5] M. Armano *et al.*, Charge-induced force noise on free-falling test masses: Results from LISA Pathfinder, Phys. Rev. Lett. **118**, 171101 (2017).

[6] H. M. Araújo, P. Wass, D. Shaul, G. Rochester, and T. J. Sumner, Detailed calculation of test-mass charging in the LISA mission, Astropart. Phys. **22**, 451 (2005).

[7] P. J. Wass, H. M. Araújo, D. N. A. Shaul, and T. J. Sumner, Test-mass charging simulations for the LISA Pathfinder mission, Classical Quantum Gravity **22**, S311 (2005).

[8] Y. Jafry and T. J. Sumner, Electrostatic charging of the LISA proof masses, Classical Quantum Gravity **14**, 1567 (1997).

[9] M. Armano *et al.*, Measuring the Galactic cosmic ray flux with the LISA Pathfinder radiation monitor, Astropart. Phys. **98**, 28 (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, Astrophys. J. **854**, 113 (2018).

[11] C. Reigber, H. Luhr, and P. Schwintzer, CHAMP mission status, Adv. Space Res. **30**, 129 (2002).

[12] B. D. Tapley, S. Bettadpur, M. Watkins, and C. Reigber, The gravity recovery and climate experiment: Mission overview and early results, Geophys. Res. Lett. **31**, L09607 (2004).

[13] R. Rummel, W. Yi, and C. Stummer, GOCE gravitational gradiometry, J. Geodes. **85,** 777 (2011).

[14] P. Touboul and M. Rodrigues, The MICROSCOPE space mission, Classical Quantum Gravity **18,** 2487 (2001).

[15] S. Buchman, T. Quinn, G. M. Keiser, D. Gill, and T. J. Sumner, Charge measurement and control for the Gravity Probe B gyroscopes, Rev. Sci. Instrum. **66,** 120 (1995).

[16] W. J. Weber, D. Bortoluzzi, A. Cavalleri, L. Carbone, M. Da Lio, R. Dolesi, G. Fontana, C. D. Hoyle, M. Hueller, and S. Vitale, Position sensors for flight testing of LISA drag-free control, in *Astronomical Telescopes and Instrumentation* (SPIE-International Society for Optical Engineering, Bellingham, WA, 2003), p. 31.

[17] D. Bortoluzzi, L. Baglivo, M. Benedetti, F. Biral, P. Bosetti, A. Cavalleri, M. Da Lio, M. De Cecco, R. Dolesi, M. Lapolla, W. Weber, and S. Vitale, LISA Pathfinder test mass injection in geodesic motion: Status of the on-ground testing, Classical Quantum Gravity **26,** 094011 (2009).

[18] H. Audley *et al.*, The LISA Pathfinder interferometry-hardware and system testing, Classical Quantum Gravity **28,** 094003 (2011).

[19] F. Antonucci, A. Cavalleri, R. Dolesi, M. Hueller, D. Nicolodi, H. B. Tu, S. Vitale, and W. J. Weber, Interaction between stray electrostatic fields and a charged free-falling test mass, Phys. Rev. Lett. **108,** 181101 (2012).

[20] E. E. Huber, Jr., The effect of mercury contamination on the work function of gold, Appl. Phys. Lett. **8,** 169 (1966).

[21] G. F. Saville, P. M. Platzman, G. Brandes, R. Ruel, and R. L. Willett, Feasibility study of photocathode electron projection lithography, J. Vac. Sci. Technol. **13,** 2184 (1995).

[22] G. Hechenblaikner, T. Ziegler, I. Biswas, C. Seibel, M. Schulze, N. Brandt, A. Schll, P. Bergner, and F. T. Reinert, Energy distribution and quantum yield for photoemission from air-contaminated gold surfaces under ultraviolet illumination close to the threshold, J. Appl. Phys. **111,** 124914 (2012).

[23] M. O. Schulte, D. N. A. Shaul, D. Hollington, S. Waschke, T. J. Sumner, P. J. Wass, L. Pasquali, and S. Nannarone, Inertial sensor surface properties for LISA Pathfinder and their effect on test mass discharging, Classical Quantum Gravity **26,** 094008 (2009).

[24] W. M. Haynes and D. R. Lide, *CRC Handbook of Chemistry and Physics* (Taylor & Francis, London, 2010).

[25] D. Hollington, Ph.D. thesis, Imperial College London, 2011.

[26] A. Cavalleri, G. Ciani, R. Dolesi, M. Hueller, D. Nicolodi, D. Tombolato, P. J. Wass, W. J. Weber, S. Vitale, and L. Carbone, Direct force measurements for testing the LISA Pathfinder gravitational reference sensor, Classical Quantum Gravity **26,** 094012 (2009).

[27] F. Antonucci, A. Cavalleri, R. Dolesi, A. Hueller, S. Perreca, S. Vitale, and W. J. Weber, Test-mass discharging test campaign on TM3 with the four-test-mass torsion pendulum, Report No. S2-UTN-TN-3080 Issue 1.0, University of Trento, 2010.

[28] P. J. Wass, L. Carbone, A. Cavalleri, G. Ciani, R. Dolesi, M. Hueller, G. Rochester, M. Schulte, T. Sumner, D. Tombolato, C. Trenkel, S. Vitale, and W. Weber, Testing of the UV discharge system for LISA Pathfinder, AIP Conf. Proc. **873,** 220 (2006).

[29] D. Tombolato, Ph.D. thesis, Università di Trento, 2008.

[30] T. Ziegler, P. Bergner, G. Hechenblaikner, N. Brandt, and W. Fichter, Modeling and performance of contact-free discharge systems for space inertial sensors, IEEE Trans. Aerospace Electron. Syst. **50,** 1493 (2014).

[31] S. Agostinelli *et al.*, GEANT4: A simulation toolkit, Nucl. Instrum. Methods Phys. Res., Sect. A **506,** 250 (2003).

[32] J. Allison *et al.*, GEANT4 developments and applications, IEEE Trans. Nucl. Sci. **53,** 270 (2006).

[33] J. Allison *et al.*, Recent developments in GEANT4, Nucl. Instrum. Methods Phys. Res., Sect. A **835,** 186 (2016).

[34] P. B. Johnson and R. W. Christy, Optical constants of the noble metals, Phys. Rev. B **6,** 4370 (1972).

[35] E. D. Palik and G. Ghosh, *Handbook of Optical Constants of Solids* (Academic, New York, 1998).

[36] K. E. Torrance and E. M. Sparrow, Theory for off-specular reflection from roughened surfaces, J. Opt. Soc. Am. **57,** 1105 (1967).

[37] W. J. Weber, L. Carbone, A. Cavalleri, R. Dolesi, C. D. Hoyle, M. Hueller, and S. Vitale, Possibilities for measurement and compensation of stray DC electric fields acting on drag-free test masses, Adv. Space Res. **39,** 213 (2007).

[38] C. C. Speake, Forces and force gradients due to patch fields and contact-potential differences, Classical Quantum Gravity **13,** A291 (1996).

[39] K.-S. Sun, S. Buchman, B. Allard, S. Williams, and R. L. Byer, LED deep UV source for charge management of gravitational reference sensors, Classical Quantum Gravity **23,** S141 (2006).

[40] K.-S. Sun, N. Leindecker, S. Higuchi, J. Goebel, S. Buchman, and R. L. Byer, UV LED operation lifetime and radiation hardness qualification for space flights, J. Phys. Conf. Ser. **154,** 012028 (2009).

[41] D. Hollington, J. T. Baird, T. J. Sumner, and P. J. Wass, Characterising and testing deep UV LEDs for use in space applications, Classical Quantum Gravity **32,** 235020 (2015).

[42] D. Hollington, J. T. Baird, T. J. Sumner, and P. J. Wass, Lifetime testing UV LEDs for use in the LISA Charge Management System, Classical Quantum Gravity **34,** 205009 (2017).