

Investigation of low-energy particle remnants in high-energy collisions at the LHC with a skipper-CCD detector

Brenda A. Cervantes-Vergara^{1,*} Santiago E. Perez^{1,2,3,†} Nicola Bacchetta^{1,4} Nuria Castelló-Mor^{1,5}
 Juan Estrada,¹ Marcos Fernandez Garcia^{1,5} Petra Merkel,¹ María Pérez Martínez^{1,5}
 Dario Rodrigues^{1,2,3} Javier Tiffenberg,^{1,3} and Rocío Vilar Cortabitarte^{1,5}

¹*Fermi National Accelerator Laboratory, Batavia, Illinois, USA*

²*Instituto de Física de Buenos Aires (IFIBA), CONICET, Buenos Aires, Argentina*

³*Universidad de Buenos Aires, Buenos Aires, Argentina and Departamento de Física, FCEN,*

Universidad de Buenos Aires, Buenos Aires, Argentina

⁴*Istituto Nazionale di Fisica Nucleare (INFN), Padova, Italy*

⁵*Instituto de Física de Cantabria, Santander, Spain*



(Received 29 August 2025; accepted 5 January 2026; published 10 February 2026)

We deployed the Mobile Skipper Testing Apparatus ~ 33 m away from the Compact Muon Solenoid collision point, the first skipper-CCD detector probing low-energy particles produced in high-energy collisions at the Large Hadron Collider. In this work, we search for beam-related events using data collected in 2024 during beam-on and beam-off periods. The dataset corresponds to integrated luminosities of 113.3 fb^{-1} and 1.54 nb^{-1} for the proton-proton and Pb-Pb collision periods, respectively. We report observed event rates in a model-independent framework across two ionization regions: $\leq 20e^-$ and $>20e^-$. For the low-energy region, we perform a likelihood analysis to test the null hypothesis of no beam-correlated signal. We found no significant correlation during proton-proton and Pb-Pb collisions. For the high-energy region, we present the energy spectra for both collision periods and compare event rates for images with and without luminosity. We observe a slight increase in the event rate following the Pb-Pb collisions, coinciding with a rise in the single-electron rate, which will be investigated in future work. Using the low-energy proton-proton results, we place 95% confidence level constraints on the mass-millicharge parameter space of millicharged particles. Overall, the results in this work demonstrate the viability of skipper-CCD technology to explore new physics at high-energy colliders and motivate future searches with more massive detectors.

DOI: [10.1103/mm7w-6s2t](https://doi.org/10.1103/mm7w-6s2t)

I. INTRODUCTION

The search for new physics at the Large Hadron Collider (LHC) has traditionally focused on high-energy signatures predicted by various extensions of the Standard Model [1–3]. However, a complementary and increasingly promising approach involves probing low-energy particle remnants produced in LHC's high-energy collisions. This unexplored regime may reveal subtle signals from weakly interacting or hidden-sector particles, such as millicharged particles (mCPs), which are predicted in well-motivated

theoretical scenarios involving kinetic mixing between Standard Model and hidden photons [4].

The skipper-CCD technology has already demonstrated world-leading sensitivity in the search for such particles. Its unique ability to resolve individual ionization electrons, combined with its low-energy threshold [5], enables the detection of faint signatures like those expected from mCP interactions. The SENSEI experiment at the MINOS cavern at Fermilab [6] used this technology to set world-leading constraints on mCPs with masses between 30 and 380 MeV using the NuMI beam [7]. Also, the CONNIE and Atucha-II Collaborations, with skipper-CCD detectors operating at nuclear power plants in South America, have placed the strongest exclusion limits for mCPs in the 1 eV to 700 keV mass range [8]. However, the production of heavier mCPs is highly suppressed in both of these environments, as the limited center-of-mass energies available for Drell-Yan production in fixed-target beams and for Compton-like production in reactors prevent efficient generation of heavier mCP pairs.

*Contact author: bcervant@fnal.gov

†Contact author: santiep.137@gmail.com

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International license](https://creativecommons.org/licenses/by/4.0/). Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI. Funded by SCOAP³.

The 13.6 TeV proton-proton collisions at the LHC open a new window to probe GeV-mass mCPs. The milliQan experiment has recently placed world-leading limits for mCPs with masses above 0.45 GeV [9]. Similar efforts are ongoing at several locations along LHC [10–13], including both upgrades to existing detectors and new proposals for the High-Luminosity LHC (HL-LHC).

To contribute to this program, we deployed MOSKITA (Mobile Skipper Testing Apparatus) at the LHC, the first skipper-CCD detector sensitive to low-energy ionization signals from high-energy collisions. MOSKITA enables the characterization of both environmental and beam-related ionization backgrounds, as well as the search for anomalous excesses that could indicate new physics.

In this work, we present the first results from MOSKITA, using data collected in 2024 during proton-proton collisions, Pb-Pb collisions, and beam-off periods. We provide details on the detector installation, performance, data acquisition, and analysis pipeline. We report the observed event rates in a model-independent framework across two ionization regions: $\leq 20e^-$ and $>20e^-$. For the low-energy region, we report p -values and 95% confidence level (CL) upper limits on the null hypothesis of no beam-correlated events, based on a profile likelihood analysis. We found no significant correlation during proton-proton and Pb-Pb collisions. For the high-energy region, we present the energy spectra for both collision periods, comparing rates for images with and without luminosity. We find an increase in the high-energy event rate following the Pb-Pb collisions, coinciding with a rise in the single-electron rate. Finally, using the low-energy results for the proton-proton collision period, we place constraints on the mass-millicharge parameter space for mCPs at a 95% CL, and project sensitivities for scenarios with a larger detector, reduced background, and higher luminosity.

II. MOSKITA AT THE LHC

A. Experimental setup

MOSKITA is a dedicated setup that has been previously used for testing skipper-CCDs for the Oscura experiment at FNAL [14,15]. It consists of a stainless steel T-shaped vacuum vessel, placed inside a 2-inch-thick lead shield, which can accommodate a CCD sensor. All the hardware needed for its operation is located on a 1.2 m² aluminum pallet.

MOSKITA was shipped to CERN and installed near LHC Interaction Point 5 in the Compact Muon Solenoid (CMS) drainage gallery ~ 70 m underground. This gallery, in which the milliQan experiment is located [16,17], is ~ 33 m away from the interaction point, separated by ~ 17 m of rock that shields the site from high-energy particles produced in collisions. According to the CMS coordinate system [18], MOSKITA is placed at an azimuthal angle ϕ of 43° and pseudorapidity η of 0.1. Figure 1




FIG. 1. MOSKITA at the CMS drainage gallery. The packaged sensor, mounted on a copper tray, hangs from the cold finger inside the stainless-steel vessel (red box). The lead shield encloses part of the vessel, fully covering the region where the sensor lies.

shows a picture of MOSKITA installed at the CMS drainage gallery and Fig. 2 shows a diagram indicating the setup's position relative to the CMS detector.

We use as detector a 6.29 MPix p-channel skipper-CCD, designed by the Lawrence Berkeley National Laboratory and fabricated at Teledyne-DALSA for the SENSEI experiment [19]. The sensor's active area is an array of 1024 rows and 6144 columns of $15 \times 15 \mu\text{m}^2$ pixels and has a thickness of 675 μm , which gives a total mass of ~ 2.2 grams. This sensor has four floating-gate amplifiers, one in each corner, that allow it to perform multiple nondestructive measurements of the charge in each pixel during readout. The number of measurements taken from the same charge packet will be referred to as N_{smp} . The sensor's packaging was done at FNAL. In this process, the sensor and a copper-Kapton flexible cable are glued to a silicon substrate and wirebonded; then, the whole assembly is placed inside a copper tray.

Once MOSKITA was fully assembled at the LHC in March 2024, we installed the packaged sensor inside the vessel. The sensor's active area is perpendicular to the ground and is not oriented toward any specific direction. During operations, the vessel holds a vacuum of $\sim 10^{-4}$ Torr, and the sensor is cooled down to ~ 140 K




FIG. 2. MOSKITA's location relative to CMS.

using a Sunpower CryoTel cryocooler to minimize thermally-generated charge. We use a low-threshold acquisition board [20] to bias and read out the sensor. The temperature is monitored at two locations: on the sensor's copper tray (CCD temperature) and on the cryocooler's cold head, next to the heater used to keep the sensor's temperature constant (heater temperature). The pressure is monitored using a MicroPirani vacuum gauge.

The sensor's performance was previously evaluated at FNAL, where three out of the four amplifiers, located in the output stages labeled 2, 3 and 4 in Fig. 3, were working as expected under operating conditions. Physically, the sensor has two serial registers, which are the last rows of pixels used to transfer the charge from the active area to the output stages. In this sensor, output stages 1 and 4, as well as 2 and 3, share a serial register. After installing the sensor in MOSKITA at the LHC, we found that half of the serial register next to output stage 2 presented problems in transferring charge with the usual operating voltages, as shown in Fig. 4. During data acquisition, each amplifier read out one quarter of the CCD (quadrant); hence, the only useful quadrants, named 1 and 2 from now on, are those read out through amplifiers in output stages 3 and 4, respectively.




FIG. 3. Diagram of the CCD installed in MOSKITA (not to scale). The two serial registers (blue) and the four output stages (red) are shown. The useful quadrants, when the CCD is read out through all four amplifiers, are labeled as quadrants 1 and 2.




FIG. 4. Image taken in MOSKITA at the LHC, rearranged as in Fig. 3. It is a 3.5-h exposure with a 15-min readout through output stages 3 and 4. The charge transfer problem in the half of the serial register next to output stage 2 is clear.

B. Data acquisition (DAQ) and operating conditions

The data acquisition cycle for one image consists of two phases: cleaning and readout. During the cleaning phase, charge in the active area of the CCD is transferred and discarded through the output stages 1 and 2 for 10 min. Thus, the readout begins with a “clean” CCD. Immediately after cleaning, we start the pixel-by-pixel readout through the four amplifiers. Each amplifier reads 515 rows and 3200 columns of pixels, measuring N_{smp} times the charge in each pixel.

Of the 515 rows, 512 correspond to the active area, that is pixels within the physical array. The remaining three correspond to the vertical “overscan” area, which consists of pixels that extend beyond the physical array clocked towards the output stage. Of the 3200 columns, the first eight correspond to the “prescan” area, which arises from eight extra physical pixels in the serial register connecting to the output stage. The next 3072 columns correspond to the active area, followed by 20 columns of horizontal overscan. The remaining 100 columns correspond to the horizontal “reverse overscan” area, that is pixels beyond the physical array clocked away from the output stage. The entire serial register is clocked to discard charge before each row of pixels in the active area is vertically transferred.

In each readout, data are stored as a two-dimensional Flexible Image Transport System (FITS) image with 515 rows and $(3200 \times N_{\text{smp}})$ columns, referred to as a “raw” image, to which a consecutive “image ID” number is assigned. In this image, every N_{smp} consecutive pixels in the same row correspond to the N_{smp} measurements of the charge in the same pixel, and their values are in analog-to-digital units (ADUs). It takes ~ 3.2 (4.8) h to acquire one raw image with $N_{\text{smp}} = 169$ (256) samples/pixel, resulting in a nonuniform exposure across pixels that increases linearly with the pixel readout time. Given this readout time frame, the image DAQ is independent of beam activity. Raw images are processed afterward for monitoring and analysis, as will be described in the following section. The pixel distributions of the processed images form a set of Gaussian peaks centered on integer multiples of the gain, a parameter that relates ADUs to the number of generated electron-hole pairs.

MOSKITA at LHC has been taking data continuously since March 2024. We divide the DAQ period into “runs,” each consisting of a set of images acquired under the same conditions. In this analysis, we use data from the first year

TABLE I. Runs during DAQ period. $V_{A/B}$ refer to the sensor's operating voltages used in each run. Thermal cycles are indicated with red horizontal lines.

Run	Dates	No. of images	Conditions
1	March 15–May 22	462	V_A ; $N_{\text{smp}} = 169$
2	May 22–July 31	500	V_B ; $N_{\text{smp}} = 169$
3	August 4–September 14	293	V_B ; $N_{\text{smp}} = 169$
4	September 14–October 17	159	V_B ; $N_{\text{smp}} = 256$
5	October 19–December 9	236	V_B ; $N_{\text{smp}} = 256$
6	December 21–31	50	V_B ; $N_{\text{smp}} = 256$

of operations, whose information is summarized in Table I. In Run 1, the sensor's operating voltages caused pixel saturation at $\sim 1400 e^-$ in quadrant 2. For subsequent runs, we changed the voltage applied to one of the gates in the sensor's output stage to achieve pixel saturation at typical levels, i.e., $\gtrsim 8000 e^-$ since a lower saturation value could impact event reconstruction at high energies. System thermal cycles to temperatures above 200 K were performed after Runs 2, 4, and 5. The operating conditions of the setup during each run were stable, maintaining a sensor temperature of 140 K and a pressure below 1×10^{-4} Torr; this is shown in Fig. 5.

We continuously monitor the readout noise, gain, and single-electron rate (SER) of each acquired image. These parameters are obtained from a three-Gaussian fit to the first three peaks of the pixel distributions of the processed images, corresponding to pixels with 0, 1, and $2 e^-$. Assuming the Gaussian fitted to the $k e^-$ peak has parameters (a_k, σ_k, μ_k) corresponding to its amplitude, standard deviation, and mean, respectively, we define the gain as $(\mu_2 - \mu_0)/2$, the SER as a_1/a_0 , and the readout noise as σ_0 .

Figures 6(a)–6(c) show the evolution of these parameters in each working quadrant over time, with the SER and the noise scaled by the gain. The gain remains stable over time, with a smaller spread during Runs 4, 5, and 6 due to improved fits resulting from the reduced noise at higher N_{smp} . Similarly, the SER is stable overall in each run, with a slight increase observed toward the end of Runs 2 and 5. As expected, the mean SER in Runs 4, 5, and 6 is higher due to




FIG. 5. Pressure and temperature of the experimental setup over eight months within the DAQ period. Spikes in the CCD temperature correspond to system thermal cycles.




FIG. 6. Evolution over time of (a) gain, (b) single-electron rate, and (c) readout noise during the DAQ period. Run periods are highlighted.

the longer image readout time associated with the increased N_{smp} ; however, the SER in Run 6 is 17% higher than in Runs 4 and 5. In contrast, the readout noise shows significant variations over time exhibiting the highest variation during Run 2, with 2.8% for quadrant 1 and 3.5% for quadrant 2. The mean value of each parameter, along with its standard deviation, obtained for each run are shown in Tables II and III for quadrants 1 and 2, respectively.

TABLE II. Gain, single-electron rate, and readout noise mean values for quadrant 1 in each run.

Run	Gain [ADU/ e^-]	SER [$10^{-3} e^-/\text{pix}$]	Noise [e^-]
1	230.4 ± 1.5	1.4 ± 0.2	0.204 ± 0.001
2	230.0 ± 1.9	1.5 ± 0.3	0.214 ± 0.006
3	230.4 ± 1.5	1.4 ± 0.2	0.208 ± 0.001
4	229.4 ± 0.8	2.0 ± 0.3	0.166 ± 0.001
5	229.6 ± 0.8	2.1 ± 0.3	0.169 ± 0.001
6	230.2 ± 0.8	2.4 ± 0.3	0.167 ± 0.001

TABLE III. Gain, single-electron rate, and readout noise mean values for quadrant 2 in each run.

Run	Gain [ADU/ e^-]	SER [$10^{-3} e^-/\text{pix}$]	Noise [e^-]
1	223.1 ± 1.4	1.2 ± 0.2	0.212 ± 0.001
2	222.9 ± 1.7	1.2 ± 0.2	0.227 ± 0.008
3	222.7 ± 1.7	1.2 ± 0.2	0.229 ± 0.001
4	221.9 ± 0.9	1.7 ± 0.2	0.183 ± 0.003
5	222.0 ± 1.0	1.8 ± 0.2	0.200 ± 0.006
6	222.4 ± 0.8	2.1 ± 0.3	0.190 ± 0.001

1. The LHC beam during DAQ

DAQ occurred during Run 3 of the LHC. In its 2024 run, the LHC conducted a proton run from March to October, followed by a lead-ion run that concluded in November. Proton-proton collisions at an energy of 13.6 TeV began on March 14, 2024. After achieving stability, the beam intensity steadily increased, reaching its maximum at the end of April with a recorded instantaneous luminosity of $\sim 2.12 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$. At full intensity, the LHC beam consisted of 2352 bunches separated by 25 ns, each containing 1.6×10^{11} protons. The beam operates in periods of ~ 10 h, during which the instantaneous luminosity maintains its initial peak for ~ 4 h before decaying exponentially until the beam is dumped. The proton run concluded on October 16, 2024, delivering a final integrated luminosity of 124.4 fb^{-1} to the CMS experiment. After the proton run, a 6-day period of proton-proton collisions at 5.36 TeV took place for calibration for the lead-ion run. Pb-Pb collisions at 5.36 TeV began on November 6, 2024. Full intensity was achieved on November 10, 2024, with 1240 bunches per beam separated by 50 ns. During this run, the beam operated in periods of ~ 6 h, maintaining its initial luminosity peak for about one hour before exponentially decaying until being dumped. The lead-ion run concluded on November 23, 2024, delivering a final integrated luminosity of 1.9 nb^{-1} to the CMS experiment. Since the DAQ is not synchronized with the beam periods, we associate to each acquired image the integral of the instantaneous luminosity over its readout time. Figure 7 shows the LHC beam instantaneous luminosity recorded by the CMS detector during the 2024 run.




FIG. 7. Instantaneous luminosity of the LHC beam recorded by the CMS detector during the 2024 run for proton-proton (black) and Pb-Pb (blue) collisions. For reference, the time periods of MOSKITA DAQ runs are highlighted.

III. DATA PROCESSING

We perform a two-step image processing. The first step is applied to each raw image, where we build its corresponding “processed” image by assigning the average of its N_{smp} measurements, excluding the first one which is noisier, to each pixel. As a result, the processed image has N_{smp} times fewer columns compared to the raw image. Additionally, the processed image includes a row-by-row subtraction of the horizontal baseline, computed as the median of the horizontal reverse overscan pixels in the row, which are considered to have zero charge. The pixel values in the processed images are in ADUs. A second processing step is applied to each processed image, where we first perform a column-by-column subtraction of the vertical baseline, computed as the median of all pixels in the column, and then we build its corresponding “calibrated” image by dividing each pixel value by the absolute gain per quadrant. Pixel values in the calibrated images are in units of electrons.

A. Energy calibration

For each run and quadrant, we sum the pixel distributions of all processed images and perform a Gaussian fit to each of the first 850 peaks. Figure 8 shows the fitted pixel




FIG. 8. Pixel distribution corresponding to quadrant 1 in Run 1 in ADUs. The peak number is equivalent to the number of collected charge carriers.

TABLE IV. Average gain in ADU/e^- and the highest percentage variation in energy calibration due to nonlinearity for each quadrant in each run.

Run	Gain [ADU/e^-]		Nonlinearity [%]	
	Quadrant 1	Quadrant 2	Quadrant 1	Quadrant 2
1	226.2	214.7	1.8	4
2	226.3	214.8	1.7	4
3	226.3	214.8	1.9	4
4	226.3	214.7	1.3	3.4
5	226.4	214.8	1.5	3.4
6	226.8	215.1	1.7	3.6

distribution corresponding to quadrant 1 in Run 1, between peak numbers 500 and 550. We plot the mean value of each of those fits against its corresponding peak number and perform a linear fit; the slope extracted from this fit corresponds to the average gain for each quadrant in each run. We find this gain to be essentially constant over time. These values are used for calibration and are shown in Table IV. To estimate the nonlinearity (N.L.) in the energy calibration, we compare the mean of each Gaussian fit, divided by the average gain for each quadrant in each run, to its corresponding peak number and extract the highest percentage variation. These values represent upper limits on the energy calibration nonlinearity, which are below 1.9% for quadrant 1 and 4% for quadrant 2 in all runs; see Table IV. The gain values obtained previously are consistent with this method when taking into account the nonlinearity.

IV. EVENT RECONSTRUCTION

From the calibrated images, we reconstruct events. We define an event as a group of neighboring pixels with charges exceeding the $1 e^-$ detection threshold, computed per image. The total charge of an event is computed as the sum of the integer charges of its constituent pixels.

A. Detection threshold

The minimum amount of charge that a pixel must have to be considered as nonempty, i.e., the $1 e^-$ detection threshold, strongly depends on the noise and the single-electron rate of the image. During the data acquisition period, the SER remained stable, but the noise exhibited some variations, as discussed during the data acquisition section. To select the $1 e^-$ detection threshold accounting for these variations, we proceed as follows. We model the $0 e^-$ and $1 e^-$ pixel distributions of each image as two Gaussians centered at 0 and 1, respectively, with the same standard deviation determined by the image noise value in units of electrons. The amplitude ratio of the $1 e^-$ Gaussian to the $0 e^-$ Gaussian equals the image SER in e^-/pix . We compute the F1 score for the $1 e^-$ signal as a function of the threshold value. We define the optimal threshold for each image as the one that maximizes the F1 score. Figure 9




FIG. 9. Optimal $1 e^-$ detection threshold per image for quadrant 1 (cyan) and 2 (purple) in each run.

shows the optimal detection threshold for each image of each quadrant over the DAQ period.

V. EVENT SELECTION

We use masks, created from the calibrated images, to exclude events from known sources. The masking algorithm is based on the approach used in dark matter searches [6,19]. Low-energy events, with charge $\leq 20 e^-$, are rejected if any of their constituent pixels are masked by any of the masks, while high-energy events, with charge $> 20 e^-$, are rejected if their energy-weighted pixel centroid is masked by the border, hot-column or serial-register-event masks.

We apply an additional selection cut to the low-energy events, beyond the masking-rejection criterion, to exclude single-row events. This cut targets events likely happening near the sensor surface, primarily causing single-pixel events, and artifacts from serial-register hits or charge-transfer inefficiencies in the output stage, which manifest as one-row multiple-pixel events.

The impact of the masks and the selection cut on the detection of low-energy events is incorporated as an efficiency, calculated for each energy bin. This efficiency is computed from simulated image, and accounts for various effects that affect the event reconstruction, including:

- (1) Diffusion: events may split into multiple clusters when diffusing, or part of their charge may fall into masked regions, resulting in only a fraction of the true deposited charge being reconstructed.
- (2) Noise: per-pixel fluctuations can result in events being reconstructed with a charge differing from the actual deposited charge.
- (3) Single-electron rate: multielectron events can form from pileup, and events overlapping with $1 e^-$ events from the SER may be reconstructed with a charge higher than originally deposited.

A. Masking algorithm

We developed and optimized a masking algorithm using only images from Run 1, in order to avoid any bias, and

then applied it to the full dataset for the final analysis. A brief discussion of the masks used in this analysis is presented here:

Border mask: this mask removes events near the image borders, as their measured energy may only be a fraction of the total energy deposited due to their proximity to the edge. Additionally, low-energy events near the borders can originate from depositions occurring outside the active area of the sensor. We mask ten pixels from the top, bottom, and left edges in the images.

Hot-column mask: this mask removes columns with an excess of $1 e^-$ events produced by defects in the sensor's active area. For each run, a single mask is generated by combining the hot-column mask of the two useful quadrants. To create the mask of each quadrant, the $1 e^-$ event rate for each column is computed using all processed images from that run, after removing all events with a charge greater than $20 e^-$ and a halo of five pixels around them. A smoothing filter is applied to the rate, the baseline is removed, and columns with a rate exceeding five median absolute deviations, along with their neighboring columns, are masked. As an example, we show in Fig. 10 the $1 e^-$ event rate per column for each quadrant (top) together with the hot-column mask for Run 2 (bottom).

Transfer-gate-trap mask: we identified a horizontal “tail” of low-energy events following high-energy events with at least one pixel in column 2952 of quadrant 1. We attribute the origin of the events forming the tail to charge emission from a charge trap in the sensor's transfer gate in column 2952. This gate enables charge transfer from the last pixel in that column to the serial register; therefore, a trap in this gate can capture charge and release it later into the serial register. To reject events originating from this trap, we mask the pixels with charge greater than $1 e^-$ in column 2952, along with all pixels to their right in the same row, in all images from quadrant 1. Figure 11 shows this mask for an image of quadrant 1.

Serial-register-event mask: this mask, created for each image, removes primarily horizontal, single-row events,




FIG. 10. Top: $1 e^-$ event rate per column for quadrants 1 (cyan) and 2 (purple). Bottom: hot-column mask for Run 2.




FIG. 11. Transfer-gate-trap mask for an image of quadrant 1. Events containing a high-energy pixel in column 2952 show a deferred charge tail. The mask (cyan) is shown superimposed on the image, and nonempty pixels inside the masked region (yellow) are highlighted.

which arise from charge depositions either directly within the serial register or diffusing into it from the surrounding inactive silicon. After applying the hot-column and the transfer-gate-trap masks, we flag events as serial-register events based on two criteria: (1) their size, and (2) the distance between events. For the size-based criteria, events are flagged if they meet either of the following conditions: (a) $l_y = 1$ pixel and $l_x > 4$ pixels, or (b) $l_y = 2$ or 3 pixels and $l_x > 7$ pixels, where l_x and l_y are the event sizes in the horizontal (row) and vertical (column) directions, respectively. These conditions are based on the event sizes of pointlike charge depositions occurring on the backside of the sensor, computed from simulations using the diffusion model described in Ref. [6]. Figure 12 shows the horizontal size distributions of simulated events with $l_y = 1$ pixel (top) and $l_y = 2$ or 3 pixels (bottom), with varying charge values.

The distance-based criteria for flagging an event with $l_y \leq 3$ pixels as a serial-register event are as follows: (1) non-single-pixel events must have a shape consisting of a single row with $l_x > 1$ pixel and immediate upper and lower rows with $l_x \leq 1$ pixel, and be within a horizontal distance of 30 pixels from another event with $l_y \leq 3$ pixels; and (2) single-pixel events require three or more of such events in the same row within a 60-pixel window. These criteria are based on the distance between events in simulated images containing uniformly distributed single-pixel events, with their total number matching the observed events with $l_y \leq 3$ pixels in Run 1 images. Figure 13 (top) shows the probability of finding at least one event at a given horizontal distance from a randomly chosen reference event, computed from the simulated images, with the 30-pixel criterion corresponding to a $\sim 5\%$ probability. Figure 13 (bottom) shows the number of single-pixel events in one-row 60-pixel windows for the simulated and Run 1 images. Once an event is flagged as a serial-register event, we mask the entire row(s) associated with it.

Crosstalk mask: electronic crosstalk between amplifiers can cause a small signal to leak from one channel with a strong signal to another. To remove events caused by crosstalk, we create a mask for each image that flags a




FIG. 12. Horizontal size distributions of simulated events with different charge values and vertical sizes: $l_y = 1$ pixel (top) and $l_y = 2$ or 3 pixels (bottom). For the plots, 10^6 events were simulated per charge value.

pixel if it was read at the same time as another pixel with a charge higher than $1000 e^-$. This threshold was set based on the crosstalk level measured from Run 1 images ($\sim 1.4 \times 10^{-4} e^-$), which would induce signals comparable to the readout noise.

Bleeding-zone mask: this mask, created for each image, removes spurious events from horizontal and vertical charge-transfer inefficiencies. For each pixel with a charge exceeding $20 e^-$, we mask 100 pixels to the right and 30 pixels above. These values were determined by analyzing the SER in Run 1 images after applying the hot-column, serial-register-event, transfer-gate-trap, and crosstalk masks. The chosen values correspond to the number of masked pixels beyond which the SER remained constant.

Halo mask: high-energy events correlate with an increased rate of low-energy events in nearby pixels, associated with secondary radiation. This mask, created for each image, flags all pixels within a 25-pixel radius of those with charges exceeding $20 e^-$. The radius was determined similarly to the parameters in the bleeding-zone mask, based on the radius at which the SER stabilized after applying the hot-column, serial-register-event, transfer-gate-trap, crosstalk, and bleeding-zone masks.




FIG. 13. Top: probability of finding an event as a function of the horizontal distance from a reference event in simulated images (blue). The 30-pixel criterion for non-single-pixel events corresponds to a $\sim 5\%$ probability (dashed red). Bottom: number of single-pixel events in 60-pixel windows along one row for simulated (red) and Run 1 (blue) images. Events to the right of the dashed black line are flagged as serial-register events.

Hot-zone mask: after applying all the previously discussed masks to a set of images, the combined event distribution is expected to be spatially uniform. However, when plotting the spatial distribution of the remaining unmasked $1 e^-$ events per quadrant for all images in Run 1, we identify regions in the sensor with an excess of events, referred to as “hot zones.” Figure 14 shows the combined spatial distribution of the remaining unmasked $1 e^-$ events from all Run 1 images for quadrant 1 (top) and quadrant 2 (bottom). Masked regions were defined from the column-projected distributions. Columns with a moving-average $1 e^-$ rate exceeding the 90th percentile are masked, and broader regions with an overall $1 e^-$ event excess are masked if their rate consistently exceeds two standard deviations above a baseline estimated from areas where the rate is minimal and nearly flat. In Fig. 14, the masked hot zones correspond to columns with most bins above ~ 70 events.

B. Detection efficiency for low-energy events

We compute the detection efficiency for events with a charge between 2 and $20 e^-$ from simulated images.




FIG. 14. Two-dimensional histograms showing the spatial distribution of $1 e^-$ events from all Run 1 images after applying the border, hot-column, serial-register-event, transfer-gate-trap, crosstalk, bleeding-zone, and halo masks for quadrant 1 (top) and quadrant 2 (bottom). Each bin represents 20×20 pixels.

To account for variations in noise and SER during the data-taking period, we construct each simulated image using the parameters and corresponding mask of a randomly selected image from all runs.

Each simulated image contains spatially uniform $1 e^-$ events, with a total number matching the SER of the selected image, and a single multielectron event with a charge between 2 and 21 e^- , simulated below unmasked pixels. The deposition depth of the multielectron event is randomly selected from a uniform distribution along the sensor depth, and its charge is diffused toward the surface using the model described in Ref. [6]. To account for readout noise, the charge in each pixel is adjusted based on the noise of the selected image. Events are then reconstructed using the energy threshold corresponding to the selected image, and the masking-rejection criterion and the selection cut are applied.

From a set of 50,000 simulated images, the detection efficiency for each low-energy bin is calculated as the ratio

of the total number of simulated multielectron events with that energy to the number of reconstructed and nonrejected multielectron events with the same energy. The dominant systematic uncertainty in the efficiency calculation arises from the diffusion model. As reported in previous dark matter searches using similar skipper-CCDs [19], this model leads to simulated low-energy clusters that are slightly wider than those observed in the data, leading to a conservative 3% underestimate of the efficiency.

VI. EVENT RATES

For the period of proton-proton collisions, we include images from Runs 1 to 4 and Run 6, resulting in a total exposure of 116.20 g-day, combining data from both quadrants, and a luminosity of 113.3 fb^{-1} . For the period of Pb-Pb collisions, we include images from Runs 5 and 6, yielding a total exposure of 31.85 g-day, combining data from both quadrants, and a luminosity of 1.54 nb^{-1} . Note that images from Run 6 are included in both periods, as they correspond to a time without collisions. The reported luminosities are calculated by summing the integrated luminosity associated with each image in the period, and are subject to a $\sim 2.5\%$ measurement uncertainty [21].

A. Low-energy region

The total post-masks exposure, ϵ_{tot} (ALL), is calculated from the unmasked regions in the images for each collision period after applying the full set of masks, yielding 39.32 g-day for proton-proton collisions and 9.45 g-day for Pb-Pb collisions. The post-masks exposure for images with zero luminosity, ϵ_{tot} ($L = 0$), is 18.24 g-day for the proton-proton collision period and 7.65 g-day for the Pb-Pb collision period.

Results by energy bin are shown in Table V for proton-proton collisions and Table VI for Pb-Pb collisions. For each period, we count the number of observed events in each bin, N , combining data from both quadrants, after applying the low-energy event selection. These numbers

TABLE V. Results for the period of proton-proton collisions. We expect 271 $2 e^-$ events and one $3 e^-$ event from the pileup of $1 e^-$ events, computed considering the SER of the images for the period.

Bin [e^-]	Observed events		Effective exposure [g-day]		Efficiency	p -value $\mu = 0$	Upper limit 95% CL [events]
	ALL	$L = 0$	ALL	$L = 0$			
2	327	161	11.40	5.29	0.29	0.31	5.88
3	17	11	19.27	8.94	0.49	0.20	1.41
4	1	0	23.59	10.94	0.60	0.12	5.20
5	2	0	26.74	12.40	0.68	0.37	3.12
6	0	0	27.92	12.95	0.71	1	2.45
7	1	0	29.10	13.50	0.74	0.23	4.15
8	2	1	29.49	13.68	0.75	0.35	1.41
9	0	0	30.67	14.23	0.78	1	1.81
10	1	0	31.85	14.77	0.81	0.53	3.73
11–20	9	3	33.42	15.50	0.85	0.41	8.17

TABLE VI. Results for the period of Pb-Pb collisions. We expect fifty-four $2 e^-$ events from the pileup of $1 e^-$ events, computed considering the SER of the images for the period.

Bin [e^-]	Observed events		Effective exposure [g-day]		Efficiency	p -value $\mu = 0$	Upper limit 95% CL [events]
	ALL	L = 0	ALL	L = 0			
2	85	72	2.74	2.22	0.29	0.52	5.80
3	6	5	4.63	3.75	0.49	0.43	2.05
4	0	0	5.67	4.59	0.60	1	1.14
5	2	2	6.43	5.20	0.68	0.57	2.02
6	0	0	6.71	5.43	0.71	1	1.50
7	2	0	6.99	5.66	0.74	0.017	5.74
8	0	0	7.09	5.74	0.75	1	1.74
9	0	0	7.37	5.97	0.78	1	1.46
10	0	0	7.65	6.20	0.81	1	2.30
11–20	6	4	8.03	6.50	0.85	0.71	4.87

are reported in Tables V and VI as total events (ALL) and as events occurring in images with zero luminosity (L = 0). Detection efficiencies for each bin, Eff, are also shown in Tables V and VI. The effective exposures for each bin, computed as $\epsilon = \text{Eff} \times \epsilon_{\text{tot}}$, are included in these tables. Note that the event rate, R , can be computed as $R = N/\epsilon$.

To investigate potential correlations between the observed events and the LHC luminosity, we perform a likelihood analysis testing the background-only hypothesis ($\mu = 0$), as described in the following subsection, for each energy bin and collision period. As we expect the low-energy background to be dominated by Compton scatterings of secondary photons from cosmic radiation and by instrumental few-electron events, both independent of luminosity, we use L = 0 data to estimate the background for L > 0 data. The corresponding p -values, $p_0 \equiv p(\mu = 0)$, for each bin are reported in Tables V and VI for the proton-proton and the Pb-Pb collision periods, respectively.

For the proton-proton collision period, the results are consistent with the background-only hypothesis. In the case of the Pb-Pb collision period, the analysis for the $7 e^-$ bin yields the lowest p -value, suggesting a possible deviation from the background-only hypothesis. However, this deviation may still be attributed to a statistical fluctuation. According to the “look-elsewhere effect,” the probability of observing at least one such fluctuation with p -value p^* in n bins is $\mathcal{P} = 1 - (1 - p^*)^n$, which in our case yields $\mathcal{P} = 0.29$.

The 95% CL upper limits on the number of signal events for each energy bin are also listed in Tables V and VI for both collision periods. Considering the results in the $4\text{--}7 e^-$ energy range, the region used for the mCP search, the corresponding p -values for testing the null hypothesis are $p_0 = 0.096$ for the proton-proton collision period and $p_0 = 0.11$ for the Pb-Pb collision period. Combining these p -values, we get an overall $p_0 = 0.059$, which approaches statistical significance (2σ). Unaccounted low-energy backgrounds that scale with beam luminosity may contribute to

the observed result. One such potential background is low-energy neutrons arriving at the detector from secondary processes from high-energy collisions.

B. Likelihood analysis

The likelihood function is expressed as the product of Poisson probability distribution functions for the total number of images, N_{img} , [22]

$$\mathcal{L}(\mathbf{n}|\mu, \mathbf{b}) = \prod_{i=1}^{N_{\text{img}}} \frac{e^{-(\mu s_i + b_i)} (\mu s_i + b_i)^{n_i}}{n_i!}, \quad (1)$$

where $\mathbf{n} = (n_1, n_2, \dots, n_{N_{\text{img}}})$ with n_i the number of observed events in the i -th image, μ is the signal strength, μs_i is the expected number of beam-related events in the i -th image, and $\mathbf{b} = (b_1, b_2, \dots, b_{N_{\text{img}}})$ with b_i the expected number of background events in the i -th image. The values of s_i and b_i are given as

$$s_i = \Delta E \times \epsilon_i \times L_i \quad \text{and} \quad b_i = \Delta E \times \epsilon_i \times R_b, \quad (2)$$

where ΔE is the energy bin width, ϵ_i is the effective exposure of the i -th image, which includes the bin’s detection efficiency from Tables V and VI, L_i is the luminosity associated to the i -th image, and R_b is the background rate computed from images with $L_i = 0$.

The profile likelihood ratio is given by

$$\lambda(\mu) = \frac{\mathcal{L}(\mathbf{n}|\mu, \hat{\mathbf{b}})}{\mathcal{L}(\mathbf{n}|\hat{\mu}, \hat{\mathbf{b}})}, \quad (3)$$

where $\hat{\mathbf{b}}(\mu)$ is the background that maximizes the likelihood function for a given μ , while $\hat{\mu}$ and $\hat{\mathbf{b}}$ are the global maximum estimators of the likelihood function.

We evaluate the compatibility of the data with different signal strengths using Monte Carlo simulations. For a given μ , we perform many pseudoexperiments. In each, we simulate a full set of images, where the number of events in each image is sampled from a Poisson distribution with mean $\mu s_i + \hat{b}_i(\mu)$, and compute the test statistic $t_\mu = -2 \ln \lambda(\mu)$. The pseudoexperiments yield the t_μ distribution, $f_\mu(t)$, from which we compute the p -value as $p(\mu) = \int_{t_{\text{obs}}(\mu)}^{\infty} f_\mu(\tilde{t}) d\tilde{t}$, with $t_{\text{obs}}(\mu)$ the test statistic computed from the actual data.

The 95% CL upper limit on the number of events is defined as the expected number of signal events corresponding to the largest value of μ for which $p(\mu) \geq 0.05$, denoted as μ_{95} . To compute this, we generate many pseudoexperiments. In each, the number of events per image is sampled from a Poisson distribution with mean $\mu_{95} s_i$, and the total number of signal events is obtained by summing over all images. The upper limit is then given by the mean of these total event counts across all pseudoexperiments.

C. High-energy region

The high-energy event rates for the proton-proton and Pb-Pb collision periods are shown in the top panels of Figs. 15 and 16, respectively. We compare the rates corresponding to images with luminosity ($L > 0$) and without luminosity ($L = 0$) in two energy ranges, from 0 to 20 keV (a) and from 0 to 1000 keV (b). In the region below 20 keV, characteristic x-ray emission peaks are visible. The most prominent ones at 8 and 8.9 keV correspond to the Cu K_α and K_β emission lines, respectively, from the CCD packaging material. In the spectra up to 1000 keV, an expected bump at ~ 220 keV corresponding to cosmic muon events is seen.

The differences in event rates between images with and without luminosity, i.e., $L > 0 - L = 0$, are shown in the bottom panels of Figs. 15 and 16 for the proton-proton and the Pb-Pb collision periods, respectively. These plots show a higher event rate below ~ 450 keV for $L = 0$ in both periods, with a more pronounced difference in the Pb-Pb collision period. Above 450 keV, where events are primarily associated to cosmic ray muons, the rates between images with and without luminosity are consistent.

After a more detailed investigation, we observe an increase in the high-energy event rate beginning after the Pb-Pb collisions and persisting through the end of Run 6, a beam-off period. Interestingly, a similar increase is seen in the single-electron rate over the same time frame; see Fig. 6(b). These correlated increases may share a common origin. One possible explanation is the production of secondary neutrons during the Pb-Pb collisions. Such neutrons could generate point defects in the CCD lattice, creating traps whose emission may contribute to the increased SER, while also activating surrounding materials




FIG. 15. Top panels: high-energy event rates for the proton-proton collision period, for images with $L = 0$ (blue) and $L > 0$ (red). Bottom panels: difference in event rates between images with and without luminosity (black), i.e., $L > 0 - L = 0$. A higher event rate below ~ 450 keV for $L = 0$ is observed. (a) 0 to 20 keV and (b) 0 to 1000 keV.

and leading to delayed high-energy events. However, this remains a hypothesis. The origin of the observed higher rate will be investigated in future work.

Figure 17 shows the differences in event rates between images with and without luminosity for both collision periods, excluding images acquired after the Pb-Pb




FIG. 16. Top panels: high-energy event rates for the Pb-Pb collision period, for images with $L = 0$ (blue) and $L > 0$ (red). Bottom panels: difference in event rates between images with and without luminosity (black), i.e., $L > 0 - L = 0$. A higher event rate below ~ 450 keV for $L = 0$ is observed. (a) 0 to 20 keV and (b) 0 to 1000 keV.

collisions. In this case, the inconsistency below ~ 450 keV vanishes for both periods.

VII. LIMITS ON mCPs

We compute the expected number of mCPs from proton-proton collisions at 13.6 TeV arriving at MOSKITA,




FIG. 17. Difference in event rates between images with and without luminosity, i.e., $L > 0 - L = 0$, from 0 to 1000 keV for the proton-proton (top) and the Pb-Pb (bottom) collision periods, excluding images acquired after the Pb-Pb collisions. Here, the difference is compatible with zero for both periods.

$\phi(\varepsilon, E_\chi, m_\chi)$, using the PYTHIA8 event generator [23], previously employed by the milliQan collaboration [17]. This simulation models mCP production from hadronic decays and their propagation through the 17 m of rock to the drainage gallery, accounting for the CMS magnet's magnetic field. It does not include mCP production via proton bremsstrahlung. For our analysis, we modify the simulation to reflect a luminosity of 113.3 fb^{-1} and the geometry of the MOSKITA sensor.

An mCP reaching the sensor would interact with electrons in the silicon according to the cross section [24]

$$\frac{d\sigma}{d\omega} = \frac{8\alpha\varepsilon^2}{n_e\beta^2} \int_0^\infty dk \left\{ \frac{1}{k} \text{Im} \left(-\frac{1}{\tilde{\varepsilon}(\omega, k)} \right) + k \left(\beta^2 - \frac{\omega^2}{k^2} \right) \text{Im} \left(\frac{1}{-k^2 + \tilde{\varepsilon}(\omega, k)\omega^2} \right) \right\}, \quad (4)$$

where $k = |\mathbf{k}|$ is the mCP three-momentum transfer to the electrons, ω is the mCP energy transfer to the electrons, ε is the mCP charge, $\beta = p_\chi/E_\chi$ is the mCP three-velocity in the detector's rest frame, n_e is the number of electrons per unit volume, α is the fine structure constant, and $\tilde{\varepsilon}(\omega, k)$ is the dielectric function of the medium. For highly-boosted mCPs, as the ones that could be generated in proton-proton collisions at the LHC, $\beta \simeq 1$. To calculate the probability of ionizing a given number of electrons in our sensor, we convolve the cross section in Eq. (4) with the probability of a recoiling electron to produce ionized electrons, as modeled in Ref. [25]; this procedure yields the detection cross section, σ_{det} . The expected number of detected events from mCP interactions is computed as

$$N(\epsilon, m_\chi) = \epsilon A \int \phi(\epsilon, E_\chi, m_\chi) [1 - e^{-dn_e \sigma_{\text{det}}}] dE_\chi, \quad (5)$$

where A and d are the detector's area and length, respectively, and ϵ is the effective exposure for the energy range considered [7,26].

Using Eq. (5), we compute $N(\epsilon, m_\chi)$ for the $4\text{--}7\ e^-$ energy range, considering ϵ as the mean of the effective exposures of the corresponding energy bins in Table V. The selected energy range corresponds to the region where interactions with bulk plasmons are expected to enhance the signal [24] and has been previously used for skipper-CCD mCP searches [7,8,26]. The 95% CL upper limit on the number of signal events in MOSKITA within this energy range, for the proton-proton collision period, computed as previously discussed, is 7.8 events, and the corresponding p -value for testing the null hypothesis, $p(\mu = 0)$, is 0.096. The corresponding 95% CL exclusion region in the millicharge-mass parameter space is defined by the condition $N(\epsilon, m_\chi) \geq 7.8$ events, and is shown in Fig. 18. Since we do not expect this limit to be competitive and it is presented mainly to illustrate skipper-CCDs' sensitivity reach, no statistical treatment of systematic uncertainties was included. However, as shown in Refs. [7,8], the impact of systematic uncertainties on the final limit is strongly reduced due to the ϵ^4 scaling of the expected event rate.




FIG. 18. 95% CL exclusion region on mCPs from MOSKITA using data collected in 2024 during the LHC proton-proton collisions (pink). Projected 95% CL exclusion regions for: (1) MOSKITA collecting 180 fb^{-1} with 3 times less background (cyan); and (2) a 200 g skipper-CCD detector, with $A = 231\text{ cm}^2$ and $d = 0.6\text{ cm}$, collecting 3000 fb^{-1} with 10 times less background (red). For comparison, constraints from previous mCP searches [7,9,17,27–30] (gray) are shown, which include SENSEI [7], using the same sensor technology, and the most recent results from milliQan [9], running at the same site. Also, the projected 95% CL exclusion region for the milliQan full bar detector collecting 3000 fb^{-1} during the HL-LHC era [16] (green) is shown.

We also show in this figure the projected exclusion regions for: (1) MOSKITA collecting 180 fb^{-1} with 3 times less background, and (2) a 200 g skipper-CCD detector, with $A = 231\text{ cm}^2$ and $d = 0.6\text{ cm}$, collecting 3000 fb^{-1} with 10 times less background. Such an increase in mass is plausible as larger, i.e., $\mathcal{O}(100\text{ g})$, skipper-CCD detectors operating with subelectron resolution have already been demonstrated [31]. Similarly, such a background reduction could be achievable by deploying a multilayered passive shield targeted for external low-energy backgrounds, particularly low-energy neutrons, as those used in skipper-CCD dark matter and reactor experiments [19,32–35].

VIII. CONCLUSIONS

We successfully installed and operated MOSKITA $\sim 33\text{ m}$ from the CMS collision point, the first skipper-CCD detector at the LHC sensitive to low-energy ionization signals from high-energy collisions. Data were acquired throughout 2024, during Run 3 of the LHC, collecting a total luminosity of 113.3 fb^{-1} during the period of proton-proton collisions and 1.54 nb^{-1} during the period of Pb-Pb collisions. Details on the detector installation, performance, data acquisition, and analysis pipeline are provided.

The event rate observed at MOSKITA is reported in two energy regions, $\leq 20\text{ e}^-$ and $> 20\text{ e}^-$, for both collision periods. In the low-energy region ($\leq 20\text{ e}^-$), we perform a likelihood analysis to test for correlations between the event rate and the LHC luminosity. No correlation was found during the period of proton-proton collisions. In contrast, during the Pb-Pb collision period, we observe a significant excess of 7 e^- events correlated with beam activity ($p = 0.017$), which may still be consistent with a statistical fluctuation.

Using the computed 95% CL upper limit on the number of $4\text{--}7\text{ e}^-$ signal events in MOSKITA for the proton-proton collision period, and comparing it to the expected number of detected events from mCP interactions, we constrain the millicharge-mass parameter space. While these constraints are not competitive due to the low mass ($\sim 2.2\text{ g}$) of the MOSKITA sensor, we project that the sensitivity reach of a more massive ($\sim 200\text{ g}$) skipper-CCD detector for the HL-LHC era would be complementary to those from similar efforts [10–13,16]. It is worth noting that even though these constraints were derived under the mCP model, MOSKITA results could be used to probe other new physics scenarios.

For the high-energy region ($> 20\text{ e}^-$), we find that the measured event rates for images with and without luminosity are inconsistent below 450 keV , for both collision periods. In particular, we observe a higher event rate after the Pb-Pb collisions occurred, coinciding with a rise in the single-electron rate. When excluding images taken after the Pb-Pb collisions, the high-energy event rates for images with and without luminosity become consistent. This increase may be attributed to secondary neutrons produced

during the Pb-Pb collisions. However, its origin remains to be investigated in future work.

Overall, the results presented here demonstrate the potential of the skipper-CCD technology to explore new physics at colliders. This work establishes a path forward for future skipper-CCD efforts at high-energy facilities. In particular at the LHC, extending data collection and/or developing a larger skipper-CCD detector will be key to improving sensitivity to rare, low-energy signals.

ACKNOWLEDGMENTS

Authors would like to thank and congratulate the CERN accelerator departments, including engineering and administrative personnel, for the excellent performance of the LHC during Run 3 and for their assistance with the MOSKITA detector installation. We gratefully acknowledge people from the CMS and milliQan Collaborations for their support, for sharing space in the drainage gallery, and

for the useful discussions held prior to the installation. Finally, authors would also like to thank Gustavo Otero y Garzon and Ricardo Piegaia for the initial talks regarding the possibility of operating a skipper-CCD detector at the LHC complex. This work was done using the resources of the Fermi National Accelerator Laboratory (Fermilab), a U.S. Department of Energy, Office of Science, Office of High Energy Physics HEP User Facility. Fermilab is managed by FermiForward Discovery Group, LLC, acting under Contract No. 89243024CSC000002.

DATA AVAILABILITY

The data that support the findings of this article are not publicly available upon publication because it is not technically feasible and/or the cost of preparing, depositing, and hosting the data would be prohibitive within the terms of this research project. The data are available from the authors upon reasonable request.

-
- [1] J. Kopp, P. A. N. Machado, M. Maltoni, and T. Schwetz, *J. High Energy Phys.* **05** (2013) 050.
 - [2] S. Schramm, *Symmetry* **14**, 1173 (2022).
 - [3] S. Biswas, E. Gabrielli, and B. Mele, *Symmetry* **14**, 1522 (2022).
 - [4] B. Holdom, *Phys. Lett.* **166B**, 196 (1986).
 - [5] J. Tiffenberg, M. Sofo-Haro, A. Drlica-Wagner, R. Essig, Y. Guardincerri, S. Holland, T. Volansky, and T.-T. Yu, *Phys. Rev. Lett.* **119**, 131802 (2017).
 - [6] L. Barak, I. M. Bloch, M. Cababie, G. Cencelo, L. Chaplinsky, F. Chierchie, M. Crisler, A. Drlica-Wagner, R. Essig, J. Estrada *et al.* (SENSEI Collaboration), *Phys. Rev. Lett.* **125**, 171802 (2020).
 - [7] L. Barak, I. M. Bloch, A. M. Botti, M. Cababie, G. Cencelo, B. A. Cervantes-Vergara, L. Chaplinsky, M. Crisler, A. Drlica-Wagner, R. Essig *et al.* (SENSEI Collaboration), *Phys. Rev. Lett.* **133**, 071801 (2024).
 - [8] A. A. Aguilar-Arevalo, J. C. D’Olivo, Y. Sarkis, N. Avalos, X. Bertou, P. Bellino, C. Bonifazi, A. Botti, G. Cencelo, J. Estrada *et al.* (CONNIE and Atucha-II Collaborations), *Phys. Rev. Lett.* **134**, 071801 (2025).
 - [9] S. Alcott, Z. Bhatti, J. Brooke, C. Campagnari, M. Carrigan, M. Citron, R. D. Santos, A. De Roeck, C. Dorofeev, T. Du *et al.*, *Phys. Rev. Lett.* **135**, 121802 (2025).
 - [10] M. Kalliokoski, V. A. Mitsou, M. de Montigny, A. Mukhopadhyay, P.-P. A. Ouimet, J. Pinfold, A. Shaa, and M. Staelens, *J. High Energy Phys.* **04** (2024) 137.
 - [11] R. Mammen Abraham, X. Ai, S. Alonso-Monsalve, J. Anders, C. Antel, A. Ariga, T. Ariga, J. Atkinson, F. U. Bernlochner, T. Boeckh, J. Boyd *et al.* (FASER Collaboration), [arXiv:2503.19775](https://arxiv.org/abs/2503.19775).
 - [12] M. Citron, F. Golf, K. Gunthoti, A. Haas, C. S. Hill, D. Imani, S. Kelly, M. Liu, S. Lowette, A. De Roeck *et al.*, [arXiv:2504.12973](https://arxiv.org/abs/2504.12973).
 - [13] F. Kling, J.-L. Kuo, S. Trojanowski, and Y.-D. Tsai, *Nucl. Phys.* **B987**, 116103 (2023).
 - [14] B. A. Cervantes-Vergara, S. Perez, J. Estrada, A. Botti, C. R. Chavez, F. Chierchie, N. Saffold, A. Aguilar-Arevalo, F. Alcalde-Bessia *et al.*, *J. Instrum.* **18**, P08016 (2023).
 - [15] B. Cervantes-Vergara, S. Perez, J. D’Olivo, J. Estrada, D. Grimm, S. Holland, M. Sofo-Haro, and W. Wong, *Nucl. Instrum. Methods Phys. Res., Sect. A* **1046**, 167681 (2023).
 - [16] A. Ball, J. Brooke, C. Campagnari, M. Carrigan, M. Citron, A. De Roeck, M. Ezeldine, B. Francis, M. Gastal, M. Ghimire *et al.*, *Phys. Rev. D* **104**, 032002 (2021).
 - [17] A. Ball, G. Beauregard, J. Brooke, C. Campagnari, M. Carrigan, M. Citron, J. De La Haye, A. De Roeck, Y. Elskens *et al.*, *Phys. Rev. D* **102**, 032002 (2020).
 - [18] S. Chatrchyan, G. Hmayakyan, V. Khachatryan, A. M. Sirunyan, W. Adam, T. Bauer, T. Bergauer, H. Bergauer, M. Dragicevic, J. Erö *et al.* (CMS Collaboration), *J. Instrum.* **3**, S08004 (2008).
 - [19] P. Adari, I. M. Bloch, A. M. Botti, M. Cababie, G. Cencelo, B. A. Cervantes-Vergara, M. Crisler, M. Daal, A. Desai, A. Drlica-Wagner *et al.* (SENSEI Collaboration), *Phys. Rev. Lett.* **134**, 011804 (2025).
 - [20] G. I. Cencelo, C. Chavez, F. Chierchie, J. Estrada, G. Fernandez-Moroni, E. E. Paolini, M. S. Haro, A. Soto, L. Stefanazzi, J. Tiffenberg *et al.*, *J. Astron. Telesc. Instrum. Syst.* **7**, 015001 (2021).
 - [21] A. J. Rádl, *Proc. Sci. LHCP2023* (2024) 283 [[arXiv:2401.10763](https://arxiv.org/abs/2401.10763)].

- [22] G. Cowan, K. Cranmer, E. Gross, and O. Vitells, *Eur. Phys. J. C* **71**, 1554 (2011).
- [23] G. Beauregard, PYTHIA mCP simulation, <https://github.com/GBeauregard/milligan-pyth-sim> (2018).
- [24] R. Essig, R. Plestid, and A. Singal, *Commun. Phys.* **7**, 416 (2024).
- [25] K. Ramanathan and N. Kurinsky, *Phys. Rev. D* **102**, 063026 (2020).
- [26] S. Perez, D. Rodrigues, J. Estrada, R. Harnik, Z. Liu, B. A. Cervantes-Vergara, J. C. D’Olivo, R. D. Plestid, J. Tiffenberg, T.-T. Yu *et al.*, *J. High Energy Phys.* **02** (2024) 072.
- [27] A. A. Prinz, R. Baggs, J. Ballam, S. Ecklund, C. Fertig, J. A. Jaros, K. Kase, A. Kulikov, W. G. J. Langeveld, R. Leonard *et al.*, *Phys. Rev. Lett.* **81**, 1175 (1998).
- [28] R. Acciarri, C. Adams, J. Asaadi, B. Baller, T. Bolton, C. Bromberg, F. Cavanna, D. Edmunds, R. S. Fitzpatrick, B. Fleming *et al.* (ArgoNeuT Collaboration), *Phys. Rev. Lett.* **124**, 131801 (2020).
- [29] S. Davidson, S. Hannestad, and G. Raffelt, *J. High Energy Phys.* **05** (2000) 003.
- [30] G. Magill, R. Plestid, M. Pospelov, and Y.-D. Tsai, *Phys. Rev. Lett.* **122**, 071801 (2019).
- [31] F. Chierchie, C. Chavez, M. Sofo Haro, G. Fernandez Moroni, B. Cervantes-Vergara, S. Perez, J. Estrada, J. Tiffenberg, S. Uemura, and A. Botti, *J. Instrum.* **18**, P01040 (2023).
- [32] A. Aguilar-Arevalo, I. Arnquist, N. Avalos, L. Barak, D. Baxter, X. Bertou, I. M. Bloch, A. M. Botti, M. Cababie, G. Cancelo *et al.* (DAMIC, DAMIC-M and SENSEI Collaborations), *Phys. Rev. D* **109**, 062007 (2024).
- [33] I. Arnquist, N. Avalos, D. Baxter, X. Bertou, N. Castelló-Mor, A. E. Chavarria, J. Cuevas-Zepeda, J. C. Gutiérrez, J. Duarte-Campderros, A. Dastgheibi-Fard *et al.* (DAMIC-M Collaboration), *Phys. Rev. Lett.* **130**, 171003 (2023).
- [34] A. A. Aguilar-Arevalo, N. Avalos, X. Bertou, C. Bonifazi, G. Cancelo, B. A. Cervantes-Vergara, C. Chavez, F. Chierchie, G. Coelho Corrêa, J. C. D’Olivo *et al.*, [arXiv:2403.15976](https://arxiv.org/abs/2403.15976).
- [35] E. Depaoli, D. Rodrigues, I. Sidelnik, P. Bellino, A. Botti, D. Delgado, M. Cababié, F. Chierchie, J. Estrada, G. F. Moroni, S. Perez, and J. Tiffenberg, *J. High Energy Phys.* **10** (2024) 155.