Search for low-mass WIMPs in a 0.6 kg day exposure of the DAMIC experiment at SNOLAB

A.E. Chavarria,* M. Settimo,† and J. Tiffenberg‡
(DAMIC collaboration)
(Dated: July 14, 2016)

We present results of a dark matter search performed with 0.6 kg-days of data acquired with charge-coupled devices (CCDs) by the DAMIC Collaboration in the SNOLAB underground laboratory. A spectral analysis is performed for energy depositions in the bulk silicon of the detectors down to an ionization signal of 60 eV electron-equivalent. The observed data are consistent with backgrounds from environmental radioactivity, allowing us to place constraints on the WIMP-nucleon scattering cross-section. The low energy threshold and the relatively low mass of the silicon nucleus make the result particularly competitive in the low-mass WIMP region ($<10\,\mathrm{GeV}/c^2$) considering the limited experimental exposure. Cross-sections relevant to the potential signal from the CDMS-II Si experiment are directly excluded for the first time.

* alvaro@kicp.uchicago.edu
† mariangela.settimo@lpnhe.in2p3.fr
‡ javiert@fnal.gov

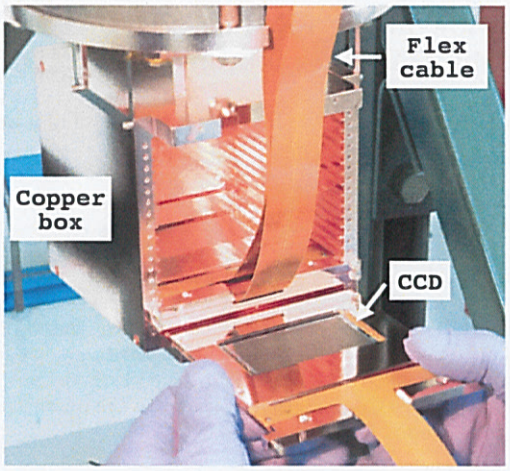

FIG. 1. Copper module holding an 8 Mpix CCD being installed in the low radioactivity copper box. Two other modules have already been installed and can be partially seen at the bottom of the box. The flex cables that carry the CCD signals are also visible.

I. INTRODUCTION

too soon discuss fish the cover set-up.

The DAMIC (Dark Matter in CCDs) experiment [?] employs the bulk silicon of scientific-grade charge-coupled devices (CCDs) to detect coherent elastic scattering of Weakly-Interacting Massive Particles (WIMPs) — putative yet-to-be-discovered particles which may explain the dark matter in the Universe [? ? ?]. By virtue of the low readout noise of the CCDs and the relatively low mass of the silicon nucleus, DAMIC is particularly sensitive to low mass ($<20\,\mathrm{GeV}/c^2$) WIMPs, which induce nuclear recoils of keV-scale energies.

or 10 ?

The upcoming DAMIC100 experiment will consist of eighteen 16 Mpix, 675 μm-thick CCDs (5.8 g each) installed in the SNOLAB laboratory. Most of the infrastructure for the experiment has been in SNOLAB since 2013. Significant improvements have been made to the shielding and inner components of the detector to decrease the radioactivity of the environment that hosts the devices. Throughout the year 2015 extensive testing was performed on the surfaces and materials immediately surrounding the CCDs. For these tests 8 Mpix CCDs were deployed, which, except for their smaller area, are identical to those planned for DAMIC100.

Dark matter search data was acquired with these test devices in designated periods to monitor their performance and to confirm the continuous capability of the detector to acquire good quality data. In this paper, we present results from a 0.6 kg·d exposure, which demonstrate for the first time the sensitivity of a silicon target to WIMPs with masses $<6\,\mathrm{GeV}/c^2$ and directly probe the parameter space corresponding to the potential WIMP signal in the CDMS II Silicon experiment [?].

II. DAMIC EXPERIMENT AT SNOLAB

block

The DAMIC CCDs are packaged in a copper module, including a silicon support frame for the CCD and a low-radioactivity flex cable to carry the signals that drive and read the device (FIG. 1). The modules are inserted in slots of a copper box that is cooled to ~120 K inside a copper vacuum vessel ($\sim 10^{-6}$ mbar). The box is shielded on all sides by lead to attenuate external γ-rays. A 18 cm-thick shield hangs immediately above the box inside the vacuum vessel and a lead ~~castle~~ of 21 cm thickness shields the copper vessel from all other sides. The innermost inch of lead comes from an ancient Spanish galleon and has negligible ^{210}Pb content, strongly suppressing the background from bremsstrahlung γs produced by ^{210}Bi decays in the outer lead shield. A 42 cm-thick polyethylene shield is used to moderate and absorb environmental neutrons. The overburden of the laboratory site (6010 m water equivalent) suppresses backgrounds from cosmic rays to a negligible level. Details of the DAMIC infrastructure at SNOLAB can be found in [?].

I would give an estimation of the level

The DAMIC CCDs were developed at Lawrence Berkeley National Laboratory MicroSystems Lab [?], starting from an existing design for the Dark Energy Survey (DES) camera (DECam) [?]. They feature a three-phase polysilicon gate structure with a buried p-channel. The pixel size is 15 μm × 15 μm and the bulk of the detector is high-resistivity (10–20 kΩ cm) n-type silicon with a thickness of 675 μm. The high-resistivity of the silicon allows for a low donor

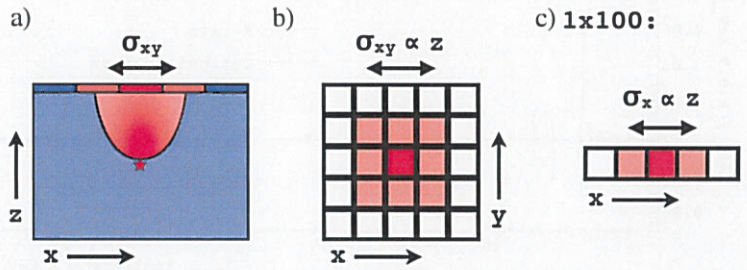

FIG. 2. a) Cross-sectional representation of the charge produced by a point-like ionization event (star) in the CCD bulk as it is drifted to the pixel array. b) The x-y distribution of charge on the pixel array following the ionization event. Due to diffusion the charge is collected in multiple pixels, with the lateral spread (σ_{xy}) being positively correlated with the depth (z coordinate) of the interaction. When the CCD is read out in the 1×1 configuration this is the pattern observed in the image. c) In the 1×100 mode, the CCD is read out in column segments one hundred pixels tall, collapsing the pixel contents along the y axis, leading to a one dimensional pattern with the charge spread out over fewer pixels. The one-dimensional lateral spread (σ_x) is positively correlated to the depth of the interaction.

density in the substrate ($\sim 10^{11}\,\mathrm{cm}^{-3}$), which leads to fully depleted operation at a substrate bias of 40 V. Ionization produced in the bulk will be drifted along the direction of the electric field (z-axis). The holes (charge carriers) will be collected and held near the p-n junction, less than 1 μm below the gates. Due to thermal motion, the ionized charge will diffuse symmetrically as it is drifted, with a spatial variance ($\sigma_x^2{=}\sigma_y^2{=}\sigma_{xy}^2$) that is proportional to the carrier transit time. Hence, there is a positive correlation between the lateral diffusion (σ_{xy}) of the collected charge and the depth of the interaction (z). FIG. 2(a) depicts a cross-sectional view of a point-like interaction in the CCD bulk, with the ionized charge being drifted to the pixel array. Due to charge diffusion, the ionized charge is collected by multiple pixels.

III. CCD READOUT

The ionized charge is held at the gates until the device is read out. During readout, the charge is transferred vertically from pixel to pixel along each column by appropriate clocking of the 3-phase gates ("parallel clocks"), while higher frequency clocks ("serial clocks") move the charge of the last row (the "serial register") horizontally to the CCD's output node, where the charge is measured by a correlated double-sampling circuit []. The inefficiency of charge transfer from pixel to pixel is as low as 10^{-6} and the readout noise for the charge collected in a pixel is $\sim 2\,e^-$ (Sec. V). The image is then reconstructed from the order in which the pixels were read out, and contains a two-dimensional stacked history (projected on the x-y plane) of all particle interactions throughout the exposure. FIG. 2(b) shows the pattern observed on the x-y plane from an ionization event in the CCD bulk. For rare-event searches, it is advantageous to take the longest possible exposures to minimize the number of readouts, and thus the number of pixels above a given threshold due to noise fluctuations. The maximum length of the exposure is limited by the CCD's dark current ($<10^{-3}\,e^-\mathrm{pix}^{-1}\mathrm{day}^{-1}$ at the operating temperature of $\sim$120 K) which starts contributing to the noise when it is the source of $\sim 0.5\,e^-$ on average to the charge measured.

The CCD can be clocked to place the charge collected by multiple adjacent pixels into the output node before the charge measurement is performed, effectively "hardware binning" the CCD. This procedure has an important impact on the response of the detector. For ionization events which produce charge that is collected in multiple pixels, fewer measurements will be performed to read out the ionized number of charge carriers, leading to a smaller contribution of the readout noise to the uncertainty in the charge measurement. This improves both the energy resolution and the energy threshold for ionization events distributed over multiple pixels. However, the distribution of the charge in the individual pixels is lost, worsening the measurement of the x, y and σ_{xy} of the event.

DAMIC data has been acquired so far with two different readout configurations: 1×1 and 1×100. The first configuration is the standard CCD readout, where the charge collected by each pixel is read out individually, offering maximum spatial resolution. In the latter configuration, one hundred CCD rows are transferred into the serial register before the charge is clocked horizontally and each column segment is read out individually. In this case, the spatial resolution in the y coordinate is lost and the events are collapsed into the x dimension, where σ_x is still positively correlated to the depth of the interaction (FIG. 2(c)). This configuration is a compromise between improving the sensitivity to ionization events that occur deep in the bulk, where lateral diffusion is significant, and retaining some

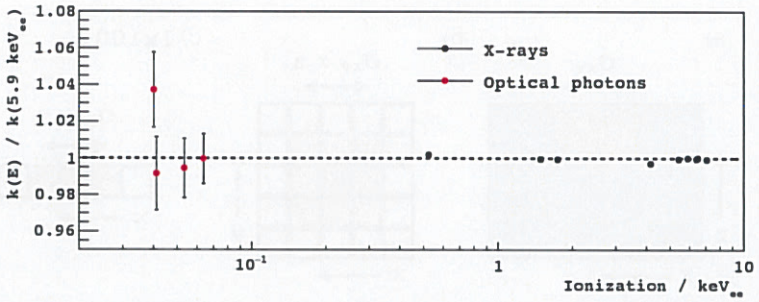

FIG. 3. Linear constant, k, relating the CCD output signal to the ionization generated in the substrate. Values are given relative to k measured at 5.9 keV$_{ee}$. Calibrations at high energies are performed with X-rays, while the lowest energy points are obtained using a optical photons, as outlined in the text. The linearity of the CCD energy response is demonstrated down to 40 eV$_{ee}$.

information on the depth of the interaction. Furthermore, as the number of charge measurements performed in the 1×1 configuration are many more, the read out time is much longer. DAMIC CCDs are read with an integration time for the correlated double-sampling of 40 μs, which leads to an image readout time of 840 s (20 s) for the 1×1 (1×100) mode.

DAMIC CCDs feature an output node at each end of the serial register. As described above, all the charge collected by the CCD pixel array is read out through one of these output nodes. The second output node, into which charge is never deposited, remains active and is also read, offering a simultaneous measurement of zero charge. This measurement is representative of the CCD noise precisely at the time of the charge measurement and allows to identify and suppress correlated electronic noise of the detector's readout chain.

IV. ENERGY AND DEPTH RESPONSE OF A CCD

The output of a CCD readout chain is an ADC value that is proportional to the number of charge carriers placed in the CCD's output node. The signal produced by electrons, which lose their energy through ionization, is also proportional to the generated number of charge carriers, with an average of one electron-hole pair produced for every 3.77 eV of deposited energy [?]. Thus, we define the electron-equivalent energy scale (in units of eV$_{ee}$) relative to the ionization produced by electrons from the photo-absorption of X-rays of known energy.

Calibrations have been performed by illuminating the CCD from the front with fluorescence X-rays from O, Al, Si, Cr, Mn and Fe. FIG. 3 summarizes the measurement of the linear calibration constant, k (ADC value / eV$_{ee}$), at different energies, which demonstrates the linear response of the CCD to electrons. From X-ray data we can also estimate the intrinsic energy response of the device. The measured resolution of 54 eV$_{ee}$ at 5.9 keV$_{ee}$ corresponds to a Fano factor [? ?] of 0.133±0.005.

To demonstrate the linearity of the CCD output to weaker signals we use optical photons, which produce a single electron-hole pair by photoelectric absorption. A light-emitting diode (LED) is installed inside the DAMIC copper vessel. The LED is turned on for a determined time interval T∼20 s throughout an image exposure. The LED is then turned off and the image is read out. The procedure is repeated for nine consecutive exposures. For a given pixel, the number of charge carriers detected in the nine images follows a Poisson distribution. The mean (μ_l) and variance (σ_l^2) of the increase in the pixel ADC values in response to the LED are related to the calibration constant (k) from the properties of the Poisson distribution as

$$k = \frac{1}{3.77\,\mathrm{eV_{ee}}}\frac{\mu_l}{\sigma_l^2}. \qquad (1)$$

Thus, from the mean and variance of the pixel values over the nine exposures it is possible to estimate k with very high statistics even at very low light levels, when only a handful of charge carriers are expected to be collected by a pixel. These results have been incorporated in FIG. 3 and demonstrate the CCD response is linear within 5% down to 40 eV$_{ee}$.

The energy depositions expected from WIMP interactions arise from recoiling silicon nuclei in the target. We rely on recent data of the response of a CCD to nuclear recoils with energies as low as 0.75 keV [?] to relate the electron-equivalent energy to the deposited energy by a nuclear recoil (in units of eV$_{nr}$). To estimate the ionization

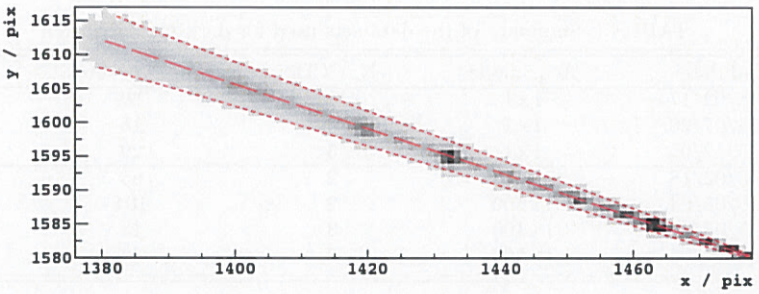

FIG. 4. Cosmic muon observed in cosmic background data acquired on the surface. Only pixels whose values are $>4\,\sigma_{\rm pix}$ are colored. Darker colors represent pixels with more collected charge. The large area of diffusion on the top left corner of the image is where the muon crosses the back of the CCD. Conversely, the narrow end on the bottom right corner is where the muon crosses the front of the device. The reconstructed track is shown by the long-dashed line. The short-dashed lines show the $3\,\sigma$ band of the charge distribution according to the best-fit diffusion model.

signal below the lowest calibration point at $60\,{\rm eV_{ee}}$ we perform a linear extrapolation, which leads to no ionization signal for nuclear recoils with energies below a cutoff of $0.38\pm0.12\,{\rm keV_{nr}}$.

The relationship between σ_{xy} and z of an interaction can be solved in one dimension from first principles given the electric field profile within the CCD substrate and the fact that the lateral variance of the carriers (σ_{xy}^2) due to diffusion is proportional to the transit time from the interaction point to the CCD pixel array [?]. The resulting relation is

$$\sigma_{xy}^2 = -A\ln|1-bz|. \tag{2}$$

The constants A and b are related to the physical properties and the operating parameters of the device and are given by

use \displaystyle n latex equation, looks better

$$A = \frac{\epsilon}{\rho_n}\frac{2k_BT}{e} \tag{3}$$

$$b = \left(\frac{\epsilon}{\rho_n}\frac{V_b}{z_D} + \frac{z_D}{2}\right)^{-1}, \tag{4}$$

where ϵ is the permittivity of silicon, ρ_n is the ~~donor charge~~ density in the substrate, k_B is Boltzmann's constant, T is the operating temperature, e is the electron's charge, V_b is the applied bias across the substrate and z_D is the thickness of the device.

In practice it is most accurate to measure the parameters A and b directly from data. We do this from cosmic background data acquired on the surface, by fitting the width of muon tracks that pierce the CCD as a function of depth. As the muon tracks follow a straight line, the depth can be calculated unambiguously from the path length on the x-y plane. FIG. 4 shows one of the muons used for calibration, where the CCD is operated at ~~same~~ nominal temperature and substrate bias ~~as those~~ at SNOLAB. The best-fit parameters to the diffusion model (Eq. 2) are A=215 μm^2 and b=1.3×10^{-3} μm^{-1}, which correspond to a maximum diffusion at the back of the device of $\sigma_{\rm max}$=21 μm=1.4 pix. The accuracy of this calibration has been validated by studying the diffusion of X-ray events that interact near the surfaces on the back and the front of the CCD [?] and from γ-ray data, which provide the σ_{xy} distribution of ionization events uniformly distributed in the bulk of the device.

Data acquired with an external ^{3}H source demonstrates that the CCD has a $\sim$2 μm dead layer on its front and back surfaces. However, there is no evidence for regions of partial or incomplete charge collection that may hinder the energy response of the device.

V. DATA SETS AND IMAGE PROCESSING

The DAMIC test setup at SNOLAB was devoted to background studies throughout the years 2013–2015, accounting to more than ten operations where the organization of the CCDs, the external detector shielding and the shielding material samples placed inside the copper box for background studies were modified. Throughout 2015, dark matter search data was acquired intermittently in both 1×1 and 1×100 acquisition modes with 8 Mpix, 675 μm-thick CCDs (2.9 g each). Table I summarizes the data runs considered for the analysis including the number of CCDs and images, and the total exposure selected for the analysis after the mask and image selection procedures discussed below.

TABLE I. Summary of the data sets used for dark matter search

Start date	End date	Acq. mode	N. CCDs	N. exposures	Total exposure / (kg·d)
2014/12/12	2015/02/17	1×1	2	225	0.235
2015/07/06	2015/07/20	1×1	3	18	0.056
2015/10/28	2015/12/05	1×1	3	29	0.091
2015/02/01	2015/02/18	1×100	2	65	0.040
2015/04/21	2015/05/04	1×100	2	104	0.065
2015/07/06	2015/07/20	1×100	3	18	0.017
2015/10/28	2015/12/05	1×100	2	44	0.082

All CCDs were individually calibrated with X-rays and cosmic backgrounds at Fermi National Accelerator Laboratory before deployment. The fluorescence copper line (8 keV) observed at SNOLAB due to radioactive particle interactions in the surrounding copper were used to confirm and, in one case correct, the energy scale of the CCDs. The observed value for $\sigma_{\rm max}$ was also monitored to validate the depth response calibrated on the surface. The radioactive background rate measured below 10 $\rm keV_{ee}$ decreased with time due to the continuos improvements in the radio-purity of the setup, with an average event rate of $29\pm2\,\rm keV_{ee}^{-1}\,kg^{-1}\,d^{-1}$ ($23\pm2\,\rm keV_{ee}^{-1}\,kg^{-1}\,d^{-1}$) for the data acquired in 1×1 (1×100) mode.

Images were taken with exposures of either 10^4 or 3×10^4 seconds, immediately followed by the acquisition of an image 'blank' whose exposure is only the readout time. Due to the small probability of a physical event occurring during readout, the blanks contain only the image noise.

may want of give an estimate

The first step in image processing is the subtraction of the pixel pedestal. This is done column-wise by evaluating the medians of the pixel values along every column and then subtracting the median from all pixel values in the column. To mitigate the effect of correlated noise, a correction is applied based on the simultaneous readings of the serial register by the two output nodes: the right-side output node that measures the charge and the left-side output node that measures only the noise (Sec. III). Since correlated noise produces a shift in the pedestal of both amplifiers, the linear relation between the left-side and right-side signals along a row is used to correct the measured pixel value.

pas claire

For each data run (Table I) we calculate the median and root mean square (RMS) of every pixel over all images in the run. These pixel variables are used to construct a "mask," which identifies pixels excluded from the analysis for having either a value which deviates more than 2 RMS from the median in at least 50% of the images or a median or RMS that is an outlier when compared to the distributions of these variables for all pixels.

Figure 5 shows an example of the distribution of pixel values after pedestal and correlated noise subtraction for a single 30 ks exposure (black) compared to its corresponding blank (blue). The values of pixels that have collected no charge ~~belong to the distribution centered at zero~~ arising from white noise, which is almost identical ~~for the 30 ks exposure~~ and its blank. A Gaussian fit to the data (red) gives the uncertainty in the individual pixel values ($\sigma_{\rm pix}$) of $1.9\,e^-$ ($\sim7\,\rm eV_{ee}$). Ionization events lead to pixel values which are significantly above the noise, constituting a tail on the positive side of the noise distribution and are only visible in the 30 ks exposure case.

bad choice of word

The consistency between each image and its blank is checked by comparing the distribution of the test statistic, ΔLL (Sec. VI), ~~of clusters arising from readout noise.~~ Images for which there is a significant discrepancy in these distributions or the clustered noise in the blank does not follow the distribution expected from white noise are excluded from the analysis. This includes some CCDs in runs acquired between February and August, 2015 where the pixel noise was relatively high ($\sim2.2\,e^-$). During this period the polyethylene shield was partially open and the CCDs installed at the top of the box were exposed to optical photons that leaked into vessel, leading to a significant amount of background charge.

?

the

VI. EVENT RECONSTRUCTION AND SELECTION

An initial search for ionization events is done with an algorithm that looks for clusters of pixels with signals larger than $4\,\sigma_{\rm pix}$. All contiguous pixels that satisfy this condition are considered part of the cluster. Variables are then computed for the cluster, which include the mean x-y position, the total number of pixels and the total energy, estimated by the addition of the values of all pixels in the cluster.

We limit this analysis to clusters with energies <10 $\rm keV_{ee}$, for which the track length of the ionizing particle is much smaller than the pixel size and thus, the energy deposition may be considered point-like. Before proceeding with the "likelihood clustering" described below, we mask in every image every cluster found in the initial search whose energy is >10 $\rm keV_{ee}$. In addition, pixels that are less than 4 pixels away from the cluster or less than 50 (200) pixels to the left of the cluster are masked in the 1×1 (1×100) data set. The latter condition is meant to prevent low energy

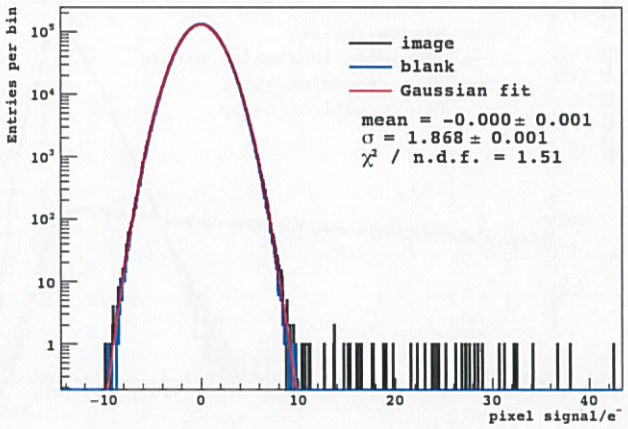

FIG. 5. Example of the pixel value distribution after image processing in one 30 ks exposure (black) and its corresponding blank (blue) acquired in December 2014. The noise in the image (uncertainty on a single pixel value) is fitted to $\sigma_{\rm pix}$=1.9 e^-.

clusters arising from stray charge due to CCD charge transfer inefficiencies. The average fraction of masked pixels in an image, including those discarded by the criteria outlined in Sec. V, is 1% (8%) in the 1×1 (1×100) data set.

On the unmasked regions of every 1×1 image we apply a likelihood clustering algorithm based on a 11×11 moving window. For every position of the window we compute two quantities: i) $\mathcal{L}_n$: the likelihood that the pixel values are described by white noise and ii) $\mathcal{L}_G$: the likelihood that the pixel values are described by a two-dimensional Gaussian function on top of white noise,

$$f_G = kN_e(3.77\,\mathrm{eV_{ee}}) \times \mathrm{Gaus}\,(x, y, \mu_x, \mu_y, \sigma_x, \sigma_y)\,, \tag{5}$$

whose integral, $kN_e(3.77\,\mathrm{eV_{ee}})$, is the addition of pixel values in the window, μ_x and μ_y are the coordinates of the center of the window and its width σ_x=σ_y=σ_{xy} is 1 pix. The log-likelihood of their ratio $-\ln\frac{\mathcal{L}_G}{\mathcal{L}_n}$ is computed, and when it becomes sufficiently negative (i.e. there is a significant preference for the Gaussian hypothesis), a cluster has been identified. The search window is further moved around to find the local minimum of this quantity. At this position the window remains fixed and a fit is performed, leaving N_e, μ_x, μ_y and σ_{xy} as free parameters and maximizing the value of $\mathcal{L}_G$. From the fit result we take N_e as the best estimate of the number of ionized charge carriers, μ_x and μ_y as the best estimate of the x-y position of the ionization event and σ_{xy} as the best estimate of the lateral spread of the charge. From N_e and σ_{xy} we can then reconstruct the energy (E) and depth of the interaction (z) (Sec. IV). The test statistic,

$$\Delta LL = -\ln\left[\frac{\mathrm{Max}(\mathcal{L}_\mathrm{G})}{\mathcal{L}_n}\right], \tag{6}$$

is also registered and represents the significance that the cluster arose from an ionization event and not from white noise.

In the 1×100 acquisition mode the clustering procedure is very similar, except that it is performed in one dimension along rows of the image. The fitting function, f_G, is reduced to a one-dimensional Gaussian with μ_x and σ_x as free parameters. The interpretations of the best fit values are analogous.

FIG. 6 shows the ΔLL distribution of all clusters in the 1×1 data set (black) and their corresponding blanks (blue). The clustered noise has the same distribution in data images and blanks, with an exponentially decreasing tail at high ΔLL values as expected for white noise. To reject clusters from readout noise, a cut at ΔLL<-28 (-25) is applied to the 1×1 (1×100) data set. The expected leakage of noise clusters is less than 0.1, as estimated from an exponential fit to the tail of the ΔLL distribution.

From the spatial distributions of events passing the ΔLL cut we notice repeating clusters that occur at the same spatial position in the CCDs. These events arise from small defects in the silicon lattice that produce an increased level of dark current at a specific spatial position. As these events are very faint, they were missed by the masking criteria outlined in Sec. V. We removed them from the final candidates with a negligible impact on the acceptance, as the probability of two uncorrelated events occurring in the same pixel is extremely low. Likewise we excluded clusters that were less than 300 μm on the x-y plane from any other cluster in the same image. These are clusters that have a common origin and unlikely to arise from WIMP interactions. Their exclusion also has a negligible impact on the

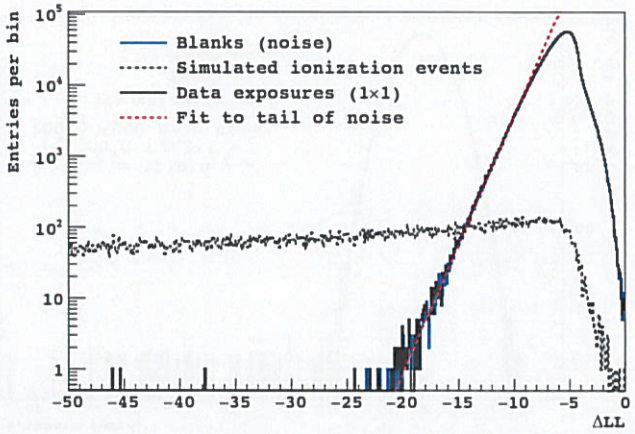

FIG. 6. ΔLL distributions for all clusters in the 1×1 data set. The blue line shows the distribution for clusters in the blanks, which are representative of the contribution from readout noise to the data set. The black-dashed line presents the expected distribution (from simulation) for ionization events that occur uniformly in the CCD bulk, assuming a constant (flat) energy distribution. The black line shows the distribution for all clusters in the 1×1 data set. The dashed red line is the fit done to the tail of the noise distribution to determine the cut used to reject readout noise. The fit is statistically consistent with the tail of the distribution. For ΔLL<-28, we expect 0.014 events from noise.

acceptance. After the application of these criteria 122 (62) final candidates clusters remain in the 1×1 (1×100) data sets. FIG. 7 shows the lateral spread versus energy distribution of the candidates.

To estimate the performance of the reconstruction algorithm for WIMP-like events we simulate the expected distribution of carriers on the pixel array from events of different energies with a uniform spatial distribution in the CCD bulk according to the diffusion model (Sec. IV). Two thousand (two hundred) simulated events were pasted on top of each of the acquired 1×1 (1×100) raw images using the corresponding calibration of pixel values to the number of charge carriers collected in the pixel. The full data processing chain was then run on each image, including the signal identification and likelihood clustering. FIG. 6 shows the ΔLL distribution of all clustered simulated events in the 1×1 data set (dashed black).

By analyzing the distributions of the reconstructed energy, we confirm that their means are within 1% of the simulated energy for events with energies greater than 100 $\mathrm{eV_{ee}}$. At lower energies there might be a slight overestimation in the reconstructed energy that is <5% at threshold. Likewise we estimate the resolution in the reconstruction of the ionization signal, arising from the uncertainty in the measurement of the number of charge carriers, to be σ_0=37 $\mathrm{eV_{ee}}$ (30 $\mathrm{eV_{ee}}$) in the 1×1 (1×100) data set. Thus, the energy response of the detector can be modeled with a resolution $\sigma^2_{\mathrm{res}}=\sigma^2_0+(3.77\,\mathrm{eV_{ee}})FE$, where F is the Fano factor.

From the fraction of simulated events that pass the event selection criteria we can also estimate the acceptance of the signal identification and of each of the data quality cuts. After all selection cuts described above, the acceptance for bulk events increases from 9% (25%) at 75 $\mathrm{eV_{ee}}$ (60 $\mathrm{eV_{ee}}$) to ~100% at 400 $\mathrm{eV_{ee}}$ (150 $\mathrm{eV_{ee}}$) in the 1×1 (1×100) data.

The better energy resolution and the higher acceptance of lower energy events in the 1×100 data set is expected due to the improved identification and reconstruction of events that originate deeper in the bulk of the device, which experience significant lateral charge diffusion.

VII. DISCRIMINATION OF SURFACE EVENTS

The selection criteria presented in Sec. VI have been defined to reliably distinguish events due to ionization produced by particle interactions from electronic noise and dark current. High-energy photons that Compton scatter in the bulk produce background with a uniform spatial distribution because the Compton scattering length is always much greater than the thickness of the CCD. Hence, ionization events from Compton scattering are only distinguishable from WIMP interactions spectrally. Nuclear recoils from WIMP interactions have a characteristic spectrum that decreases exponentially with increasing energy, while the Compton scattering spectrum is almost flat throughout the WIMP search energy region.

There is also a contribution to the selected clusters from low energy electrons and photons radiated by the surfaces facing the devices, and from electrons that start in the silicon and exit the device after only depositing a small fraction

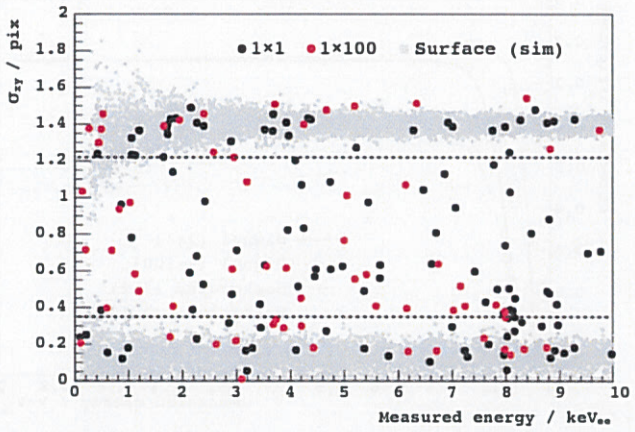

FIG. 7. Lateral spread (σ_{xy}) versus measured energy of the clusters that pass the selections cuts outlined in Sec. VI. Black (red) markers correspond to candidates in the 1×1 (1×100) data set. Gray markers show the expected distribution of energy deposits near the front and back surfaces of the device. The dashed lines represent the fiducialization cuts described in Sec. VII.

of their energy. These events occur tens of μm or less from the surface of the devices and can be rejected by a cut on σ_{xy}.

We select events with $0.35<\sigma_{xy}<1.22$, leading to 51 (28) candidates surviving in the 1×1 (1×100) data set. The dashed lines in FIG. 7 represent these fiducial cuts. The group of events at 8 $\mathrm{keV_{ee}}$ corresponds to Cu fluorescence X-rays produced by radioactive background interactions in the surrounding copper structure. Due to the relatively long X-ray absorption length at this energy (65 μm), some of the events leak into the fiducial region. We avoid this background by restricting the WIMP search to clusters with energies <7 $\mathrm{keV_{ee}}$.

Simulations allow us to estimate the signal acceptance and background leakage in the fiducial region. The gray markers in FIG. 7 show the σ_{xy} versus energy for one of these simulations, where the interactions were simulated to occur <15 μm from the front and back surfaces of the device in the 1×1 data.

The signal acceptance is estimated to be ~75%. The discrimination is >95% for surface electrons with energy depositions >1.5 $\mathrm{keV_{ee}}$ and for external photons with incident energies 1.5–4.5 $\mathrm{keV_{ee}}$. For higher energy photons the rejection decreases down to 85% at 6.5 $\mathrm{keV_{ee}}$ due to their longer absorption length. Below 1.5 $\mathrm{keV_{ee}}$ the σ_{xy} reconstruction worsens leading to significant leakage into the fiducial region for which we must account. To estimate the leakage we develop a model of the radioactive background in which we consider contributions from both bulk and surface events. From the σ_{xy} distribution of clusters with energies in the range 4.5–7.5 $\mathrm{keV_{ee}}$, where the expected contribution from a WIMP signal is smallest, we estimate the relative fractions of surface and bulk events in the background. We perform this estimate using all available data, including data acquired with a higher threshold devoted to background studies and excluded from the WIMP search, and consider the specific contributions of the particular CCDs to the WIMP search data exposure. This leads to a background composition that is 65±10% (60±10%) from the bulk, 15±5% (25±5%) from the front and 20±5% (15±5%) from the back in the 1×1 (1×100) data set. This composition is assumed to be constant with energy, which is justified by the fact that the background continuum of both bulk and surface events is expected and observed to be approximately constant in energy intervals the size of the WIMP search region.

FIG. 8 shows the final detection efficiency after fiducialization for events with spatial distributions as expected for the WIMP signal and for the background. The turn-on of the efficiency curves near threshold is due to the selection cuts to reject white noise described in Sec. VI. At high energies the signal acceptance is almost constant with a background leakage that is 65% (60%) of the signal in the 1×1 (1×100) data due to the contribution of Compton events to the background. Toward low energies there is an increase in the detection efficiency for background relative to the signal due to leakage of ionization events originating from the surface.

VIII. LIKELIHOOD ANALYSIS OF THE SPECTRUM

After event selection, 31 (23) final candidates remain in the fiducial region with energies <7 $\mathrm{keV_{ee}}$ in the 1×1 (1×100) data set. Each reconstructed candidate is characterized by its measured electron-equivalent energy, E_i. We

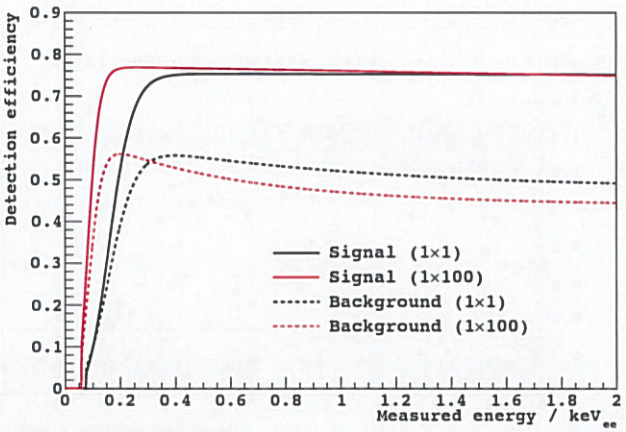

FIG. 8. Detection efficiency of events as a function of reconstructed energy for the 1×1 (black) and 1×100 (red) data sets after the selection and fiducial cuts outlined in Sec. VI and the fiducial cuts in Sec. VII. Solid lines present the acceptance of the WIMP signal while dashed lines present the leakage of background events considering both bulk and surface contributions.

use this observable to define an extended Likelihood function for the signal+background model:

$$\mathcal{L}_{s+b}(s,b,M|\vec{E}) = e^{-(s+b)} \times \Pi_{i=1}^{N}[sf_s(E_i|M) + bf_b(E_i)], \tag{7}$$

where s and b are the expected number of signal and background events in the fiducial region, $f_s(E|M)$ and $f_b(E|M)$ are the PDFs for the signal and background, and $N = 54$ is the total number of selected events. The PDF for the expected WIMP spectrum, $f_s(E|M)$, depends on the WIMP mass M, the standard halo parameters and the detector response (ionization efficiency, detection efficiency and energy resolution):

$$f_s(E|M) = C(\sigma_0)\epsilon_{\rm det}(E)\int \frac{\mathrm{d}R(E_{\rm nr},M,\sigma_{\chi-\rm n}=\sigma_0)}{\mathrm{d}E_{\rm nr}}\left|\frac{\mathrm{d}E_{\rm nr}}{\mathrm{d}E_{ee}}\right| \mathrm{Gaus}(E-E_{ee},\sigma_{\rm res})\,\mathrm{d}E_{ee}, \tag{8}$$

where $\mathrm{d}R(E_{\rm nr},M,\sigma_{\chi-\rm n}=\sigma_0)/\,\mathrm{d}E_{\rm nr}$ is the predicted WIMP energy spectrum for a reference WIMP-nucleon cross section σ_0 and $C(\sigma_0)$ is such that the integral of f_s in the search region is normalized to 1. The ionization efficiency, $\mathrm{d}E_{\rm nr}/\,\mathrm{d}E_{ee}$, is used to convert the WIMP energy spectrum, which is a function of the the nuclear recoil energy $E_{\rm nr}$, to the observable variable E_{ee}, the ionization produced by the nuclear recoil. To account for the finite energy resolution of the detector we compute the convolution between the ionization energy spectrum and a Gaussian distribution with variance $\sigma_{\rm res}^2$ as modeled in Sec. VI. As a last step, the spectrum is multiplied by the detector efficiency, $\epsilon_{\rm det}(E)$, computed in Sec. VII.

To account for performance differences between the 1×1 and 1×100 data sets, we define a joint likelihood function:

$$\mathcal{L}_{\rm joint}(s_{\rm tot},\vec{b},M|\vec{E}) = \Pi_{k=1}^{\#\rm datasets}\mathcal{L}_k(\alpha_k(M)s_{\rm tot},b_k,M|\vec{E}), \tag{9}$$

where the index k runs over the different data sets and $\mathcal{L}_k$ is the corresponding likelihood function. Note that the functional forms of f_s and f_b depend on the data set as the efficiencies differ between data sets (FIG. 8). The total number of expected signal events, $s_{\rm tot}$, relates to the expected number of events on the k-th dataset through the multiplicative factor, α_k, that depends on the relative size of the exposure and the signal spectrum from a WIMP of mass M.

To quantify the statistical significance of a discovery and to compute an upper limit on the WIMP interaction rate (if no discovery is made), we perform a hypothesis test based on the profile likelihood ratio statistic q. This test compares the goodness of fit of two models, one of which, $\mathcal{L}_{\rm restricted}$, is a special case of the other, $\mathcal{L}_{\rm free}$.

For this discovery test the q statistic can be written as

$$q = -\ln\left[\frac{\mathrm{Max}\{\mathcal{L}_{\rm restricted}(\vec{b}|\vec{E},s_{\rm tot}=0)\}}{\mathrm{Max}\{\mathcal{L}_{\rm free}(s_{\rm tot},\vec{b},M|\vec{E})\}}\right], \tag{10}$$

where the numerator $\mathcal{L}_{\rm restricted}$ is maximized while keeping $s_{\rm tot}$=0. The denominator $\mathcal{L}_{\rm free}$ corresponds to the global maximum obtained from the fit to the data with all parameters free. The statistic q is positive by construction and values closer to zero indicate that the restricted fit has a likelihood similar to the unconstrained (free) case. On the

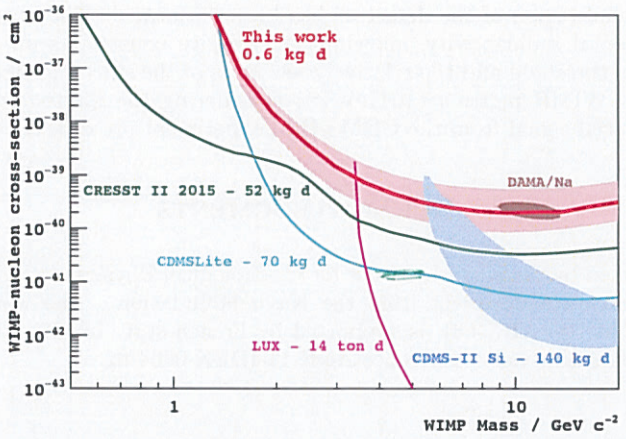

FIG. 9. Upper limit (90% C.L.) on the WIMP cross-section $\tilde{\sigma}_{\chi-\mathrm{n}}$ derived from this analysis (red line). The expected sensitivity $\pm 1\sigma$ is shown by the red band. For comparison we also include 90% C.L. exclusion limits from other experiments [? ? ?] and the 90% C.L. contours corresponding to the potential WIMP signals of the DAMA [?] and CDMS-II Si [?] experiments.

other hand, large values reflect that the null (background-only) hypothesis is unlikely. To quantify how likely is a particular value of q, the corresponding PDF is required. To compute this distribution we use a fully Frequentist approach and obtain it by performing the estimation of q outlined above for a large number of Monte Carlo samples generated from the background-only model (s_{tot}=0).

We perform the discovery test on the joint data set assuming the standard halo parameters: Galactic escape velocity of 544 m s^{-1}, most probable Galactic WIMP velocity of 220 m s^{-1}, mean orbital velocity of Earth with respect to the Galactic Center of 232 m s^{-1} and local dark matter density of 0.3 GeV cm^{-3}. We find the recorded events to be compatible with the background-only hypothesis with a p-value of 0.6. We thus proceed to set a 90% confidence level upper limit on the WIMP-nucleon elastic scattering cross-section, $\tilde{\sigma}_{\chi-\mathrm{n}}$.

To compute the upper limit we follow an analog approach where, for each value of M, we perform a scan on s to find a $\tilde{s}$ such that the test based on the corresponding $q(\tilde{s})$,

$$q(\tilde{s}) = -\log\left[\frac{\mathrm{Max}\{\mathcal{L}_{\mathrm{restricted}}(\vec{b}\,|\,\vec{E}, M, s_{\mathrm{tot}}=\tilde{s})\}}{\mathrm{Max}\{\mathcal{L}_{\mathrm{free}}(s_{\mathrm{tot}}, \vec{b}\,|\,\vec{E}, M)\}}\right], \tag{11}$$

rejects the hypothesis $s_{\mathrm{tot}} \geq \tilde{s}$ with the desired 90% C.L. Note that for each of the scanned points we generate the corresponding $q(s)$ distribution from Monte Carlo.

The limit on the WIMP cross-section $\tilde{\sigma}_{\chi-\mathrm{n}}$ is computed from $\tilde{s}$, the total exposure of the experiment, $\mathcal{E}$, and the normalization constant C (Eq. 8) as

$$\tilde{\sigma}_{\chi-\mathrm{n}} = C\frac{\tilde{s}}{\mathcal{E}}. \tag{12}$$

The 90% exclusion limit obtained from our data is shown by the red line in FIG. 9. The wide red band presents the expected sensitivity of our experiment, generated from the distribution of outcomes of 90% C.L. exclusion limits from a large set of Monte Carlo background-only samples. The good agreement between the expected and achieved sensitivity confirms the consistency between the likelihood construction and experimental data.

IX. CONCLUSION

We have presented the results of a dark matter search performed with 0.6 kg-days of data acquired with charge-coupled devices (CCDs) in the SNOLAB underground laboratory. The energy response of the detectors has been characterized down to a threshold of 60 eV$_{\mathrm{ee}}$. In particular, the linearity of the ionization signal and the ionization efficiency of nuclear recoils have been directly measured. The readout noise of the devices has been shown to be remarkably stable and predictable, allowing for the efficient selection of ionization events. The reconstruction of the depth of particle interactions in the silicon substrate from the measurement of the charge diffusion has been calibrated. We exploit this information to reject events occurring near the surfaces of the devices. A spectral analysis

of the candidates in the fiducial region of the detectors has been performed. The observed data are consistent with backgrounds from environmental radioactivity, allowing us to derive constraints on the WIMP-nucleon scattering cross-section. The low energy threshold and the relatively low mass of the silicon nucleus make the result particularly competitive in the low-mass WIMP region ($<10\,\mathrm{GeV}/c^2$) considering the limited experimental exposure. Cross-sections relevant to the potential signal from the CDMS-II Si experiment are directly excluded for the first time.

ACKNOWLEDGMENTS

This work has been supported by the Kavli Institute for Cosmological Physics at the University of Chicago through grant NSF PHY-1125897 and an endowment from the Kavli Foundation. The work of M.S., made in the ILP LABEX (under reference ANR-10-LABX-63), is supported by French state funds managed by the ANR within the Investissements d'Avenir programme under reference ANR-11-IDEX-0004-02.