
GPD workbook

Release 0.11.0

The IXPE Collaboration

Apr 18, 2021

BASICS:

1	GPD Overview	3
1.1	Absorption	4
1.2	Primary Ionization	6
1.3	Drift	8
1.4	Multiplication	9
1.5	Readout	9
2	Readout ASIC	11
2.1	Overview	11
2.2	Geometrical Layout	11
2.3	Electrical Characteristics	13
3	Triggering and ROI definition	17
3.1	Hardware implementation	17
3.2	ROI definition	18
3.3	Effective threshold and threshold dispersion	20
3.4	Trigger efficiency	21
3.5	Trigger Threshold and Containment	23
4	Event readout	25
4.1	Event capture timing	25
4.2	Serial Event Readout	26
4.3	Pedestal Subtraction	27
4.4	Analog Reset	29
4.5	Deadtime budget	29
5	Back-End Electronics	31
5.1	The DAQ board	31
6	The Gas Electron Multiplier	33
6.1	Design and challenges	33
6.2	Gain characteristics	33
7	GPD Inventory	35
8	Event Reconstruction	37
8.1	Clustering	37
8.2	Moments analysis	40
9	Monte Carlo simulation	49

10 The Role of the Gas Pressure	51
10.1 Absorption Efficiency	51
10.2 GEM Gain	51
10.3 Track Properties	54
11 Secular Pressure Variations	57
11.1 Short Summary and Basic Facts	57
11.2 Historical Note	57
11.3 Proxies for the Pressure	58
11.4 Basic Modeling	59
11.5 Fitting	61
11.6 Summary of the Results	63
11.7 Complementary Evidence	65
11.8 Impact on High-Level Scientific Performance	68
11.9 Appendix: Summary Plots for all Detectors	69
12 Charging	81
12.1 The Basic Equation	81
12.2 A Simple One-Species Model	82
12.3 Science Implications	85
12.4 The Microscopic Picture	85
12.5 A Realistic Two-Species Model	86
13 Alpha Particles	87
13.1 Basic Phenomenology	87
13.2 Stopping Power	87
13.3 Correcting the saturation	90
13.4 Track Profiles	92
13.5 Alpha particles and Charging	94
13.6 A Monte Carlo Simulation	95
14 Workbook Development	99
14.1 Structure of the Repository	99
14.2 Editing Contents	99
14.3 About the Figures	100
14.4 Compiling the Documentation	100
14.5 Changing the Style	100
15 Release Notes	103
16 Indices and tables	105

This page is meant to be a living entry point for technical information about the polarization-sensitive Gas Pixel Detectors (GPD) for the Imaging X-ray Polarimetry Explorer (IXPE) mission.

Warning: An initial version of this page was discussed in a face-to-face meeting of the Instrument Team at MSFC on January 15, 2020, and it was decided that the underlying technology (i.e., sphinx) was a sensible mean to share within the Collaboration the kind of technical information that might eventually converge into the end-user documentation associated with the public data release after launch. At this point in time the content is still fluid, and the details of the release process have to be sorted out, but from a technical standpoint the material is intended to be 100% correct.

If you are only here to read the documentation, go ahead and have fun. We appreciate feedback, and you are very welcome to report issues or propose changes on the project's [issue tracker](#).

If you are planning to contribute to the workbook instead, you should find all the relevant information on our [Workbook Development](#) page. Feel free to open an issue on our [tracker](#) if that is not the case.

GPD OVERVIEW

Fig. 1.1 illustrates the conceptual design of the IXPE Gas Pixel Detectors (GPD), as well as the basic detection principle. In a nutshell: photons enter the active gas volume through a thin Beryllium window, and have some probability of being absorbed in the gas. The primary ionization generated by the photoelectron (and, possibly, the Auger electron) drifts toward the Gas Electron Multiplier (GEM) under the action of the electric field in the absorption gap. The GEM provides the necessary gas gain, and the charge generated in the avalanche is (partially) collected on the readout ASIC, acting as a finely segmented anode.




Fig. 1.1: Conceptual design of the GPD. The volume of the gas cell is ideally divided in to parts: the (upper) absorption gap, between the entrance window and the GEM top, and the (lower) transfer gap, between the GEM bottom and the readout ASIC.

Table 1 summarizes some of the relevant figures of the GPD design for IXPE.

Parameter	Value
Typical V_{drift}	-2800 V
Typical V_{top}	-880 V
Typical V_{bottom}	-400 V
V_{ASIC}	0 V
Thickness of the absorption gap	10 mm
Thickness of the transfer gap	0.6 mm
GEM thickness	50 μm
Typical electric field in the absorption gap	2 kV/cm
Typical electric field in the GEM	100 kV/cm
Typical electric field in the transfer gap	5 kV/cm
Gas filling	Pure DME at 800 mbar

See also:

The Gas Electron Multiplier, Readout ASIC, Event readout, Triggering and ROI definition.

1.1 Absorption

X-rays impinge on the active gas volume through a 50 μm Beryllium window coated with Aluminum on the inner face (the thickness of the Al deposition is of the order of 50 nm). The transparency of the window is around 45% at 2 keV and greater than 95% above 5 keV, as shown in Fig. 1.2.




Fig. 1.2: Transparency of the GPD entrance window as a function of the photon energy. In addition to the 50 μm slab of Beryllium, the calculation includes the effect of the window aluminization on the inner side and the contaminants of the bulk materials, from the [data sheet](#).

Once the photon has passed through the entrance window, the probability of being absorbed in the active gas volume (1 cm of pure DME at 800 mbar and 20 degrees C) is shown in Fig. 1.3. Over the IXPE energy band, it ranges from about 45% at 2 keV to just below 1% at 8 keV.




Fig. 1.3: Photo-absorption efficiency of 1 cm of pure DME at 800 mbar and 20 degrees C as a function of the photon energy.

Roughly speaking, weak energy dependencies aside, photons are absorbed by the DME Oxygen atom two thirds of the time and in one of the two Carbon atoms the remaining one third of the times, the Hydrogen being irrelevant for all practical purposes. The fluorescence yields for all the three elements are negligible, and the atomic relaxation following a photoelectric absorption proceeds via the emission of an Auger electron in virtually 100% of the cases. Note that the energy of the Auger electron is different for O and C and, being its emission of direction spatially isotropic, this is one of the limiting factors for the polarimetric capabilities of the GPD at low energy.

Element	Absorption yield	K edge	Fluorescence yield
O	62.3%	525 eV	0.6%
C	37.7%	277 eV	0.2%
H	0.02%	13.6 eV	0

1.1.1 Passive Conversions

It is important to emphasize that *not all the X-ray events triggering the readout are absorbed in the active volume of the gas*. The contribution from X-rays absorbed in either the entrance window or in the top copper layer of the GEM and producing photoelectrons entering the gas cell with enough energy to cause a trigger is not negligible—and produces the characteristic low-energy tail in the pulse-height distribution from mono-energetic sources. The effect is more pronounced at relatively high energy, and is relevant since the modulation factor for events absorbed in the passive material is significantly lower than that for genuine gas absorptions.




Fig. 1.4: Relative probability of absorption in the active gas volume, the entrance window and the GEM top copper layer for all the events triggering the readout ASIC, as a function of the photon energy.

It is maybe interesting to note that Monte Carlo simulations clearly show that most of the triggering window conversions actually take place in the thin Aluminum coating on the inner face of the window. This is due to the fact that Aluminum has a much higher photoelectric cross section relative to Beryllium, and the coating is thin enough (50 nm) that the escape probability for the photoelectron is not negligible.

Note: The passive conversions represent an additive term to the expected event rate on top of that show in Fig. 1.3.

On a related note, the Copper line at 8.9 keV creates an abrupt increase of the GEM conversions above this energy. This effect does not show up in Fig. 1.4 as it happens to be outside the canonical 2–8 keV IXPE energy range, but would need to be taken into account in order to expand the energy reach of the detector beyond 8 keV.

1.2 Primary Ionization

Note: The photoelectron energy is, in general, smaller than the original photon energy by an amount close to the energy of the K shell of the element by which the photon itself is absorbed (525 or 277 eV in DME).

The average energy necessary to produce an electron-ion pair in DME is of the order of 28 eV. This means that the photoelectron (and the Auger electron) emitted when a 5.9 keV photon is absorbed produce about 210 electrons of primary ionization. More properties of the DME relative to the generation of the primary ionization are listed in the following table.

Parameter	Value
Chemical composition	(CH ₃) ₂ O
Density at 1 atm, 0 degrees C	2.115 mg/cm ³
Average energy per electron/ion pair	28 eV
Average number of primary pairs @ 5.9 keV	210
Fano factor	0.30

We note, incidentally, that these figures set the scale for the best possible energy resolution achievable in practice

$$\text{FWHM} \left(\frac{\sigma_E}{E} \right) = 2.35 \sqrt{\frac{F}{N_p}} = 14.5\% \text{ @ } 5.9 \text{ keV}$$

Our Geant4 Monte Carlo simulation indicates that the electron range in the Continuous Slow-Down Approximation (CSDA) is of the order of 0.1 mm at 2 keV and 1.1 mm at 8 keV. In practice, due to the Coulomb scattering, the *effective range* (i.e., the distance between the beginning and the end of the track) is significantly shorter than that, the *detour factor* being 0.55–0.65 depending on the energy, as shown in Fig. 1.5.




Fig. 1.5: Track length in the GPD as a function of the photoelectron energy. The dashed lines represent the expected photoelectron range in the Continuous Slow-Down Approximation (CSDA) and the effective range (accounting for the multiple scattering), respectively. The points represent a proxy for the reconstructed track length, based on the full Geant4 Monte Carlo simulation of the detector, and are affected by additional Physical processes, such as the charge diffusion in the gas.

At the reconstruction level we have a good proxy of the track length in the square root of the second moment along the longitudinal axis of the track. (Since most of the charge is included within 2 sigma from the barycenter, a scale factor of the order of 4 allows to compare this metrics directly with the range predicted by the simulation.)

See also:

Event Reconstruction

It goes without saying that, in real life, there are other effects affecting the reconstructed track length (primarily the transverse diffusion in the gas) and the latter can be dominant, especially at low energy.

1.3 Drift

The primary ionization generated by the photoelectron (and by the Auger electron) drifts toward the Gas Electron Multiplier under the effect of the electric field in the absorption gap. The process is essentially a random walk governed by the transverse diffusion coefficient, which is function of the drift electric field.




Fig. 1.6: Transverse diffusion coefficient as a function of the drift electric field in pure DME at 800 mbar (calculated with [Garfield](#)).

In the simplest formulation the distribution of the transverse displacement of the charge after a given drift length, in any of the two orthogonal directions, can be assumed to be a Gaussian with mean 0 and a standard deviation of

$$\sigma = \sigma_t \sqrt{d}.$$

The detector is typically operated setting the drift field at the minimum of the transverse diffusion—that for DME at 800 mbar is around 2 kV/cm.

Parameter	Value
Minimum transverse diffusion	$68 \mu\text{m cm}^{-\frac{1}{2}}$
Minimum diffusion field	2 kV cm^{-1}

1.4 Multiplication

The GEM provides the gas multiplication of the primary ionization, which is necessary to generate an electrical signal detectable by the readout ASIC.

In the simplest possible treatment of the problem, the probability density function for the number of secondary electron/ion pairs generated by a single primary electron is exponential, and the mean of the distribution is what we general refer to as *gain*:

$$p(n) = \frac{1}{G} \exp \left\{ -\frac{n}{G} \right\}$$

Simulations show that, depending on the exact microscopic configuration of the electric field, a significant fraction of the charge generated in the avalanche is immediately lost on the GEM bottom face, and it is therefore necessary to distinguish the *avalanche gain* (related to the avalanche charge) and the *effective gain* (related to the charge available at the readout). At the typical working point, roughly speaking, the former is of the order of 400 and the latter about half that much.

Parameter	Value
Average avalanche gain	400
Average effective gain	200
Average readout charge at 5.9 keV	6.5 fC

See also:

The Gas Electron Multiplier

1.5 Readout

The track image is derived by the signal induced on a finely segmented readout anode, implemented as the top Aluminum layer of a custom ASIC. All the details about the readout chip are covered in the [Readout ASIC](#) section.

See also:

Readout ASIC

READOUT ASIC

2.1 Overview

The XPOL ASIC constitutes the multiple anode of the Gas Pixel Detectors at the focal plane of the IXPE mission. The CMOS VLSI chip integrates more than 16.5 million transistors and is organized as a matrix of 105600 pixels with a 15 mm x 15 mm active area. Each pixel is composed of a hexagonal electrode (top layer of metal) connected to a charge-sensitive amplifier followed by a shaping circuit.

The chip's honeycomb matrix is organized in 16 identical clusters of 6600 pixels (22 rows of 300 pixels each) with an independent differential analog readout buffer. Each cluster features a customizable self-triggering capability. Upon the activation of an external digital control input, or from on-chip wired-OR combination of each cluster self-triggering circuit, the maximum of the shaped pulse is stored inside each pixel cell for subsequent readout. The latter is accomplished by sequentially connecting the output of each cell to the analog bus common to each cluster. The self-triggering function also includes an on-chip signal processing for automatic localization of the event. This allows reducing the readout time, as in the corresponding readout mode only the pixels within the region of interest are scanned.

2.2 Geometrical Layout

The pixel matrix is organized in 300 columns at 50.00 μm pitch and 352 rows at 43.30 μm pitch, as show in the figure below.

Parameter	Value	Description
N_{col}	300	Number of columns
N_{row}	352	Number of rows
p_{col}	50.00 μm	Column (horizontal) pitch
p_{row}	43.30 μm	Row (vertical) pitch

Each pixel is identified, in the electronics reference system, by a unique address $\langle X, Y \rangle$, $\langle 0, 0 \rangle$ corresponding to the upper left pixel and $\langle 299, 351 \rangle$ to the lower right pixel of the matrix. This is the very reference system used by the chip to store the coordinates of corners of the the region of interest to be readout upon trigger.

The physical reference system on the readout plane is defined according to the following specifications:

1. the x axis is directed along the matrix rows, in the direction of increasing X;
2. the y axis is directed along the matrix readout columns, in the direction of decreasing Y (or row number);
3. the origin of the system is such that the center of the leftmost (uppermost) pixel and the center of the the rightmost (lowermost) pixel in the matrix are equidistant from the center itself along the x (y) coordinates.




Fig. 2.1: Geometrical layout of the readout ASIC for the IXPE focal-plane detectors. The different colors (white and grey) used for the pixels has no other meaning than visually distinguishing even and odd columns.

This last property is handy because it allows to specify fiducial cuts in physical space in the simplest possible way (i.e., x and y can be used directly to define circles and/or rectangles centered within the readout plane).

The system is uniquely determined and the basic transformations between pixel coordinates and physical coordinates read:

$$x = p_{col} \left[X - \frac{1}{2} \left(N_{col} - \frac{3}{2} + Y \bmod 2 \right) \right]$$

$$y = p_{row} \left[\frac{1}{2} (N_{row} - 1) - Y \right]$$

(For completeness, the azimuthal track angle is measured counter-clock-wise from the x axis as shown in the figure. We also note that the center of the coordinate system does not coincide with the center of any of the pixels in the matrix.)

2.3 Electrical Characteristics

Before we start discussing some highly technical aspects of the GPD operation and performance, it is useful to set the stage and write down a series of fundamental figures that provide context for the following chapters.

Parameter	Value
Reference voltage (V_{ref})	800 mV
Shaping time	4 μ s
Pixel gain	400 mV/fC
Pixel Noise	22 electrons ENC
Dynamic range	1 V

Upon power on, the readout chip generates a *reference voltage* (typically referred to as V_{ref}), of the order of 800 mV, that is distributed to all the pixels in the matrix and acts like a fundamental reference for the pedestals and the trigger threshold (although, technically, both the pedestals and the thresholds have a finite dispersion and, generally speaking, none of them is *exactly* equal to V_{ref}).

Note: Although, strictly speaking, this is related more to the back-end electronics than to the readout chip, it is useful to remind that the ADC on the DAQ board has 14 bits of resolution over a 2.4 V full dynamic range, or a voltage resolution of 0.146 mV/ADC count. Coupled to the pixel gain, this translates into a charge characteristics for the whole system of 2.3 electrons/ADC count.

It follows that the 1 V dynamic range of the readout ASIC corresponds to a saturation of about 6800 ADC counts.

2.3.1 Gain

The XPOL specifications call for an average gain for the pixel amplifier of 400 mV/fC, or 64 μ V/electron. While it is less than trivial to perform a precise measurement of this figure, mainly for the fact that in the GPD the ASIC and GEM gain are hard to disentangle from each other, the number is not unreasonable, by assuming an effective gain of the GEM of about 200 at the nominal working point.

We emphasize that the internal charge-injection system of the ASIC is not suitable for a precise, pixel-by-pixel gain measurement, as each pixel has a separate test capacitor, and the corresponding capacitance dispersion is of the same order of magnitude or larger than the gain dispersion.

See also:

Back-End Electronics, The Gas Electron Multiplier.

2.3.2 Noise

The noise of the system can be effectively measured exploiting the internal charge injection system, by stimulating one at a time a set of pixels across the entire active area of the detector and calculating the standard deviation of the PHA values in the ROIs being generated.




Fig. 2.2: Typical distribution of the pixel noise. The three square-like structures visible in the color plot are three distinct ROIs around noisy pixels that tend to appear relatively more often than the average in the output data stream.

The average measured pixel noise is of the order of 9.5 ADC counts, that, when rescaled for the charge characteristics of the system, translates into an equivalent noise charge of about 22 electrons.

Note: When operating with an event-by-event pedestal subtraction with a single pedestal sample, the *effective* noise is really a factor of square root of 2 larger, or 31 electrons ENC.

Although there is a relatively long tail in the flattened noise distribution, 95% of the pixels features a noise RMS smaller than 14.3 ADC counts (or 1.5 times the average value) and the vast majority of the matrix below 20 ADC counts RMS. In addition, the noise distribution is fairly uniform across the active area of the ASIC.

2.3.3 Pedestals

Pixel pedestals are measured in the same fashion as the pixel noise is measured—as a matter of fact the two things are performed contextually with the very same pixel scan. Fig. 2.3 shows a typical pedestal distribution across the active area of the readout ASIC.




Fig. 2.3: Typical distribution of the pixel pedestals. The slight regular non-uniformity across the active are of the chip is due to a proximity effect to the single output buffer used for the readout.

TRIGGERING AND ROI DEFINITION

Although the XPOL ASIC can be externally triggered, one of its peculiar features is its self-triggering capability, which is exploited in the standard IXPE operating mode and is therefore the main focus of this section.

3.1 Hardware implementation

At the very fundamental level every 4 pixels are logically OR-ed together to contribute to a local trigger with a dedicated fast (relatively to that used for pulse-height analysis) shaping amplifier, as shown in the Figure below. This basic building block of 2 by 2 pixels is called a *trigger minicluster*, and is central to the entire machinery of event triggering and readout.




Fig. 3.1: Geometrical layout of a 2 by 2 trigger minicluster, with the aluminum charge-collecting pads of the four pixels visible on the top layer. The XPOL analog part of the XPOL readout ASIC is essentially a 150 by 176 matrix of exact replicas of this 100 x 86.6 μm minicluster.

The reasoning behind this design choice was the need to have sufficient margin in the signal-to-noise ratio to guarantee full trigger efficiency at all the energies of interest (the signal in the four pixels sums up linearly, while the noise, being incoherent, sums up in quadrature).

A few remarks on the trigger circuit layout:

- the top-left pixel of a minicluster is always even-even (and the parity of the other three pixels follow from this), e. g., the minicluster on the top-left corner of the ASIC is made of pixels at address $\langle 0, 0 \rangle$, $\langle 0, 1 \rangle$, $\langle 1, 0 \rangle$ and $\langle 1, 1 \rangle$;
- each minicluster has an aspect ratio different from unity, i. e., its width is intrinsically larger than its height; as we shall see in the following, this has implications for the definition of the ROI;

- each pixel includes a masking function to disable its contribution to the trigger function (note that the masking bits are to be set at each power-up of the circuit).

As it turns out, in normal operating conditions, the part of the track that participates in the trigger is mainly the Bragg peak. Capturing the parts of the track with a relatively low ionization density relies on a proper padding of the region of interest (ROI) to be read out, as we shall see in the next section.

3.2 ROI definition

Upon trigger, the event is automatically localized by the ASIC core logic in a rectangle containing *all triggered miniclusters plus a padding of 4 / 5 additional miniclusters (along the X / Y coordinates)*. More specifically, the chip calculates the region of interest defined by the <Xmin, Ymin> and <Xmax, Ymax> coordinates of the upper-left and lower-right corners of the rectangle, respectively. (It goes without saying that the ROI is truncated to the physical size of the readout matrix when the event happens to be at the edges of the active area.)

See also:

Event readout

The fact that the padding (in logical coordinates) is different along the two directions is meant to compensate as much as possible for the aspect ratio of the underlying minicluster, as in physical space this figures translate into $4 \times 100 = 400$ μm in X and $5 \times 86.6 = 433$ μm in Y, i.e., the additional padding is roughly the same in all four directions. *As long as the trigger efficiency is close enough to 100% and the tracks are fully contained in the ROIs of the events, this trigger/readout strategy does not introduce, per se, any preferential direction.*

It is important to notice that the information about the exact set of miniclusters participating to the trigger is not directly carried over by the chip upon the definition of the region of interest. Therefore the two basic rules for interpreting the output data stream are:

- *we don't know the exact combination of miniclusters within the core of the ROI that participated to the trigger*—all we know is that this combination generated a particular rectangle, e.g., at least two opposite corners in the core did trigger;
- *we know that none of the miniclusters outside the core, i. e., in the padding border did trigger.*

Note: That all said, we can always obtain a (more or less accurate) proxy of the precise trigger minicluster topology by summing up the signals in groups of 2 by 2 sets of pixels and comparing the result offline with an estimate of the trigger threshold with each specific minicluster. Being the electronics paths for the trigger and for the spectroscopy physically different, this strategy will never achieve a perfect precision, but the experience shows that this approach provides useful information in practice.

The smallest possible region of interest in normal condition is that where a single minicluster is triggering (e.g., this is the typical situation when we inject charge in a single pixel). In this case the resulting ROI is $(4 + 1 + 4) = 9$ by $(5 + 1 + 5) = 11$ miniclusters, or $18 \times 22 = 396$ pixels. On the edges of the active area, where the window is truncated to the physical boundaries of the chip, the region of interest can be as small as $(1 + 4) = 5$ by $(1 + 5) = 6$ miniclusters, or $10 \times 12 = 120$ pixels, if the triggering minicluster happen to be on one of the four physical corners of the ASIC.

At any given energy the routine Science data acquisition generally involves a variety of ROI configurations. The following figure shows the distribution of the most frequent regions of interests at 2.6 keV. At higher energies, the average ROI size tends to be larger, and the relative spread in the window size also tends to be larger (both of which is to be expected).

Warning: It cannot be stressed enough that, at any given energy, electrons emitted at different angles will tend to trigger preferentially different, well defined regions of interest. Horizontal events will preferentially result in




Fig. 3.2: Geometrical arrangement of a $22 \times 26 = 572$ region of interest, which is quite typical for low-energy photons (for completeness, this ROI topology accounts for about half of the tracks at 2.6 keV). The white pixels constitute the *core* of the ROI, defined by the combination of miniclusters participating in the trigger, while the padding pixels (8 rows on the left and the right, and 10 columns on the top and the bottom) are depicted in gray. The indices running along the X and Y axes outside the ROI matrix are relative to the $\langle X_{\min}, Y_{\min} \rangle$ coordinates of the top-left corner.




Fig. 3.3: Distribution of the ROI topology for 2.6 keV, unpolarized X-rays. The most frequent region of interest (22 by 26) accounts for about a half of the events.

ROIs that are large and short, and vertical events will tend to trigger windows that are narrow and tall. A basic consequence of this is that the modulation curves for any specific ROI configuration are, in general, significantly modulated, and quite often highly so. This is a well-known effect that is understood and reproduced by the Monte Carlo simulation of the detector, and the fact can be simply rephrased as: it is never a good idea to perform selection cuts on the ROI size.

3.3 Effective threshold and threshold dispersion

The trigger threshold is physically determined by a dedicated DAC on the back-end electronics, and is conveniently measured with respect to the V_{ref} of the ASIC, since the latter sets the approximate level for the pedestals of the pixels amplifiers. Converting the trigger threshold to a physically meaningful value (e. g., electrons or equivalent ADC counts) requires a precise knowledge of the gain values (in mV/fC) for both the fast (trigger) and the slow (spectroscopy) amplifier, as well as the relative offset, on a pixel by pixel basis. This is impractical to do, based on first principles, as we have essentially no direct mean to measure these figures and compare with the design values.

That all said, the fact that the miniclusters in the padding ring are guaranteed not to have contributed to the trigger (which in turn implies that the signal was lower than the trigger threshold) provides an opportunity for a sensible measurement of the threshold itself *in equivalent ADC counts*, i. e., in physically interesting units. A generic algorithm to achieve that could be spelled as:

- for a given event, select all the pixels in the padding ring of the region of interest;
- sum the PHA values for the four pixels within each trigger minicluster, and accumulate the corresponding PHA histogram (again, for each minicluster);
- evaluate the trigger threshold as the endpoint (e. g., estimated via a given quantile of the corresponding cumulative distribution) of such histogram.

This is illustrated in the following figure.




Fig. 3.4: Pulse-height distribution for a specific minicuster, accumulated for the events where the minicuster itself does not participate to the trigger. The endpoint of the distribution is a measure of the effective trigger threshold for that very minicuster.

A deep flat field with sufficient statistics (e.g., any of the ones that we use for the characterization of the spurious modulation) allows therefore for a direct measurement of the trigger threshold on a minicuster by minicuster basis, as illustrated in the following figure.

The two numbers to keep in mind are essentially:

- *the average effective trigger threshold with the nominal detector configuration is of the order of 450 ADC counts (or 1000 electrons of signal charge);*
- *the relative minicuster-by-minicuster dispersion of the values is of the order of 10%.*

The first figure is important to understand the trigger threshold (see next section), while the second is in line with the expectations, based on the uniformity of components of the underlying technology.

3.4 Trigger efficiency

As mentioned before, in normal operating conditions, the part of the track that participates in the trigger is mainly the Bragg peak, whose morphology is largely independent from the energy of the incoming photon.

Roughly speaking, an electron propagating in DME at 0.8 bar releases some 1.5 keV in the final 100 μm of its track. In ideal conditions, at the nominal effective gain of 3000 ADC counts per keV, this would in principle translate into 4500 ADC counts, independently on energy. In practice the effects of the transverse diffusion through the drift length of the primary ionization and the charge sharing onto the readout plane are such that the minicuster with the highest signal is significantly lower, typically of the order of 2000 ADC count. As discussed in the previous section, *this is still a factor of 4 larger than the effective trigger threshold in the nominal science data-taking configuration*, which implies that the GPD trigger efficiency is nominally 100% with a fairly large margin.

All of this is qualitatively consistent with the fact that, as illustrated in the following figure, the trigger rate measured in a given setup is, to first order, independent from the trigger threshold up to 3 to 4 times the nominal value of 200–250 mV.




Fig. 3.5: Distribution of the effective trigger threshold for a sub-sample of 583 minicluster, obtained with a deep flat field over a 3 mm diameter circular spot. With a deep flat field over the full detector the measurement can be performed over the entire ASIC.




Fig. 3.6: Event rate as a function of the trigger threshold measured for the GPD EM (GPD 23), irradiated with 5.9 keV X-rays. As expected the rate is roughly independent from the threshold over a large range, up to 3 or 4 times the nominal value of 200–250 mV. For completeness, a linear fit to the plateau region provides a slope of about $3e-5$, corresponding to a relative drop of 1% going from 200 to 500 mV.

3.5 Trigger Threshold and Containment

Although the trigger efficiency is, to a large extent, threshold independent, this is not the case for most of the event-by-event topological values, as the size of the region of interest tends to get smaller as the threshold increase, up to the point where the padding might become insufficient to contain the entire track (i.e., the track itself can be truncated).

Somewhat counter-intuitively, this more than an issue at relatively high energies, where the tracks are longer, and the average specific ionization is smaller. For reference the basic numerology at 5.9 keV, when increasing the threshold from 200 to 500 mV (which we emphasize is unrealistically high by any reasonable metrics), is as follow:

- the average ROI size decreases by some 10%;
- the average track size decreases by about 2.5% (see the figure below);
- the average pulse height decreases by about 1.5%;
- the fraction of truncated events rises from zero to slightly above 1%.

As mentioned above, the situation is even more favorable at lower energies, where the peak of the IXPE sensitivity is.




Fig. 3.7: Average track size at 5.9 KeV as a function of the trigger threshold measured for the GPD EM (GPD 23).

Note: While the decrease of the ROI size is potentially beneficial in terms of necessary data volume for bright sources (and, to a much lesser extent, in terms of dead-time per event) the track truncation effect is potentially detrimental (and more so), as it is generally the initial part of the track (with the lowest specific ionization, but with most of the polarimetric information) being lost. The most sensible strategy for setting the trigger threshold is to set it to mask all the pixels out of family and then set it to the smallest possible value compatible with the noise floor. The sweet spot between 200 and 300 mV is satisfactory from all the standpoints.

EVENT READOUT

Although the XPOL readout ASIC provides four different readout modes (including the direct addressing of a single pixel and the parallel readout of the entire frame with either 8 or 16 output buses), the IXPE back-end electronics is designed in such a way that only the so called *windowed* readout mode, in which the chip is configured for self-triggering and automatic definition of the ROI for the serial readout, can be exercised.

As explained in details in the [Triggering and ROI definition](#) page, in windowed mode the event is localized upon trigger in a rectangle containing all triggered mini-clusters plus a pre-defined padding for successive serial readout(s).

See also:

[Triggering and ROI definition](#)

4.1 Event capture timing

The basic XPOL event capture timing is illustrated in the following diagram, which we shall describe with some details.

Note: The left and right part of the diagram, and particularly the two red lines on the top, are not meant to be intended as happening one after the other in time, but rather in parallel, the left part showing the output of the slow shaping amplifier for a given pixel, and the right the output of the sample and hold for the same pixel. (T_{pk} is the shaping time of the aforementioned amplifier, or 4 us.)

The `TrgWindow` LVDS input is a logical signal that inhibits the trigger, i.e., it has to be asserted in order for the chip to trigger at all, whatever the readout mode is. In windowed mode the `TrgWindow` is typically asserted, except for a small period of time around the analog reset, as we shall see later in the section.

In self-trigger mode the `TrigOut` signal is generated by the chip whenever one or more minicluster outputs cross the threshold, and prevents additional triggers to happen until it is externally reset. T_{dh} is the time delay between the arrival of the event (assuming for the sake of the discussion that the input to the charge-sensitive amplifier is a delta) and the threshold crossing. Since the shaper equipping the trigger chain is faster than the one connected to the pixel amplifier, this is smaller than T_{pk} , and typically of the order of 1 us.

The `MaxHold` control input is really relevant when the self-trigger functionality is disable (i.e., in the first three readout modes), but needs to remain asserted even in windowed mode until the readout sequence has been completed.

Going back to the standard windowed readout mode, as soon as the `TrigOut` signal is asserted, the chip asserts the `Track` signal, and after a period of peak detection of 10 us the `Hold` signal freezes the analog output data and, contextually, the `DataReady` LVDS output is asserted, signaling that the actual serial readout sequence can start.




Fig. 4.1: XPOL event capture timing. Note that, since the diagram is meant to act as a reference for all the four XPOL readout modes, not all the parts are directly relevant for the windowed mode in which IXPE operates, as we shall clarify in the main text.

4.2 Serial Event Readout

As stated multiple times already, upon an event trigger in windowed readout mode, the pre-defined window padding is added to the minimum rectangle area containing all the triggering miniclusters. The coordinates of the corresponding corners $\langle X_{min}, Y_{min} \rangle$ and $\langle X_{max}, Y_{max} \rangle$ are automatically stored into an internal register and available as four 9-bit data outputs as soon as the `DataReady` output is asserted, ready to be used to insert the line and column tokens at the right place to scan the corresponding area of interest as one dynamically-defined cluster.

Note: The aforementioned window coordinates do survive until they are explicitly reset (via the so-called `AnalogReset` signal), and can be used for additional, multiple readouts of the same ROI. This is extremely important, as it is the basic mechanism through which we perform the event-by-event pedestal subtraction.

Upon definition of the region of interest, the serial readout proceeds driven by the readout clock, and the analog output of each pixel is sequentially routed to the differential output buffer connected to the external ADC. The most critical parameter in this context is T_{asu} , or the analog output settling time, i.e., the time that is necessary for the output buffer to reach the target value for the ADC to sample. Extensive experience in the characterization of the spurious modulation shows that this is of the order of 200 ns, which limits the maximum clock frequency for serial readout to 5 MHz.

Note: The readout clock is provided by the IXPE back-end-electronics by prescaling the 50 MHz system clock by a fixed, integer demultiplier. For the nominal 5 MHz readout clock and a ROI size of 600 pixels, which is fairly typical at the energy where the IXPE sensitivity peaks, the time which is necessary for a complete event readout is 120 μ s.




Fig. 4.2: XPOL serial readout timing.

(Technically T_{ahd} , or the analog buffer output time, is also potentially relevant, as the associated discharge time is not infinite, but since we're interested in reading as quickly as possible, this is not a concern in practice.)

When looked at in the two-dimensional, physical space, the sequential readout of pixels proceeds left-to-right and top-to-bottom with a recirculating token driven by the readout clock.

Warning: When switching from one row to the next one, the settling time for the signal on the first pixel of the row is significantly longer because of the time needed for the analog bus lines to charge up. In practice the data from the first column are largely unusable and, as a matter of fact, are identically set to zero by the firmware running on the FPGA of the DAQ board, in order to optimize the compression factor of the data packing algorithm. As the ROI padding, with the nominal science data-taking settings, is typically more than enough to fully contain the events, this has no measurable practical implication.

4.3 Pedestal Subtraction

In nominal science data taking mode each region of interest is read out twice (or more), the first readout sampling the physical signal and the second (and, possibly, all the subsequent ones) sampling the pedestal. The FPGA on the DAQ board then performs the pixel-by-pixel subtraction and proceeds to the zero suppression and the event compression.

The main advantage of the event-by-event pedestal subtraction is that it effectively cancels out all time-dependent effect. And the effective noise with a single pedestal sample is only a factor of square root of two worst than the optimal situation with a large number of samples.

Note: One interesting implication of the geometrical arrangement of the serial readout is the fact that the time, relative to the trigger signal, at which each pixel is read out is determined by its position in the region of interest. *Any possible transient generated by the digital readout activity and affecting the analog output might affect different pixels in slightly different ways, and the same pixel might be affected in slightly different ways depending on the exact size and shape of the ROI for a given event.* This is one of the factors that makes the on-board (or offline) subtraction of pre-determined pedestals significantly more challenging than one might imagine, and makes the event-by-event pedestal subtraction so appealing (for the simple fact that when the ROI is read out multiple times in exactly the same fashion these effects tend to nicely cancel out.)

One important drawback of the event-by-event approach to the pedestal subtraction is the impact on the deadtime per event. This is not only due to the fact that one has to go through the serial readout twice (or more). The most important limitation is actually constituted by the fact that, in order to obtain a pedestal sampling statistically independent from




Fig. 4.3: Sample 22 x 26 = 572 region of interest. The readout proceeds left-to-right and top-to-bottom, as indicated by the sequential number labeling each pixel.

the signal readout, one has to wait a time of the order of 500 μs between two subsequent readouts. As a matter of fact, this is the single larger contribution the deadtime in nominal data taking.

4.4 Analog Reset

At the end of the readout sequence, the reset of analog signals and all trigger-related variables (e.g., the ROI coordinates) is performed by applying a pulse (about 5 μs wide) on the `AnaResetbar` input, as show in the diagram below.




Fig. 4.4: XPOL analog reset timing.

Warning: In order to avoid false triggering due to possible ringing immediately after the analog reset, it is necessary to inhibit the trigger for a time of the order of 150 μs after the analog reset. In standard self-trigger, windowed mode, this is achieved by temporarily de-asserting the `TrigWindow` LVDS signal.

4.5 Deadtime budget

At this point we are ready to flash out the most important contributions to the GPD deadtime budget. The timing involved in the typical readout sequence (with a single pedestal sample) roughly reads:

- trigger threshold crossing time (T_{dh}): 1 μs ;
- peak-detection time: 10 μs ;
- signal serial readout: 120 μs (for 600 pixels at 5 MHz);
- pedestal delay, necessary for the signal in the sample and hold to decay: 500 μs ;
- peak-detection time (repeated to make the timing of the second readout identical to the first): 10 μs ;
- pedestal readout: 120 μs (for 600 pixels at 5 MHz);
- pedestal subtraction: 120 μs (for 600 pixels at 50 MHz system clock, assuming 10 clock cycles per pixel for the arithmetic);
- analog reset: 5 μs ;
- analog reset delay, necessary to avoid re-triggering due to possible ringing induced by the analog reset: 150 μs .

Note: The pedestal subtraction cannot be easily parallelized and executed after the analog reset, as it needs exclusive access to the memory bank, just like the event readout. The data transmission through the science data interface, on

the other end, is not blocking.

Actual mileage might vary depending on the precise settings of the back-end electronics related to the readout sequence, but what we are looking at is essentially a constant term of the order of 700 us, plus a term proportional to the ROI size, with a slope of 200×2 (readout) + 200 (subtraction arithmetics) = 600 ns/pixel. For a *typical* ROI of 600 pixels, this back-of-the-envelope calculation provides an estimate of about 1 ms, which is the right ball-park.

BACK-END ELECTRONICS

The IXPE Back-End Electronics (BEE) is responsible for all the commanding and control of the readout chip, as well as the generation of low- and high-voltages for the detector and the handling of the science data and telemetry interfaces.

At the hardware level, the BEE consists of three distinct electronic boards plugged onto a common backplane:

- the Low-Voltage Power Supply (LVPS);
- the High-Voltage Power Supply (HVPS);
- the DAQ board.

It goes without saying that the latter is the most relevant for the GPD operation and is therefore the main focus of this section.

5.1 The DAQ board

From a functional standpoint, the DAQ board incorporates two fundamental components of the GPD operation:

- the analog-to-digital converter for the serial readout;
- the FPGA controlling the ASIC configuration, control and readout, as well as the data formatting and the communication with the detector service unit.

5.1.1 The Readout Sequence

One of the main functions of the DAQ board is the management of the ASIC readout sequence. This functionality is implemented with a fair degree of flexibility, some of the most important adjustable parameters being:

- *Readout clock demultiplier*: the master clock prescale factor determining the clock frequency for the serial readout;
- *Number of pedestals samples*: the number of event to be readout for the online pedestal subtraction;
- *Pedestal sample delay*: the delay between the signal readout and each successive readout for pedestal subtraction.

The DAQ board has the capability of interrupting the readout sequence, based on the size of the region of interest, right after the ASIC has latched the coordinates of the latter. There are two separate adjustable parameters determining the minimum and maximum number of pixels in the ROI for any given event to be readout. This functionality is intended to inhibit the readout, without inducing un-necessary dead time, for events that are extremely unlikely to be genuine X-Ray, e.g., noise triggers of very large tracks induced by charged particles.

See also:

Readout ASIC.

5.1.2 Analog to Digital Conversion

The ADC on the DAQ board has 14 bits of resolution over a 2.4 V full dynamic range, or a voltage resolution of 0.146 mV/ADC count. Coupled to the pixel gain, this translates into a charge characteristics for the whole system of 2.3 electrons/ADC count.

Parameter	Value
Dynamic range	2.4 V (-1.2 V to 1.2 V)
Resolution	14 bit
ADC voltage resolution	0.146 mV/fC
System charge characteristics	2.3 electrons/ADC count

5.1.3 Event Formatting

In addition to the event-by-event pedestal subtraction, the DAQ board is responsible for formatting the event data and transferring them to the detector service unit. This operation includes:

- a zero suppression with an adjustable threshold;
- a compression with a dedicated algorithm.

A subset of the events, with an adjustable prescale factor, is passed along without zero suppression and compression, with a constant of 256 ADC counts added. The intent of these *diagnostic events* is to have access to the negative part of the noise distribution in the ROI for diagnostic purposes.

THE GAS ELECTRON MULTIPLIER

The Gas Electron Multiplier provides the gain stage for the GPD. Being intrinsically two-dimensional, it is particularly suited for applications, such as photoelectric X-ray polarimetry, where preserving the morphology of the primary charge cloud is crucial.

6.1 Design and challenges

Gas Electron Multipliers are customarily used in high-energy physics and X-Ray detection and imaging. Compared with those used in different applications, the main peculiarity of the IXPE GEM is their fine pitch, which in turn is dictated by the necessity to preserve as much as possible the morphology of the photoelectron tracks.

Parameter	Value
Hole pitch	50 μm
Hole diameter	30 μm
Thickness	50 μm
Substrate	Liquid Crystal Polymer (LPC)
Metal coating	Copper

Roughly speaking, the manufacturing process for the GEMs can be broken up into three main steps:

- the holes are drilled into the top and bottom Copper layers by means of standard lithographic techniques;
- the holes are drilled into the dielectric with laser etching process;
- a wet-etching post-processing polishes the Copper surfaces in order to ensure the necessary robustness against micro-discharges.

It is worth highlighting that the process, fine-tuned through the breadboard development phase of the mission, is pushing the manufacturing technology to the limits in all the areas.

See also:

Charging

6.2 Gain characteristics

The GEM gain is known to scale exponentially with the voltage across the top and bottom faces over a relatively wide voltage range below the discharge regime:

$$G(V) = G_0 e^{\gamma V}. \quad (6.1)$$

(We emphasize that, given the memory-less nature of the exponential, this relation holds no matter if the GEM voltage is measured in absolute terms or relative to a generic working point.)




Fig. 6.1: Typical GEM gain-voltage characteristics. The uniformity of the GEMs, in this respect, is excellent across the flight production, and this plot is representative of any of the IXPE flight GEMs.

There are two interesting quantities that can be calculated starting from this basic model, namely the derivative of the relative gain increase with respect to the input voltage (which, incidentally, coincides with the gamma parameter in the model)

$$\frac{1}{G} \frac{dG}{dV} = \gamma \quad [\text{V}^{-1}],$$

and the voltage increase determining an increase in the gain of a factor of 2

$$V_{double} = \frac{\log 2}{\gamma} \quad [\text{V}].$$

Either of these two figures essentially provides a complete characterization of the GEM characteristics, at least in this simple model. (As we said before, if the voltage change is measured relative to a given working point, neither one depends on the specific value of the working point itself.)

Note: For the IXPE flight GEMs the gain increases by about 3% per V, and doubles every 25 V.

GPD INVENTORY

Below is a succinct table summarizing the GPD inventory for the flight detector units, including the serial identifiers and the filling times. (The latter are relevant for *Secular Pressure Variations*.)

DU #	GPD #	Filling date	GEM #
FM1	GPD 32	Mar 18, 2019	FM3-c2 3-2-4
FM2	GPD 31	Feb 11, 2019	FM3-c1 3-3-5
FM3	GPD 34	May 28, 2019	FM3-c2 2-5-5
FM4	GPD 36	Jul 11, 2019	FM3-c2 3-2-3

For completeness, the following are the spare flight detectors that have been assembled and filled but not used for any of the detector units. They act as a control sample for the long-term monitoring of the GPDs and are referred to several times throughout the document.

DU #	GPD #	Filling date	GEM #
N/A	GPD 30	Jan 24, 2019	FM3-c2 3-2-4
N/A	GPD 33	Apr 04, 2019	Techtra v2.0-4
N/A	GPD 35	Jun 18, 2019	Techtra v2.0-11
N/A	GPD 37	Jul 31, 2019	FM3-c1 3-3-7
N/A	GPD 38	Jan 31, 2020	FM3-c1 1-5-9

And, since we are at it, here are some other non-flight models with some technical and/or historical relevance.

Use	GPD #	Filling date	GEM #
Long-term studies	GPD 22	Sep 18, 2017	Jun17 1450
GPD EM	GPD 23	May 23, 2018	Nov17 4-4-1
DU EM	GPD 28	Nov 22, 2018	Techtra v1.0-2
Spurious modulation studies	GPD 29	Dec 11, 2018	FM3-c1 3-3-8

EVENT RECONSTRUCTION

With ‘reconstruction’ we refer to all the operations necessary to extract, from raw data, the physical information of interest for polarimetry, imaging and spectroscopy.

A fundamental aspect of this process is that it is performed on a event-by-event basis: as such, it is sometimes referred to as the *event reconstruction*.

8.1 Clustering

The very first step of the reconstruction is a fixed-threshold zero suppression, where pixels whose pedestal-subtracted signal q_i is below (or equal to the) threshold are simply masked and ignored in all the subsequent steps. It is important to note that this zero-suppression stage is additional to any eventual zero suppression that may take place on-board. Though this is mostly an implementation detail, we note also that, for the sake of efficiency, we execute the zero suppression together with the clustering stage described below; thus, the zero suppression threshold is often referred to as the ‘clustering threshold’. Such threshold can be expressed either as an absolute PHA values or in terms of standard deviations of noise. An example of an event after the zero suppression stage is shown in Fig. 8.1.

As said, after the zero suppression a clustering stage is executed, with the aim of grouping together all the contiguous pixels above threshold. Among the different possible choices, our implementation of the clustering is based on the DBSCAN algorithm. In DBSCAN, the definition of a cluster relies on the abstract concepts of *core point*, *directly reachable point*, *reachable point* and *density-connected points* (with ‘point’ translating concretely into ‘pixel’ for our purposes).

A *core point* is a point with at least a certain density - that is a certain number - of other points in its neighborhood, with such density being an adjustable parameter (see Fig. 8.3). In the original DBSCAN algorithm the size of the neighborhood is adjustable as well, while in our case the neighborhood of a pixel is simply made by its 6 adjacent pixels on the hexagonal grid, among which we count how many are above the zero suppression threshold. We note that, differently from the original algorithm, we do not count the central pixel itself for the purpose of reaching the density threshold.

A pixel q is said to be *directly reachable* from a pixel p if p is a core pixel and q is in its neighborhood - that is if p and q are adjacent. This relation is not strictly symmetric, because q may not be a core pixel. A pixel q is *reachable* from a pixel p if there exists a path starting from p and reaching q in which every intermediate pixel is directly reachable from the previous one. This means that every pixels along the path must be a core pixel, except possibly q (the last one). Finally, two pixels are said to be *density-connected* if exists at least one third pixel from which they are both reachable. A cluster is a *complete* set of mutually density-connected pixels. Here ‘complete’ means that every pixel which is reachable from a pixel belonging to the cluster must be part of the cluster as well. Non-core pixels form the edge of the cluster: they still have to be above threshold but, differently from core pixels, do not have to satisfy the density requirement.

Technically speaking, the clustering algorithm is able to find an arbitrary number of clusters in a single ROI; most of the time, however, only one of these clusters will contain the actual charge deposited by the x-ray, while the rest are just random noise fluctuations. We thus discard any cluster that does not include at least a certain number n of pixels,




Fig. 8.1: A real zero-suppressed event from a run with an unpolarized ^{55}Fe source. For completeness the entire readout window is shown, with the boundary column and row numbers indicated in the corners, while only the pixels above the zero-suppression threshold (set to 9 ADC counts in this particular case) are displayed in colors. For the latter, the small number within each hexagon represents the corresponding ADC values.

with n currently being set to 6. It is also theoretically possible that the real charge deposit is split into two different clusters. In practice, this is almost never the case for x-rays in the energy range of interest, unless an extremely high clustering threshold is used. Thus, for the vast majority of our purposes, we will take into account only the main cluster, or simply the ‘track’, defined as the one with the highest total pulse height Q .

Note: In the unlikely event of more than one cluster with the exact same pulse height, we pick the one with the largest number of pixels.

From the reconstruction standpoint, a track is essentially a list of pixels, each fully characterized by a position in space (x_i, y_i) and a signal pulse height q_i . Note that, from a purely logical point of view, the words ‘track’ and ‘cluster’ are essentially equivalent in this context, as they refer to the same set of pixels. However, for clarity of exposition, we will adopt the following distinction: tracks are defined in the physical space (XPE coordinates system), while clusters are defined in the logical space (XPOL coordinates system). While the clustering happens in logical space, all the reconstruction downstream happens in physical space, with no further reference to the pixel address (i.e., column and row) within the readout window.




Fig. 8.2: Main cluster corresponding to the event shown in Figure Fig. 8.1. Note that the reconstruction reference system is rotated by 90° with respect to the orientation of the readout ASIC.




Fig. 8.3: Effect of the density parameter on the clustering: two clusters built from the same raw event and using the same zero-suppression threshold, but with density parameter equal to 2 (left) and 4 (right) respectively. Increasing the value removes the unidimensional tails of pixels that spawn from the main cluster.

8.2 Moments analysis

After the clustering stage, the vast majority of the reconstruction flow deals with the calculation of the basic track properties, as detailed in this section.

8.2.1 Pulse Height

The pulse height Q of the cluster is simply the sum of all the signals in the pixels within the cluster

$$Q = \sum_i q_i.$$

(Obvious as it is, here and in the following the index i is running over the list of pixels in the cluster.)

8.2.2 Baricenter calculation

Next, the baricenter (x_b, y_b) of the cluster is calculated as

$$x_b = \frac{1}{Q} \sum_i x_i q_i$$

$$y_b = \frac{1}{Q} \sum_i y_i q_i,$$

as illustrated in Fig. 8.4.

8.2.3 1-dimensional moments along a given axis

The pixel coordinates in a generic reference frame rotated by an angle ϕ (with respect to the original one) around a generic point (x_0, y_0) read

$$x'(\phi) = (x - x_0) \cos \phi + (y - y_0) \sin \phi$$

and the n -th order moment of the corresponding one dimensional charge distribution projected on the axis passing by (x_0, y_0) and oriented at an angle ϕ with respect to the original x -axis is therefore

$$M_n(\phi) = \frac{1}{Q} \sum_i [(x - x_0) \cos \phi + (y - y_0) \sin \phi]^n q_i.$$

8.2.4 Digression: generalized moments analysis

The basic reconstruction flow exploits a moments analysis technique in at least two flavors and we summarize here the relevant formalism.

The basic un-weighted moments analysis provides the direction of the principal axis of the cluster via the maximization of the second moment of the one-dimensional projected charge distribution along a generic direction ϕ

$$M_2(\phi) = \frac{1}{Q} \sum_i [(x - x_b) \cos \phi + (y - y_b) \sin \phi]^2 q_i.$$

The corresponding condition

$$\frac{\partial M_2(\phi)}{\partial \phi} = 0$$

is easily inverted into the algebraic equation

$$\phi_{\max/\min} = -\frac{1}{2} \arctan \left(\frac{2 \sum_i (x - x_b)(y - y_b) q_i}{\sum_i [(y - y_b)^2 - (x - x_b)^2] q_i} \right),$$

which has two distinct solutions $\phi_{\max}$ and $\phi_{\min}$ in the interval $[-\pi/2, \pi/2]$. $\phi_{\max}$ and $\phi_{\min}$ differ by $\pi/2$ and identify the two directions in which the variance of the projected one-dimensional charge distribution is maximum of minimum, respectively. Equivalently, $\phi_{\max}$ and $\phi_{\min}$ indicate the directions of the principal axes of the ellipsoid of inertia of the two-dimensional charge distribution. The longitudinal and transverse moments of the latter can be readily evaluated

$$M_2^{\text{long}} = M_2(\phi_{\max})$$

$$M_2^{\text{trans}} = M_2(\phi_{\min})$$

and they are essentially the square of the of the major and minor semi-axis of the aforementioned ellipsoid of inertia. For completeness, their ratio related to the eccentricity e of such ellipsoid by the relation

$$\varepsilon = \frac{M_2^{\text{long}}}{M_2^{\text{trans}}} = \frac{1}{\sqrt{1-e}}.$$

Fig. 8.4 illustrates the concept.

This basic algorithm can be generalized by using a generic *pivot* point (x_p, y_p) instead of the baricenter and introducing a set of multiplicative weights w_i (matching the length of the pixel vector). The formulae above are modified through

$$\begin{aligned} (x_b, y_b) &\rightarrow (x_p, y_p) \\ q_i &\rightarrow q_i w_i \\ Q &\rightarrow \sum_i q_i w_i \end{aligned}$$

and the weighted moments analysis can in principle produce a large variety of different interesting outputs, depending on the choice of the pivot and the weights.

8.2.5 First-pass moments analysis

In the first part of the reconstruction an un-weighted moments analysis is performed using the baricenter of the cluster as the pivot point, as illustrated in Fig. 8.4. This provides the angle ϕ_{max} of the principal axis of the cluster and the longitudinal and transverse moments M_2^{long} and M_2^{trans} of the charge distribution.

8.2.6 Third moment and track skewness

Next we project the charge distribution on the principal axis of the cluster, using the baricenter as the origin of the reference system, and calculate the skewness of the resulting one-dimensional probability density function

$$M_3^{\text{long}} = \frac{1}{Q} \sum_i [(x - x_b) \cos \phi_{\text{max}} + (y - y_b) \sin \phi_{\text{max}}]^3 q_i.$$

Needless to say, M_3^{long} has the physical dimensions of a length cubed. Although, for our purposes, most of the information is encoded in its sign, as we shall see in a second, there's at least two interesting quantities that one can construct starting from M_3^{long} . The first is simply its cubic root, which has the physical dimension of a length, and can be represented on the longitudinal track profile, as illustrated in Fig. 8.5. The second is the *skewness* of the one-dimensional charge distribution projected on the principal axis of the cluster

$$\gamma_1 = \frac{M_3^{\text{long}}}{(M_2^{\text{long}})^{3/2}}$$

which is essentially the third moment normalized to the power $3/2$ of the major semi-axis of the ellipsoid of inertia, and is therefore a pure number (i.e., it carries no physical units).

Since the longitudinal development of the ionization is on average asymmetric, with the largest signal residing in the Bragg peak at the end of the track, the sign of the skewness is a solid indicator of the direction in which the original photon conversion point is lying, relative to the baricenter of the charge distribution.

8.2.7 Second-pass moments analysis

We exploit the knowledge of the sign of the skewness to define a search region for finding the photon conversion point. The search region is defined by selecting all those pixels whose centers are located




Fig. 8.4: First moments analysis for the sample event shown in Fig. 8.1.




Fig. 8.5: One-dimensional projection of the charge distribution along the principal axis of the cluster and corresponding skewness for the event shown in Fig. 8.1. The Bragg peak on the right, along with a smaller peak due to the Auger electron on the opposite side, are clearly visible, and the sign of the skewness—negative in this case—is indicating the location of the start of the track relative to the baricenter.

- at a distance $r_{\min} < d < r_{\max}$ from the baricenter of the full cluster;
- on the side of the baricenter of the full cluster indicated by the sign of the skewness.

As illustrated in Fig. 8.4 this is effectively defining a half-ring where the pixel closest to the photon conversion point are located. $r_{\min}$ and $r_{\max}$ are formally free parameters for the algorithm, but in fact we empirically set them to

$$\begin{aligned} r_{\min} &= 1.5 \times M_2^{\text{long}} \\ r_{\max} &= 3.5 \times M_2^{\text{Tms}} \end{aligned}$$

We then proceed to calculate the barycenter (x_c, y_c) of all the pixels in the search region. We use this point to define a set of exponentially decreasing weights for all the pixels in the clusters (both inside and outside the search region) based on their distance d from it

$$w(d) = e^{-\frac{d}{d_0}},$$

where d_0 is typically set to the pitch of the readout plane, or $50 \mu\text{m}$. More precisely, the value w_i of the weight function is calculated for the i -th pixel, and a new baricenter is estimated

$$\begin{aligned} x'_b &= \frac{1}{Q} \sum_i x_i q_i w_i \\ y'_b &= \frac{1}{Q} \sum_i y_i q_i w_i, \end{aligned}$$

which is used as our best estimate for the photon conversion point. Eventually, the conversion point is used as the pivot for a second iteration of the moments analysis, this time weighted with the above defined weights w_i . The direction of the principal axis in this new moments analysis is used as the final photoelectron direction.




Fig. 8.6: Identification of the photon conversion point for the sample event shown in Fig. 8.1. The white (i.e., non-hatched) half-ring defines the pixel selection region for the calculation of the new baricenter.




Fig. 8.7: Second pass moments analysis for the sample event shown in Fig. 8.1.

Note: In the legacy reconstruction software, *Pixy*, the best estimate of the conversion point used to be simply the unweighted barycenter of the search region.

MONTE CARLO SIMULATION

THE ROLE OF THE GAS PRESSURE

The gas filling pressure affects the GPD performance and the characteristics of the photoelectron tracks in several different ways, and this section includes a concise summary of the basic facts about this very topic.

Note: The filling pressure is nominally 800 mbar at the reference temperature of 20 degrees C, and is measured to within a few % at the filling facility. Since it has been suggested that the pressure might undergo changes after the sealing of the detector (e.g., through anomalous absorption by the ceramic walls of the gas cell), the following of this section focuses on what means we have to detect and quantify any such possible change experimentally.

10.1 Absorption Efficiency

Fig. 10.1 shows the impact on the gas absorption efficiency of the filling pressure, relative to the nominal filling pressure of 800 mbar. At high energy, in the optically-thin regime, the efficiency reduction is approximately equal to the reduction in the pressure.

10.2 GEM Gain

The scaling of the GEM gain as a function of the gas pressure, all other conditions unchanged, is not trivial to calculate.

The scaling used in [Tamagawa et al. \(2009\)](#)

$$G(P, T) = G(T_0, P_0) \exp \left\{ C \left(\frac{T}{P} - \frac{T_0}{P_0} \right) \right\}$$

was customarily used to correct the gain during the incoming tests of the GEM (in a test setup with a gas system) for the environmental conditions around the atmospheric pressure, but is not suited to scale the gain much far from 1 atm.

On the other hand, the scaling used in [Takeuchi et al. \(2014\)](#), where the first Townsend coefficient

$$\alpha = \frac{\ln G}{x}$$

depends on the pressure

$$\frac{\alpha}{P} = A \exp \left\{ -B \frac{P}{E} \right\}$$

with A and B being two constants depending on the geometry of the detector and the gas mixture, and E is the GEM electric field.

We have three data points to anchor the aforementioned model to the IXPE detector, namely:




Fig. 10.1: Absorption efficiency of 1 cm of DME at different pressures, relative to the nominal value of 800 mbar.

- the flight GPD production, filled at 800 mbar;
- GPD 27, filled at 400 mbar for the purpose of studying the spurious modulation;
- a *dry* GPD assembly, equipped with a DME flow system at 1 atm.

Figure Fig. 10.2 show a fit to these three data points, after the corresponding gain values have been rescaled to a common voltage working point, using the measured GEM gain characteristics.

See also:

The Gas Electron Multiplier

As we see, the two models are in reasonable agreement close to the atmospheric pressure, but Tamagawa et al. (2009) grossly overestimate the GEM gain at low pressure. Assuming the best fit model to our data using Takeuchi et al. (2014), the following table include a few interesting values for reference.

Pressure [mbar]	Relative gain
800	1.00
775	1.20
750	1.49
725	1.82
700	2.25
675	2.77

From a purely phenomenological point of view, the scaling of the gain vs. pressure is approximately exponential (at least locally), and the

$$G(p) = \exp \left\{ -\frac{(p - p_{\text{ref}})}{p_{\text{scale}}} \right\} \quad \text{with} \quad p_{\text{scale}} = 123 \text{ mbar}$$




Fig. 10.2: Relative gain in pure DME for the IXPE GEM foils, rescaled to a common working point, measured for three configurations (400 mbar, 800 mbar and 1 atm). The normalization is conventionally set to 1 at the nominal pressure of 800 mV. The lines are the two models described in the text, that are in good agreement at near atmospheric pressure.

describes the full model to within 1% (around the reference pressure of 800 mbar) between 650 and 850 mbar. This formula can be used effectively for quick calculations.

10.3 Track Properties

Even assuming that the GEM gain remains unchanged (e.g., if the GEM voltage is properly rescaled), any change in the gas filling pressure impacts all the track properties through two distinct mechanisms:

- the photoelectron range increases with the pressure;
- the transverse diffusion coefficient decreases as

$$\sigma_t \propto \frac{1}{\sqrt{P}}.$$

In practical terms this means that, as the gas pressure decreases:

- tracks interest more pixels;
- tracks are longer;
- tracks are wider (i.e., more extended in the transverse direction).

This is illustrated in the following figures, compiled using monochromatic simulations at discrete pressure values. (Note that in all cases we have selected events in the photo-peak, in order to make any comparison with real data easier.)




Fig. 10.3: Average reconstructed track size as a function of the filling pressure at 2.7 and 5.9 keV. At 5.9 keV tracks in DME at 700 mbar interest 10% more pixels than at 800 mbar, on average.




Fig. 10.4: Average reconstructed track length as a function of the filling pressure at 2.7 and 5.9 keV. For the purpose of the plot the track length is defined in terms of the longitudinal second moment M_2^l as $l = 4\sqrt{M_2^l}$. The increase when going from 800 mbar to 700 mbar is slightly in excess of 10% at 5.9 keV.




Fig. 10.5: Average reconstructed track width as a function of the filling pressure at 2.7 and 5.9 keV. For the purpose of the plot the track width is defined in terms of the transverse second moment M_2^t as $t = \sqrt{M_2^t}$.

SECULAR PRESSURE VARIATIONS

With the expression *secular pressure variation* we indicate the characteristic trend of pressure inside the gas cell for the IXPE gas pixels detectors to decrease with time, with typical time-scales of months. This section is devoted to a quantitative analysis of the phenomenon and its implications on the calibration, operation and scientific performance of the detectors.

11.1 Short Summary and Basic Facts

We have compelling evidence that the pressure in the GPD gas cell decreases with time after the filling, over time scales of months, with an overall asymptotic deficit of about 100–150 mbar, compared to the nominal 800 mbar at filling time.

While we have no direct pressure measurements to support this claim, we have several independent indirect indications making up for a fairly consistent picture, namely:

- the track length tends to increase with time;
- the gas gain, at the same operating voltage, tends to increase with time;
- the absolute quantum efficiency, or the relative event rate (measured in a reproducible test setup) tends to decrease with time;
- the modulation factor tends to increase with time.

In addition, we have direct metrological measurements of the vertical displacement of the GPD beryllium window for most of the detectors, quantitatively supporting the evidence for a pressure decrease inside the gas cell.

While the root cause of the effect, at this time, is unknown, it is clear that the phenomenon is not due to a real leak between the gas cell and the external world, as the energy resolution of all the flight detectors does not show any sign of worsening with time. (In addition, a real leak would be at odd with all the results of the leak tests performed at various stages of the GPD assembly). Instead, the scenario of internal DME ab/adsorption by some of the material in the gas cell seems to be the most natural one.

Finally, we have clear indications that the effect is mostly physical (as opposed to chemical) in nature, and that it is, at least partially, reversible.

What follows is meant to be a comprehensive review of the available data, of the basing modeling strategy, and of the impacts on the high-level Science performance of the mission.

11.2 Historical Note

Historically, the first indirect hint of this very phenomenon was a slow change with time of the average gain of the detectors, that was observed (although with significant differences in terms of the magnitude of the effect) as early

as a decade ago, during the R&D program leading to the mission proposal, and then through the IXPE breadboard development phase, as well as in all the flight models.

This *secular gain variation*, was extensively studied and characterized, but not directly linked to a change in the internal pressure of the gas cell until about February 2020, when the connection of this phenomenon with a deficit of quantum efficiency measured during the calibration of flight detector units was finally made. Following the first quantitative analyses, over the following couple of months, we have accumulated an overwhelming evidence that a pressure change is the root cause of the secular gain variations.

Note: This section is organized in strict *logical* (rather than *chronological*) order, in an attempt of providing a coherent interpretative framework. For this reason, the presentation is much less gain-centric than the one historically used for much of the material accumulated and presented during the investigation phase.

11.3 Proxies for the Pressure

In absence of direct measurements we use mainly three proxies for inferring the internal pressure of the gas cell: the (absolute or relative) quantum efficiency, the average track length and the gas gain (mostly at 5.9 keV).

See also:

The Role of the Gas Pressure.

The quantum efficiency is, in principle, the most straightforward measurement of the pressure. Measurements of the *absolute* quantum efficiency are somewhat complex to perform, and are typically systematic-limited, which makes it impractical to use them as a mean for a fine, long-term monitoring of the detector performance. (On the other hand, they are crucial as an independent cross check.) The quantum efficiency can be measured in *relative* terms, e.g., by measuring the event rate with a reference radioactive source in a standard holder. This is a much simpler measurement to do and has been indeed systematically exploited in order to investigate the secular pressure variations.

Any change of the pressure in the gas cell will cause changes in the track topology, e.g., through changes in the electron range and transverse diffusion. In principle there are several (not independent) topological track quantities that can be used for the purpose. Among them the track length (i.e., the square root of the second moment along the main axis of the track) is the most robust against changes in the environmental conditions and the data acquisition settings.

Warning: Inferring the pressure from the track length requires input from the Geant 4 Monte Carlo simulation. Since the simulation itself is not perfect, we typically have to allow for a multiplicative normalization factor for the Monte Carlo predictions of the order of 1 (and typically within a few %). We emphasize that most of the fine tuning of our Geant 4 model was performed before we discovered the secular pressure variations, i.e., by trying and matching simulations performed at 800 mbar with calibration data presumably taken at a smaller pressure—therefore it is natural to expect that the current version of the simulation would tend to underestimate the pressure.

It follows that, until a comprehensive revision of the physical parameters of the simulation is performed, any attempt to derive the internal pressure from the vanilla Monte Carlo output on a point-by-point basis is moot.

The gas gain depends on the gas pressure at a fixed composition, and is a third, independent proxy that we have customarily used to investigate the secular variations. We emphasize that, compared to the former two quantities, the gain is somewhat more difficult to use, as its dependence on the pressure is complex, and the pulse-height must be properly re-scaled to a common high-voltage working point in order to compare different detectors. In addition, the gain is known to show other kind of variations (e.g., due to the GEM charging, or induced by changes in the environmental conditions) that need to be carefully controlled.

11.3.1 Systematic Errors

Even with a moderately long data acquisition (say a few tens of thousands of events) any of the aforementioned pressure proxies will be measured to a statistical accuracy well below the systematic uncertainties due to the limited reproducibility of the setup, e.g., the accuracy with which the radioactive source and holder are places on the titanium frame of the GPD.

Generally speaking, initial tests specifically performed on the flight control sample of GPDs to address this very question, indicate that the average track length and the event rate can be measured to an accuracy of 0.5%. The gain, on the other hand, is more sensitive to the environmental conditions, and is hard to measure to better than some 2%. We shall come back on this in a moment, while describing our combined fit procedure.

11.4 Basic Modeling

In the simplest possible model, we shall assume that the pressure changes with time with according to an exponential law:

$$p(t; \tau, \Delta_p) = p_0 - \Delta_p(1 - e^{-(t-t_0)/\tau}).$$

This is a two-parameter model: the characteristic time constant of the exponential and the asymptotic pressure loss. Time is naturally referred to the GPD filling date, and (although a more refined modeling is possible) the corresponding pressure can, to first order, be assumed to be the nominal 800 mbar value.

From a purely phenomenological point of view, the scaling of the average track length with the pressure is reasonably well described by a power law

$$L(p) = L_0 \left(\frac{p}{p_0} \right)^{-\alpha_L}$$

while that of the gain vs. pressure is, at least locally, approximately exponential

$$G(p) = G(p_0) \exp \left\{ -\frac{(p - p_0)}{p_{\text{scale}}} \right\}.$$

For completeness, the quantum efficiency of the detector (which can be inferred, in relative terms, by means of rate measurements in a reference setup) is proportional to the pressure and therefore varies with the very same time dependence.

It follows that we are in the interesting situation where our three proxies of the internal gas pressure vary in time each one with their own peculiar scaling:

$$\begin{aligned} L(t; \tau, \Delta_p) &= L_0 \left(\frac{p(t)}{p_0} \right)^{-\alpha_L} \\ G(t; \tau, \Delta_p) &= G(p_0) \exp \left\{ -\frac{(p(t) - p_0)}{p_{\text{scale}}} \right\} \\ Q(t; \tau, \Delta_p) &\propto p(t) \end{aligned}$$

It is worth emphasizing that the scalings for the track length and the gain, rigorously speaking, are *not* simple exponential functions. The situation is, in general, fairly complicated, but it is instructive to examine the limiting case in which the time constant of the process is infinite and the pressure in fact scales linearly with time

$$p(t; m) = p_0 - m(t - t_0).$$

In this situation the time dependence of the gain simplifies in

$$G(t; m) = G(p_0) \exp \left\{ \frac{m(t - t_0)}{p_{\text{scale}}} \right\},$$

i.e., it is a pure *increasing* exponential.

Note: In a model where the pressure decreases linearly with time, we expect the gain to increase exponentially with time. In other words: in this scenario the gain increase is *not* linear, and shows a characteristic (positive) concavity.

11.4.1 Model scalings

Before we move forward it is instructive to take a quick look at the relevant scaling with time of the quantities we use to infer the pressure. Fig. 11.1 shows a basic toy test setup where we let the pressure drop exponentially by 150 mbar over different time scales, ranging from 1 month to 2 years.




Fig. 11.1: Time dependence of the pressure for a toy test setup where the pressure itself drops exponentially by 150 mbar over different time scales, ranging from 1 month to 2 years.

Calculating the average track length as a function of time is a trivial matter of function composition and can be simply done analytically, as shown in Fig. 11.2 at the reference energy of 5.9 keV.

As we said, the result is not rigorously exponential, but we empirically find that it is approximately so: the dotted lines (barely visible) in the figure represent the best-fit simple exponential for each of the different values of τ . Interestingly, this *effective* time constant is consistently larger than that of the underlying physical process, and more markedly so when the baseline for the measurement is short with respect to the time constant itself: over 18 months the relative difference is only 10% for a τ of one month, rising to some 35% for a τ of two years.

Things get even more interesting with the gain, as shown in Fig. 11.3. The situation is qualitatively similar to that of the track length, except for the fact that the now the effective time scale for the gain change can be up to an order of magnitude larger than that of the underlying pressure change. (This is a basic consequence of the different scalings with the pressure: a power law for the track length and an exponential for the gain.)




Fig. 11.2: Average track length at 5.9 keV as a function of time for the very same toy setup shown in Fig. 11.1.

Additionally a new spurious phenomenon emerges: generally speaking the exponential fit to exact gain scaling is fairly poor (up to several tens of %) near the origin, as shown in the inset. (While this might seem an obscure technical detail, this was in fact one of the main issues in early attempts to fit to the gain profile, extrapolating the inferred pressure to the filling time, explaining a significant part of the DU FM4 mystery.)

Note: In a setup where the pressure, assumed to change exponentially with time, is inferred by proxies such as the track length, the gain and the quantum efficiency, one should be well aware of the different scalings with time of the quantities involved:

- the quantum efficiency (and, as a consequence, the event rate, when measured in a reproducible setup) scales exactly like the pressure;
- the track length scales approximately exponentially, with a time constant that is generally longer than the underlying pressure evolution;
- the gain scales in a way that somewhat resembles an exponential, but the effective associated time constant, depending on the particular setup and boundary conditions, can differ by factors, or even orders of magnitude, from the time constant of the underlying physical process.

We emphasize again that if the gain changes linearly with time, this does not imply that the pressure decrease is not saturating.

11.5 Fitting

In order to infer the time evolution of the pressure, we perform a combined fit of all the available measurements of our proxies, i.e., the track length, the gain, and the rate measurement in a reference setup where available. More




Fig. 11.3: Gain as a function of time for the very same toy setup shown in Fig. 11.1..

specifically we minimize the quantity:

$$\chi^2(\tau, \Delta_p, N_L, N_G, N_q) = \sum_{i=1}^{n_L} \frac{(L_i - L(ti; \tau, \Delta_p))^2}{\sigma_{L_i}^2 + \sigma_{L_{sys}}^2} + \sum_{j=1}^{n_G} \frac{(G_j - G(tj; \tau, \Delta_p))^2}{\sigma_{G_j}^2 + \sigma_{G_{sys}}^2} + \sum_{k=1}^{n_Q} \frac{(Q_k - Q(tk; \tau, \Delta_p))^2}{\sigma_{Q_k}^2 + \sigma_{Q_{sys}}^2}.$$

As stated earlier on, the baseline model assumes that the pressure at the filling time coincides with the nominal value of 800 mbar, and the time dependence is parametrized in terms of two numbers: the characteristic time of the exponential and the overall asymptotic pressure decrease.

Being most of the measurement systematic-limited, a suitable choice of the scale for the systematic uncertainties on the track length, gain and quantum efficiency/rate measurements provides an effective mean to control the weight of each component in the combined fit.

11.5.1 Fitting By-Products

We fit up to three auxiliary parameters in the process—namely three multiplicative normalization factors for the proxies involved. These provide additional useful information and provide a mean to verify the internal consistency of the model. Specifically:

- the normalization factor for the track length is necessary to account for imperfections in the Monte Carlo simulation of the detector, and should be of the order of unity (seen from a different prospective, this can be fed back into the simulation to fine tune the underlying physical parameters); in the limit where all the detectors are identical and operated under the same conditions, the inferred normalization factor should be the same for all the GPDs within the experimental uncertainties—any significant difference might point to room for improvement in the modeling;

- the normalization factor for the gain is largely irrelevant, as it is determined by our choice of the GEM HV working point; differences in the fitted values across different detectors reflect possible small differences in the gain-voltage characteristics of the GEMs that are not critical for the secular pressure variations;
- the normalization factor for the rate measurements are typically undetermined *a priori*, but a comparison of the best-fit values between different detectors *tested in the same setup* provide useful information about possible intrinsic differences.

11.6 Summary of the Results

In this section we provide a top-level summary of the results of the modeling and fitting described above.

In order to illustrate the process, we focus on the GPDs in the *flight control sample*, i.e., those detectors assembled during the flight production that were not selected to be installed in one of the four detector units—the main reasons being that they have been regularly monitored for the specific purpose of characterizing the secular pressure variations, and therefore provide the richest data sets. Particularly, Fig. 11.4 shows the basic plot for one of the control GPDs.

It goes without saying that we are primarily concerned about the behavior of the detectors installed on the detector units, and a complete compilation of the summary plots can be found further below at the end of this section.

The best-fit parameters for the flight GPDs are summarized in Fig. 11.5.

Note: The estimated asymptotic residual pressure is typically between 630 mbar and 660 mbar, with a characteristic time constant of the process between 6 months and a year. (Note that GPD 35 is the detector that has been heated at high temperature for two days, and the corresponding fit has been performed by excluding the data points around the experiment.)

We emphasize that, while there is a significant variability across different detectors, especially in terms of the timescale of the process, there are no clear outliers, as a naive exponential fit to the gain curves would erroneously suggest.

11.6.1 Evolution Convexity and Saturation of the Effect

One of the very fundamental questions is whether we have solid evidence that the pressure drops saturates for all the GPDs.

Figure Fig. 11.6 shows the result of a linear fit to the secular data for GPD 30—the one with the longest time constant among the entire flight production. While, at a first glance, the gain increase with time is approximately linear, the flattening of the effect is clearly visible in the track length. (It is unfortunate that we don't have rate data for the entire time span for this detector.) We emphasize one more time that, should the pressure decrease linearly with time, we would expect the gain to show a definite concavity, as evident in the third panel from the top.

The overall chisquare for the linear fit, compared with that for the exponential fit shown in Fig. 11.4, corresponds to a preference for the second by a statistical significance of more than 40 sigma.

Although the measured statistical significance with which the exponential model is preferred over the linear model is not an intrinsic property of any given detector (as it also depends on the temporal baseline that we have available for the measurement, as well as the regularity of the sampling) it is instructive to look at this very metric for all the detector, as shown in Fig. 11.7.

Note: Even for the GPDs with the slowest evolution within the entire flight production, the convexity of the pressure variation with time is well established, and the exponential model is favored over the linear one with a high statistical significance for all the detectors, with no exceptions. In other words we have compelling evidence for all the detectors that the pressure is *not* decreasing linearly with time.




Fig. 11.4: Secular pressure variation summary plots for GPD 30. The blue line on the top is the pressure time profile calculated with the combined fit, while the three orange lines are the corresponding, derived models for the three measured proxies. The gain measurements, acquired at ~fixed average pulse-height, have been re-scaled to a common GEM working point. The measured event rates have been corrected for the radioactive source decay.




Fig. 11.5: Top-level summary of the best-fit parameters for the pressure models for a compilation of detectors, including the three flight detector units and the control sample.

Given that our is a physically-motivated model that seems to fit reasonably well with a handful of parameters hundreds of data points form a large number of different detectors, it is fair to say that also the fact that the pressure in the gas cell is asymptotically reaching a stable value of of about 650 mbar is well supported by the data at hand.

11.7 Complementary Evidence

We do have an independent, quantitative evidence that the pressure in the gas cell is decreasing through the standard detector metrology that is performed to support the alignment of the detector units with the mirror-module assemblies in the integration stage.

Prior to filling, the vertical position of the center of the Be window is measured for each detector with respect to bottom plane of the titanium frame holding it. In this configuration the differential pressure between the two sides of window is zero, and therefore the window itself is, in principle, perfectly aligned with the titanium bottom plane. This first measurement constitutes the zero for the following ones.

After the gas filling, the window is subjected to a ~200 mbar differential pressure (that, as we have seen, is changing with time), which causes a measurable movement of the vertical position of the window center. Based on FEM simulations and actual tests, this shift is expected to be linear with the differential pressure up several hundreds mbar:

$$\Delta z \propto (p_{\text{out}} - p_{\text{in}}).$$

It is therefore possible, at least in principle, to gauge the internal pressure in the gas cell from a direct metrological measurement. In practice this is achieved with a CMM with an optical head—in a nutshell: four reference points are measured on the four outer angles of the window to define the reference plane, and the vertical position of the center




Fig. 11.6: Linear fit to the secular pressure variation for GPD 30—this should be compared directly with Fig. 11.4.




Fig. 11.7: Top-level summary of the measured statistical significance with which the exponential model for the pressure change is preferred over the linear model. For the vast majority of the detectors this is more than 20 sigma; even for those with the slowest evolution (e.g., GPD 29) and/or the shortest temporal baseline for the measurement (e.g., GPD 37) it is larger than 30 sigma, and increasing with time.

is then measured with respect to that. The nominal reproducibility of the CMM positioning is 2–3 μm , but the actual error on our measurement is somewhat larger, due to the uncertainties in the manual positioning of the head, as well as the systematic uncertainty on the pre-fill zero-measurement. In the following we shall assume an error of 5 μm .

In Fig. 11.8 the window displacement measurement is shown, for each of the GPD in the flight control sample, and at different moments in time, as a function of the pressure in the gas cell, estimated from the combined fit described above. Different points in the plot with the same color track directly the pressure evolution of any given detector. Overall, the data points are consistent, as expected, with a straight line, for an overall excursion of 50 μm , over an estimated maximum pressure range in excess of 100 mbar. We emphasize that the intercept of the best-fit straight line is consistent with 0, as predicted by our naive linear model.




Fig. 11.8: Compilation of window displacement measurements for the GPD in the flight control sample, at different moments in time, as a function of the pressure in the gas cell, estimated from our combined fit. The excellent linearity provides quantitative support for the goodness of our the entire fitting process.

11.8 Impact on High-Level Scientific Performance

When discussing the impact of the pressure variations on the polarimetric sensitivity we should keep in mind that there are two competing effect at work:

- the quantum efficiency of the detector is, to first order, proportional to the pressure, and therefore the overall effective area of the telescopes decreases as the pressure decreases;
- tracks are more elongated at relatively low pressure, and therefore the modulation factor of the GPDs increases as the pressure decreases.

The polarization sensitivity (i.e., the minimum detectable polarization for a given source spectrum and integration time over the nominal energy band) scales linearly with the modulation factor and with the inverse of the square root of the effective area. As we shall see, when one folds all these factors in, the net result is that the pressure dependence of the MDP is fairly mild.

The pressure dependence of the modulation factor can be gauged by means of the full Geant 4 detector simulation, as illustrated in Fig. 11.9. This is quantitatively supported by the analysis of calibration data. As a rule of thumb, between 2 and 8 keV, and over the 600–800 mbar pressure range, the modulation factor scales linearly with the pressure, increasing by 0.015 (absolute) for every 100 mbar of pressure decrease. Just to illustrate this basic scaling: if the nominal modulation factor is 0.3 at a given energy (this particular value being representative of the peak of the IXPE sensitivity), the actual value increases to 0.315 going from 800 to 700 mbar, and the corresponding 5% (relative) increase is almost entirely compensating the (square root of) 13% decrease in effective area.




Fig. 11.9: Compilation of the pressure dependence of the modulation factor at different energies, as predicted by the Geant 4 Monte Carlo simulation of the detector.

The argument can be made rigorous by integrating the relevant quantities over the IXPE energy range and calculating the actual broadband MDP for different values of the pressure, as illustrated in Fig. 11.10. The scaling is non linear, but the typical degradation of the MDP as the pressure decreases is about 0.15% (relative) every 10 mbar of pressure loss, or less than 3% when going from 800 to 660 mbar.

Note: As the GPD internal pressure decreases, the loss of effective area is largely compensated by the increase of the modulation factor, the latter running at about 0.015 (absolute) every 100 mbar. Accordingly, the MDP depends only mildly on the pressure—about 0.15% (relative) every 10 mbar.

11.9 Appendix: Summary Plots for all Detectors

We collect here for completeness all the summary plots for all the detectors we have useful data from.







Fig. 11.10: Relative scaling of the MDP as a function of the GPD internal pressure.


11.9.1 Flight production


11.9.2 Prototypes and Qualification Models


























CHARGING

By the term *charging* or *charge build up* we generically refer to the fact that, when the GPD is irradiated, part of the charge from the avalanche in the GEM can be temporarily deposited onto specific regions of the dielectric substrate, which in turns tend to modify the configuration of the electric field in the holes, potentially causing measurable changes in the gain.

This effect is typical of many classes of gas detector, and in the case of the IXPE GPDs is exacerbated by the fine pitch of the GEM, which makes the ratio between the hole volume and the exposed surface particularly unfavorable. The charging has been extensively studied and characterized through the mission phase A and, in parallel, a vigorous program for mitigating the effects by fine-tuning the GEM production process was put in place and pursued.

The basic features of the charging-induced temporal gain variations *for the (SciEnergy) flight GEMs* are as follows:

- the GEM gain tends to decrease when the GPD is irradiated at a sufficiently high rate;
- the effect scales with the energy flux per unit area;
- the typical time scale for the gain decrease (charging) decreases with increasing input energy flux;
- the typical time scale for recovery (discharging), in absence of radiation, is constant and significantly longer than the charging time scale, even at a moderate input energy flux;
- the overall fractional gain decrease increases with the input energy flux up to a maximum asymptotic value that is characteristic of the specific GEM;
- the effect is qualitatively similar across different GPDs, but each one will require its specific calibration constants.

Note: As far as the charging-induced gain variations, the dithering of the observatory is beneficial, as it reduces significantly the average energy flux per unit area. Nonetheless, the effect will be measurable for IXPE observations of bright sources, and therefore we need to be fully equipped to model and correct it.

12.1 The Basic Equation

At the very fundamental level any charging effect can be described with a capacitor-like differential equation

$$\frac{dq(t)}{dt} = \alpha q(t) \quad (12.1)$$

whose solution is a simple exponential (either increasing or decreasing depending on the sign of alpha).

In our case this has to be modified in order to incorporate the basic features that we listed above, and precisely:

- we need two separate terms, one for the charging and one for the discharging;

- the dependence on the input energy flux must be incorporated in the charging term;
- the maximum accumulated charge must be incorporated in the charging term, as well.

That all said, the simplest equation that we can possibly write is

$$\frac{dq(t)}{dt} = \overbrace{R\alpha_c \left(1 - \frac{q(t)}{q_{\max}}\right)}^{\text{charge}} - \overbrace{\alpha_d q(t)}^{\text{discharge}},$$

where:

- $q(t)$ is the accumulated charge per unit area as a function of time, measured in fC mm^{-2} ;
- R is the input charge flux per unit area of the source, measured in $\text{fC s}^{-1} \text{ mm}^{-2}$;
- α_c is an (adimensional) charging constant;
- α_d is a discharge constant, with the physical dimensions of a rate.

When R is constant the model can be solved analytically in closed form. If only the discharge was at play, the accumulated charge would reach zero asymptotically (for any boundary condition) with a time constant of

$$\tau_d = \frac{1}{\alpha_d} \quad \text{and} \quad \lim_{t \rightarrow \infty} q(t) = 0.$$

On the other hand, if only the charging is at play, the accumulated charge tends to $q_{\max}$ with a different time constant

$$\tau_c = \frac{q_{\max}}{R\alpha_c} \quad \text{and} \quad \lim_{t \rightarrow \infty} q(t) = q_{\max}.$$

In the general case, when both effects are present, the general solution reads

$$\tau = \left(\frac{1}{\tau_c} + \frac{1}{\tau_d} \right)^{-1} \quad \text{and} \quad \lim_{t \rightarrow \infty} q(t) = q_{\max} \left(\frac{\tau_d}{\tau_c + \tau_d} \right).$$

That is to say: the system tends asymptotically to a partially charged state, with a time constant that is the inverse of the sum of the inverses of the charge and discharge time constants.

12.2 A Simple One-Species Model

To zero order, the *modified-capacitor* differential equation outlined in the previous subsection is adequate for a basic understanding of the charging, and reproduces all the basic phenomenology observed in the data. The only tiny additional piece of information that is needed to make use of it is the link between the accumulated charge and the gain variation, the reason behind this being that we only observe and measure the gain, never the charge directly.

In the simplest possible picture in which the change in the effective electric field is proportional to the accumulated charge, and the fractional variation of the electric field itself is small (what we're looking at is really a relative gain variation of 10%, which is a variation of a few V in terms of effective potential on the GEM, or less than 1% at the 450–500 V working point), we can approximate the exponential dependence of the gain on the voltage with a linear relation and assume that

$$\Delta g(t) = \frac{G(t)}{G_0} - 1 = -\beta q(t) = -\delta_{\max} \frac{q(t)}{q_{\max}}$$

(Here $\Delta g(t)$ is the *variation of the relative* gain, with zero corresponding to the unperturbed state of the system, where there is no accumulated charge and $G(t) = G_0$. Note that in our model $\Delta g(t)$ is always negative, i.e., the charging systematically lowers the gain.)

In recasting our model in terms of observable quantities, we shall also replace the charge flux with the energy flux for the source term, which is possible since the charge produced in the avalanche is proportional to the original photon energy. We therefore have

$$\frac{d\Delta g(t)}{dt} = -\frac{\delta_{\max}}{q_{\max}} \frac{dq(t)}{dt} = -\overbrace{\frac{F}{k_c} (\delta_{\max} + \Delta g(t))}^{\text{charge}} - \overbrace{\frac{\Delta g(t)}{\tau_d}}^{\text{discharge}},$$

where:

- F is the input energy flux per unit area of the source, measured in $\text{keV s}^{-1} \text{mm}^{-2}$;
- $\delta_{\max}$ is maximum gain excursion, reached when the accumulated charge saturates to its maximum value $q_{\max}$;
- k_c , measured in keV mm^{-2} , is a charging constant incorporating α_c and $q_{\max}$;
- τ_d is the time constant for the discharging.

This equation is the basis for the numerical solution of the gain evolution with time. The system evolves with an overall time constant that relates to the input energy flux according to

$$\tau(r) = \left(\frac{F}{k_c} + \frac{1}{\tau_d} \right)^{-1}$$

and the asymptotic relative gain variation (in the case of constant input energy flux) is given by

$$\delta(r) = -\frac{\delta_{\max}}{\left(1 + \frac{k_c}{F\tau_d} \right)}.$$

We notice explicitly that there exists a critical energy flux for the input source of radiation

$$F_{\text{critical}} = \frac{k_c}{\tau_d}$$

that separates two fundamentally distinct regimes:

- when $F \gg F_{\text{critical}}$ the system is dominated by the charging processes and $\delta(r) \rightarrow -\delta_{\max}$ with a time constant $\tau(r) \rightarrow k_c/F$;
- when $F \ll F_{\text{critical}}$, on the other hand, the discharge is dominating and $\delta(r) \rightarrow 0$ with a time constant $\tau(r) \rightarrow \tau_d$.
- (And, incidentally, $\delta(F_{\text{critical}}) = -\delta_{\max}/2$.)

We emphasize that one of the main post-dictions of the model is that the charging is faster at high rate—and, more precisely, the charging time constant is inversely proportional to the input energy flux.

From a practical standpoint, the three model parameters can be unambiguously determined from a series of calibrations at different energy fluxes. The following table summarizes some typical values for the figures of interest.

Parameter	Typical value
$\delta_{\max}$	~ 0.1
k_c	$\sim 10^5 \text{ keV mm}^{-2}$
τ_d	$\sim 10^5 \text{ s}$
r_{critical}	$\sim 1 \text{ keV mm}^{-2} \text{ s}^{-1}$




Fig. 12.1: Caption goes here.




Fig. 12.2: Caption goes here.




Fig. 12.3: Caption goes here.

12.3 Science Implications

12.4 The Microscopic Picture

Although a complete and self-consistent microscopic modeling of the effect is not trivial, it is worth to examine how the orders of magnitude play out in the overall budget.

Warning: The figures and arguments in this section should be taken *cum grano salis*, as they are largely qualitative and should probably not be trusted beyond the orders of magnitude.

In the following we shall make the following assumptions about the inner working of the detector. The average energy to create an electron-ion pair in the gas is of the order of 30 eV—which means that a 3 keV photon creates some 100 electron-ion pairs of primary ionization. The primary ionization drifts toward the GEM under the influence of the drift electric field, and is multiplied by the GEM itself. While the effective gain of the GEM is of the order of 200, simulations show that a significant fraction of the avalanche charge (up to a factor of 2) is lost on the GEM bottom, so that the gain corresponding to the full charge available in the avalanche (which is relevant for the charging) can be assumed to be of the order of 400. In a nutshell, a 3 keV photon produces an avalanche of about 40,000 electrons, or 6.5 fC.

Parameter	Value	Description
W	~ 30 eV	Average energy per pair
G_{eff}	~ 200	GEM effective gain
$G_{\text{avalanche}}$	~ 400	GEM total avalanche gain
C_{GEM}	~ 250 pF	GEM capacitance
A_{GEM}	~ 250 mm ²	GEM surface
V_0	~ 500 V	Nominal GEM working point

When seen as a parallel-plane capacitor, the GEM at the nominal working point of 500 V accumulates a static charge per unit area of

$$Q_0 = C_{\text{GEM}} V_0 / A_{\text{GEM}} \sim 5 \times 10^5 \text{ fC mm}^{-2}$$

At the critical energy flux of $\sim 1 \text{ keV mm}^{-2} \text{ s}^{-1}$ we expect an asymptotic gain drop of $\delta_{\text{max}}/2 \sim 5\%$. When seen in voltage space, at the typical characteristics of 3% per V, this corresponds to a drop of the GEM working point of just below a couple of V—see the section about the *The Gas Electron Multiplier* for more details—or less than 0.5% of the nominal 500 V.

When seen in charge space, we are looking after a maximum accumulated charge per unit area of the order or

$$q_{\text{max}} \sim 0.005 Q_0 \sim 2.5 \times 10^3 \text{ fC mm}^{-2}$$

The question is whether is reasonable for the GEM avalanche to provide the necessary charge to sustain the process. We can tie our microscopic naive modeling of the phenomenon to the input energy flux by noting that, if we assume for a second that the the system is completely discharged at the time $t = 0$, our basic equation reads

$$\frac{dq(0)}{dt} = R\alpha_c = \frac{q_{\text{max}}}{\tau_c} =$$

In other words, the maximum accumulated charge per unit area q_{max} is related to the derivative of the charge accumulating *at the beginning of the charging process*:

$$q_{\text{max}} = \tau_c \frac{dq(0)}{dt}$$

Now, our critical energy flux of $\sim 1 \text{ keV mm}^{-2} \text{ s}^{-1}$ translates into an avalanche current per unit area of $\sim 2 \text{ fC mm}^{-2} \text{ s}^{-1}$. If *all* the avalanche charge went into charging the GEM, and assuming a charging time constant of

$$\tau_c = \frac{k_c}{F_{\text{critical}}} \sim 10^5 \text{ s}$$

this would imply a maximum accumulated charge of about $\sim 2 \times 10^5 \text{ fC mm}^{-2}$, or almost 50% of the static charge accumulated by the GEM capacitor, i.e., 100 times what we need.

Note: $\sim 1\%$ of the avalanche charge initially going into charging is all we need to reduce the effective GEM potential by 2 V (and the gain by 5%).

12.5 A Realistic Two-Species Model

ALPHA PARTICLES

Tracks of highly ionizing charged particles (most likely alphas) are observed in the GPD with a measurable, although small, rate and typically attributed to alpha decays of radioactive contaminants in the beryllium window.

The realization that single particle tracks could cause a measurable gain and leave a persistent imprinting on the GEM over time-scales of a day or more came originally as a surprise, but fits reasonably well, in fact, with our basic understanding of the charging process in the GEM.

In this section we briefly review the basic facts about alpha particles in the GPD and possible implications for Science observations in orbit.

13.1 Basic Phenomenology

The easiest way to familiarize with this sort of highly ionizing event is to look at single-track displays—there are two of them shown below.

When looked superficially, all these highly-ionizing tracks seem to be fairly similar to each other in terms of the basic characteristics:

- they are typically 5–10 pixels (i.e., a few hundreds μm) wide;
- they are fairly straight;
- they have a total pulse height of the order of a few millions of ADC counts (i.e., 1–1.5 MeV);
- the vast majority of the pixels in the track is saturated.

Note: For completeness: the 1 V dynamic range of the pixel amplifiers in the readout ASIC, when coupled to the IXPE back-end electronics (and, particularly, with the 0.146 mV/ADC counts voltage resolution of our analog to digital conversion) results into a saturation value of about 6500 ADC counts (or just over 2 keV) at the single pixel level. With a pixel area of $2.17 \times 10^{-3} \text{ mm}^2$, this corresponds to an energy density of about 1 MeV mm^{-2} .

These similarities are at odds with the early observation that *not all the highly-ionizing tracks do release a long-lasting footprint on the GEM*. We shall see in fact that, once the saturation is taken into account, significant differences emerge.

13.2 Stopping Power

The stopping power and effective range in the Continuous Slow-Down Approximation (CSDA) can be readily be calculated through the [ASTAR](#) database.




Fig. 13.1: Highly ionizing event in the GPD (real data from GPD 29). We shall refer to this particular event as *track 1* in the rest of this section.




Fig. 13.2: Highly ionizing event in the GPD (real data from GPD 29). We shall refer to this particular event as *track 2* in the rest of this section.




Fig. 13.3: Range as a function of the alpha particle kinetic energy in DME at 0.8 bar.

The range for a 1 MeV alpha is a few mm, reaching 1 cm around 3 MeV. This means, ideally, that 3 MeV is the maximum total energy that a vertical alpha can release in the GPD. For inclined tracks this can be larger by up to roughly a factor of two.

On a related note, the stopping power at the Bragg peak is of the order of 4–5 MeV/cm, which for our typical track width (determined by the transverse diffusion in the gas) is well above the 1 MeV mm⁻² saturation value.

13.3 Correcting the saturation

The basic idea to handle the saturation is fairly simple: fit the *un-saturated* transverse tails of the track and use the fit to correct the *saturated* core.

In reality there is a few subtleties connected with this, but a simplified algorithm to achieve the goal might be sketched as follows:

- align the principal axis of the track with one of the orthogonal axis—say y (this facilitates the projections enormously);
- project the charge content of the hexagonal pixels onto a rectangular grid with the proper scales, e.g., 100 μm on the longitudinal axis and 10 μm on the transverse axis (this can be achieved by brute force via a Monte Carlo re-sampling);
- fit the un-saturated part of the track in each of the transverse slices with a gaussian model (dictated by the physics of charge diffusion in gas);
- use the best-fit model to correct the charge in the core and recalculate the energy.

The procedure is illustrated in the figure above. We emphasize that the correction factors are not small, and in fact quite the opposite—say between 1 and 10 depending on the track and the position along the longitudinal development.




Fig. 13.4: Bragg peak for a 10 MeV alpha particle in DME at 0.8 bar.




Fig. 13.5: Example of a gaussian fit to one of the transverse slices of the track shown in the second event display above. As expected, the saturation takes place at around 1 MeV mm^{-2} .

(This, however, should hardly be surprising in events where 90% of the pixels are saturated.)

That all said, when we look at the final products of the process we find out that the two tracks, essentially indistinguishable in the raw images, actually look quite different after the saturation correction, with the second one featuring a ionization density a factor of 5 larger.




Fig. 13.6: Raw and corrected track image for the track in the first event display shown at the beginning of the section.

13.4 Track Profiles

It is perhaps instructive to look at the *corrected* longitudinal and transverse track profile, and compare with the raw data, as (although fairly noisy) they do provide a direct view into the correction.

Warning: It is not straightforward to reconcile a 9 MeV total energy with a maximum stopping power of 15 MeV/cm with alpha the cross-section data. This might be largely irrelevant in practice, but it is something that we might want to think through in details.

The transverse profile of the track is interesting in that it can, at least in principle, be calculated from first principles. In the basic hypothesis that the alpha track is downward-going, and assuming that it crosses the entire gas cell, from the beryllium window to the readout ASIC, we do expect the track to be narrower in correspondence of the Bragg peak and larger at the beginning—since the primary electrons drift over a larger distance.

More precisely we do expect the transverse dimension of the track to be the sum (in quadrature) of a constant term (representing the effect of the transfer gap) and a term proportional to the square root of the drift length in the absorption gap, which is in turn proportional to the longitudinal coordinate l :

$$\sigma^2(l) = d_t \sigma_t^2 + d_a(l) \sigma_l^2.$$




Fig. 13.7: Raw and corrected track image for the track in the second event display shown at the beginning of the section.




Fig. 13.8: Reconstructed longitudinal profile of the highly-ionizing track shown in the first event display.

This is indeed what it is observed in reality, and the fitted parameters for the two transverse diffusion coefficients are in the right ball-park.




Fig. 13.9: Reconstructed transverse profile of the highly-ionizing track shown in the first event display.

13.5 Alpha particles and Charging

We can now go back to the main item for discussion, namely the connection between alpha particle and charging.

See also:

Charging

We are used to think about the phenomenon as a dynamical process, where two competing effects (charging and discharging) operate at the same time with different time-scales. The basic prediction of the model, i.e., that the time-scale for charging is inversely proportional to the input energy flux, is the key to understand the process in this new context, where we have instead an enormous amount of charge produced instantaneously in the avalanche. In other words, the basic mechanism for this new, single-track charging is

- the avalanche develops *instantaneously*, i.e., the input energy flux is infinite and the charging time constant is zero;
- if the charge is enough to produce an effect, the result is a (small) step function in the gain;
- the discharge starts immediately and proceeds with the usual time constant ($\sim 10^5$ s).

The real question is whether the energy density provided by a low-energy alpha particle is enough or not to charge the GEM significantly, and we have now all the ingredients in our hands to answer this question.

We are observing (after the correction for saturation) up to 25 MeV mm^{-2} of energy density in the Bragg peak of the largest energy deposits. In charge space this translates into 10^6 e mm^{-2} of primary ionization, becoming $4 \times 10^8 \text{ e mm}^{-2}$ (or $6 \times 10^4 \text{ fC mm}^{-2}$) after the multiplication.

When the GEM is seen as a parallel-plane capacitor, the static charge per unit area that it holds is given by

$$Q_0 = C_{\text{GEM}} V_0 / A_{\text{GEM}} \sim 5 \times 10^5 \text{ fC mm}^{-2},$$

which means that the total charge density available in the avalanche for the highest-ionization tracks is of the order of 10% of that in the GEM. If all of this went into charging (which is clearly un-physical—after all we would see no tracks, if that was the case) that would produce a 10% change in the electric field, or a 50 V drop in the effective potential, which would result in a factor of 4 gain reduction—remember: the gain of the GEM doubles every 25 V.

What we are looking after is more like a 2 V reduction of the effective potential, resulting into a gain drop by 6%—remember: the gain increases by 3% per V. In order to explain this we need to assume that 2–3% of the avalanche charge goes into charging, which is not unrealistic.

See also:

The Gas Electron Multiplier

13.6 A Monte Carlo Simulation

Our basic understanding of the self-generated alpha particles in the GPD is not particularly prone to a full Geant4 simulation, as the exceedingly large number of secondary particles produced make the process very computationally-intensive.

On the other hand, it is not terribly hard to setup a simplified simulation capturing the relevant geometrical (i.e., the detector layout) and physical (i.e., the stopping-power cross sections) aspects of the problem. More specifically, the basic simulation flow goes as follows:

- generate the alpha particles (at a given energy) uniformly over the volume of the Be window, with an isotropic direction distribution;
- follow the particles through the Beryllium, the Aluminum layer and the DME;
- calculate the deposited energy and the track length in the active gas volume (for those particles that get there).

We emphasize that, based on the projected track length on the readout plane, rough estimates of the expected ROI size and track size (in pixels) can be calculated under reasonable assumptions.

A few relevant figures help visualizing the orders of magnitude involved in the process, here:

- the alpha range in Beryllium is 24.7 μm , or about half the thickness of the GPD window;
- about 12% of the alpha particles generated in the window do escape down to the Al layer;
- the average energy loss in the Aluminum is of the order of 25 keV, i.e., negligible;
- about 0.6% of the alpha particles generated in the window do leave a triggerable track in the GPD active gas volume.

The figure below shows the *geometrical* simulation of a small number of tracks, to give a rough idea of their distribution across the active area of the detector. The emerging picture is qualitatively in agreement with what we observe in real data.

Warning: Our default cut on large regions of interests might prevent us from reading out the longest alpha tracks.

For completeness we show below the distribution of the energy deposited in the GPD, as well as the length of the tracks, projected on the readout plane.




Fig. 13.10: Simulation of a small number of tracks in the GPD from 5 MeV alpha particles generated into the Beryllium window. The blue square represents the GPD active area. Tracks generating ROIs larger than 5000 (that would be discarded with the nominal data-taking settings) are shown with light-gray dashed lines. We notice that *short* tracks tend to cluster at the borders of the active area.




Fig. 13.11: Distribution of the energy deposited in the GPD by 5 MeV alpha particles generated in the Beryllium window.




Fig. 13.12: Distribution of the track length (projected onto the readout plane) for 5 MeV alpha particles generated in the Beryllium window.

WORKBOOK DEVELOPMENT

Warning: This page is intended for people contributing to the workbook, and contains all the relevant information for modifying, building and deploying it. If you are only interested in reading the document, chances are that you don't care about this page.

These pages have been created with [sphinx](#). (Yes: if you do want to contribute to the workbook you will have to install it.) Originally developed as the documentation system for the Python programming language, sphinx is the de-facto standard for documenting Python packages and is becoming increasingly popular as a mean fo publishing documentation of any sort on the web (from software packages to user manuals and technical blogs).

sphinx is based on the [reStructuredText](#) markup language and, among other things, features:

- several different output formats, including static HTML, LaTeX, pdf, ePub, plain and rich text;
- extensive cross-reference facilities
- support for clean hierarchical structures
- automatic indices.

14.1 Structure of the Repository

The source code for this document is hosted on the private IXPE software repository at [this location](#).

Though this might change in the future, the structure of the repository, at this point in time, is fairly simple:

- *contents*: contains all the .rst files with the actual pages;
- *figures*: contains all the figures used in the documentation—we might want to add some structure to the folder at some point in the future;
- *release*: contains the tools for the workbook release—see further down for more details.
- *gpdworkbook*: contains all the Python macros to create the plots, the basic idea being that we can make the package self-consistent and implement the capability of regenerating all the documentation plots in a self-consistent fashion.

14.2 Editing Contents

To edit the workbook output you can either modify existing .rst files or create new .rst files—the [reStructuredText](#) primer and the [sphinx](#) documentation are your friends in this job.

If you are adding new files do not forget to add them to one of the toc trees in the master (index.rst) file—sphinx will complain with a warning if you don't.

14.3 About the Figures

We take pride in making an effort to have good-quality figures in the workbook.

Asymptotically we would like to have each single figure in the document (or at least most of them) generated by a small snippet of code living in the workbook repository. Since we don't want to duplicate tons of code, it is ok for these snippets to import `ixpeobssim` and/or `gpdsw` Python modules, but if all of a sudden we need to change a label in a plot, we ought to be able to do it from within `gpdworkbook`.

As a general rule we want to include `.png` figures in the html output and `.pdf` figures in the pdf output. The code snippets should be set up in such a way two versions of each plot are saved next to each other, and the figure should be included with a wildcard

```
.. figure:: ../figures/file_name.*
```

so that each sphinx engine (e.g., html or LaTeX) picks the most appropriate file format.

14.4 Compiling the Documentation

There is a top-level Makefile with targets for the generation of the html and pdf output. (The original sphinx Makefile lives in the `contents` folder.) Compiling the documentation should be as easy as `cd`-ing into the local repository and running `make`:

```
make html
make pdf
```

By default the top-level Makefile places all the output into a `_build` directory placed at the top level.

Again: in order to compile the documentation you will need to have `sphinx` installed. Additionally, since the pdf is created through the LaTeX typographical engine, you will need a (fairly complete) LaTeX installation. Any recent `TeXLive` will do it.

14.5 Changing the Style

This requires digging somewhat deep into the sphinx internals, but there is large room to improve on the current style of the document, in both the html and the pdf flavors.

The workbook is currently using the (ubiquitous) `alabaster` sphinx theme, the reasons for that being that alabaster is nice-looking, extensively used and mobile-ready (among other things). We're not really looking forward for changing the base theme, although that is surely an option that can be discussed.

That all said, there are two places where the alabaster theme can be effectively customized, both of which are covered [here](#):

- the `conf.py` file;
- a custom css style sheet, which in our case lives in `contents/_static/ixpe_style.css`.

Needless to say, as they affect all developers, in the long term changes to the style should be done in a coordinated fashion.

Warning: We have not found a suitable location for the static html output, yet, but we are working on it. In case you wonder, the main issue is to have it in a location that is password-protected without having all the users create yet another account. For the time being the pdf is the preferred output product to be distributed.

RELEASE NOTES

Release 0.11.0, Sun, 18 Apr 2021 07:46:03

Release 0.10.0, Mon, 15 Feb 2021 07:00:38

Release 0.9.3, Thu, 08 Oct 2020 15:43:32

Release 0.9.2, Thu, 08 Oct 2020 10:25:16

Release 0.9.1, Wed, 07 Oct 2020 11:02:45

Release 0.9.0, Wed, 07 Oct 2020 10:58:12

Release 0.8.0, Fri, 05 Jun 2020 15:05:22

Release 0.7.0, Tue, 11 Feb 2020 01:01:31

- Macro and section on the secular gain variation expanded.
- GPD inventory expanded.

Release 0.6.0, Mon, 10 Feb 2020 01:43:23

- ASIC section refactored.
- BEE section added.

Release 0.5.3, Sun, 09 Feb 2020 20:35:48

- Updated release notes.

Release 0.5.2, Sun, 09 Feb 2020 20:34:21

- Yet another minor fix for LaTeX compilation.

Release 0.5.1, Sun, 09 Feb 2020 20:33:32

- Minor fix for LaTeX compilation.

Release 0.5.0, Sun, 09 Feb 2020 20:26:47

- GEM section expanded.

Release 0.4.0, Sun, 09 Feb 2020 19:46:56

- Section on the role of the gas pressure for the GPD performance added.
- General introduction about the GPD vastly expanded.
- Development section updated.

Release 0.3.0, Tue, 04 Feb 2020 17:04:42

- Minor clarification in the event reconstruction section.

- Added a section about the simulation of the alpha track properties.

Release 0.2.0, Fri, 31 Jan 2020 14:43:23

- Section on charging updated.
- Section on the secular gain variations updated.
- Added material on the alpha particles.
- Added gpdworkbook folder as a seed for a mechanism to be able to recreate all plots programmatically.
- Development workflow documentation updated.

Release 0.1.0, Thu, 23 Jan 2020 15:37:55

- Master file revised (issue #1).
- Workbook development description added (issue #5).
- Charging chapter factored out the GEM chapter.
- Added GPD inventory (issue #10).
- Figure sizes changed across the document (issue #2).

Release 0.0.3, Sun, 19 Jan 2020 23:15:15

- pdf compilation through LaTeX fixed in the top-level Makefile.

Release 0.0.2, Sun, 19 Jan 2020 22:58:20

- Debugging the release manager :-)

Release 0.0.1, Sun, 19 Jan 2020 22:55:20

- Rudimentary release manager in place.

Release 0.0.0

- Initial release based on the discussion at the January 2020 face-to-face meeting in Huntsville.

INDICES AND TABLES

- `genindex`
- `modindex`
- `search`