Quantitative Characterisation of SiPM Radiation Damage using Emission Microscopy

Author: Lai Gui

Supervisors: Dr Esteban Currás Rivera
Dr Guido Haefeli
Prof. Olivier Schneider

Laboratoire de Physique des Hautes Énergies
École Polytechnique Fédérale de Lausanne

July 13, 2026
PHYS-599(b) - Projet de master en physique

Contents

Acknowledgements

I would like to thank my supervisors, Dr Esteban Currás Rivera, Dr Guido Haefeli and Prof. Olivier Schneider, for their guidance and support throughout this project. I am especially grateful to Dr Esteban Currás Rivera for his mentorship over the course of this thesis, and for his support through a difficult period towards the end of my studies. I thank Dr Guido Haefeli for his valuable feedback, and Prof. Olivier Schneider for making this project possible — and, more importantly, without whom I would not have had the opportunity to work with the laboratory from the very beginning of my time at EPFL.

Special thanks go to my fellow Master's students Maxime Nemoz, Samuel Morand, Guillaume Arnold, Giacomo Belleri, Bogdan Boucher and Antoine Sauvion, for their companionship and collaboration throughout my Master's project. I am also deeply grateful to the other members of LPHE for their support and kindness during my studies and my time in Lausanne. Finally, I would like to acknowledge the LHCb collaboration and the researchers whose pioneering work laid the foundations for this project.

Abstract

Silicon Photomultipliers (SiPMs) are the photodetector of choice for the LHCb Scintillating-Fibre (SciFi) tracker and are increasingly ubiquitous across High Energy Physics. Their principal weakness is the dark-count rate (DCR), which grows with the neutron fluence accumulated in the collider environment and will become a limiting factor for the High-Luminosity LHC (HL-LHC) upgrade and the LHCb Upgrade II. This work applies emission microscopy (EMMI), a spatially resolved defect-localisation technique borrowed from semiconductor failure analysis, to characterise the radiation damage of SiPMs from two manufacturers, FBK (W3m, 31 µm pitch, $V_{\mathrm{bd}} \approx 31\,\mathrm{V}$) and Hamamatsu (H2024, 50 µm pitch, $V_{\mathrm{bd}} \approx 51\,\mathrm{V}$), each studied with a non-irradiated reference and samples at four neutron fluences of 3×10^{11}, 1×10^{12}, 3×10^{12} and $1 \times 10^{13}\,\mathrm{n_{eq}/cm^2}$, the highest of which exceeds the $\sim 3 \times 10^{12}\,\mathrm{n_{eq}/cm^2}$ expected at the SciFi SiPMs over the full HL-LHC programme by a factor of three. Each sensor is biased above breakdown and imaged through a microscope onto a cooled Hamamatsu ORCA-Quest 2 qCMOS camera ($-20°$C, 10-minute exposures), recording the faint near-infrared light emitted by dark-count avalanches while the leakage current is logged simultaneously with a Keithley 2450 SourceMeter.

A reproducible analysis pipeline is developed to turn these images into quantitative, cell-level observables. The breakdown voltage of each channel is first extracted with the inverse-logarithmic-derivative (ILD) method, the microcell grid is then detected automatically, and the emission hotspots within each cell are reconstructed as 2D Gaussians using percentile thresholding, local-maximum detection and watershed segmentation. This makes it possible to follow a range of per-cell and per-hotspot quantities as a function of overvoltage and fluence, from the emission intensity and hotspot density down to the brightness, size, current contribution and in-cell position of each individual hotspot. The extracted distributions are well described by Exponentially Modified Gaussian (EMG) fits and are proposed both as benchmarks for detector simulation and as candidate non-destructive metrics for inferring the accumulated fluence, establishing emission microscopy as a cell-level diagnostic for SiPMs deployed in the increasingly radiative environments of future experiments.

1 Introduction

1.1 High Energy Physics and CERN

High Energy Physics (HEP), also called particle physics, is the study of the most fundamental constituents of matter and the forces through which they interact. Our current understanding is encapsulated in the *Standard Model* (Figure 1), a quantum field theory that classifies all known elementary particles into the spin-$\frac{1}{2}$ fermions that make up matter, comprising six quarks and six leptons organised in three generations, and the spin-1 gauge bosons that mediate the fundamental interactions: the photon for electromagnetism, the $W^{\pm}$ and Z bosons for the weak force, and the gluons for the strong force. The scalar Higgs boson, discovered at CERN in 2012, completes the model by giving mass to the elementary particles through the Brout-Englert-Higgs mechanism. The Standard Model has been confirmed to remarkable precision, yet it leaves open questions, such as the nature of dark matter, the matter-antimatter asymmetry of the universe, and the absence of gravity from the framework, which motivate continued experimental study.

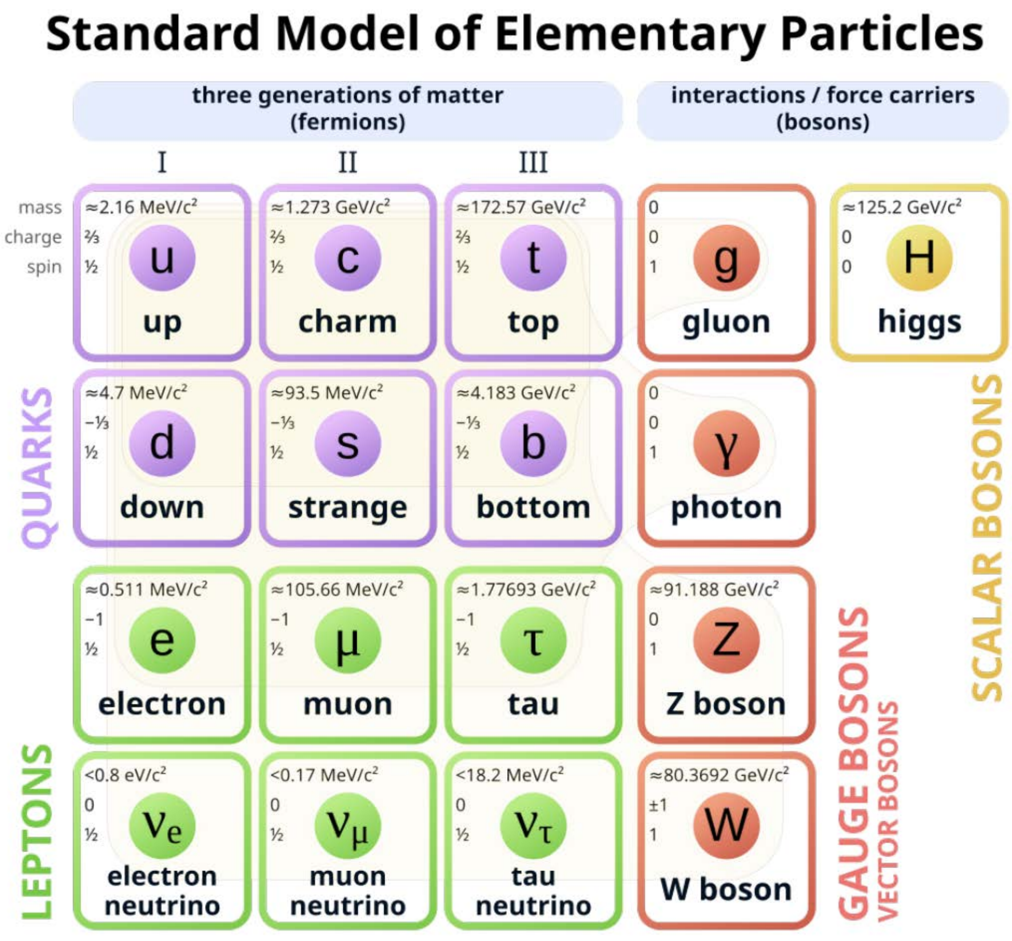

Figure 1: The Standard Model of particle physics [1], comprising the six quarks and six leptons (three generations of spin-$\frac{1}{2}$ fermions) that constitute matter, the spin-1 gauge bosons that mediate the electromagnetic, weak and strong interactions, and the scalar Higgs boson responsible for electroweak symmetry breaking.

Because the heavy particles of interest are unstable and were abundant only in the hot, dense conditions of the early universe (Figure 2), they cannot simply be found in nature; they must be produced by accelerating beams of particles to relativistic energies and colliding them head-on, converting kinetic energy into the rest mass of new particles according to $E = mc^2$. The higher the collision energy, the heavier the states that can be created and the earlier the cosmological epoch that can be probed, which is why HEP relies on ever more powerful particle accelerators and colliders.

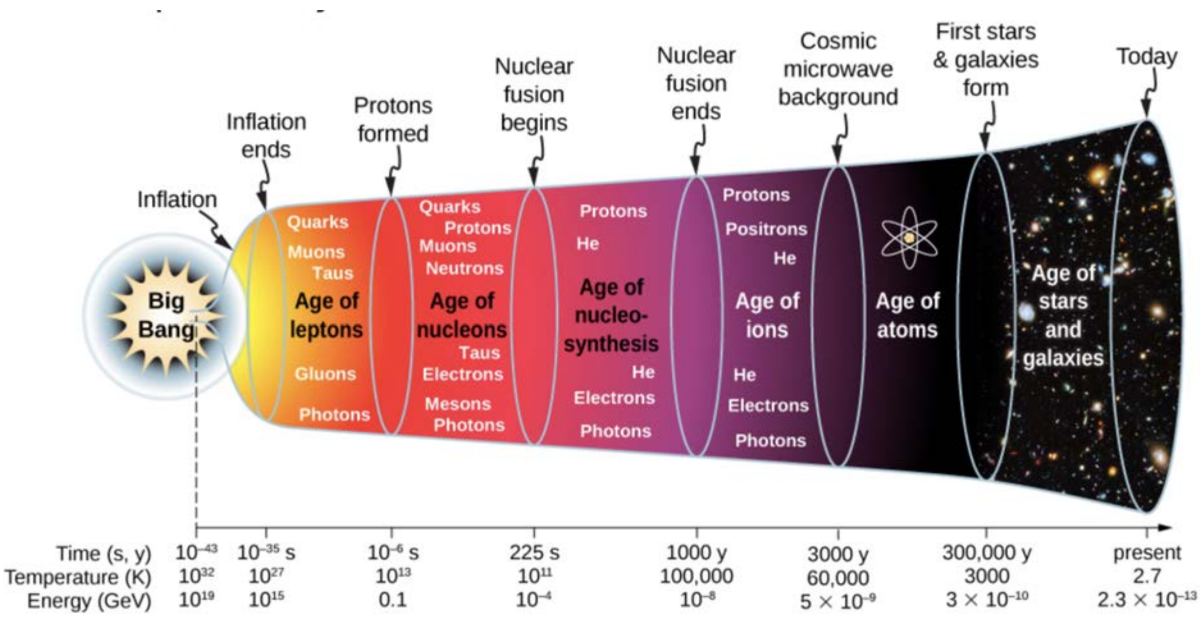

Figure 2: Timeline of the Big Bang [2]. The heavy, short-lived particles studied at colliders were abundant only in the earliest, hottest moments after the Big Bang; recreating them today requires concentrating comparable energy densities in head-on collisions.

CERN, the European Organization for Nuclear Research (from the French *Conseil Européen pour la Recherche Nucléaire*), is the world's largest particle-physics laboratory. Founded in 1954 and located on the Swiss-French border near Geneva (Figure 3), it operates a chain of accelerators that culminates in the Large Hadron Collider and hosts some of the largest scientific experiments ever built. CERN brings together tens of thousands of scientists, engineers and students from across the globe in a long-standing model of international collaboration, with the common goal of probing the structure of matter and the fundamental forces. The laboratory has been the site of landmark discoveries, including the W and Z bosons and, most recently, the Higgs boson.

Figure 3: Aerial view of the CERN site near Geneva, with the ∼27 km LHC ring and the locations of its four main detectors (ATLAS, CMS, ALICE, LHCb) overlaid [3].

1.2 LHC and LHCb

The Large Hadron Collider (LHC) is the highest-energy particle collider and the largest machine ever built (Figure 4). Housed in a tunnel about 27 km in circumference, roughly 100 m below ground, it accelerates two counter-rotating beams of protons in opposite directions and brings them into collision at four interaction points. The protons are guided and focused by superconducting magnets operating at 1.9 K and are grouped into thousands of bunches that cross every 25 ns. In its current Run 3, the LHC collides protons at a centre-of-mass energy of $\sqrt{s} = 13.6$ TeV, the highest ever achieved in the laboratory. The rate of collisions is characterised by the *instantaneous luminosity*; integrated over time it determines the size of the recorded data set and hence the reach for rare processes. The forthcoming HL-LHC upgrade will raise the integrated luminosity by roughly an order of magnitude, and with it the radiation fluence to which the detectors, and their SiPM photodetectors in particular, are exposed.

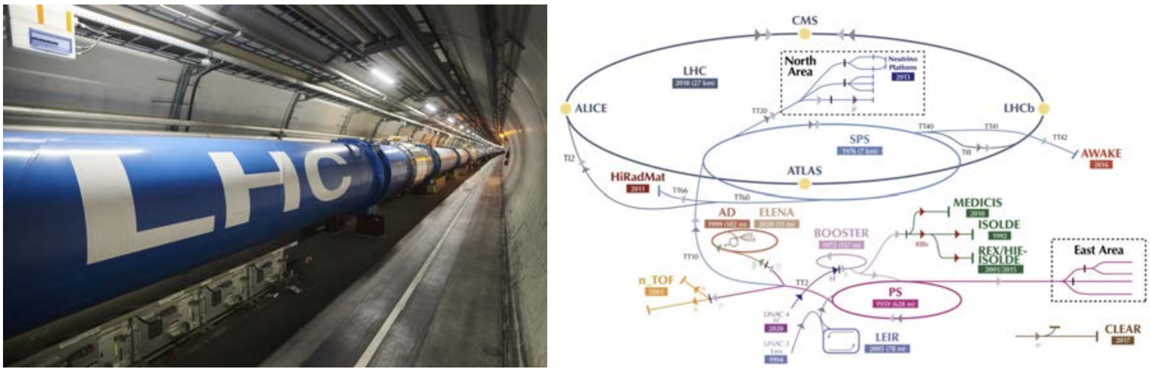

Figure 4: **Left**: the LHC accelerator tunnel, showing the LHC magnets [4]. **Right**: schematic of the CERN accelerator complex, in which a chain of injectors feeds the LHC, the newest and largest ring [5].

A large detector is installed at each of the four collision points to record the particles produced. ATLAS (A Toroidal LHC ApparatuS) and CMS (Compact Muon Solenoid) are general-purpose detectors that independently discovered the Higgs boson and search broadly for new physics; ALICE (A Large Ion Collider Experiment) is optimised for the study of the quark-gluon plasma in heavy-ion collisions; and LHCb (Large Hadron Collider beauty) specialises in the physics of hadrons containing beauty (b) and charm (c) quarks. Here at the Laboratoire de Physique des Hautes Énergies (LPHE) at EPFL, we participate heavily in the LHCb experiment.

Unlike the cylindrical, hermetic designs of ATLAS and CMS, LHCb is built as a single-arm forward spectrometer (Figure 5), instrumenting only the cone around the beam axis in which the b- and c-hadrons are predominantly produced.

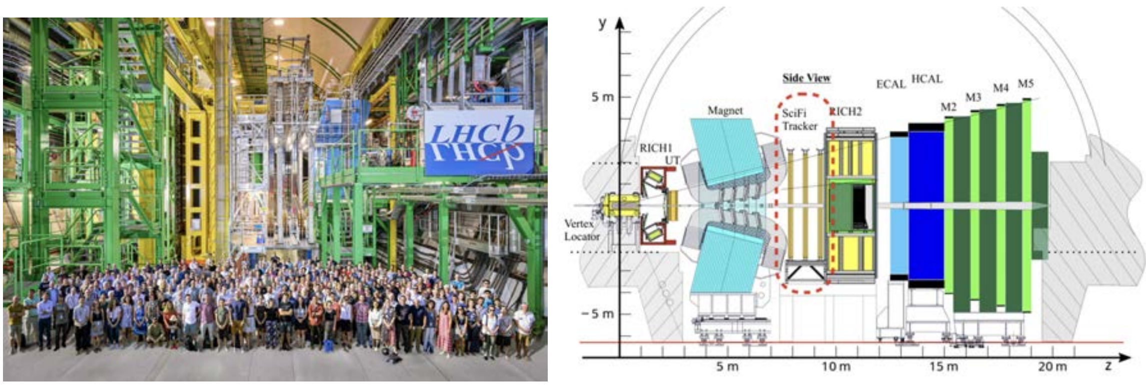

Figure 5: **Left**: the LHCb detector with members of the collaboration for scale [6]. **Right**: schematic cross-section of the single-arm forward spectrometer, with the SciFi tracker (read out by the SiPMs studied here) located downstream of the dipole magnet [7].

Its subdetectors include a high-precision vertex locator, tracking stations on either side of a dipole magnet, ring-imaging Cherenkov detectors, calorimeters and muon chambers, which together reconstruct the decay vertices and momenta of these short-lived hadrons (Figure 6). This makes LHCb a precision laboratory for studying CP violation and rare decays, where tiny deviations from the Standard Model could reveal the effects of new particles. The downstream tracking stations are based on scintillating-fibre (SciFi) technology read out by SiPMs, which is the context for the present work.

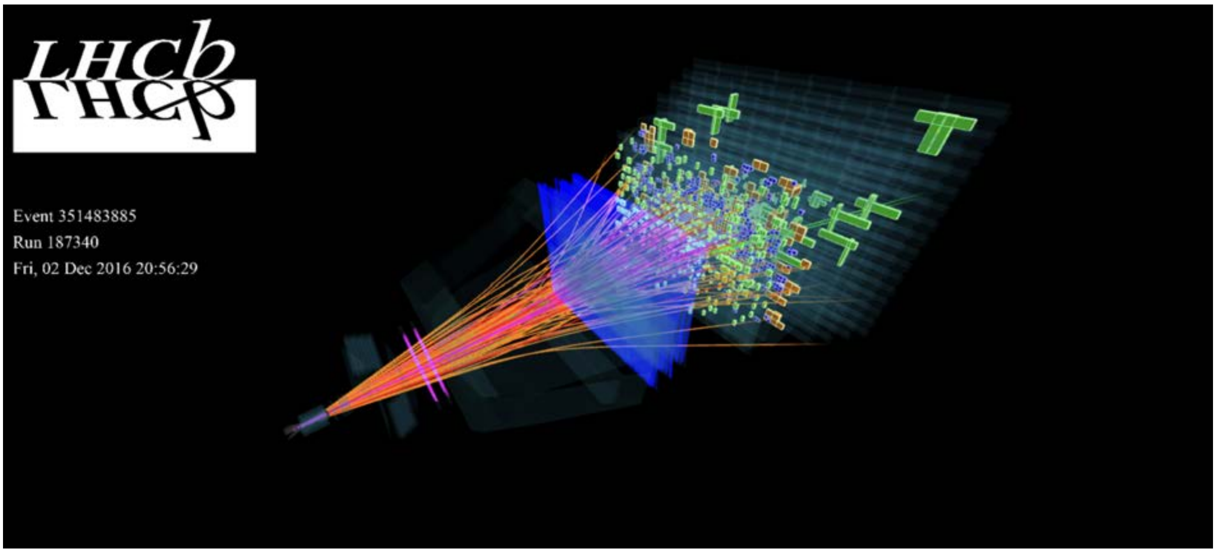

Figure 6: Illustration of the detector response and track reconstruction for a single collision event at LHCb, in which the charged-particle trajectories are reconstructed from hits in the successive tracking layers [8].

1.3 Silicon Photomultiplier

Silicon Photomultipliers (SiPMs) are the main devices used to detect the light emitted by scintillators in subdetectors such as the LHCb SciFi tracker, where they are coupled directly to the ends of the scintillating fibres (Figure 7). Their radiation tolerance therefore sets a hard constraint on the operational lifetime of the tracker, which is the motivation for characterising their damage in detail.

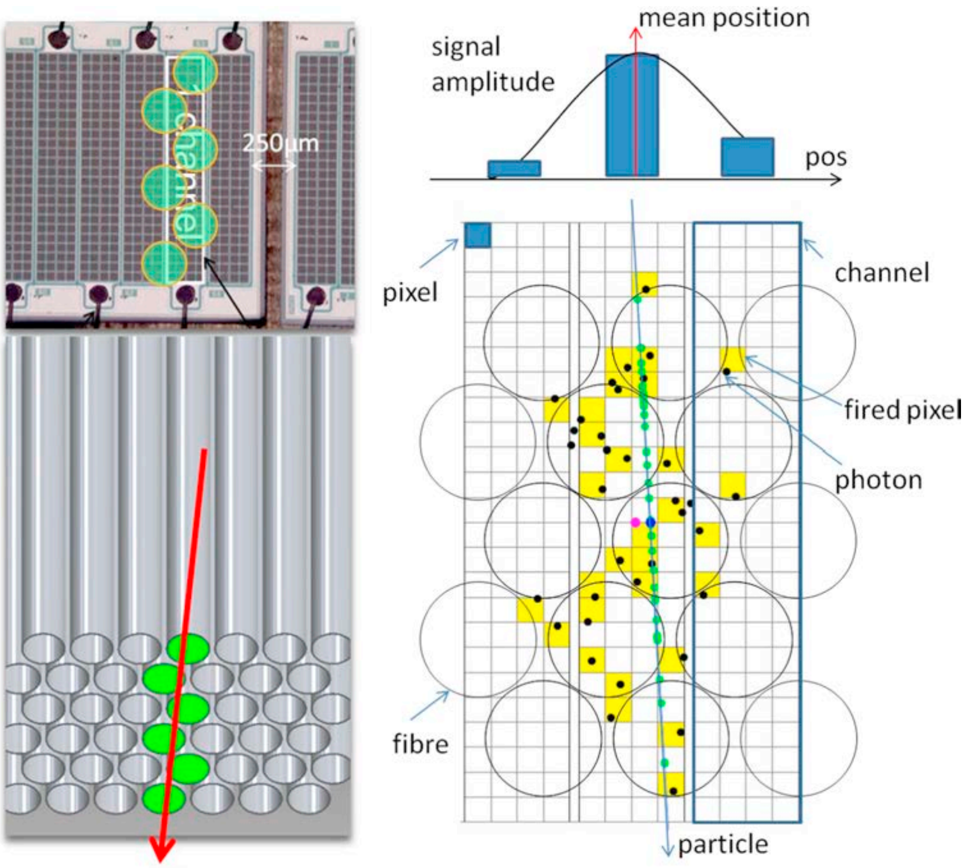

Figure 7: Optical coupling between the scintillating fibres of the SciFi tracker and the SiPM arrays that read them out [9]. Scintillation light produced in the fibres by traversing charged particles is guided to, and detected by, the SiPMs at the fibre ends.

A SiPM is a solid-state photodetector consisting of a dense, two-dimensional array of Single-Photon Avalanche Diodes (SPADs) (Figure 8), each connected in series with its own quenching resistor and read out in parallel on a common bias line [10]. Each SPAD, or *microcell*, is a reverse-biased p-n junction operated in *Geiger mode*, i.e. biased above its breakdown voltage V_{bd} (Figure 9). The amount by which the applied bias exceeds breakdown is the *overvoltage*, $\text{OV} = V_{\text{bias}} - V_{\text{bd}}$, and it sets the operating point of the device. When a free charge carrier, generated by an absorbed photon or thermally, enters the high-field depletion region, the electric field is strong enough that a self-sustaining avalanche develops through impact ionisation. The resulting current pulse is essentially independent of the number of carriers that triggered it, so a single microcell acts as a binary "fired/not-fired" element. The total charge released per avalanche, expressed in units of the elementary charge, defines the *gain* of the cell, which grows approximately linearly with overvoltage. The quenching resistor R_q then limits and stops the avalanche by dropping the voltage across the junction below breakdown, after which the cell recharges with a recovery time

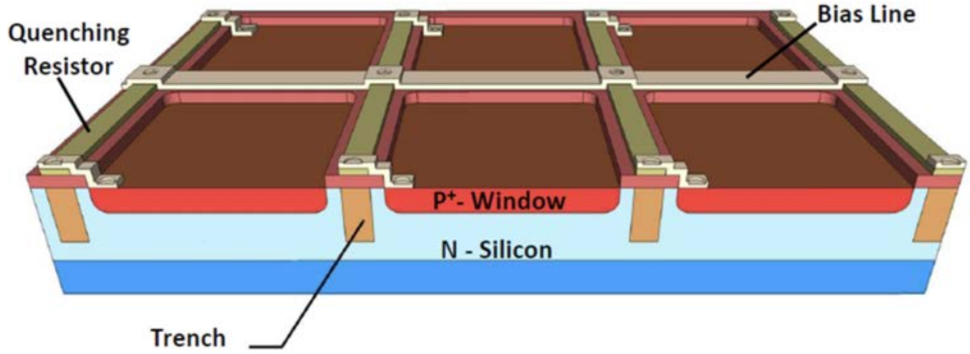

Figure 8: Schematic structure of a SiPM composed of an array of SPAD (Single-Photon Avalanche Diode) microcells [10]. Each microcell contains a quenching resistor, a bias line, and typically a p^+ window on an n-type silicon substrate. Trenches between microcells provide optical isolation to suppress crosstalk.

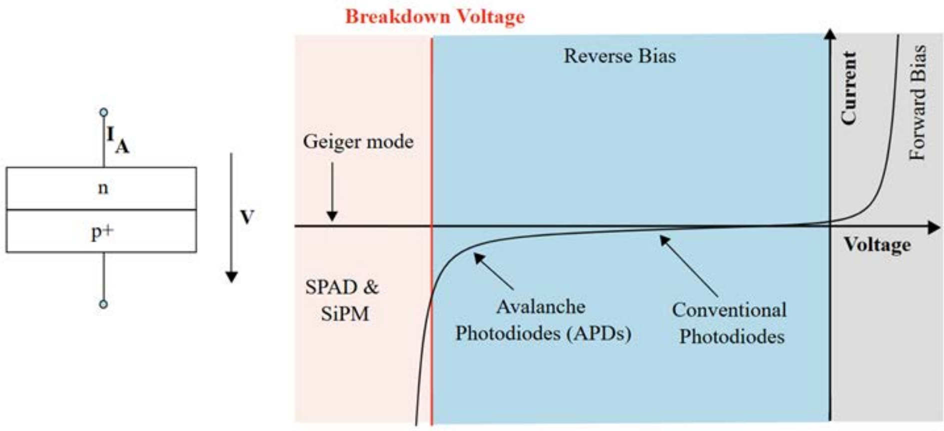

Figure 9: Current-voltage characteristic of a SPAD microcell [11]. In Geiger mode the SPAD is reverse-biased above its *Breakdown Voltage*, so that the absorption of a single photon, or the thermal generation of a single carrier, can trigger a full avalanche. The overvoltage $\mathrm{OV} = V_{\mathrm{bias}} - V_{\mathrm{bd}}$ sets the field strength and hence the gain and trigger probability.

constant $\tau_{\mathrm{rec}} = R_q C_{\mathrm{cell}}$. Because the microcells fire independently and their outputs are summed, the SiPM signal is proportional to the number of cells fired, giving the device photon-number-resolving capability at low occupancy.

Compared with the vacuum Photomultiplier Tubes (PMTs) traditionally used in high-energy physics, SiPMs offer single-photon sensitivity and a gain ($\sim 10^6$) comparable to

PMTs while being compact, mechanically robust, insensitive to magnetic fields, and operable at low bias voltages ($\mathcal{O}$(10–50 V)) [10]. These properties make them well suited to the tightly packed, high-field environments of modern collider detectors, where they can be segmented finely and placed close to the active medium.

The principal limitation of SiPMs is their dark-count behaviour. Imperfections in the silicon lattice arise both from manufacturing and, crucially, from radiation-induced defects. They introduce energy levels within the bandgap that act as generation centres for thermally generated electron-hole pairs. These spontaneous carriers can trigger avalanches indistinguishable from genuine photon signals, producing the Dark Count Rate (DCR). In the radiation environment of the LHC, displacement damage from neutrons and other hadrons steadily increases the defect density, and therefore the DCR, degrading the signal-to-noise performance of the detector over time.

This concern becomes more acute with the upcoming High-Luminosity LHC (HL-LHC) upgrade, during which the integrated fluence seen by the SciFi SiPMs will rise substantially (the expected levels are quantified in Figure 12 and the accompanying discussion below). Maintaining adequate performance over the full operational lifetime therefore requires SiPM designs and operating strategies that remain radiation-tolerant up to these fluences. Characterising how key figures of merit such as DCR and time resolution evolve under irradiation is essential both for validating candidate sensors for the LHCb Upgrade II and for guiding the development of more radiation-hard SiPM technologies. This work contributes to that effort by quantifying, with emission microscopy, where and how strongly irradiated SiPMs emit secondary photons as a spatially resolved proxy for their dark-count activity.

1.3.1 Dark Count Rate (DCR)

The dominant source of intrinsic noise in a Silicon Photomultiplier is the spontaneous, light-independent firing of microcells known as the Dark Count Rate (DCR). Its microscopic origin lies in defects of the silicon lattice. In a perfect crystal, promoting an electron from the valence to the conduction band requires the full bandgap energy $E_g \approx 1.12\,\mathrm{eV}$, so thermal generation at the operating temperature is strongly suppressed. Lattice imperfections, whether introduced during manufacturing (impurities, dislocations) or, more importantly here, by radiation-induced displacement damage, create localised energy levels within the forbidden gap. These mid-gap states act as stepping stones for carrier generation: through the Shockley-Read-Hall (SRH) mechanism [12] a carrier can be thermally excited from the valence band to a trap level and then to the conduction band in two smaller steps, dramatically increasing the generation rate relative to the defect-free case. A carrier generated in this way inside the high-field depletion region can trigger a full avalanche that is electrically indistinguishable from a genuine photon-induced count.

Because each defect provides such a generation centre, the DCR scales with the density of mid-gap states and therefore increases monotonically with accumulated radiation fluence. This direct correspondence between lattice defects and dark counts is precisely what makes the spatial distribution of dark-count activity a useful proxy for the distribution of radiation damage across the sensor, a distribution that is mapped in this work by emission microscopy.

Several factors contribute to the DCR:

1. **Temperature:** Higher temperature increases the likelihood of thermal activation of carriers within the depletion region. The DCR exhibits an exponential dependence on temperature:

$$\text{DCR} \propto \exp\left(-\frac{E_a}{kT}\right), \tag{1}$$

   where E_a is the activation energy for carrier generation, k is the Boltzmann constant, and T is the absolute temperature in *Kelvin (K)*.

2. **Overvoltage:** Increasing the overvoltage enhances the electric field within each microcell, thereby increasing the probability that a thermally generated carrier will trigger an avalanche.

3. **Pixel Size:** Larger pixels have a greater active volume for thermal carrier generation, leading to a proportionally higher DCR.

4. **Radiation Damage:** Exposure to ionising radiation, particularly neutron or proton fluence, introduces defects in the silicon lattice that act as trap states. These traps serve as trap-assisted tunnelling (TAT) centres, enhancing both dark counts and afterpulsing. Because tunnelling is only weakly temperature-dependent whereas thermal (SRH) generation is exponentially suppressed by cooling (Eq. 1), TAT becomes the relatively dominant dark-count mechanism at the reduced operating temperatures used to mitigate DCR in irradiated detectors such as the SciFi tracker SiPMs.

Under standard operating conditions, the dark-count generation in a SiPM follows a Poisson distribution:

$$P(n;\lambda,t) = \frac{(\lambda t)^n e^{-\lambda t}}{n!} \quad \Rightarrow \quad P(n=0;\lambda,t) = e^{-\lambda t}, \tag{2}$$

where λ represents the rate of dark counts. However, in the presence of severe radiation damage, or when the pixel recovery time becomes comparable to the average inter-event spacing, deviations from Poisson statistics occur.

1.3.2 Afterpulsing

Radiation-induced traps within the silicon lattice can capture charge carriers during an avalanche. These carriers may later be released after a characteristic delay, potentially retriggering an avalanche in the same pixel. This phenomenon, known as afterpulsing, introduces correlated noise following an initial event. The probability of afterpulsing depends on factors such as the trap density and the lifetime of the intermediate states introduced by the traps, both of which increase with fluence.

1.3.3 Recovery Time

After an avalanche event, the quenching resistor and pixel capacitance determine how quickly the electric field within the microcell recovers. The recovery process follows an exponential behaviour:

$$V(t) = V_{\text{bias}} \left(1 - e^{-t/\tau_{\text{rec}}}\right), \tag{3}$$

where $\tau_{\text{rec}} = R_q C_{\text{cell}}$ is the recovery time constant, defined by the quenching resistor R_q and the pixel capacitance C_{cell}. A longer recovery time reduces the effective detection efficiency for subsequent photons but also decreases the probability of correlated avalanches.

In heavily irradiated SiPMs, where the dark-count rate is high, the field inside a cell may not have time to recharge fully before the next avalanche is triggered, resulting in a lower effective overvoltage than the nominal applied value. This saturation effect is revisited in the current-voltage measurements of Section 4.1.

1.3.4 Secondary Photon Emission of SiPM

Charge-carrier multiplication takes place in the high-field region of the SiPM. There, the accelerated carriers dissipate part of their energy through the emission of photons (Figure 10). It is precisely this secondary emission, one faint optical photon for roughly every 10^5–10^6 carriers, that emission microscopy detects and uses as an optical signature of avalanche (and hence dark-count) activity.

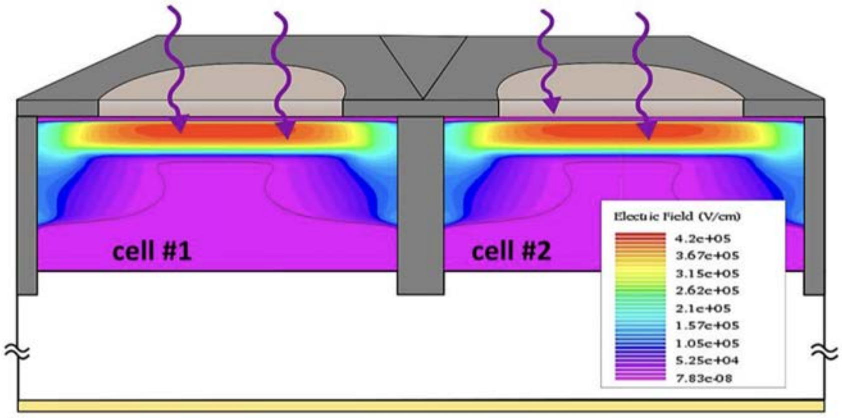

Figure 10: Electric-field profile through a SiPM microcell [13]. Carrier multiplication is confined to the narrow high-field region, where both electrons and holes gain enough energy to sustain the avalanche and to emit the secondary photons imaged in this work. The lower-field depletion region around it in blue (see also Figure 56) drifts carriers into this multiplication volume.

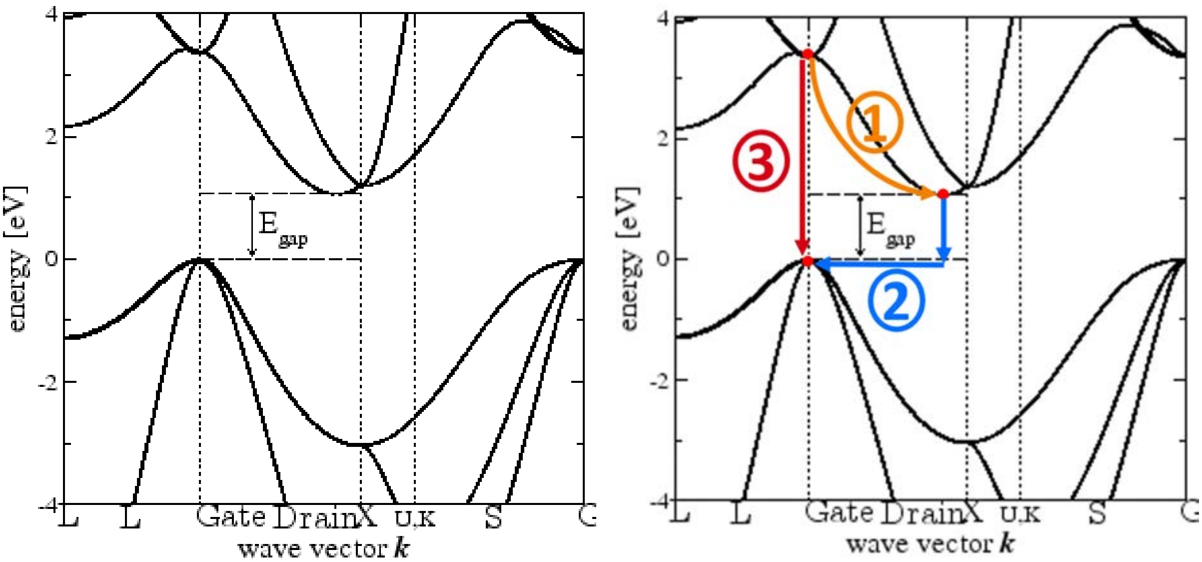

Figure 11: **Left**: band structure of silicon; the band above 0 eV is the conduction band and the band below is the valence band [14]. **Right**: the same structure annotated with the three photon-emission mechanisms relevant to a SiPM avalanche: (1) intraband bremsstrahlung of hot carriers in the conduction band, giving the broad visible-infrared continuum; (2) indirect interband (phonon-assisted) recombination near the $\sim$1.12 eV gap, giving near-infrared emission; and (3) direct interband recombination at the Γ-point ($\sim$3.4 eV), contributing weakly in the UV.

Photon emission proceeds mainly through three processes, indicated on the silicon band structure of Figure 11:

1. **Intraband Bremsstrahlung:** When a SiPM is biased above breakdown, carriers (electrons and holes) are accelerated by the strong electric field ($\sim$3-5 $\times$ 10^5 V/cm in the avalanche region) and become hot carriers with kinetic energies far exceeding the thermal-equilibrium value. These hot carriers undergo intraband transitions, decelerating by emitting a photon without changing band (no recombination involved); this is essentially carrier bremsstrahlung. The emitted spectrum is broad and continuous, extending from the visible into the short-wave infrared ($\sim$500-1600 nm) because the hot-carrier energy distribution is broad. This is the primary source of the observed emission continuum.

2. **Indirect interband transition:** Silicon is an indirect-bandgap semiconductor, so radiative electron-hole recombination requires a phonon to conserve crystal momentum. The emitted photon energy is slightly below E_g, giving a near-infrared range of $\sim$1100–1200 nm ($E_g \approx 1.12$ eV at 300 K, corresponding to $\sim$1100 nm). The efficiency is extremely low ($\sim 10^{-6}$ per carrier pair) owing to the indirect nature of the transition. However, in a Geiger-mode avalanche $\sim 10^5$–10^6 carriers are generated per fired microcell, so even this low efficiency produces a measurable photon flux, despite the camera's sensitivity being very low at this wavelength.

3. **Direct interband transition:** Silicon also possesses direct transitions at the Γ-point direct gap ($\sim$3.4 eV, i.e. UV at $\sim$365 nm), where hot carriers with sufficient energy can undergo quasi-direct recombination without requiring a phonon. This contributes to the UV portion of the emission spectrum, but at these energies the absorption length in silicon is so short that few of these photons escape the device to be detected.

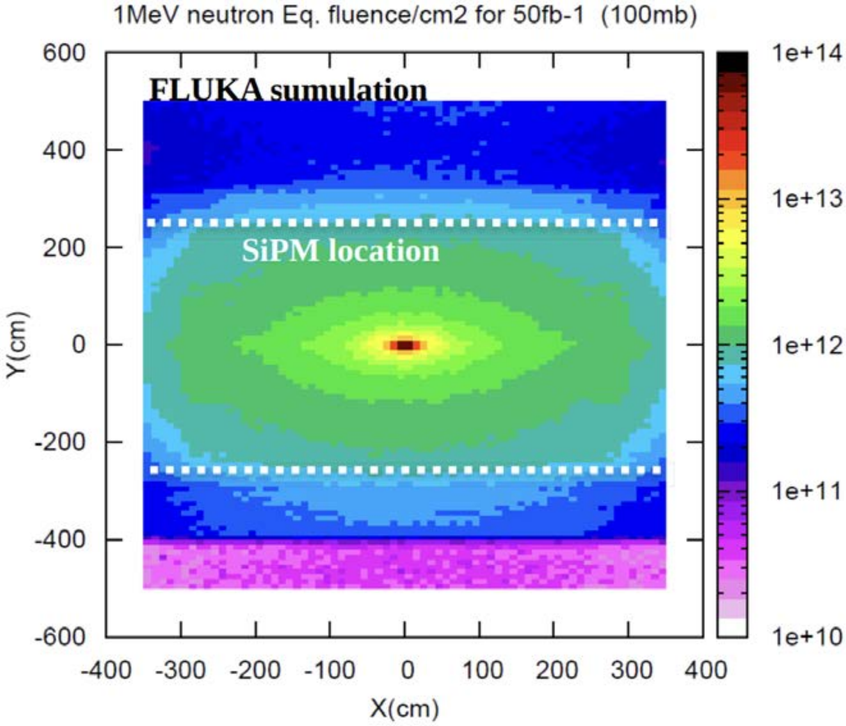

Figure 12: Simulated 1-MeV-neutron-equivalent fluence in the x–y plane of the LHCb detector at the end of the HL-LHC operation [7]. The SciFi SiPMs sit away from the beam axis, in a region of reduced but still significant fluence; integrated over the full programme this reaches $\sim 3 \times 10^{12}\,\mathrm{n_{eq}/cm^2}$, motivating the fluence range studied here.

1.3.5 Neutron Flux in the HL-LHC Environment

In the SciFi detector system the SiPM arrays are coupled to the scintillating fibres and positioned away from the beam centre, so they are exposed to a reduced neutron flux compared with on-axis detector components (Figure 12). Nevertheless, over the full High-Luminosity LHC (HL-LHC) operation cycle the accumulated fluence is expected to reach approximately $3 \times 10^{12}\,\mathrm{n_{eq}/cm^2}$. Such exposure significantly alters the SiPM characteristics, raising the dark-count rate and degrading reliability.

In this study the sensors are characterised up to $1 \times 10^{13}\,\mathrm{n_{eq}/cm^2}$, roughly three times the expected HL-LHC SciFi dose, so that the damage is probed with a comfortable safety margin. The objective is to identify and quantify the dark-signal activity cell by cell, in preparation for the LHCb Upgrade II programme and, more broadly, to establish a diagnostic that remains applicable in the even harsher radiation environments of future experiments.

2 Methodology

2.1 Experimental Setup

The experiment is carried out in the dark room on the second floor of the Cubotron building. The setup consists of an optical microscope, a light source, a power supply, a cooled camera, and a vacuum chuck that holds the SiPM in place. A labelled photograph of the bench is shown in Figure 13.

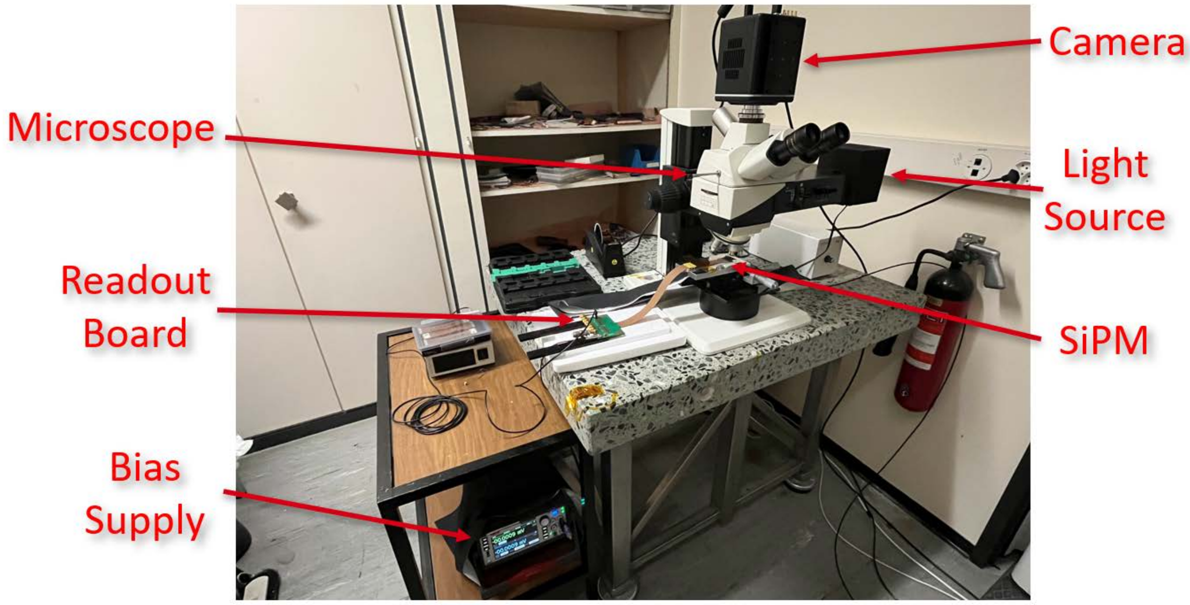

Figure 13: The emission-microscopy bench used in this project. The cooled qCMOS camera is mounted on the microscope, beneath which the SiPM is held flat by a vacuum chuck; the bias supply is connected to the SiPM through a readout board, while a halogen lamp provides the illumination used only for the short reference exposure.

The SiPM is biased through a Keithley 2450 SourceMeter, illuminated (for the reference image) by a Leica halogen lamp, and imaged through interchangeable microscope objectives onto the camera sensor. Before any image is taken, the camera is allowed to cool until its sensor reaches its operating temperature of −20°C, which takes about two minutes and is required to keep the dark current of the sensor itself negligible over the long emission exposures.

The image sensor is a Hamamatsu ORCA-Quest 2 qCMOS camera (model C15550-22UP). Its acquisition settings, kept fixed throughout the campaign, are summarised in Tab. 1. The optical and electrical components of the bench are listed in Tab. 2.

Table 1: Acquisition settings of the Hamamatsu ORCA-Quest 2 qCMOS camera (C15550-22UP), held fixed throughout the measurement campaign. The 10-minute emission exposure and −20°C sensor temperature are chosen to accumulate enough of the faint secondary-emission signal while keeping the sensor's own dark current negligible; the ∼3 ms reference exposure captures the illuminated cell structure for later superposition.

Parameter	**Setting**
Sensor type	qCMOS, monochrome
Operating temperature	−20°C
Exposure time	10 min (emission image); ∼3 ms (reference image)
Binning	1×1
Frame size	4096×2304 pixels
Data depth	16 bit
Readout speed	Standard
Capture mode	Area, Internal trigger

Table 2: Optical and electrical components of the emission-microscopy bench. The 20×/0.40 objective (highlighted in the acceptance-angle calculation of Section 4.4) is used for the emission measurements as a compromise between field of view and light collection; the Keithley 2450 both biases the sensor and logs the leakage current used for plotting the intensity-current correlation.

Component	**Model**	**Specification**
Objectives (Leica)	566077, N Plan EPI 5×/0.12 BD	5×, inf/–/B
	566049, N Plan L 20×/0.40	20×, inf/0/c
	566036, N Plan L 50×/0.50	50×, inf/0/c
Source meter	Keithley 2450	SMU, bias and current logging
Light source	Leica 12 V / 100 W halogen	set to 3 V

The camera also offers a photon-number-resolving mode, which was not used in this work. This mode has an upper limit in number of photon counts that can be resolved per bin (200 photons/pixel under 4096×2304 frame size), which is not well-suited to the long (10-minute) exposures used here, and even less so for the highly irradiated samples, where the accumulated dark-count signal greatly exceeds the single-photon scale and saturates the available dynamic range. The standard high-dynamic-range readout is therefore preferred.

2.2 SiPMs under Investigation

The SiPMs under investigation come from two suppliers, FBK and Hamamatsu, and their breakdown voltages and fluences are summarised in Tab. 3.

Table 3: Fluence and breakdown voltage (V_{bd}, channel 118) of the Hamamatsu (H2024, 50 µm) and FBK (W3m/W1m, 31 µm) sensors across the five radiation levels studied. The V_{bd} values, extracted with the ILD method of Section 3.1, are stable to within ∼0.2 V across fluence for each manufacturer (∼51 V for Hamamatsu, ∼31 V for FBK), confirming that irradiation shifts the dark-count rate rather than the breakdown point and allowing overvoltages to be compared consistently between samples.

Producer	**Name**	**Pixel Size**	**Fluence**	$\boldsymbol{V}_{\mathrm{bd}}$ (ch118)
	H2024	50 µm	1×10^{13} $\mathrm{n_{eq}/cm^2}$	51.69 ± 0.02 V
	H2024	50 µm	3×10^{12} $\mathrm{n_{eq}/cm^2}$	50.97 ± 0.02 V
Hamamatsu	H2024	50 µm	1×10^{12} $\mathrm{n_{eq}/cm^2}$	51.45 ± 0.07 V
	H2024	50 µm	3×10^{11} $\mathrm{n_{eq}/cm^2}$	51.31 ± 0.02 V
	H2024	50 µm	0 $\mathrm{n_{eq}/cm^2}$	51.10 ± 0.11 V
	W3m	31 µm	1×10^{13} $\mathrm{n_{eq}/cm^2}$	31.07 ± 0.02 V
	W3m	31 µm	3×10^{12} $\mathrm{n_{eq}/cm^2}$	31.00 ± 0.02 V
FBK	W3m	31 µm	1×10^{12} $\mathrm{n_{eq}/cm^2}$	31.11 ± 0.02 V
	W3m	31 µm	3×10^{11} $\mathrm{n_{eq}/cm^2}$	31.19 ± 0.02 V
	W1m	31 µm	0 $\mathrm{n_{eq}/cm^2}$	30.94 ± 0.04 V

The breakdown voltages are obtained with the *Inverse Logarithmic Derivative* (ILD) method described in Section 3.1.

The Hamamatsu SiPMs are shown in Figure 14. Each device has 128 channels in total; a channel consists of 5 rows and 32 columns, giving 160 cells of 50 µm × 50 µm, all with the same orientation. The readouts are done on opposite sides for alternating channels.

The FBK SiPMs, shown in Figure 15, also consist of 128 channels, each containing 8 rows and 55 columns of 31 µm × 31 µm cells. The two end columns of each channel are inactive, leaving 53 active columns, in total 424 active cells. All readout is done on the same side of the channel, and the cell orientation is flipped vertically on every second row.

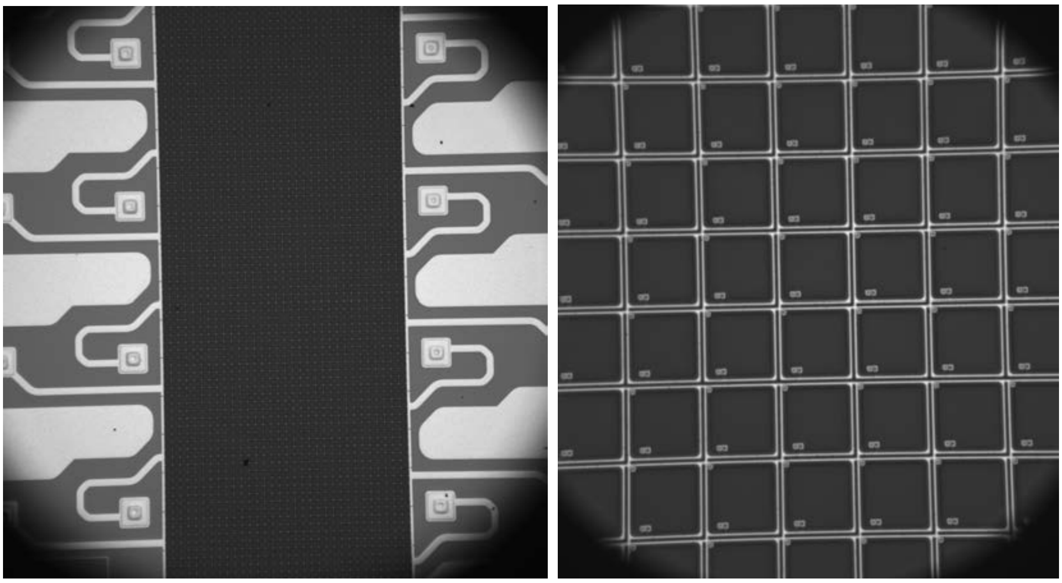

Figure 14: Hamamatsu H2024 SiPM imaged under 5× (**left**) and 50× (**right**) magnification. The 50× view resolves the individual 50 µm × 50 µm microcells, which share a common orientation and are separated by the bright cell walls used later for automatic grid detection.

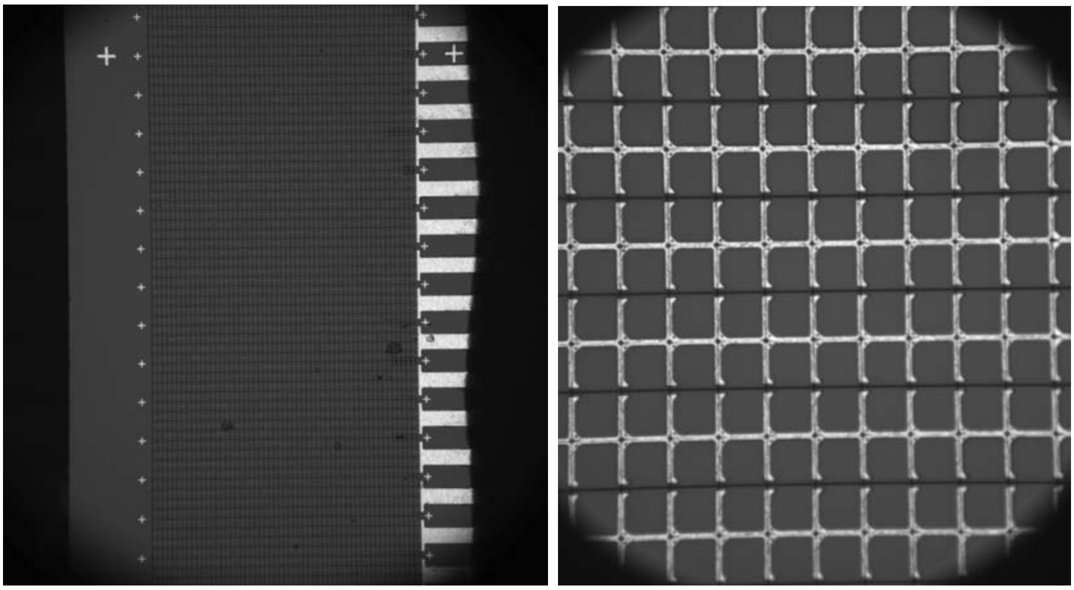

Figure 15: FBK W3m SiPM imaged under 5× (**left**) and 50× (**right**) magnification. The finer 31 µm × 31 µm cells are visible at 50×; note the alternating cell orientation on successive rows and the dim (rather than bright) inter-row gaps, both of which are handled explicitly in the grid-identification step.

2.3 Emission Microscopy

Emission microscopy (EMMI) consists of superimposing a reference image (taken under illumination with a short exposure of ∼3 ms) onto an emission image (taken over an extended period in the dark, always 10 minutes in this work) to visualise the precise location and intensity of the photon emission, as shown in Figure 16. The colour scheme is a customised "plasma" colormap in which the pixel value is coupled to the alpha (opacity) channel, so that dim pixels are rendered transparent and only genuine emission stands out against the reference structure. This provides a spatially resolved map of the radiation damage, or structural defects, within the SiPM.

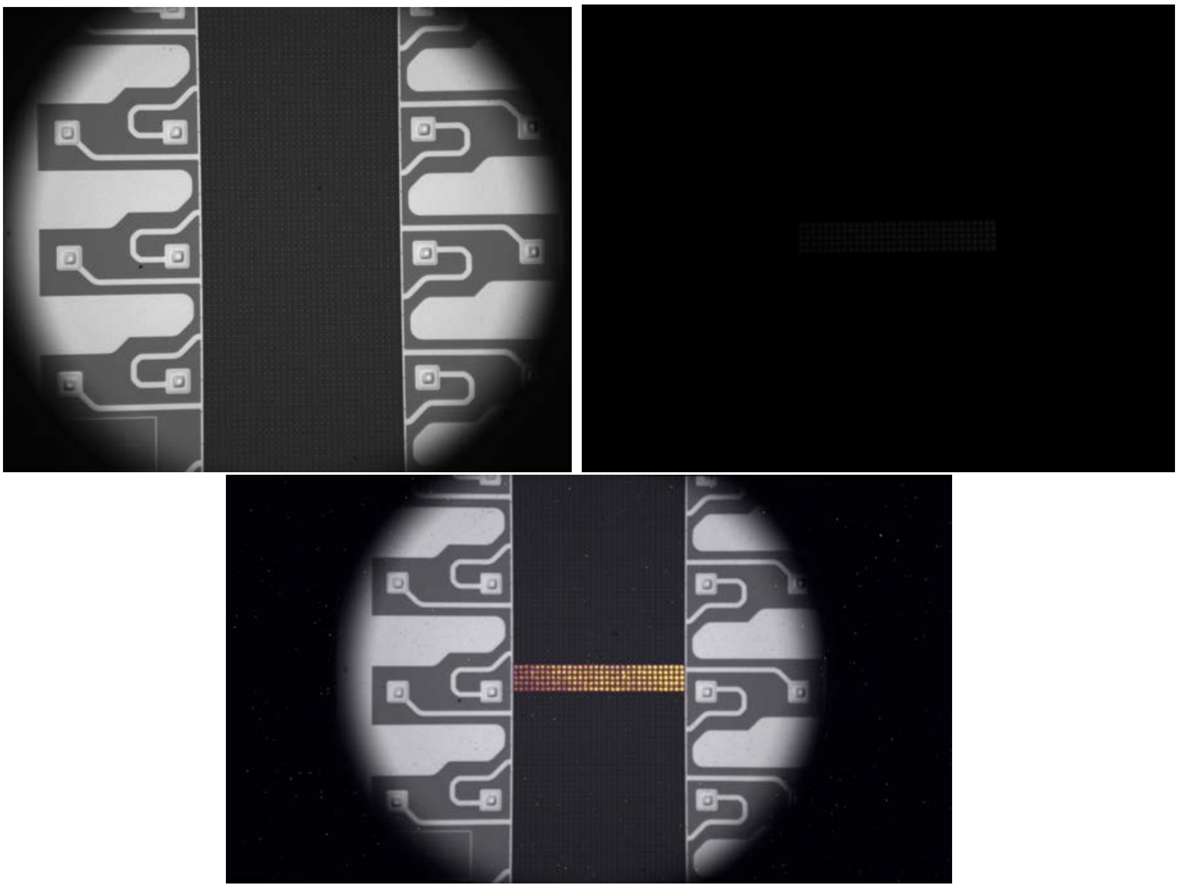

Figure 16: Construction of an EMMI image. **Top left**: reference image, a short illuminated exposure recording the cell structure. **Top right**: emission image, a 10-minute dark exposure recording only the secondary photons from avalanches. **Bottom**: the superimposed EMMI image, in which the alpha-coupled plasma colormap overlays the emission intensity on the cell structure, localising the dark-count activity to specific positions within each cell.

Emission microscopy was first developed and widely used in the semiconductor industry as a failure-analysis tool as early as the 1990s [15]: integrated circuits and power devices emit faint visible and near-infrared light at sites of abnormal current flow, including leakage currents, gate-oxide breakdown, latch-up and hot-carrier injection points. A cooled, high-sensitivity camera can pinpoint these defects non-destructively. The same physical effect, light emission by hot carriers in a high electric field, underlies the photon emission of a SiPM operated above breakdown. Applying EMMI to SiPMs therefore transfers a mature defect-localisation method to the problem of mapping radiation-induced dark-count sites, allowing damage to be studied cell by cell rather than only as a bulk current, and making it a natural candidate diagnostic for the more heavily irradiated sensors of future detectors.

3 Experiment Procedures

3.1 Breakdown Voltage Measurement

The breakdown voltage, V_{bd}, was determined using the Inverse Logarithmic Derivative (ILD) method [16]. In this method the current of an individual SiPM channel is measured as a function of the applied bias voltage in the dark. For the non-irradiated SiPMs, whose dark current is very small, the sensor is instead illuminated to obtain a more precise V_{bd} measurement.

The ILD is defined as

$$\mathrm{ILD}(I(V)) = \left(\frac{d}{dV}\ln I\right)^{-1} = I\frac{dV}{dI}. \tag{4}$$

In the voltage region just above breakdown, and assuming a negligible crosstalk contribution, the SiPM current can be approximated by a power-law dependence on the overvoltage $\Delta V = V - V_{\mathrm{bd}}$:

$$I(V) \propto (V - V_{\mathrm{bd}})^{1+\kappa}, \tag{5}$$

where κ is a constant describing the voltage dependence of the photon detection efficiency (PDE).

Applying the ILD to this relationship yields a linear dependence:

$$\mathrm{ILD}(I(V)) = \frac{1}{1+\kappa}(V - V_{\mathrm{bd}}). \tag{6}$$

Thus the ILD varies linearly with the applied voltage and becomes zero at $V = V_{\mathrm{bd}}$. This property allows for a precise determination of the breakdown voltage by extrapolating the linear region of the ILD curve to the voltage axis. The method enhances accuracy by identifying the inflection point of the logarithmic current curve, corresponding to the

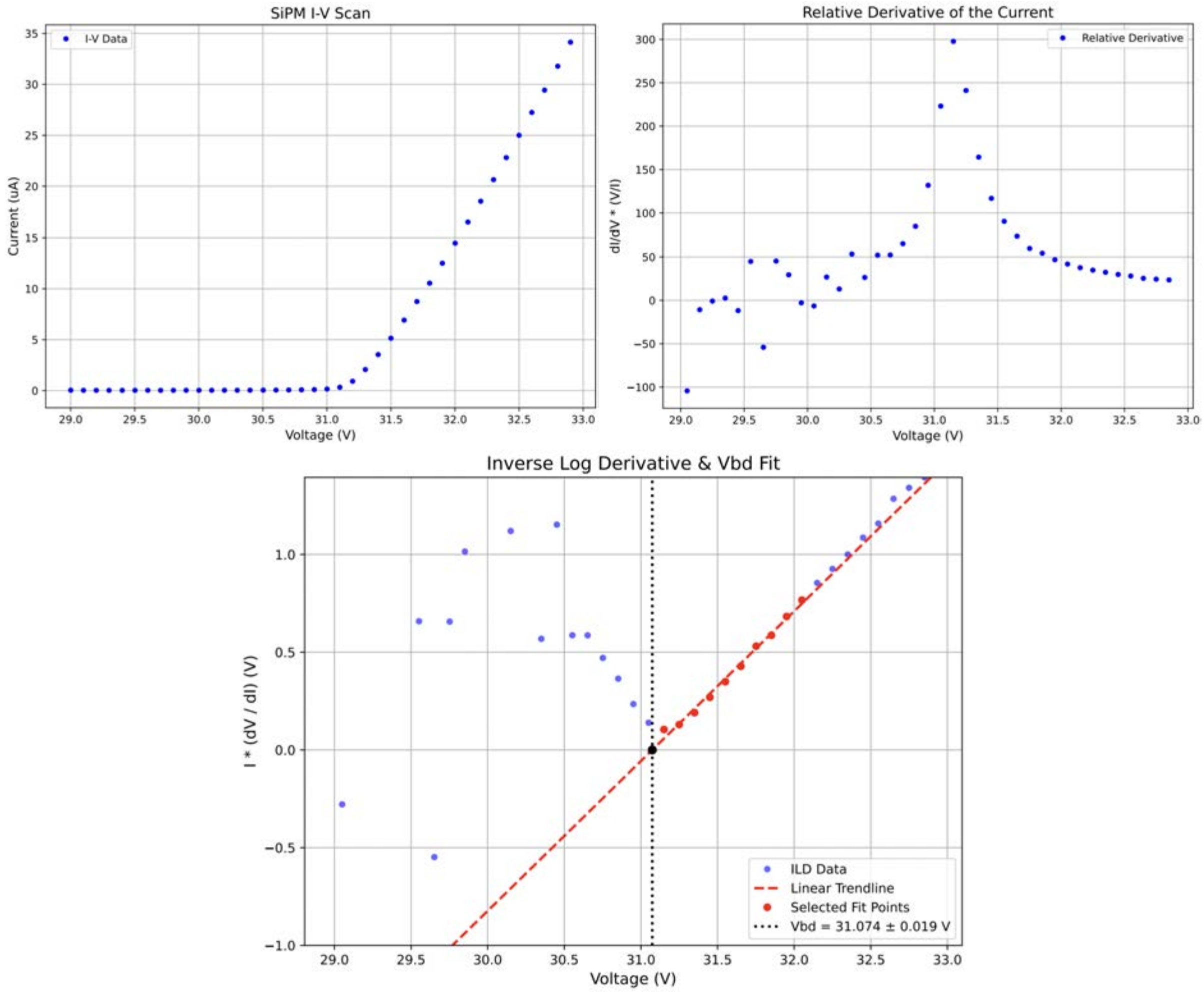

Figure 17: Breakdown-voltage extraction with the ILD method, shown for FBK $1 \times 10^{13}\,\mathrm{n_{eq}/cm^2}$ (channel 118). **Top left**: dark I–V scan. **Top right**: relative derivative $(\mathrm{d}I/\mathrm{d}V) \times (V/I)$ of that scan. The peak of this plot is used to identify the starting point of the ILD linear fit. **Bottom**: the ILD, obtained by inverting the relative derivative and multiplying by V; a linear fit to the region above breakdown extrapolates to V_{bd} at the voltage intercept, giving $31.07 \pm 0.02\,\mathrm{V}$ for this sample.

onset of avalanche multiplication, as illustrated for the FBK $1 \times 10^{13}\,\mathrm{n_{eq}/cm^2}$ sensor in Figure 17. The breakdown voltages obtained for all sensors are listed in Tab. 3.

3.2 Image and Current Taking

Images are acquired with the HCImage Live software driving the qCMOS camera, using the acquisition settings listed in Tab. 1. First the reference image is taken with the light

Figure 18: A photo showing a black cloth covering between the objective lens and the vacuum chuck to block potential faint light produced inside the dark room from electronics such as the power supply and the camera LED.

source on and an exposure of $\mathcal{O}(1\,\text{ms})$; then the emission image is taken with all lights switched off and a black cloth draped over the SiPM and its surroundings, to block stray light from the camera LED or the power supply from reaching the setup, as shown in Figure 18.

While the emission image is being integrated, a parallel program records the current and voltage read back from the Keithley SourceMeter into a CSV file at 5-second intervals. This time series provides the basis for the current (and hence charge) measurement associated with each image, which is later correlated with the emission intensity in Section 4.3. Throughout data taking, the source-meter compliance is capped at $50\,\mu\text{A}$ to protect the SiPMs from potential damage caused by excessive current.

To balance image quality against data-collection efficiency, the $20\times$ objective is used. This requires taking two images per channel, one for the left- and one for the right-hand side, which are then merged to form the image of the whole channel. This process is detailed in Section 3.4.

3.3 Camera Noise Subtraction

The camera contains a number of persistently bright pixels arising from manufacturing defects. To remove this fixed-pattern noise, a dark image is taken with no bias applied to the SiPM, using the same intended exposure time as the measurement (10 minutes), and subtracted from each emission image. The markedly cleaner result in Figure 19 confirms both the effectiveness of the method and the assumption that the defective pixels produce a fixed additive offset.

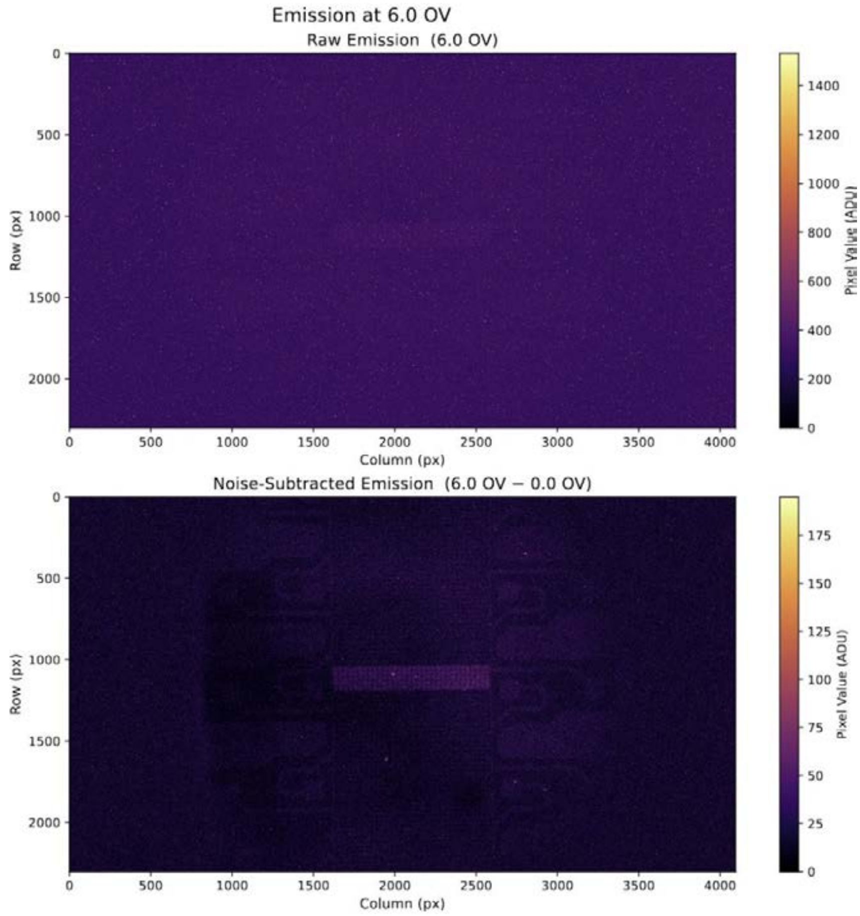

Figure 19: Effect of fixed-pattern-noise subtraction. **Top**: raw 10-minute emission image, dominated by persistently bright defect pixels. **Bottom**: the same image after subtracting an unbiased dark frame of equal exposure; the underlying cell structure and genuine emission become clearly visible, validating the fixed-offset assumption.

The second panel here retains some visible surrounding structure because of ambient light from the camera LED; during the actual data-taking an additional cover is added to block this contribution, as shown in Figure 18.

3.4 Grid Identification

Grid identification enables cell-by-cell hotspot extraction and a systematic characterisation of the in-cell position distribution. It also breaks the very large emission image into small per-cell batches, avoiding computational overload.

As shown in Figure 20, the process involves two steps:

1. **Rotation correction.** The image is rotated so that the variance of the column-summed pixel profile is maximised, which aligns the bright and dark grid lines with the image axes.

2. **Grid-line identification.** The columns and rows with the largest or smallest summed brightness are located. For the bright borders of the Hamamatsu sensors, bright strips are identified as the cell borders, whereas for the dim borders of the FBK sensors, dark strips are identified.

Two parameters control the second step: the cell side length (in pixels), and the brightness prominence (how much brighter or darker a strip must be relative to its surroundings to be accepted as a border, which sets the detection sensitivity).

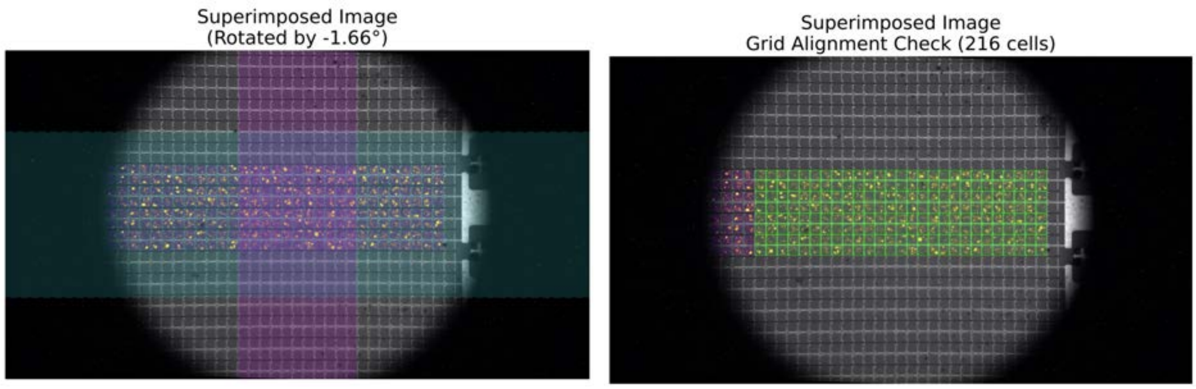

Figure 20: Automatic grid identification. **Left**: a rotated EMMI image; only the reference image is used for the rotation correction, and the highlighted bands are the row/column regions summed to locate the vertical and horizontal grid lines. **Right**: the recovered grid superimposed on the EMMI image. This represents one half of the channel; the opposite half is imaged separately and merged (Figure 21).

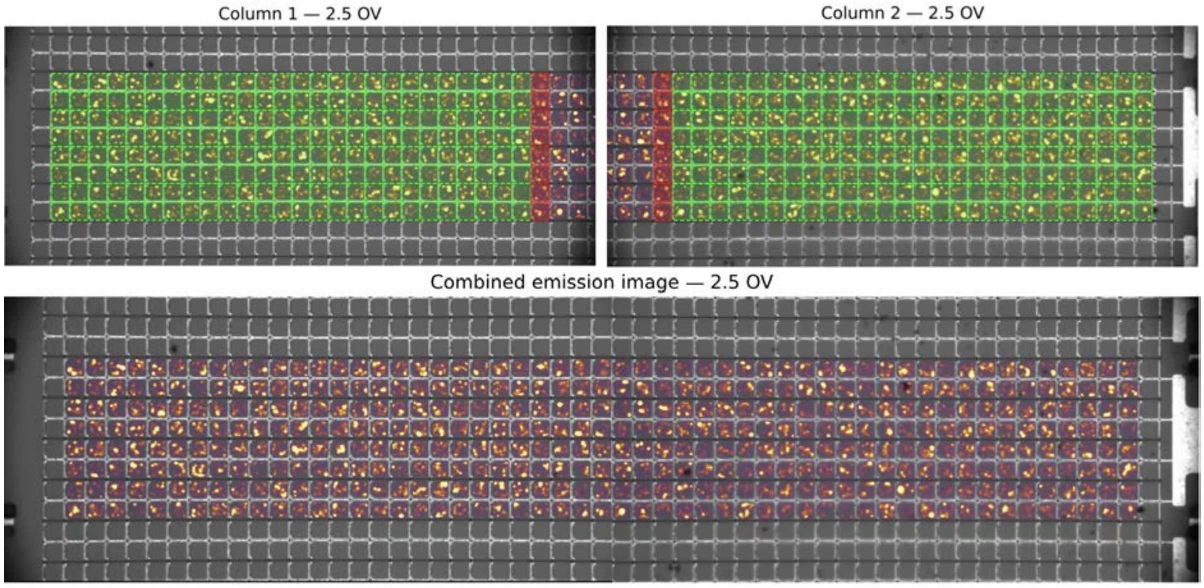

Figure 21: Merging the two halves of a channel, shown for FBK $3 \times 10^{12}\,\mathrm{n_{eq}/cm^2}$ (channel 118) at 2.5 OV. **Top**: the left- and right-hand images of the channel, each acquired separately with the 20× objective. **Bottom**: the two halves stitched into the complete channel image on which the analysis is performed.

After grid identification, the two halves of the channel are merged into the complete emission image, as shown in Figure 21. For the FBK sensors there is an additional dark gap every two rows (Figure 15), the row identification is therefore performed twice, once detecting the bright strip every two rows and once the dim strip every two rows, and the two sets of rows are combined to form the complete set of horizontal grid lines.

Owing to the finite thickness of the cell walls and the optical distortion of the camera, the borders are not perfectly aligned for every cell. The residual misalignment is approximately 2 pixels, which is small compared with the cell-wall thickness of about 10 pixels; we therefore consider this error negligible for the purposes of this project.

3.5 Hotspot Extraction

Zooming in on each individual cell, the emission hotspots can be extracted and reconstructed as 2D Gaussians [17]. This approach allows the separation of two closely spaced hotspots without missing dim hotspots.

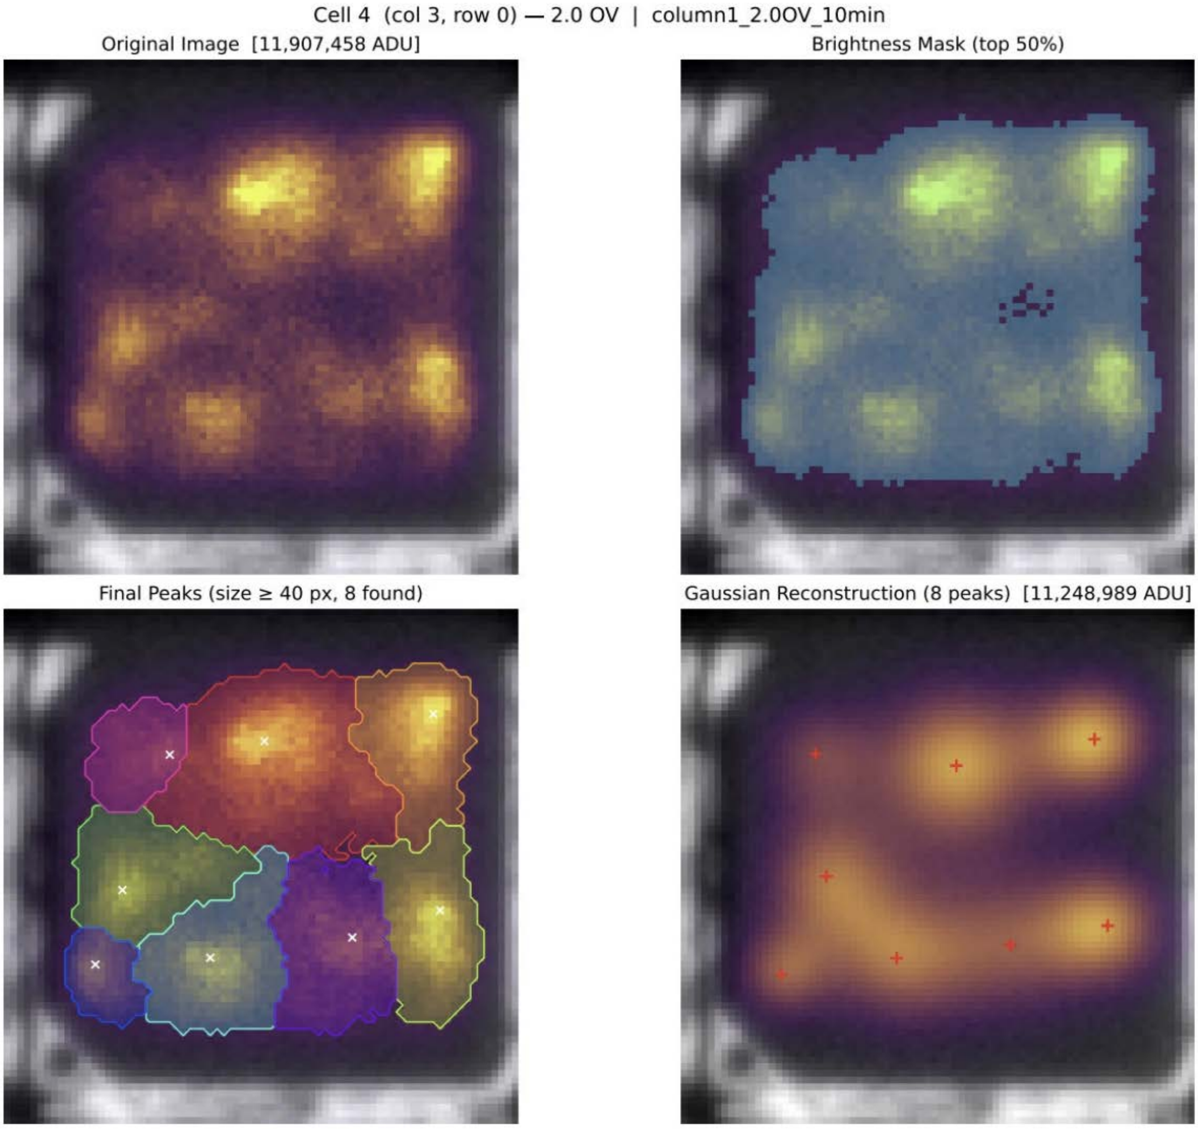

Figure 22: The hotspot-identification pipeline. **Top left**: original per-cell emission image. **Top right**: active-pixel mask from an adaptive percentile threshold (here the top 50% of pixels). **Bottom left**: local-maximum detection followed by watershed segmentation, assigning each above-threshold pixel to a peak. **Bottom right**: the emission reconstructed from the resulting 2D-Gaussian parameters (location, amplitude, width), which reproduces the original while yielding a quantitative description for every hotspot.

As shown in Figure 22, our method extracts the hotspot information in the following steps:

1. **Adaptive thresholding.** A percentile threshold is used instead of a hard-coded

value, so that even in dim-hotspot scenarios such as low overvoltage the mask still extracts the hotspot features reliably.

2. **Peak detection and segmentation.** Two functions are applied to the masked pixels: first `peak_local_max` from `skimage.feature` identifies the local maxima. This function requires a `MIN_DISTANCE` input that must be tuned to avoid over- or under-capturing the maxima. The identified peaks are then passed to the `watershed` function from `skimage.segmentation` to assign the area belonging to each peak.

3. **Conversion to 2D-Gaussian parameters.** The pixels of each peak are converted into the parameters of a 2D Gaussian:

   - **Location:** the intensity-weighted mean of the peak pixels;
   - **Amplitude:** the mean of the three brightest pixels of the peak;
   - **Sigma:** tuned so that the integrated intensity of a circle of the same area as the original peak equals the total intensity contained within that peak.

One drawback of this method is that each reconstructed 2D Gaussian contributes to its neighbours, so the summed reconstruction is slightly brighter than the original image. A final scaling is therefore applied so that the total intensity of the reconstructed image within the masked "active cell region" matches that of the original; in practice this lowers the output amplitudes by about 5% relative to the true peak values. In principle the bias could be removed by simultaneously fitting all 2D Gaussians in a cell, but given the large number of cells and the limited computing power available, this is computationally impractical, therefore a more efficient analysis method is left for future work.

Figure 23 compares the original and reconstructed emission images for the non-irradiated Hamamatsu sensor, and Figure 24 does the same for the $1 \times 10^{13}\,\mathrm{n_{eq}/cm^2}$ sensor. In both the sparse (low-density) and the crowded (high-density) hotspot regimes the method captures and reconstructs the hotspot information well. A systematic figure of merit for the reconstruction quality has not yet been developed; a natural candidate would be the pixel-wise squared-intensity difference between the original and reconstructed images. This is left for future work to validate.

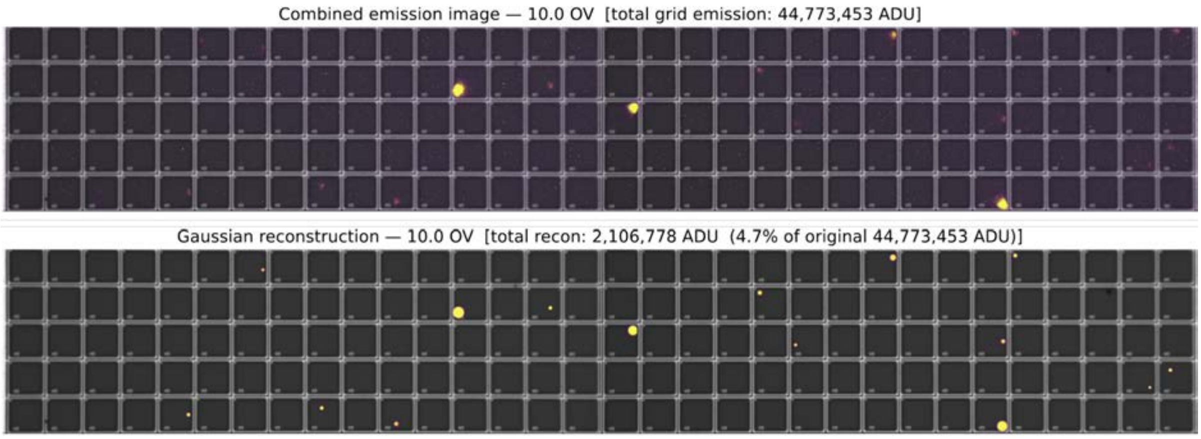

Figure 23: Reconstruction validation in the low-density regime: non-irradiated Hamamatsu sensor at 10.0 OV. **Top**: original emission image, with well-separated hotspots. **Bottom**: the corresponding 2D-Gaussian reconstruction, which reproduces the sparse hotspots faithfully.

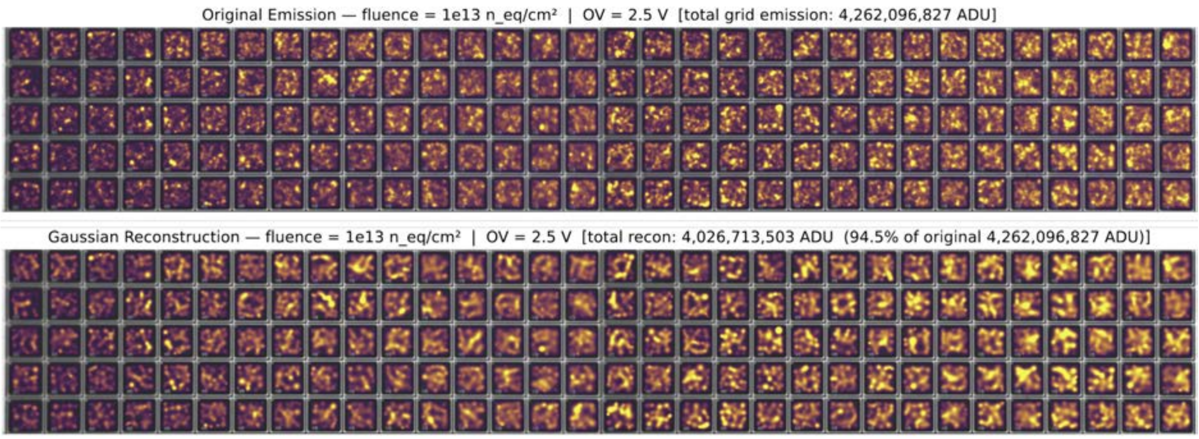

Figure 24: Reconstruction validation in the high-density regime: Hamamatsu $1 \times 10^{13}\,\mathrm{n_{eq}/cm^2}$ sensor at 2.5 OV. **Top**: original emission image, densely populated with overlapping hotspots. **Bottom**: the 2D-Gaussian reconstruction, which still recovers the individual hotspots despite the crowding, though this regime is where residual peak-merging is most likely.

Four parameters are involved in the method:

1. `DETECTION_THRESHOLD_PCT`: the percentile of the brightest pixels used to set the active-pixel threshold;
2. `MIN_DISTANCE`: the minimum distance between local-maximum peaks;

3. N_PEAK_BRIGHTNESS: the number of brightest pixels used to determine the amplitude of the 2D Gaussians;

4. MIN_PEAK_SIZE: the minimum peak size for a peak to be reconstructed as a 2D Gaussian.

Given a metric for the goodness of reconstruction and sufficient computing power, these parameters could also be optimised to deliver the best reconstruction. Achieving a better reconstruction directly improves the fidelity of the extracted damage-site information, and hence the precision of all the analyses that follow.

4 Results and Discussion

4.1 Current-Voltage Relationship

Figures 25 and 26 show the dark I–V scans for the FBK and Hamamatsu sensors at all fluences. The current at a given overvoltage rises with fluence, as expected from the increasing density of generation centres, but the rise is sub-linear in fluence, pointing to a saturation associated with the microcell recovery time. In the unsaturated regime each avalanche releases the full charge corresponding to the gain at that overvoltage, and since the avalanche rate scales with the defect density (and hence fluence), the current should be linear in fluence. Saturation occurs when cells are re-triggered before the field has fully recharged over the recovery time τ_{rec}, so that less than the nominal gain charge is released per avalanche. The practical consequence is that the **effective overvoltage** seen by a heavily irradiated cell is lower than the nominal applied value, a caution that must be kept in mind when comparing quantities across fluences at fixed nominal OV.

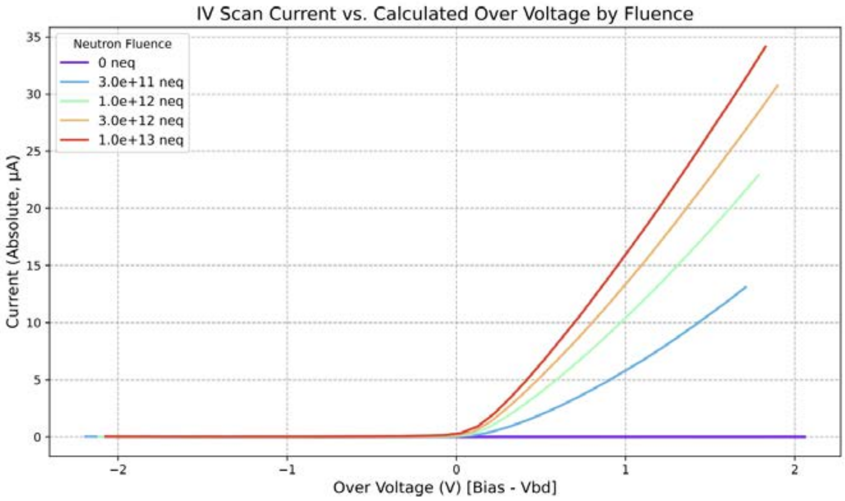

Figure 25: Dark current-voltage relationship of the FBK sensors at the five radiation levels. The current at fixed overvoltage grows with fluence but not linearly in it: the compression at high fluence is the signature of recovery-time saturation, in which cells re-fire before fully recharging, lowering the effective overvoltage.

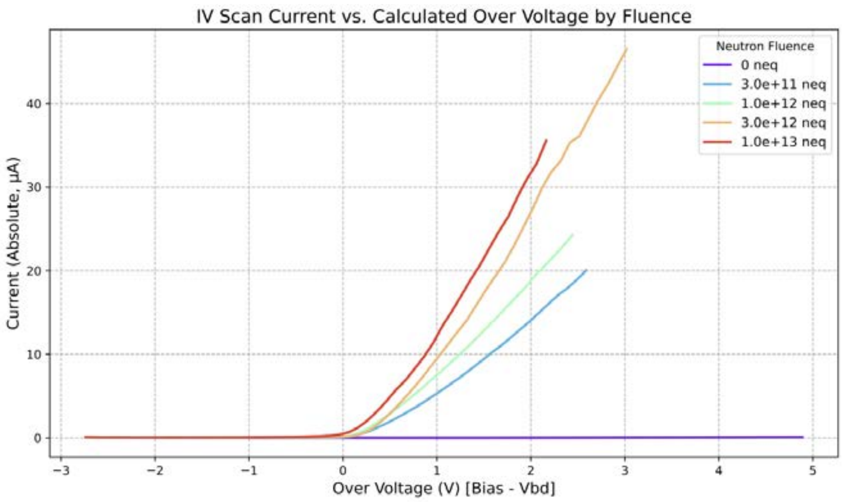

Figure 26: Dark current-voltage relationship of the Hamamatsu sensors at the five radiation levels. As for FBK (Figure 25), the current increases with fluence sub-linearly, indicating the same recovery-time saturation and a correspondingly reduced effective overvoltage at the highest fluences.

4.2 Averaged Emission Intensity by Cell

A useful first-order observable is the emission intensity averaged over each cell. For the non-irradiated FBK sensor at 10.0 OV, two cells stand out as 4 to 5 times brighter than the rest (Figure 27), and the per-cell intensity histogram confirms these two outliers sitting well above the bulk of the distribution. Such isolated bright cells are the optical signature of localised manufacturing defects.

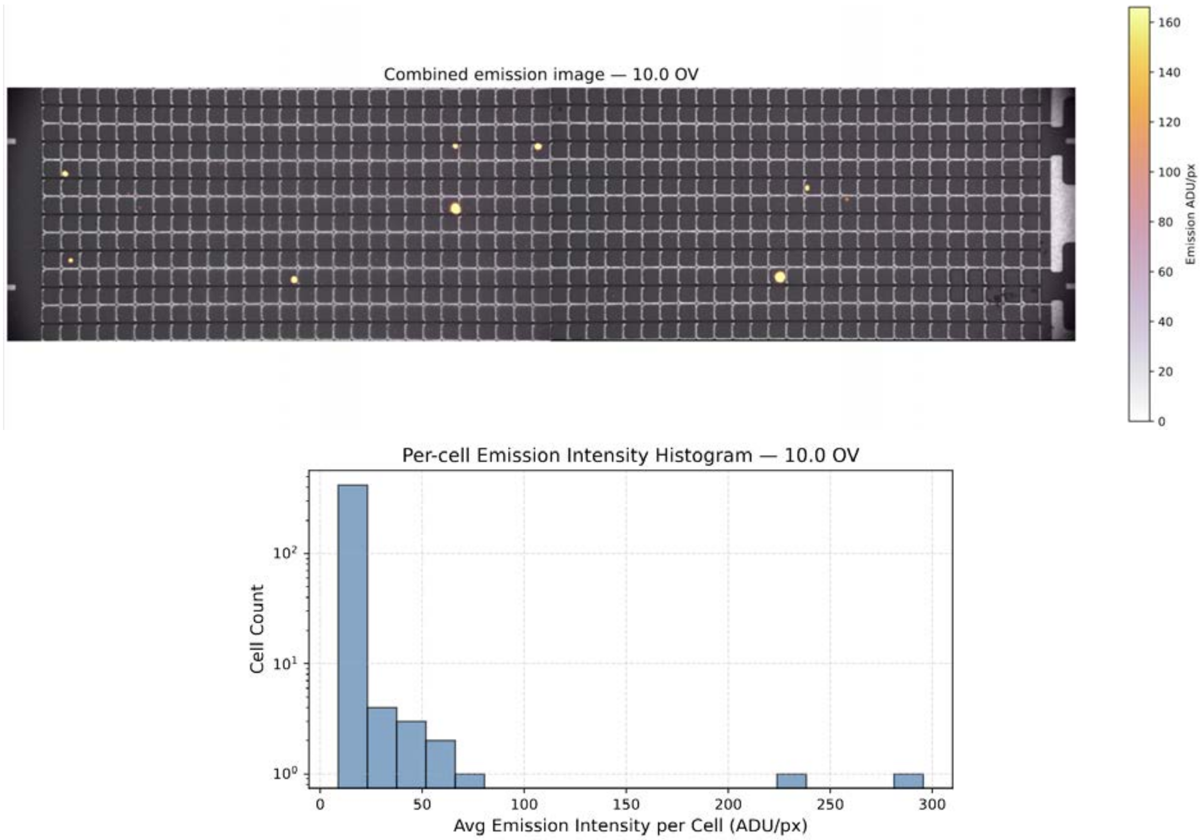

Figure 27: Non-irradiated FBK sensor at 10.0 OV. **Top**: EMMI image, in which two cells emit 4–5× more strongly than their neighbours. **Bottom**: histogram of the per-cell averaged intensity, where these two cells appear as clear outliers above the otherwise narrow bulk distribution, identifying them as isolated severe manufacturing defects.

By contrast, for the heavily irradiated FBK $1 \times 10^{13}\,\mathrm{n_{eq}/cm^2}$ sensor at 2.5 OV the damage is spread far more uniformly across the cells, as seen in Figure 28.

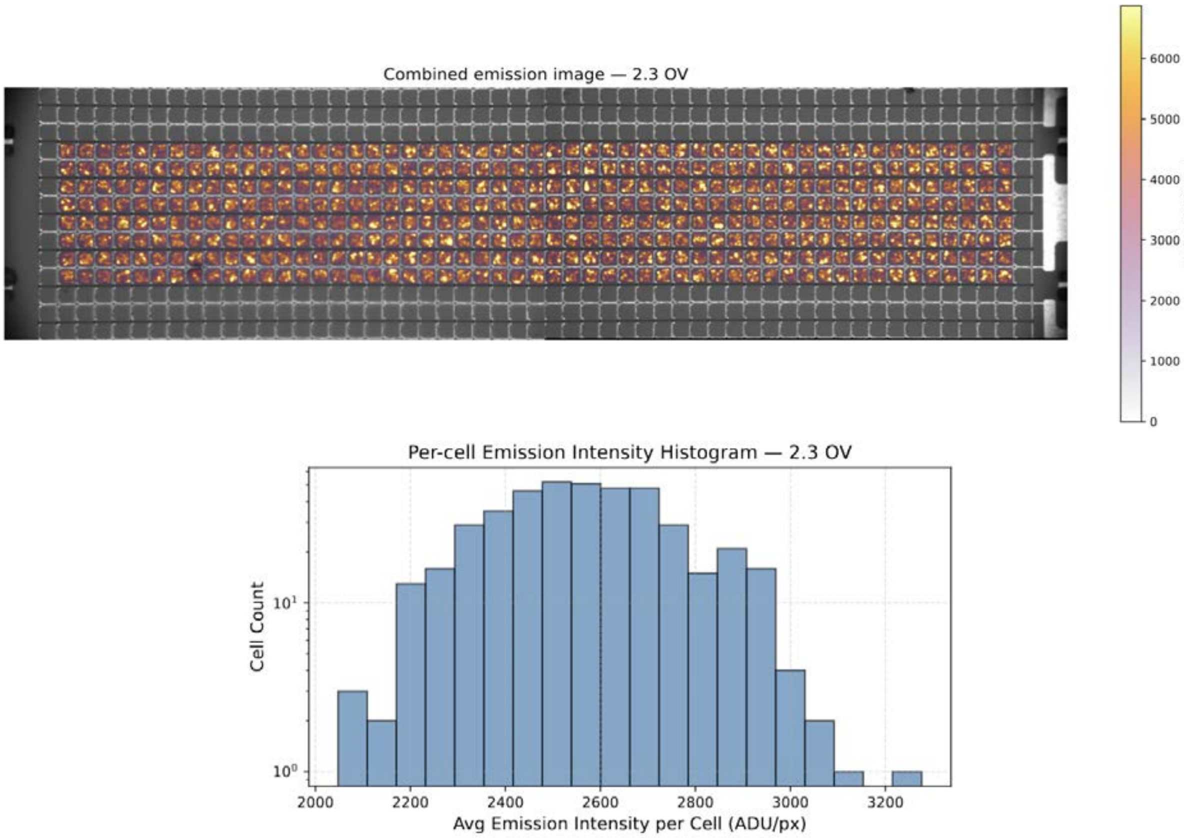

Figure 28: FBK $1 \times 10^{13}\,\mathrm{n_{eq}/cm^2}$ sensor at 2.5 OV. **Top**: EMMI image, showing emission distributed almost uniformly across every cell. **Bottom**: the corresponding per-cell intensity histogram, now a single peak with no outliers, reflecting the spatially uniform neutron-induced damage that dominates over any manufacturing defect.

In summary, the non-irradiated sensors contain a few cells markedly brighter (by a factor of 4 to 5) than the rest due to isolated manufacturing defects, whereas after irradiation this contrast is washed out: the uniform neutron exposure introduces bulk damage everywhere that dominates the localised defects, and the higher dark current of the irradiated sensors causes them to reach the 50 µA compliance limit at a much lower overvoltage, preventing them from being operated at the high overvoltage (10.0 V) at which the manufacturing outliers stand out most prominently.

4.3 Emission Intensity vs Current

Both FBK and Hamamatsu show a linear relationship between the averaged emission intensity and the mean current, as seen in Figures 29 and 30 respectively. For each manufacturer this linearity is universal across the different fluences: the emitted light tracks the total avalanche charge irrespective of how the damage that produced that charge was distributed. The FBK slope is steeper than the Hamamatsu one, signalling a higher number of photons emitted per unit charge (and hence per charge carrier).

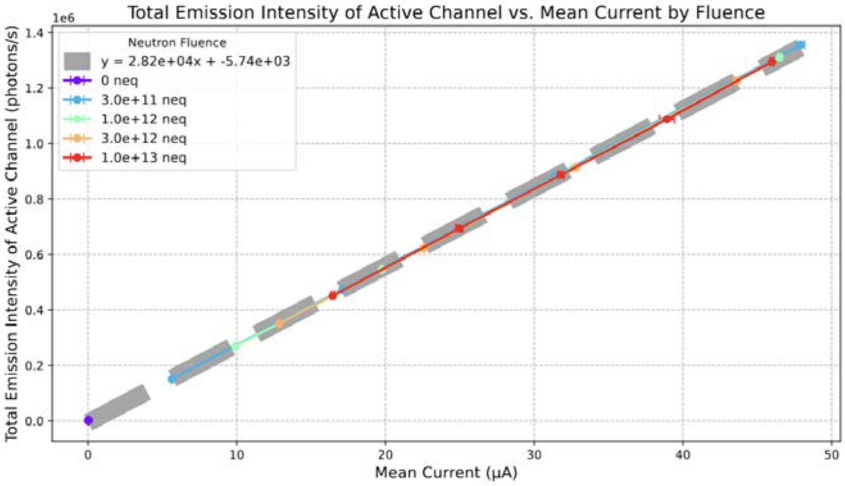

Figure 29: Averaged emission intensity versus mean current for the FBK sensors. All fluences fall on a single line of slope 2.82×10^4 photons/µC, corresponding to 4.51×10^{-9} detected photons per charge carrier. The fluence-independence of the slope confirms that the emitted light is a direct proxy for the avalanche charge.

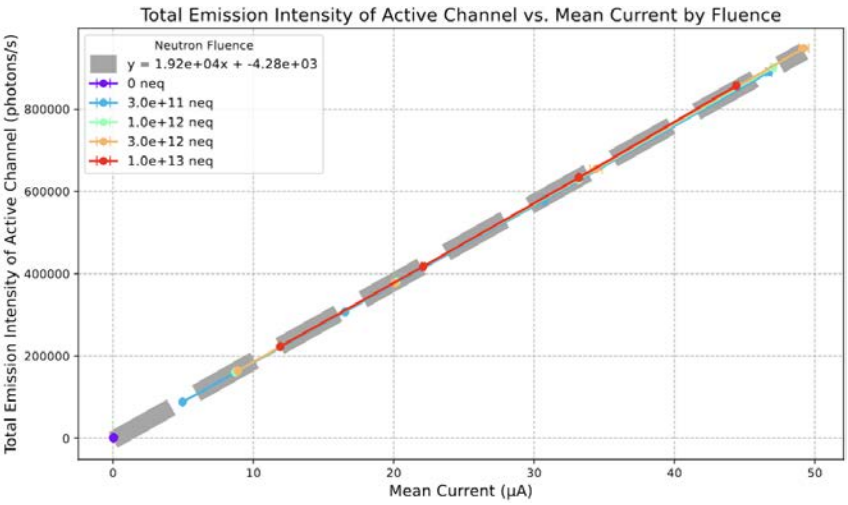

Figure 30: Averaged emission intensity versus mean current for the Hamamatsu sensors. As for FBK, all fluences share a single linear relationship, here with the shallower slope of 1.92×10^4 photons/µC (3.07×10^{-9} detected photons per charge carrier), indicating a lower detected-photon yield per unit charge than the finer-pitched FBK device.

From the two fits we obtain the number of photons detected per charge carrier; converting this to the number of photons produced per charge carrier (Section 4.4) then allows a comparison with the work of McLaughlin et al. [18]. The measured detected yields are

$$\begin{aligned}
\text{FBK:}\quad & 2.82 \times 10^4 \text{ photons/µC} \;\rightarrow\; 4.51 \times 10^{-9} \text{ photons detected per charge carrier,}\\
\text{HPK:}\quad & 1.92 \times 10^4 \text{ photons/µC} \;\rightarrow\; 3.07 \times 10^{-9} \text{ photons detected per charge carrier.}
\end{aligned}$$

4.4 Number of Photons per Charge Carrier Estimation

Several factors limit the fraction of produced photons that reach the camera:

1. The quantum efficiency of the camera (Figure 31). Since our setup cannot measure the fill factor or the avalanche-trigger probability of the camera, the manufacturer's quantum efficiency is used directly as the photon detection efficiency.

2. The SiPM emission spectrum. Lacking a spectrometer, we adopt the spectrum measured in the work of McLaughlin et al. [18] as an approximation, shown in Figure 32.

3. The maximum emission angle from within the silicon that still reaches the microscope's field of view, set by the refractive indices of Si and SiO_2, and the numerical aperture of the objective.

4. Attenuation as the photon travels through the media, set by the depth of the multiplication region.

5. Reflection at the interfaces between different materials, set by the refractive indices of each medium.

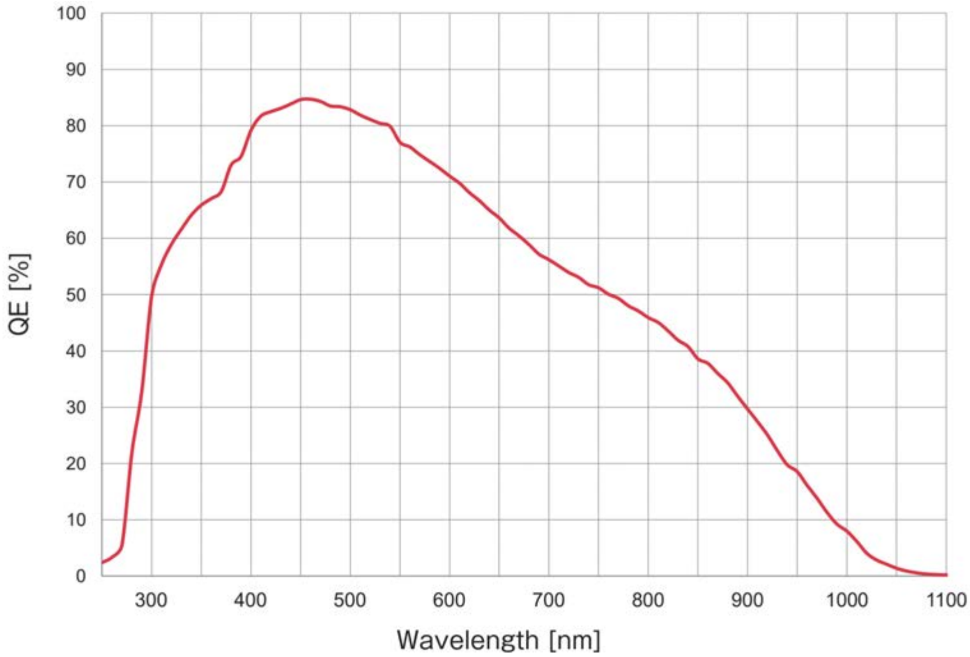

Figure 31: Quantum efficiency of the Hamamatsu ORCA-Quest 2 camera versus wavelength [19], used here as the photon detection efficiency. The efficiency falls steeply towards the near-infrared, strongly suppressing the detected fraction of the silicon emission that peaks in that range and thereby dominating the overall detection efficiency.

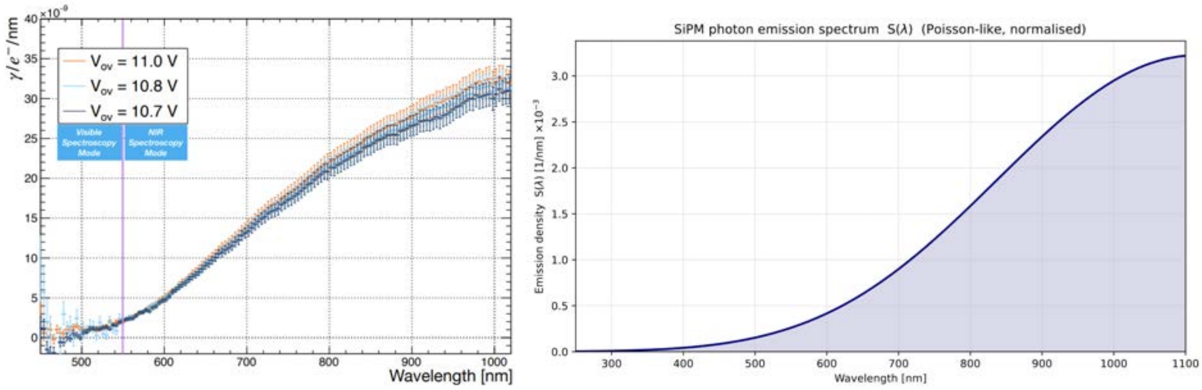

Figure 32: **Left**: emission spectrum of a Hamamatsu VUV4 SiPM as measured by McLaughlin et al. [18]. **Right**: the probability density function of the SiPM emission spectrum adopted from the left measured spectrum and used to weight the wavelength-dependent detection factors.

The first two factors listed above are accounted for using the figures above; the last three are combined into a wavelength-dependent Photon Acceptance Correction Factor $A(\lambda)$, through the expression [18]:

$$
\begin{aligned}
A(\lambda) &= \frac{\int_0^{2\pi} d\phi}{\int_0^{2\pi} d\phi \int_0^{\pi} \sin\theta \, d\theta} \int_0^{\theta_{\mathrm{Si}}} e^{-\frac{d_{\mathrm{P}}}{\cos\theta \, \mu(\lambda)}} \left(1 - R_{\mathrm{Si,SiO_2}}(\lambda,\theta)\right) \left(1 - R_{\mathrm{SiO_2,Atm}}(\lambda,\theta')\right) \sin\theta \, d\theta \\
&= \frac{1}{2} \int_0^{\theta_{\mathrm{Si}}} e^{-\frac{d_{\mathrm{P}}}{\cos\theta \, \mu(\lambda)}} \left(1 - R_{\mathrm{Si,SiO_2}}(\lambda,\theta)\right) \left(1 - R_{\mathrm{SiO_2,Atm}}(\lambda,\theta')\right) \sin\theta \, d\theta,
\end{aligned}
\tag{7}
$$

where the symbols denote:

- θ, ϕ: the polar and azimuthal emission angles of the photon inside the silicon, with the prefactor $\frac{1}{2}$ normalising the solid angle of an isotropic point source;
- θ_{Si}: the maximum polar angle inside the silicon that can still be collected by the objective (the acceptance half-angle);
- d_{P}: the depth of the photon-production (multiplication) region below the surface;
- $\mu(\lambda)$: the absorption length of the photon in silicon, so that $\exp[-d_{\mathrm{P}}/(\cos\theta \, \mu(\lambda))]$ is the attenuation along the slant path to the surface;
- $R_{\mathrm{Si,SiO_2}}$ and $R_{\mathrm{SiO_2,Atm}}$: the reflectivities at the Si/SiO_2 and SiO_2/atmosphere interfaces, removing the fraction of light lost to reflection at each boundary;
- θ': the refraction angle inside the SiO_2 layer, related to θ through Snell's law.

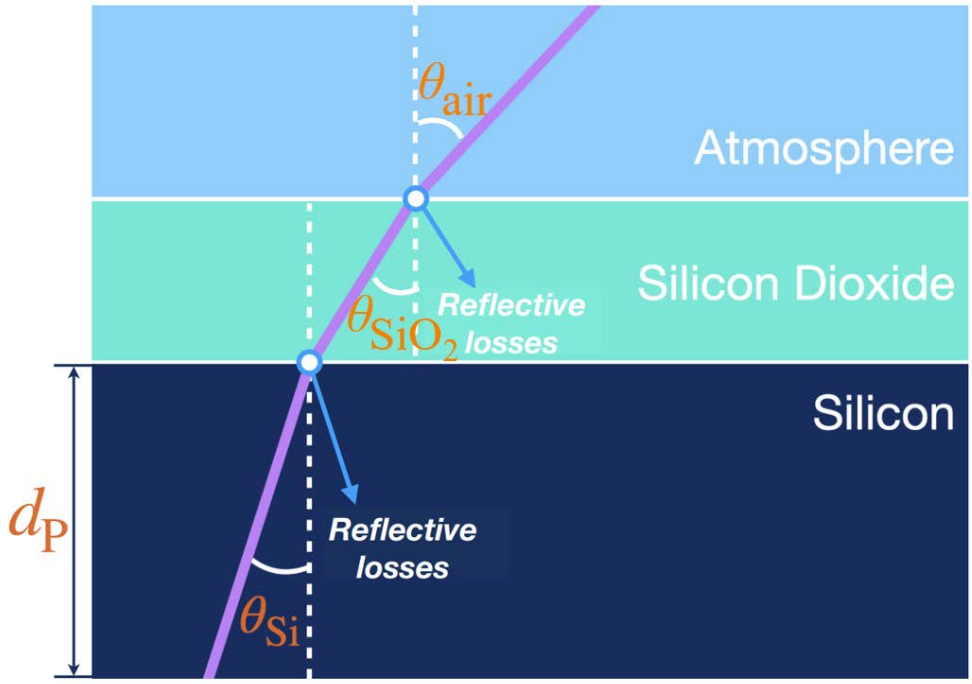

Figure 33: Geometry of the refraction and reflection of a secondary photon as it propagates from its production depth in the silicon, through the SiO_2 layer, to the objective [18]. The angle $\theta_{\rm Si}$ corresponds to the maximum of the emission angle θ, and $\theta_{\rm SiO_2}$ corresponds to the maximum of θ'. The interface reflectivities entering Eq. 7 are also illustrated on this diagram.

The refraction and reflection of photons are visualised in Figure 33. The acceptance angle $\theta_{\rm Si}$ follows from Snell's law applied across the Si/SiO_2/air stack [18]:

$$\mathrm{NA} \equiv n_{\rm air} \sin\theta_{\rm air} = n_{\rm SiO_2} \sin\theta_{\rm SiO_2} = n_{\rm Si} \sin\theta_{\rm Si}, \tag{8}$$

$$\therefore \quad \theta_{\rm Si} = \sin^{-1}\left(\frac{\mathrm{NA}}{n_{\rm Si}}\right), \tag{9}$$

where the numerical aperture $\mathrm{NA} = n_{\rm air} \sin\theta_{\rm air}$ is fixed by the 20× collection objective (Tab. 2), whose angular opening is measured to be $\theta_{\rm air} = 26.57°$.

Assuming unpolarised light, Fresnel's formula gives the reflectivity at each interface:

$$R = \frac{1}{2}\left(|r_s|^2 + |r_p|^2\right),$$

where the reflection coefficients for s-polarised (perpendicular) and p-polarised (parallel) light are

$$r_s = \frac{n_1 \cos\theta_1 - n_2 \cos\theta_2}{n_1 \cos\theta_1 + n_2 \cos\theta_2}, \quad r_p = \frac{n_2 \cos\theta_1 - n_1 \cos\theta_2}{n_2 \cos\theta_1 + n_1 \cos\theta_2}.$$

Figure 34 shows the resulting reflectivity at the highest collectable angle for the two interfaces as a function of wavelength.

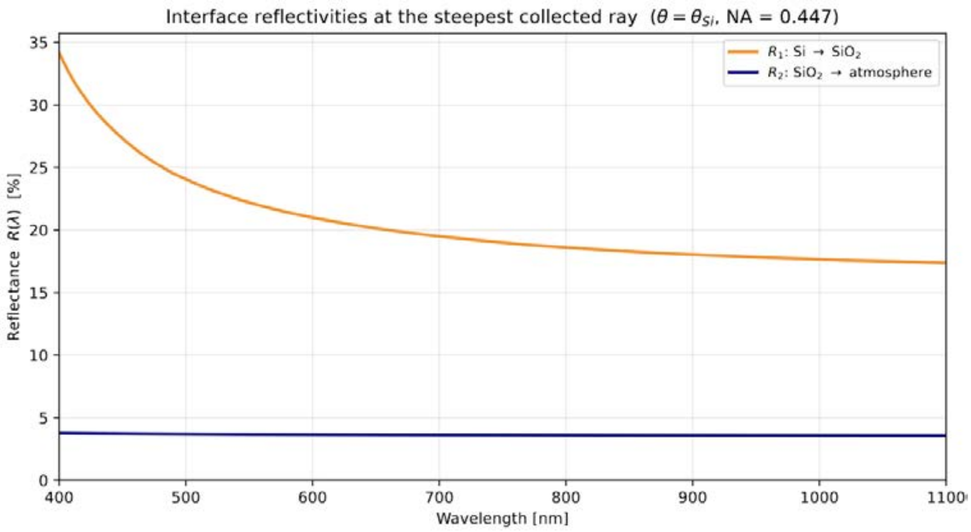

Figure 34: Fresnel reflectivity at the maximum collectable angle for the Si/SiO_2 and SiO_2/atmosphere interfaces as a function of wavelength, computed for unpolarised light. These curves supply the $(1 - R)$ transmission factors in Eq. 7 and quantify the fraction of emission lost to reflection at each boundary.

The production depth d_P is taken from the paper of Gallina et al. [20], which reports 1.8 ± 0.1 µm for the Hamamatsu H2017 SiPM and 0.145 ± 0.01 µm for the FBK LF SiPM. Because of the short absorption length of SiO_2 (Figure 35), absorption within the SiO_2 itself is neglected.

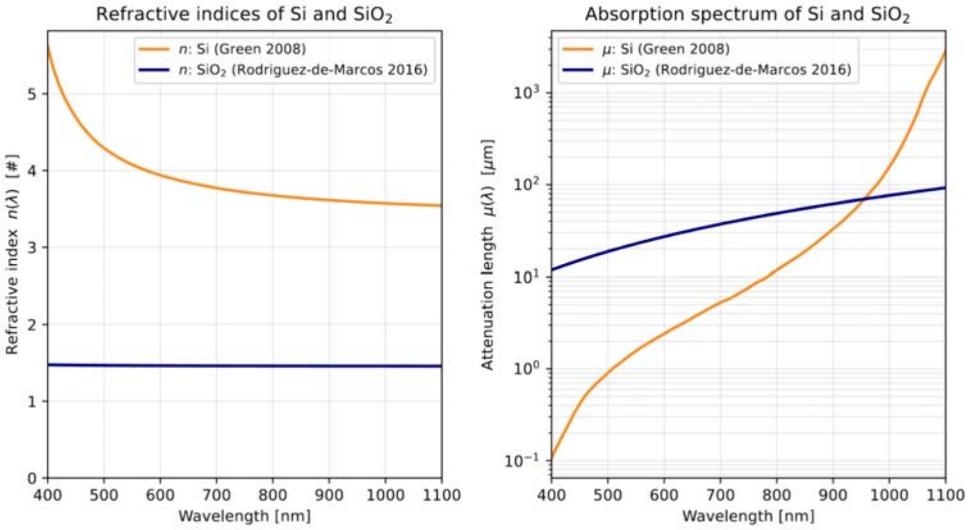

Figure 35: Optical properties of the two relevant media. **Left**: refractive indices of Si [21] and SiO_2 [22]. **Right**: photon absorption length in the two materials; the very short absorption length in silicon in the visible/UV explains why only near-infrared emission from the production depth escapes to be detected.

Combining the attenuation, refraction and reflection terms, the photon-acceptance correction factor of Eq. 7 can be evaluated; the result is shown in Figure 36.

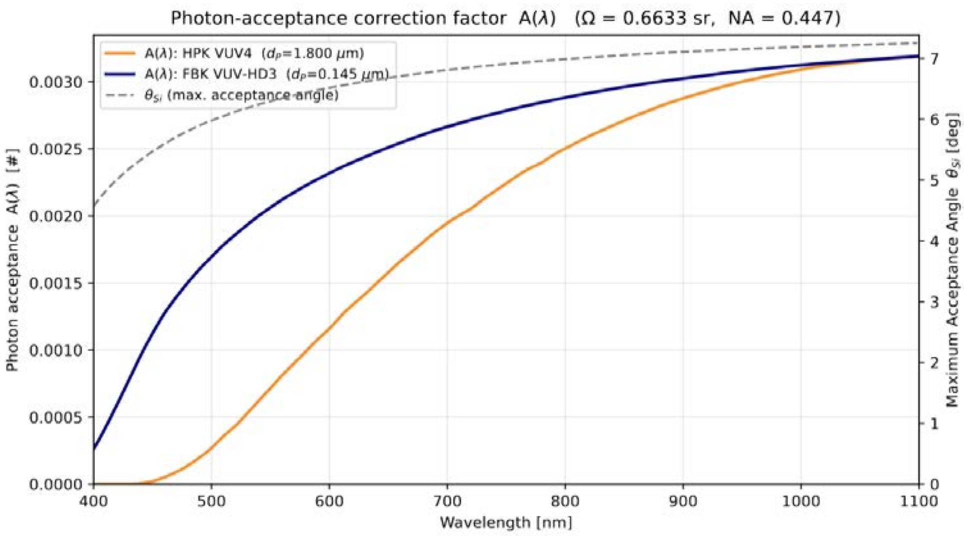

Figure 36: The photon-acceptance correction factor $A(\lambda)$ of Eq. 7, combining the geometric acceptance, path attenuation in silicon, and the Fresnel reflection losses at the Si/SiO_2 and SiO_2/air interfaces. It represents the wavelength-dependent probability that a photon produced at depth d_P escapes into the acceptance cone of the objective.

Multiplying this factor by the camera quantum efficiency (Figure 31) and the SiPM emission spectrum (Figure 32) yields the overall detection density, shown in Figure 37.

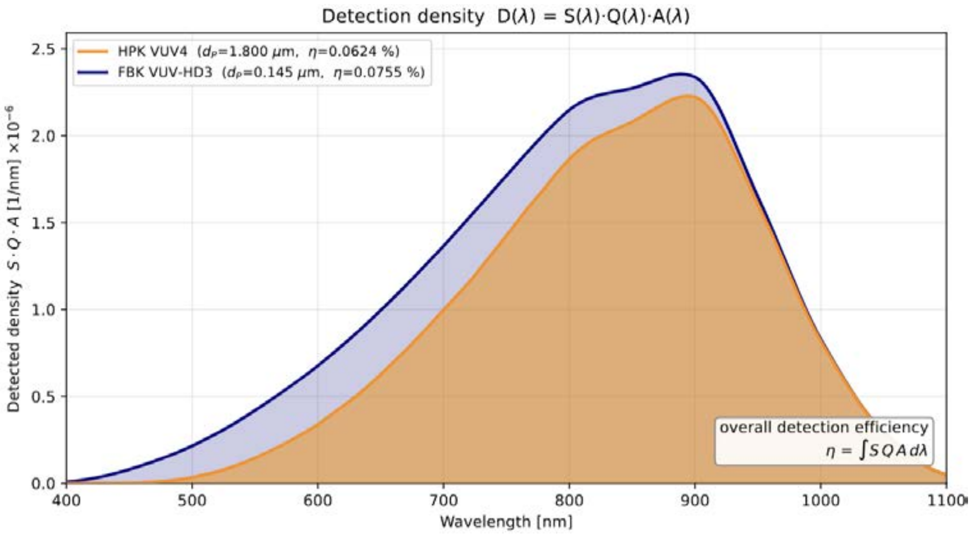

Figure 37: Overall photon detection density versus wavelength, formed as the product of the emission spectrum (Figure 32), the camera quantum efficiency (Figure 31) and the acceptance factor $A(\lambda)$ (Figure 36). Integrating this curve over wavelength gives the overall detection efficiency used to convert detected into produced photons: 0.0755% for FBK and 0.0624% for Hamamatsu.

Integrating the detection density over wavelength gives an overall detection efficiency of 0.0755% for FBK and 0.0624% for Hamamatsu. Dividing the measured detected-photon yields per charge carrier by these efficiencies gives the number of photons produced per charge carrier:

$$\text{FBK:} \quad \frac{4.512 \times 10^{-9}}{0.0755\%} = 5.98 \times 10^{-6} \text{ photons produced per charge carrier,}$$

$$\text{HPK:} \quad \frac{3.072 \times 10^{-9}}{0.0624\%} = 4.92 \times 10^{-6} \text{ photons produced per charge carrier.}$$

FBK has a slightly higher photon-production rate per charge carrier, which may be attributed to its smaller, more densely packed cells, which recharge more readily and sustain a higher local field at a given overvoltage. This stronger electric field produces more energetic (hotter) carriers and therefore more secondary photons per avalanche.

Compared with the work of McLaughlin et al. [18], which quotes about 5×10^{-6} for FBK and 9×10^{-6} for Hamamatsu, our values reproduce the same order of magnitude despite the limitations of our setup, supporting the validity of the analysis method. The relative ordering of the two manufacturers differs from that reference, which is plausible given the different sensor generations and the approximations made in the spectrum and acceptance modelling.

4.5 Hotspot Number

We next study the number of hotspots per cell for the two manufacturers, shown in Figure 38 for FBK and Figure 39 for Hamamatsu. Both plots are capped at 5.0 OV because the non-irradiated ($0\,n_{eq}/cm^2$) sample stays near zero counts until 10.0 OV. The count is only indicative, as it depends on the peak-separation distance chosen during peak identification.

The number is stable across overvoltages at a fixed fluence, as expected, since the number of damage sites is set by the irradiation rather than the operating electric field. As overvoltage increases, the count tends to decrease slightly: the hotspots grow in area and brightness and progressively merge, so the peak-finder function becomes less effective in resolving them individually.

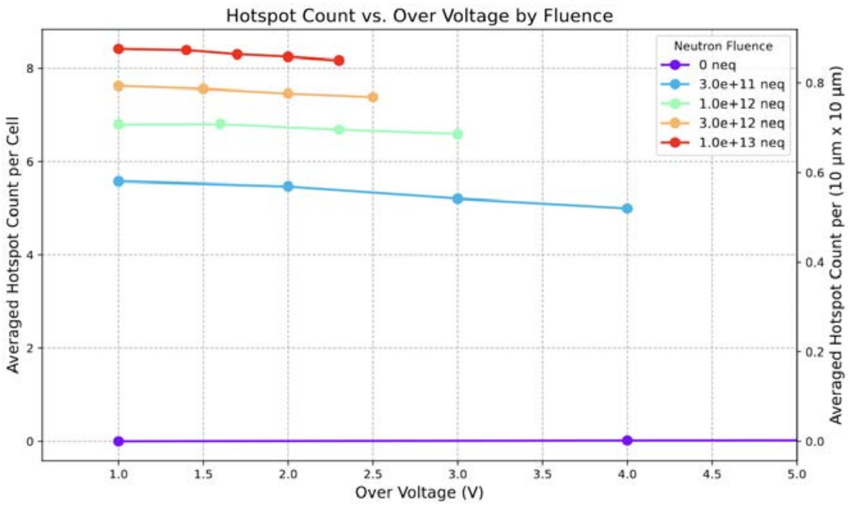

Figure 38: Number of hotspots per cell for the FBK sensors versus overvoltage, at each fluence. The count is roughly constant with overvoltage at fixed fluence, confirming that the number of damage sites is fixed by the irradiation; the slight decrease towards high overvoltage reflects the merging of neighbouring hotspots rather than a loss of damage sites.

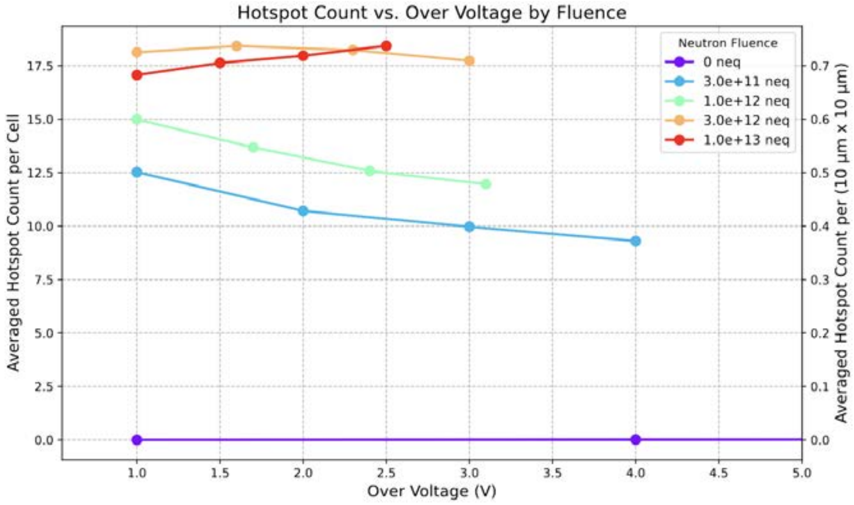

Figure 39: Number of hotspots per cell for the Hamamatsu sensors versus overvoltage, at each fluence. At $1 \times 10^{13}\,\mathrm{n_{eq}/cm^2}$, the hotspot count increases with overvoltage instead of decreasing. This could be a peak-finding limitation.

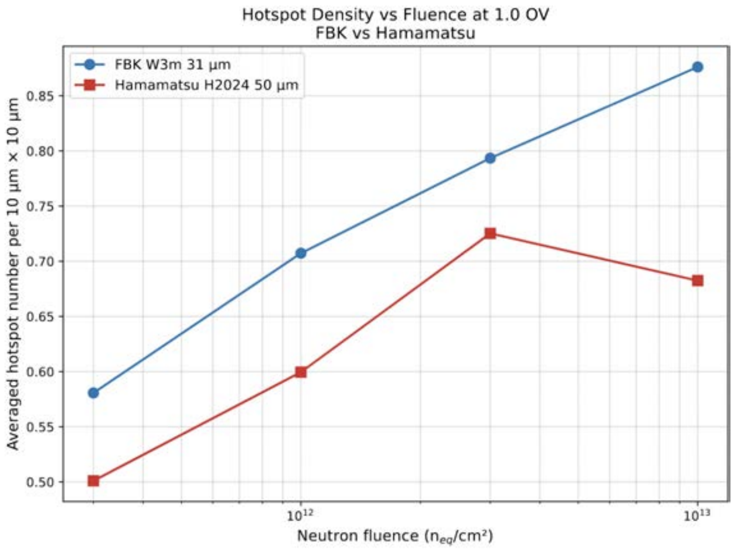

Figure 40: Area-normalised hotspot density of the FBK and Hamamatsu sensors at 1.0 OV across the four irradiated fluences. The finer-pitched FBK density rises monotonically from ∼0.58 to ∼0.88 hotspots per 10 µm × 10 µm, whereas the Hamamatsu density saturates and turns over slightly (from ∼0.73 to ∼0.68) at the highest fluence, where the crowded, bright hotspots in its larger cells begin to merge, an algorithmic resolution limit rather than a real reduction in damage sites.

To compare the two manufacturers on an equal footing, the hotspot count is normalised by cell area to give a hotspot density, shown at 1.0 OV in Figure 40.

Two trends are visible. First, the FBK W3m (31 µm) sensors show a systematically higher hotspot density than the Hamamatsu H2024 (50 µm) sensors at every fluence, rising monotonically from about 0.58 to 0.88 hotspots per 10 µm × 10 µm. Second, while the Hamamatsu density first increases with fluence, it turns over slightly at $1 \times 10^{13}\,\mathrm{n_{eq}/cm^2}$ (from ∼0.73 down to ∼0.68): at the highest damage level the individual hotspots in the larger Hamamatsu cells become so numerous and bright that they merge, and the apparent count falls. This is an algorithmic resolution limit rather than a genuine reduction in damage sites.

4.6 Hotspot Intensity

We fit the hotspot-intensity distribution with an *Exponentially Modified Gaussian* (EMG),

$$f(x; \mu, \sigma, \lambda) = \frac{\lambda}{2} \exp\left[\frac{\lambda}{2}\left(2\mu + \lambda\sigma^2 - 2x\right)\right] \mathrm{erfc}\left(\frac{\mu + \lambda\sigma^2 - x}{\sqrt{2}\sigma}\right) , \tag{10}$$

where the complementary error function is

$$\mathrm{erfc}(x) = 1 - \mathrm{erf}(x) = \frac{2}{\sqrt{\pi}} \int_x^\infty e^{-t^2}\, dt \ . \tag{11}$$

This distribution resembles a Gaussian but carries a longer tail on the positive side, as illustrated in Figure 41, which makes it well matched to the intensity spectra: a roughly Gaussian bulk of typical hotspots plus a high-intensity tail from the brightest damage sites.

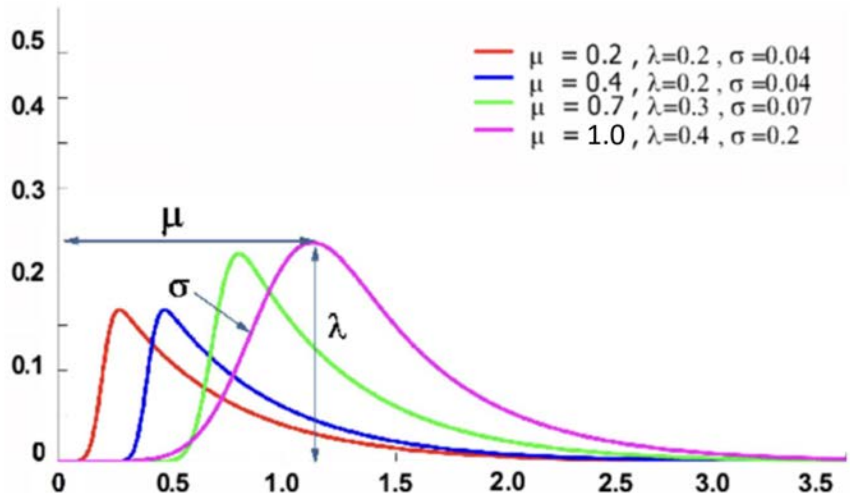

Figure 41: Exponentially Modified Gaussian (EMG) distributions for varying μ, λ and σ, which control the peak position, the exponential decay rate (and hence the length of the positive tail) and the peak width, respectively [23]. The positive tail is what allows the EMG to capture the population of unusually bright hotspots.

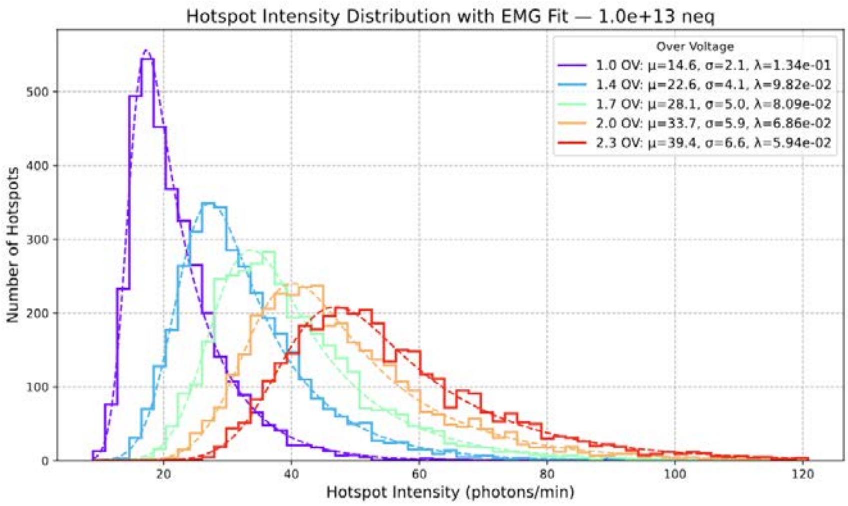

Figure 42: Hotspot-intensity distributions of the FBK $1 \times 10^{13}\,\mathrm{n_{eq}/cm^2}$ sensor at several overvoltages, each fitted with an EMG (solid curves). As overvoltage increases the distribution both shifts to higher intensity and broadens, reflecting the growing gain and secondary-photon yield per avalanche.

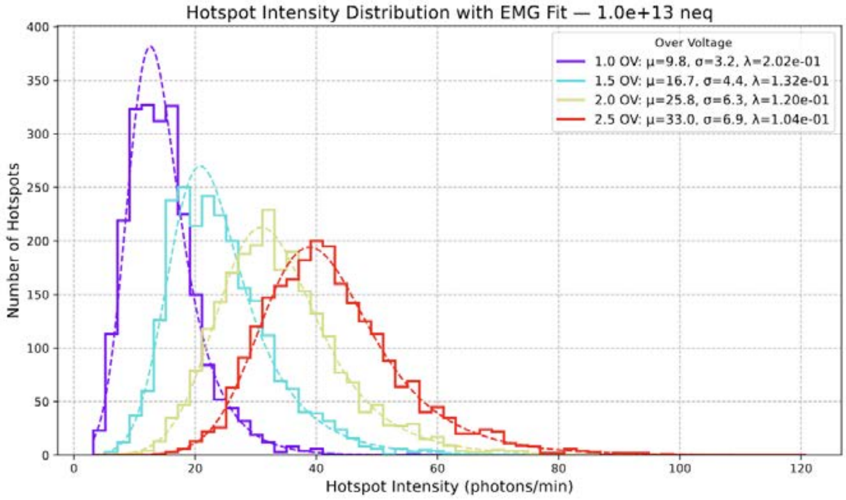

Figure 43: Hotspot-intensity distributions of the Hamamatsu $1 \times 10^{13}\,\mathrm{n_{eq}/cm^2}$ sensor at several overvoltages, with EMG fits. The same shift-and-broaden behaviour with overvoltage is seen as for FBK (Figure 42), but the distributions peak at lower intensity.

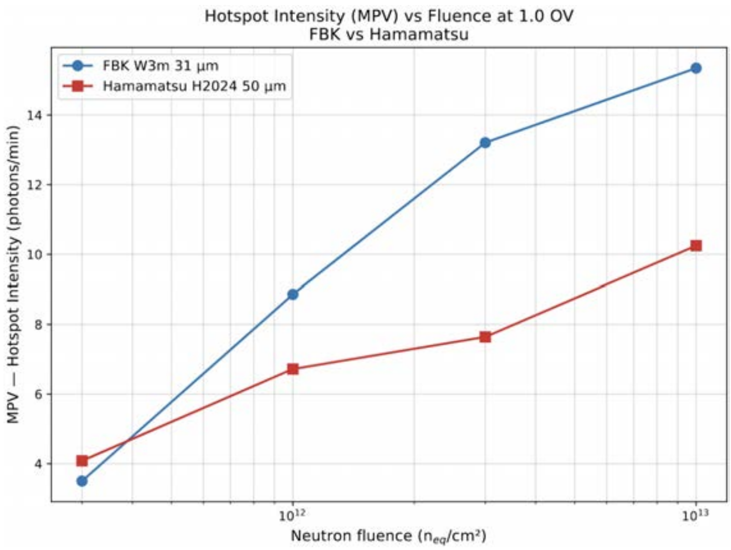

Figure 44: Most probable value (MPV) of the EMG fits to the hotspot intensity for the FBK and Hamamatsu sensors at 1.0 OV across the four irradiated fluences. The two manufacturers agree at the lowest fluence (∼4 photons/min) but diverge as damage accumulates, the FBK intensity rising more steeply to ∼15 photons/min at $1 \times 10^{13}\,\mathrm{n_{eq}/cm^2}$ against ∼10 photons/min for Hamamatsu.

The EMG fits describe the data well across overvoltage and fluence (Figure 42, Figure 43; parameters in Tab. **??**). Such a parameterised distribution can serve as a benchmark for the quality of simulations. The most probable value (MPV) can further be used as a measure of the degree of damage, which would in principle allow the accumulated fluence to be inferred from it.

The MPVs of the fits are shown in Figure 44. At the lowest fluence the two manufacturers produce hotspots of comparable brightness (MPV $\approx$ 4 photons/min), but the FBK intensity rises much more steeply with fluence, reaching $\sim$15 photons/min at $1 \times 10^{13}\,\mathrm{n_{eq}/cm^2}$ against $\sim$10 photons/min for Hamamatsu. This is the same ordering already seen in the per-charge photon yield of Section 4.3, where the FBK slope of 2.82×10^4 photons/µC exceeds the Hamamatsu 1.92×10^4 photons/µC: the stronger field of the smaller FBK cells gives hotter carriers and a higher secondary-photon yield per avalanche, which compounds with the higher hotspot density to make each FBK damage site appear brighter.

4.7 Hotspot Sigma

The hotspot spatial size σ can likewise be fitted with an EMG. In general the hotspot σ increases with overvoltage, while the width of its EMG distribution stays relatively constant, unlike the intensity distribution, whose EMG width varies substantially (Figure 45, Figure 46).

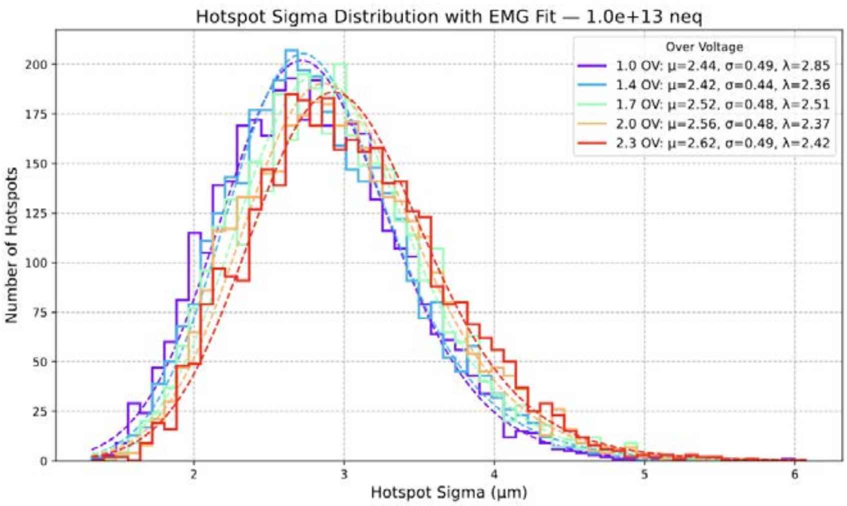

Figure 45: Hotspot-size (σ) distributions of the FBK $1 \times 10^{13}\,\mathrm{n_{eq}/cm^2}$ sensor at several overvoltages, with EMG fits. The distribution shifts only modestly to larger σ with overvoltage while retaining an almost constant shape, in contrast to the strongly broadening intensity distributions.

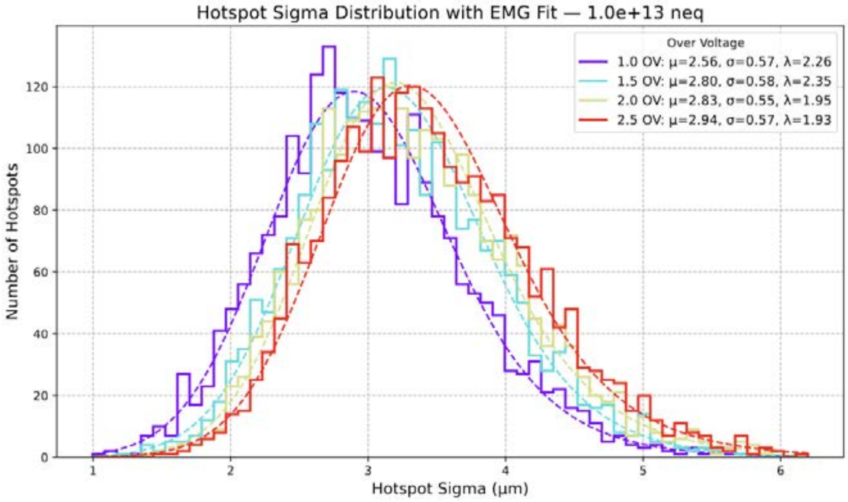

Figure 46: Hotspot-size (σ) distributions of the Hamamatsu $1 \times 10^{13}\,\mathrm{n_{eq}/cm^2}$ sensor at several overvoltages, with EMG fits, showing the same weak overvoltage dependence of the size as the FBK sensor.

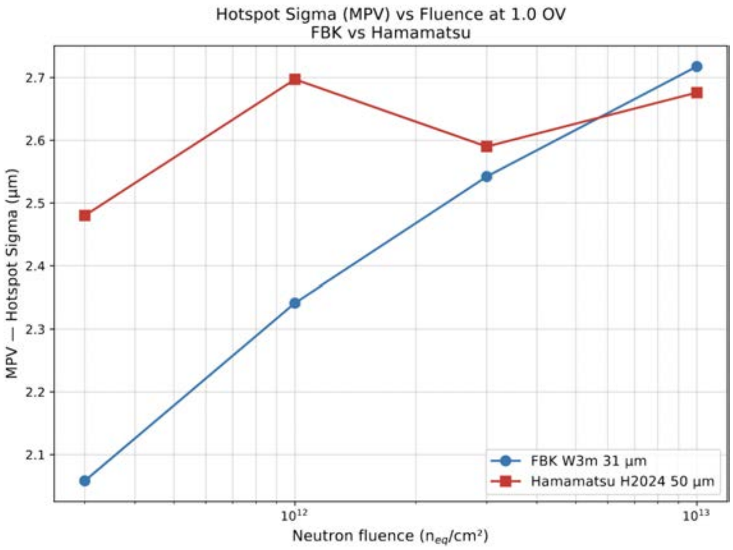

Figure 47: Most probable value (MPV) of the EMG fits to the hotspot spatial size σ for the FBK and Hamamatsu sensors at 1.0 OV across the four irradiated fluences. The ordering is opposite to that of the intensity (Figure 44): the larger Hamamatsu cells give bigger hotspots at low fluence, but the FBK size grows monotonically while the Hamamatsu one stays roughly constant, so the two converge to $\sim$2.7 µm at the highest fluence.

The MPVs of the size fits are shown in Figure 47. The size ordering is the opposite of the intensity ordering. At low fluence the Hamamatsu hotspots are spatially larger (MPV

$\sigma \approx 2.5\,\mu\text{m}$ versus $\approx 2.1\,\mu\text{m}$ for FBK), as expected from the larger 50 µm Hamamatsu cell pitch, which gives each emission site more room to spread. The FBK size then grows steadily with fluence while the Hamamatsu value stays roughly constant ($\sim$2.5–2.7 µm), so that by $1 \times 10^{13}\,\text{n}_{\text{eq}}/\text{cm}^2$ the two are comparable ($\sim$2.7 µm). In other words, increasing damage makes FBK hotspots both brighter and broader, whereas for Hamamatsu it is mainly the brightness that grows.

4.8 Hotspot Current Contribution

Combining the amplitude and the size gives the integrated intensity of each hotspot, and hence a distribution of its current contribution to the whole channel. This is obtained by taking the fraction of the total emission intensity contributed by a given hotspot and multiplying by the mean current. These distributions are again fitted with an EMG (Figure 48, Figure 49), with the fit parameters and MPVs listed in Tab. **??**.

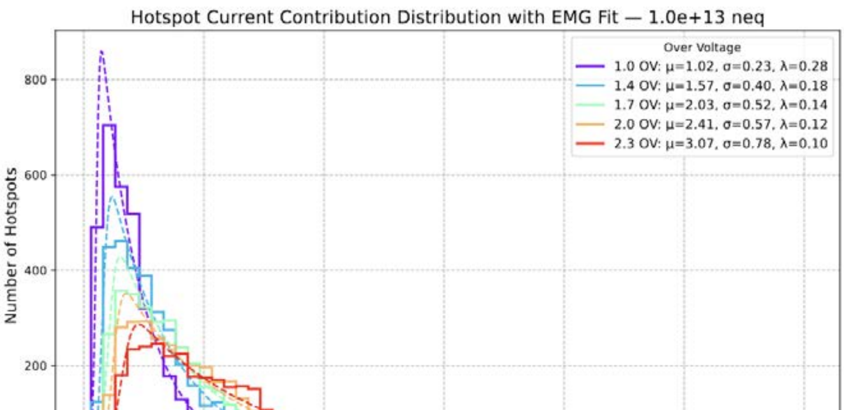

Figure 48: Per-hotspot current-contribution distributions of the FBK $1 \times 10^{13}\,\text{n}_{\text{eq}}/\text{cm}^2$ sensor at several overvoltages, with EMG fits. As overvoltage rises the distribution moves to larger per-hotspot currents, since each site both brightens and (weakly) broadens.

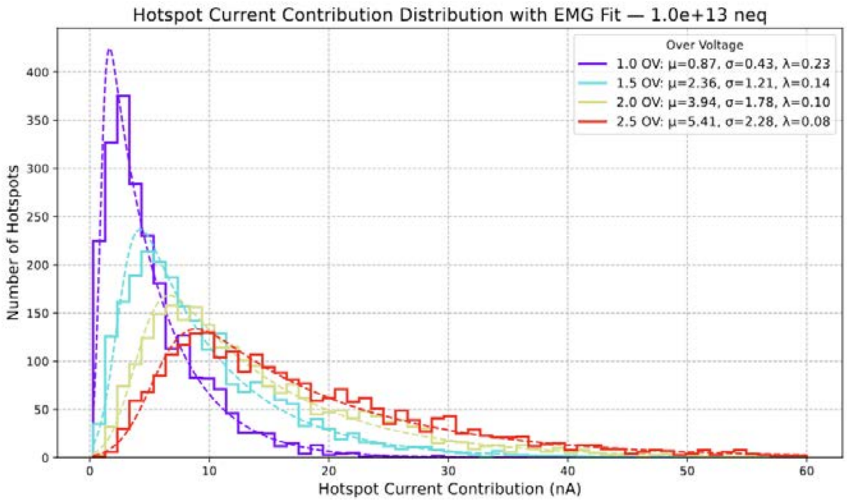

Figure 49: Per-hotspot current-contribution distributions of the Hamamatsu $1 \times 10^{13}\,\mathrm{n_{eq}/cm^2}$ sensor at several overvoltages, with EMG fits, following the same overvoltage trend as the FBK sensor.

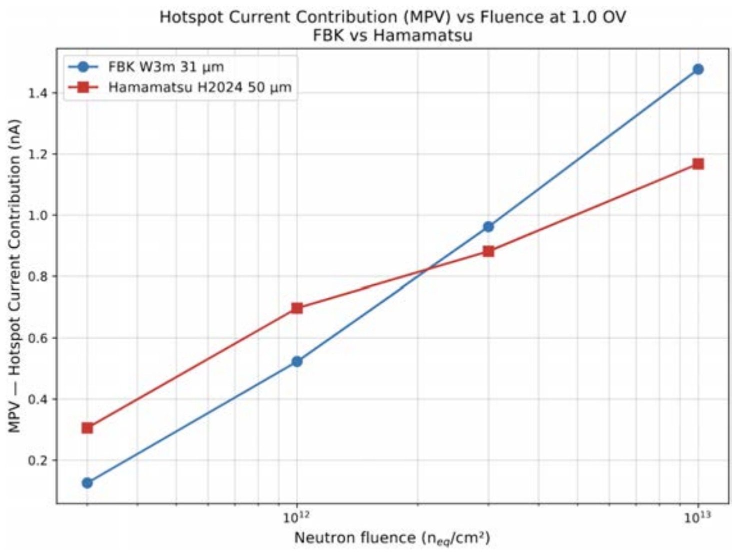

Figure 50: Most probable value (MPV) of the EMG fits to the per-hotspot current contribution for the FBK and Hamamatsu sensors at 1.0 OV across the four irradiated fluences. The two track each other closely and cross at intermediate fluence: Hamamatsu sites carry marginally more current where their larger size dominates (low fluence), while the steeper FBK growth makes FBK the larger contributor above $\sim 3 \times 10^{12}\,\mathrm{n_{eq}/cm^2}$, where their brightness dominates.

The MPVs of the current-contribution fits are shown in Figure 50. The per-hotspot current contribution combines brightness and size and is therefore the most direct single-site measure of dark activity. The two manufacturers track each other closely, both rising from $\sim$0.1–0.3 nA at $3 \times 10^{11}\,\mathrm{n_{eq}/cm^2}$ to $\sim$1.2–1.5 nA at $1 \times 10^{13}\,\mathrm{n_{eq}/cm^2}$. The curves cross at intermediate fluence: Hamamatsu sites carry slightly more current at low fluence (where their larger size dominates), while the steeper FBK growth makes FBK the larger current contributor at the highest fluence (where its brightness dominates). This crossing mirrors the opposing intensity and size trends discussed above.

4.9 Hotspot Location Distribution

For the non-irradiated SiPMs the hotspots gather at the cell borders (Figure 51). We attribute this to two effects: manufacturing defects are more likely to occur at the cell borders, seeding electron-hole generation centres there, and at sufficiently high overvoltage the borders develop a field strong enough for avalanches to be triggered.

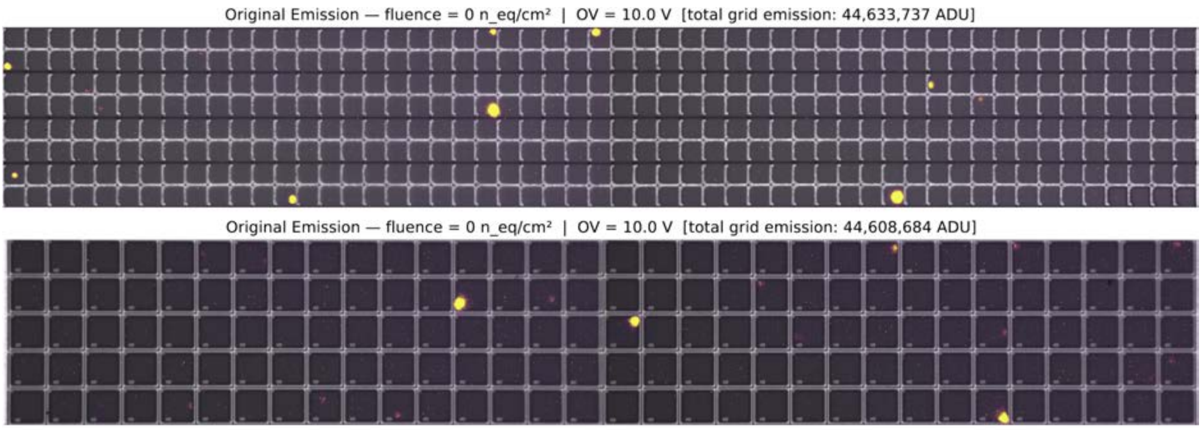

Figure 51: EMMI images of the non-irradiated sensors at 10.0 OV: FBK (**top**) and Hamamatsu (**bottom**). In the absence of bulk radiation damage the sparse emission concentrates along the cell borders, consistent with manufacturing defects and edge fields.

For the irradiated SiPMs, we analyse the hotspot locations using the intensity-weighted centre of all of the detected peaks,

$$\bar{x} = \frac{\sum_i x_i I_i}{\sum_i I_i}, \quad \bar{y} = \frac{\sum_i y_i I_i}{\sum_i I_i}.$$

The weighted centres cluster neither at the centre nor at the border of the cell, but in a specific square region surrounding the centre, as the one-dimensional x- and y-projections that Figures 52 (FBK) and 53 (Hamamatsu) show.

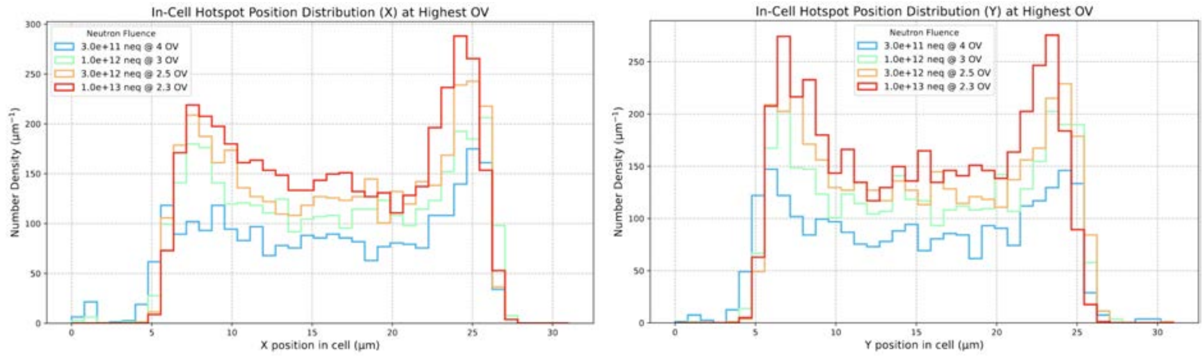

Figure 52: Hotspot weighted centre distributions of the FBK sensors at the highest overvoltage for all fluences, projected onto the x (**left**) and y (**right**) axes. Both projections are double-peaked, the hotspots avoiding both the cell centre and the extreme edges, which is the one-dimensional signature of the square-ring clustering made explicit in Figure 54.

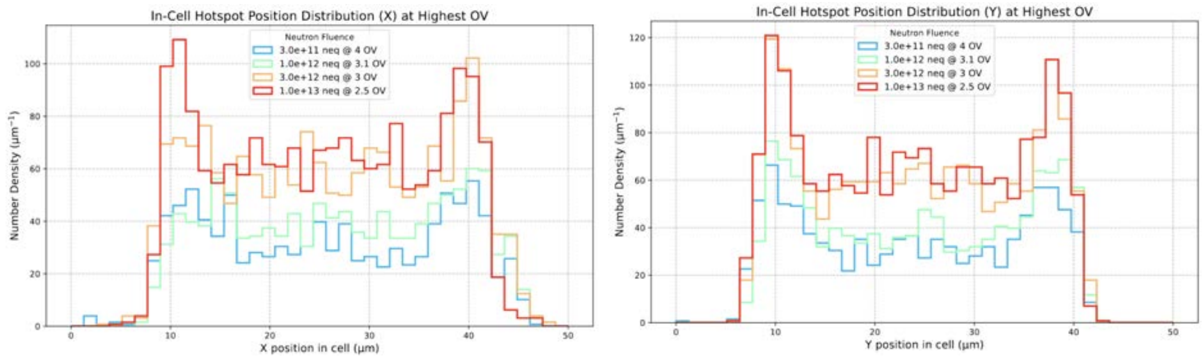

Figure 53: In-cell hotspot-position distributions of the Hamamatsu sensors at the highest overvoltage for all fluences, projected onto the x (**left**) and y (**right**) axes. The same double-peaked, centre-avoiding structure as FBK (Figure 52) is present.

This is particularly clear when the location distribution is plotted as a 2D histogram overlaid on the reference image, as shown below in Figure 54.

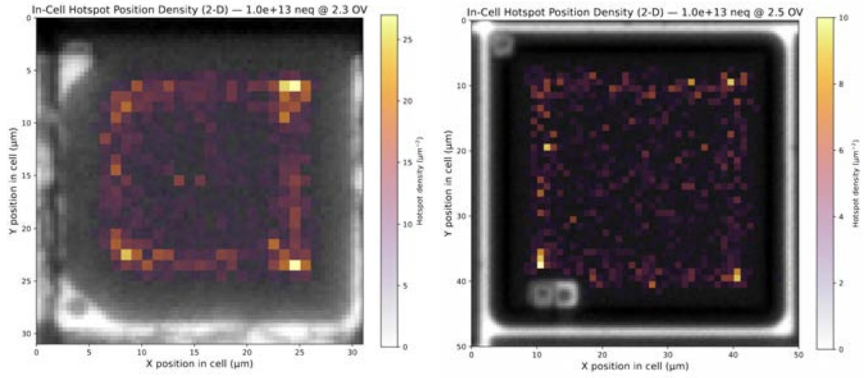

Figure 54: Two-dimensional in-cell hotspot-position histograms at $1 \times 10^{13}\,\mathrm{n_{eq}/cm^2}$, overlaid on the reference cell image: FBK at 2.3 OV (**left**) and Hamamatsu at 2.5 OV (**right**). Both reveal a bright square ring of emission set in from the cell edges.

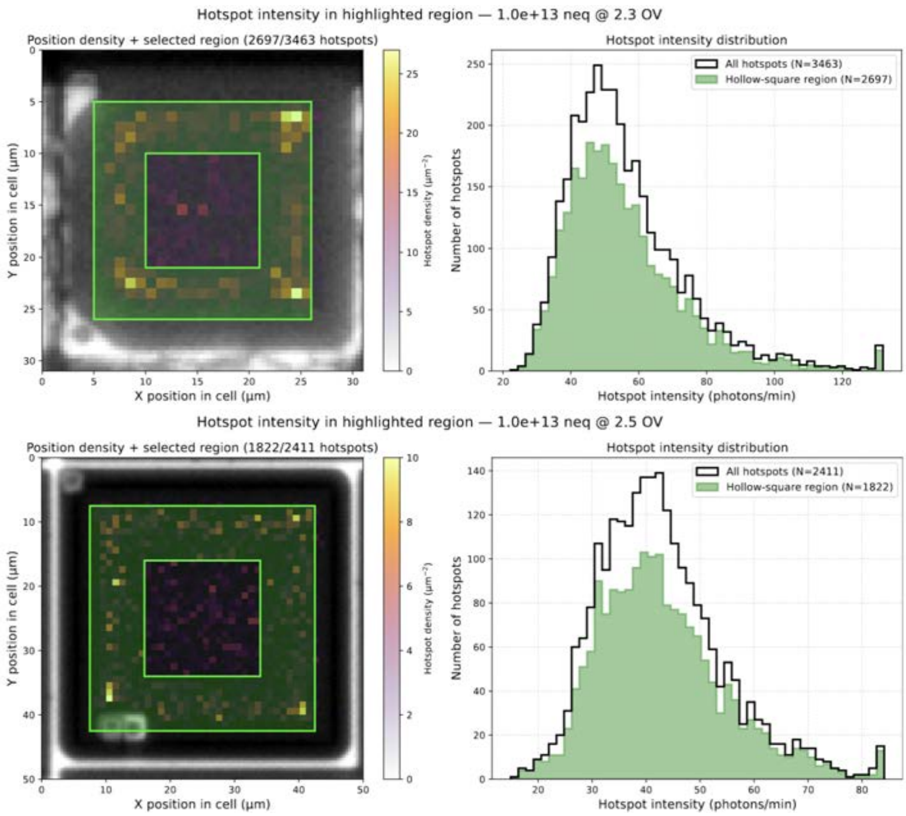

Figure 55: Amplitude uniformity around the square ring. **Top**: a hollow-square region selected on the FBK 2D position distribution (left) and the amplitude distribution of the hotspots falling inside it (right). **Bottom**: the same for Hamamatsu.

If we examine the amplitude profile of the hotspots within this ring, we find no strong intensity bias around it (Figure 55). Although the hotspots preferentially sit on the ring, their brightness within it shows no marked positional bias, indicating that the ring reflects where avalanches are triggered rather than where the strongest emission occurs.

The most likely cause of this clustering is the extent of the depletion region. The depletion region lies above the multiplication region and carries a field strong enough to drift electron-hole pairs but too weak to trigger an avalanche (Figure 10, Figure 56). Under neutron damage, pairs generated in this region drift towards the high-field multiplication volume and trigger avalanches at its edge, which is exactly the square ring observed in the 2D histogram.

Another explanation is that this positional bias stems from the reconstruction algorithm which tends to group hotspots near the cell borders and assign them a weighted centre in the square ring region. Disentangling these two contributions would require further study.

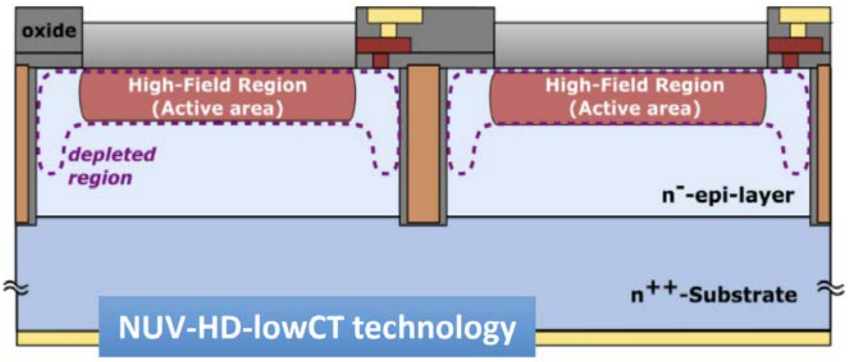

Figure 56: Schematic cross-section of the internal structure of a SiPM microcell [13], complementing the field profile of Figure 10. The depletion region expands beyond the narrow high-field multiplication region towards the cell border; carriers generated by radiation damage in the former drift down into the latter, triggering avalanches at its perimeter and producing the square-ring emission pattern of Figure 54.

4.10 Fill Factor Estimation

The fill factor of a SiPM can be estimated from the fraction of the cell area that emits light, since these are the regions with a field strong enough to sustain an avalanche. For each manufacturer the averaged EMMI image of a cell at the highest fluence ($1 \times 10^{13}\,\mathrm{n_{eq}/cm^2}$) and highest overvoltage is used, and the active area is defined as the contiguous region emitting more than 30% of the peak intensity (Figure 57).

This method gives a fill factor of 61.4% for the FBK 31 µm sensor and 58.9% for the Hamamatsu 50 µm sensor. These comparable values, with the more densely packed FBK device slightly ahead, are consistent with the smaller dead-area fraction expected for a finer cell pitch. They should be read as lower bounds: as discussed for the in-cell hotspot distribution, carriers generated in the depletion region above the multiplication volume can drift in and trigger an avalanche, so the effective photon-detection area is potentially larger than the light-emitting area measured here.

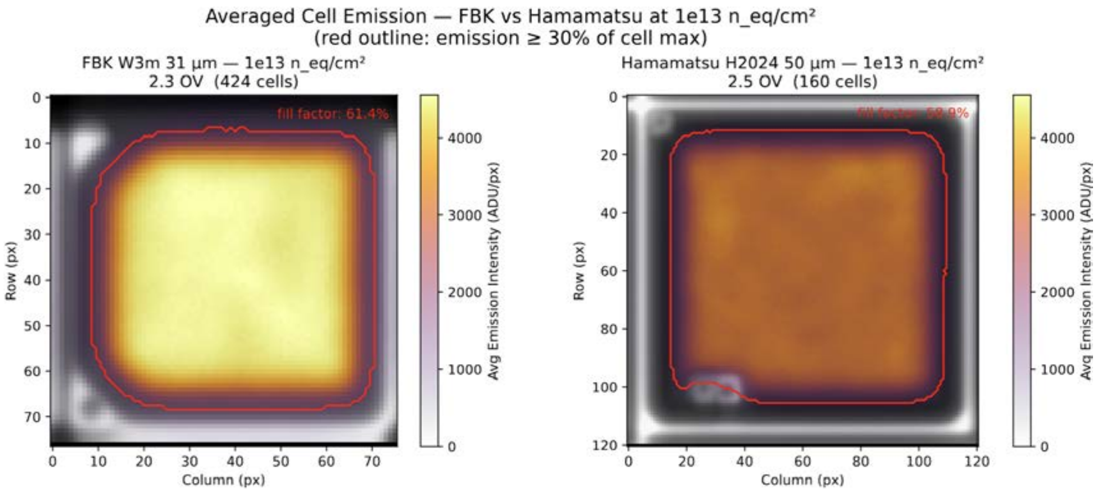

Figure 57: Fill-factor estimation from the averaged cell emission of the FBK (**left**) and Hamamatsu (**right**) sensors at $1 \times 10^{13}\,\mathrm{n_{eq}/cm^2}$ and the highest overvoltage. The red contour marks the region emitting at least 30% of the cell's peak intensity, taken as the light-emitting active area; the enclosed fractions give fill factors of 61.4% (FBK) and 58.9% (Hamamatsu).

5 Conclusion

In this project, emission microscopy has been used to quantitatively characterise the radiation damage of SiPMs from FBK (W3m, 31 µm, $V_{\mathrm{bd}} \approx 31\,\mathrm{V}$) and Hamamatsu (H2024, 50 µm, $V_{\mathrm{bd}} \approx 51\,\mathrm{V}$) across one non-irradiated reference and four neutron fluences, up to $1 \times 10^{13}\,\mathrm{n_{eq}/cm^2}$, a level that exceeds the $\sim 3 \times 10^{12}\,\mathrm{n_{eq}/cm^2}$ expected at the LHCb SciFi SiPMs over the full HL-LHC programme by a factor of three. By combining the 10-minute emission images with the simultaneously logged leakage current, and by developing an image-processing pipeline that reconstructs individual emission hotspots as 2D Gaussians, we have extracted and analysed the per-cell emission intensity, the number of photons produced per charge carrier, and the hotspot number, intensity, size, current contribution, spatial distribution and fill factor.

Several physically consistent trends emerge. The dark I–V curves rise sub-linearly with fluence, the signature of recovery-time saturation which leads to a reduced effective overvoltage in the most heavily irradiated cells. The total emission intensity scales linearly with the mean current for both manufacturers (slopes 2.82×10^4 and 1.92×10^4 photons/µC for FBK and Hamamatsu), and after correcting for the camera quantum efficiency, the emission spectrum and the Snell-Fresnel photon-acceptance factor, this yields overall detection efficiencies of 0.076% (FBK) and 0.062% (Hamamatsu) and photon-production rates of 5.98×10^{-6} and 4.92×10^{-6} photons per charge carrier. For FBK these are in good agreement with the literature value of $\sim 5 \times 10^{-6}$, while the Hamamatsu result falls somewhat below the quoted $\sim 9 \times 10^{-6}$; this shortfall could stem from the limitations of the setup, and potentially from genuine differences between the specific SiPMs measured here and those in the literature. The more densely packed FBK cells, with their stronger peak field, show higher hotspot densities (rising monotonically from 0.58 to 0.88 per 10 µm × 10 µm) and brighter hotspots whose intensity MPV grows more steeply with fluence (to ∼15 versus ∼10 photons/min at $1 \times 10^{13}\,\mathrm{n_{eq}/cm^2}$), whereas the larger Hamamatsu cells produce spatially broader hotspots at low fluence; the two size distributions converge to ∼2.7 µm and the per-hotspot current contributions cross over, both rising from ∼0.1–0.3 nA to ∼1.2–1.5 nA. In the irradiated sensors the hotspots cluster in a square ring within each cell rather than at the centre or the borders, consistent with carriers generated in the depletion region drifting into the high-field multiplication volume; from the light-emitting area we estimate fill factors of 61.4% (FBK) and 58.9% (Hamamatsu).

The main limitation is the hotspot-extraction algorithm, which still depends on hand-tuned parameters (most notably the peak-separation distance) and can therefore introduce bias, particularly at high overvoltage where neighbouring hotspots merge and the apparent hotspot count falls. A systematic figure of merit for the reconstruction quality (for example the pixel-wise squared-intensity residual), together with simultaneous multi-Gaussian fitting under greater computing power, would allow these parameters to be optimised and the residual ∼5% amplitude bias to be removed.

Nonetheless, the method provides a reproducible, spatially resolved way to quantify SiPM radiation damage cell by cell, a capability that bulk I–V or DCR measurements cannot offer. The extracted Exponentially Modified Gaussian (EMG) distributions of hotspot intensity, size and current contribution can serve as benchmarks for damage simulations, and their most-probable values, in particular the hotspot intensity MPV, are promising candidate metrics for inferring the accumulated fluence non-destructively. Establishing this MPV-fluence calibration is the natural next step, and would turn EMMI into a quantitative diagnostic: a tool for screening and qualifying radiation-hard SiPM candidates, for cross-checking simulated damage maps, and for estimating the dose absorbed by sensors already in the field. As the LHCb Upgrade II and, further ahead, future colliders push SiPM photodetectors into ever more radiative environments well beyond $1 \times 10^{13}\,\mathrm{n_{eq}/cm^2}$, such a cell-level optical diagnostic offers a valuable complement to the bulk electrical figures of merit currently used to certify their radiation tolerance.

References

[1] Wikipedia contributors. *Standard Model — Wikipedia, The Free Encyclopedia*. [Online; accessed 25-June-2026]. 2026. URL: https://en.wikipedia.org/wiki/Standard_Model.

[2] Atlearner. *What is Big Bang Theory? Explained*. [Online; accessed 25-June-2026]. 2020. URL: https://www.atlearner.com/2020/02/what-is-big-bang-theory.html.

[3] European Southern Observatory (ESO). *Aerial view of the CERN site near Geneva*. [Online; accessed 25-June-2026]. 2026. URL: https://supernova.eso.org/germany/exhibition/images/cern-aerial-cc/.

[4] Welt der Physik. *Der Large Hadron Collider am CERN*. [Online; in German; accessed 25-June-2026]. 2026. URL: https://www.weltderphysik.de/gebiet/teilchen/experimente/teilchenbeschleuniger/cern-lhc/.

[5] Y. Nie, R. Schmidt, and V. Chetvertkova. "Numerical simulations of energy deposition caused by 50 MeV—50 TeV proton beams in copper and graphite targets". In: *Physical Review Accelerators and Beams* 20.8 (2017). DOI: 10.1103/PhysRevAccelBeams.20.081001. URL: https://doi.org/10.1103/PhysRevAccelBeams.20.081001.

[6] LHCb Collaboration. *The 2025 Breakthrough Prize in Fundamental Physics awarded to the LHCb collaboration*. [Online; accessed 25-June-2026]. 2025. URL: https://lhcb-outreach.web.cern.ch/2025/04/07/the-2025-breakthrough-prize-in-fundamental-physics-awarded-to-the-lhcb-collaboration/.

[7] LHCb Collaboration. *LHCb Tracker Upgrade Technical Design Report*. Tech. rep. 2014. DOI: 10.17181/CERN.OOOE.909T. URL: https://cds.cern.ch/record/1647400.

[8] LHCb Collaboration. *End of 2016 data taking period*. [Online; accessed 25-June-2026]. 2016. URL: https://lhcb-outreach.web.cern.ch/2016/12/05/end-of-2016-data-taking-period/.

[9] Thomas Kirn and LHCb Collaboration. "SciFi – A large scintillating fibre tracker for LHCb". In: *Nuclear Instruments and Methods in Physics Research Section A* 845 (2017), pp. 481–485. DOI: 10.1016/j.nima.2016.06.057. URL: https://www.sciencedirect.com/science/article/pii/S016890021630599X.

[10] S. Vinogradov. "The Silicon Photomultiplier (SiPM) Concept and Design Development". In: *APPEC News* (Mar. 2015). Available at: https://www.appec.org/news/the-silicon-photomultiplier-sipm-concept-and-design-development/.

[11] *Single-Photon Avalanche Diode (SPADs)*. MeetOptics Academy. Available at: https://www.meetoptics.com/academy/single-photon-avalanche-diode.

[12] W. Shockley and W. T. Read. "Statistics of the Recombinations of Holes and Electrons". In: *Physical Review* 87.5 (1952), pp. 835–842. DOI: 10.1103/PhysRev.87.835. URL: https://doi.org/10.1103/PhysRev.87.835.

[13] F. Acerbi. *SiPM rad-hard R&D: AIDAinnova production at FBK*. Presentation at Fondazione Bruno Kessler (FBK). Presented on 2025-05-14. May 2025.

[14] S.-E. Ungersböck. "Advanced Modeling of Strained CMOS Technology". [Online; accessed 27-June-2026]. PhD thesis. Institute for Microelectronics, TU Wien, 2007. URL: https://www.iue.tuwien.ac.at/phd/ungersboeck/node28.html.

[15] J. Kölzer et al. "Quantitative emission microscopy". In: *Journal of Applied Physics* 71.11 (1992), R23–R41. DOI: 10.1063/1.350466. URL: https://doi.org/10.1063/1.350466.

[16] V. Chmill et al. "Study of the breakdown voltage of SiPMs". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 845 (2017), pp. 56–59. DOI: 10.1016/j.nima.2016.04.047. URL: https://doi.org/10.1016/j.nima.2016.04.047.

[17] E. Engelmann, E. Popova, and S. Vinogradov. "Spatially resolved dark count rate of SiPMs". In: *The European Physical Journal C* 78.11 (2018), p. 971. DOI: 10.1140/epjc/s10052-018-6454-0. URL: https://doi.org/10.1140/epjc/s10052-018-6454-0.

[18] J. B. McLaughlin et al. "Characterisation of SiPM Photon Emission in the Dark". In: *Sensors* 21.17 (2021), p. 5947. DOI: 10.3390/s21175947. URL: https://doi.org/10.3390/s21175947.

[19] Hamamatsu Photonics K.K. *ORCA-Quest 2 qCMOS® camera C15550-22UP Product Brochure*. Tech. rep. [Online; accessed 25-June-2026]. 2025. URL: https://www.hamamatsu.com.

[20] G. Gallina et al. "Characterization of SiPM Avalanche Triggering Probabilities". In: *IEEE Transactions on Electron Devices* 66.10 (2019), pp. 4228–4234. DOI: 10.1109/TED.2019.2935690. URL: https://doi.org/10.1109/TED.2019.2935690.

[21] Martin A. Green. "Self-consistent optical parameters of intrinsic silicon at 300K including temperature coefficients". In: *Solar Energy Materials and Solar Cells* 92 (2008), pp. 1305–1310. DOI: 10.1016/j.solmat.2008.06.009.

[22] L. V. Rodríguez-de Marcos et al. "Self-consistent optical constants of SiO_2 and Ta_2O_5 films". In: *Optical Materials Express* 6.11 (2016), pp. 3622–3637. DOI: 10.1364/OME.6.003622. URL: https://doi.org/10.1364/OME.6.003622.

[23] A. Tripathi and M. Shukla. "Physically inspired synthesis of Santoor and Sarod". In: *2015 International Conference on Information and Communication Technology for Intelligent Systems (ICIINFS)*. Dec. 2015. DOI: 10.1109/ICIINFS.2015.7399004. URL: https://doi.org/10.1109/ICIINFS.2015.7399004.