PWO 2025 Simulation

Data-constrained beam-monitor alignment from Geant4 to native, unclipped circuit ADC counts

A source-audited, equation-by-equation account of the nine-energy simulation: raw-data-constrained effective beam-monitor placement, latent beam-profile forward folding, beam generation, analytic detector geometry, Geant4 energy and optical transport, custom wavelength-dependent photon detection, finite-cell SiPM response, the GTM-EIC/CITIROC-equivalent input network and slow shaper, ideal positive-peak HOLD, and the native floating ideal ADC-input-equivalent coordinate used by downstream work.

Nominal momentum settings	98, 197, 297, 395, 492, 584, 706, 739, 796 MeV/c
Aligned campaign sample	Exactly 10,000 generated events per setting
Endpoint	`direct_5V_channel_adc_unclipped`
Document status	final—audited all-nine equal-10k campaign
Build date	22 August 2026

This endpoint is a floating ideal ADC-input-equivalent count. It precedes fitted pedestal, noise, the data-to-simulation residual scale, threshold, clipping, hardware-word encoding, event selection, Stage-4 weighting, and reconstructed energy.

Abstract

This report documents the forward PWO simulation used for the 2025 nine-energy beam-test study, including a data-constrained effective BM placement and a latent-Gaussian strip forward fold. Its target naming convention is crystal 20 / GID 34+40 / $R(2,2)$ / software C14 / GDML `PWO_15`. The BM result is an effective analysis-coordinate placement, not a mechanical survey. The report distinguishes the nominal momentum setting from the generated positron kinetic energy and the audited GDML specification from the compiled analytic geometry that the production executable actually used. Optical photons are transported by Geant4, tested against a custom 5-V SiPM PDE, and converted to time-binned effective fired cells with correlated noise, dark counts, finite occupancy, and recovery. Forty-four independently read SiPM channels are then propagated through an exact linear nodal input-network kernel, a CR–RC^2 slow shaper, a declared 200-ns positive-peak HOLD assumption, and a fixed charge-to-count conversion. The final stored quantity is the native, unclipped floating ideal ADC-input-equivalent response `direct_5V_channel_adc_unclipped`. The package supplies the equations, empirical nine-energy plots, source hashes, limitations, and a rebuildable report source. Detector-level BM closure passed at all nine energies. The aligned all-nine equal-10k optical campaign completed 9/9; all ROOT sidecars and the seven-tree, 10,000-unique-event schema audit passed. The Condor cluster-998 digitizer audit passed for all nine 10,000-by-44 native arrays, and the regenerated figures, tables, numerical closure, equation coverage, and source/asset manifests passed their release audits. Every numerical result in this report is therefore derived from the single aligned equal-10k campaign.

Contents

Executive summary

> **Release status:** the data-constrained BM method and target crosswalk are frozen, and detector-level BM closure passed at all nine energies. The aligned all-nine equal-10k optical campaign completed 9/9 and passed its ROOT checksum/schema audit. Native digitization passed its cluster-998 all-nine audit, and every dependent figure, table, equation index, and provenance manifest was regenerated and audited from that same campaign.

This report freezes the forward Monte Carlo chain that starts with one generated positron and ends with the native, floating, unclipped circuit prediction

$$S_{ec} \equiv \texttt{direct\_5V\_channel\_adc\_unclipped}(e, c), \quad (1)$$

for generated event e and independently read channel c. The equality in Eq. (1) is a definition: S_{ec} is the precise array field used as the report endpoint, not an inferred energy and not the integer ADC word recorded by the beam-test hardware. "Unclipped" means that no 11,000-count ceiling is applied to this field. The endpoint has no fitted pedestal, electronic noise, channel threshold, fitted data-to-simulation residual calibration, integer quantization, signed-payload encoding, or event selection.

The aligned campaign represents nine nominal beam-momentum settings—98, 197, 297, 395, 492, 584, 706, 739, and 796 MeV/c—with 10,000 optical events at each setting. Each label denotes a momentum setting, not a monoenergetic kinetic energy. The generator samples a positive momentum and a correlated transverse beam profile, Geant4 transports the positron and its shower through two beam monitors and a 6×6 $PbWO_4$ array, and optical photons propagate through the crystal, grease, window, and active-silicon model. A wavelength-dependent Bernoulli test supplies the custom SiPM photon-detection efficiency (PDE). Prompt crosstalk, afterpulses, dark pulses, saturation, and exponential recovery then produce sparse 0.25-ns effective-fired-cell histories. The later sections of this report propagate those histories through the audited input network, CITIROC-equivalent shaper, 200-ns HOLD, and native ADC conversion.

The principal forward-model contracts are:

- Production used `FTFP_BERT` plus `G4OpticalPhysics` in Geant4 11.4.2. It used the compiled `analytic_from_audited_gdml` backend; the supplied GDML was the frozen geometry authority but was not parsed at runtime.
- The simulated detector has 36 crystals and 72 SiPM active layers. Downstream native ADC arrays retain 44 independently read channels in 22 crystals. Three differently scaled combined-readout pixels are not folded into that channel coordinate.
- The jobs used the audited seven-point effective-isotropic PWO refractive-index curve and the scaled 5-V PDE curve. A later local 91-point Sellmeier-derived curve did not enter these events.
- The stated SiPM gain 6.0×10^6 is metadata inside the Geant4 executable. Geant4 writes photons and effective fired cells; only the downstream circuit stage turns a fired-cell history into charge and then ADC counts.
- The ROOT branch `energy_MeV` is the sampled positron kinetic energy, whereas `beam_momentum_MeV_c` is its sampled momentum. Neither should be silently replaced by the nominal label.
- All primary scientific plots use exactly 10,000 events per energy from the aligned campaign. Every energy has one direct optical ROOT with unique event IDs 0–9999; no legacy or mixed-population sample enters the result set.

The alignment adds a separate contract. The report target naming convention is crystal 20 / GID 34+40 / $R(2, 2)$ / software C14 / GDML `PWO_15`. These aliases come from different numbering systems but identify the same target. Raw BM data constrain an effective transverse strip placement, and a latent correlated Gaussian is forward-folded through the finite, gapped, quantized

32-strip grid before its parameters are supplied to the gun. The non-optical all-nine closure gate, aligned optical ROOT campaign, cluster-998 native digitization, and report-asset rebuild have all passed. The release manifests bind every plotted endpoint to the same direct 10,000-event population at each energy.

This is a reproducibility report for the forward model, not a claim that every material or electronics constant is exact. The simplified geometry omits the full mechanical assembly, the optical curves contain stated approximations, the incident direction has no angular divergence, and several Geant4 process settings inherit version-specific defaults. Those limitations are preserved as part of the model definition rather than concealed by downstream fitting.

1 Scope and frozen forward-model contract

This report preserves the frozen electronics endpoint while applying the data-constrained BM placement and incident transverse-profile inference. The aligned all-nine equal-10k optical campaign is complete and its ROOT checksum/schema audit passes. Its nine native-ADC products and all dependent report assets have also been regenerated and audited. The equations and numerical products now describe the same frozen equal-10k population end to end.

The purpose of this report is to make every transformation in the following forward chain inspectable:

$$\underbrace{(p,x,y,e^{+})}_{\text{generator}} \longrightarrow \underbrace{(E_{\text{dep}},N_\gamma,t_\gamma)}_{\text{Geant4}} \longrightarrow \underbrace{N_{\text{PDE}}(t)}_{\text{custom PDE}} \longrightarrow \underbrace{N^{\text{aug}}_{\text{fired}}(t)}_{\text{SiPM response}} \longrightarrow \underbrace{A^{\text{hold}} = \max_{t\le t_{\text{HOLD}}} V_{\text{shape}}(t)}_{\text{circuit}} \longrightarrow \underbrace{S_{ec}}_{\text{native ADC}} . \tag{2}$$

The arrows in Eq. (2) mean "is used to compute", not equality. Here p, x, y are generated beam momentum and transverse position, E_{dep} is deposited energy, N_γ and t_γ are optical-photon counts and times, $N_{\text{PDE}}(t)$ is the detected-photon history, $N^{\text{aug}}_{\text{fired}}(t)$ is the history after SiPM nuisance processes and saturation, and A^{hold} is the ideal positive shaped maximum retained up to the HOLD closing time rather than the instantaneous voltage at that time. The physical polarity convention is positive SiPM input, negative low-gain preamplifier output, and positive slow-shaper peak. No absolute-value rectifier or minimum-peak selector is part of this circuit endpoint. Writing the intermediate objects prevents a common category error: deposited energy is not converted directly to ADC by a single fitted scale.

1.1 Included evidence

The simulation half uses the aligned production executable and source, the frozen GDML and effective BM transform, the latent beam-profile table, the seven transferred optical CSV inputs, the campaign runner and Condor submit description, and nine checksummed optical ROOT files. Condor cluster 995 ran one 168-thread optical job at each energy. Every job returned zero and wrote exactly 10,000 unique event IDs 0–9999 with the required seven-tree schema. The electronics half uses the audited production digitizer and readout map; its aligned native NPZ products pass the Condor cluster-998 audit with shape $10{,}000 \times 44$, finite floating values, exact fixed-noise replay, and the required unclipped endpoint. Primary figures use those same direct 10,000-event populations so every nominal energy contributes an equal event count.

1.2 Explicit exclusions

The endpoint S_{ec} is before all beam-data calibration and selection. This report therefore does *not* apply a beam-monitor zero-angle requirement, the two-SiPM crystal selection, crystal multiplicity cleaning, Stage 4 smooth weights, a fitted pedestal, pedestal noise, a fitted threshold, a per-channel data-to-simulation residual scale, an 11,000-count ceiling, integer rounding, a signed 14-bit payload transform, DSCB spectrum fitting, or energy-resolution extraction. Those operations answer different questions and must not be retroactively described as part of the native Monte Carlo endpoint.

The compatibility field `direct_5V_channel_adc`, which imposes an early 11,000-count cap, is retained only for a failure-mode audit. It is never the primary output. Likewise, combined-readout channels have a different scale; they may be documented as mapping metadata but are excluded from the 44-column independent-channel ADC array and its diagnostic sum.

1.3 Authority order

When a descriptive filename, a later local file, and the production evidence disagree, this report follows the following authority order:

1. the checksummed executable/source and the literal production command;
2. the checksummed ROOT or NPZ field actually consumed downstream;
3. the frozen geometry and optical-input hashes transferred to the jobs;
4. campaign notes and ancillary reports.

This order matters in two known cases. First, the log proves that runtime geometry was `analytic_from_audited_gdml`, despite the executable's GDML-oriented name. Second, the optical jobs used the seven-point PWO index file whose SHA-256 begins `44eb6e`; the newer 91-point desktop curve is not a permissible substitute without an optical rerun.

1.4 What "ADC count" means here

Native ADC is a floating ideal ADC-input-equivalent coordinate defined by the digitizer's held voltage and the external ADC voltage scale. It is numerically count-like, but it is not a hardware word and need not be an integer. The model assumes unity, settled transfer from the held internal shaper excursion through peak memory, follower, multiplexer, and `OUT_LG` to that external input. Later in the report the circuit definition has the form

$$S_{ec} = \frac{V_{ec}^{\text{peak}}}{V_{\text{LSB}}}, \tag{3}$$

where V_{ec}^{peak} is the ideal held peak for event e, channel c, and $V_{\text{LSB}} = 125\ \mu\text{V/count}$ is the nominal ideal external-ADC code width for the declared 2.048-V reference. Division by volts per count yields counts. Equation (3) does not add acquisition pedestal, electronic noise, saturation, or quantization; those omissions are deliberate properties of S_{ec}.

2 Mathematical notation and equation-reading guide

This section defines the mathematical language used throughout the report. Every later displayed equation is accompanied by prose that states (i) what each symbol represents, (ii) what operation is performed, (iii) why the operation exists in the forward model, and (iv) what it does not model.

Indices and random variables. The index e labels a generated event, c an independently digitized channel, $s \in \{0, \ldots, 71\}$ a simulated SiPM, $j \in \{0, \ldots, 35\}$ a $PbWO_4$ crystal, b a time bin, and k a scintillation component or spectral grid point. Upper-case symbols such as N denote counts or random variables; lower-case values denote realized or configured quantities as stated locally. The notation $X \sim \mathcal{D}$ means that X is sampled from distribution $\mathcal{D}$; it does not assert that a finite histogram exactly equals its parent density.

Normal and Poisson distributions. The standard-normal density and cumulative distribution function are

$$\phi(z) = \frac{e^{-z^2/2}}{\sqrt{2\pi}}, \qquad \Phi(a) = \int_{-\infty}^{a} \phi(z)\, dz. \tag{4}$$

The density ϕ assigns relative probability per unit z, while $\Phi(a)$ is the probability that a standard-normal draw does not exceed a. For a Poisson random variable,

$$N \sim \text{Poisson}(\lambda), \qquad P(N = n) = e^{-\lambda} \frac{\lambda^n}{n!}, \quad n = 0, 1, 2, \ldots. \tag{5}$$

Here $\lambda \geq 0$ is both the mean and variance. The model uses Poisson draws for discrete nuisance-pulse counts; Eq. (5) is not a fit to the observed ADC spectra.

Bernoulli trials and expectations. A Bernoulli variable $B \sim \text{Bernoulli}(q)$ equals one with probability q and zero otherwise, so $\mathbb{E}[B] = q$. The expectation $\mathbb{E}[X]$ is a model average over repeated random realizations, not necessarily the arithmetic mean of one finite production shard. Conditional notation $P(A \mid B)$ means the probability of A given that B holds.

Time series and sums. $N_s[b]$ denotes the content of time bin b for sensor s. A sum such as $\sum_s N_s[b]$ combines sensors in the same bin; a sum over b integrates a history over its declared time window. The fine-bin width is $\Delta t = 0.25$ ns and the simulated window is $T = 1000$ ns. Sparse storage omits rows whose declared inputs are zero; omission of a row is therefore not the same as an unknown or rejected event.

Units. Momentum is reported in MeV/c, kinetic or deposited energy in MeV, length in mm, time in ns, charge in pC, voltage in V, and native ideal ADC-input-equivalent response in counts. Natural particle-physics notation sometimes suppresses explicit factors of c in relativistic equations; this report restores the units in prose. A spectral quantity tabulated versus wavelength λ is passed to Geant4 versus photon energy $E_\gamma = hc/\lambda$, so increasing wavelength corresponds to decreasing photon energy.

Densities, interpolation, and normalization. A probability density $f(x)$ is normalized when $\int f(x)\, dx = 1$ over its stated support. A tabulated material property such as refractive index, absorption length, reflectivity, or PDE is not a probability density and is not normalized to unit area. "Linear interpolation" between tabulated points (u_k, v_k) means

$$v(u) = v_k + \frac{u - u_k}{u_{k+1} - u_k} \left(v_{k+1} - v_k\right), \qquad u_k \leq u \leq u_{k+1}. \tag{6}$$

The fraction in Eq. (6) specifies where u lies between adjacent grid points; multiplying by the endpoint difference then interpolates the property. This is a piecewise-linear numerical rule, not a claim that the underlying material response is physically linear.

Table 1: Aligned simulation identities. A SHA-256 digest changes if even one input byte changes; matching the digest therefore establishes byte identity, not merely a matching filename.

Item	Frozen value
Executable	167,600 bytes; `d1c70788ec93709b75b7a766dbdd537f91c38a1fd7a9eaaac43fa3b67351bda1`
C++ source	58,179 bytes; `ef8839692415165940f1b2f4e9e4cca12718125aee0a6c631f1ffcf4852becd6`
Frozen GDML	94,331 bytes; `596ab478cbf9d25841465298a99c8c8028199b5eaedea5e288d77e4b600e7f93`
BM alignment JSON	490 bytes; `198d39f60368688de1dca30a300aba2167ae838155b212c06fbcd8811478a265`
Latent beam table	2,205 bytes; `6b5f78c1d96d306ba4a85ff0093732de1de116ca52826f66ced6efce04894ba0`
Geant4	`geant4-11-04-patch-02 [MT]`, GCC 11.4.1 CMS EL9 toolchain
Random engine	CLHEP `MixMaxRng`, reset from a deterministic event seed

Floating counts. The saturation model returns an expected effective number of newly fired cells and therefore permits non-integer values. The circuit convolution and ADC division also remain floating point. Integer-valued words enter only in the real acquisition decoder, which is outside this report's endpoint.

3 Aligned all-nine campaign and positron generator

3.1 Production identity

The aligned optical campaign completed on 21 August 2026. Its checksum-bound run identity is recorded in the source manifest. Production used one `zdc-gdml-beamtest` executable compiled in C++20 release mode with `-O3 -DNDEBUG` against multithreaded Geant4 11.4.2. The base physics list was `FTFP_BERT`; optical mode additionally registered `G4OpticalPhysics`. One direct 10,000-event ROOT job was run for each of the nine energies. No merge or mixed-population sample defines the aligned sample. Table 1 records the common highest-level identities; the package provenance binds the runner, helper, Condor submit, aligned inputs, ROOT sidecars, schema audit, and downstream digitizer handoff.

The aligned production build did not define `EIC_HAS_GEANT4_GDML`. Every job consequently printed `GEOMETRY backend=analytic_from_audited_gdml`. The same GDML was transferred to all jobs and remains the audited geometry specification, while Geant4 constructed the source-coded analytic copy and applied the frozen effective transform to all 128 BM strips. The PWO placement and all z coordinates remained unchanged. This distinction is part of the reproducibility contract: a build with runtime GDML enabled is a model variation even if it reads a file with the same name.

3.2 One campaign, one equal-10k analysis sample

The campaign helper starts from seed base 2026082201, adds 10^5 for optical mode, and then adds the zero-based energy index. This produces the nine seeds in Table 2. Thread count controls execution parallelism and does not change a fixed event because the engine is reset at event start; changing the campaign seed does change the generated event and is therefore recorded explicitly. All nine return codes were zero, every ROOT SHA-256 sidecar verified, and the schema audit found 10,000 unique event IDs and all seven required trees in every file.

Table 2: Execution identity of the aligned optical ROOT at each energy. Each row is one direct job from Condor cluster 995.

Nominal p [MeV/c]	Events	IDs	Threads	Seed	Result
98	10,000	0–9999	168	2026182201	PASS
197	10,000	0–9999	168	2026182202	PASS
297	10,000	0–9999	168	2026182203	PASS
395	10,000	0–9999	168	2026182204	PASS
492	10,000	0–9999	168	2026182205	PASS
584	10,000	0–9999	168	2026182206	PASS
706	10,000	0–9999	168	2026182207	PASS
739	10,000	0–9999	168	2026182208	PASS
796	10,000	0–9999	168	2026182209	PASS

Table 3: Frozen nine-energy campaign and the equal-statistics report sample. The nominal label is a momentum setting; the ROOT energy branch is positron kinetic energy.

Nominal p [MeV/c]	Current [A]	μ_p [MeV/c]	σ_p [MeV/c]	Seed	Threads	Events
98	50	98.19	4.92	2026182201	168	10,000
197	100	197.94	4.91	2026182202	168	10,000
297	150	297.30	5.29	2026182203	168	10,000
395	200	395.90	5.55	2026182204	168	10,000
492	250	492.50	5.83	2026182205	168	10,000
584	300	584.59	6.17	2026182206	168	10,000
706	375	706.61	6.58	2026182207	168	10,000
739	400	739.16	6.71	2026182208	168	10,000
796	450	796.60	6.89	2026182209	168	10,000

3.3 Positive-truncated momentum

The nominal energy names label magnet/beam settings. For each event the code repeatedly drew a Gaussian momentum until the draw was positive. The resulting normalized density is

$$f_p(p \mid \mu_p, \sigma_p) = \frac{\phi((p-\mu_p)/\sigma_p)}{\sigma_p\,\Phi(\mu_p/\sigma_p)}\,\mathbf{1}[p>0], \tag{7}$$

where p is the momentum magnitude, μ_p and σ_p are the configured Gaussian mean and width, ϕ and Φ are defined in Eq. (4), and $\mathbf{1}[p>0]$ equals one only for a positive draw. The denominator is the probability that the original Gaussian lies above zero, so it renormalizes the retained half-line to unit probability. The code implements Eq. (7) by rejection sampling, not by evaluating the density. Because every μ_p/σ_p in this campaign is very large, rejection is practically negligible; the rule nevertheless forbids an unphysical negative momentum exactly.

The sampled momentum was converted to positron kinetic energy by

$$E_{\text{kin}} = \sqrt{p^2c^2 + m_e^2c^4} - m_ec^2. \tag{8}$$

The square root is the relativistic total energy and the subtraction removes the electron rest energy, $m_ec^2 \simeq 0.511$ MeV. Thus the ROOT field `energy_MeV` contains E_{kin} from Eq. (8); it is neither p nor the nominal integer label. At these energies pc and total energy are numerically close, but treating them as identical would discard the explicit generator contract.

3.4 Correlated transverse profile

The aligned campaign uses the latent pre-aperture center, widths, and correlation derived in Section 4; their nine numerical rows are in Table 4. Copying the measured post-strip widths directly would double-apply the BM aperture.

For independent standard-normal draws $Z_x, Z_y \sim \mathcal{N}(0,1)$, the gun coordinates were

$$x = \mu_x + \sigma_x Z_x, \tag{9}$$

$$y = \mu_y + \sigma_y \left(\rho Z_x + \sqrt{1-\rho^2}\, Z_y\right). \tag{10}$$

Here (μ_x, μ_y) is the energy-specific beam center in GDML coordinates, (σ_x, σ_y) are marginal beam widths, and ρ is the configured GDML x–y correlation. The combination in Eq. (10) has variance $\rho^2+(1-\rho^2) = 1$, so y retains width σ_y; sharing Z_x gives $\mathrm{Cov}(x, y) = \rho\sigma_x\sigma_y$. Therefore the equations produce the requested marginal widths and correlation without rotating or resampling the detector geometry. They model a Gaussian core only, with no non-Gaussian halo or event-by-event angular correlation.

For auditing the data-to-generator derivation, the intermediate PWO-local/readout convention is related to GDML coordinates by

$$x_{\text{PWO local}} = -x_{\text{GDML}} + 20.5\ \text{mm}, \qquad y_{\text{PWO local}} = y_{\text{GDML}} + 11.5\ \text{mm}. \tag{11}$$

The inverse of Eq. (11) supplies the final GDML gun center after the BM prediction is translated into the PWO-local frame. The aligned runner consumes those final GDML values directly. Only the x coordinate is mirrored, so $\rho_{\text{GDML}} = -\rho_{\text{PWO local}}$ while both marginal widths remain unchanged. The translations alter means but not widths or correlations.

Why target labels 20, 14, and 15 can all appear. The report target naming convention is crystal 20 / GID 34+40 / $R(2,2)$ / software C14 / GDML `PWO_15`. These names identify one sensor pair in different numbering systems. In the crystal-number convention, $6(5-i_Y)+i_X = 20$ at $R(2,2)$. In the readout coordinate map it is $R(i_X, i_Y) = R(2,2)$, whose software row-major label $C = 6i_Y + i_X$ is C14. Those channels map to simulated sensors 30 and 31 and hence to Geant4/GDML crystal copy 15, whose physical name is `PWO_15`. These are aliases from different numbering systems, not different beam targets. In particular, the readout map reflects the horizontal index, $i_X = 5 - r_x$, while preserving the row, $i_Y = r_y$. All nine requested beam means lie inside `PWO_15`. The frozen map does *not* reflect i_Y: positive readout i_Y points toward positive GDML y (upward in Figure 1), while positive readout i_X points toward negative GDML x (leftward).

3.5 Vertex, direction, and truth plane

Every primary began at $(x, y, -50\ \text{mm})$ with direction exactly $(0, 0, +1)$. There was no generated angular divergence and no upstream magnet, collimator, converter, air path, or beamline material between the gun and the first beam-monitor station in this simplified world. Transport can subsequently change the primary direction through interactions.

A virtual plane at $z_\star = 125.0$ mm recorded the first forward crossing of the primary track. If a step runs from positions $\mathbf{r}_0$ to $\mathbf{r}_1$ and brackets that plane, the recorded crossing is the linear step interpolation

$$\alpha = \frac{z_\star - z_0}{z_1 - z_0}, \qquad \mathbf{r}_\star = \mathbf{r}_0 + \alpha(\mathbf{r}_1 - \mathbf{r}_0). \tag{12}$$

The dimensionless fraction α identifies where within the Geant4 step the plane is crossed. Equation (12) avoids adding a physical scoring volume and therefore does not perturb transport. It assumes the Geant4 step is represented by the straight segment between its endpoints for this bookkeeping purpose. The plane lies 0.45 mm upstream of the PWO front face at 125.45 mm.

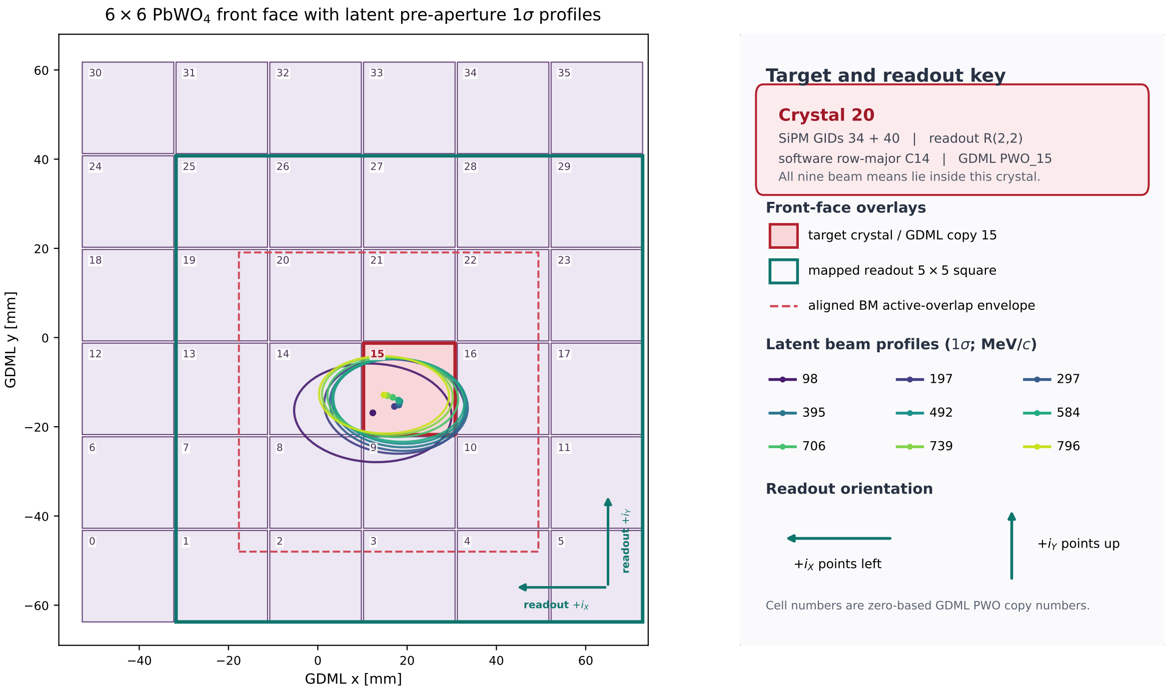

Figure 1: Measured energy-dependent Gaussian beam cores transformed to the GDML coordinate convention and shown on the $PbWO_4$ front face. The numerical means, widths, and correlations are the data-derived latent pre-aperture inputs; these are actual fit outputs, not illustrative beam spots. Box numbers are zero-based GDML copy numbers; the solid teal square encloses the mapped hardware/readout 5×5 footprint, and the small teal labels give readout $R(i_X, i_Y)$. The teal arrows show that $+i_X$ points left and $+i_Y$ points up in this GDML view. The dashed red rectangle is the BM active-overlap envelope after the common runtime transform in Eq. (22). The red target box is crystal 20 / GID 34+40 / $R(2, 2)$ / software C14 / GDML `PWO_15`. Every plotted beam mean lies inside it. The legend labels denote nominal momentum settings in MeV/c.

3.6 Schedule-independent event seeds

The global event identity is $e = o + \ell$, where o is the configured event offset and ℓ is the local Geant4 event number. Before generating each event, the program computes

$$u_0 = C + \texttt{0x9e3779b97f4a7c15}\,(10000i + o + \ell + 1) \pmod{2^{64}}, \tag{13}$$

$$u_1 = (u_0 \oplus (u_0 \gg 30))\texttt{0xbf58476d1ce4e5b9} \pmod{2^{64}},$$

$$u_2 = (u_1 \oplus (u_1 \gg 27))\texttt{0x94d049bb133111eb} \pmod{2^{64}},$$

$$u_3 = u_2 \oplus (u_2 \gg 31), \qquad s = 1 + (u_3 \bmod 2147483561). \tag{14}$$

Here C is the campaign seed, $i \in \{0, \ldots, 8\}$ is the global energy index, $\oplus$ is bitwise exclusive-or, $\gg$ is an unsigned right shift, and multiplication wraps modulo 2^{64}. The SplitMix64-style sequence diffuses nearby integer event identifiers into unrelated bit patterns; the final modulo maps them to the positive seed range accepted by the engine. Resetting CLHEP to s at event start makes each event reproducible regardless of which worker thread processes it. This guarantees deterministic identity, not statistical independence in the mathematical sense. In the aligned campaign, every job has $o = 0$, and $C = 2026182201 + i$ for energy index $i = 0, \ldots, 8$. Thus every energy contains the direct event-ID range 0–9999 from one job; there is no merged population or alternate event-offset branch in the analysis sample.

4 Data-constrained effective beam-monitor placement

This section defines the data-constrained placement method. The objective is detector-level closure: after Geant4 transports the beam through the simulated strip planes, applying the same reconstruction to Monte Carlo and data should return the same beam-profile core. The result is an *effective analysis-coordinate placement*. It is not a claim that the mechanical BM pose has been surveyed uniquely in three translations and three rotations.

4.1 One target, five naming conventions

The report target naming convention is crystal 20 / GID 34+40 / $R(2,2)$ / software C14 / GDML `PWO_15`. These names come from five numbering systems but identify the same mapped sensor pair. In the crystal-number convention, the zero-based top-down index is

$$C_{\text{top down}} = 6(5 - i_Y) + i_X = 6(5 - 2) + 2 = 20. \tag{15}$$

Equation (15) defines the crystal-number label; it does not add a geometry constraint. In the saved mapping the pair is at readout coordinate $R(i_X, i_Y) = R(2,2)$, for which the software row-major label $C = 6i_Y + i_X$ is C14. Reflecting the horizontal readout index into the GDML array gives copy number 15 and physical name `PWO_15`. Thus the report naming chain is

$$\boxed{\text{crystal 20} \longleftrightarrow \text{GID 34+40} \longleftrightarrow R(2,2) \longleftrightarrow \text{software C14} \longleftrightarrow \text{GDML PWO\_15}} \tag{16}$$

The arrows in Eq. (16) mean that all five names select the same mapped sensor pair; they do not assert that their integers should be equal. In particular, C14 is a software row-major label, not the crystal-number label. Changing which alias is printed on a plot does not move the target or the beam.

The front-face plot encloses the mapped hardware/readout region with one square 5×5 boundary. Its readout axes are drawn explicitly: $+i_Y$ points upward, along positive GDML y, whereas $+i_X$ points leftward, along negative GDML x. The target box itself carries crystal 20, GID 34+40, $R(2,2)$, software C14, and GDML `PWO_15` so that no color or copy number has to carry the crosswalk implicitly.

4.2 Frozen raw target and the residual that fixes the origin

The raw target is not re-decoded for this analysis. It is the existing 594,572-event sample for which the straightest candidate has exactly zero reconstructed angle in both transverse projections. For event e, let

$$\mathbf{b}_e^{\text{raw}} = \begin{pmatrix} x_e^{\text{raw}} \\ y_e^{\text{raw}} \end{pmatrix} \tag{17}$$

denote its selected BM position in the frozen raw global millimetre map. Equation (17) is a coordinate definition. It contains no fit and no Geant4 quantity. The station names require care: simulated station 1 corresponds to raw BM2/T2 and simulated station 2 corresponds to raw BM1/T1. Within each simulated station, plane 0 is long in x and therefore measures y, while plane 1 is long in y and therefore measures x.

The effective transverse origin comes from a separate, event-matched calorimeter relation. If $\mathbf{p}_e$ is the reconstructed PWO position and $\mathbf{b}_e$ is the matched BM position in their common data frame, the residual is defined as

$$\mathbf{d}_e = \mathbf{p}_e - \mathbf{b}_e, \qquad \overline{\mathbf{d}}_{\text{equal energy}} = \begin{pmatrix} +4.312389986 \\ -3.306717162 \end{pmatrix} \text{mm}. \tag{18}$$

The subtraction order in Eq. (18) is PWO minus BM. The reported vector is an equal-energy average, so a high-statistics energy cannot dominate merely because it has more events. It says how far an event-matched PWO response centroid lies from the BM track in the data coordinate system. It is *not* the measured beam center, not the direct runtime translation applied to a GDML strip, and not a mechanical survey.

Let the nominal PWO center in the raw data frame be $\mathbf{c}_{\text{PWO}}^{\text{raw}} = (20.90, 132.62)$ mm. The BM-implied local PWO coordinate is

$$\mathbf{q}_e = \mathbf{b}_e^{\text{raw}} + \overline{\mathbf{d}}_{\text{equal energy}} - \mathbf{c}_{\text{PWO}}^{\text{raw}}. \tag{19}$$

The residual first translates the BM prediction to the PWO response frame; subtracting $\mathbf{c}_{\text{PWO}}^{\text{raw}}$ then expresses it relative to the PWO origin. This is why a local tracked point near $(4, -3)$ mm and a raw global point near $(21, 133)$ mm may describe the same event. Data and Monte Carlo must be compared only after both are expressed in the same named frame.

The PWO-local to GDML convention reflects x but preserves y:

$$x_{\text{GDML}} = 20.5 \text{ mm} - q_x, \qquad y_{\text{GDML}} = q_y - 11.5 \text{ mm}. \tag{20}$$

Substituting Eq. (19) into Eq. (20) gives the single frozen affine map

$$\boxed{x_{\text{GDML}} = -x_{\text{raw}} + 37.087610014 \text{ mm}, \qquad y_{\text{GDML}} = y_{\text{raw}} - 147.426717162 \text{ mm}.} \tag{21}$$

Translations change a mean but not a width. The x reflection changes the sign of a covariance, so $\rho_{xy}^{\text{GDML}} = -\rho_{xy}^{\text{raw}}$, while σ_x and σ_y retain their magnitudes. Equation (21), rather than the residual vector alone, is the coordinate map used for the beam generator.

4.3 Runtime transform of the strip lattice

The current source-coded strip coordinates use a different transverse origin and handedness from Eq. (21). Relative to that coded lattice, the effective aligned placement is

$$x_{\text{BM}}^{\text{new}} = -x_{\text{BM}}^{\text{old}} + 18.917660014 \text{ mm}, \qquad y_{\text{BM}}^{\text{new}} = y_{\text{BM}}^{\text{old}} - 11.406742162 \text{ mm}. \tag{22}$$

Equation (22) is applied after constructing all 128 BM strip volumes, whether their source is the analytic fallback or a runtime-parsed GDML. The identical transform is used for both stations and both planes. It changes no BM z coordinate, no PWO placement, no strip dimension, and no gun direction. The corresponding identity-default command line values are

```
--bm-x-sign -1 \
--bm-offset-x-mm 18.917660014 \
--bm-offset-y-mm -11.406742162
```

The word "offset" in these switches refers to the post-reflection coded-GDML offset. It therefore must not be replaced by the $(+4.312, -3.307)$ mm PWO-minus-BM residual from Eq. (18); the two vectors belong to different coordinate operations.

4.4 Why the measured width cannot be copied into the gun

The measured BM sample has already passed through a finite detector. Each axis has 32 strips of width $w = 2.000$ mm on a pitch of approximately 2.099 mm. The roughly 0.099-mm gaps and the finite 32-strip aperture remove some incident particles, and reconstruction replaces a continuous hit by a strip-center coordinate. Consequently the measured core width is a truncated and quantized observable. Using it directly as the incident Gaussian width would let the simulation aperture truncate the profile a second time.

For one energy, the latent incident coordinate $\mathbf{r} = (x, y)^\mathsf{T}$ is modeled as

$$\mathbf{r} \sim \mathcal{N}_2(\boldsymbol{\mu}, \boldsymbol{\Sigma}), \qquad \boldsymbol{\Sigma} = \begin{pmatrix} \sigma_x^2 & \rho\sigma_x\sigma_y \\ \rho\sigma_x\sigma_y & \sigma_y^2 \end{pmatrix}. \tag{23}$$

Here $\boldsymbol{\mu}$ is the continuous pre-aperture center, $\sigma_x, \sigma_y > 0$ are its marginal standard deviations, and ρ is its correlation. The covariance matrix in Eq. (23) retains both marginal widths and the tilt of the two-dimensional Gaussian core.

Let strip centers be x_i and y_j. The active rectangle associated with a coincident strip pair is

$$\mathcal{C}_{ij} = [x_i - w/2, x_i + w/2] \times [y_j - w/2, y_j + w/2]. \tag{24}$$

Because the pitch is wider than w, the rectangles in Eq. (24) do not tile the plane: the uninstrumented gaps are retained. The unconditioned probability of pair (i, j) is

$$p_{ij}(\boldsymbol{\theta}) = \int_{x_i - w/2}^{x_i + w/2} f_X(x) \left[\Phi\left(\frac{y_j + w/2 - \mu_{Y|x}}{\sigma_{Y|x}} \right) - \Phi\left(\frac{y_j - w/2 - \mu_{Y|x}}{\sigma_{Y|x}} \right) \right] dx, \tag{25}$$

where

$$\mu_{Y|x} = \mu_y + \rho\frac{\sigma_y}{\sigma_x}(x - \mu_x), \qquad \sigma_{Y|x} = \sigma_y\sqrt{1 - \rho^2}. \tag{26}$$

In Eq. (25), f_X is the marginal normal density and Φ is the standard-normal CDF. Equation (26) is the conditional normal distribution of Y at a fixed x. It turns the inner y integral into a difference of two CDF values. The remaining x integral is evaluated deterministically with 16-point Gauss–Legendre quadrature in every strip; no random Monte Carlo integration noise enters the fit.

The modeled ideal active-area acceptance and the conditional strip-pair probabilities are

$$A(\boldsymbol{\theta}) = \sum_{i=1}^{32}\sum_{j=1}^{32} p_{ij}, \qquad q_{ij}(\boldsymbol{\theta}) = \frac{p_{ij}}{A(\boldsymbol{\theta})}. \tag{27}$$

The first quantity is the fraction of the latent Gaussian lying in an active X–Y strip intersection. Dividing by it makes the q_{ij} sum to one among ideal accepted hits. This conditioning models aperture and gaps, but it does not claim that A equals the observed raw efficiency: thresholds, dead/noisy channels, multi-strip response, and candidate choice are not in this ideal geometrical factor.

4.5 Moment fit and uncertainty

Both data counts and forward-folded probabilities are quantized to the strip centers (x_i, y_j). Each marginal starts from a weighted median and 1.4826 times its weighted median absolute deviation; mean and standard deviation are then recomputed five times inside a symmetric $\pm 2.5\sigma$ window. Correlation is computed on the joint surviving X–Y core. Denote this complete deterministic operation by $\mathbf{h}(\boldsymbol{\theta})$.

The five-dimensional fit coordinate is

$$\mathbf{t} = \left(\mu_x^{\mathrm{core}} \quad \mu_y^{\mathrm{core}} \quad \log \sigma_x^{\mathrm{core}} \quad \log \sigma_y^{\mathrm{core}} \quad \operatorname{atanh} \rho_{xy}^{\mathrm{core}}\right)^{\mathsf{T}}. \tag{28}$$

Taking logarithms makes a negative fitted width impossible. The Fisher transform atanh ρ maps the physical interval $(-1, 1)$ to the real line and treats positive and negative correlations symmetrically near zero. Four hundred deterministic multinomial bootstrap replicas of the observed 32×32 table give the full covariance matrix $\mathbf{V}$ of Eq. (28). A tiny eigenvalue floor is used only to make this empirical covariance numerically positive definite; it is not a physics regularizer on the fitted beam.

The per-energy fit minimizes the correlated squared residual

$$\chi^2(\boldsymbol{\theta}) = [\mathbf{h}(\boldsymbol{\theta}) - \mathbf{t}_{\mathrm{data}}]^{\mathsf{T}}\, \mathbf{V}^{-1}\, [\mathbf{h}(\boldsymbol{\theta}) - \mathbf{t}_{\mathrm{data}}] = \left\| \mathbf{L}^{-1}\, [\mathbf{h}(\boldsymbol{\theta}) - \mathbf{t}_{\mathrm{data}}] \right\|_2^2, \tag{29}$$

where $\mathbf{V} = \mathbf{L}\mathbf{L}^{\mathsf{T}}$ is its Cholesky factorization. Multiplication by $\mathbf{L}^{-1}$ measures residuals in units of their correlated bootstrap uncertainty; it prevents two strongly related moment discrepancies from being counted as independent evidence. This is a five-moment χ^2 closure objective, not a per-event negative log-likelihood and not a per-strip electronics-efficiency fit. Nine deterministic starting points probe convergence. The local covariance $(\mathbf{J}^{\mathsf{T}}\mathbf{J})^{+}$ uses the fitted whitened-residual Jacobian $\mathbf{J}$ and a pseudoinverse; it is a local linear uncertainty approximation, while the saved bootstrap percentiles expose nonlinearity.

4.6 Independent detector-level closure gate

The latent moment fit can close its own deterministic forward model and still fail in Geant4 because transport, strip energy deposition, station ordering, or reconstruction may differ. The non-optical 10k pilot therefore reconstructs MC using every finite positive-energy strip, considers every station-pair candidate, minimizes absolute angle with deposited energy as an exact-angle tie-break, retains exact-zero X and Y tracks, converts selected copy IDs into the frozen raw strip-center map, and applies the same five-iteration $\pm 2.5\sigma$ core estimator.

For a mean coordinate $u \in \{x, y\}$, the combined standard error and hard allowance are

$$s_{\Delta\mu_u} = \sqrt{s_{\mu_u,\mathrm{MC}}^2 + s_{\mu_u,\mathrm{data}}^2}, \qquad |\mu_{u,\mathrm{MC}} - \mu_{u,\mathrm{data}}| \le \max(1.05\ \mathrm{mm}, 2 s_{\Delta\mu_u}). \tag{30}$$

Independent finite-sample errors add in quadrature in Eq. (30). The fixed 1.05-mm floor is one half of the approximately 2.099-mm pitch, so a statistically precise but sub-strip-scale offset is not rejected solely because its standard error is tiny.

For a width, the analyzer combines the two absolute width standard errors and then normalizes by the frozen raw width:

$$s_{\Delta\sigma_u/\sigma_u} = \frac{\sqrt{s_{\sigma_u,\mathrm{MC}}^2 + s_{\sigma_u,\mathrm{data}}^2}}{\sigma_{u,\mathrm{data}}}, \qquad \frac{|\sigma_{u,\mathrm{MC}} - \sigma_{u,\mathrm{data}}|}{\sigma_{u,\mathrm{data}}} \le \max(0.05, 2 s_{\Delta\sigma_u/\sigma_u}). \tag{31}$$

The two terms under the square root in Eq. (31) are the independent MC- and data-width uncertainties. The raw width is the declared reference scale rather than another fitted denominator inside the

Table 4: Data-derived latent pre-aperture generator profile for the aligned campaign. Means, widths, and correlation are in GDML coordinates. The ideal active fraction is $A(\boldsymbol{\theta})$ from Eq. (27); it is not a predicted raw electronics efficiency.

Nominal p (MeV/c)	μ_x (mm)	μ_y (mm)	σ_x (mm)	σ_y (mm)	ρ_{xy}	A (%)
98	12.3229	-16.8759	17.6240	11.0544	-0.04906	84.93
197	17.1671	-15.4458	15.7344	10.5922	-0.05289	87.62
297	18.1432	-15.1258	15.4412	10.2890	-0.06188	87.84
395	18.2836	-14.6493	15.1039	9.9533	-0.07086	88.17
492	18.3880	-14.2772	14.9061	9.5536	-0.06134	88.36
584	18.0645	-14.0236	14.7503	9.3461	-0.06399	88.55
706	16.7916	-13.3991	14.6431	8.9020	-0.04873	88.77
739	15.5893	-13.0325	14.6544	8.7713	-0.04342	88.78
796	14.8302	-12.8534	14.5563	8.6670	-0.02580	88.84

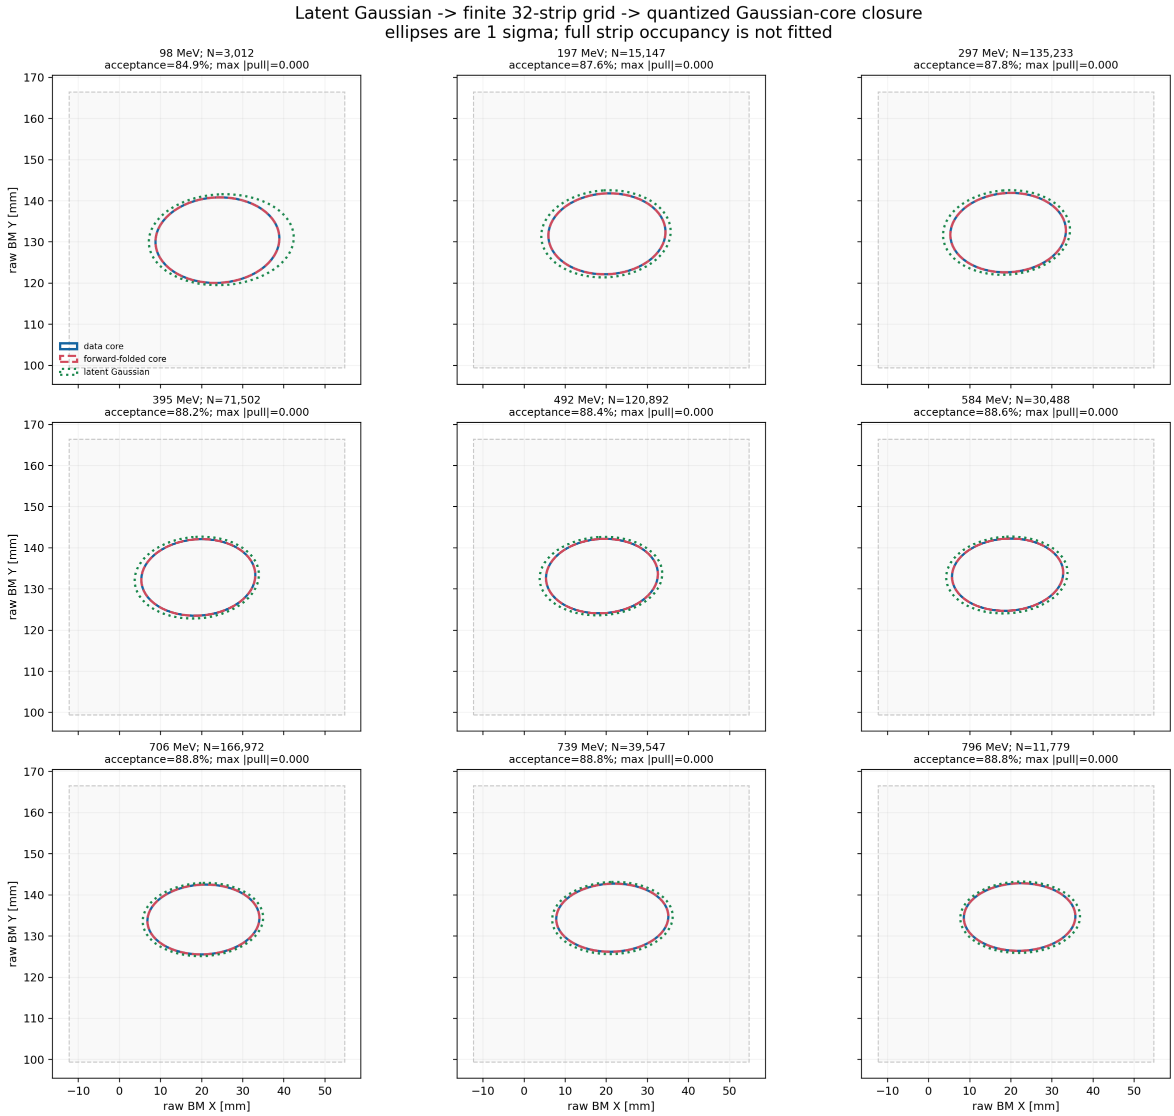

Figure 2: Frozen raw BM Gaussian-core targets and their latent-Gaussian forward folds through the finite, gapped, quantized strip lattice. These are real data-derived fit products, not illustrative curves. The independent detector-level Geant4 closure is evaluated in the following subsection.

uncertainty propagation. The 5% floor represents the declared core-profile tolerance rather than an electronics model.

The correlation gate is

$$s_{\Delta\rho} = \sqrt{s^2_{\rho,\mathrm{MC}} + s^2_{\rho,\mathrm{data}}}, \qquad |\rho_{\mathrm{MC}} - \rho_{\mathrm{data}}| \le \max(0.03, 2s_{\Delta\rho}). \tag{32}$$

Equation (32) tests the signed ellipse tilt separately from its two marginal widths. Acceptance required all three gate families to pass at all nine energies. The gate uses only the selected core moments: Geant4 supplies $E_{\mathrm{dep}} > 0$, whereas the raw acquisition includes thresholds and other electronics effects that this geometry closure does not emulate.

The completed non-optical 10k analysis passed every hard gate at every energy. The exact-zero reconstructed MC population ranges from 6,291 to 8,563 events per 10,000 generated events. Across all nine settings, the largest absolute mean residual is 0.243 mm, the largest relative width residuals are 2.28% in x and 2.92% in y, and the largest absolute correlation residual is 0.0201. These values are all inside their per-energy limits in Table 5. The retained fraction is reported for diagnosis only; it is not interpreted as raw-electronics efficiency.

4.7 Selected-position diagnostics before the exact-zero cut

Four plane-resolved diagnostics compare the selected strip position in BM1-X, BM1-Y, BM2-X, and BM2-Y across all nine nominal momenta. Their population is the four-plane-trackable sample: an event must contain enough accepted strip information to form the straightest candidate in both stations and both projections. The histograms are filled immediately after that candidate is chosen and *before* the exact-zero-angle requirement used by the hard closure sample. They therefore expose the position population removed by the later angular cut rather than silently inheriting its sculpting.

The two reconstruction paths use observables native to their respective detectors. Data candidates are formed from fired strips, with ADC used to break an exact-angle tie. Monte Carlo candidates are formed from strips with positive deposited energy, with deposited energy used for the corresponding tie-break. Neither procedure treats ADC and deposited energy as numerically equivalent; each quantity only orders otherwise angle-degenerate candidates within its own sample.

For every momentum and every plane, the data and Monte Carlo histograms are normalized independently to unit area. Thus their overlay compares selected position *shape*, not event yield. A difference can reveal station- or plane-dependent acceptance, strip gaps and quantization, angular-population differences, or raw-electronics threshold and efficiency effects absent from the positive-deposit Monte Carlo definition. Because the normalization removes yield information and the detector observables are intentionally not identical, these four figures are diagnostics only. They do not modify the effective placement, latent beam parameters, or any mean, width, or correlation gate defined above.

Table 5: Detector-level BM closure against the frozen raw exact-zero target. Each residual is followed by its hard limit in parentheses. The table is generated from `provenance/bm_closure_metrics.csv`; occupancy and cut flow are diagnostics, not gates.

E (MeV)	$N_{\rm MC}^{0^\circ}$	$\Delta\mu_x$ (limit) (mm)	$\Delta\mu_y$ (limit) (mm)	$\lvert\Delta\sigma_x\rvert/\sigma_x$ (limit) (%)	$\lvert\Delta\sigma_y\rvert/\sigma_y$ (limit) (%)	$\lvert\Delta\rho\rvert$ (limit)	Verdict
98	6,291	-0.243 (1.050)	-0.125 (1.050)	0.86 (5.00)	0.09 (5.00)	0.0195 (0.0446)	PASS
197	7,652	-0.168 (1.050)	-0.042 (1.050)	0.30 (5.00)	0.26 (5.00)	0.0176 (0.0300)	PASS
297	8,082	+0.059 (1.050)	-0.009 (1.050)	1.10 (5.00)	1.38 (5.00)	0.0061 (0.0300)	PASS
395	8,246	+0.160 (1.050)	-0.100 (1.050)	0.48 (5.00)	0.21 (5.00)	0.0085 (0.0300)	PASS
492	8,400	+0.028 (1.050)	-0.093 (1.050)	0.81 (5.00)	0.44 (5.00)	0.0021 (0.0300)	PASS
584	8,444	+0.187 (1.050)	-0.051 (1.050)	2.28 (5.00)	0.73 (5.00)	0.0007 (0.0300)	PASS
706	8,558	-0.144 (1.050)	+0.067 (1.050)	0.12 (5.00)	0.88 (5.00)	0.0006 (0.0300)	PASS
739	8,563	-0.050 (1.050)	-0.144 (1.050)	0.58 (5.00)	2.92 (5.00)	0.0090 (0.0300)	PASS
796	8,549	+0.175 (1.050)	-0.033 (1.050)	0.63 (5.00)	0.04 (5.00)	0.0201 (0.0300)	PASS

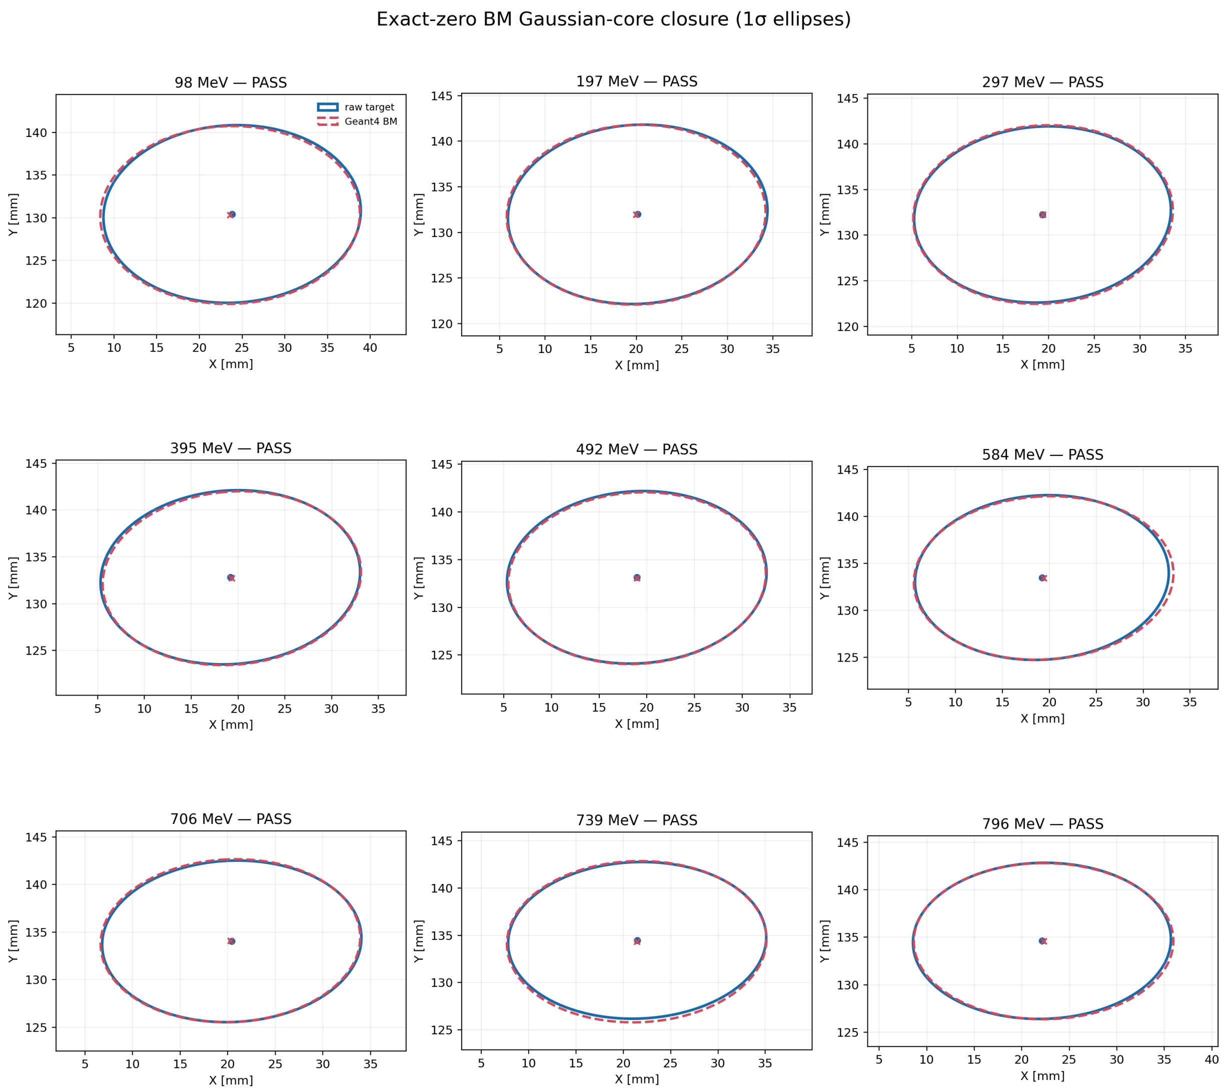

Figure 3: All-nine raw-versus-MC exact-zero BM core ellipses after the effective runtime placement. The accompanying machine-readable verdict is `provenance/bm_closure_summary.json`.

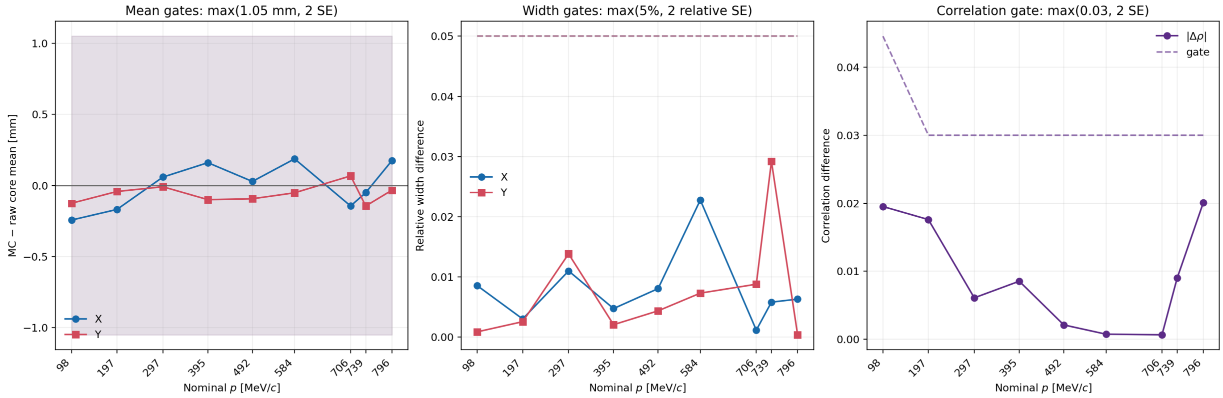

Figure 4: Per-energy mean, width, and correlation residuals against their hard closure limits. This plot and `provenance/bm_closure_metrics.csv` must be generated by the same analyzer invocation.

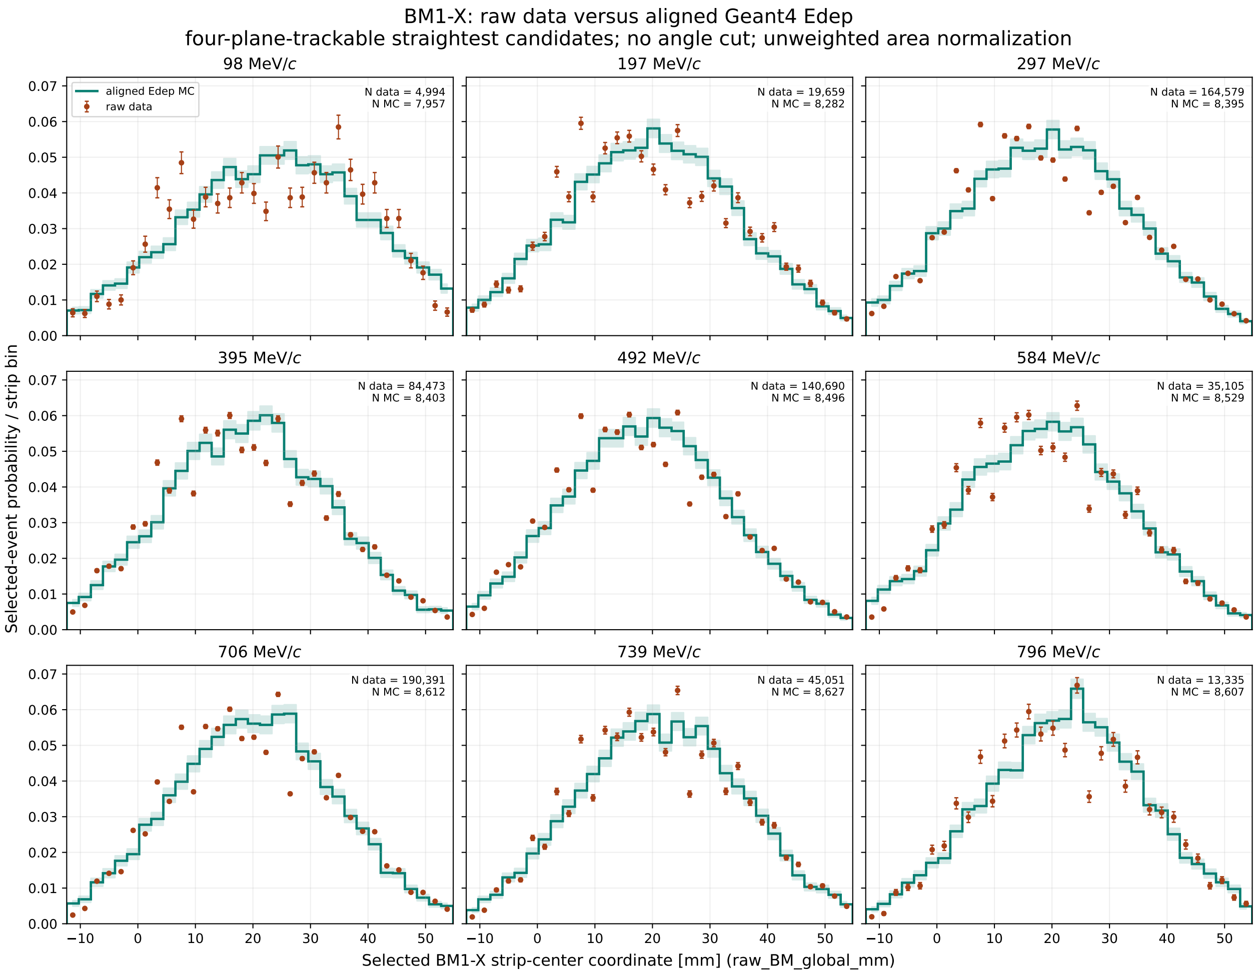

Figure 5: BM1-X selected strip-center positions for the four-plane-trackable straightest candidates before the exact-zero cut. Each panel is one nominal momentum setting in MeV/c. The raw-data fired-strip distribution and the aligned positive-deposit Monte Carlo distribution are normalized separately to unit area; ADC and deposited energy are used only for their respective exact-angle tie-breaks. This is a position-shape diagnostic, not a closure gate or an efficiency measurement.

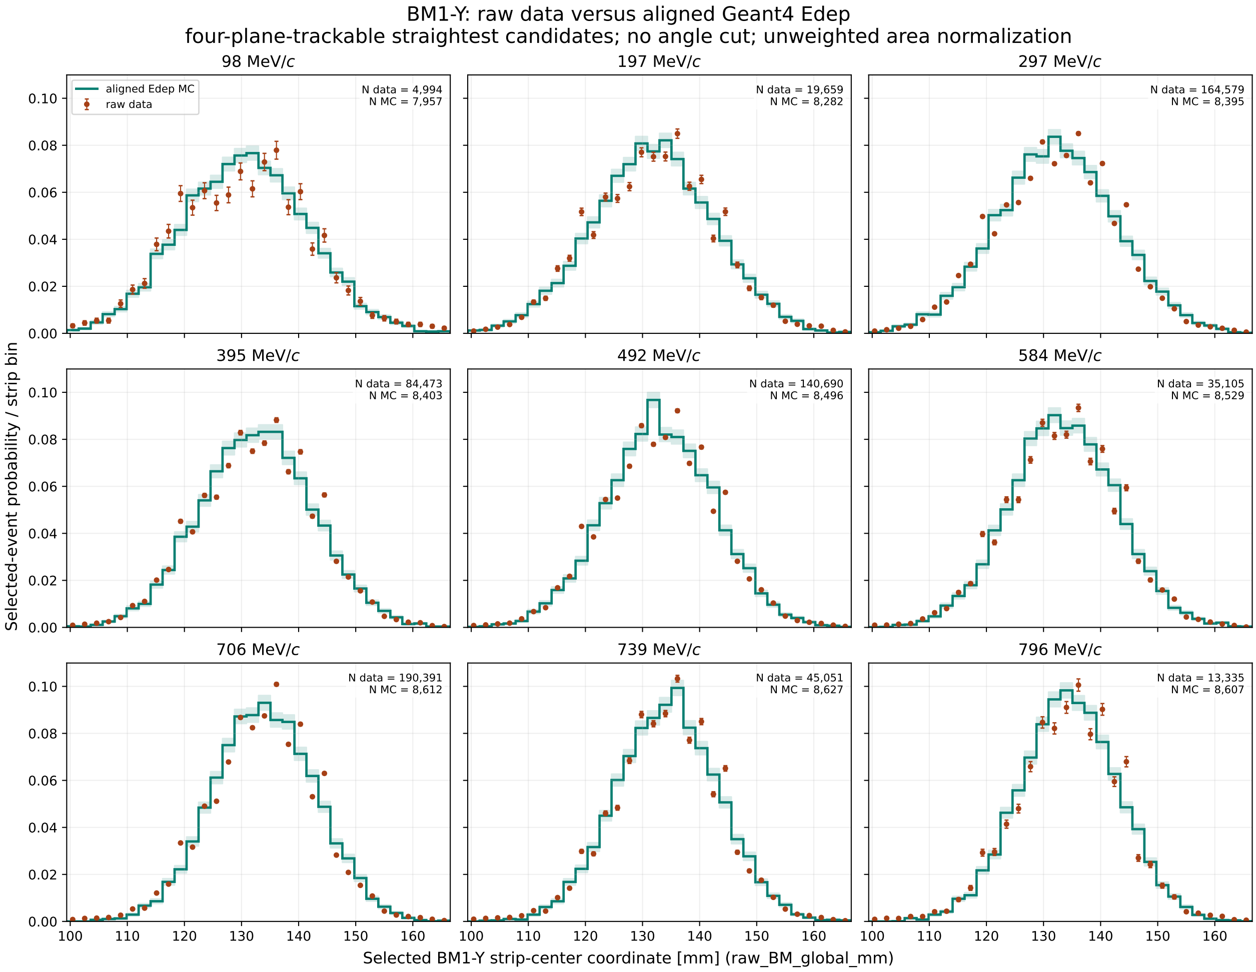

Figure 6: BM1-Y selected strip-center positions for the same four-plane-trackable pre-zero-cut population. Data use fired strips with an ADC tie-break; Monte Carlo uses positive-deposit strips with a deposited-energy tie-break. The independent unit-area normalization at each momentum compares shape only and does not encode relative acceptance or efficiency.

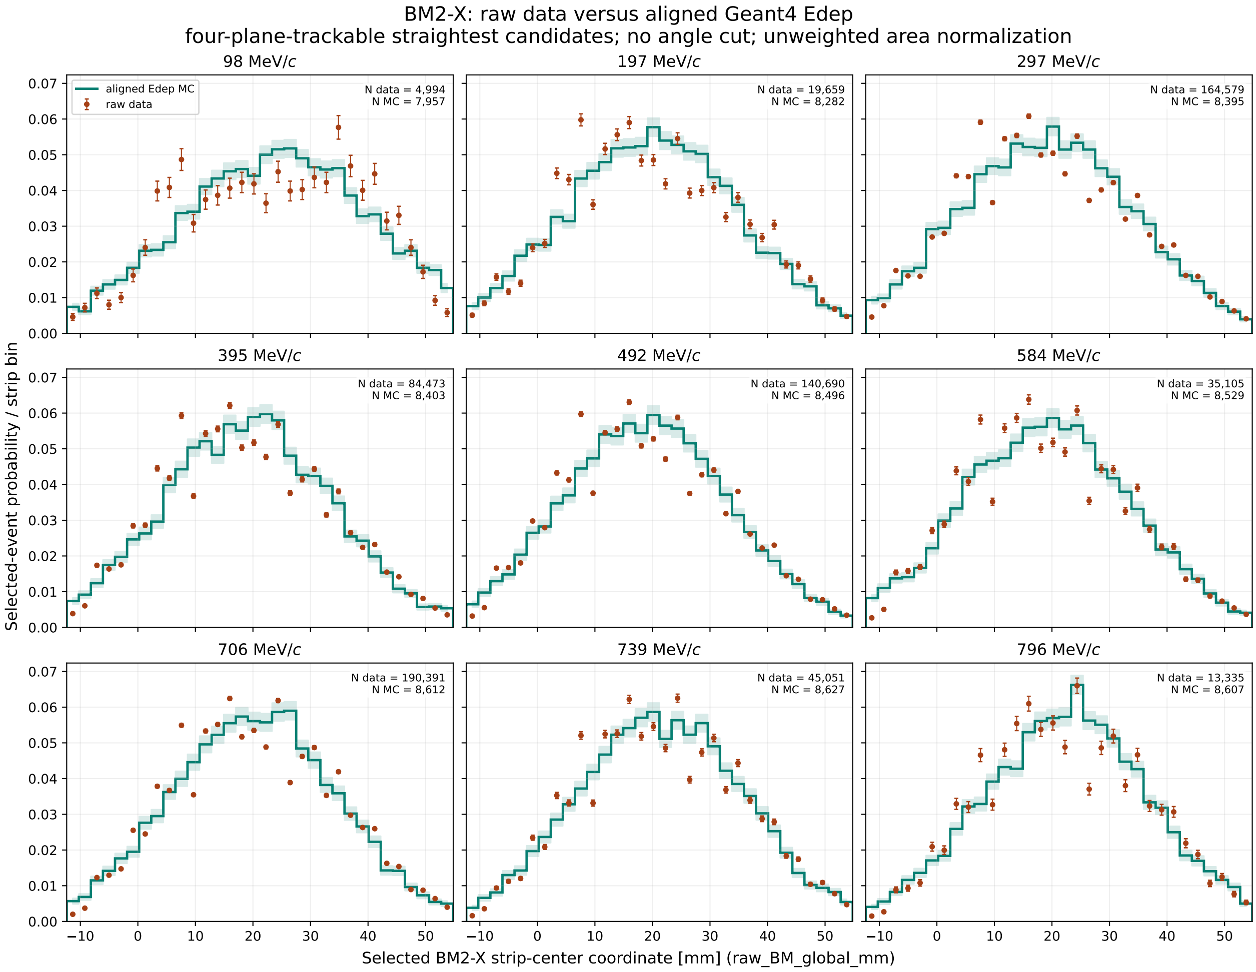

Figure 7: BM2-X selected strip-center positions for the four-plane-trackable straightest candidates before the exact-zero cut. The independently normalized data and Monte Carlo curves compare the selected-position shape after their detector-specific ADC and deposited-energy tie-breaks. Yield and efficiency are intentionally outside this diagnostic.

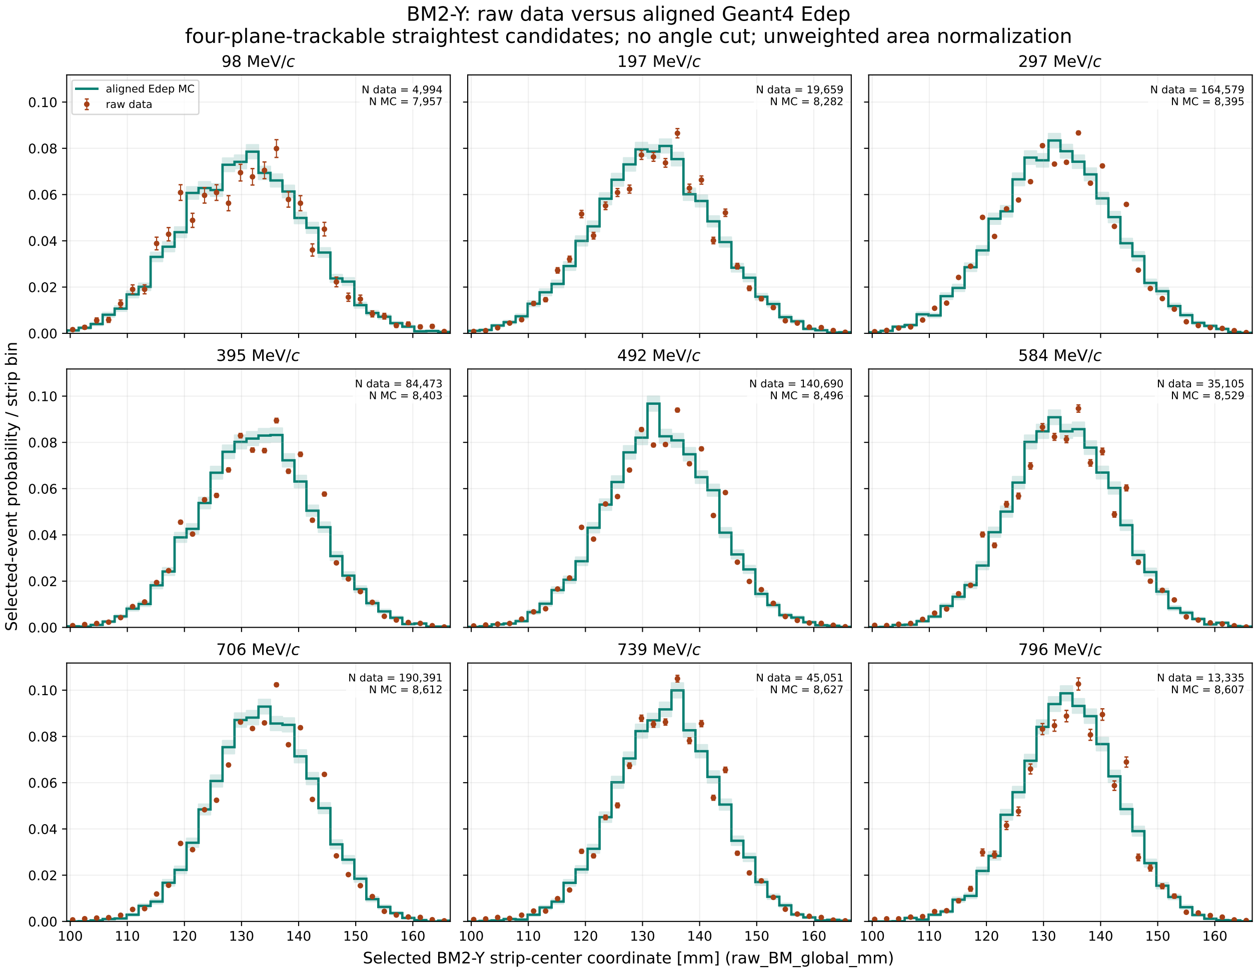

Figure 8: BM2-Y selected strip-center positions for the same four-plane-trackable population before the exact-zero cut. Separate unit-area normalization makes this a shape overlay between fired-strip data and positive-deposit Monte Carlo; differences can reflect detector response, angular population, strip gaps, or efficiency, but do not alter the alignment fit or its hard gates.

The closure source manifest binds the plotted BM result to the raw-target hash, latent-fit hash, runtime binary and source, alignment JSON, runner, Condor submit description, ROOT inputs, and closure analyzer. All nine hard gates passed. The final source manifest extends this chain through the nine audited optical ROOT inputs, the cluster-998 digitizer audit, and the native NPZ products used by the complete Geant4-to-native-ADC report.

5 Analytic geometry and runtime materials

The production geometry is an analytic, source-coded reproduction of the frozen GDML. All placements are unrotated daughters of the world. The inventory closure is

$$N_{\text{place}} = 1_{\text{ world}} + 36_{\text{ PWO}} + 128_{\text{ BM}} + 72_{\text{ grease}} + 72_{\text{ window}} + 72_{\text{ active}} = 381. \tag{33}$$

Equation (33) is a count audit: it adds one world, 36 crystals, four planes of 32 beam-monitor strips, and three rear layers for each of 72 sensors. It establishes that expected components are present; by itself it does not prove absence of overlap or physical fidelity to the full mechanical detector.

5.1 World and $PbWO_4$ array

The world is a $400 \times 400 \times 500$ mm^3 box centered at the origin and filled with `G4_Galactic`. The calorimeter contains 36 $PbWO_4$ crystals in six rows and six columns. Each crystal is $20.5 \times 20.5 \times 53.4$ mm^3 on a 21.0-mm pitch, leaving a nominal 0.5-mm transverse gap. Its center is at $z = 152.15$ mm, so the front and rear faces are at 125.45 and 178.85 mm.

For zero-based row r and column c, the analytic placement is

$$j = 6r + c, \tag{34}$$

$$x_j = -42.5000005 \text{ mm} + c(21.0 \text{ mm}),$$
$$y_j = -53.500025 \text{ mm} + r(21.0 \text{ mm}). \tag{35}$$

Equation (34) is the row-major copy-number mapping. Equation (35) starts at the lower-left analytic center and adds one pitch per column or row. The 0.5-nm x offset and 25-nm y offset inherited from the source are preserved as placement data, not interpreted as survey precision. With the runtime $PbWO_4$ radiation length $X_0 = 8.9245$ mm, the modeled longitudinal depth is

$$d/X_0 = 53.4 \text{ mm}/(8.9245 \text{ mm}) \simeq 5.98. \tag{36}$$

The units cancel in Eq. (36), leaving the number of radiation lengths. This describes the simulated crystal depth; it does not include upstream material or longitudinal leakage corrections.

5.2 Beam-monitor stations

Each of two stations has two 32-strip planes. The source-labelled X-plane strips are $80 \times 2 \times 2$ mm^3, long along GDML x and stepped along GDML y; the source-labelled Y-plane strips are $2 \times 80 \times 2$ mm^3, long along GDML y and stepped along GDML x. Thus a plane's source label is not, by itself, a mathematical statement about its indexed coordinate.

The nominal station midplanes are 0 and 35 mm, hence their separation is 35 mm. Strip centers are read literally from the audited 32-element placement arrays and have approximately 2.0992-mm pitch. The small printed nonuniformity comes from frozen input rounding. The nominal geometry does not contain the fitted BM/PWO transform. Production applies the common transverse effective transform in Eq. (22) to all 128 strips; it leaves these z positions and the PWO geometry unchanged.

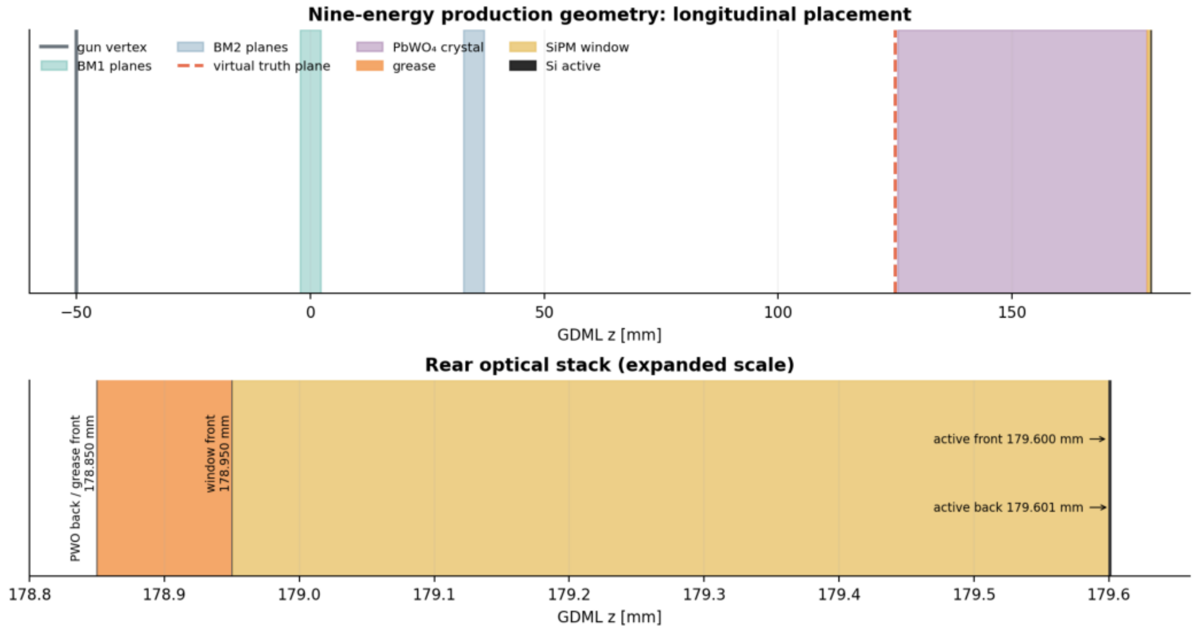

Figure 9: Actual longitudinal layout reconstructed from the frozen placement inventory. The virtual truth plane is a scorer, not a physical slab.

Table 6: Beam-monitor plane positions. Each strip is 2 mm thick in z.

Plane	Center z (mm)	z extent (mm)	Copy numbers
BM1 Y	-1.2	-2.2 to -0.2	32–63
BM1 X	+1.2	+0.2 to +2.2	0–31
BM2 Y	33.8	32.8 to 34.8	32–63
BM2 X	36.2	35.2 to 37.2	0–31

Table 7: Simplified rear coupling chain for each simulated sensor.

Layer	Full size (mm^3)	Center z (mm)	z extent (mm)
Optical grease	$7 \times 7 \times 0.100$	178.9000	178.850–178.950
SiPM window	$7 \times 7 \times 0.650$	179.2750	178.950–179.600
Active silicon	$6 \times 6 \times 0.001$	179.6005	179.600–179.601

5.3 Rear optical stack and sensor indexing

Each crystal has two SiPM coupling chains. For local sensor $\ell \in \{0, 1\}$,

$$s = 2j + \ell, \qquad x_s = x_j + (2\ell - 1)(5.25\ \mathrm{mm}), \qquad y_s = y_j. \tag{37}$$

The first equation maps crystal j to consecutive sensor identifiers; the factor $2\ell - 1$ evaluates to -1 for the left sensor and $+1$ for the right sensor. Therefore the two 6-mm-wide active areas are centered 10.5 mm apart. This indexing describes 72 simulated active layers and must not be confused with the later 44-channel independent-readout map.

The grease begins exactly at the PWO rear face in the nominal placement and the three layers meet consecutively. There is no explicit PCB, SiPM package body, cable, air gap, aluminum, Teflon, support structure, or finite-thickness ESR volume. These omissions make the geometry a controlled optical model, not a complete CAD representation.

5.4 Runtime material construction

The frozen GDML has an empty `materials` block. The C++ detector construction resolves every material reference at runtime as summarized in Table 8.

The BM material name is descriptive but its literal mass fractions are unusually hydrogen-rich and must not be silently replaced by a conventional BC408 formula. Moreover, no optical material-properties table is attached to it. The strips therefore deposit energy and scatter the beam, but do not generate or transport BC408 scintillation light; their ROOT output is deposited energy per strip. Grease and window share density and composition but have different optical constants. No temperature, radiation-damage, crystal-lot, or channel-to-channel material variation is modeled.

5.5 Transport settings that geometry alone does not encode

The logged electromagnetic table range is 100 eV to 100 TeV with seven bins per decade. Electron/-positron energy-loss fluctuations and lateral multiple-scattering displacement are enabled; the step function is $(0.2, 1\ \mathrm{mm})$, with Urban multiple scattering below 100 MeV and WentzelVI above it. The runtime log reports a $PbWO_4$ Birks coefficient of $0.0333333\ \mathrm{mm/MeV}$. No custom production cuts, regions, fields, step limiter, optical-photon stacking policy, or maximum-Cherenkov-photon setting was specified. Unlisted behavior therefore inherits Geant4 11.4.2 defaults and is version-dependent rather than fully encoded by the GDML.

Table 8: Runtime materials. "Optical MPT" indicates whether a material properties table was attached in optical mode.

Material	Density (g/cm^3)	Composition/model and use	Optical MPT
G4_Galactic	10^{-25}	NIST galactic vacuum; world	yes
G4_PbWO4	8.28	NIST $PbWO_4$; 36 crystals	yes
BeamTestBC408	1.034	custom C/H mass fractions 0.47524752/0.52475248; BM strips	no
BeamTestGrease	1.06	stoichiometric C_2H_6OSi; 72 coupling layers	yes
BeamTestWindow	1.06	stoichiometric C_2H_6OSi; 72 windows	yes
G4_Si	$\simeq 2.33$	NIST silicon; 72 active layers	yes

6 Optical production: scintillation, transport, boundaries, and PDE

Optical mode registers `G4OpticalPhysics` on top of `FTFP_BERT`. The relevant configured processes are scintillation, Cherenkov generation, bulk absorption, and optical-boundary transport. Rayleigh, Mie, and wavelength-shifting processes have no supplied material response in this model. No macro overrides the Geant4 optical defaults, so details not stated below inherit Geant4 11.4.2 behavior.

6.1 From wavelength files to Geant4 photon energy

Every optical CSV contains pairs (λ_k, v_k). The loader converts the wavelength to photon energy,

$$E_{\gamma,k} = \frac{hc}{\lambda_k}, \tag{38}$$

sorts the resulting pairs into increasing E_γ, and adds them to a Geant4 material-properties table. In Eq. (38), h is Planck's constant, c is the speed of light, and λ_k is the tabulated vacuum wavelength. The division maps short wavelength to high photon energy; sorting is necessary because Geant4 expects an increasing energy axis. It does not alter the tabulated property values or create measurements between points; interpolation supplies those intermediate values.

6.2 $PbWO_4$ scintillation

The configured total light yield is $Y = 200$ photons/MeV with Geant4 `RESOLUTIONSCALE=1.0`. At the expectation level, the light-production scale may be written

$$\mathbb{E}[N_\gamma \mid E_{\text{vis}}] \simeq Y E_{\text{vis}}, \qquad Y = 200 \text{ photons/MeV}. \tag{39}$$

Here N_γ is the number of scintillation photons and E_{vis} is the energy visible to Geant4's scintillation process after its configured quenching treatment. Multiplying photons per MeV by MeV gives photons. Equation (39) explains the mean scale only; actual photon generation occurs step by step in Geant4 and includes its statistical fluctuations and process implementation. It is not a direct formula for detected photons or ADC.

All three scintillation components use the same wavelength spectrum. Their fractions are $(a_1, a_2, a_3) = (0.73, 0.23, 0.04)$, and their decay constants are $(\tau_1, \tau_2, \tau_3) = (5, 14, 110)$ ns. With no explicit rise time, the idealized emission-delay density is

$$f_t(t) = \mathbf{1}[t \geq 0] \sum_{k=1}^{3} a_k \frac{e^{-t/\tau_k}}{\tau_k}, \qquad \sum_{k=1}^{3} a_k = 1. \tag{40}$$

Each exponential integrates to one on $t \geq 0$, and the fractions sum to one, so Eq. (40) is normalized. The mixture selects a prompt, intermediate, or slow component and then an exponential delay. It describes the configured scintillation timing distribution, not the final sensor waveform: photon flight paths, boundary interactions, PDE, microcell effects, and circuit shaping all broaden or alter later histories.

The emission spectrum has relative intensity 0.01, 0.08, 0.35, 0.85, 1.00, 0.90, 0.55, 0.18, 0.04, and 0.01 at 350, 370, 390, 410, 420, 440, 460, 500, 550, and 600 nm. "Relative" means that only the shape matters; its maximum being one is not a 100% photon-production probability.

Table 9: Optical inputs actually used by production. These hashes identify the files, including the seven-point PWO refractive-index curve.

Input	SHA-256
pwo_rindex.csv	44eb6eea38f477dbdc30dd8e013fde86fc83f029ce453c56293b09831fe6420a
pwo_absorption_mm.csv	14662563278edce10fd31cac3a32326be1ed6cbeab7e386d08dcb36bd5479dee
pwo_emission.csv	8a1c8044f769e30851f47953a4d3d29e8aab05a7c50a2974beb1e7665ea9e255
grease_rindex.csv	e61568e1a90eec1f08f5414ec88033fe820915502750b00ba3ada2c280296ffb
grease_absorption_mm.csv	4507c241eff836a908392d79681c18a5702215ba5c93fb045980ce67c4764072
esr_reflectivity.csv	09da111d237795502ea04d15bc7b6ddaab9044f9ecd87119bf12bf5ba708c53c
onsemi_microfc_60035_cseries_pde_5v.csv	0411dab984bc9eb730c532e9fc06e8eae2da07035f83bad3d7e9adb8d1d996d9

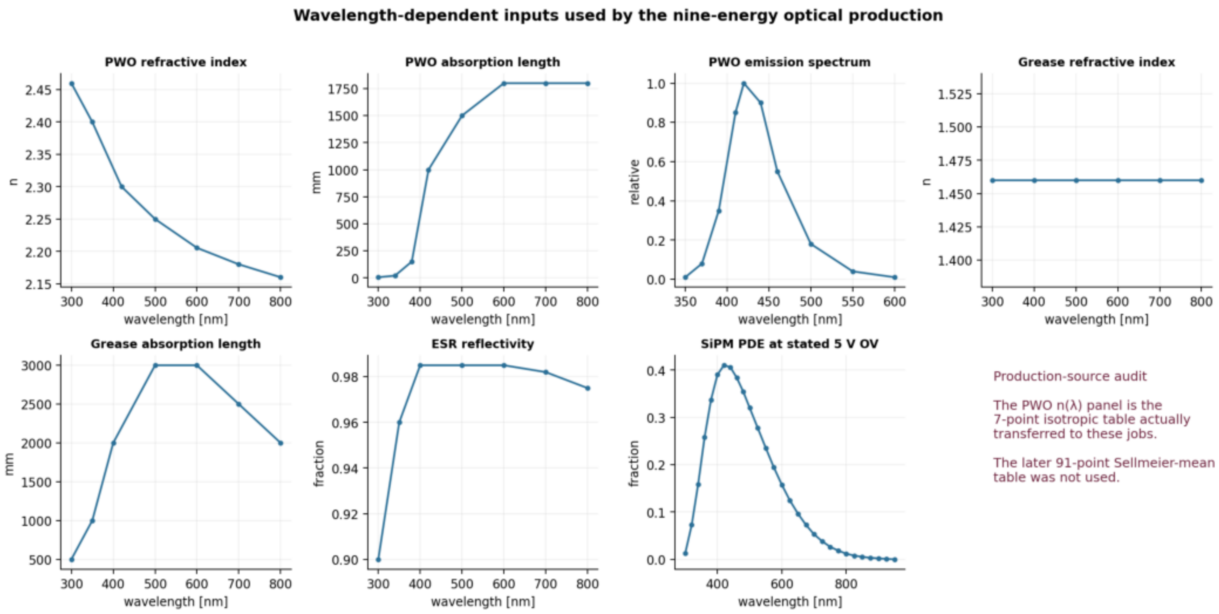

Figure 10: Wavelength-dependent refractive indices, absorption lengths, $PbWO_4$ emission, ESR reflectivity, and 5-V SiPM PDE reconstructed from the exact production CSVs. The curves are empirical/model inputs, not fitted by the ADC report.

6.3 Bulk refractive index and absorption

Production used the seven-point effective-isotropic PWO refractive index

λ (nm)	300	350	420	500	600	700	800
n_{PWO}	2.46	2.40	2.30	2.25	2.2055	2.18	2.16

with piecewise interpolation in photon energy. This is not the later 91-point Sellmeier-derived scalar mean, and it does not model birefringence or polarization-dependent ordinary and extraordinary indices.

The PWO absorption length is 5, 20, 150, 1000, 1500, 1800, 1800, and 1800 mm at 300, 340, 380, 420, 500, 600, 700, and 800 nm. In a homogeneous segment of path length ℓ, the corresponding Beer–Lambert survival law is

$$P_{\text{survive}}(\ell \mid \lambda) = \exp\left[-\frac{\ell}{L_{\text{abs}}(\lambda)}\right]. \tag{41}$$

$L_{\text{abs}}(\lambda)$ is the wavelength-dependent absorption length. The ratio ℓ/L_{abs} is dimensionless; longer paths or smaller absorption lengths reduce survival. Geant4 samples this attenuation along the actual piecewise photon trajectory. Equation (41) does not include boundary loss, PDE, or wavelength re-emission, and the input is a representative unirradiated reference curve rather than a measurement of each beam-test crystal.

The world has $n = 1$ over 1.30–6.50 eV. Grease has $n = 1.46$ from 300–800 nm and absorption lengths of 0.5–3.0 m across the tabulated range. The effective window has constant $n = 1.59$ and 10-m absorption length, while active silicon has constant $n = 3.5$ and 1-μm absorption length over 1.30–6.50 eV. These window and silicon constants are simplified effective values, not wavelength-resolved vendor measurements.

6.4 Dielectric boundaries and ESR wrapping

No explicit optical surface is assigned at the PWO/grease or grease/window boundaries. Geant4 therefore applies ordinary dielectric-boundary transport from the adjacent bulk indices at those two interfaces. The final window/active-silicon encounter is different: the shared active-silicon logical volume carries the explicit `SiPMActiveAbsorber` `dielectric_metal` skin described below, and the stepping action applies the custom PDE whenever the post-step physical volume is an active sensor. It must not be described as a third unmodified dielectric interface. For the two ordinary interfaces, the unpolarized normal-incidence reflectance reduces to

$$R_{12} = \left(\frac{n_1 - n_2}{n_1 + n_2}\right)^2. \tag{42}$$

Here n_1 and n_2 are the refractive indices on the two sides. The squared amplitude ratio gives reflected intensity and shows why index matching changes collection. Equation (42) is a pedagogical normal-incidence limit for those ordinary boundaries; Geant4 uses the appropriate angle- and polarization-dependent treatment and can produce total internal reflection. It is not the rule used to count a photon at the active-silicon skin.

The ESR is a zero-thickness `unified`, `polished`, `dielectric_metal` optical surface with zero surface efficiency and reflectivity 0.900, 0.960, 0.985, 0.985, 0.985, 0.982, and 0.975 at 300, 350, 400, 500, 600, 700, and 800 nm. For a photon incident on that surface,

$$B_R \sim \text{Bernoulli}(R_{\text{ESR}}(E_\gamma)) \tag{43}$$

summarizes the reflect-or-absorb probability controlled by the table: $B_R = 1$ denotes reflection and R_{ESR} is interpolated at the photon energy. Geant4 executes the boundary process; the equation explains the role of `REFLECTIVITY`, not all angular details of the `unified` model.

There are 36 directional border surfaces from each PWO physical volume to the world. Exposed PWO-to-world boundaries therefore use ESR. The two grease footprints are PWO-to-grease boundaries, so the PWO-to-world border is not selected and those rear apertures remain optically open. The model contains no ESR thickness, aluminum, Teflon, or air layer beneath a film.

6.5 Custom SiPM detection probability

Every active-silicon logical volume has a polished `dielectric_metal` absorbing skin with Geant4 `REFLECTIVITY`=0 and surface `EFFICIENCY`=0. Zero Geant4 efficiency is intentional: detection is performed explicitly in the stepping action. When an optical photon enters `SiPMActive_NN`, the code first increments the incident "ideal photon" count, evaluates the custom PDE, and kills the photon whether or not it passes.

The adopted 5-V curve was constructed from a digitized 2.5-V shape by

$$\text{PDE}_{5\text{V}}(\lambda) = \frac{0.41}{0.31}\,\text{PDE}_{2.5\text{V}}(\lambda), \tag{44}$$

so its 420-nm maximum equals the 41% datasheet value. The ratio rescales the curve vertically without changing its wavelength shape. It is a representative approximation, not an independent 5-V wavelength scan of the installed devices.

For incident photon q with energy $E_{\gamma q}$, detection is

$$D_q = \mathbf{1}[U_q < \text{PDE}_{5\text{V}}(E_{\gamma q})]\,, \qquad U_q \sim \text{Uniform}(0,1), \tag{45}$$

and a sensor-bin detected count is $N_{\text{PDE}} = \sum_q D_q$. Because a uniform draw satisfies $\Pr[U_q < \text{PDE}_{5\text{V}}(E_{\gamma q})] = \text{PDE}_{5\text{V}}(E_{\gamma q})$, Eq. (45) is a Bernoulli trial whose expectation is the interpolated PDE. The table is interpolated linearly in photon energy and PDE is exactly zero outside its 300–950-nm coverage. This test includes only the adopted PDE probability; transport and coupling losses have already acted before a photon reaches the active layer. It does not apply electronics threshold or SiPM saturation, which enter later.

7 SiPM microcell response and the ROOT handoff contract

The custom PDE result is still a detected-photon history. Before handing that history to the downstream circuit digitizer, the Geant4 executable adds an explicit SiPM response model. It stores both a signal-only history and an augmented history that includes prompt crosstalk, afterpulses, and dark pulses. Neither history is charge or ADC.

7.1 Frozen operating point and time grid

The 5-V setting selects the PDE file. Because matching 5-V nuisance values were not available in the adopted source, the crosstalk, afterpulse, and dark parameters retain the documented 2.5-V values. The gain is validated and printed but is not used inside Geant4 to convert cells to charge. Its first operative use is in the downstream circuit stage.

Detected-photon arrival time t is mapped to a fine bin by

$$b(t) = \left\lfloor \frac{t}{\Delta t} \right\rfloor, \qquad \Delta t = 0.25\text{ ns}, \qquad 0 \le t < T = 1000\text{ ns}. \tag{46}$$

Table 10: SiPM response parameters supplied to all nine optical jobs.

Parameter	Production value
Model label	onsemi MICROFC-60035-SMT
Active area	6×6 mm^2
Physical microcells N_{cell}	18,980 per SiPM
Stated overvoltage	5.0 V
Stated avalanche gain	6.0×10^6
Recovery constant τ_r	95 ns
Prompt-crosstalk probability p_{XT}	0.07
Afterpulse probability p_{AP}	0.002
Dark rate R_{dark}	1.2×10^6 s^{-1} per SiPM

The floor in Eq. (46) assigns each time to the bin whose left edge is $b\Delta t$. The grid has $T/\Delta t = 4000$ bins. Photons at or after 1 μs are counted in explicit late counters but do not enter the fine history. A second view sums groups of 100 fine bins into sixteen 25-ns bins over the first 400 ns; it is a coarser representation of the same response, not a second simulation.

7.2 Prompt crosstalk

Let N_{sb}^{sig} be the number of PDE-detected signal photons in sensor s, bin b. The number of prompt-crosstalk additions is sampled as

$$C_{sb} \sim \text{Poisson}\left(N_{sb}^{\text{sig}} \frac{p_{\text{XT}}}{1 - p_{\text{XT}}}\right), \qquad p_{\text{XT}} = 0.07. \tag{47}$$

The Poisson mean grows linearly with the number of signal seeds. Adding the original seeds gives expected prompt input $N_{sb}^{\text{sig}}/(1 - p_{\text{XT}})$; this is why the factor $p_{\text{XT}}/(1 - p_{\text{XT}})$ is used instead of simply p_{XT}. The code draws all additions at the signal bin. Equation (47) is a compact cluster-size approximation: added crosstalk pulses do not recursively seed another explicit generation of crosstalk.

7.3 Afterpulses

For each signal-populated bin, the number of afterpulses is

$$A_{sb} \sim \text{Poisson}\left(p_{\text{AP}} N_{sb}^{\text{sig}}\right), \qquad p_{\text{AP}} = 0.002, \tag{48}$$

and each afterpulse delay is sampled by

$$\delta t = -\tau_r \ln U, \qquad U \sim \text{Uniform}(0, 1), \qquad \tau_r = 95 \text{ ns}. \tag{49}$$

Equation (48) makes the mean afterpulse count proportional to detected signal. In Eq. (49), the inverse transform of a uniform draw produces the exponential density $e^{-\delta t/\tau_r}/\tau_r$ for positive delay. The delay is truncated to the 0.25-ns bin grid and pulses landing beyond 1 μs are not stored. Only the original PDE-detected signal seeds this calculation; crosstalk, dark pulses, and afterpulses do not recursively create further afterpulses in this implementation.

7.4 Dark pulses

For each sensor and 1-μs event window,

$$D_s \sim \text{Poisson}(R_{\text{dark}} T), \qquad R_{\text{dark}} T = (1.2 \times 10^6 \text{ s}^{-1})(10^{-6} \text{ s}) = 1.2. \tag{50}$$

Thus each SiPM receives 1.2 dark pulses on average per simulated window. Each drawn pulse is placed uniformly among the 4000 fine bins. The multiplication in Eq. (50) converts a rate to an expected count; it does not force every event to contain 1.2 pulses. No temperature dependence, channel variation, or burst-correlated dark noise is modeled.

7.5 Recovery and instantaneous saturation

Signal-only and augmented inputs are propagated through identical but separate occupancy states. Immediately before a populated bin b, the state is decayed from the preceding processed bin:

$$O_b^- = O_{b'}^+ \exp\left[-\frac{(b-b')\Delta t}{\tau_r}\right], \tag{51}$$

$$N_b^{\text{avail}} = \max(1, N_{\text{cell}} - O_b^-), \tag{52}$$

$$F_b = N_b^{\text{avail}}\left[1 - \exp\left(-\frac{N_b^{\text{in}}}{N_b^{\text{avail}}}\right)\right], \tag{53}$$

$$O_b^+ = O_b^- + F_b. \tag{54}$$

O_b^- and O_b^+ are the effective occupied-cell populations just before and after bin b; b' is the previous nonempty augmented bin; N_b^{avail} is the number of available cells; N_b^{in} is either the signal-only or augmented pulse input; and F_b is the corresponding effective newly fired count. Equation (51) implements exponential recovery: longer elapsed time reduces occupancy. The maximum in Eq. (52) protects the numerical expression from zero or negative availability.

Equation (53) is a Poissonized, large-availability mean-field approximation to distinct-cell occupancy. For an exact fixed integer input N choosing uniformly among A independent cells, the expected number of distinct cells would instead be $A\{1-[1-(1/A)]^N\}$; replacing $[1-(1/A)]^N$ by $\exp(-N/A)$ gives the implemented smooth expression. For $N_b^{\text{in}} \ll N_b^{\text{avail}}$, expanding the exponential gives $F_b \simeq N_b^{\text{in}}$; for very large input, F_b approaches N_b^{avail} and saturates. Because it is a smooth mean-field quantity rather than a draw of discrete cell identities, a stored fired-cell value need not be an integer. Finally, Eq. (54) adds newly occupied cells to the surviving occupancy for the next interval.

Because Eq. (51) permits cells to recover and fire again, the time integral

$$N_{s,[0,T)}^{\text{fired}} = \sum_{b=0}^{3999} F_{sb} \tag{55}$$

can exceed 18,980. This does not violate the instantaneous inventory: Eq. (53) limits each bin using the current availability, while Eq. (55) counts repeated firings after recovery across a microsecond. The model is mean-field rather than a discrete identity-tracking simulation of all physical microcells.

7.6 ROOT ntuple contract

Each of the nine aligned jobs writes seven ntuples with Geant4's in-job multithreaded ntuple merging. Every file contains exactly 10,000 unique events with local identifiers 0–9999 and configured offset zero. The global event key is nevertheless defined generally as

$$e = \texttt{event\_offset} + \texttt{local\_event}. \tag{56}$$

Equation (56) makes the identity independent of ROOT row order. Multithreaded output may be interleaved, so consumers must join on the stored `event` key (and sensor/crystal/bin as appropriate). The aligned sample uses the direct files and requires no cross-job merge.

Table 11: Geant4 ROOT handoff schema.

Ntuple	Expected granularity	Contract
`events`	1 row/event	Event/local identifiers; sampled kinetic energy, momentum, and gun coordinates; BM/PWO deposited-energy totals; photon and fired-cell totals; late counters; truth-plane position and direction.
`crystals`	36 rows/event	Crystal identifier and deposited energy; ideal, PDE, signal-fired, and augmented-fired responses summed over its two simulated SiPMs.
`sensors`	72 rows/event	Sensor, crystal, and local-sensor identifiers; dense per-SiPM ideal/PDE counts and 400-ns/1-μs fired-cell integrals.
`bm_strips`	sparse	One row only for a BM strip with positive deposited energy; station, plane, strip, and deposited energy.
`timing_0p25ns`	sparse aggregate	Event/bin/start time plus all-sensor PDE input, augmented input, signal-fired, and augmented-fired response through 1 μs. Rows with both aggregate inputs zero are omitted.
`gate_25ns`	16 rows/event	The same aggregate quantities rebinned into 25-ns intervals through 400 ns, including empty bins.
`sensor_timing_0p25ns`	sparse per SiPM	Event/sensor/crystal/local sensor/bin keys and per-sensor PDE input, augmented input, signal-fired, and augmented-fired response. This is the time-resolved source for downstream circuit digitization.

"Ideal photons" are optical photons entering an active-silicon volume before the custom PDE decision. "PDE detected" counts the successful Bernoulli decisions. "Signal fired" applies saturation and recovery to PDE detections alone. "Augmented input" adds crosstalk, afterpulses, and dark pulses before saturation; "augmented fired" is the result after saturation and recovery. These names identify distinct stages and are not interchangeable normalization choices.

Table 12: ROOT-to-NPZ field contract used in this report.

Container	Field	Shape	Unit	Meaning
ROOT events	`event`	10000	index	Generated event identity; required to be exactly 0 through 9999
ROOT events	`energy_MeV`	10000	MeV	Positron kinetic energy, not the nominal momentum label
ROOT events	`pwo_edep_MeV`	10000	MeV	Energy deposited in all PWO crystals
ROOT events	`ideal_photons`	10000	count	Optical photons reaching active silicon before the PDE decision; scintillation and possible Cherenkov contributions are included
ROOT events	`pde_detected`	10000	count	Photons accepted by the wavelength-dependent PDE decision
ROOT events	`fired_signal_1us`	10000	cell	Signal-fired microcells in the one-microsecond window

Container	Field	Shape	Unit	Meaning
ROOT events	fired_augmented_1us	10000	cell	Signal plus crosstalk, afterpulse, and dark contributions
ROOT sensor timing	bin, fired_augmented	sparse	0.25 ns, cell	Per-sensor impulse train used for circuit convolution
Aligned native NPZ	direct_5V_channel_adc_unclipped	10000 x 44	native count	Floating, unclipped circuit-equivalent held peak; sole scientific endpoint
Aligned native NPZ	direct_5V_channel_adc	10000 x 44	legacy count	Legacy 11k-clipped compatibility field; present but explicitly excluded from every scientific result
Aligned native NPZ	fixed_standard_normal	10000 x 44	dimensionless	Pre-drawn standard-normal variates reserved for a downstream noise realization; deliberately unused in this report

The per-event and per-crystal records obey additive closure relations such as

$$N_e^{\mathrm{PDE}} = \sum_{s=0}^{71} N_{es}^{\mathrm{PDE}}, \qquad E_e^{\mathrm{PWO}} = \sum_{j=0}^{35} E_{ej}^{\mathrm{dep}}, \tag{57}$$

up to the exact stored data types and floating summation. The first equality adds detected photons from all sensors; the second adds deposited energy from all crystals. They are integrity checks between ntuple granularities, not new physics models. Likewise, summing sparse `sensor_timing_0p25ns` fired values over bins reproduces the stored 1-μs sensor integral to the declared numerical tolerance. A missing sparse row represents zero declared input for that key, not a removed event.

The `events` fields `truth_plane_recorded` and `truth_plane_*` must be interpreted together. When no forward primary crossing is recorded, the flag is zero and the coordinates/direction remain NaN. NaN means unavailable; replacing it with zero would fabricate a central impact.

7.7 Aligned equal-10k source identity

The source population is the nine direct optical ROOT files recorded in the source manifest, one file per energy and 10,000 unique event IDs 0–9999 in each file. All nine ROOT SHA-256 sidecars verified, and the frozen schema audit passed for all seven required ntuples. No legacy or mixed-population optical file belongs to this source contract.

The regenerated Table 12 and source manifest bind each aligned optical ROOT digest to its native NPZ digest and audited event count. The canonical Condor cluster-998 audit, `native_digitizer_cluster998_audit.json`, has SHA-256 `9a4a279d3222a1c3a35c6503aef9f577573466d920bfe94cbf79498b2f44f2c9` and reports PASS for all nine $10{,}000 \times 44$ native products. The report asset audit independently rehashes the ROOT/NPZ pairs and admits only `direct_5V_channel_adc_unclipped` as the scientific endpoint.

No beam-test event selection or ADC transformation occurs in the full Geant4 ROOT files. Their exact handoff to the circuit stage is the time-resolved augmented fired-cell history plus its event and sensor identity. Charge conversion, channel mapping, circuit convolution, HOLD, and native ADC produce the separate compact NPZ endpoint; no beam-data calibration enters it. Consequently, calling any ROOT fired-cell branch *ADC* would collapse several auditable stages into an incorrect single scale factor.

8 From 72 simulated SiPMs to 44 native circuit channels

Geant4 records two SiPMs on every crystal of the 6×6 array, hence $36\times2=72$ sensor identities. The native circuit product does not digitize an arbitrary geometrical sum of those sensors. It follows the frozen 2025 readout map and retains only SiPMs that have their own, independently scaled readout. This section makes that reduction explicit before any waveform calculation is introduced.

8.1 Geant4 sensor identity

Let $r_y \in \{0,\ldots,5\}$ be the Geant4 crystal row, let $r_x \in \{0,\ldots,5\}$ be its column, and let $a \in \{0,1\}$ identify the left or right rear SiPM. The C++ production code assigns the crystal and sensor integers

$$j = 6r_y + r_x, \qquad s = 2j + a. \tag{58}$$

Here $j \in \{0,\ldots,35\}$ is the row-major $PbWO_4$ crystal number and $s \in \{0,\ldots,71\}$ is the sensor number stored in the ROOT trees. Thus integer division gives $j = \lfloor s/2 \rfloor$, while $s \bmod 2$ recovers the local SiPM side. Equation (58) is an identity assignment, not a detector-response model: it says which physical sensitive volume produced a record.

The frozen readout convention reflects the Geant4 horizontal index: $i_X = 5 - r_x$ and $i_Y = r_y$. The report target naming convention is crystal 20 / GID 34+40 / $R(2,2)$ / software C14 / GDML `PWO_15`. In the crystal-number convention, $6(5-i_Y)+i_X = 20$. The frozen map places the pair at $R(2,2)$. If the software row-major label is $C = 6i_Y + i_X$, this coordinate is C14; it maps to Geant4 $j = 15$, sensors 30 and 31, and physical name `PWO_15`. Thus "crystal 20," "software C14," and "GDML `PWO_15`" are different numbering conventions for the GID 34+40 sensor pair, not different beam targets. The mapped hardware 5×5 footprint is therefore Geant4 columns $r_x = 1,\ldots,5$ and rows $r_y = 0,\ldots,4$. Readout i_Y is not reflected: it increases with positive GDML y; only i_X is reflected.

The sparse ROOT tree `sensor_timing_0p25ns` stores only occupied event–sensor–time bins. For a requested sensor set $\mathcal{S}$ and a 1 µs window divided into $N_t = 4000$ bins, the dense input reconstructed by the digitizer is

$$x_{es}[n] = \sum_{\ell} f_\ell^{\text{aug}}\, \mathbf{1}(e_\ell = e)\, \mathbf{1}(s_\ell = s)\, \mathbf{1}(b_\ell = n), \qquad \Delta t = 0.25\,\text{ns}. \tag{59}$$

In this expression, e labels an event, s labels a sensor, n labels a time bin, and row ℓ of the sparse tree supplies event e_ℓ, sensor s_ℓ, bin b_ℓ, and the effective augmented fired-cell amount f_ℓ^{aug}. The indicator $\mathbf{1}(\cdot)$ equals one when its condition is true and zero otherwise. The sum therefore places every sparse entry into exactly one dense sensor time series. After the mapping below, x_{ec} denotes this same series with native channel c replacing its unique sensor identity s; no second physical response is introduced.

As a lossless consistency test, summing the reconstructed bins must reproduce the per-event sensor summary:

$$F_{es}^{\text{aug}} = \sum_{n=0}^{N_t-1} x_{es}[n], \qquad \max_{e,s} \left| F_{es}^{\text{aug}} - F_{es}^{\text{aug,ROOT}} \right| \le 2\times10^{-3}. \tag{60}$$

The left side is the reconstructed one-microsecond effective fired-cell total; the reference on the right is the value in the ROOT `sensors` tree. The tolerance allows floating-point accumulation of fractional occupancy-corrected cells. Failure of Eq. (60) means that a timing row, event offset, or channel identity was lost, and the digitization must stop rather than continue with a distorted waveform.

8.2 Independent and combined readouts

Each row u of the frozen readout map contains a list $\mathcal{S}_u$ of one or more simulated sensor identities. The map classifies a row by

$$I_u^{\text{ind}} = \mathbf{1}(|\mathcal{S}_u| = 1), \qquad \mathcal{C}_{44} = \{u : I_u^{\text{ind}} = 1\}. \tag{61}$$

The cardinality $|\mathcal{S}_u|$ is the number of simulated SiPMs represented by row u. A row is independent only when that number is exactly one, and $\mathcal{C}_{44}$ is the resulting set of native channels. The audited map contains 47 rows: 44 singleton rows and three rows that each combine two SiPMs. Consequently $|\mathcal{C}_{44}| = 44$. This is a hard identity rule; it is not an ADC threshold or an event selection.

The 44 singleton rows form 22 crystal positions. If $p = (i_X, i_Y)$ denotes a 2025 crystal identity, define

$$\mathcal{C}(p) = \{c \in \mathcal{C}_{44} : (i_X(c), i_Y(c)) = p\}, \qquad \mathcal{P}_{22} = \{p : |\mathcal{C}(p)| = 2\}. \tag{62}$$

The set $\mathcal{C}(p)$ contains the independently stored channels at crystal p, and $\mathcal{P}_{22}$ retains only positions with exactly two of them. The source audit finds $|\mathcal{P}_{22}| = 22$ and $|\mathcal{C}(p)| = 2$ for every retained position. This equality documents the topology only: the report does *not* require either SiPM to fire and does not reject an event when one or both responses are zero.

For reproducible array columns, the singleton rows are sorted lexicographically by (i_Y, i_X, GID). If $\pi(c) \in \{0, \ldots, 43\}$ is that sorted position, the native matrix is defined by

$$S_{e,\pi(c)} = S_{ec}, \qquad c \in \mathcal{C}_{44}. \tag{63}$$

Equation (63) fixes which physical channel occupies each NPZ column. It performs no numerical transformation of the response S_{ec}, which will be defined in Section 11. Sorting by sensor number alone would not reproduce this contract.

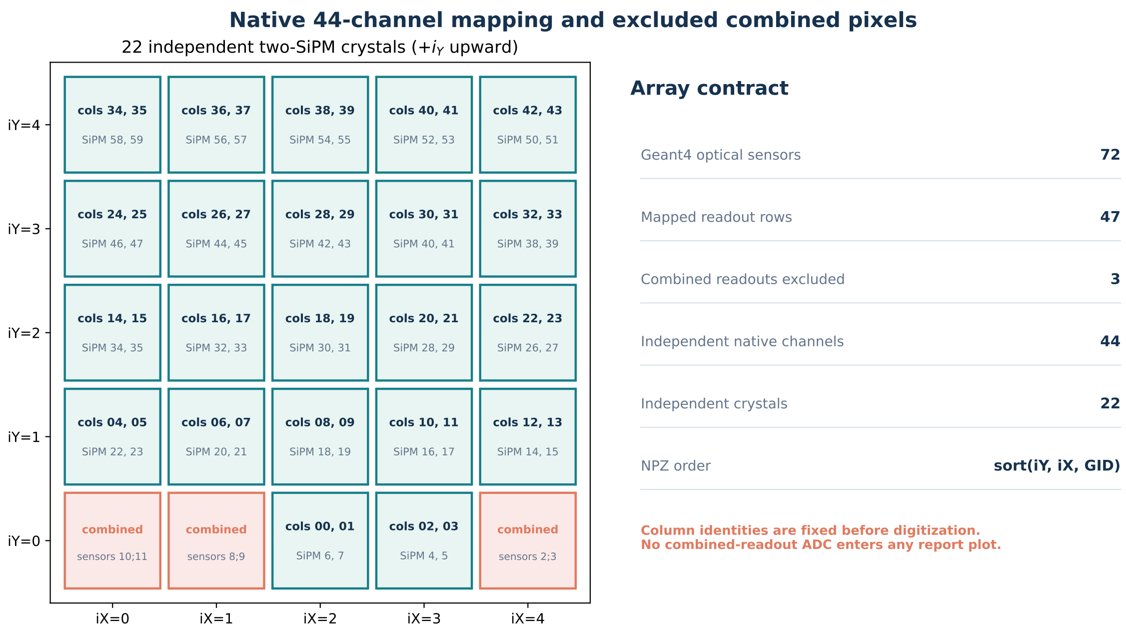

Figure 11: Frozen mapping from the simulated 72-SiPM array to the 44 native independent circuit channels. The three combined-readout crystals are identified but excluded; they are not rescaled or split into invented independent channels. The grid now follows the physical Cartesian front-face orientation, with $+i_Y$ upward; this avoids a row-0-at-top table convention that can make i_Y appear reflected.

Table 13: Ordered identity of all 44 independent NPZ columns. Three combined-readout rows are excluded before this ordering is formed.

NPZ column	iX	iY	Geant4 sensor	map readout	hardware GID
0	2	0	6	4	32
1	2	0	7	5	38
2	3	0	4	2	25
3	3	0	5	3	26
4	0	1	22	18	54
5	0	1	23	19	59
6	1	1	20	16	44
7	1	1	21	17	49
8	2	1	18	14	33
9	2	1	19	15	39
10	3	1	16	12	24
11	3	1	17	13	27
12	4	1	14	10	13
13	4	1	15	11	15
14	0	2	34	30	55
15	0	2	35	31	60
16	1	2	32	28	45
17	1	2	33	29	50
18	2	2	30	26	34
19	2	2	31	27	40
20	3	2	28	24	23
21	3	2	29	25	28
22	4	2	26	22	12
23	4	2	27	23	19
24	0	3	46	42	56
25	0	3	47	43	61
26	1	3	44	40	46
27	1	3	45	41	51
28	2	3	42	38	35
29	2	3	43	39	41
30	3	3	40	36	22
31	3	3	41	37	29
32	4	3	38	34	11
33	4	3	39	35	18
34	0	4	58	54	57
35	0	4	59	55	62
36	1	4	56	52	47
37	1	4	57	53	52
38	2	4	54	50	36
39	2	4	55	51	42
40	3	4	52	48	21
41	3	4	53	49	30
42	4	4	50	46	10
43	4	4	51	47	17

The complete table is the authority for the nontrivial geometrical orientation and GID association. Figure 11 is a visual summary of that table. In particular, the 44-channel product covers 22 crystals in the mapped 5×5 region; it is neither a 72-channel full-array digitization nor a 47-column emulation of the mixed hardware readouts.

8.3 The channel boundary used below

Every equation in Sections 9–11 is evaluated independently for each $c \in \mathcal{C}_{44}$. The exact same single-SiPM circuit kernel is used for all 44 columns. There is no channel merging before the endpoint and no crystal-level acceptance condition after it. Crystal pairs remain available as metadata, but the scientific endpoint is the per-channel matrix S_{ec}.

9 Exact three-node input-network model

Each retained SiPM is propagated through the same linear equivalent of the MicroFC–60035 terminal, the GTM input network, and the CITIROC capacitive input load. "Exact" here means that the differential equations of this declared three-node linear circuit are solved through their matrix exponential; it does not mean that unpublished transistor-level ASIC behavior is simulated.

9.1 Topology and constants

The model uses one SiPM per native channel. Its terminal capacitance is $C_S = 3.400\,\text{nF}$ and its intrinsic recovery time is $\tau_q = 95.0\,\text{ns}$. The equivalent recovery resistance is therefore

$$R_q = \frac{\tau_q}{C_S} = \frac{95.0\,\text{ns}}{3.400\,\text{nF}} = 27.941{,}176\,\Omega. \tag{64}$$

The quotient has units of ohms because an RC time constant obeys $\tau = RC$. In this equivalent, R_q summarizes the aggregate recovery of the complete SiPM terminal; it is not the resistance of one microscopic quench resistor. Because combined readouts were removed in Section 8, no parallel-SiPM multiplier is applied to C_S or R_q.

The high- and low-gain CITIROC inputs load the same board node. Their capacitances add in parallel:

$$C_I = C_{\text{HG}} + C_{\text{LG}} = 15.0\,\text{pF} + 1.5\,\text{pF} = 16.5\,\text{pF}. \tag{65}$$

Here C_I is the total small-signal capacitive load seen at the ASIC input. Although both input capacitances affect the waveform, only the low-gain voltage path is read out later. Adding the capacitances is appropriate because both terminate at the same AC reference in this equivalent circuit.

The board branch consists of a termination resistor $R_T = 51.1\,\Omega$ from the input node to a third node, followed by $C_B = 220\,\text{nF}$ from that third node to the HV/AC ground. Thus R_T and C_B form a *shunt branch*; they are not inserted in series between the SiPM and CITIROC input. This topology distinction materially changes the prompt amplitude and long time constant.

9.2 Nodal differential equation

Let $\boldsymbol{v}(t) = (v_a(t), v_i(t), v_b(t))^{\mathsf{T}}$ contain the avalanche node, CITIROC input node, and block-capacitor node voltages. Voltages in this report use the physical readout polarity: the MicroFC standard output is its anode, the GTM/CITIROC input receives a positive-going signal for positive SiPM charge, and the common SiPM cathode is biased positively. After the prompt charge has established the initial condition, Kirchhoff's current law gives

$$\mathbf{C}\,\frac{d\boldsymbol{v}}{dt} + \mathbf{G}\,\boldsymbol{v} = \mathbf{0}, \tag{66}$$

where **C** is the capacitance matrix and **G** is the conductance matrix. The first term contains currents proportional to changing voltage; the second contains resistive currents; their sum must vanish at every node when no later external current source is present. The zero on the right is therefore a statement about the post-injection evolution, not an assertion that the injected avalanche charge was zero.

For the declared topology the matrices are

$$\mathbf{C} = \begin{pmatrix} C_S & -C_S & 0 \\ -C_S & C_S + C_I & 0 \\ 0 & 0 & C_B \end{pmatrix}, \qquad \mathbf{G} = \begin{pmatrix} R_q^{-1} & 0 & 0 \\ 0 & R_T^{-1} & -R_T^{-1} \\ 0 & -R_T^{-1} & R_T^{-1} \end{pmatrix}. \tag{67}$$

The $\pm C_S$ entries express a capacitor between v_a and v_i, so its current depends on the derivative of their voltage difference. The C_I entry loads v_i to AC ground, and C_B loads v_b. In **G**, R_q^{-1} drains the avalanche node, while the symmetric $\pm R_T^{-1}$ block describes current through the resistor between v_i and v_b. This term-by-term interpretation is also a topology audit: moving R_T into the first row would describe a different physical network.

A prompt avalanche-charge magnitude $q_0 > 0$ is represented by an oriented initial voltage across the sensor capacitance,

$$\boldsymbol{v}(0^+) = \begin{pmatrix} -q_0/C_S & 0 & 0 \end{pmatrix}^\mathsf{T}. \tag{68}$$

The notation 0^+ means immediately after the idealized instantaneous avalanche. Dividing charge by capacitance gives voltage; the minus sign fixes the capacitor orientation so that the ensuing anode/input-node pulse v_i is positive. The first entry therefore places the magnitude q_0 on C_S while the two external nodes initially remain at their reference voltage. The unit kernel sets $q_0 = 1\,\mathrm{pC}$; this is only a convenient normalization and is not yet the charge assigned to one fired microcell. Reversing every voltage sign would be a mathematically equivalent global coordinate choice, but it would not describe the literal board polarity used here.

Because **C** is invertible, Eq. (66) has the state-transition solution

$$\boldsymbol{v}(t) = \exp\left(-\mathbf{C}^{-1}\mathbf{G}\, t\right) \boldsymbol{v}(0^+). \tag{69}$$

The matrix exponential evolves all three coupled voltages simultaneously. It includes every mode of the specified linear network, rather than replacing the network with a hand-chosen single exponential. Numerical eigenvectors are used to evaluate this expression on the 0.25 ns sample grid.

If λ_m is a negative real eigenvalue of $-\mathbf{C}^{-1}\mathbf{G}$, its decaying time constant is

$$\tau_m = -\frac{1}{\mathrm{Re}\,\lambda_m}. \tag{70}$$

A negative eigenvalue produces $\exp(\lambda_m t)$, so the minus sign in Eq. (70) makes τ_m positive. The numerical circuit has two finite decaying modes, $\tau_{\text{fast}} = 0.297{,}448\,\mathrm{ns}$ and $\tau_{\text{slow}} = 265.170{,}235\,\mathrm{ns}$, plus the DC mode expected from the AC-coupled topology. The later CR stage rejects the DC component.

9.3 Low-gain preamplifier coordinate

The production model represents the CITIROC low-gain branch with the inverting voltage gain $g_{\text{LG}} = 4.0$. This is the declared campaign configuration recorded by the “x4” run tag and is kept fixed; it is not a fit degree of freedom. The optical ROOT files carry detector observables rather than a repeated copy of the electronics slow-control record. The datasheet relation is $g_{\text{LG}} = C_{\text{LG}}/C_{f,\text{LG}}$; with its literal $C_{\text{LG}} = 1.5\,\mathrm{pF}$ input capacitance, the assumed factor four corresponds to the derived feedback capacitance $C_{f,\text{LG}} = 375\,\mathrm{fF}$. The modeled output before slow shaping is

$$v_{\text{pre}}(t) = -g_{\text{LG}}\, v_i(t) = -4 v_i(t). \tag{71}$$

The positive-going input-node voltage v_i is the second component of Eq. (69); the factor four is the selected low-gain amplification in the declared configuration, and the minus sign denotes inversion. No per-channel or beam-derived scale is introduced here.

The numerical unit-charge audit makes the meaning of this factor explicit: for a 1 pC initial charge, max $v_i = 0.188{,}345{,}492$ mV and min $v_{\text{pre}} = -0.753{,}381{,}968$ mV, so the ratio of their magnitudes is exactly four. Thus $g_{\text{LG}} = 4$ is the voltage gain applied to the already solved input node v_i. It is not the shortcut $q/C_{f,\text{LG}}$ and not an overall charge-to-ADC gain: the SiPM capacitance, shunt branch, parallel input load, slow shaper, peak-memory path, and ADC-input conversion remain separate stages.

The signed analogue sequence is therefore $v_i > 0 \longrightarrow v_{\text{pre}} = -4v_i < 0 \longrightarrow v_{\text{shape}} > 0$ at the principal slow-shaper lobe. The second inversion is the minus sign in the slow-shaper transfer function derived in Section 10. This polarity is essential because the CITIROC peak detector stores the highest positive slow-shaper value; it is not an absolute-value device.

10 Slow shaping, convolution, trigger, and HOLD

The nodal solution supplies the low-gain preamplifier waveform produced by a prompt unit charge. The digitizer next applies the declared CITIROC-equivalent slow shaper, convolves that response with every fired-cell time series, and closes an ideal peak detector at a fixed delay from an explicit event-level trigger reference. All operations remain linear until the positive-maximum search and HOLD closure.

10.1 CR–RC^2-equivalent slow shaper

For the selected shaper control and common RC time constant $\tau_s = 62.5$ ns, the continuous transfer function is

$$H_{\text{p15}}(s) = -\frac{R_f C_1 s}{(1 + sR_0C_0)(1 + sR_1C_1)(1 + sR_fC_f)} = -\frac{10 s\tau_s}{(1 + s\tau_s)^3}. \tag{72}$$

The first line writes the literal component values shown in the CITIROC1A datasheet page-15 slow-shaper schematic:

$$(R_0, C_0) = (25\,\text{k}\Omega, 2.5\,\text{pF}), \qquad (R_1, C_1) = (10\,\text{k}\Omega, 6.25\,\text{pF}), \qquad (R_f, C_f) = (100\,\text{k}\Omega, 625\,\text{fF}). \tag{73}$$

All three products are exactly $R_0C_0 = R_1C_1 = R_fC_f = 62.5\,\text{ns} = \tau_s$, while $R_fC_1/\tau_s = 10$. Substitution therefore gives the second line of Eq. (72); the factor ten is already the dimensionless component-ratio coefficient implied by those page-15 literals and has not been omitted. It is not a flat gain: the response vanishes at DC and at infinite frequency, its AC magnitude peaks at 3.849, and its transient amplitude ratio depends on the input waveform. Here s is the Laplace-frequency variable. The numerator $s\tau_s$ is the CR high-pass zero: it makes $H_{\text{p15}}(0) = 0$ and therefore rejects a constant DC level. The three equal denominator factors give poles at $s = -1/\tau_s$, which determine the pulse width, and the minus sign gives the inversion convention. This is a circuit response, not an energy-calibration curve.

The notation H_{p15} identifies the literal page-15 component response; it does *not* denote a unit-peak normalization. With the physical polarity fixed in Section 9, v_i is positive, the low-gain preamplifier output is negative, and this second inversion makes the principal shaped lobe positive.

This non-unit-normalized interpretation is independently consistent with the high-level CITIROC model of Marano et al., which derives $H(s) = -(R_3/R_2)s\tau/(1+s\tau)^3$ and obtains a frequency-dependent AC maximum near 3.85 for its factor-ten component ratio.[1] That work is an earlier ideal-op-amp, high-level PSPICE model. It corroborates the component-derived, frequency-dependent shaper normalization but is neither a transistor-level representation nor evidence for the exact register state of the present CITIROC1A acquisition.

10.2 Shaper normalization and coordinate freedom

The factor ten in Eq. (72) is an amplitude normalization fixed by the page-15 components, not a unit-peak convention. To see the distinction analytically, drive the transfer function with a unit step. Multiplication by $1/s$ and inverse Laplace transformation give

$$y_{\text{step}}(t) = -5\left(\frac{t}{\tau_s}\right)^2 e^{-t/\tau_s}, \qquad \max_{t\geq 0}|y_{\text{step}}(t)| = 20e^{-2} = 2.706706 \quad \text{at } t = 2\tau_s. \tag{74}$$

The maximum follows by differentiating u^2e^{-u} with $u = t/\tau_s$: its nonzero stationary point is $u = 2$. Thus a unit-peak coordinate could instead define

$$\begin{aligned} H_{\text{norm}}(s) &= \frac{H_{\text{p15}}(s)}{A}, \qquad\qquad A = 20e^{-2} = 2.706706, \\ G_{\text{norm}} &= AG_{\text{abs}} = (2.706706)(0.387995) = 1.05019. \end{aligned} \tag{75}$$

Here H_{p15} is the literal page-15 response used to produce the native simulation coordinate, and $G_{\text{abs}} = 0.387995$ is the fitted empirical data-to-simulation ADC ratio attached to that coordinate. In shorthand, $\text{ADC}_{\text{data}} \simeq G_{\text{abs}}\text{ADC}_{\text{sim,abs}}$, with channel deviations and the pedestal/noise terms included in the full calibration model. The fixed low-gain $\times 4$ response and the complete page-15 shaper are already inside $\text{ADC}_{\text{sim,abs}}$; therefore G_{abs} is not a circuit gain and is not applied to the native arrays. The "abs" suffix on these pre-existing calibration-coordinate names means the literal, non-unit-normalized amplitude convention; it is not an absolute-value operation on the waveform. Neither the CITIROC1A datasheet nor the Marano high-level model prescribes division by A; H_{norm} is optional bookkeeping, while H_{p15} remains the production response.

The signal prediction is unchanged because $G_{\text{norm}}S_{\text{norm}} = G_{\text{abs}}S_{\text{abs}}$ when $S_{\text{norm}} = S_{\text{abs}}/A$. For the pre-clipping response inside the high-CDF forward model $D_c = G_{c,\text{abs}}(S_{\text{abs}} + P_{c,\text{abs}} + \sigma_{P,\text{abs}}z)$, exact invariance requires

$$\begin{aligned} S_{\text{norm}} &= S_{\text{abs}}/A, & G_{c,\text{norm}} &= AG_{c,\text{abs}}, \\ P_{c,\text{norm}} &= P_{c,\text{abs}}/A, & \sigma_{P,\text{norm}} &= \sigma_{P,\text{abs}}/A, \\ T_{c,\text{norm}} &= T_{c,\text{abs}}, & \tau_{T,\text{norm}} &= \tau_{T,\text{abs}}. \end{aligned} \tag{76}$$

The data-ADC threshold T_c, its softness τ_T, and the log-scale deviations about the global ratio stay unchanged. Substitution gives $G_{c,\text{norm}}(S_{\text{norm}} + P_{c,\text{norm}} + \sigma_{P,\text{norm}}z) = G_{c,\text{abs}}(S_{\text{abs}} + P_{c,\text{abs}} + \sigma_{P,\text{abs}}z)$ exactly. The later 11,000-count ceiling also stays in the same output coordinate because its pre-clipping input is invariant. This is only a reparameterization of the amplitude coordinate; it neither repairs nor proves any physical gain and does not refit the data. It rescales the intermediate waveform coordinate but changes neither the waveform shape nor the calibrated output prediction.

The step factor is also not the gain of the actual input pulse. For the modeled 1 pC initial charge, the negative low-gain preamplifier waveform has magnitude 0.753,381,968 mV, while the complete slow-shaper kernel has a positive maximum of 1.626,387,557 mV at 105.5 ns. The ratio 2.158782 divides those two separately attained magnitudes; it is not a same-time voltage-transfer ratio. The

[1] D. Marano et al., "CITIROC High-Level Analog Front-End Model Implementation and Simulations," 2014, online paper.

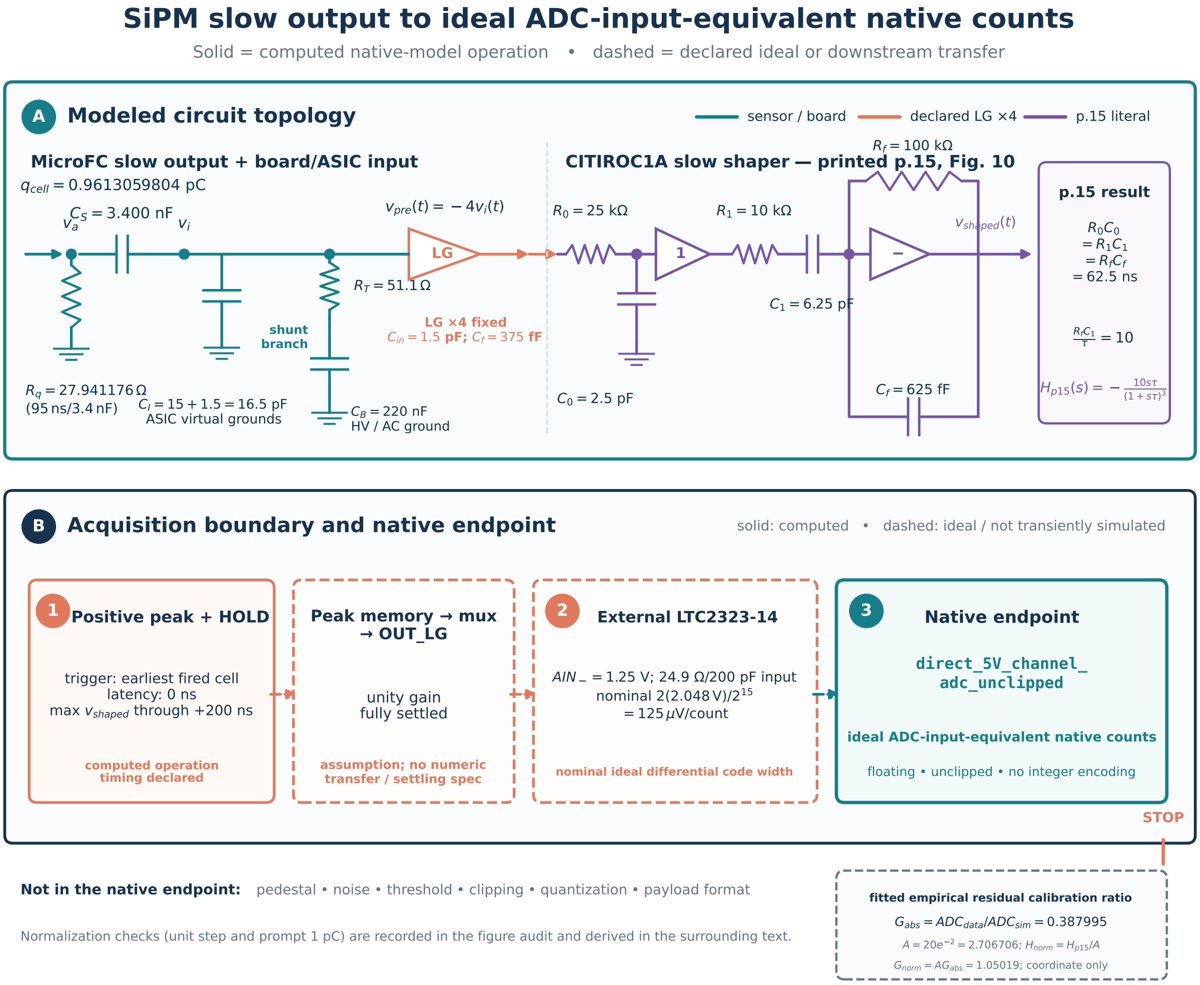

Figure 12: End-to-end circuit-equivalent path used by the native digitizer. The schematic separates three kinds of information: sensor/board values, literal CITIROC1A datasheet components, and declared modeling assumptions. Low-gain ×4 is the fixed declared run configuration; the first-hit trigger with a 200-ns ideal positive-peak HOLD and unity settled transfer through the peak-memory/multiplexer path are modeling assumptions. The page-15 slow-shaper values already generate the factor ten in Eq. (72). The last box is the floating, unclipped native endpoint, before pedestal, a fitted data-to-simulation residual scale, threshold, clipping, rounding, or selection.

input-network waveform is not a step, so its spectrum weights the shaper differently. This pulse-specific check is the relevant end-to-end response, whereas 2.706706 describes only the unit-step normalization. A known injected-charge measurement, with the actual register state and peak-path settling recorded, is needed to separate the physical gain factors.

Four distinct meanings of time and "peak." The datasheet's 62.5 ns setting is the common value of the three RC products and the control label; it is not the observed waveform peak time. The literal transfer function's unit-step response peaks at $2\tau_s = 125$ ns. The complete modeled 1 pC input-network plus shaper kernel has a different input shape and peaks at 105.5 ns. Finally, an event's ideal HOLD value is the largest positive shaped value found from its first-hit trigger through trigger plus 200 ns; it is a searched maximum inside that event window, not a sample forced to any of the three times above.

10.3 Sampled slow shaper and unit-charge kernel

The sampled implementation uses the zero-order-hold recursion for one RC pole. If $x[n]$ is its input and $u[n]$ its output, define

$$u[n] = (1-\alpha)x[n] + \alpha u[n-1], \qquad \alpha = \exp\left(-\frac{\Delta t}{\tau_s}\right), \tag{77}$$

with a zero initial state and $\Delta t = 0.25$ ns. The coefficient α is the fraction of the previous state that survives one sample; the factor $1-\alpha$ supplies the new piecewise-constant input. Because the exponential is evaluated analytically, this is the exact endpoint update of one isolated RC pole for a piecewise-constant input, rather than a forward- Euler update.

For zero initial state, the corresponding discrete transfer function is

$$L_d(z) = \frac{U(z)}{X(z)} = \frac{1-\alpha}{1-\alpha z^{-1}}. \tag{78}$$

Here $X(z)$ and $U(z)$ are the z-transforms of the sampled input and output, and z is the discrete-frequency variable. Equation (78) is the transfer actually applied at finite Δt; it should not be replaced algebraically by the continuous function $1/(1+s\tau_s)$. Denote the input-to-output operation in Eq. (77) by $\mathcal{L}_{\tau_s}$.

Starting with $u_0[n] = v_{\text{pre}}(n\Delta t)$, the code forms

$$u_1 = \mathcal{L}_{\tau_s}\{u_0\}, \qquad u_2 = \mathcal{L}_{\tau_s}\{u_1\}, \qquad u_3 = \mathcal{L}_{\tau_s}\{u_2\}, \qquad y[n] = -10\big(u_2[n] - u_3[n]\big). \tag{79}$$

The first three relations apply the same sampled RC recursion successively. Their finite-step transfer is

$$H_d(z) = -10\left[L_d(z)^2 - L_d(z)^3\right]. \tag{80}$$

$H_d(z)$ is the exact transfer of the implemented discrete cascade. It is the sampled construction associated with the analogue target Eq. (72), but the two are not algebraically identical at finite Δt. Thus u_1 is needed internally to build the later pole states even though the reported output is the difference between u_2 and u_3.

For the unit initial charge $q_0 = 1$ pC in Eq. (68), define the complete sampled kernel

$$k[n] = \frac{y(n\Delta t \mid q_0 = 1\,\text{pC})}{1\,\text{pC}}, \tag{81}$$

so $k[n]$ has units V/pC. It includes the three-node input response, the low-gain factor, and the slow shaper. In the physical polarity convention, its numerically audited positive maximum is $1.626387557 \times 10^{-3}$ V/pC at 105.5 ns. That maximum is a prediction of the frozen equivalent circuit; it has not been aligned to beam data.

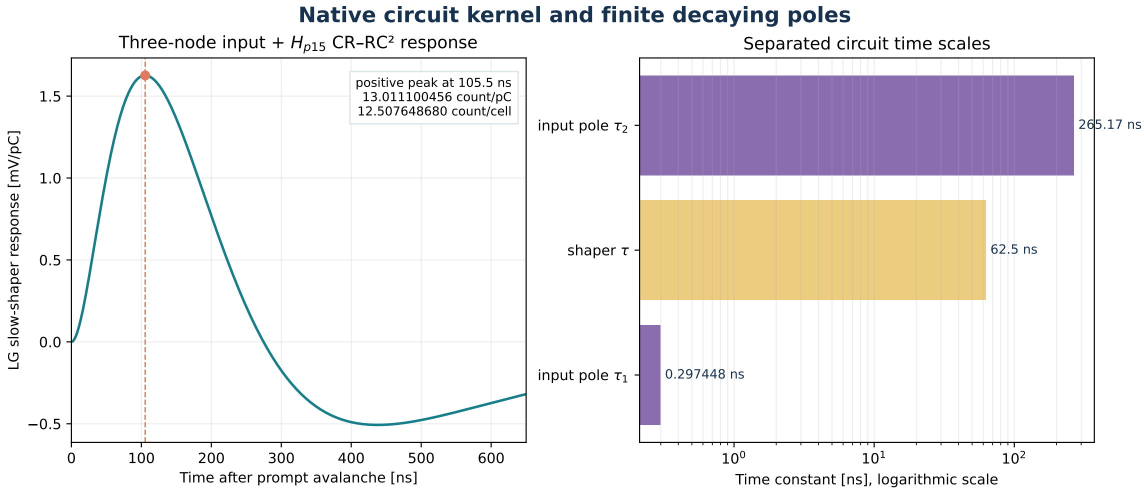

Figure 13: Numerical modes of the exact nodal circuit and the resulting unit-charge low-gain slow-shaper kernel. The plotted waveform is computed from the stated matrices and filter recursion, rather than fitted to an ADC spectrum.

10.4 Finite grid and linear convolution

The ROOT timing input spans $T_{\text{in}} = 1{,}000\,\text{ns}$. To retain the response of late cells, the waveform is padded by

$$T_{\text{pad}} = \max(4\tau_s,\, 2.5\tau_{\text{slow}})\,, \qquad N_{\text{out}} = N_t + \left\lceil \frac{T_{\text{pad}}}{\Delta t} \right\rceil = 6652. \tag{82}$$

The first term retains four shaper time constants and the second retains two and a half time constants of the slower 265.170,235 ns network decay; the larger is used. With $N_t = 4000$ and 0.25 ns bins this adds 2652 samples. Padding changes how much of the causal tail is evaluated, not the one-microsecond fired-cell content against which Eq. (60) is checked.

For event e and native channel c, let $x_{ec}[m]$ be the augmented effective fired-cell amount in input bin m. The unit-charge-normalized shaped waveform is the causal convolution

$$\widetilde{v}_{ec}[n] = \sum_{m=0}^{N_t-1} x_{ec}[m]\, k[n-m], \qquad k[r] = 0 \text{ for } r < 0. \tag{83}$$

Each fired-cell contribution places a shifted copy of the same circuit kernel at its arrival time; linearity allows those copies to be summed. The tilde emphasizes that physical fired-cell charge has not yet been multiplied in: $\widetilde{v}$ carries "kernel voltage per pC times cell" units. The production implementation evaluates the same linear convolution with zero-padded real FFTs and keeps its first N_{out} causal samples. FFT evaluation changes computational cost, not Eq. (83).

Plain-language reading of the convolution. For one event and one channel, write $N[n] \equiv x_{ec}[n]$. Then $N[n]$ is the effective fired-cell count versus time, and $k[n]$ is the voltage response per unit charge to one ideal pulse after it passes through the input network and slow shaper. The product $N[m]k[n-m]$ is one shifted copy of that pulse response, weighted by how many effective cells arrived in bin m; the sum in Eq. (83) adds all such copies to make the channel waveform. Multiplication by q_{cell} assigns physical charge to each effective cell. The model then follows the signed physical waveform from the event trigger through the declared 200 ns HOLD window, retains its largest positive value, and divides the resulting physical voltage by $L_{\text{ADC}} = 125\,\mu\text{V/count}$ to obtain the native count coordinate, idealized as an ADC-input-equivalent response, described in Section 11.

This is a linear time-invariant (LTI), circuit-equivalent model: every pulse uses the same kernel, and overlapping responses add linearly. It is not a transistor-level SPICE simulation. Pedestal, electronic noise, discriminator threshold, clipping, integer rounding, and event selection are downstream of this endpoint and are deliberately excluded.

10.5 Trigger reference and ideal HOLD

The archived optical ROOT does not contain a measured fast-shaper crossing or a hardware HOLD edge. The model therefore declares its timing reference explicitly. For an event with at least one augmented fired cell among the 44 channels, the trigger bin is

$$n_e^{\text{trig}} = \min\left\{ n : \sum_{c \in \mathcal{C}_{44}} x_{ec}[n] > 0 \right\}. \tag{84}$$

The set asks for the earliest bin in which any retained channel has positive augmented fired-cell content. It is one common event-level reference, not a separate trigger for every channel. Signal, crosstalk, afterpulse, or dark activity can provide this first entry because all are already represented in x_{ec}.

With declared trigger latency $t_{\text{lat}} = 0$ and HOLD delay $t_H = 200\,\text{ns}$, the closing time and sample are

$$t_e^{\text{hold}} = n_e^{\text{trig}} \Delta t + t_{\text{lat}} + t_H, \qquad n_e^{\text{hold}} = \left\lfloor \frac{t_e^{\text{hold}}}{\Delta t} \right\rfloor. \tag{85}$$

The first relation starts the fixed delay at the trigger reference; the floor selects the latest sampled instant not later than the closing time. The index is clipped to the evaluated waveform grid as a numerical guard. The 200 ns value is an explicit modeling assumption, not an event-by-event quantity recovered from the ROOT file.

The ideal peak detector returns the largest positive shaped value seen up to and including that sample:

$$A_{ec}^{\text{hold}} = \max_{0 \le n \le n_e^{\text{hold}}} \widetilde{v}_{ec}[n]. \tag{86}$$

The upper limit implements the HOLD closure. “Ideal” means that the model remembers the exact running positive maximum with no droop, aperture jitter, pedestal noise, or comparator error. An inactive channel is assigned zero.

This operation follows the documented analogue polarity and peak-memory semantics. The CITIROC slow-control choice selects either a positive-maximum peak detector or a signed track-and-hold path. It provides no minimum detector, absolute-value rectifier, or programmable analogue polarity flip. The separately documented *Trigger Polarity* bit changes the digital trigger-output edge convention; it does not invert the low-gain analogue waveform and cannot turn the peak detector into an absolute-value device.

For diagnostics, the open-window peak and HOLD retention are

$$A_{ec}^{\text{true}} = \max_{0 \le n < N_{\text{out}}} \widetilde{v}_{ec}[n], \qquad R_{ec}^{\text{hold}} = \frac{A_{ec}^{\text{hold}}}{A_{ec}^{\text{true}}} \quad (A_{ec}^{\text{true}} > 0). \tag{87}$$

A^{true} is the largest positive value available if peak sensing never closed, and R^{hold} is the fraction retained by the declared HOLD. This ratio diagnoses premature closure but does not rescale the stored endpoint to recover the missed amplitude.

Finally, let

$$K = \max_{0 \le n < N_{\text{out}}} k[n], \qquad P_{ec} = \frac{A_{ec}^{\text{hold}}}{K}. \tag{88}$$

K is the peak produced by one prompt unit-charge kernel and P_{ec} is the held response in prompt-peak-equivalent fired-cell units. If M equal cells arrived simultaneously and HOLD occurred after their peak, linearity would give $P_{ec} = M$. Dispersed photon timing, afterpulses, the long input tail, and HOLD can make P_{ec} differ from the arithmetic one-microsecond fired- cell total. No empirical correction is applied to force them to agree.

The representative event is selected deterministically from the 98-MeV/c generated IDs 0–9999. The eligible sample requires a recorded truth-plane crossing (`truth_plane_recorded=1`) and a finite $S_e^{(44)} > 0$. Within that finite, positive sample, the event sum closest to the sample median is chosen, with the lowest event ID breaking an equal-distance tie; the displayed channel is then the one with the largest native S_{ec} in that event.

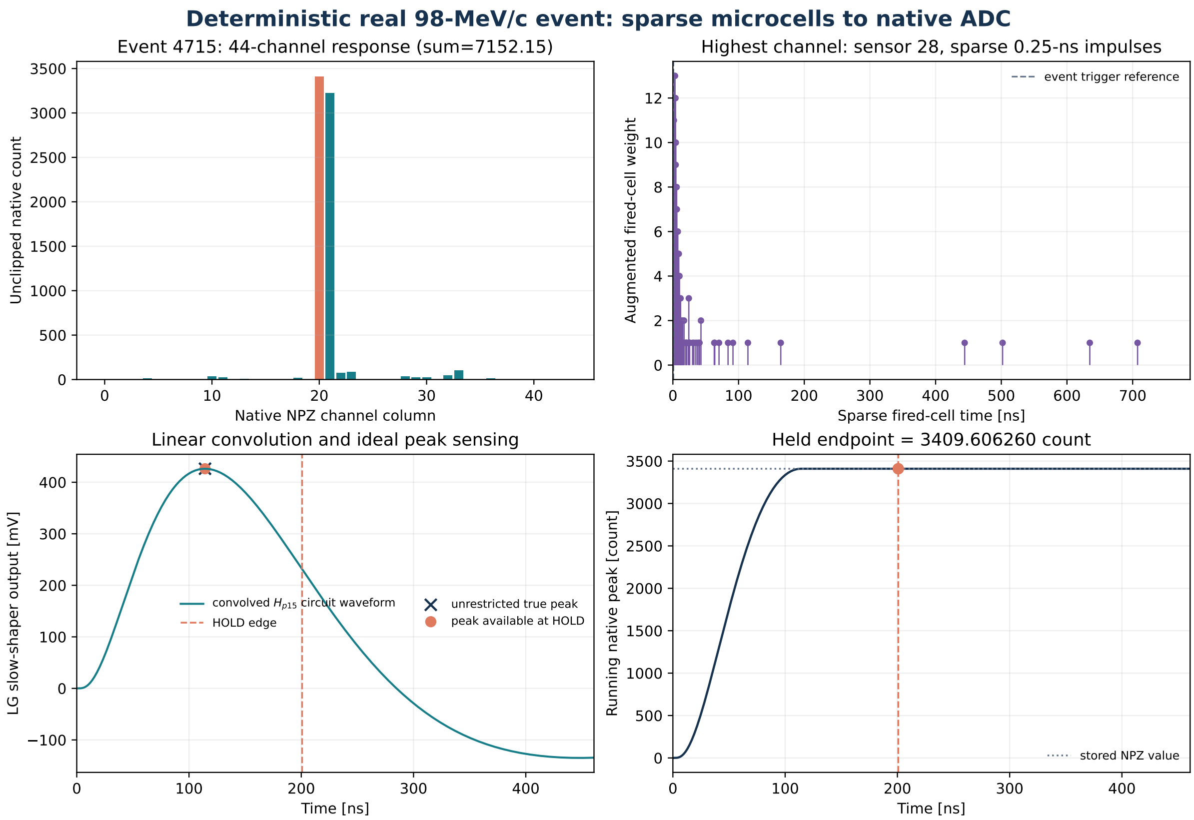

Figure 14: Deterministic worked 98-MeV/c event. Among event IDs 0–9999 with `truth_plane_recorded=1` and finite positive $S_e^{(44)}$, choose the sum closest to that eligible sample's median, break equal-distance ties with the lowest event ID, and then choose the highest-native-ADC channel. Its sparse augmented cells are followed through convolution, trigger, HOLD, and held peak. This choice is illustrative only and removes no event from the reported distributions.

11 Charge normalization and the native ADC endpoint

The held peak P_{ec} in Eq. (88) is still expressed in prompt-peak-equivalent cells. Only two fixed conversions remain: the charge assigned to one fired cell and the voltage-to-count scale implied by the declared ADC LSB. This section performs those conversions and then stops. It does not add an empirical detector calibration.

The later raw-data calibration fits a global empirical data-to-simulation ADC ratio $G_{\text{abs}} = 0.387995$, together with channel variations. In that fit, the simulated operand is the native endpoint defined in this report, so the relationship is schematically $\text{ADC}_{\text{data}} \simeq G_{\text{abs}}\text{ADC}_{\text{sim,abs}}$ before the remaining pedestal, noise, threshold, and selection terms are evaluated. The suffix "abs" identifies the literal page-15 amplitude convention used for $\text{ADC}_{\text{sim,abs}}$; it does not claim that G_{abs} is an absolute hardware gain and does not denote an absolute-value peak detector. The simulated operand uses the physical signed chain $+v_i \to -4v_i \to +v_{\text{shape}}$ and retains the positive shaper maximum.

The simulation operand already contains the fixed low-gain $\times 4$ preamplifier response and the literal page-15 shaper response. Consequently, G_{abs} is a downstream residual calibration ratio, not an extra CITIROC circuit factor, and it is not applied anywhere in the native arrays reported here. It absorbs a degenerate combination of optical yield, PDE, fired-cell charge, held-output/multiplexer transfer, external-ADC scale, and data/Monte Carlo selection effects; its value cannot isolate which term causes a mismatch and is not evidence for an omitted page-15 factor.

11.1 What the residual ratio diagnoses, and the proposed decision test

The cause of $G_{\text{abs}} = 0.387995$ remains unresolved. The completed audits rule out several tempting bookkeeping explanations for such a large ratio:

- low-gain $\times 4$ is present, with the independent $|\min v_{\text{pre}}|/\max v_i$ check equal to four;
- the literal page-15 factor ten is present in H_{p15} rather than omitted;
- an independent sampled state-space kernel agrees with the production implementation to 2.65×10^{-17} V, while the report's representative endpoint numerical closure is 9.1×10^{-13} count, excluding a consequential convolution or implementation error;
- the continuous ideal kernel peak is only 0.206% above the production sampled peak, so the finite-grid current-sample convention cannot explain a factor near 0.388; and
- dividing H_{p15} by the unit-step peak merely changes coordinates: $G_{\text{norm}} = AG_{\text{abs}}$ exactly preserves the predicted data value when the simulated signal, pedestal, and noise coordinates are divided by A.

The largest unmeasured electronics assumption is the unity, fully settled transfer from the internal shaped peak through the peak-memory cell, follower, multiplexer, and `OUT_LG` to the external ADC pin. Even so, the fitted ratio is also degenerate with the fired-cell charge, optical yield, PDE, correlated-noise model, readout mapping, event selection, and external-ADC reference. The present evidence therefore does not identify one cause.

A matched known-charge excitation is the proposed decision test. The quoted full-chain slope belongs specifically to the model initial condition in which an avalanche charge q_0 is placed across the SiPM capacitance C_S. For an actual optically pulsed SiPM whose avalanche charge is measured independently, or an electrical SiPM emulator that reproduces that same q_0/C_S topology, the complete model predicts

$$m_{\text{model,full}} = 13.01110046 \text{ counts/pC}, \qquad G_{\text{abs}} m_{\text{model,full}} = 5.04824192 \text{ counts/pC}. \tag{89}$$

Equivalently, the corresponding fired-cell values are 12.50764868 and 4.85290515 counts per declared fired-cell. The second slope is a *hypothetical full-chain* value if the entire fitted ratio were caused by hardware response; it is not automatically the expected slope of an internal test-capacitor injection.

Two complementary measurements are required because the slow-shaper peak is waveform-dependent.

1. Use an optically pulsed SiPM with independently known avalanche charge, or an electrical SiPM emulator that realizes the modeled q_0 initial condition on C_S and traverses the modeled R_T–C_B and ASIC input network. Only this model-equivalent full-chain excitation can be compared directly with $m_{\text{model,full}} = 13.01110046$ counts/pC and with the hypothetical 5.04824192-count/pC value. A generic pulse generator connected at the board or ASIC input has a different node and waveform and therefore requires its own modeled reference slope.
2. For the approximately 3 pF internal test capacitor, use $Q = C_{\text{test}}\Delta V$ only after measuring or otherwise characterizing its actual capacitance, drive waveform, and injection node, then model

that path to obtain a distinct $m_{\text{model,Ctest}}$. That injection bypasses some or all of the external SiPM/input network, so comparing it directly with $m_{\text{model,full}}$ would be invalid. Its purpose is to localize transfer downstream of the test node.

Both scans should preserve and record the full register state, including the exact low-gain $\times 4$ and 62.5 ns shaper configuration. Where the board exposes it, record the pre-memory slow-shaper monitor `Out_ssh_LG`, then held `OUT_LG`, ADC code, HOLD delay, and SELECT-to-CNV timing versus charge. Held `OUT_LG` plus ADC alone can measure their common transfer but cannot localize a preamplifier/shaper error against peak-memory/multiplexer behavior. The extra node and timing records, or measured voltage ratios between all accessible nodes, provide that localization.

For matched model and measurement slopes, define

$$K_{\text{full}} = \frac{m_{\text{meas,full}}}{m_{\text{model,full}}}, \qquad K_{\text{down}} = \frac{m_{\text{meas,Ctest}}}{m_{\text{model,Ctest}}}. \tag{90}$$

If the internal scan establishes that K_{down} is the missing transfer on the common held-output path, the actionable correction and remaining beam-data residual are

$$\text{ADC}_{\text{sim,corr}} = K_{\text{down}}\,\text{ADC}_{\text{sim,native}}, \qquad G_{\text{physics}} = \frac{\text{ADC}_{\text{data}}}{\text{ADC}_{\text{sim,corr}}} = \frac{G_{\text{abs}}}{K_{\text{down}}}. \tag{91}$$

An external matched slope near 5.04824 counts/pC would show that the complete hardware response carries approximately the fitted 0.388 ratio, but the internal scan is still needed to decide whether that loss occurs downstream or in the external input network. An external matched slope near 13.0111 counts/pC would leave the full modeled electronics scale intact and direct the investigation toward optics, PDE, fired-cell charge, correlated noise, mapping, and selection. No 5.048-versus-13.011 threshold is assigned to the internal test-capacitor scan; it must be judged against $m_{\text{model,Ctest}}$. These are proposed measurements and a conditional correction workflow, not completed charge-injection results or a claim that the cause has been found.

11.2 Fired-cell charge at the production operating point

The aligned nine-job optical campaign was run with the 5-V PDE curve and a declared SiPM avalanche gain $G_{\text{SiPM}} = 6.0 \times 10^6$ electrons per fired cell. Using the exact elementary-charge magnitude $q_e = 1.602176634 \times 10^{-19}$ C, the charge per cell is

$$q_{\text{cell}} = q_e G_{\text{SiPM}} = (1.602176634 \times 10^{-19}\ \text{C})(6.0 \times 10^6) = 0.9613059804\ \text{pC}. \tag{92}$$

The gain is the mean number of avalanche electrons represented by one effective fired cell, so multiplying by the electron charge gives coulombs; the final factor of 10^{12} converts coulombs to picocoulombs. This q_{cell} is the production digitizer's explicit gain-based charge. It must not be silently replaced by a different capacitance-times-overvoltage estimate, even though that alternative is useful as a separate hardware systematic study.

The same value appears in two conceptually different places in the frozen flow. Upstream, the declared gain labels the 5-V SiPM operating point used with the PDE and microcell settings. Here, it assigns physical charge to the already generated effective fired-cell series. The digitizer does not rerun the PDE decision, crosstalk, occupancy, or recovery when applying Eq. (92).

11.3 Circuit voltage to ideal ADC-input-equivalent count coordinate

Let $L_{\text{ADC}} = 125\ \mu\text{V/count}$ be the nominal ideal code width at the external LTC2323-14 differential input for the declared REFOUT $= 2.048$ V coordinate, and let $K = 1.626387557 \times 10^{-3}$ V/pC be the positive unit-charge kernel peak. The model assumes that the baseline-subtracted held positive slow-shaper maximum reaches that ADC input through the CITIROC peak-memory cell, follower, channel multiplexer, and `OUT_LG` path with unity, fully settled transfer. CITIROC1A datasheet Figure 15 establishes that topology but publishes neither a numeric end-to-end transfer nor a settling guarantee for this operating sequence. Accordingly, the conversion below defines an ideal ADC-input-equivalent coordinate rather than a measured internal-to-external gain. The nominal external-ADC scale follows from that reference and signed 15-bit span:

$$L_{\mathrm{ADC}} = \frac{2(2.048\,\mathrm{V})}{2^{15}} = 125\,\mu\mathrm{V/count}. \tag{93}$$

The numerator is the full differential span 2 REFOUT; dividing it into 2^{15} signed-code intervals gives volts per count. The device's specified 2.044 V to 2.053 V REFOUT range corresponds to 124.756 μV to 125.305 μV/count, a variation smaller than about 0.25%. This tolerance belongs in a hardware systematic but is far too small to explain the fitted ratio 0.387995. Within the declared nominal ideal coordinate, the charge mapping is

$$\gamma_{\mathrm{ADC}} = \frac{10^6 K}{L_{\mathrm{ADC}}} = 13.01110046\,\frac{\mathrm{counts}}{\mathrm{pC}}. \tag{94}$$

The factor 10^6 converts the kernel amplitude from volts to microvolts; division by microvolts per count then yields counts per picocoulomb. Every quantity on the right comes from the fixed circuit model or external ADC LSB, subject to the unity/settling and baseline-subtraction assumption above. No spectrum peak, beam energy, or data/MC alignment enters this conversion.

Combining the two fixed scales gives

$$\begin{aligned}\kappa_{\mathrm{cell}} &= q_{\mathrm{cell}}\gamma_{\mathrm{ADC}} \\ &= (0.9613059804\ \mathrm{pC/cell})(13.01110046\ \mathrm{counts/pC}) \\ &= 12.50764868\ \frac{\mathrm{counts}}{\text{peak-equivalent cell}}.\end{aligned} \tag{95}$$

Picocoulombs cancel between the factors. The remaining unit means that one prompt cell, observed after its complete peak and under the declared linear circuit, would contribute 12.50764868 native counts. A delayed or dispersed set of cells first changes P_{ec} through convolution and HOLD; the fixed factor κ_{cell} does not compensate for that timing effect.

11.4 Definition and storage contract

The endpoint for event e and native channel c is

$$\boxed{S_{ec} = P_{ec}\,q_{\mathrm{cell}}\,\gamma_{\mathrm{ADC}} = P_{ec}\,\kappa_{\mathrm{cell}}} \quad \longleftrightarrow \quad \texttt{direct\_5V\_channel\_adc\_unclipped}[e,c]. \tag{96}$$

P_{ec} is the ideal held peak in prompt-cell equivalents, q_{cell} converts cells to charge, and γ_{ADC} converts that charge to the ideal ADC-input-equivalent count coordinate. Their product S_{ec} is stored as a floating-point array with 44 columns ordered by Eq. (63). The boxed equality is the terminal scientific definition of this report.

The word "ADC" in S_{ec} specifies the ideal ADC-input-equivalent voltage coordinate set by L_{ADC}. It does *not* imply that S_{ec} is a hardware word or that the intervening peak-memory/multiplexer transfer has been measured. At this boundary there is

- no pedestal or added-noise draw;
- no fitted channel-dependent data-to-simulation residual scale;
- no discriminator threshold or zero suppression;
- no upper clipping, saturation cap, or valid-range operation;
- no rounding, integer quantization, signed-payload encoding, or offset;
- no two-SiPM crystal requirement or event rejection.

Consequently zero means that the modeled channel had no held response, while positive non-integer values are valid. There is no imposed upper bound.

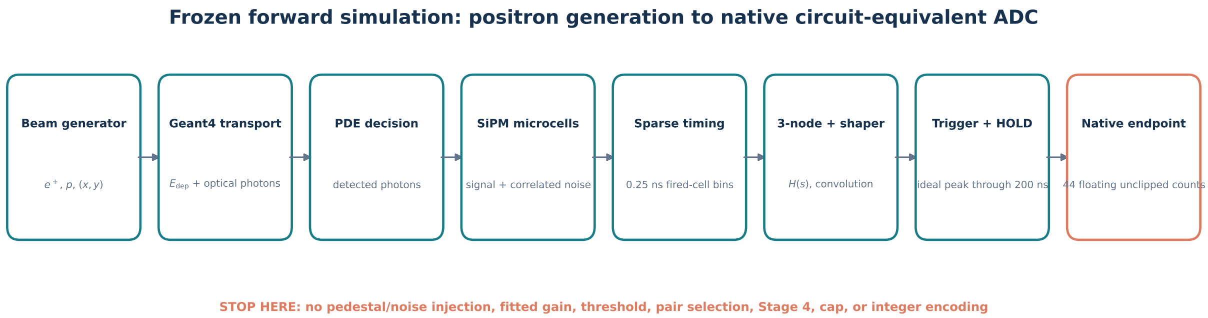

Figure 15: Frozen end-to-end flow from the generated positron through the per-channel native endpoint. The heavy terminal boundary after `direct_5V_channel_adc_unclipped` marks where this report stops; operations beyond it are not part of the native circuit prediction.

11.5 Unselected 44-channel diagnostic sum

To display one spectrum per event without inventing a crystal selection, the report uses the arithmetic sum over all native channels,

$$S_e^{(44)} = \sum_{c \in \mathcal{C}_{44}} S_{ec}. \tag{97}$$

Every one of the 44 columns participates. A zero channel simply adds zero; there is no fired-channel multiplicity test, coincidence requirement, beam- monitor cut, truth-crystal cut, or threshold. Therefore $S_e^{(44)}$ is a transparent pipeline diagnostic, not the selected calorimeter observable used in a later data analysis.

For energy setting E, let histogram bin B_b contain values satisfying $a_b \le s < a_{b+1}$. Its plotted probability is

$$h_{Eb} = \frac{1}{N_E} \sum_{e=1}^{N_E} \mathbf{1}\left(S_e^{(44)} \in B_b\right). \tag{98}$$

$N_E = 10{,}000$ is the number of generated events at that energy and the indicator counts events in bin b. Division by N_E makes the bin heights probabilities. Over bins covering every event, they obey $\sum_b h_{Eb} = 1$; if a finite displayed range omits events, the sum instead equals the displayed event fraction. Because h_{Eb} is not divided by bin width, it is probability per bin rather than a probability density and should be *summed*, not integrated. The panels are normalized independently; their heights do not encode relative beam flux.

The empirical cumulative distribution for an energy setting is

$$\widehat{F}_E(s) = \frac{1}{N_E} \sum_{e=1}^{N_E} \mathbf{1}\left(S_e^{(44)} \le s\right). \tag{99}$$

$\widehat{F}_E(s)$ is the observed fraction at or below s. For exact reproducibility, the plotted quantile uses linear interpolation between the zero-indexed ordered values $S_{E,[0]} \le \cdots \le S_{E,[N_E-1]}$:

$$h = (N_E - 1)q, \qquad i = \lfloor h \rfloor, \qquad \lambda = h - i, \qquad Q_E(q) = (1-\lambda) S_{E,[i]} + \lambda S_{E,[i+1]}, \tag{100}$$

with $Q_E(1) = S_{E,[N_E-1]}$ at the upper endpoint. The index i selects the ordered value immediately below fractional rank h, and λ gives the fractional distance to the next ordered value. This is the standard "linear" sample-quantile convention; stating it prevents software-default changes from silently shifting tail summaries. The report follows $q \in \{0.05, 0.50, 0.95, 0.999\}$ and also reports the arithmetic mean. These statistics describe the complete simulated distribution, including its low and high tails; none is a fitted parametric peak or resolution parameter.

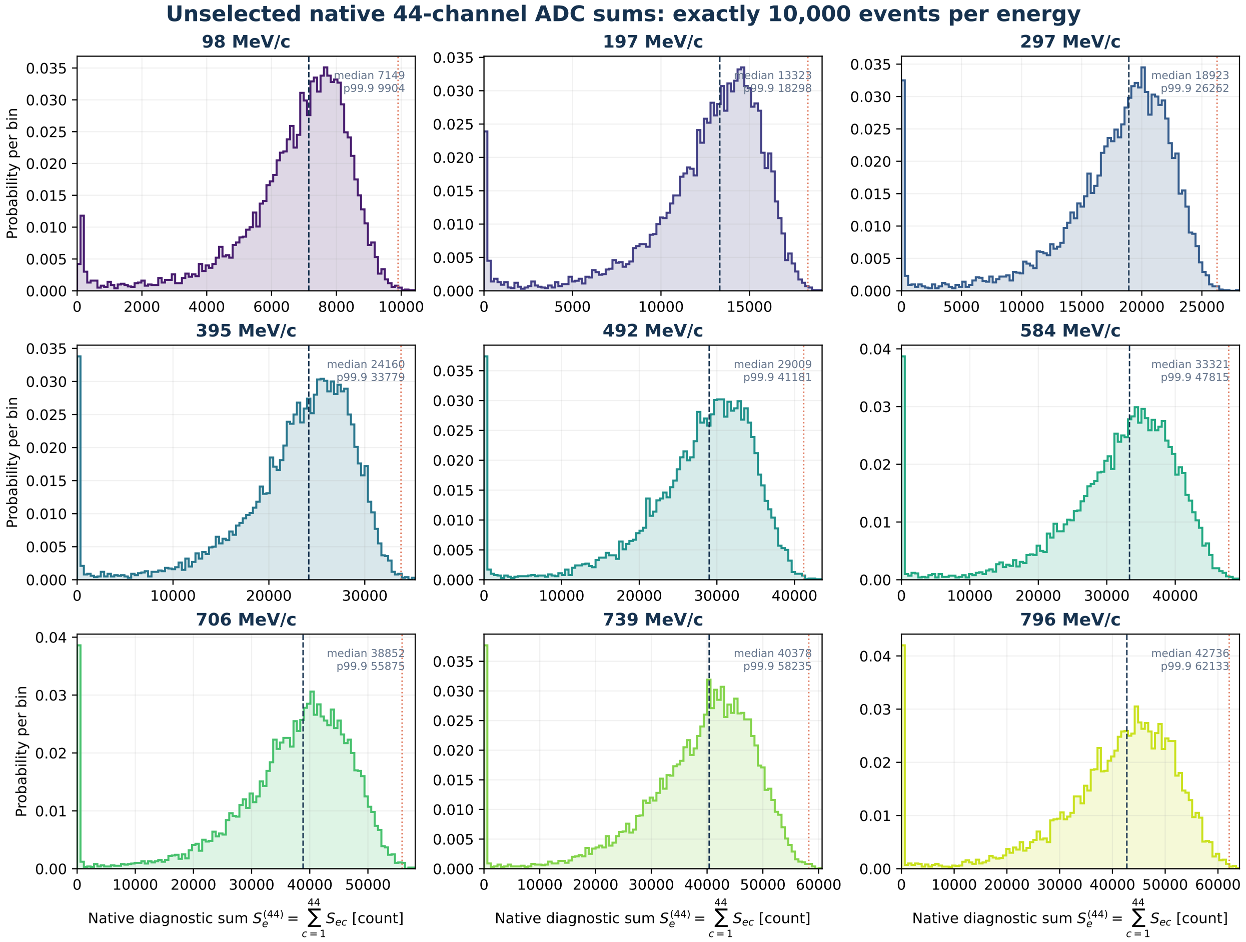

Figure 16: Full native $S_e^{(44)}$ spectra for generated event IDs 0–9999 at all nine momentum settings. Each panel is independently probability normalized. No threshold, event selection, or empirical calibration has been applied.

Table 14: Unselected native 44-channel diagnostic-sum statistics. Quantiles use NumPy's linear interpolation convention.

p [MeV/c]	N	p05	median	mean	SD	p95	p99.9	events with any $S_{ec} > 11$k [%]
98	10,000	3059.52	7149.34	6748.70	1797.78	8800.20	9903.68	0.000
197	10,000	4935.52	13322.52	12500.16	3593.99	16489.74	18298.06	0.000
297	10,000	5981.47	18922.84	17677.78	5297.32	23629.65	26261.62	2.300
395	10,000	7422.72	24159.95	22621.44	6836.24	30306.43	33779.06	32.490
492	10,000	8438.66	29009.46	27191.03	8348.95	36739.30	41180.79	58.290
584	10,000	8400.73	33320.58	31238.19	9787.83	42681.69	47815.04	71.820
706	10,000	10970.94	38851.61	36493.21	11435.31	50022.58	55875.35	80.820
739	10,000	13107.51	40378.09	37863.97	11737.21	51766.94	58234.53	82.560
796	10,000	7639.89	42735.79	40002.84	12998.39	55426.84	62132.59	83.990

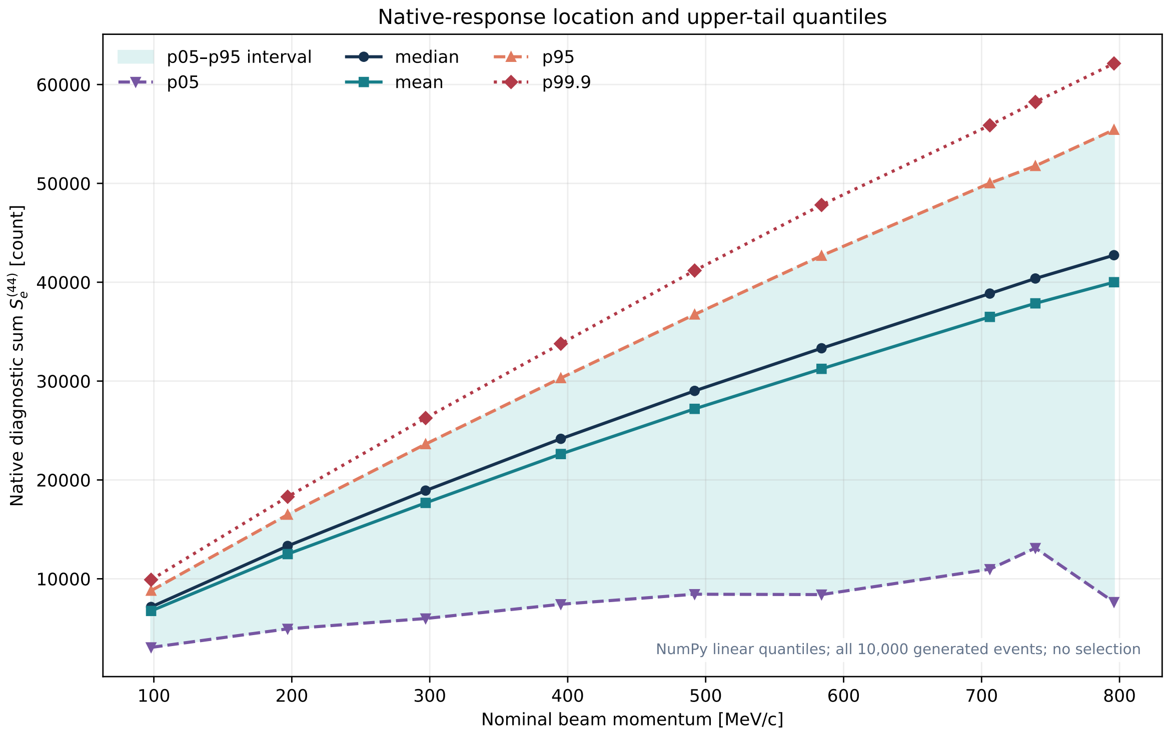

Figure 17: Mean and empirical quantiles of the unselected native 44-channel sum versus nominal beam-momentum setting. Connecting lines guide the eye; they are not an energy-calibration fit.

11.6 Why the clipped compatibility field is not the endpoint

The clipped compatibility field is formed by the channel-wise operation

$$S_{ec}^{\text{legacy}} = \min(S_{ec}, 11000) \quad \longleftrightarrow \quad \texttt{direct\_5V\_channel\_adc}[e, c]. \tag{101}$$

Values below 11,000 are unchanged, whereas every larger value is mapped to the same number. Equation (101) is therefore many-to-one and irreversibly removes upper-tail information. It is retained solely to audit the cap's information loss; every report figure or summary that uses a channel-ADC field uses S_{ec} from Eq. (96) rather than this legacy field.

For each energy and channel, the affected fraction is measured before any cap as

$$f_{Ec}^{>11\text{k}} = \frac{1}{N_E} \sum_{e=1}^{N_E} \mathbf{1}(S_{ec} > 11000). \tag{102}$$

This fraction answers how often the legacy operation would alter a native channel. It does not clip the values used to compute the fraction, and it is not interpreted as a measured hardware-saturation probability.

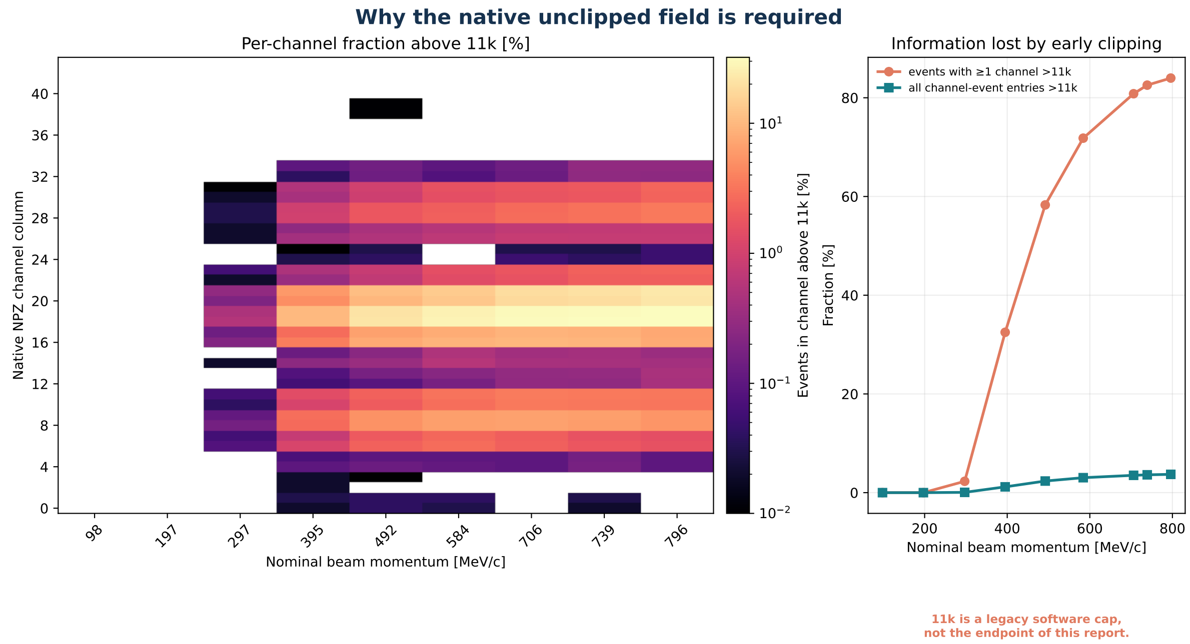

Figure 18: Per-channel fraction of native responses above the nominal 11,000-count reference boundary. Nonzero entries demonstrate why the unclipped field is required to preserve the simulated upper tail.

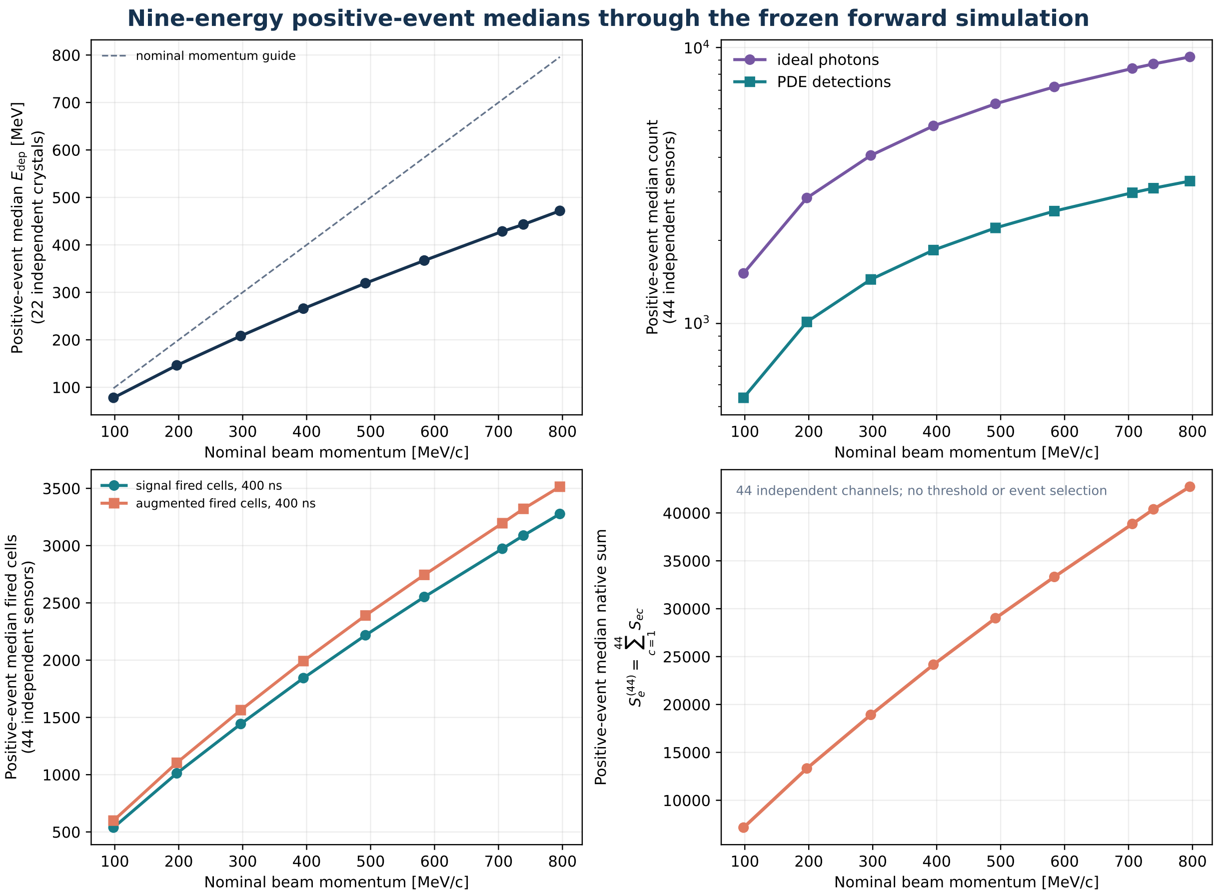

Figure 19: Nine-energy progression through frozen simulated quantities: deposited energy, photons reaching active silicon and PDE detections, signal and augmented effective fired cells, and the terminal native circuit sum. Each ordinate retains its own physical units; the panel is a flow audit, not a claim that the heterogeneous stages share one linear scale. All plotted points are positive-event medians over the 22 independent crystals and their 44 mapped independent sensors. Every pre-circuit point is regenerated from the same aligned set of nine direct 10,000-event ROOT files; fired-cell stages use the first 400 ns. The terminal native medians are computed from the regenerated aligned unclipped NPZ at every setting. Those 400-ns fired totals are a diagnostic only: the native circuit endpoint consumes the full 1-μs sparse per-sensor history.

12 Interpretation, limitations, and conclusions

12.1 What the native spectra do and do not measure

The rise of the native response with beam setting is a forward-simulation result: a higher-momentum positron usually deposits more energy, produces more transported scintillation photons, fires more effective microcells, and hence produces a larger circuit peak. The relationship need not be exactly linear. Longitudinal and transverse leakage, photon-transport fluctuations, finite microcell occupancy, correlated pulses, the event time profile, and fixed HOLD all act before $S_e^{(44)}$ is formed.

Accordingly, the ordinate in Figures 16 and 17 is a simulated ideal ADC-input-equivalent count coordinate, while the abscissa label is the nominal beam-momentum setting. It is not an event's true deposited energy and not a reconstructed energy. A connecting line between energy settings summarizes the frozen samples only; it neither fits a calibration law nor supplies an inverse map from counts to MeV.

The 44-channel arithmetic sum has one deliberate virtue: every endpoint column enters with coefficient one. It therefore exposes changes introduced by the forward pipeline without entangling them with a discriminator or crystal-acceptance rule. Its corresponding limitation is equally important: it covers the 22 independently read crystals, not every sensor of the full 6×6 Geant4 array and not the three combined-readout crystals.

12.2 Model limitations and uncertainty boundary

The forward calculation is reproducible but not a transistor-level or data-calibrated electronics simulation. Its principal limitations are:

1. **Effective BM placement, not a survey.** The common transverse transform is inferred from the raw BM profile and an event-matched PWO-minus-BM residual. It does not determine a unique mechanical six-degree-of-freedom pose, and the high-energy residual plateau indicates a small energy-dependent origin systematic that is not absorbed into separate per-energy detector placements.
2. **Gaussian-core BM closure.** The latent profile fit matches five reconstructed core moments after an ideal strip/gap/aperture fold. It does not fit thresholds, dead/noisy efficiency, raw ADC response, or multi-strip response. Plane-resolved selected-position shapes remain detector-level diagnostics rather than hidden fit constraints.
3. **Linear common circuit.** All 44 channels use the same three-node, low-gain, slow-shaper kernel. Component tolerances, channel-to-channel gain variation, pedestal drift, analogue nonlinearity, and real saturation are not modeled.
4. **Declared trigger/HOLD timing.** The optical ROOT has fired-cell times but no measured fast-shaper crossing or HOLD edge. The first augmented cell among the retained channels defines the trigger, the latency is zero, and HOLD closes after exactly 200 ns. Dark activity can therefore move the modeled trigger earlier.
5. **Ideal positive-peak sensing.** The running positive maximum has no droop, aperture width, clock jitter, comparator error, or pickup. The model does not posit a minimum detector, rectifier, or programmable analogue polarity flip; the digital Trigger Polarity control is unrelated to this analogue peak-memory operation. These omissions affect the relationship between the model coordinate and a real acquisition word, not the internal definition of S_{ec}.
6. **Fixed charge convention.** One effective fired cell is assigned 0.9613059804 pC from the declared gain 6.0×10^6. Uncertainty or temperature dependence of that gain is not propagated. A separate capacitance-times-overvoltage estimate is intentionally not substituted.

7. **Finite optical statistics.** The aligned study uses one direct 10,000-event sample at every setting. Finite Monte Carlo sampling therefore remains in every distribution and tail quantile; this report does not claim a larger-sample extrapolation.
8. **Restricted channel topology.** The endpoint covers the 44 independently read SiPMs in 22 crystals. The three differently scaled combined readouts are excluded rather than approximated, and the other simulated sensors outside the mapped readout region are not silently added.
9. **No downstream detector effects.** Pedestal, electronic noise, thresholding, range enforcement, quantization, payload encoding, and event cuts are absent. Therefore agreement or disagreement between this native count and a measured ADC branch cannot be interpreted without an explicitly separate detector-output model.

No error band in this report claims to cover all of these effects. The empirical distributions quantify event-to-event Monte Carlo variation only. They do not marginalize over finite-sample uncertainty, circuit constants, optical-property uncertainty, geometry tolerances, or SiPM operating-point uncertainty.

12.3 Conclusions

For a frozen ROOT event and channel map, the pipeline is deterministic. It reconstructs each retained sensor's 0.25 ns augmented fired-cell series, applies the exact matrix solution of the declared three-node input network, the CR–RC^2-equivalent shaper, and ideal 200-ns HOLD, and converts the held peak with fixed charge and voltage scales. The terminal result is the $N_E \times 44$ floating array `direct_5V_channel_adc_unclipped`.

For the aligned campaign, all nine native products and their dependent figures, tables, manifests, and audits were regenerated successfully. The completed optical ROOT set and its corresponding native arrays therefore form one checksum-bound end-to-end report population.

The endpoint is native in a narrow and useful sense: its amplitude scale follows from the circuit model and declared LSB, with no fitted beam-data to simulation-ADC residual scale applied. The upstream BM geometry does use the separately documented beam-derived effective placement. The endpoint is also intentionally incomplete as a model of recorded data. Preserving the unclipped float retains the information needed for any later, separately documented detector-output transformation, while preventing that later transformation from being mistaken for Geant4 or circuit physics.

A Circuit-to-ADC derivation and validation ledger

This appendix expands the two algebraic steps most easily obscured by a software implementation and lists the invariants that make the numerical product auditable.

A.1 Scalar Kirchhoff equations behind the nodal matrices

Writing Kirchhoff's current law separately at the avalanche, input, and block-capacitor nodes gives

$$C_S \frac{d(v_a - v_i)}{dt} + \frac{v_a}{R_q} = 0, \tag{103}$$

$$C_S \frac{d(v_i - v_a)}{dt} + C_I \frac{dv_i}{dt} + \frac{v_i - v_b}{R_T} = 0, \tag{104}$$

$$C_B \frac{dv_b}{dt} + \frac{v_b - v_i}{R_T} = 0. \tag{105}$$

In Eq. (103), the sensor-capacitor current and recovery-resistor current leave the avalanche node. In Eq. (104), the opposite sensor-capacitor current, the CITIROC input-load current, and the termination-branch current sum to zero. Equation (105) balances current into the block capacitor against current through the termination resistor. Collecting coefficients of $d\boldsymbol{v}/dt$ and $\boldsymbol{v}$ produces exactly the matrices in Eq. (67); no branch is introduced by the matrix notation.

A.2 Discrete pole cascade and its analogue target

The one-pole recursion in Eq. (77) has the discrete transfer $L_d(z)$ in Eq. (78). The cascade used by the code factors as

$$H_d(z) = -10\big[L_d(z)^2 - L_d(z)^3\big] = -10 L_d(z)^2 \big[1 - L_d(z)\big]. \tag{106}$$

The first equality mirrors the computed $-10(u_2 - u_3)$ and the second factors the two common discrete poles. Equation (106) therefore describes the implemented sampled filter exactly.

Its connection to the continuous target is a small-step limit, not a finite- step identity:

$$\lim_{\Delta t \to 0} L_d\big(e^{s\Delta t}\big) = \frac{1}{1 + s\tau_s}, \qquad \lim_{\Delta t \to 0} H_d\big(e^{s\Delta t}\big) = -\frac{10 s\tau_s}{(1 + s\tau_s)^3} = H(s). \tag{107}$$

Substituting $z = e^{s\Delta t}$ relates discrete and continuous frequency variables. As Δt tends to zero, expanding the exponentials in $\alpha = e^{-\Delta t/\tau_s}$ gives the analogue one-pole response and then the target $H(s)$. At the production value $\Delta t = 0.25\,\text{ns}$, however, the calculation uses $H_d(z)$ and makes no claim of exact equality to the continuous transfer.

A.3 Prompt-cell normalization check

Suppose M effective cells arrive in one bin n_0 and HOLD closes after the kernel maximum. The input and convolution become

$$x[n] = M\,\delta_{nn_0}, \qquad \widetilde{v}[n] = M\,k[n - n_0], \qquad P = \frac{M\,\max_n k[n]}{K} = M. \tag{108}$$

The Kronecker delta δ_{nn_0} is one only in the arrival bin. Linearity then produces M shifted unit kernels, and division by their unit peak $K = \max k$ returns exactly M. This is why P is called a peak-equivalent-cell count. For a distributed time series the first equality does not hold, so P need not equal the arithmetic fired-cell total.

A.4 HOLD bound

Because the samples available before HOLD form a subset of the complete waveform samples,

$$0 \le A_{ec}^{\text{hold}} \le A_{ec}^{\text{true}}, \qquad 0 \le R_{ec}^{\text{hold}} \le 1 \quad (A_{ec}^{\text{true}} > 0). \tag{109}$$

Both peaks are maxima of the positive physical-polarity waveform and are therefore nonnegative for an active channel. A maximum over a subset cannot exceed the maximum over the full set, which gives the first inequality; dividing by the positive true peak gives the second. Values outside this interval would expose an indexing or numerical error.

A.5 Dimensional closure of the endpoint

The complete unit cancellation is

$$\underbrace{P_{ec}}_{\text{peak-equivalent cells}} \; \underbrace{q_{\text{cell}}}_{\text{pC/cell}} \; \underbrace{\gamma_{\text{ADC}}}_{\text{counts/pC}} = \underbrace{S_{ec}}_{\text{native counts}} . \tag{110}$$

Peak-equivalent cells cancel cells, and picocoulombs cancel picocoulombs, leaving counts. This dimensional check also shows why a fitted data-to-simulation residual scale would be a new transformation rather than part of the endpoint defined here.

A.6 Required numerical invariants

The production digitizer and the report rebuild collectively treat the following as failures rather than warnings. Their audit scopes are stated explicitly so a report-level spot check is not confused with a new exhaustive redigitization:

- the event identifiers in an input shard are not contiguous and unique;
- in production digitization, any channel/event sparse timing total differing from its ROOT one-microsecond summary by more than 2×10^{-3} effective cells; the report rebuild independently repeats this closure test for the plotted representative event across all 44 retained channels;
- the frozen map does not contain exactly 47 rows, three combined rows, 44 unique singleton sensors, and 22 two-singleton crystals;
- a circuit matrix parameter is nonpositive, a finite pole fails to decay, or the unit kernel has a nonpositive physical-polarity peak;
- an output energy does not contain the declared primary event IDs 0–9999 or the endpoint does not have exactly 44 ordered columns;
- any primary spectrum or summary is built from the clipped compatibility field rather than `direct_5V_channel_adc_unclipped`.

Together with the source hashes in the package provenance, these invariants connect the mathematical definition to a reproducible array rather than only to a schematic description.

A.7 Frozen source manifest

Table 15 is the human-readable index of the files that define or attest this report. The table abbreviates digests for page layout; `provenance/source_manifest.csv` stores every absolute source path, byte size, full SHA-256 digest, and verification mode.

Table 15: Frozen sources used by this report. Full paths, byte sizes, complete digests, and verification modes are in `provenance/source_manifest.csv`.

Role	Source path	SHA-256
report asset builder	REPORT_BUILD_ROOT/build_report_assets.py	fce8e8ebb1…62dbeacd
aligned campaign 47-row channel map	BM_ALIGNED_CAMPAIGN/input/mc-readout-to-real-gid-map.csv	de79f6b94e…979db9c3
aligned campaign effective BM transform	BM_ALIGNED_CAMPAIGN/input/bm_effective_alignment.json	198d39f603…1478a265
aligned campaign audited GDML	BM_ALIGNED_CAMPAIGN/input/mdpwo-bm-up2-sensitive-aligned.gdml	596ab478cb…600e7f93
aligned campaign optical ROOT schema audit	BM_ALIGNED_CAMPAIGN/provenance/optical_root_schema_audit.json	f3ce4c0791…4a14c8ce
aligned campaign optical sidecar audit	BM_ALIGNED_CAMPAIGN/provenance/optical_root_sidecar_verification.txt	5d36217202…6f9ede67
aligned campaign tooling checksum inventory	BM_ALIGNED_CAMPAIGN/provenance/campaign_tooling_SHA256SUMS	c7004cd318…7a25ef01
aligned campaign binary checksum inventory	BM_ALIGNED_CAMPAIGN/provenance/new_binary_SHA256SUMS	314bba699d…01b602a5
canonical cluster-998 native digitizer audit	BM_ALIGNED_CAMPAIGN/provenance/native_digitizer_cluster998_audit.json	9a4a279d32…2f44f2c9
canonical cluster-998 native digitizer audit completion sidecar	BM_ALIGNED_CAMPAIGN/provenance/native_digitizer_cluster998_audit.json.sha256	87c936a395…aab156d7
aligned optical 168-core submit description	BM_ALIGNED_CAMPAIGN/condor/optical-tidc04-168core.sub	27d8375a39…4136316c
aligned native digitizer submit description	BM_ALIGNED_CAMPAIGN/condor/native-digitizer-tidc04.sub	c66e1e7992…fe632d1e
aligned optical campaign runner	ASGC_PROJECT/tools/gdml_beamtest/bm_alignment_campaign/run_bm_aligned_all9_10k_tidc04.sh	b542d06fa9…d534aaf6
aligned native digitizer campaign runner	ASGC_PROJECT/tools/gdml_beamtest/bm_alignment_campaign/run_bm_aligned_all9_10k_native_digitizer_tidc04.sh	6d905353f9…961def71
aligned native runtime requirements	BM_ALIGNED_CAMPAIGN/analysis/python-site-bookworm-py311.requirements.txt	21b473c862…4ca6145c
aligned native resolved runtime	BM_ALIGNED_CAMPAIGN/analysis/python-site-bookworm-py311.resolved.txt	bd1229a481…084d3b59
aligned native runtime manifest	BM_ALIGNED_CAMPAIGN/analysis/python-site-bookworm-py311.runtime.json	7a7640d062…d3af8318
phase-1 simulation methods	REPORT_DEPENDENCY_ROOT/output/generated/nine-energy-geant4-methods-phase1-draft-20260820/REPORT_PHASE1_GEANT4_SETUP.md	efa6d14456…df45f80c
CITIROC1A datasheet, printed p.15 slow-shaper source	CITIROC1A_DATASHEET	e952fe3348…99a18391
packaged CITIROC1A V2.53 datasheet, positive-signal and peak-detector source	REPORT_DEPENDENCY_ROOT/studies/citiroc_input_network/reference/CITIROC1A-Datasheet.53.pdf	178eaf2a8d…92eb93dd
MicroFC sensor datasheet, terminal capacitance source	REPORT_DEPENDENCY_ROOT/studies/citiroc_input_network/reference/MICROC-SERIES-D.pdf	278c09d1d2…ee41a3f7
Marano 2014 supporting ideal-op-amp PSPICE CITIROC model	REPORT_DEPENDENCY_ROOT/studies/citiroc_input_network/reference/CITIROC_High-Level_Analog_Front-End_Model_Marano_2014.pdf	a93014f733…7f389c0a
LTC2323-14 official datasheet, nominal reference and tolerance	REPORT_DEPENDENCY_ROOT/studies/citiroc_input_network/reference/LTC2323-14.pdf	acb684448f…8ce4c341
GTM board and external-ADC hardware audit	REPORT_DEPENDENCY_ROOT/output/generated/pwo-gtm-circuit-audit-20260816/gtm-circuit-audit.json	debb5d8772…afced9c2
independent ideal full-chain circuit validation audit	REPORT_DEPENDENCY_ROOT/output/generated/citiroc1a_page15_full_chain_circuit_simulation_20260822/audit.json	43ac9f1f5c…bafe5430
independent ideal full-chain circuit validation metrics	REPORT_DEPENDENCY_ROOT/output/generated/citiroc1a_page15_full_chain_circuit_simulation_20260822/metrics.csv	784c81e358…66fd4fda

Role	Source path	SHA-256
independent ideal full-chain circuit simulator source	REPORT_DEPENDENCY_ROOT/output/generated/citiroc1a_page15_full_chain_circuit_simulation_20260822/provenance/simulate_citiroc_page15_circuit.py	bbcf556949…55c4b7cb
all-nine physical-polarity endpoint-equivalence audit	BM_ALIGNED_CAMPAIGN/analysis/citiroc-physical-polarity-equivalence-20260822/citiroc_physical_polarity_endpoint_equivalence_all9_10k.json	59c3485970…7a1e34c9
all-nine physical-polarity endpoint-equivalence completion sidecar	BM_ALIGNED_CAMPAIGN/analysis/citiroc-physical-polarity-equivalence-20260822/citiroc_physical_polarity_endpoint_equivalence_all9_10k.json.sha256	8972cac671…95353579
all-nine physical-polarity endpoint-equivalence strict package inventory	BM_ALIGNED_CAMPAIGN/analysis/citiroc-physical-polarity-equivalence-20260822/SHA256SUMS	4c378ca1b6…278d6865
all-nine physical-polarity endpoint-equivalence run log	BM_ALIGNED_CAMPAIGN/analysis/citiroc-physical-polarity-equivalence-20260822/run.log	3dcd84437f…f5bccac1
frozen raw-coincidence exact-zero BM event sample	REPORT_DEPENDENCY_ROOT/output/generated/bm-raw-coincidence-zero-degree-hit-sigma-nine-energy-20260818/BM_raw_coincidence_zero_degree_event_sample.csv.gz	3537bf346b…324fbfba
frozen raw BM Gaussian-core target summary	REPORT_DEPENDENCY_ROOT/output/generated/bm-beam-profile-for-alignment-review-nine-energy-20260818/BM_beam_profile_center_sigma_by_energy.csv	ccc2a10a52…2681f4b5
event-matched PWO-minus-BM mean-position residual	REPORT_DEPENDENCY_ROOT/output/generated/pwo-bm-residual-event-matched-mean-positions-nine-energy-20260818/PWO_BM_residual_event_matched_mean_positions.csv	f73d879e97…13353152
latent pre-aperture all-nine generator parameters	BM_ALIGNED_CAMPAIGN/input/bm_latent_beam_parameters.csv	6b5f78c1d9…04894ba0
latent strip forward-fold fit summary	REPORT_DEPENDENCY_ROOT/output/generated/bm-latent-gaussian-strip-forward-fold-nine-energy-20260821/summary.json	0c7d1a7e1d…a627abbe
BM detector-level closure analyzer	REPORT_DEPENDENCY_ROOT/tools/gdml_beamtest/analyze_bm_mc_raw_closure.py	be900d37df…9fe61a13
BM closure campaign input manifest	BM_ALIGNED_CAMPAIGN/mc_manifest_edep_10k.csv	1312860176…4bf51f3f
BM detector-level closure verdict	BM_ALIGNED_CAMPAIGN/analysis/bm_raw_closure/bm_closure_summary.json	2c865596df…725894f7
BM detector-level closure metrics	BM_ALIGNED_CAMPAIGN/analysis/bm_raw_closure/bm_closure_metrics.csv	38e12bbd79…c70bf5ce
BM detector-level ROOT source manifest	BM_ALIGNED_CAMPAIGN/analysis/bm_raw_closure/bm_mc_source_manifest.csv	f77b459a45…eb12753d
authoritative BM selected-position data-vs-aligned-MC package: README.md	BM_ALIGNED_CAMPAIGN/analysis/bm-plane-selected-position-data-vs-aligned-edep-mc-all9-20260822/README.md	1003ee4357…2064598e
authoritative BM selected-position data-vs-aligned-MC package: SHA256SUMS	BM_ALIGNED_CAMPAIGN/analysis/bm-plane-selected-position-data-vs-aligned-edep-mc-all9-20260822/SHA256SUMS	451701e95d…1254b499
authoritative BM selected-position data-vs-aligned-MC package: audit.json	BM_ALIGNED_CAMPAIGN/analysis/bm-plane-selected-position-data-vs-aligned-edep-mc-all9-20260822/audit.json	615612ab8c…7b4ebfdf
authoritative BM selected-position data-vs-aligned-MC package: bm1_x_selected_position_data_vs_aligned_edep_mc_3x3.pdf	BM_ALIGNED_CAMPAIGN/analysis/bm-plane-selected-position-data-vs-aligned-edep-mc-all9-20260822/bm1_x_selected_position_data_vs_aligned_edep_mc_3x3.pdf	24f0ef0d16…158c1b31
authoritative BM selected-position data-vs-aligned-MC package: bm1_x_selected_position_data_vs_aligned_edep_mc_3x3.png	BM_ALIGNED_CAMPAIGN/analysis/bm-plane-selected-position-data-vs-aligned-edep-mc-all9-20260822/bm1_x_selected_position_data_vs_aligned_edep_mc_3x3.png	f9b44f0f47…ed0498bf

Role	Source path	SHA-256
authoritative BM selected-position data-vs-aligned-MC package: bm1_y_selected_position_data_vs_aligned_edep_mc_3x3.pdf	BM_ALIGNED_CAMPAIGN/analysis/bm-plane-selected-position-data-vs-aligned-edep-mc-all9-20260822/bm1_y_selected_position_data_vs_aligned_edep_mc_3x3.pdf	2a71ae77c6…82380d83
authoritative BM selected-position data-vs-aligned-MC package: bm1_y_selected_position_data_vs_aligned_edep_mc_3x3.png	BM_ALIGNED_CAMPAIGN/analysis/bm-plane-selected-position-data-vs-aligned-edep-mc-all9-20260822/bm1_y_selected_position_data_vs_aligned_edep_mc_3x3.png	022438db56…97f58b48
authoritative BM selected-position data-vs-aligned-MC package: bm2_x_selected_position_data_vs_aligned_edep_mc_3x3.pdf	BM_ALIGNED_CAMPAIGN/analysis/bm-plane-selected-position-data-vs-aligned-edep-mc-all9-20260822/bm2_x_selected_position_data_vs_aligned_edep_mc_3x3.pdf	5f3523a82b…7592b637
authoritative BM selected-position data-vs-aligned-MC package: bm2_x_selected_position_data_vs_aligned_edep_mc_3x3.png	BM_ALIGNED_CAMPAIGN/analysis/bm-plane-selected-position-data-vs-aligned-edep-mc-all9-20260822/bm2_x_selected_position_data_vs_aligned_edep_mc_3x3.png	df23df412c…ad07b2b0
authoritative BM selected-position data-vs-aligned-MC package: bm2_y_selected_position_data_vs_aligned_edep_mc_3x3.pdf	BM_ALIGNED_CAMPAIGN/analysis/bm-plane-selected-position-data-vs-aligned-edep-mc-all9-20260822/bm2_y_selected_position_data_vs_aligned_edep_mc_3x3.pdf	0c653f4b6e…03d9dac3
authoritative BM selected-position data-vs-aligned-MC package: bm2_y_selected_position_data_vs_aligned_edep_mc_3x3.png	BM_ALIGNED_CAMPAIGN/analysis/bm-plane-selected-position-data-vs-aligned-edep-mc-all9-20260822/bm2_y_selected_position_data_vs_aligned_edep_mc_3x3.png	11335f7956…0f6fbc01
authoritative BM selected-position data-vs-aligned-MC package: bm_plane_data_mc_bins.csv	BM_ALIGNED_CAMPAIGN/analysis/bm-plane-selected-position-data-vs-aligned-edep-mc-all9-20260822/bm_plane_data_mc_bins.csv	dcf27e1408…79d9eebe
authoritative BM selected-position data-vs-aligned-MC package: bm_plane_data_mc_summary.csv	BM_ALIGNED_CAMPAIGN/analysis/bm-plane-selected-position-data-vs-aligned-edep-mc-all9-20260822/bm_plane_data_mc_summary.csv	c22acf23fb…7f04c8e8
authoritative BM selected-position data-vs-aligned-MC package: source_manifest.csv	BM_ALIGNED_CAMPAIGN/analysis/bm-plane-selected-position-data-vs-aligned-edep-mc-all9-20260822/source_manifest.csv	bf7cf540e1…fb4a6dce
audited GDML placement inventory used for front-face drawing	REPORT_DEPENDENCY_ROOT/output/generated/nine-energy-geant4-methods-phase1-draft-20260820/geometry_placements_all_381_volumes.csv	39d7b6a4df…78447374
phase-1 optical properties as used	REPORT_DEPENDENCY_ROOT/output/generated/nine-energy-geant4-methods-phase1-draft-20260820/wavelength_dependent_properties_as_used.csv	928950140e…e610d01f
audited phase-1 geometry figure source	REPORT_DEPENDENCY_ROOT/output/generated/nine-energy-geant4-methods-phase1-draft-20260820/geometry_longitudinal_layout.png	1fcfeef111…ad0e9112
audited phase-1 optical-curves figure source	REPORT_DEPENDENCY_ROOT/output/generated/nine-energy-geant4-methods-phase1-draft-20260820/wavelength_dependent_properties_as_used.png	41aabf868e…f306a2b5

Role	Source path	SHA-256
aligned native digitizer entry point	ASGC_PROJECT/tools/gdml_beamtest/prepare_single_true5v_digitizer_observables.py	c17af5a076…aa3b74e6
archived aligned native circuit projection	ASGC_PROJECT/tools/gdml_beamtest/compare_low3_circuit_adc_to_data.py	c451f69f00…ac33a079
archived ballistic convolution and HOLD	ASGC_PROJECT/tools/gdml_beamtest/analyze_ballistic_adc_calibration.py	9f9bf4e97f…1e24459c
archived three-node circuit implementation recorded by native summaries	ASGC_PROJECT/tools/gdml_beamtest/citiroc_input_network.py	fe7dc026b1…4fb47fc6
current physical-polarity three-node circuit implementation	REPORT_DEPENDENCY_ROOT/tools/gdml_beamtest/citiroc_input_network.py	b1f9233c73…7a12f6e1
current physical-polarity native circuit projection	REPORT_DEPENDENCY_ROOT/tools/gdml_beamtest/compare_low3_circuit_adc_to_data.py	a0a3d7489b…6cafb1e9
current positive-maximum ballistic convolution and HOLD	REPORT_DEPENDENCY_ROOT/tools/gdml_beamtest/analyze_ballistic_adc_calibration.py	106666e07f…caf83c5a
all-nine physical-polarity endpoint-equivalence comparator	REPORT_DEPENDENCY_ROOT/tools/gdml_beamtest/audit_citiroc_polarity_endpoint_equivalence.py	0c93a606d2…be6181b8
aligned production executable	BM_ALIGNED_CAMPAIGN/build/eic-geant4-bookworm/zdc-gdml-beamtest	d1c70788ec…7351bda1
aligned optical input: pwo_rindex.csv	ASGC_PROJECT/optical_data/pwo_rindex.csv	44eb6eea38…1fe6420a
aligned optical input: pwo_absorption_mm.csv	ASGC_PROJECT/optical_data/pwo_absorption_mm.csv	1466256327…d5479dee
aligned optical input: pwo_emission.csv	ASGC_PROJECT/optical_data/pwo_emission.csv	8a1c8044f7…5ea9e255
aligned optical input: grease_rindex.csv	ASGC_PROJECT/optical_data/grease_rindex.csv	e61568e1a9…80296ffb
aligned optical input: grease_absorption_mm.csv	ASGC_PROJECT/optical_data/grease_absorption_mm.csv	4507c241ef…c4764072
aligned optical input: esr_reflectivity.csv	ASGC_PROJECT/optical_data/esr_reflectivity.csv	09da111d23…a708c53c
aligned optical input: onsemi_microfc_60035_cseries_pde_5v.csv	ASGC_PROJECT/optical_data/onsemi_microfc_60035_cseries_pde_5v.csv	0411dab984…d1d996d9
98 MeV/c aligned optical ROOT	BM_ALIGNED_CAMPAIGN/results/optical-E098MeV-10000events.root	ed19c1221a…21525bbc
98 MeV/c aligned native NPZ	BM_ALIGNED_CAMPAIGN/digitized-native/optical-E098MeV-10000events.native-true5v.npz	83498fd315…4f87c8ef
98 MeV/c aligned optical ROOT checksum sidecar	BM_ALIGNED_CAMPAIGN/results/optical-E098MeV-10000events.root.sha256	73a3ea54af…24032c9f
98 MeV/c aligned optical run log	BM_ALIGNED_CAMPAIGN/logs/optical-E098MeV-10000events.attempt-manual-0-1193302.log	49f9dffe03…f319d0df
98 MeV/c aligned native digitizer summary	BM_ALIGNED_CAMPAIGN/digitized-native/optical-E098MeV-10000events.native-true5v.summary.json	3cc1d12c0e…c696635f
98 MeV/c aligned native completion sidecar	BM_ALIGNED_CAMPAIGN/digitized-native/optical-E098MeV-10000events.native-true5v.sha256	23174016c5…377dfb37
197 MeV/c aligned optical ROOT	BM_ALIGNED_CAMPAIGN/results/optical-E197MeV-10000events.root	a816105597…3f8c87d0
197 MeV/c aligned native NPZ	BM_ALIGNED_CAMPAIGN/digitized-native/optical-E197MeV-10000events.native-true5v.npz	eb1dddcd7d…44ad3f5d
197 MeV/c aligned optical ROOT checksum sidecar	BM_ALIGNED_CAMPAIGN/results/optical-E197MeV-10000events.root.sha256	8224b72bb1…d771aa94
197 MeV/c aligned optical run log	BM_ALIGNED_CAMPAIGN/logs/optical-E197MeV-10000events.attempt-manual-1-1283681.log	0156ef42c2…af251f6d
197 MeV/c aligned native digitizer summary	BM_ALIGNED_CAMPAIGN/digitized-native/optical-E197MeV-10000events.native-true5v.summary.json	eb93171244…d213a6ae
197 MeV/c aligned native completion sidecar	BM_ALIGNED_CAMPAIGN/digitized-native/optical-E197MeV-10000events.native-true5v.sha256	4b795283f0…1456c4b8
297 MeV/c aligned optical ROOT	BM_ALIGNED_CAMPAIGN/results/optical-E297MeV-10000events.root	9806c49f1f…3fc5bb5f

Role	Source path	SHA-256
297 MeV/c aligned native NPZ	BM_ALIGNED_CAMPAIGN/digitized-native/optical-E297MeV-10000events.native-true5v.npz	83939002d7…18ea30aa
297 MeV/c aligned optical ROOT checksum sidecar	BM_ALIGNED_CAMPAIGN/results/optical-E297MeV-10000events.root.sha256	13aded6628…79912485
297 MeV/c aligned optical run log	BM_ALIGNED_CAMPAIGN/logs/optical-E297MeV-10000events.attempt-manual-2-1251204.log	e65b9d8249…d4167eda
297 MeV/c aligned native digitizer summary	BM_ALIGNED_CAMPAIGN/digitized-native/optical-E297MeV-10000events.native-true5v.summary.json	03acd878d4…e90759e7
297 MeV/c aligned native completion sidecar	BM_ALIGNED_CAMPAIGN/digitized-native/optical-E297MeV-10000events.native-true5v.sha256	6d8cdbadd5…d2442d09
395 MeV/c aligned optical ROOT	BM_ALIGNED_CAMPAIGN/results/optical-E395MeV-10000events.root	43a38d2cec…0e70ff5b
395 MeV/c aligned native NPZ	BM_ALIGNED_CAMPAIGN/digitized-native/optical-E395MeV-10000events.native-true5v.npz	d5ca94f623…c9002255
395 MeV/c aligned optical ROOT checksum sidecar	BM_ALIGNED_CAMPAIGN/results/optical-E395MeV-10000events.root.sha256	14d7627966…25252c9b
395 MeV/c aligned optical run log	BM_ALIGNED_CAMPAIGN/logs/optical-E395MeV-10000events.attempt-manual-3-1315416.log	ba05734440…af95e533
395 MeV/c aligned native digitizer summary	BM_ALIGNED_CAMPAIGN/digitized-native/optical-E395MeV-10000events.native-true5v.summary.json	cea803ad3b…533124ec
395 MeV/c aligned native completion sidecar	BM_ALIGNED_CAMPAIGN/digitized-native/optical-E395MeV-10000events.native-true5v.sha256	4f0cf88bc3…89795e6a
492 MeV/c aligned optical ROOT	BM_ALIGNED_CAMPAIGN/results/optical-E492MeV-10000events.root	dc06cf7083…1d961e7e
492 MeV/c aligned native NPZ	BM_ALIGNED_CAMPAIGN/digitized-native/optical-E492MeV-10000events.native-true5v.npz	b973cfbf8b…f123f803
492 MeV/c aligned optical ROOT checksum sidecar	BM_ALIGNED_CAMPAIGN/results/optical-E492MeV-10000events.root.sha256	64eb380e74…c872ea37
492 MeV/c aligned optical run log	BM_ALIGNED_CAMPAIGN/logs/optical-E492MeV-10000events.attempt-manual-4-1223121.log	e3c019e350…aaf46a6e
492 MeV/c aligned native digitizer summary	BM_ALIGNED_CAMPAIGN/digitized-native/optical-E492MeV-10000events.native-true5v.summary.json	40f14d725b…e4c5df14
492 MeV/c aligned native completion sidecar	BM_ALIGNED_CAMPAIGN/digitized-native/optical-E492MeV-10000events.native-true5v.sha256	d38ac39ab4…cd459f06
584 MeV/c aligned optical ROOT	BM_ALIGNED_CAMPAIGN/results/optical-E584MeV-10000events.root	97eda7bccf…cebd6d7f
584 MeV/c aligned native NPZ	BM_ALIGNED_CAMPAIGN/digitized-native/optical-E584MeV-10000events.native-true5v.npz	69bb3ef240…5e0ef3c6
584 MeV/c aligned optical ROOT checksum sidecar	BM_ALIGNED_CAMPAIGN/results/optical-E584MeV-10000events.root.sha256	2321421ef3…7afc82c6
584 MeV/c aligned optical run log	BM_ALIGNED_CAMPAIGN/logs/optical-E584MeV-10000events.attempt-manual-5-1320875.log	8d049c91ed…f7f1cbd0
584 MeV/c aligned native digitizer summary	BM_ALIGNED_CAMPAIGN/digitized-native/optical-E584MeV-10000events.native-true5v.summary.json	1e1a5a719d…55291223
584 MeV/c aligned native completion sidecar	BM_ALIGNED_CAMPAIGN/digitized-native/optical-E584MeV-10000events.native-true5v.sha256	ed838e3fe2…98f99f63
706 MeV/c aligned optical ROOT	BM_ALIGNED_CAMPAIGN/results/optical-E706MeV-10000events.root	96de7b0ac2…559b6b95
706 MeV/c aligned native NPZ	BM_ALIGNED_CAMPAIGN/digitized-native/optical-E706MeV-10000events.native-true5v.npz	761e3b9612…0d030365
706 MeV/c aligned optical ROOT checksum sidecar	BM_ALIGNED_CAMPAIGN/results/optical-E706MeV-10000events.root.sha256	1f262a9b9c…e0927080
706 MeV/c aligned optical run log	BM_ALIGNED_CAMPAIGN/logs/optical-E706MeV-10000events.attempt-manual-6-1199870.log	e590267017…f391c644
706 MeV/c aligned native digitizer summary	BM_ALIGNED_CAMPAIGN/digitized-native/optical-E706MeV-10000events.native-true5v.summary.json	bb3ddf368e…8c0a2496
706 MeV/c aligned native completion sidecar	BM_ALIGNED_CAMPAIGN/digitized-native/optical-E706MeV-10000events.native-true5v.sha256	be50886a10…b552f6a9
739 MeV/c aligned optical ROOT	BM_ALIGNED_CAMPAIGN/results/optical-E739MeV-10000events.root	8769a374a9…1c986532
739 MeV/c aligned native NPZ	BM_ALIGNED_CAMPAIGN/digitized-native/optical-E739MeV-10000events.native-true5v.npz	6200eb4ff4…af129789
739 MeV/c aligned optical ROOT checksum sidecar	BM_ALIGNED_CAMPAIGN/results/optical-E739MeV-10000events.root.sha256	957ba947f6…9412c0cd

Role	Source path	SHA-256
739 MeV/c aligned optical run log	BM_ALIGNED_CAMPAIGN/logs/optical-E739MeV-10000events.attempt-manual-7-1192977.log	dd5a573a57…f995fb81
739 MeV/c aligned native digitizer summary	BM_ALIGNED_CAMPAIGN/digitized-native/optical-E739MeV-10000events.native-true5v.summary.json	c3ec32d69b…61d93fa8
739 MeV/c aligned native completion sidecar	BM_ALIGNED_CAMPAIGN/digitized-native/optical-E739MeV-10000events.native-true5v.sha256	9f240f00d0…d478a6b3
796 MeV/c aligned optical ROOT	BM_ALIGNED_CAMPAIGN/results/optical-E796MeV-10000events.root	6392eb9289…a55348d9
796 MeV/c aligned native NPZ	BM_ALIGNED_CAMPAIGN/digitized-native/optical-E796MeV-10000events.native-true5v.npz	f683207c65…489b04e5
796 MeV/c aligned optical ROOT checksum sidecar	BM_ALIGNED_CAMPAIGN/results/optical-E796MeV-10000events.root.sha256	d5cd7a88d7…5ea38ee1
796 MeV/c aligned optical run log	BM_ALIGNED_CAMPAIGN/logs/optical-E796MeV-10000events.attempt-manual-8-1163343.log	ec833cbe51…85235ad2
796 MeV/c aligned native digitizer summary	BM_ALIGNED_CAMPAIGN/digitized-native/optical-E796MeV-10000events.native-true5v.summary.json	4a2d1e172d…ce68d238
796 MeV/c aligned native completion sidecar	BM_ALIGNED_CAMPAIGN/digitized-native/optical-E796MeV-10000events.native-true5v.sha256	081a5dde13…021f032c