

Scintillator–SiPM Detection of Cosmic-Ray Muons: Physics, Detector Architectures, Coincidence Techniques, Acceptance, and the Comparability of Reported Flux and Angular Measurements

Raja Babu¹, Nirmal Kumar¹, Malay Neerav¹, Ashim Roy^{1*}

¹University Department of Physics, Bhupendra Narayan Mandal University, Madhepura, Bihar, India-852113

*Corresponding author:

ORCID (A. Roy): 0000-0002-5622-4260

DOI: <https://doi.org/10.51244/IJRSI.2026.1309000068>

Received: 17 September 2026; Accepted: 22 September 2026; Published: 06 October 2026

ABSTRACT

Plastic scintillators read out by silicon photomultipliers (SiPMs) are now widely used in compact cosmic-ray muon telescopes, from desktop counters to square-metre muography hodoscopes. The measurements these instruments report, however, are difficult to compare: vertical intensity, integral flux, coincidence rate, and acceptance are not always distinguished, and the quantities that control them — geometric acceptance, plane efficiency, coincidence window, energy threshold, and environmental conditions — are reported inconsistently. This review is organized around that measurement chain rather than around individual instruments. It summarizes the muon physics that fixes what a surface telescope can measure, treats the scintillator and SiPM properties that set the photoelectron budget and the singles rate, classifies published detector architectures, and analyses coincidence background, acceptance, and efficiency quantitatively. Published zenith-angle exponents are then compared: acceptance-corrected tracking and pulse-charge measurements give $n = 2.00 \pm 0.04$ (stat) ± 0.14 (syst) and 1.90 ± 0.06 (stat) ± 0.10 (syst), whereas tilt scans of wide-acceptance counting telescopes fitted directly with a cosine-power law give 1.39 ± 0.01 and 1.44 ± 0.06 . A numerical acceptance calculation performed for this review, and clearly separated from literature-derived values, shows that for two square planes of width a separated by d the apparent exponent obtained from uncorrected tilt-scan rates falls from 1.99 at $d/a = 10$ to 1.62 at $d/a = 1$ and 1.28 at $d/a = 0.5$ when the true exponent is 2, and that forward folding the same model through the acceptance recovers the input exponent exactly. Finite acceptance can therefore account for a substantial fraction of the published spread without invoking atmospheric physics, although the geometries of the published telescopes differ from the idealization used here and no correction is derived for them. Its principal practical outcome is a minimum reporting standard for compact scintillator–SiPM muon measurements, set out in Section 15.5 and supported by an uncertainty framework and an assessment of the gaps that currently limit the comparability and the imaging performance of these detectors.

Keywords: cosmic-ray muons, plastic scintillator, silicon photomultiplier, coincidence counting, geometric acceptance, zenith-angle distribution, muography

INTRODUCTION

Scope and Rationale

Cosmic-ray muons are abundant penetrating charged particles at sea level and are available continuously, without an artificial radiation source or radiological control. For a scintillation detector they are therefore both a convenient commissioning signal and, in muography, the measurement probe itself. The combination of plastic scintillator with silicon photomultiplier (SiPM) readout has made such detectors compact, low-

voltage, and inexpensive enough to be built by school groups [1],[2] and, at the other extreme, to instrument tens of square metres in the field [3],[4],[5].

Two early experiments established the pattern that much of the later work has followed: the measurement of tunnel overburden from the attenuation of the muon flux [6] and the search for hidden chambers in the Second Pyramid of Giza with spark chambers placed in a chamber at its base [7]. Both compared a measured directional rate with an expected open-sky rate, and both therefore depended on knowing the detector acceptance, its efficiency, and its background. Modern reviews document how widely the technique has since spread [8],[9],[10],[11],[12], and the replacement of photomultiplier tubes (PMTs) by SiPMs has changed the engineering of these detectors substantially: bias voltages of a few tens of volts, insensitivity to magnetic fields, low power, and mechanical robustness, at the cost of a large thermally generated dark count rate and a gain that varies with temperature [13],[14],[15].

Gap addressed by This Review

Existing reviews are organized by application or by imaging algorithm [8],[9],[10],[11],[16]. What is missing is a critical treatment of the measurement chain of the compact scintillator–SiPM telescope itself, and in particular of why its published results do not compare well. The zenith-angle exponent illustrates the problem.

Acceptance-corrected measurements give values close to 2 [17],[18], and compilations of low-altitude results cluster between about 1.96 and 2.2 across a wide range of geomagnetic cut-offs [17]; yet two recent tilt scans with counting telescopes, fitted directly with a cosine-power law, report 1.39 ± 0.01 [19] and 1.44 ± 0.06 [20]. A discrepancy of this size exceeds the spread reported among low-altitude sites, and instrumental acceptance effects can contribute substantially to it; quantifying that contribution for an individual experiment requires its specific geometry and analysis to be modelled, which the literature has not done.

This review therefore (i) collects the muon physics that determines what a surface telescope can measure, (ii) treats the scintillator and SiPM properties that set the photoelectron budget and the singles rate, (iii) classifies detector architectures and their front-end and trigger implementations, (iv) analyses coincidence background, acceptance, and efficiency quantitatively, (v) compares published flux and angular results, and (vi) presents an original numerical acceptance calculation that quantifies the bias introduced by fitting an uncorrected cosine-power law. Throughout, results taken from the literature and calculations performed by the authors are separated explicitly, and Appendix A states the basis of every figure.

Contributions of This Review

Four elements are new relative to existing reviews. First, an explicit treatment of the measurement chain of compact scintillator–SiPM telescopes, in which intensity, integral flux, coincidence rate, acceptance and efficiency are kept distinct (Section 1.3) and their interdependence is made quantitative (Sections 8–10).

Second, an original numerical acceptance calculation, performed by the authors and based on no experimental data, showing that a direct cosine-power fit to uncorrected tilt-scan rates of two thin square planes returns an apparent exponent of about 1.62 at $d/a = 1$ and 1.28 at $d/a = 0.5$ when the true exponent is 2, and that forward folding the same model through the acceptance recovers $n = 2$ (Section 11.2). Third, a comparison of published systems that separates construction from operating parameters and identifies which reported quantities can and cannot be compared between instruments (Sections 12 and 15.4, Tables 6, 7 and 10). Fourth, a minimum reporting standard for compact-telescope measurements (Section 15.5), which is the principal practical outcome of the review.

Terminology, Method, and Conventions

Terms that are frequently conflated are used here as follows. Differential intensity is the flux per unit area, solid angle, time, and energy; intensity $I(\theta)$ is its integral above a stated energy or momentum threshold, in

$\text{m}^{-2} \text{s}^{-1} \text{sr}^{-1}$; integral flux J is the intensity integrated over the upper hemisphere through a horizontal surface, in $\text{m}^{-2} \text{s}^{-1}$; count rate and coincidence rate are detector observables in s^{-1} ; the geometric factor or acceptance G ($\text{m}^2 \text{sr}$) converts an intensity into a rate for a given geometry; and efficiency-corrected flux is a rate divided by acceptance and by trigger efficiency. A count rate normalized by detector area alone is none of these and is not comparable between instruments.

Sources were located on publisher platforms for the instrumentation and astroparticle journals in which this work appears — principally Nuclear Instruments and Methods A, the Journal of Instrumentation, IEEE Transactions on Nuclear Science, the Journal of Applied Physics, the Journal of Cosmology and Astroparticle Physics and Physical Review D — on the arXiv and NASA ADS records of specific detector platforms, and by following the reference lists of those platform papers; the Particle Data Group review supplied the reference values [21]. Coverage extends from the first muon-attenuation measurements of the 1950s to mid-2026, with the detector-platform material concentrated in the last decade, during which SiPM readout became common.

A source was retained when it reported a quantity that can be checked by a reader — a detector geometry, coincidence window, singles rate, threshold, efficiency, acceptance or fitted exponent — or when it provided a reference value or a formalism used here; sources offering only qualitative descriptions of a detector were not used for the comparison tables. Detectors with photomultiplier or gaseous readout are included only where they supply benchmarks that the scintillator–SiPM literature lacks. Extensive-air-shower physics, underground muon physics and scattering tomography are outside the scope except where they bear directly on a surface telescope. Bibliographic details were checked against publisher or arXiv records, and preprints, technical reports and manufacturer documents are identified as such in the reference list. No formal systematic-review protocol was applied and no claim of exhaustive coverage is made: the aim is a representative and verifiable rather than a complete sample. Where a source does not report a quantity, the tables state “not reported” rather than an estimate.

Cosmic-Ray Muons at the Surface

Production and Intensity Scale

Primary protons and helium nuclei interact in the upper atmosphere and produce charged pions and kaons whose decays supply almost all ground-level muons (Fig. 1). Production peaks near 15 km; a muon loses about 2 GeV by ionization on the way down and arrives with a mean energy near 4 GeV [21]. Survival is a relativistic effect: with a proper lifetime of $2.197 \mu\text{s}$, only time dilation allows the flight to the ground. The vertical intensity above 1 GeV/c is about $70 \text{ m}^{-2} \text{s}^{-1} \text{sr}^{-1}$, and the flux through a horizontal detector is of order $1 \text{ cm}^{-2} \text{min}^{-1}$ [21]. Table 1 collects the reference values used in this review.




Fig. 1. Production and propagation of cosmic-ray muons relevant to a surface telescope: meson production near 15 km, decay to muons, ionization loss and decay during transit, the increase of slant depth with zenith

angle, and the overburden that sets the effective energy threshold of an indoor detector. Original schematic, not to scale.

Table 1. Reference properties of cosmic-ray muons at the surface used in this review. The integral flux entry is computed from Eq. (2) by the authors using the vertical intensity and exponent in the rows above it.

Quantity	Value	Basis
Mass; mean lifetime	105.66 MeV/c ² ; 2.197 μs	[21]
Typical production altitude; ionization loss to ground	≈15 km; ≈2 GeV	[21]
Mean energy at sea level	≈4 GeV	[21]
Vertical intensity, $p > 1$ GeV/c	≈70 m ⁻² s ⁻¹ sr ⁻¹	[21]
Flux through a horizontal detector	of order 1 cm ⁻² min ⁻¹ (order-of-magnitude figure)	[21]
Integral flux from Eq. (2) for $n = 2$, $I_0 = 70$ m ⁻² s ⁻¹ sr ⁻¹	≈110 m ⁻² s ⁻¹ (0.66 cm ⁻² min ⁻¹)	Authors' calculation
Zenith exponent n	≈2 for $E_\mu \approx 3$ GeV; steeper at lower energy, flatter at higher	[21]
Charge ratio	1.2766 ± 0.0032 (stat) ± 0.0032 (syst) below 100 GeV/c; 1.268 world average above 10 GeV/c	[22],[23]
Energy loss in plastic (minimum ionizing)	≈2 MeV cm ² g ⁻¹	[21],[24]

Measured vertical intensities scatter around the reference value by considerably more than the quoted statistical errors, because they refer to different momentum thresholds, altitudes, and geomagnetic cut-offs. An acceptance-corrected tracking measurement at 160 m with a threshold near 0.11 GeV/c gives $(7.007 \pm 0.002 \text{ stat} \pm 0.526 \text{ syst}) \times 10^{-3} \text{ cm}^{-2} \text{ s}^{-1} \text{ sr}^{-1}$, and earlier low-altitude results compiled in the same work span about $(6.2\text{--}9.1) \times 10^{-3} \text{ cm}^{-2} \text{ s}^{-1} \text{ sr}^{-1}$ [17]; a measurement at 2268 m with a cut-off rigidity of 8.2 GV gives 101.2 ± 1.8 (stat) ± 5.5 (syst) m⁻² s⁻¹ sr⁻¹ [18]. Altitude and threshold, not detector quality, dominate this spread.

Angular Dependence and Integral Flux

Muons arriving at larger zenith angle traverse more atmosphere, lose more energy, and have more time to decay, so the intensity falls with θ . Near sea level it is conventionally written as

$$I(\theta) = I_0 \cos^n \theta \quad (1)$$

where I_0 is the vertical intensity and n an effective exponent. The value $n \approx 2$ is characteristic of muons near 3 GeV; the distribution steepens at lower energies and flattens towards a $\sec \theta$ dependence for energies well above the pion critical energy at $\theta < 70^\circ$ [21]. Equation (1) is a flat-atmosphere approximation and fails towards the horizon, where spherical-atmosphere treatments are required [25],[26]. Figure 2(b) shows Eq. (1) for the exponents discussed in this review. Integrating over the upper hemisphere gives the flux through a horizontal surface,

$$J = \int_{\Omega} I(\theta) \cos\theta \, d\Omega = 2\pi I_0 / (n + 2) \quad (2)$$

For $n = 2$ and $I_0 = 70 \, \text{m}^{-2} \, \text{s}^{-1} \, \text{sr}^{-1}$, Eq. (2) gives $J \approx 110 \, \text{m}^{-2} \, \text{s}^{-1}$ ($0.66 \, \text{cm}^{-2} \, \text{min}^{-1}$). The familiar figure of about $1 \, \text{cm}^{-2} \, \text{min}^{-1}$ [21] is therefore an order-of-magnitude statement, and comparisons at the 10–50% level require the threshold and the angular model to be specified.

Energy Spectrum and Charge Ratio

For the design of a telescope the differential spectrum matters mainly through the threshold: the rate is dominated by muons of a few GeV, while the muons that penetrate substantial overburden lie in the tail. The standard parametrization is

$$dI_{\mu}/dE_{\mu} = 0.14 (E_{\mu}/\text{GeV})^{-2.7} [(1 + 1.1E_{\mu}\cos\theta/115 \, \text{GeV})^{-1} + 0.054 (1 + 1.1E_{\mu}\cos\theta/850 \, \text{GeV})^{-1}] \quad (3)$$

with E_{μ} in GeV, the two terms describing pion and kaon parents [27]. It is valid for $E_{\mu} > 100/\cos\theta \, \text{GeV}$ and $\theta < 70^{\circ}$, because muon decay and the curvature of the atmosphere are neglected [26]; extensions replace $\cos\theta$ by an effective $\cos\theta$ and add the low-energy behaviour [25],[26],[28]. Figure 2(a) evaluates Eq. (3) within its stated domain. Equation (3) is therefore a high-energy reference, not a description of the few-GeV muons that dominate the counting rate of a surface telescope: expected rates near the vertical should be obtained from the integral intensity of Table 1 together with the angular law of Eq. (1), and a differential spectrum at few-GeV energies requires the low-energy parametrizations instead [25],[26]. World-average vertical spectra above 10 GeV are accurate to about 7–17% depending on momentum [23], which sets a floor on how well any surface telescope can be expected to agree with a model prediction.

The charge ratio $N(\mu^{+})/N(\mu^{-})$ is constant at low momentum and rises above about 100 GeV/c as kaon decays contribute; the most precise surface measurement gives 1.2766 ± 0.0032 (stat) ± 0.0032 (syst) below 100 GeV/c [22], consistent with the world average of earlier data [23]. A scintillator–SiPM telescope without a magnetic field cannot separate the charges, so the ratio enters only indirectly, through stopping-muon and decay studies.




Fig. 2. (a) Differential muon intensity from Eq. (3) for three values of $\cos\theta$, plotted only within the stated validity domain $E_{\mu} > 100/\cos\theta \, \text{GeV}$. (b) Relative intensity from Eq. (1) for the exponents reported in the literature and compared in Section 11. Analytical reconstructions from the cited equations; no experimental data are plotted (Appendix A).

Environmental Modulation

Two environmental effects modulate the rate at the percent level and are relevant to any measurement lasting longer than a day. The barometric effect reduces the rate when the overlying air mass increases; the temperature effect acts through the density of the production layer, changing the competition between meson decay and interaction. Both have been characterized with dedicated ground-based detectors, and the mass-weighted temperature method reproduces the observed seasonal variation of muon detector rates [29],[30]. For a compact telescope these effects are usually smaller than its systematic uncertainties, but they set the limit for long-term stability studies and for muography, where the reference and the measurement exposures may be separated by weeks.

Propagation Through Matter and Energy Threshold

Muon transport determines three things for a surface telescope: the signal it produces in the scintillator, the particles it can reject by absorption, and the effective threshold imposed by its surroundings. The mean energy loss is

$$\langle -dE/dx \rangle = a(E) + b(E) E \quad (4)$$

where $a(E)$ is the electronic (ionization) term and $b(E)$ the energy-scaled radiative term from bremsstrahlung, pair production, and photonuclear interactions [24]. At the few-GeV energies that dominate the surface flux the radiative term $b(E)E$ in Eq. (4) is negligible and muons are close to minimum ionizing, depositing about $2 \text{ MeV cm}^2 \text{ g}^{-1}$ [21]. In a 1 cm thick plastic scintillator of density near 1.02 g cm^{-3} this corresponds to roughly 2 MeV, which is the quantity that fixes the photoelectron budget in Section 4 and makes muon pulses separable from single-photoelectron dark pulses.

Angular deflection is described by the Highland–Lynch–Dahl form of multiple Coulomb scattering,

$$\theta_0 = (13.6 \text{ MeV} / \beta cp) z \sqrt{(x/X_0) [1 + 0.038 \ln(x z^2/(X_0 \beta^2))]} \quad (5)$$

with p the momentum, x/X_0 the thickness in radiation lengths, and z the charge [21]. For a few-GeV muon crossing a thin plastic plane the deflection is a fraction of a milliradian and is irrelevant to a counting telescope; it becomes the measured observable in scattering tomography [31],[32] and a resolution limit in transmission muography through thick objects.

Because $a(E)$ is slowly varying, the range in matter is approximately proportional to energy, and the opacity of the overburden — the integrated density along the line of sight — defines the minimum energy a detected muon must have had at the surface. This is why an indoor telescope does not measure the open-sky flux: the building above it removes the softest part of the spectrum and steepens the observed angular distribution, an effect explicitly treated as a systematic in tracking measurements by varying the modelled concrete thickness by 20% [17]. It is also why absorption checks that compare configurations with different plane separations are ambiguous: separation changes acceptance as well as path length (Section 8).

Scintillator Physics Relevant to Muon Detection

Sections 4 to 9 follow the chain of Fig. 3, from energy deposit to recorded coincidence. The treatment of scintillator and SiPM physics here is deliberately restricted to the properties that determine muon-telescope performance; comprehensive accounts of the materials and of the devices are available elsewhere [13],[15],[33].




Fig. 3. Detection chain of a scintillator–SiPM muon telescope, with the section of this review that treats each stage. One analogue channel per plane feeds the trigger; environmental logging and live-time recording are part of the chain rather than optional additions. Original schemati

Light Production and the Photoelectron Budget

Organic scintillators convert a small fraction of the deposited energy into optical photons through excitation of the polymer base and non-radiative transfer to primary and secondary dyes. The response to densely ionizing particles is quenched according to Birks' relation,

$$dL/dx = S (dE/dx) / [1 + kB (dE/dx)] \quad (6)$$

where S is the absolute yield and kB the quenching parameter [33]. For minimum-ionizing muons the quenching correction is negligible, so the light output is proportional to path length, and the pulse-height distribution of a plate carries the track-length information that can be used to infer the angular distribution without tilting the detector [18]. The quantity that matters for triggering is the photoelectron yield

$$N_{pe} = Y \cdot \Delta E \cdot \epsilon_{coll} \cdot PDE \quad (7)$$

where Y is the light yield per unit deposited energy, ΔE the deposit, ϵ_{coll} the light-collection efficiency, and PDE the photon detection efficiency of the SiPM at the emission wavelength. Five distinct quantities are involved here and are often conflated: the deposited energy, fixed for a minimum-ionizing muon by path length at about 2 MeV per centimetre of plastic; the scintillation photons produced, $Y\Delta E$, of order 10^4 per MeV; the collected photons reaching the photosensor, reduced by bulk attenuation, wrapping, fibre trapping and coupling losses; the photoelectrons N_{pe} that actually trigger microcells; and the pulse amplitude, which further depends on SiPM gain and on the front-end transfer function. Only the last is compared with the discriminator threshold, and the trigger efficiency is the probability that this amplitude exceeds the threshold for a given impact position, incidence angle and temperature. Because the losses multiply, a high-yield scintillator with poor collection can give a lower trigger efficiency than a modest scintillator with efficient fibre coupling, and quoting N_{pe} in photoelectrons, together with the impact position, coupling and overvoltage at which it was measured, provides a substantially more transferable basis for comparison between detectors than any of the preceding quantities.

Material and Geometry Choices

Commercial cast plastics differ mainly in light output, decay time, and bulk attenuation length (Table 2) [34]. For muon telescopes the relevant trade is between attenuation length, which governs response uniformity over a large plate, and decay time, which governs timing. Polyvinyltoluene-based material yields about 20% more light than polystyrene, and extruded polystyrene doped with PPO and POPOP reaches

roughly 80–85% of the light output of cast material at a small fraction of its cost [35]; this economic argument is what made square-metre muography planes practical [3],[36]. Light output also depends on temperature: coefficients below about $0.6\% \text{ }^{\circ}\text{C}^{-1}$ have been measured for several plastics and fibres over 0–40 $^{\circ}\text{C}$, with one PVT-based material showing no measurable dependence in that range [37]. This is smaller than the temperature dependence of SiPM gain (Section 5) but acts in the same direction and is not separable from it without an independent gain monitor.

Table 2. Properties of representative cast plastic scintillators from the manufacturer data sheet [34], with extruded polystyrene for comparison. Light-yield temperature coefficients measured for several plastics and fibres over 0–40 $^{\circ}\text{C}$ are below about $0.6\% \text{ }^{\circ}\text{C}^{-1}$, and one PVT-based material showed no measurable dependence in that range [37].

Material	Light output (% anthracene)	Photons per MeV	λ_{max} (nm)	Decay (ns)	Attenuation length (cm)	Characteristic
EJ-200 (PVT)	64	10 000	425	2.1	380	Long attenuation, large plates
EJ-204 (PVT)	68	10 400	408	1.8	160	Highest yield, fastest
EJ-208 (PVT)	60	9 200	435	3.3	400	Uniformity over large area
EJ-212 (PVT)	65	10 000	423	2.4	250	Thin sheets
Extruded polystyrene with PPO/POPOP	$\approx 80\text{--}85\%$ of cast equivalent	not reported	not reported	not reported	not reported	Low cost; enables square-metre planes [35]

For muon measurements the implication is that the scintillator governs the stability of the trigger efficiency more than its absolute value: a light-yield temperature coefficient of a few tenths of a percent per degree moves the pulse-height spectrum across a fixed threshold in the same direction as SiPM gain drift, so long angular scans and muography exposures need either temperature control or a monitored photoelectron scale. Light collection is the weakest link in large detectors. Direct coupling of a small tile to a single SiPM gives a compact and uniform detector but does not scale, because the fraction of light reaching a few- mm^2 sensor falls rapidly with plate area. Wavelength-shifting (WLS) fibres embedded in grooves or extrusion holes solve the scaling problem by absorbing blue light and re-emitting it into a trapped mode that is guided to a small photosensor; it is widely used in the large-area scintillator–SiPM systems surveyed in Section 6 [3],[4],[5],[36],[38]. The price is a reduced and position-dependent yield: photoelectron yields measured along and across extruded counters with embedded fibres vary systematically with the impact point, fibre diameter, and reflective coating [36], and bar-to-bar response differences of up to 50% have been reported for a muography tracker, requiring bars to be grouped by counting rate [38]. Response non-uniformity translates directly into a position-dependent efficiency, which enters the acceptance calculation of Section 9 and is rarely reported.

Silicon Photomultipliers in Muon Telescopes

Operating Principle and Parameters

A SiPM is a dense array of single-photon avalanche diodes, each with its own quenching resistor, biased a few volts above breakdown and connected in parallel (Fig. 4). Each fired cell delivers the same charge $Q_{\text{cell}} = C_{\text{cell}}\Delta V$, where $\Delta V = V_{\text{bias}} - V_{\text{bd}}$ is the overvoltage, so the output amplitude is proportional to the number of fired cells, and the gain, defined as the number of carriers produced per detected photon, is

$$G = Q_{\text{cell}} / e = C_{\text{cell}} \Delta V / e, \quad \Delta V = V_{\text{bias}} - V_{\text{bd}} \quad (8)$$

with C_{cell} the microcell capacitance and e the elementary charge [13],[39]. Overvoltage is the master control: gain rises linearly with it, PDE rises and saturates, while dark count rate, optical crosstalk, and afterpulsing all increase [13],[39],[40]. Table 3 summarizes how each parameter propagates into telescope performance.




Fig. 4. Silicon photomultiplier. (a) Geiger-mode microcells in parallel: each fired cell contributes the same charge, so the output is proportional to the number of fired cells, and the gain follows Eq. (8). (b) Qualitative dependence of gain, photon detection efficiency, dark count rate and optical crosstalk on overvoltage. Panel (b) is a schematic of the trends described in Refs. [13],[39],[40] and contains no measured data.

Table 3. SiPM parameters, their physical origin, their effect on a muon telescope, and the established control. Numerical entries are quoted only where a cited source reports them.

Parameter	Physical origin	Effect on a muon telescope	Control and reported values
Gain	Charge per fired cell divided by the elementary charge, Eq. (8)	Sets pulse amplitude relative to a fixed discriminator threshold	Bias setting; $1.7\text{--}3.0 \times 10^6$ measured for 1 mm^2 devices in a three-fold telescope [20]
Photon detection efficiency	Fill factor $\times$ quantum efficiency $\times$ trigger probability	Enters the photoelectron budget, Eq. (7); determines separation of the muon peak from noise	Overvoltage; spectral match to the scintillator or fibre [13],[39]
Dark count rate	Thermally generated carriers	Dominates the singles rate and therefore the accidental rate, Eqs. (9)–(10)	Threshold in photoelectrons: 51–75 kHz at 0.5 p.e. falling to 0.5–1.2 kHz at 3 p.e. for 1 mm^2 devices [20]
Optical crosstalk and afterpulsing	Secondary photons; trapped carriers	Distort the photoelectron scale and add correlated pulses within the window	Overvoltage; device technology [39],[40]
Saturation and dynamic range	Finite number of microcells	Not limiting for minimum-ionizing muons, unlike calorimetry	Cell count [14],[39]

Single-photon time resolution	Avalanche and transit statistics	Sub-nanosecond; not the limiting term for the coincidence window	Front-end walk and cabling dominate instead [15],[39],[41]
Breakdown-voltage temperature coefficient	Bandgap temperature dependence	Gain and efficiency drift with ambient temperature	Bias compensation: gain held to $\approx 2\%$ over 60 °C by a compensating driver [42]; dark-pulse feedback reduced a 10 °C variation from 22% to 0.78% [43]
Radiation damage	Bulk defects raising the dark rate	Negligible at the surface; relevant underground and in space-adjacent use	Device selection, cooling [44]

Three consequences are specific to muon telescopes. First, dark counts set the singles rate, and therefore the accidental coincidence rate of Section 8: dark count rates of 51–75 kHz at a 0.5 photoelectron threshold were measured for 1 mm² devices in a three-fold telescope, falling to 0.5–1.2 kHz when the threshold was raised to 3 photoelectrons [20]. Raising the threshold is thus the cheapest background reduction available, provided the photoelectron budget of Eq. (7) leaves the muon peak comfortably above it. Second, saturation is not a limitation here: a minimum-ionizing muon produces far fewer photoelectrons than the cell count of a typical device, so the response is linear, unlike in calorimetry [14],[39]. Third, the timing limitation is at system level rather than sensor level: SiPM single-photon time resolution is well below a nanosecond [15],[39], but the achievable resolving time is set by the whole chain — photoelectron statistics at threshold, amplitude-dependent walk of leading-edge discriminators, front-end bandwidth, cable-delay matching and TDC binning — so windows used in practice range from tens of nanoseconds in ASIC-based systems [45] to 200 ns in oscilloscope-triggered ones [20].

Temperature Dependence and Stability

The breakdown voltage rises roughly linearly with temperature, so at fixed bias the overvoltage, and with it the gain and the PDE, drift [13],[15],[40]. For a detector operated outdoors or in an unconditioned room this is the dominant source of efficiency drift, because a gain change moves the pulse-height spectrum across a fixed discriminator threshold. Two remedies are established: compensating the bias against a temperature measurement, which manufacturers implement in dedicated drivers and which holds the gain constant to within about 2% over a 60 °C range [42], and closed-loop compensation using the dark pulses of a reference SiPM, which reduced the dark-pulse amplitude variation over a 10 °C span from about 22% to 0.78% [43]. Radiation damage, which raises the dark count rate substantially [44], is not a limitation at the surface but matters for long underground or space-adjacent deployments.

Comparison with Photomultiplier Tubes

Against PMTs, SiPMs offer low bias voltage, compactness, magnetic-field immunity, mechanical robustness, and low power — the last being decisive for solar-powered field telescopes [3],[4]. PMTs retain advantages in single-channel active area, in dark rate per unit area, and in the absence of correlated noise, which is why the largest school networks and several reference angular measurements still use them [18],[46],[47]. For a compact telescope the SiPM disadvantage — a dark count rate of tens of kHz per mm² at low threshold — is entirely manageable by the combination of threshold setting and coincidence, which is the subject of Section 8.

Detector Architectures

A Taxonomy

Published scintillator–SiPM muon detectors fall into four classes that differ in what they can measure, not merely in size (Fig. 5, Table 4). A single counter gives a rate and no direction, and cannot separate muons

from dark counts, ambient radioactivity, and the soft component. A two-plane telescope adds a coincidence requirement and a cone-shaped acceptance, which makes rate measurements meaningful and permits tilt scans. A three- or four-fold stack suppresses accidentals by orders of magnitude and, importantly, allows the efficiency of an internal plane to be measured in situ from the tracks defined by the others. A tracking hodoscope with segmented orthogonal planes reconstructs a direction for every event, which removes the need for tilt scanning and is the prerequisite for imaging.




Fig. 5. Taxonomy of scintillator-SiPM muon detector architectures, with the readout-coupling and deployment axes that cut across it. Capability, cost and calibration burden all increase from left to right. Original schematic.

Table 4. Architectures of scintillator-SiPM muon detectors and their measurement capabilities.

Architecture	Typical fold	Direction information	Strengths	Limitations
Single counter	1	None	Simplest; highest rate per channel	Cannot separate muons from dark counts, ambient radiation or the soft component
Two-plane telescope	2	Acceptance cone; tilt scan	Simple, cheap, adequate for rate work	Wide acceptance biases uncorrected angular fits (Section 11.2); accidentals significant at high singles rate
Three- or four-fold stack	3–4	Acceptance cone; tilt scan	Accidental suppression by orders of magnitude; in-situ efficiency by Eq. (13)	Lower rate; more channels and cabling
Tracking hodoscope	Hit in every plane	Per-event direction	Imaging; no tilt scanning; acceptance computable per direction	High channel count; response non-uniformity; large calibration burden

Representative Implementations

At the compact end, the CosmicWatch design couples a $5 \times 5 \times 1$ cm plastic tile to a single SiPM and forms a coincidence by cabling two independent units together [1],[48]; the current revision adds on-board storage, environmental sensors, and a component cost below US\$100 [2], and the same hardware has been repurposed for laboratory β -source measurements after calibration against the muon flux [49]. A recent three-fold telescope couples $25 \times 2.5 \times 1.28$ cm bars with embedded 1 mm WLS fibres to 1 mm^2 multi-pixel photon counters and triggers on an oscilloscope [20]. Field and imaging systems scale this up with fibres: MU-RAY built 1 m^2 and $2 \times 2 \text{ m}^2$ X–Y planes from triangular extruded strips with WLS fibres, SiPMs, and a SPIROC front-end chosen for low power in a volcanic environment [3],[50]; SiRO uses six 1 m^2 bar layers grouped pairwise into three X–Y planes at variable separation [5]; the ScIDEP project operates two complementary trackers, one based on $61 \times 61 \times 2$ cm PVT plates with orthogonal 2 mm WLS fibres at 1 cm pitch read by SiPMs, and one with four layers of 36 polystyrene bars [38]. A different route to area is a monolithic plate with segmented light guides: MuGrid-v2 divides a $30 \times 30 \times 1$ cm plastic scintillator into 27 channels using a 3D-printed light-guide grid holding 1 mm wavelength-shifting fibres read by S13360-1350 MPPCs, reaching an intrinsic spatial resolution of 4.6 mm without cutting the scintillator into strips [45],[51]. The largest deployed scintillator–SiPM muon system is the buried AMIGA array, whose detectors comprise three 10 m^2 scintillation modules at 2.3 m depth read out through a dedicated SiPM front end [4]. School-network detectors with PMT readout and a four-channel discriminator and TDC card remain the benchmark for distributed operation [46],[52],[53].

Two design regularities emerge. Architecture follows the measurement, not the budget: counting telescopes remain two- or three-fold and compact, while anything intended to image is segmented and fibre-coupled. And channel count is a major driver of electronics cost and system complexity, particularly for segmented readout, which is why portable imaging systems are converging on modest numbers of large, segmented planes rather than on many small counters [45],[54],[55].

Front-End Electronics and Data Acquisition

The chain from SiPM to stored event (Fig. 3) is short but contains most of the parameters that later prove difficult to compare. Bias must be stable and, in uncontrolled environments, temperature-compensated (Section 5.2). The SiPM signal is then amplified and shaped; the shaping time trades pulse-pair resolution against signal-to-noise, and determines the walk of a leading-edge discriminator, which dominates the timing spread of a low-cost telescope [41]. Discrimination may be performed by a comparator with a fixed threshold, expressed either in millivolts or, more usefully, in photoelectrons, since the photoelectron scale provides a substantially more transferable basis for comparison between detectors.

Coincidence logic is implemented in three ways in the literature: discrete logic or a dedicated counting module, an ASIC such as SPIROC in large channel-count systems [3], and FPGA-based logic, which is now standard for hodoscopes because it provides programmable majority conditions, time stamping, and data formatting in one device [4],[38]. School-network cards occupy an intermediate position, with four discriminator inputs, time-to-digital conversion of about 1 ns, and programmable trigger logic [46]. A fourth route, now common in imaging systems, is a commercial timing ASIC: MuGrid-v2 reads 27 channels with a TOFPET2 evaluation kit whose per-channel discriminators and 30 ps time binning allow the threshold to be set at about 2.5 photoelectrons from the measured dark-count step structure and the coincidence window to be 15 ns [45],[56], about 13 times narrower than the oscilloscope-based case above. At the smallest scale, microcontrollers [2] and digital oscilloscopes [20] serve as both trigger and DAQ; the oscilloscope route is convenient but couples the measurement to instrument-specific settings, and a reported coincidence window of 200 ns is set largely by the readout arrangement rather than by the timing of the muon signal itself [20].

Waveform digitization is worth its cost when pulse height carries information, as in track-length methods [18], in threshold optimization, and in diagnosing correlated noise; simple scaler readout suffices for rate work provided live time and dead time are recorded. Two practices are still uncommon and should not be: logging temperature and pressure alongside the counts, which the latest desktop design does by default [2],

and recording the singles rate of each channel continuously, without which the accidental rate of Eq. (10) cannot be reconstructed after the fact.

Coincidence Detection and Background Rejection

Background Taxonomy

Four backgrounds must be distinguished because they scale differently. SiPM dark counts are uncorrelated between channels and are controlled by threshold and by coincidence fold. Ambient gamma and beta radiation deposits energy in single planes and is similarly uncorrelated between separated planes, but it is not reduced by raising the threshold as effectively as dark counts, since its deposits can approach those of a muon in a thin plate.

The soft electromagnetic component of the shower is correlated with cosmic rays but is absorbed in a few centimetres of material, so it is rejected by requiring penetration through the stack. Air showers produce genuinely correlated multi-particle events, which are background for a single-track telescope and signal for a distributed array [46],[53].

Accidental Coincidences

Conventions differ between authors, so they are fixed here. The resolving time τ is defined such that two discriminated pulses are counted as coincident when their arrival times differ by less than τ , which corresponds to a total overlap window of width 2τ ; where a window width $w = 2\tau$ is quoted instead, the factors of two below are absorbed into w . With this convention, for two independent channels with singles rates R_1 and R_2 , chance coincidences occur at [41]

$$R_{\text{acc}}^{(2)} = 2\tau R_1 R_2 \quad (9)$$

and for n independent channels, all pulses being required to overlap and $\tau R_i \ll 1$, the generalization consistent with Eq. (9) at $n = 2$ is

$$R_{\text{acc}}^{(n)} = n \tau^{n-1} \prod_{i=1}^n R_i \quad (10)$$

Adding a plane therefore changes the accidental rate by a factor of about $[(n+1)/n]\tau R \sim 1.5 \times 10^{-5}$ at $\tau = 10$ ns and $R = 1$ kHz — while the genuine coincidence rate R_{true} , given by Eq. (11), falls only by the efficiency of the added plane and by the change in acceptance. Figure 6 evaluates Eqs. (9) and (10) over the window range used in practice. The published numbers make the point concretely: with singles rates of 1.2, 0.5, and 0.92 kHz at a 3 photoelectron threshold and $\tau = 200$ ns, the three-fold accidental rate is 6.6×10^{-5} Hz against a measured vertical three-fold rate of 0.152 Hz [20].

Applying Eq. (9) to the two outer channels of the same detector gives about 0.44 Hz, which exceeds the muon signal; at those singles rates a two-fold telescope with that window would have been dominated by accidentals. Tracking systems with far more channels use much shorter effective windows for the same reason, for example a 100 ns four-layer trigger in a large RPC stack [17].




Fig. 6. Coincidence background. (a) Definition of the resolving time. (b) Two-fold accidental rate from Eq. (9) for equal singles rates of 0.1, 1 and 10 kHz. (c) n-fold accidental rate from Eq. (10) at $R = 1$ kHz. The horizontal line is the genuine coincidence rate $R_{\text{true}} = 0.09$ Hz computed from Eq. (11) for two thin 5×5 cm planes 5 cm apart; dash-dotted lines mark windows of 100 ns and 200 ns used in published systems. Analytical calculation (Appendix A).

Window and Threshold Optimization, and the Limits of the Formulae

A relativistic muon crosses 1 m in 3.3 ns, so true time differences in a laboratory telescope are a few nanoseconds. The window must nevertheless cover the measured timing spread from discriminator walk, front-end jitter, and cable mismatch, and should be set from the measured time-difference distribution rather than adopted from another instrument. Threshold and window are not independent: raising the threshold lowers R_i quadratically or faster in Eq. (9) but also reduces the plane efficiency, so the optimum is the pair that minimizes the accidental-to-true ratio at acceptable efficiency, which requires a threshold scan with the coincidence rate measured simultaneously [20].

Equations (9) and (10) describe uncorrelated pulses only. Electronic pickup common to several channels, particles produced in the material near the telescope, and air-shower events are not described by them, and in every case the correlated component makes the true background larger than the formula predicts. The reliable procedure is therefore to measure the accidental rate directly by delayed coincidence, inserting a delay much longer than τ into one channel, and to use Eqs. (9) and (10) as a cross-check [41]. This measurement is inexpensive and is reported in only a minority of compact-telescope papers.

Acceptance and Efficiency

Geometric Factor

For a telescope whose axis is tilted by θ_0 from the vertical, the true coincidence rate is

$$R(\theta_0) = \varepsilon I_0 \int_{\Omega} A_{\text{eff}}(\omega) \cos^n \theta(\omega) d\Omega \equiv \varepsilon I_0 G_n(\theta_0) \quad (11)$$

where $A_{\text{eff}}(\omega)$ is the projected overlap area of the planes for arrival direction ω , $\theta(\omega)$ the zenith angle of that direction, and ε the combined trigger efficiency; G_n is the acceptance weighted by the angular model. Only in the limit of planes small compared with their separation does this factorize into the familiar expression

$$G_n(\theta_0) \approx (A^2/d^2) \cos^n \theta_0, \quad d \gg a \quad (12)$$

and A^2/d^2 is an upper bound on the isotropic geometric factor [57]; exact expressions for circular and rectangular elements are available [57],[58]. Compact telescopes are far from this limit. A numerical evaluation performed for this review gives, for square planes, an isotropic geometric factor equal to 0.63, 0.86 and 0.96 of A^2/d^2 at $d/a = 1, 2$ and 4 ; Eq. (12) therefore overestimates the geometric factor by about 60% at $d/a = 1$, by 16% at $d/a = 2$ and by 4% at $d/a = 4$, and is accurate to about 5% only for $d \gtrsim 4a$. This criterion is derived here for thin, identical, parallel square planes and should not be transferred without recalculation to elongated bars, thick plates, unequal planes or misaligned geometries. Evaluating Eq. (11) with the $\cos^2 \theta$ weighting for two thin 5×5 cm planes 5 cm apart gives $G_2(0) \approx 13 \text{ cm}^2 \text{ sr}$, a factor 0.52 of the $25 \text{ cm}^2 \text{ sr}$ obtained from Eq. (12), the additional reduction coming from the angular weighting that suppresses inclined directions retained by the thin-telescope formula; with $I_0 = 70 \text{ m}^{-2} \text{ s}^{-1} \text{ sr}^{-1}$ and $\varepsilon = 1$ the expected rate is 0.09 Hz.

Equation (11) factorizes the rate into an efficiency, an intensity and a purely geometric acceptance only under restrictive conditions: the efficiency must be independent of impact position, incidence angle, deposited energy and time. None holds exactly. Fibre-coupled planes have position-dependent yield (Section 4.2); inclined tracks deposit more energy and are therefore detected with slightly higher probability than vertical ones at the same threshold, partially compensating the geometric fall-off; and gain drift makes the efficiency time-dependent. Where these dependences are weak they can be absorbed into an effective constant ε and the factorization is adequate, as it is for thin, directly coupled tiles operated at stable temperature with a threshold well below the muon peak. Where they are not, the efficiency belongs inside the integral, $\varepsilon = \varepsilon(\omega, r, \Delta E, T)$, and the acceptance must be evaluated by Monte Carlo using a measured response map rather than by geometry alone. The distinction matters for reporting, because a single quoted efficiency implicitly asserts that the factorization holds.

Because the angular half-width scales as $\arctan(a/d)$, rate and angular resolution are controlled by the same ratio. Enlarging the planes at fixed separation raises the rate but widens the acceptance; scaling a and d together preserves the angular response and increases the rate in proportion to area. Thickness, plate-to-plate misalignment, and position-dependent efficiency (Section 4.2) all perturb A_{eff} and are the reasons that Monte Carlo evaluation, typically in GEANT4 [59], is preferred over closed-form expressions once a detector departs from two identical parallel rectangles [17],[18].

Efficiency

Plane efficiency must be measured, not assumed, because it depends on light yield, impact position, SiPM overvoltage, and threshold. In a three-plane stack whose middle plane covers the acceptance of the outer pair,

$$\varepsilon_2 = N_{123} / N_{13} \quad (13)$$

where the counts are corrected for accidentals; outer planes require a fourth plane or an interchange of positions. Tracking systems generalize this by estimating layer efficiencies from events triggered by other layers, and additionally map them pixel by pixel, which also exposes correlated inefficiencies [17]. Reported per-plane efficiencies for a well-characterized scintillator tracking system reach about 97% with a hit resolution of 2.5 mm [60]. Efficiency must be re-measured whenever threshold, bias, or operating temperature changes, and this is the point at which the SiPM properties of Section 5 enter the flux measurement quantitatively. The implication for flux work is direct: without an in-situ measurement the efficiency term in Eq. (14) is an assumption, and a 5% error in it propagates in full into the quoted intensity, comparable to the entire systematic budget of a careful tracking measurement [17].

Flux Measurements and Their Normalization

Converting counts into an intensity requires every factor identified above. With N_c coincidences recorded in live time T_{live} ,

$$I_0 = [N_c - R_{\text{acc}}^{(n)} T_{\text{live}}] / [\epsilon_{\text{tot}} G_n(\theta_0) T_{\text{live}}] \quad (14)$$

where ϵ_{tot} is the product of the plane efficiencies, G_n is evaluated for the actual geometry and an assumed exponent, and the accidental term is measured or computed from Eq. (10). If n is itself being measured, G_n and n are determined iteratively. Tracking analyses use the same factorization, separating trigger, selection, and data-acquisition efficiencies and inserting a simulated acceptance [17]. The statistical uncertainty is $1/\sqrt{N_c}$, so a 1% measurement needs about 10^4 coincidences, roughly 30 h at the 0.09 Hz computed in Section 9.1 for a compact pair — a useful reminder that for compact telescopes the statistical uncertainty often remains the dominant term until exposures of many days are accumulated, after which the systematic terms of Section 14 take over.

Three normalization errors recur in the compact-telescope literature. Dividing the coincidence rate by plane area alone produces a quantity in $\text{cm}^{-2} \text{s}^{-1}$ that depends on the telescope geometry and is not the integral flux of Eq. (2); it is reported in this form in at least one recent study [20] and cannot be compared with I_0 or J . Comparing a measured rate with the “ $1 \text{ cm}^{-2} \text{ min}^{-1}$ ” rule of thumb ignores both the momentum threshold to which that figure refers and the $\cos^n\theta$ weighting of the acceptance. And quoting I_0 without the threshold makes the result incomparable with published values that span $(6.2\text{--}9.1) \times 10^{-3} \text{ cm}^{-2} \text{s}^{-1} \text{sr}^{-1}$ largely because their thresholds and sites differ [17],[18]. Table 5 collects measurements that do report the necessary information.

Table 5. Reported vertical intensity and zenith-angle exponent, with the method used to obtain them. "Not reported" indicates that the quantity does not appear in the cited source.

Exponent n	Vertical intensity I_0	Threshold and site	Method	Source
≈ 2 (reference)	$\approx 70 \text{ m}^{-2} \text{s}^{-1} \text{sr}^{-1}$	$p > 1 \text{ GeV}/c$, sea level	Reference values	[21]
1.95 ± 0.08	not reported	$E_\mu > 1 \text{ GeV}$, $\theta = 0\text{--}70^\circ$	Geant4 simulation validated against laboratory detector data	[61]
2.00 ± 0.04 (stat) ± 0.14 (syst)	$(7.007 \pm 0.002 \pm 0.526) \times 10^{-3} \text{ cm}^{-2} \text{s}^{-1} \text{sr}^{-1}$	$p \gtrsim 0.11 \text{ GeV}/c$, 160 m a.s.l.	12-layer RPC tracking stack; fit weighted by simulated acceptance	[17]
1.90 ± 0.06 (stat) ± 0.10 (syst)	$101.2 \pm 1.8 \pm 5.5 \text{ m}^{-2} \text{s}^{-1} \text{sr}^{-1}$	2268 m a.s.l., cut-off 8.2 GV	Stationary thick scintillator bar; pulse-charge (track-length) method with simulation	[18]
1.39 ± 0.01	not reported	276 m a.s.l., fit for $\theta \leq 75^\circ$	QuarkNet counters (PMT); tilt scan, direct fit of Eq. (1) to rates	[19]
1.44 ± 0.06	not reported (rate normalized by plane area only)	not reported	Three-fold MPPC telescope; tilt scan in 10° steps, direct fit of Eq. (1)	[20]
not fitted	not reported	266 m a.s.l.	Desktop SiPM pair on a rotating frame; compared with the $\cos^2\theta$ expectation	[48]

Angular Distribution and the Acceptance Bias

Measurement Approaches

Three methods are used. Tilt scans rotate a counting telescope and record the rate at a series of zenith angles: two desktop units 52 mm apart were rotated at 266 m altitude with about 10 h per angle and compared with the $\cos^2\theta$ expectation [48]; QuarkNet counters were stepped in about 15° increments with about 3.1 days per step [19]; a three-fold MPPC telescope was stepped in 10° increments with 1 h per angle and a longer vertical run [20]. Event-by-event tracking instead reconstructs θ for each muon and fits the observed distribution folded with a simulated acceptance [17]. A third approach exploits Eq. (6): in a thick stationary bar the pulse charge measures track length and therefore inclination, so the angular distribution and the vertical intensity can be extracted without moving the detector [18]. Reported exponents are compared in Table 5.

11.2 Acceptance bias: a calculation performed for this review

The following calculation is the authors' own and uses no experimental data. Equation (11) was integrated numerically by two-dimensional quadrature over the direction tangents for two thin square planes of width a separated by d , with both traversal senses accepted, an input intensity following Eq. (1) with $n_{\text{true}} = 2$, unit efficiency, and no background. The resulting rates were normalized to the vertical and fitted with $\cos^m\theta_0$ over $0-75^\circ$ in 5° steps with equal weights, exactly as a tilt scan is usually analysed; the fitted exponent m is the apparent exponent n_{app} , to be distinguished from the input value n_{true} . Figure 7 shows the outcome. The procedure is documented here so that it can be reproduced. The integration variables are the direction tangents t_x and t_y of the arrival direction in the telescope frame, each running over $\pm a/d$, beyond which the projected overlap area vanishes; the integrand is the overlap area of the two planes multiplied by $\cos^4\psi \, dt_x \, dt_y$, which accounts for the projection of the plane onto the arrival direction and for the solid-angle element, and by the intensity of Eq. (1) evaluated at the zenith angle of that direction for the given tilt. Quadrature is by the trapezoidal rule on an $N \times N$ grid with $N = 401$ for the fitted values and $N = 301$ for the curves of Fig. 7(b), and the fits are ordinary non-linear least squares with equal weights. Two checks were run. Grid refinement changes the fitted exponent by less than 10^{-4} between $N = 101$ and $N = 801$, so the values quoted below are converged at the precision reported. In the thin-telescope limit the computed isotropic geometric factor approaches the analytic upper bound A^2/d^2 of Ref. [57] from below, reaching 0.993 of it at $d/a = 10$ and 0.998 at $d/a = 20$, which is the expected behaviour and serves as an independent validation of the integrator.

The apparent exponent falls monotonically as the telescope becomes more compact: $m = 1.99$ at $d/a = 10$, 1.94 at 3, 1.86 at 2, 1.62 at 1, 1.28 at 0.5, and 0.93 at 0.2, against the input value of 2. The mechanism is averaging: at each tilt the telescope accepts a range of zenith angles of half-width $\arctan(a/d)$, so the measured curve is the true law convolved with the acceptance, which flattens it. The bias is not a small correction for compact geometries; at $d/a = 1$, typical of desktop pairs, it removes about 0.4 from the exponent.

Three properties of the bias matter for interpretation. It depends on the input exponent: at $d/a = 1$ the naive fit returns 1.30, 1.62, and 1.92 for inputs of 1.5, 2.0, and 2.5, so it cannot be removed by a single additive correction. It depends on the fit range, since the flattening grows with tilt angle; restricting the fit to small angles reduces the bias at the cost of statistics. And it is fully recoverable: when the same model is forward folded through Eq. (11) for the actual geometry and fitted to the same simulated rates, the input exponent is recovered to four significant figures for every case tested (Fig. 7(d)). Forward folding, not post-hoc correction, is therefore the appropriate analysis.

The limitations of the calculation should be stated. The planes are treated as thin, identical, perfectly aligned rectangles with uniform efficiency; thickness, misalignment, and the position-dependent response of fibre-coupled planes (Section 4.2) are not included, and each would modify the effective acceptance. The model is therefore an idealization, and its numerical values apply to the square thin-plane case alone; they should not be transferred to elongated bars, thick plates, unequal or misaligned planes, or fibre-coupled planes with

non-uniform response without repeating the calculation for that geometry. Azimuthal effects, backgrounds, and threshold effects are excluded by construction. The model therefore establishes the size and the sign of the geometric contribution for a well-defined idealization; it does not reproduce the geometry of any published experiment, and the values quoted here must not be used as corrections to the results of Refs. [19] or [20], whose telescopes are elongated rather than square and whose analyses differ in fit range and background treatment.




Fig. 7. Acceptance bias in tilt-scan measurements: a calculation performed by the authors for two thin square planes with a true $\cos^2 \theta$ intensity. (a) Geometry. (b) Normalized coincidence rate versus tilt for several separation-to-width ratios; the dashed curve is the true intensity. (c) Apparent exponent from a direct fit of Eq. (1) over 0–75°, with the band marking exponents reported by counting telescopes, shown for scale only because their geometries differ from this model. (d) Recovered exponent for three input values: the direct fit is biased low, while forward folding through Eq. (11) recovers the input. Original numerical calculation (Appendix A).

Separating Instrumental from Physical Effects

Four instrumental effects push a naively fitted exponent below the physical value: acceptance averaging, as quantified above; the directional ambiguity of an untimed telescope, which must be distinguished from both acceptance and physical flux: without time-of-flight information a two-plane telescope registers a coincidence for a track crossing its axis in either sense, so near $\theta_0 = 90^\circ$ its acceptance consists of two opposing lobes that both view the near-horizontal sky, and the recorded rate corresponds to the sum of two symmetric acceptance regions rather than to the intensity along one direction. Because the upward-going muon flux at the surface is negligible, this reflects the directional acceptance of the instrument, which is symmetric under reversal of the track direction when no time-of-flight information is available, rather than any upward-going component of the physical flux; the choice of how to normalize the near-horizontal point is therefore an analysis convention; one desktop measurement handled it by halving the horizontal point [48]. Resolving the ambiguity event by event is a system-level timing requirement, well below the 3.3 ns flight time across 1 m, rather than a demand on the sensor (Section 5.1); angle-independent backgrounds from accidentals, ambient radioactivity, and the soft component, which add a constant term and which at least one analysis absorbed into an empirical offset parameter [20]; and unmodelled overburden, whose azimuthal asymmetry distorts the scan. Only one effect pushes the other way: a lower energy threshold

steepens the true distribution [21]. It follows that instrumental effects can contribute significantly to exponents well below 2 obtained from wide-acceptance counting telescopes, and that their contribution has to be quantified for the specific detector before an atmospheric interpretation is adopted. How much of any individual result is instrumental cannot be stated without repeating the calculation for that telescope's own geometry, fit range and background model, which the published information does not always permit; what can be said is that the geometric term alone has the right sign and a magnitude comparable to the discrepancy. The agreement of acceptance-corrected measurements with $n \approx 2$ across very different technologies [17],[18],[61] remains the stronger evidence about the physics.

Comparative Synthesis of Published Scintillator–SiPM Detectors

Tables 6 and 7 compare published systems, separating what they are built from and how they are operated from what they report. Entries are taken only from the cited sources; “not reported” means the quantity does not appear there, and no value has been inferred. Two things are immediately visible. First, construction parameters are reported well: scintillator type, dimensions, photosensor, layer count, and readout coupling are available for essentially every system, which makes design comparison possible. Second, operating and result parameters are reported poorly and inconsistently. Of the eleven systems in Tables 6 and 7, the threshold, the coincidence window and the timing performance each appear for two, a per-channel singles rate for one, and a quantitative measurement result for five; a measured per-plane efficiency appears for one system [45], and no entry states the geometric acceptance, the overburden above the detector or the environmental conditions during the reported run. The summary rows of Tables 6 and 7 give the counts column by column. The consequence is that the systems cannot be ranked on any single performance figure, and this review deliberately does not attempt such a ranking.

Table 6. Construction parameters of published scintillator–SiPM muon detectors, with a PMT-based school-network detector as a benchmark. "Not reported" means the quantity is not stated in the cited sources; no values have been inferred. The final row counts, for each column, the systems for which the quantity is reported.

System	Scintillator and dimensions	Photosensor	Layers / planes	Separation	Coupling	Ref.
CosmicWatch (v2)	Plastic tile, $5 \times 5 \times 1$ cm	One SiPM per unit	Two independent units	52 mm in the reported angular scan	Direct	[1],[48]
CosmicWatch v3X	Plastic scintillator, dimensions not reported	SiPM	One or two units	not reported	Direct	[2]
Three-fold MPPC telescope	Bars, $25 \times 2.5 \times 1.28$ cm	Hamamatsu S10362-11, 1 mm^2	Four bars, three in the trigger	not reported	≈ 1 mm WLS fibre	[20]
MU-RAY	Triangular extruded strips; planes of 1 m^2 and $2 \times 2 \text{ m}^2$	SiPM per strip	X–Y planes (three in the 1 m^2 prototype)	not reported	WLS fibre	[3],[50]
SiRO	Scintillator bars, six 1 m^2 layers	SiPM per bar	Six layers grouped into three X–Y	Variable	WLS fibre / direct, per bar	[5]

			planes			
ScIDEP plates	PVT plates, 61 × 61 × 2 cm	Hamamatsu S14160 per fibre	Two plates, orthogonal fibres in each	not reported	2 mm WLS fibres at 1 cm pitch	[38]
ScIDEP μ 36	Polystyrene bars, 36 per layer	SiPM per bar	Four layers	Top pair separated from bottom pair; value not reported	WLS fibre	[38]
Cosmic Muon Explorer	Two plastic scintillator panels, 70 × 70 mm	Hamamatsu S13360-2050VE	Two planes	not reported	WLS fibre	[62]
MuGrid-v2	Monolithic plastic scintillator, 30 × 30 × 1 cm, segmented into 27 light-guide channels	Hamamatsu S13360-1350 MPPC per channel	Single plane; three planes in the efficiency test	15 cm in the efficiency test	Kuraray Y-11(200) 1 mm WLS fibres in a 3D-printed light-guide grid	[45],[51]
AMIGA (buried)	Scintillation modules, 10 m ² each	SiPM front end	Three modules per detector, 2.3 m underground	n.a.	WLS fibre	[4]
QuarkNet (PMT benchmark)	Scintillator counters, dimensions not reported	PMT	Four counters	User-set	not reported	[46],[52]
Reporting completeness	9 of 11	11 of 11	11 of 11	5 of 11	10 of 11	—

Table 7. Operating parameters and reported results for the systems of Table 6. The final row counts, for each column, the systems for which the quantity is reported; the prevalence of "not reported" in the middle columns is itself a finding of this review (Section 12). Geometric acceptance, overburden and environmental conditions are reported by none of these systems and therefore have no column here.

System	Fold	Threshold	Window	Timing	Singles rate	Reported result	Application
CosmicWatch (v2)	2	not reported	not reported	not reported	not reported	Angular scan compared with $\cos^2\theta$	Education, portable measurements

CosmicWatch v3X	1–2	not reported	not reported	not reported	not reported	not reported	Education, outreach
Three-fold MPPC telescope	3	3 p.e.	200 ns	not reported	0.5–1.2 kHz at 3 p.e.	0.152 Hz vertical; $n = 1.44 \pm 0.06$	Angular distribution
MU-RAY	Multi-plane tracking	not reported	not reported	not reported	not reported	not reported	Volcano muography
SiRO	Track through three planes	not reported	not reported	not reported	not reported	not reported	Muography; laboratory and salt mine
ScIDEP plates	X–Y hit pattern	not reported	not reported	not reported	not reported	Open-sky zenith distribution over ≈ 10 days	Pyramid muography
ScIDEP $\mu 36$	Hit in each of four layers	not reported	not reported	not reported	not reported	not reported	Pyramid muography
Cosmic Muon Explorer	2	not reported	not reported	not reported	not reported	not reported	Portable flux measurements, outreach
MuGrid-v2	3-fold in the efficiency test; single-plane reconstruction in operation	≈ 2.5 p.e.	15 ns	TOFPET2 ASIC, 30 ps TDC binning	not reported	Layer efficiencies $86.12 \pm 0.09\%$, $90.83 \pm 0.11\%$, $87.72 \pm 0.08\%$; spatial resolution 9.2 mm (centroid), 4.6 mm time-weighted; 769 718 events in a 48 h open-sky run	Building muography
AMIGA (buried)	Buried muon counting	not reported	not reported	not reported	not reported	not reported	Air-shower muon density

QuarkNet (PMT benchmark)	2–4, programmable	not reported	not reported	TDC $\approx$ 1 ns	not reported	$n = 1.39 \pm$ 0.01 from a tilt scan [19]	School networks, flux and shower studies
Reporting completeness	11 of 11	2 of 11	2 of 11	2 of 11	1 of 11	5 of 11	11 of 11

What can be compared meaningfully are parameters that are physically equivalent across designs. Photoelectron yield per minimum-ionizing particle is one, and it separates directly coupled tiles from fibre-coupled planes; it is reported systematically only in dedicated counter studies [36]. Accidental-to-true ratio at the stated window and singles rate is another, and where the necessary numbers are given it can be recomputed from Eqs. (9) and (10) [20]. Angular acceptance, expressed as the ratio d/a or as a solid angle, is a third, and it is the parameter that Section 11 shows to be decisive for angular measurements; it can be reconstructed from geometry for most systems. Absolute count rate, by contrast, is not comparable between detectors of different acceptance, and neither is any “flux” normalized by area alone.

Four design conclusions follow from the comparison itself rather than from any single system. First, coincidence fold tracks the singles rate at the operating threshold rather than the budget: detectors whose channels run at kilohertz singles rates with windows of order 100 ns use three-fold logic or require a hit in every layer [20],[38], whereas a nanosecond-scale window permits fewer planes for the same accidental rate [45]. Second, readout coupling is dictated by area: direct tile–SiPM coupling is used at the 10^1 – 10^2 cm² scale [1],[2], while the systems larger than roughly 10^3 cm² in Table 6 predominantly use wavelength-shifting fibres and accept the response non-uniformity that follows [3],[4],[5],[38],[45]. Third, among the systems surveyed here it is channel count rather than area that sets the electronics: beyond a few tens of channels, the surveyed systems use ASIC readout with per-channel discrimination and time stamping [3],[45],[56]. Fourth, the better-characterized systems share one practice, in-situ efficiency measurement by the three-plane method of Eq. (13) with the layer order exchanged, which yields per-layer efficiencies of 86–91% for a monolithic segmented plane [45] and about 97% for an extruded-bar tracker [60]; a system without such a measurement cannot support an efficiency-corrected flux. What the comparison cannot support is a ranking: the operating parameters needed to compare performance are missing for most systems, and a composite score would therefore encode reporting practice rather than detector quality.

Viewed across the classes of Table 4, the trade-off is consistent. Compact two-fold counters achieve high rate per channel and low cost but give biased angular information unless forward folded. Three- and four-fold stacks buy background suppression and in-situ efficiency measurement for a modest loss of rate. Fibre-coupled segmented planes buy per-event direction and imaging capability at the cost of channel count, response non-uniformity, and a much larger calibration burden [3],[5],[38]. The largest deployed systems accept that burden because the science requires directional information, and they invest correspondingly in front-end integration and monitoring [4].

Muography: How Detector Characteristics Set Imaging Performance

Muography is treated here only insofar as it is driven by the detector properties analysed above; comprehensive surveys of the applications exist [8],[9],[10],[11]. In transmission muography the observable is the ratio

$$T(\theta,\varphi) = I(\theta,\varphi) / I_0(\theta,\varphi) \quad (15)$$

of the flux measured through the object to an open-sky reference in the same direction, from which the opacity along that line of sight is inferred. Because both terms are measured with the same instrument, acceptance and efficiency cancel to first order provided they are stable between the two exposures; what does not cancel is drift. This reverses the priorities of Sections 9 and 10: absolute acceptance matters less

than its reproducibility, and SiPM gain stability (Section 5.2) becomes a leading systematic. In scattering tomography, where the observable is the deflection distribution of muons crossing the volume between two trackers [31],[32],[63], angular resolution rather than absolute normalization is decisive, and the Coulomb scattering of Eq. (5) sets the physical limit.




Fig. 8. Transmission muography and the detector properties that limit it. The measured flux through the object is compared with an open-sky reference in the same direction (Eq. (15)); acceptance and efficiency cancel to first order only if they are stable between exposures. Original schematic.

Table 8. Muography performance requirements and the detector characteristics that set them.

Imaging requirement	Detector characteristic that sets it	Consequence for design	Example systems
Angular resolution	Segmentation pitch and plane separation	Segmented X–Y planes rather than plain counters; separation traded against acceptance	[3],[5],[38]
Statistics through high opacity	Acceptance × exposure time; spectrum tail above the required energy	Large area and long campaigns; open-sky reference of comparable duration	[38],[64]
Background rejection	Coincidence fold; requirement of a hit in every layer; accidental rate, Eqs. (9)–(10)	Fake events mimic low-opacity paths, so fold and window are imaging parameters	[20],[38]

Stability between exposures	SiPM gain and efficiency drift; environmental conditions	Acceptance cancels in Eq. (15) only if stable; compensation and logging required	[2],[42],[43]
Field robustness and power	Sensor technology and front-end design	Low-voltage, low-power SiPM readout replaced PMTs in field instruments	[3],[4],[50],[54]
Scattering-mode resolution	Position resolution of the trackers; multiple scattering, Eq. (5)	Sets the achievable material discrimination	[31],[32],[60],[63]

The mapping from detector parameter to imaging performance is direct (Table 8). Angular resolution follows from segmentation pitch and plane separation, which is why imaging systems are built as segmented X–Y planes rather than as plain counters [3],[5],[38]. Statistics follow from acceptance times exposure: because the useful flux through a large opacity is a small tail of the spectrum, area and time are the only levers, and field campaigns run for weeks to months, with open-sky reference distributions measured over comparable periods [38]. Background suppression matters more than in laboratory counting, because a small fake rate mimics a low-opacity path; this is the argument for higher coincidence fold and for requiring hits in every layer [38]. Finally, robustness and power dictate the technology in the field, which is the reason SiPM readout displaced PMTs in volcano and borehole instruments [3],[50],[54].

The same considerations appear in adjacent applications with different weightings: archaeological and civil-structure surveys emphasize portability and stability over long unattended runs [7],[65],[66],[67],[68]; geoscience and mining emphasize ruggedness and underground operation [5],[69]; nuclear-material and cargo inspection emphasize position resolution and short measurement times, and use scintillator trackers with dedicated spectrometry to add momentum information [60],[70]. In every case the limiting quantities are those of Sections 8–10: acceptance, efficiency stability, and background.

14. CALIBRATION, UNCERTAINTY, AND REPRODUCIBILITY

A compact telescope has a small statistical uncertainty only after long exposures, so its systematic budget is usually dominant. Table 9 lists the terms, their propagation path, and the mitigation established in the literature. Three deserve emphasis because they are frequently omitted.

Acceptance is a model, not a measured quantity, and its uncertainty must be propagated. The thin-telescope formula of Eq. (12) overestimates the geometric factor by about 60% at $d/a = 1$, 16% at $d/a = 2$ and 4% at $d/a = 4$ (Section 9.1), so a compact telescope with d/a of order unity that uses Eq. (12) without recalculation carries a normalization error of tens of percent, while the error becomes small once $d \gtrsim 4a$. Efficiency is a measurement that must be repeated: because it depends on threshold, bias, and temperature, an efficiency determined once at commissioning does not describe a detector operated later at a different temperature, and gain compensation reduces but does not eliminate the drift [42],[43]. Overburden is an unmeasured absorber: it modifies both normalization and the angular distribution, and the only published treatment that quantifies it does so by varying the modelled concrete thickness by 20% in simulation [17].

Table 9. Sources of uncertainty in scintillator–SiPM muon measurements, their propagation path, and mitigation. Magnitudes are quoted only where a cited source reports them.

Source	How it enters	Treatment and reported	Mitigation
--------	---------------	------------------------	------------

		magnitude	
Counting statistics	$1/\sqrt{N_c}$ on every rate	≈ 30 h for 1% at 0.09 Hz (authors' calculation, Section 10)	Longer exposure; larger acceptance
Geometric acceptance	Multiplies the rate in Eqs. (11), (14)	Eq. (12) overestimates the geometric factor by $\approx 60\%$ at $d/a = 1$ (ratio 0.63) and by $\approx 4\%$ at $d/a = 4$ (authors' calculation)	Numerical integration or Monte Carlo for the real geometry [57],[59],[71],[72]
Alignment and plane thickness	Perturb $A_{\text{eff}}(\omega)$	Rarely reported	Survey tolerance; include thickness in the acceptance model
Plane efficiency	Multiplies the rate; depends on threshold, bias, position	$\approx 97\%$ per plane with 2.5 mm hit resolution reported for one tracker [60]	Measure by Eq. (13); re-measure after any change
Threshold definition	Sets efficiency and singles rate	Reported in photoelectrons in some studies [20], in millivolts or not at all in others	Quote in photoelectrons from the single-photoelectron spectrum [40]
SiPM temperature drift	Gain shift moves the spectrum across the threshold	22% dark-pulse amplitude change per 10 °C uncompensated; 0.78% compensated [43]	Bias compensation [42]; temperature logging [2]
Accidental coincidences	Additive background, Eqs. (9)–(10)	6.6×10^{-5} Hz against 0.152 Hz for one three-fold telescope [20]	Delayed-coincidence measurement [41]
Correlated background	Not described by Eqs. (9)–(10)	Rarely quantified in compact-telescope reports	Measure directly; separate air-shower events [46],[53]
Dead time	Reduces live time in Eq. (14)	Explicitly included in one tracking analysis [17]	Record live time in hardware
Atmospheric pressure and temperature	Percent-level rate modulation	Characterized by dedicated networks [29],[30]	Log and correct for long runs

Overburden	Raises the effective threshold; distorts the angular distribution	Treated by varying the modelled concrete thickness by 20% [17]	Model the building; state it in reports
Angular model in the acceptance	G_n depends on the assumed n	Iterate, or forward fold (Section 11.2)	Forward folding rather than post-hoc correction

Open simulation chains make the acceptance term auditable: cosmic-ray generators, transport codes and detector simulations are now combined to predict directional rates for a specific geometry and site, and are increasingly published alongside the measurement [45],[59],[71],[72].

Reproducibility requires calibration procedures that are stated in enough detail to be repeated: the photoelectron scale from the single-photoelectron dark spectrum [20],[40], the threshold in photoelectrons rather than millivolts, the accidental rate by delayed coincidence [41], plane efficiency by Eq. (13), and the acceptance by numerical integration or Monte Carlo with the geometry stated [59]. None of these is expensive; their systematic absence, rather than their difficulty, is what limits comparability.

Critical Synthesis

What is established

Several results are supported by independent measurements with different technologies and are not currently in dispute. The sea-level intensity scale and its dependence on threshold, altitude, and geomagnetic cut-off are known to within the spread documented in compilations [17],[21],[23]. The angular distribution of few-GeV muons is close to $\cos^2\theta$, and acceptance-corrected measurements with entirely different technologies agree on this [17],[18],[61]. Plastic scintillator coupled to SiPMs provides ample signal for minimum-ionizing muons, and a threshold of a few photoelectrons combined with two- or three-fold coincidence reduces uncorrelated background to negligible levels [20],[36]. SiPM gain drift with temperature is understood and correctable [13],[42],[43]. Fibre-coupled segmented planes scale to square metres and support imaging [3],[4],[5],[38].

Where results disagree, and why

The disagreements are concentrated in derived quantities rather than in raw rates. Zenith exponents from wide-acceptance counting telescopes lie well below those from acceptance-corrected measurements, and the calculation performed for this review, which uses no experimental data, shows that for an idealized square-plane geometry the geometric contribution alone has the right sign and a magnitude able to account for a substantial fraction of the gap (Section 11.2); angle-independent background and the directional ambiguity of untimed telescopes act in the same direction. Apportioning the gap for a specific experiment would require its own geometry and analysis to be modelled. Absolute intensities differ between sites for physical reasons — threshold, altitude, cut-off rigidity — that are quantified in the literature [17],[18], and between instruments for instrumental reasons that usually are not: unstated acceptance models, unmeasured efficiencies, and undocumented overburden. The distinction matters because it identifies which differences carry information. A rate difference between two telescopes may be physical — altitude, geomagnetic cut-off, overburden and energy threshold all change the true intensity — or instrumental, through acceptance, efficiency, dead time and background; it cannot be assigned to either category unless those quantities are reported, which is why Table 10 lists them as prerequisites for comparison. A difference in exponent between a low-threshold and a high-threshold measurement made with the same geometry and analysis, by contrast, is interpretable as physical.

Comparisons that require additional information

Three comparisons are difficult to interpret from the information usually published and should be treated with caution unless the additional quantities noted below are available. Count rates between telescopes of different geometry, because acceptance is rarely given in a form that permits normalization. “Fluxes” normalized by area alone, because they are geometry-dependent quantities with flux-like units [20]. And exponents fitted to uncorrected tilt-scan data, whether compared with each other or with the canonical value, because the acceptance bias is comparable to the reported uncertainties, which are usually statistical only: the quoted 1.39 ± 0.01 [19] carries a statistical error nearly forty times smaller than the bias identified here.

What can and cannot be compared

Table 10 separates the quantities that different groups can meaningfully compare from those that cannot be compared as published, and states what would be required in each case. The pattern is systematic: quantities defined at the detector level — count rates, area-normalized fluxes, thresholds in millivolts, exponents from uncorrected fits — are instrument-specific, whereas quantities defined after acceptance and efficiency correction, or expressed in intrinsic units such as photoelectrons, transfer between instruments.

Table 10. Quantities reported by compact muon telescopes: which can be compared directly between instruments, and what information the remainder would require. Section references indicate where each requirement is derived.

Reported quantity	Directly comparable?	Reason	Information required for comparison
Coincidence count rate	No	Depends on acceptance, efficiency, threshold and site	Plane sizes, separation and thickness; efficiency; threshold in p.e.; altitude, cut-off and overburden (Sections 9, 10)
Rate normalized by plane area only	No	Retains the geometry dependence; neither an intensity nor an integral flux	Acceptance G_n and efficiency, i.e. Eqs. (11) and (14) (Section 10)
Vertical intensity I_0	Conditionally	Physically dependent on threshold, altitude and geomagnetic cut-off	Momentum or energy threshold, altitude, cut-off rigidity, acceptance method (Sections 2.1, 10)
Integral flux J	Conditionally	Follows from I_0 and the angular model through Eq. (2)	Angular model and exponent n assumed, together with the threshold and site of the underlying I_0 (Section 2.2)
Exponent n from an uncorrected tilt-scan fit	No	Acceptance bias is comparable to or larger than the quoted uncertainties	d/a, fit range and background model; preferably a forward-folded refit (Section 11.2)
Exponent n from tracking or track-length methods	Conditionally	Acceptance enters the fit explicitly, but the result still depends on threshold and site	Acceptance model and its validation, energy or momentum threshold, altitude, cut-off rigidity and overburden

			(Sections 11.1, 14)
Accidental coincidence rate	Conditionally	Recomputable from Eqs. (9) and (10) only if the resolving-time convention is stated, and checkable by delayed coincidence	Resolving time τ with its convention and per-channel singles rate at the operating threshold (Section 8)
Discriminator threshold	Only when quoted in photoelectrons	A voltage threshold depends on gain, overvoltage and front-end response	Photoelectron scale from the single-photoelectron spectrum, overvoltage, temperature (Sections 5, 14)
Plane or trigger efficiency	Conditionally	Depends on threshold, impact position and temperature	Measurement method (Eq. (13)), threshold, temperature, position dependence (Section 9.2)
Photoelectron yield per minimum-ionizing particle	Yes, with the impact position stated	Intrinsic to the scintillator, coupling and photosensor	Impact position, coupling scheme, overvoltage (Section 4.1)
Spatial or angular resolution	Partially	Definitions differ: intrinsic, reconstructed, or effective after tracking	Method of determination and reference used (Sections 12, 13)

A Minimum Reporting Standard for Compact Scintillator–SiPM Muon Measurements

The requirements scattered through Sections 8 to 11 are collected here as a checklist, which is the principal practical outcome of this review. A measurement that reports these items can be compared with others; one that omits them generally cannot, whatever its internal precision.

- **Geometry.** Plane dimensions and thickness, separation, alignment tolerance, and the resulting ratio d/a or acceptance solid angle.
- **Scintillator and coupling.** Material and wrapping, direct or wavelength-shifting-fibre coupling, and photoelectron yield per minimum-ionizing particle with the impact position at which it was measured.
- **Photodetector.** SiPM type and active area, overvoltage, operating temperature, and the temperature-compensation scheme if any.
- **Trigger.** Threshold in photoelectrons together with the calibration used to set it, coincidence fold, resolving time τ with its convention, and the measured time-difference distribution.
- **Background.** Per-channel singles rates at the operating threshold, the accidental rate measured by delayed coincidence, and any correlated background identified.
- **Efficiency.** Per-plane values, the method used (Eq. (13) or equivalent), and any position or angle dependence.
- **Exposure.** Live time, dead time, and the number of recorded coincidences.
- **Site.** Altitude, geomagnetic cut-off rigidity, and a description of the overburden above the detector.
- **Environment.** Logged temperature and pressure for runs longer than a day, with any correction applied.

- **Analysis.** Acceptance model and how it was computed, angular model assumed, fit range, and whether the model was forward folded through the acceptance (Section 11.2).

Research Gaps and Challenges

The gaps that follow from the preceding sections are less about device performance than about measurement practice and scaling. There is no common reporting convention, and Tables 6 and 7 show how much this costs: the field cannot presently assemble a coherent dataset from its own literature. The effective energy threshold of a surface telescope is undefined in most reports, being set by an unmodelled building, which makes threshold-dependent quantities such as the exponent ambiguous. Acceptance modelling is often reduced to a closed-form expression outside its domain of validity. Long-term stability of SiPM telescopes in uncontrolled environments is sparsely documented, with most characterizations spanning hours to weeks, although stability over months is precisely what muography and environmental monitoring require.

Response uniformity of large fibre-coupled planes remains a calibration burden rather than a solved problem [36],[38]. Directional discrimination between upward- and downward-going tracks is a system-level timing requirement rather than a sensor one: the sensors already resolve well below a nanosecond [39], but the flight time across 1 m is 3.3 ns and the achievable resolution is set by photoelectron statistics at threshold, discriminator walk, cable matching and TDC binning, which remains demanding for low-cost chains even though commercial timing ASICs now offer tens of picoseconds of binning [45],[56]. Finally, a substantial fraction of the compact-telescope literature consists of preprints, educational reports and technical documentation, whose data are rarely published in reusable form.

Future Directions

Several developments would address these gaps directly. On the sensor side, devices with lower dark count rate and higher photon detection efficiency relax the threshold–efficiency trade of Section 8, and digital SiPMs, which integrate cell-level discrimination and time stamping, can reduce dependence on parts of the analogue timing chain that currently limit resolution [13],[73]. On the scintillator side, cheaper extruded material and improved fibre coupling continue to lower the cost per square metre [35],[36]. On the electronics side, FPGA-based triggering with per-channel time stamping makes majority logic, delayed-coincidence background measurement, and environmental logging routine rather than optional [4],[38]. Distributed networks of inexpensive, environmentally logged units are the natural route to studying atmospheric and long-term effects with compact hardware [2],[46],[53],[54].

Reconstruction is moving towards learned methods in muography [16], which raises the calibration bar rather than lowering it, since such methods require well-characterized detector responses. The most valuable near-term development, however, is mundane: openly published detector models and raw datasets, of the kind already provided with at least one desktop design [49], would allow acceptance and threshold to be recomputed consistently across groups and would convert the present collection of incomparable measurements into a usable body of data.

CONCLUSIONS

Scintillator–SiPM muon telescopes are simple to build and difficult to compare. Their physics inputs are well established: the sea-level intensity and its angular distribution are known within a spread set by threshold, altitude, and geomagnetic cut-off, and the relevant scintillator and SiPM properties are characterized in the device literature. What limits the field is the measurement chain between them. Dark counts at the operating threshold determine whether two planes suffice or a third is required, and the accidental rate follows steeply from window and fold. The thin-telescope formula overestimates the geometric factor by about 60% when $d \approx a$, the exact value being 0.63 of the approximation, and by only about 4% once $d \gtrsim 4a$.

Fitting an uncorrected cosine-power law to tilt-scan rates lowers the apparent zenith exponent from 2 to about 1.6 at $d/a = 1$ and 1.3 at $d/a = 0.5$, as the calculation presented here shows, while forward folding the

same model through the acceptance recovers the input exactly. Instrumental acceptance effects can therefore contribute substantially to the published spread in zenith exponents; quantifying their contribution for an individual experiment requires that experiment's geometry and analysis to be modelled. Numerical acceptance modelling, measured plane efficiencies, delayed-coincidence background estimates, forward-folded angular analysis, and reporting to the standard set out in Section 15.5, the principal practical outcome of this review, would make compact-telescope results comparable, and would put the same instruments on a sound footing for muography, where stability rather than absolute normalization is the limiting quantity.

APPENDIX A. PROVENANCE OF FIGURES

Table 11 states, for every figure, whether it is an original schematic, an analytical reconstruction from a cited equation, or a numerical calculation performed by the authors, together with the parameters used. No figure in this review reproduces digitized experimental data.

Table 11. Provenance of every figure: original schematics, analytical reconstructions from cited equations, and calculations performed by the authors.

Figure	Classification	Basis	Parameters and sources
Fig. 1	Original schematic	Concept diagram of production and propagation	Altitudes and energy losses as in Table 1 [21]; not to scale
Fig. 2(a)	Analytical reconstruction	Eq. (3)	$\cos \theta = 1, 0.5, 0.25$; plotted only for $E_\mu > 100/\cos \theta$ GeV [26],[27]
Fig. 2(b)	Analytical reconstruction	Eq. (1)	$n = 1.39, 1.44, 1.90, 2.00$ from Table 5 [17],[18],[19],[20],[21]
Fig. 3	Original schematic	Block diagram of the detection chain	No quantitative content
Fig. 4(a)	Original schematic	Microcell array and Eq. (8)	No quantitative content
Fig. 4(b)	Schematic of qualitative trends	Overvoltage dependences described in [13],[39],[40]	Arbitrary units; contains no measured data
Fig. 5	Original schematic	Architecture taxonomy of Table 4	No quantitative content
Fig. 6(a)	Conceptual illustration	Coincidence within $\pm\tau$	Pulse shapes and times arbitrary [41]
Fig. 6(b), (c)	Analytical calculation	Eqs. (9), (10); reference line from Eq. (11)	$R = 0.1, 1, 10$ kHz; $n = 2-4$; $\tau = 1-1000$ ns; reference line for 5×5 cm planes at $d = 5$ cm, $I_0 = 70 \text{ m}^{-2} \text{ s}^{-1} \text{ sr}^{-1}$, $\varepsilon = 1$; windows marked from [17],[20]
Fig. 7	Original numerical calculation by the authors	Eq. (11) with Eq. (1) by two-dimensional quadrature; least-squares fits of $\cos^m \theta_0$ and of	Thin square planes; $d/a = 0.2-10$; input $n = 1.5, 2.0, 2.5$; $\theta_0 = 0-75^\circ$ in 5° steps, equal weights; both traversal senses accepted;

		the forward-folded model	acceptance formalism after [57]
Fig. 8	Original schematic	Transmission muography, Eq. (15)	No quantitative content

REFERENCES

1. S. N. Axani, K. Frankiewicz, and J. M. Conrad, "The CosmicWatch Desktop Muon Detector: a self-contained, pocket sized particle detector," J. Instrum., vol. 13, art. P03019, 2018.
2. Desktop Muon Detector (v3X)," J. Instrum., vol. 20, art. P10040, 2025.
3. G. Ambrosi, F. Ambrosino, R. Battiston, et al., "The MU-RAY project: volcano radiography with cosmic-ray muons," Nucl. Instrum. Methods Phys. Res. A, vol. 628, no. 1, pp. 120–123, 2011.
4. Aab et al. (Pierre Auger Collaboration), "Design, upgrade and characterization of the silicon photomultiplier front-end for the AMIGA detector at the Pierre Auger Observatory," J. Instrum., vol. 16, art. P01026, 2021.
5. Niculescu-Oglintzanu, D. Stanca, A. Bălăceanu, M. Dobre, A. Gherghel-Lascu, A. Saftoiu, R. Smău, and C. Vancea, "SiRO, a scintillator-based hodoscope for muography applications," J. Appl. Phys., vol. 136, no. 17, art. 174501, 2024.
6. E. P. George, "Cosmic rays measure overburden of tunnel," Commonwealth Engineer, vol. 42, pp. 455–457, 1955.
7. L. W. Alvarez, J. A. Anderson, F. El Bedwei, et al., "Search for hidden chambers in the pyramids," Science, vol. 167, no. 3919, pp. 832–839, 1970.
8. S. Procureur, "Muon imaging: principles, technologies and applications," Nucl. Instrum. Methods Phys. Res. A, vol. 878, pp. 169–179, 2018.
9. G. Bonomi, P. Checchia, M. D'Errico, D. Pagano, and G. Saracino, "Applications of cosmic-ray muons," Prog. Part. Nucl. Phys., vol. 112, art. 103768, 2020.
10. Lechmann, D. Mair, A. Ariga, T. Ariga, A. Ereditato, R. Nishiyama, C. Pistillo, P. Scampoli, F. Schlunegger, and M. Vladymyrov, "Muon tomography in geoscientific research – a guide to best practice," Earth-Sci. Rev., vol. 222, art. 103842, 2021.
11. H. K. M. Tanaka, C. Bozza, A. Bross, et al., "Muography," Nat. Rev. Methods Primers, vol. 3, art. 88, 2023.
12. L. Bonechi, R. D'Alessandro, and A. Giammanco, "Atmospheric muons as an imaging tool," Rev. Phys., vol. 5, art. 100038, 2020.
13. C. Piemonte and A. Gola, "Overview on the main parameters and technology of modern silicon photomultipliers," Nucl. Instrum. Methods Phys. Res. A, vol. 926, pp. 2–15, 2019.
14. F. Simon, "Silicon photomultipliers in particle and nuclear physics," Nucl. Instrum. Methods Phys. Res. A, vol. 926, pp. 85–100, 2019.
15. S. Gundacker and A. Heering, "The silicon photomultiplier: fundamentals and applications of a modern solid-state photon detector," Phys. Med. Biol., vol. 65, no. 17, art. 17TR01, 2020.
16. S. Luo, C. Feng, G. Zeng, et al., "Image reconstruction techniques in muography: a review of algorithms and physical principles," J. Appl. Phys., vol. 138, no. 4, art. 041101, 2025.
17. S. Pethuraj, V. M. Datar, G. Majumder, N. K. Mondal, K. C. Ravindran, and B. Satyanarayana, "Measurement of cosmic muon angular distribution and vertical integrated flux by 2 m × 2 m RPC stack at IICHEP-Madurai," J. Cosmol. Astropart. Phys., vol. 2017, no. 09, art. 021, 2017.
18. B. Olmos Yáñez and A. A. Aguilar-Arevalo, "A method to measure the integral vertical intensity and angular distribution of atmospheric muons with a stationary plastic scintillator bar detector," arXiv:1912.05462 (preprint), 2019.
19. R. C. Venterea and U. Ekka, "An analysis of muon flux from angle variation of the QuarkNet cosmic ray detector," arXiv:2306.13689 (preprint), 2023.
20. S. Manithottathil, A. Gupta, M. Kumar, and N. Poonthottathil, "Characterization of an MPPC-based scintillator telescope and measurement of cosmic muon angular distribution," arXiv:2602.16297 (preprint), 2026.

21. S. Navas et al. (Particle Data Group), "Review of particle physics," *Phys. Rev. D*, vol. 110, no. 3, art. 030001, 2024.
22. V. Khachatryan et al. (CMS Collaboration), "Measurement of the charge ratio of atmospheric muons with the CMS detector," *Phys. Lett. B*, vol. 692, no. 2, pp. 83–104, 2010.
23. T. Hebbeker and C. Timmermans, "A compilation of high energy atmospheric muon data at sea level," arXiv:hep-ph/0102042 (preprint), 2001.
24. D. E. Groom, N. V. Mokhov, and S. I. Striganov, "Muon stopping power and range tables, 10 MeV–100 TeV," *At. Data Nucl. Data Tables*, vol. 78, no. 2, pp. 183–356, 2001.
25. P. Shukla and S. Sankrith, "Energy and angular distributions of atmospheric muons at the Earth," arXiv:1606.06907 (preprint), 2016.
26. M. Guan, M.-C. Chu, J. Cao, K.-B. Luk, and C. Yang, "A parametrization of the cosmic-ray muon flux at sea-level," arXiv:1509.06176 (preprint), 2015.
27. T. K. Gaisser, *Cosmic Rays and Particle Physics*. Cambridge, U.K.: Cambridge Univ. Press, 1990.
28. D. Chirkin, "Fluxes of atmospheric leptons at 600 GeV – 60 TeV," arXiv:hep-ph/0407078 (preprint), 2004.
29. R. R. S. de Mendonça, C. R. Braga, E. Echer, et al., "The temperature effect in secondary cosmic rays (muons) observed at the ground: analysis of the Global Muon Detector Network data," *Astrophys. J.*, vol. 830, no. 2, art. 88, 2016.
30. Riádigos, D. García-Castro, D. González-Díaz, and V. Pérez-Muñuzuri, "Atmospheric temperature effect in secondary cosmic rays observed with a 2 m² ground-based tRPC detector," *Earth Space Sci.*, vol. 7, art. e2020EA001131, 2020.
31. N. Borozdin, G. E. Hogan, C. Morris, W. C. Priedhorsky, A. Saunders, L. J. Schultz, and M. E. Teasdale, "Radiographic imaging with cosmic-ray muons," *Nature*, vol. 422, p. 277, 2003.
32. J. Schultz, K. N. Borozdin, J. J. Gomez, et al., "Image reconstruction and material Z discrimination via cosmic ray muon radiography," *Nucl. Instrum. Methods Phys. Res. A*, vol. 519, no. 3, pp. 687–694, 2004.
33. J. B. Birks, *The Theory and Practice of Scintillation Counting*. Oxford, U.K.: Pergamon Press, 1964.
34. Eljen Technology, "General purpose plastic scintillators EJ-200, EJ-204, EJ-208, EJ-212," manufacturer data sheet, accessed 2026.
35. Pla-Dalmau, A. D. Bross, and K. L. Mellott, "Low-cost extruded plastic scintillator," *Nucl. Instrum. Methods Phys. Res. A*, vol. 466, no. 3, pp. 482–491, 2001.
36. Artikov, V. Baranov, G. C. Blazey, et al., "Photoelectron yields of scintillation counters with embedded wavelength-shifting fibers read out with silicon photomultipliers," *Nucl. Instrum. Methods Phys. Res. A*, vol. 890, pp. 84–95, 2018.
37. Peralta, "Temperature dependence of plastic scintillators," *Nucl. Instrum. Methods Phys. Res. A*, vol. 883, pp. 20–23, 2018.
38. Gamal, S. Aly, D. Geeraerts, et al., "The ScIDEP muon radiography project at the Egyptian Pyramid of Khafre," *J. Appl. Phys.*, vol. 138, no. 4, art. 044901, 2025.
39. F. Acerbi and S. Gundacker, "Understanding and simulating SiPMs," *Nucl. Instrum. Methods Phys. Res. A*, vol. 926, pp. 16–35, 2019.
40. R. Klanner, "Characterisation of SiPMs," *Nucl. Instrum. Methods Phys. Res. A*, vol. 926, pp. 36–56, 2019.
41. G. F. Knoll, *Radiation Detection and Measurement*, 4th ed. Hoboken, NJ, USA: Wiley, 2010.
42. Hamamatsu Photonics, "How does temperature affect the gain of an SiPM?," manufacturer technical note, accessed 2026.
43. F. Licciulli and C. Marzocca, "An active compensation system for the temperature dependence of SiPM gain," *IEEE Trans. Nucl. Sci.*, vol. 62, no. 1, pp. 228–235, 2015.
44. E. Garutti and Y. Musienko, "Radiation damage of SiPMs," *Nucl. Instrum. Methods Phys. Res. A*, vol. 926, pp. 69–84, 2019.
45. Y. Tao, Y. Ning, Y. Yuan, et al., "MuGrid-v2: a novel scintillator detector for multidisciplinary applications," *J. Appl. Phys.*, vol. 138, no. 2, art. 024501, 2025.

46. S. Hansen, T. Jordan, T. Kiper, D. Claes, G. Snow, H. Berns, T. H. Burnett, R. Gran, and R. J. Wilkes, "Low-cost data acquisition card for school-network cosmic ray detectors," *IEEE Trans. Nucl. Sci.*, vol. 51, no. 3, pp. 926–930, 2004.
47. H. Sogarwal and P. Shukla, "Measurement of atmospheric muon angular distribution using a portable setup of liquid scintillator bars," *arXiv:2202.01445* (preprint), 2022.
48. S. N. Axani, "The physics behind the CosmicWatch Desktop Muon Detectors," *arXiv:1908.00146* (preprint), 2019.
49. J. Masias, F. Delgado, L. Coll, A. Pérez, J. Bazo, and A. Gago, "Using a portable muon detector for radioactive source measurements and identification," *Nucl. Instrum. Methods Phys. Res. A*, vol. 944, art. 162588, 2019.
50. Anastasio, F. Ambrosino, D. Basta, et al., "The MU-RAY experiment: an application of SiPM technology to the understanding of volcanic phenomena," *Nucl. Instrum. Methods Phys. Res. A*, vol. 718, pp. 134–137, 2013.
51. H. Yang, G. Luo, T. Yu, et al., "MuGrid: a scintillator detector towards cosmic muon absorption imaging," *Nucl. Instrum. Methods Phys. Res. A*, vol. 1042, art. 167402, 2022.
52. QuarkNet, "Cosmic Ray Muon Detector User's Manual," Fermi National Accelerator Laboratory, Batavia, IL, USA, technical report, accessed 2026.
53. outreach event for astroparticle physics," in *Proc. 35th Int. Cosmic Ray Conf. (ICRC2017)*, PoS(ICRC2017)405, 2017.
54. Y. Ning, Y. Yuan, T. Yu, et al., "Development of portable cosmic-ray muon detector array for muography," *arXiv:2503.18800* (preprint), 2025.
55. Z. Liang, T. Hu, X. Li, Y. Wu, C. Li, and Z. Tang, "A cosmic ray imaging system based on plastic scintillator detector with SiPM readout," *J. Instrum.*, vol. 15, art. C07033, 2020.
56. Di Francesco, R. Bugalho, L. Oliveira, et al., "TOFPET2: a high-performance ASIC for time and amplitude measurements of SiPM signals in time-of-flight applications," *J. Instrum.*, vol. 11, art. C03042, 2016.
57. J. D. Sullivan, "Geometrical factor and directional response of single and multi-element particle telescopes," *Nucl. Instrum. Methods*, vol. 95, no. 1, pp. 5–11, 1971.
58. G. R. Thomas and D. M. Willis, "Analytical derivation of the geometric factor of a particle detector having circular or rectangular geometry," *J. Phys. E: Sci. Instrum.*, vol. 5, no. 3, p. 260, 1972.
59. S. Agostinelli et al. (GEANT4 Collaboration), "GEANT4 – a simulation toolkit," *Nucl. Instrum. Methods Phys. Res. A*, vol. 506, no. 3, pp. 250–303, 2003.
60. V. Anghel, J. Armitage, F. Baig, et al., "A plastic scintillator-based muon tomography system with an integrated muon spectrometer," *Nucl. Instrum. Methods Phys. Res. A*, vol. 798, pp. 12–23, 2015.
61. Bektasoglu and H. Arslan, "Investigation of the zenith angle dependence of cosmic-ray muons at sea level," *Pramana – J. Phys.*, vol. 80, no. 5, pp. 837–846, 2013.
62. Y. Elangovan, S. Kakkad, and B. Satyanarayana, "Cosmic Muon Explorer: a portable detector for cosmic muon flux measurements and outreach," *arXiv:2602.23674* (preprint), 2026.
63. Checchia, "Review of possible applications of cosmic muon tomography," *J. Instrum.*, vol. 11, art. C12072, 2016.
64. K. Nagamine, M. Iwasaki, K. Shimomura, and K. Ishida, "Method of probing inner-structure of geophysical substance with the horizontal cosmic-ray muons and possible application to volcanic eruption prediction," *Nucl. Instrum. Methods Phys. Res. A*, vol. 356, no. 2, pp. 585–595, 1995.
65. G. Liu, X. Luo, H. Tian, et al., "High-precision muography in archaeogeophysics: a case study on Xi'an defensive walls," *J. Appl. Phys.*, vol. 133, no. 1, art. 014901, 2023.
66. R. Han, Q. Yu, Z. Li, et al., "Cosmic muon flux measurement and tunnel overburden structure imaging," *J. Instrum.*, vol. 15, art. P06019, 2020.
67. K. Morishima, M. Kuno, A. Nishio, et al., "Discovery of a big void in Khufu's Pyramid by observation of cosmic-ray muons," *Nature*, vol. 552, pp. 386–390, 2017.
68. S. Procureur et al., "Precise characterization of a corridor-shaped structure in Khufu's Pyramid by observation of cosmic-ray muons," *Nat. Commun.*, vol. 14, art. 1144, 2023.

69. Z. X. Zhang, T. Enqvist, M. Holma, and P. Kuusiniemi, “Muography and its potential applications to mining and rock engineering,” *Rock Mech. Rock Eng.*, vol. 53, pp. 4893–4907, 2020.
70. G. Jonkmans, V. N. P. Anghel, C. Jewett, and M. Thompson, “Nuclear waste imaging and spent fuel verification by muon tomography,” *Ann. Nucl. Energy*, vol. 53, pp. 267–273, 2013.
71. N. Zurlo, G. Bonomi, A. Donzella, D. Pagano, A. Zenoni, and G. Zumerle, “A new Monte Carlo muon generator for cosmic-ray muon applications,” *PoS CompTools2021*, art. 019, 2022.
72. V. Niess, “The PUMAS library,” *Comput. Phys. Commun.*, vol. 279, art. 108438, 2022.
73. T. Frach, G. Prescher, C. Degenhardt, R. de Gruyter, A. Schmitz, and R. Ballizany, “The digital silicon photomultiplier – principle of operation and intrinsic detector performance,” in *IEEE Nucl. Sci. Symp. Conf. Rec.*, 2009, pp. 1959–1965.