

MuoVis

Breakthrough muon tomography system for non-destructive scanning of large objects

1. Clear explanation — what it is and how it works

MuoVis — this is a device that "sees through" large objects without destruction, using natural cosmic radiation.

What exactly is used?

Particles fall from space to Earth constantly — **muons**.

A muon is an elementary particle similar to an electron, but ~200 times heavier.

Thanks to high energy (1–10 GeV), muons pass through:

- mountain rocks
- concrete
- metal
- buildings
- containers

This property is used for scanning.

From where → to where the signal goes

1. Space → Earth

Muons are born from the interaction of cosmic rays with the atmosphere.

2. Object → Detector

Muons pass through the investigated object.

If there is a void or dense material inside — the trajectory changes slightly.

3. Detector → Electronics

Scintillators detect a flash of light.

Silicon Photomultiplier registers photons.

FPGA and TDC measure time and coordinates.

4. Electronics → Data Processing

Algorithms reconstruct trajectories and build a 3D map of densities.

5. Processing → User

The system shows:

- hidden voids
- dense inclusions
- potentially dangerous zones
- structural anomalies

2. What does the system do in the end?

MuoVis allows:

- detecting hidden chambers in archaeological objects
- finding voids in mines
- checking containers for dense materials
- analyzing the condition of bridges and tunnels
- working in hard-to-reach zones with drones

Muon Tomograph

Problem

Research of archaeological monuments, geological objects, infrastructure, nuclear safety, and industry faces barriers:

- Non-destructive approach: cannot drill/destroy in mounds, pyramids, nuclear containers or bridges — risks of damage/contamination.
- Limited access: objects in mountains, underground, tropics or dangerous zones (mines, reactors), without heavy equipment.
- Shortcomings of traditional methods: radars (GPR) weak in dense rocks; X-ray/CT not for large scales; detectors vulnerable in extreme conditions.
- High cost/complexity: existing muon systems (e.g., from KEK or Ideon) bulky, expensive (10k+ USD) and not universal for real-time monitoring in security/mining.
- Lack of dynamics: analogs are static, without AI-predictions or mobile (drone) integration.

Solution: MuoVis — Breakthrough universal muon tomograph for archaeology, geology, nuclear safety and industry

The system detects cosmic muons in modular Micromegas/RPC-stacks, synchronizes time, reconstructs tracks and builds 3D-map of densities/anomalies.

Breakthrough improvements:

- Hybrid with GPR/LiDAR/SAR for data fusion; MST for nuclear detection.
- Real-time AI (U-Net/GANs) with predictive analytics (e.g., "risk of collapse").
- Drone/robot integration for dynamic scanning.
- Open-source with crowdsourcing data for global network.
- Improvements: dual readout, FPGA-coincidences + TDC, LoRa-module with SD-buffer, atomic clocks/IRIG-B, IP67-case, LED self-test, Cherenkov-optics for energy.

Architecture (layers)

- Detector module (hit): Micromegas/RPC-panel (16×16 cm) → gas mixture → TIA → discriminator → TDC → FPGA (readout, decoder, filter, LoRa/SD).
- Synchronization: Atomic clocks + IRIG-B (indoor/outdoor), GPS/INS for drone.
- Data collection: FPGA/MCU pack events → gateway RPi Zero writes to SQLite → cloud for crowdsourcing.
- Reconstruction: Real-time backprojection → SIRT/Kalman → AI-denoising/prediction (U-Net/GANs on real+synthetic).
- Mobility: Drone-mount (UAV with RPC) or crawling robots for dynamics.
- Mechanics/energy: Modular cassettes IP67, drone-adapters, LiFePO4 50 Ah + solar/MPPT, DC-DC (12→5V), HV for Micromegas.

Details of the block (expanded)

A) Muon detector (panel)

- Sensitive element: Micromegas for superaccuracy ($<1^\circ$) + RPC fallback for portability; gas mixture + Cherenkov-optics for muon energy. Dual readout for positions/MST (deviation for nuclear materials).
- Readout: Electrodes \rightarrow TIA \rightarrow discriminator \rightarrow ToT for energy.
- Temperature compensation: Thermosensor + DAC.
- Self-test: LED-injector.
- Power: Local 12 \rightarrow HV boost, monitoring.

B) Mini-stack / geometry and TDC

- Mini-stack: 2–3 layers Micromegas/RPC, baseline 50 cm for $\leq 1^\circ$ resolution.
- Coincidence: FPGA iCE40 with MST-algorithms.
- TDC: TDC7200 (1 ns).
- Format: {ts, id, hit_mask, pos, ToT, energy, deviation}.

C) Data collection and communication

- MCU/gateway: RP2040 + RPi Zero.
- Format: Protobuf + CRC, blockchain for data security.
- Transport: LoRa for field (5–10 km), Zigbee fallback, SD-buffer; API for crowdsourcing.
- Network: Offline/online, batch to cloud (AWS/GCP).

D) Reconstruction and tomography

- Pipeline: Real-time backprojection \rightarrow SIRT/Kalman \rightarrow AI (U-Net/GANs) for denoising/predictions (e.g., risk of collapse).
- Parameters: Adaptive voxels 0.5–5 cm.
- Fusion: GPR/LiDAR/SAR + MST for hybrid (improvement 50–70%).
- UI: Unity AR (interactive 3D, AR-overlay, predictive alerts on tablet).

E) Power and thermal

- Power pack: LiFePO₄ 50 Ah, 24–48 h, solar with MPPT; drone-batteries for UAV.
- DC-DC: 12 \rightarrow 5V, filters.
- Monitoring: PMIC + forecast.
- Heat: Passive radiators.

Mechanics and deployment

- Cassette: Aluminum 20 \times 20 \times 5 cm, IP67, weight <5 kg; drone-mount adapters (DJI-compatible).
- Frame: Foldable tripods + robot-chassis for crawling.
- Placement: 4–8 stacks, dynamic with drones/robots; borehole-option for mining.

Synchronization / calibration

- Mode: Atomic clocks/IRIG-B + GPS/INS for drone.
- Autocalibration: LED + muons (5–10 min).
- Auto-threshold: <50 Hz/channel.

- Alignment: Software + hybrid data.

Forecast of flow and shooting

- Area: $\sim 256 \text{ cm}^2 \rightarrow \sim 500\text{--}1000$ tracks/min (improved Micromegas).
- Coincidences: 100–200 events/min/stack.
- Demo: Void/anomaly in minutes; real — real-time monitoring.

Software and algorithms

- Firmware: Verilog for FPGA, C++ for MCU; edge-AI on RPi.
- Protocol: Protobuf, retransmit; API for crowdsourcing.
- Server/GUI: Python + Unity AR; GANs for scenarios.
- ML: U-Net/GANs, training on Geant4 + real (e.g., volcano/mining datasets).

Safety and EMC

- HV isolation, indicators; dosimetry for operators.
- Decoupling, ferrites, EMC-tests; blockchain for data.

Reasons for improvements (briefly)

- Micromegas/Cherenkov: Superaccuracy + energy for MST<nuclear security="">.
- AI predictive: Real-time alerts, novelty for stability.
- Drone/robots: Dynamics in inaccessible zones.
- Open-source: Global network, like in AGEMERA.

BOM (for one 2-layer stack, breakthrough version)

- Micromegas/RPC-panels: 2 pcs (~ 300 USD/pc).
- Readout-boards + Cherenkov: 2 pcs (~ 150 USD).
- FPGA iCE40 + RP2040: 1 set (~ 50 USD).
- LoRa (SX1276): 1 pc (~ 10 USD).
- Atomic clocks/IRIG-B: 1 pc (~ 100 USD).
- RPi Zero: 1 pc (~ 20 USD).
- LiFePO4 50 Ah + MPPT: 1 kit (~ 200 USD).
- IP67-case + drone-adapter: ~ 150 USD.
- GPR/LiDAR module: ~ 500 USD.
- AI-chip (edge): ~ 100 USD. Total: $\sim 1500\text{--}2000$ USD/stack (savings vs analogs, but + for breakthrough).

Success metrics

- Angular resolution: $\leq 1^\circ$.
- Stability: $\pm 1 \text{ mV}/^\circ\text{C}$.
- Throughput: ≥ 1000 events/min.
- Autonomy: 24–48 h (drone — 1–2 h flight).
- Demo: Anomaly/risk in minutes with prediction.

Alternative applications

- Archaeology: Hidden chambers in pyramids.
- Geology/Mining: Voids/mines with borehole.
- Nuclear security: Detection of materials in containers/MST.
- Infrastructure: Bridges/tunnels for stability.
- Underwater/extreme: Robots for bases.

Risks and mitigation

- Statistics: Increase stacks/drone, AI-acceleration.
- Temperature: Compensation + monitoring.
- Communication: SD + retransmit.
- EMC: Ferrites, tests.
- Scale: Open-source + partnerships.
- Regulations: IAEA certification for nuclear.

Explanation of operation

muons penetrate through the object, Micromegas/RPC detect hits/deviations, FPGA/TDC set marks, LoRa sends to RPi/cloud. AI reconstructs 3D real-time with predictions, hybrid GPR/LIDAR/SAR improves accuracy. Drone/robots add dynamics. Open-source network — global monitoring. For archaeology — non-destructive; for security — detection of nuclear threats; for mining — mapping ores.

Scintillator: Detailed technical description

Scintillator — this is a key component of the muon tomography detector, converting the energy of ionizing particles (in this case, cosmic muons) into visible light (photons). In the muon tomograph system, the scintillator detects the passage of a muon, generating a light signal, which is then captured by photodetectors (for example, SiPM). The main mechanism is scintillation: excitation of electrons in the material by the muon leads to their relaxation with the emission of photons. Efficiency is determined by light yield, decay time, emission wavelength, density, and radiation resistance.

In the original project architecture, it is recommended to use plastic scintillators (for example, BC-408 or EJ-200) for portability and cost, or crystalline (NaI(Tl), CsI(Tl)) for high efficiency. Below is a detailed technical description based on standard specifications and applications in muon tomography. I rely on data from manufacturers (Eljen Technology, Saint-Gobain) and scientific publications on muon detectors. For muon tomography, key parameters are: light yield ~8-70% of anthracene (standard), response time 1-5 ns for plastics and 230-250 ns for crystals, and energy resolution depending on size and geometry.

1. Purpose and principle of operation

The scintillator converts the kinetic energy of the muon (typically 1-10 GeV in surface applications) into photons in the visible range. Process:

- Ionization and excitation: the muon loses energy through ionization of material atoms, creating electron-hole pairs. The energy loss is described by the Bethe-Bloch formula:

$$-\frac{dE}{dx} = \frac{4\pi z^2 e^4 N_A \rho}{m_e c^2 \beta^2} \left[\ln \left(\frac{2m_e c^2 \beta^2 \gamma^2}{I(1-\beta^2)} \right) - \beta^2 \right]$$

where:

- z — muon charge (=1),
- e — electron charge,
- N_A — Avogadro's number,
- ρ — material density,
- m_e — electron mass,
- $\beta = v/c, \gamma = 1/\sqrt{1-\beta^2}$,
- I — mean ionization potential (for plastic ~65 eV, for NaI ~450 eV).

For muons with energy ~4 GeV:

$$\frac{dE}{dx} \approx 2 \text{ MeV}/(\text{g}/\text{cm}^2)$$

This corresponds to the generation of order

$$10^4 - 10^5 \text{ photons/MeV}$$

- **Emission:** Excited electrons relax, emitting photons. Light yield (LY) is measured in photons/MeV or % of anthracene (standard: 100% ~38 000 photons/MeV for neutrons).
- **Registration:** Photons are collected by fiber light guides (WLS) or directly by photodetectors. In

- **Registration:** Photons are collected by fiber light guides (WLS) or directly by photodetectors. In muon tomography, it is important to minimize dead zones and ensure angular resolution $<5^\circ$ for tracking.
- **Application in tomography:** Scintillators are used in layers (mini-stacks) for registering muon trajectories. Muon flux density on Earth $\sim 10^2$ muons/m²/s, so for a panel 0.3 m² (~ 300 cm²) ~ 30 – 300 tracks/min is expected, depending on geometry and altitude.

2. Types of scintillators

Two classes are recommended: plastic (for compactness and cost) and crystalline (for high efficiency). Choice depends on balance between light yield, cost, and durability.

2.1. Plastic scintillators (for example, BC-408, EJ-200, EJ-212)

Material: polystyrene (PST) or polyvinyltoluene (PVT) with additives (p-terphenyl + POPOP for fluorescence). Density ~ 1.05 – 1.06 g/cm³, effective Z ~ 5.8 (low, but sufficient for muons >1 GeV). Non-hygroscopic, mechanically strong (Young's modulus ~ 3 GPa), resistant to radiation up to 10^4 – 10^5 rad.

Technical characteristics (EJ-200 as example):

- Light yield: 10 000–10 400 photons/MeV (64% of anthracene), ~ 8 – 10% of theoretical maximum (for 1 MeV electron). Efficiency formula:

$$\eta = \frac{LY}{E_{dep}}$$

where E_{dep} — deposited energy. For muons ~ 10 – 15% of NaI(Tl) due to low Z, but sufficient for tracking (proxy-energy through ToT).

- Decay time: main component ~ 2.1 – 3 ns (fast response for high fluxes), with tail ~ 10 – 20 ns.

Full pulse shape:

$$I(t) = I_0 e^{-t/\tau}, \quad \tau \approx 2 \text{ ns}$$

Suitable for coincidences in <10 ns window.

- Emission wavelength: peak ~ 423 nm (blue), range 380–490 nm. Refractive index $n = 1.58$, attenuation length >3.8 m (minimal losses in plates up to 1 m).
- Density and sizes: 1.023 g/cm³. Typical geometries: plates 300×100×10 mm or bars 300×20×10 mm (segmented for positioning). Coverage area: 10–15 cm²/detector for optimal sensitivity; for 1000 m² system — ~ 10 – 20 m² material.
- Durability: radiation hardness $> 10^4$ rad (without significant LY degradation). Temperature range: -20 ... $+50$ °C (LY drops by $\sim 1\%$ /°C above 20 °C). Mechanics: Shore D hardness=75, not brittle.
- Energy efficiency: $\sim 10\%$ for muons 1–10 GeV (due to low Z, muons are minimally ionizing particles, MIP).

Cost: 50–150 USD/m² for basic plates.

Comparison with other plastics (for example, BC-408):

- BC-408: LY = 10 000 photons/MeV (64% anthracene), $\tau = 2.1$ ns, peak 425 nm, attenuation length 2.1 m, density 1.032 g/cm³. Better for large panels (>1 m), but shorter light attenuation.
- EJ-212 (thin): LY = 9 200 photons/MeV, $\tau = 2.4$ ns, for fast applications (for example, timing in tomography).

Applications in muon tomography: plastics are preferred for portable systems due to lightness (weight ~1 kg/panel) and cost. In projects like COSMIC (muon tracker), segmented bars with WLS-fibers (Kuraray Y-11, \varnothing 1 mm) are used for light collection, increasing uniformity by 2–3 times.

Light collection: increase up to 3× with WLS (wavelength-shifting fibers). In systems with triangular bars (triangular strips), position resolution ~1 mm, for 3D-reconstruction by backprojection or SIRT.

2.2. Crystalline scintillators (for example, NaI(Tl), CsI(Tl))

Material: inorganic crystals with dopants (Tl for activation). High Z (NaI: $Z_{\text{eff}} \approx 50$, CsI: $Z_{\text{eff}} \approx 54$), density 3.67–4.51 g/cm³, hygroscopic (require hermetization). Best efficiency for low-energy gamma/muons.

NaI(Tl) (sodium iodide with thallium):

- Light yield: 38 000–41 000 photons/MeV (100% anthracene), ~30–60% theoretical. Efficiency: $\eta \approx 45$ photons/keV. Non-proportionality: LY drops by ~20% below 100 keV due to quenching.
- Decay time: main component ~230–250 ns (exponential: $\tau = 250$ ns), with components ~1 ms and 12 ms (temperature-dependent). Formula:

$$I(t) = \sum I_i e^{-t/\tau_i}, \quad \tau_1 = 250 \text{ ns (90\%)}, \tau_2 = 1 \text{ ms (10\%)}$$

- Wavelength: peak 415 nm (compatible with bialkali PMT), range 300–550 nm.

$n = 1.85$, attenuation coefficient ~0.1–0.2 cm⁻¹.

- Density and sizes: 3.67 g/cm³. Crystals 3×3 – 20×20 cm, thickness 5–10 mm.

For tomography: plates 10×10 cm², ~10–100 m² for large systems (cost 500–1500 USD/m²).

- Durability: radiation hardness ~10³ rad (LY degradation >10%).

Temperature range: -20...+50 °C (LY maximum at 25 °C, drop ~0.3%/°C).

Hygroscopic (NaCl-like), requires aluminum housing.

- Energy efficiency: ~60% for gamma/muons (high Z).

Energy resolution: 5.4–7% FWHM at 662 keV (with SiPM).

CsI(Tl) (cesium iodide with thallium):

- Light yield: 54 000–65 000 photons/MeV (~165% NaI(Tl)), peak 550 nm (green, optimal for SiPM/PD).

Non-proportionality: varies by ionization (fast/slow components for PSD).

- Decay time: 3–5 ns (fast component), main ~680 ns + 3.3 μ s (depends on particle: τ short for gamma, long for alpha).

Afterglow: 0.05–0.39% at 20 ms (low for tomography).

- Wavelength: peak 550 nm, range 350–700 nm.

Photodetectors (SiPM): Detailed technical description

Photodetectors of the SiPM type (Silicon Photomultipliers, silicon photomultipliers) — these are solid-state devices based on avalanche photodiodes (APD), operating in Geiger mode (Geiger-mode APD, GM-APD). They are designed for registering weak optical signals from scintillators, converting photons into an electrical signal with high gain.

SiPM consist of an array of micropixels (SPAD — single-photon avalanche diodes), connected in parallel, which provides high sensitivity to single photons. In muon tomography, SiPM capture light from scintillators (for example, 10^3 – 10^4 photons per muon hit), generating a proportional pulse for energy and time estimation.

SiPM outperform traditional photomultiplier tubes (PMT) in compactness, low supply voltage (30–50 V versus 1–2 kV for PMT), magnetic resistance, and cost. Drawbacks — high dark noise (dark count rate, DCR) and optical crosstalk.

As of August 2025, the SiPM market is growing: volume — 146–152 mln USD (2024–2025), forecast — 208–240 mln USD by 2030 (CAGR 6–8 %). Development is driven by improvement in radiation hardness (for LHC/CMS), time resolution (<50 ps), and application in LiDAR, nuclear safety, and medical tomography.

1. What is it?

SiPM represent an array of thousands of micropixels (SPAD), each of which is an avalanche photodiode in Geiger mode.

- Active area: from 1×1 mm² to 6×6 mm² (or larger in arrays).
- Material: silicon (Si), with n-p or p-n junction.
- Key characteristics:
- photon detection efficiency (PDE) — up to 50–65 % (improved to 70 % in 2025 models for LiDAR);
- gain $G = 10^5 - 10^7$;
- time resolution — <50 ps (in 2025 achieved 20–30 ps for ToF tasks);
- dark count (DCR) — 10^5 – 10^6 s⁻¹·mm⁻² at 25 °C.

2. What is it used for?

- Converting light from scintillators into an electrical signal (muons, gamma, neutrons).
- High-precision photon counting: PDE >50 % in range 420–450 nm (scintillator BC-408).
- In muon tomography: determining position, time, and energy of hit by pulse amplitude.
- In CMS projects (CERN): work in radiation conditions $>10^{14}$ n·cm⁻² (upgrade 2025).
- Other applications: LiDAR (autonomous transport, mining exploration), PET/MRT (medicine), nuclear safety.

3. How does SiPM work?

Structure.

- Micropixel size — 10–100 μm.
- Fill factor — 50–80 %.

- Each SPAD equipped with quenching resistor.
- Number of subpixels — from 100 to 10^4 .

Geiger mode.

Bias voltage:

$$V_{\text{bias}} = V_{\text{breakdown}} + V_{\text{over}},$$

where $V_{\text{over}} = 1 - 5$ V, and $V_{\text{breakdown}} \approx 25 - 30$ V.

Upon absorption of a photon ($h\nu > E_g$

$$\approx 1.1 \text{ eV}),$$

an electron-hole pair arises, which initiates an avalanche.

Gain:

$$G = \frac{C_{\text{cell}} V_{\text{over}}}{q},$$

where $C_{\text{cell}} \approx 50 - 100$ fF, q — electron charge.

Typical value G

$\approx 10^6$ at $V_{\text{over}} = 3$ V.

Pulse summation.

Output charge:

$$Q_{\text{out}} = N_{\text{ph}} \cdot G \cdot q \cdot \text{PDE},$$

where N_{ph} — number of photons.

For $N_{\text{ph}} = 10^3$: $Q_{\text{out}} \approx 10^9 - 10^{10}$ electrons (amplitude 1-10 mV on 50 Ohm).

Time characteristics.

- Recovery time: $\tau_{\text{rec}} = R_q \cdot C_{\text{cell}} \approx 10 - 100$ ns.

- Jitter:

$$\sigma_t = \sqrt{\sigma_{\text{walk}}^2 + \sigma_{\text{noise}}^2},$$

value < 50 ps (2025 — up to 20-30 ps).

Noise.

- DCR — $10^5-10^6 \text{ s}^{-1}\cdot\text{mm}^{-2}$.
- Optical crosstalk (OCT) — 5-20 %.
- Afterpulsing — 1-10 %.

Efficiency.

$$PDE = \frac{QE}{\text{FF}} \cdot P_{\text{trigger}}$$

where QE

approx 80 – 90

in range 400-500 nm.

Final value PDE = 40-65 % (in 2025 — up to 70 %).

4. What technologies are used?

- Silicon photodiodes (N-on-P or P-on-N structure).
- CMOS-compatible processes (0.13-0.35 μm), antireflective coatings.
- Trench-isolation technology (OCT <5 %).
- Active and passive avalanche quenching (AQ <10 ns, PQ \approx 50 ns).
- Supply voltage — 25-70 V.
- Power consumption — 30-50 $\text{mW}\cdot\text{mm}^{-2}$.
- Manufacturers: Hamamatsu (S13360), ON Semiconductor (MicroFJ), KETEK (PM33xx), AdvanSiD.

5. Number of SiPM in the project

- One scintillator bar (3-5 cm^2): 4-16 SiPM.
- Mini-stack (3 layers, 5 bars): 30-120 SiPM.
- Total coverage area: 100-150 mm^2 .
- Full system (3-6 stacks): 100-500 SiPM.

6. What does SiPM consist of?

- Array of SPAD with p-n junction (depth 10-20 μm).
- Amplification system (quenching resistors or transistors).
- Electronic circuits (analog/digital output, TDC/ADC in digital SiPM).
- Housing: ceramic or plastic, with glass window.
- Sizes: 3×3 – 6×6 mm, thickness \approx 1 mm.

7. Cost of SiPM

- Hamamatsu S13360 (3×3 mm): 50–150 USD.
- ON Semi MicroFJ-30035 (3×3 mm): 60–120 USD.
- KETEK PM3350 (3.5×3.5 mm): 40–100 USD.
- Bulk purchases: reduction by 20–50 %.

8. Additional calculations and features

Noise.

$$DCR = DCR_{\text{thermal}} + DCR_{\text{tunnel}}, DCR_{\text{thermal}} \sim \exp\left(-\frac{E_g}{2kT}\right).$$

Dark count doubles every 7–10 °C.

Response time.

- Rise time: 1–5 ns.
- Pulse duration: 10–50 ns.
- For ToT method:

$$\text{ToT} = \frac{\tau_{\text{tm}}}{\text{cdot}} \ln(N_{\text{tph}} + 1).$$

Example calculation.

For muon hit ($N_{\text{tph}} = 1000$, PDE = 50 %, $G = 10^6$):

$$Q \approx \frac{N_{\text{tph}}}{\text{cdot}} \text{PDE} \cdot G \cdot q \approx 10 \text{ pC}.$$

Signal-to-noise ratio:

$$SNR = \frac{\sqrt{N_{\text{tph}} \cdot \text{PDE}}}{\sqrt{DCR_{\text{tint}} + OCT \cdot N_{\text{tph}}}} > 10.$$

Positioning.

For two SiPM on fiber:

$$Pos = \frac{L}{2\alpha \cdot \text{cdot}} \ln\left(\frac{S_1}{S_2}\right),$$

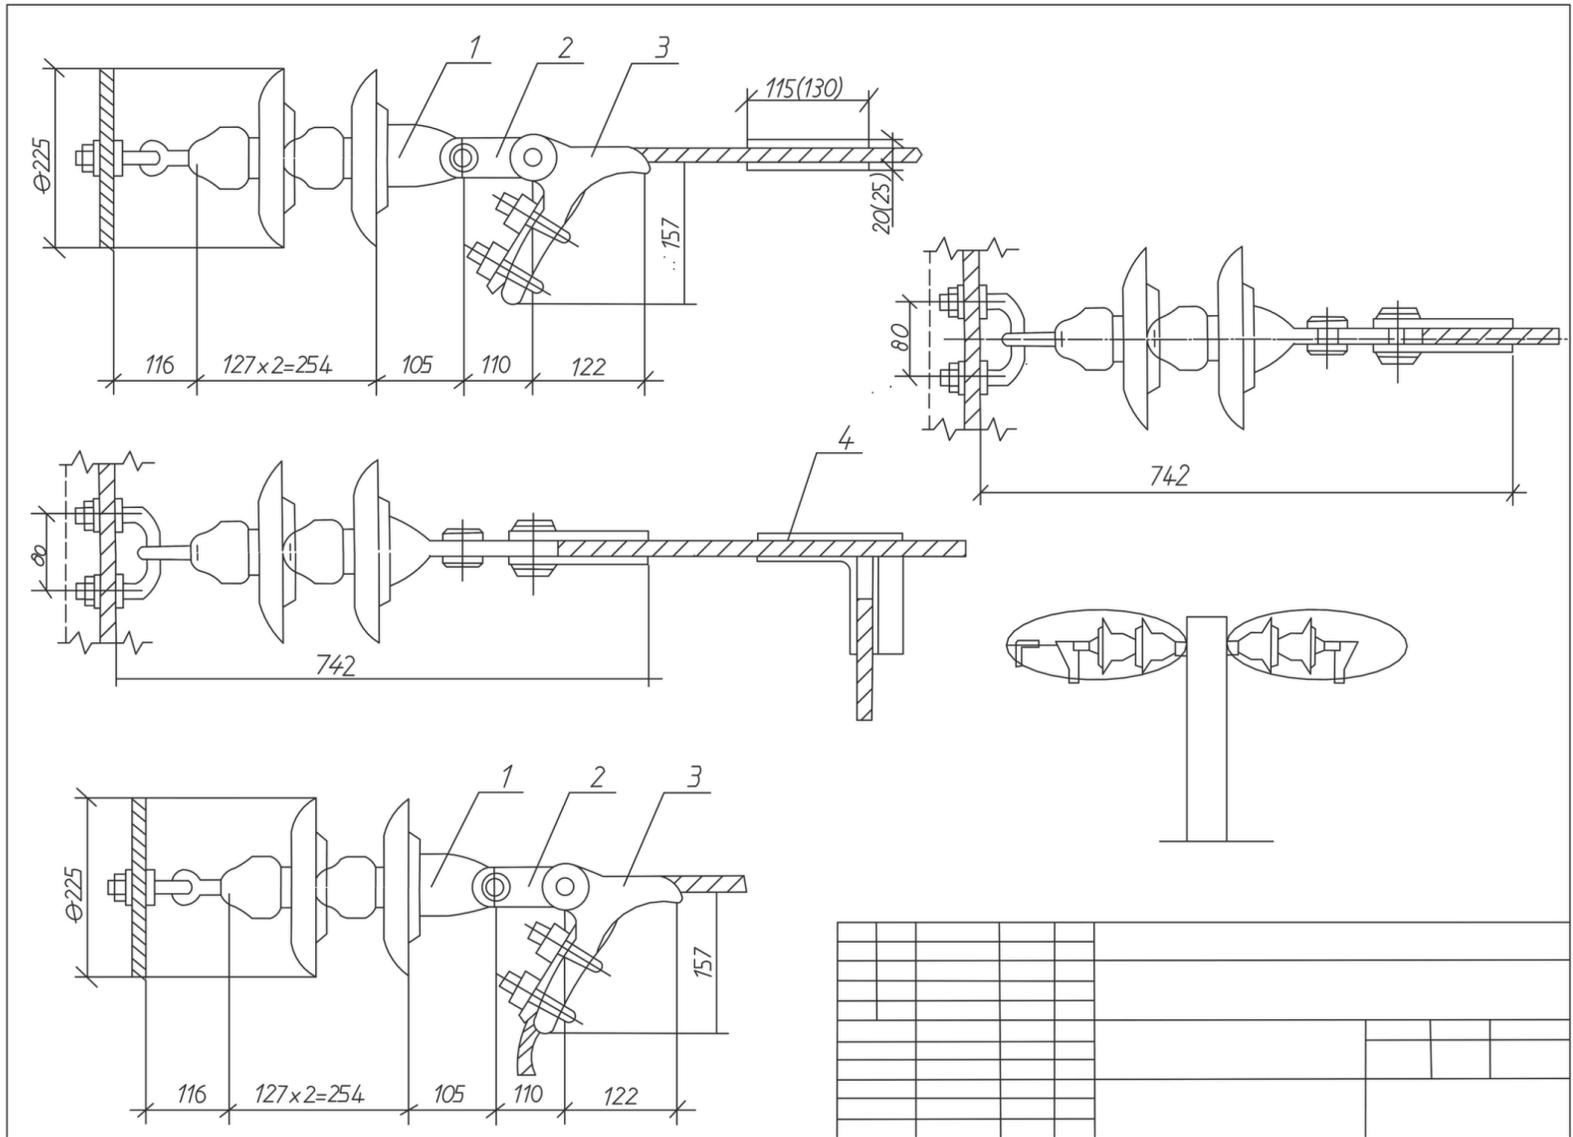
where

$\alpha \approx 0.01 - 0.05 \text{ cm}^{-1}$ — fiber attenuation coefficient.

Example parameters for project:

- Size: 3×3 mm².
- PDE: 50–60 % (420–450 nm).
- Gain: $G = 10^6$.

- DCR: $<1 \text{ MHz} \cdot \text{mm}^{-2}$.
- Response time: 10-50 ns.
- Price: 50-150 USD.



Analog tract: Detailed technical description

Analog tract (analog front-end, AFE) — this is a chain of electronic components processing the raw signal from photodetectors (SiPM) before digitization. It includes amplifiers (for example, transimpedance amplifiers, TIA), comparators/discriminators (for threshold detection) and high-voltage circuit (HV-scheme) for SiPM power. Goal — amplify weak signal (millivolts/nanoseconds), reduce noise, measure amplitude/time and prepare for TDC/ADC. In muon tomography, AFE is critical for accurate tracking: time resolution <10 ns, dynamic range 10^0 - 10^4 photons, low noise (SNR >20 dB). As of August 2025, advancements focus on low-power ASIC (application-specific integrated circuits) for portable detectors, with integration into SiPM-arrays for muon imaging in nuclear safety and geology. Noise reduced to <1 fA/ $\sqrt{\text{Hz}}$, power consumption <10 mW/channel in new designs for TOF (time-of-flight) systems.

1. What is it?

Analog tract — sequence of circuits: TIA for converting current to voltage, preamp/shaper for filtering, comparator for trigger and HV-scheme for stable SiPM power ($V_{bias} = 25 - 70$ V). This AFE minimizes jitter ($\sigma_t < 50$ ps), compensates temperature drifts and provides linearity for ToT (time-over-threshold) as energy proxy. In portable systems (e.g., muon detectors) AFE is integrated into ASIC for compactness, as in Cosmo ArduSiPM (2024-2025: PCB with AFE for SiPM, low-noise <1 pA/ $\sqrt{\text{Hz}}$).

2. What is it used for?

- **Amplifiers:** Increase signal amplitude from SiPM (current ~ 1 - 10 mA per avalanche) to levels for ADC/TDC (0.1-5 V). Reduce noise (electronic noise influence) for single-photon resolution.
- **Comparators:** Detect events above threshold, generating TTL/LVDS pulses for coincidences. Used for CFD (constant fraction discriminator) for timing <10 ns.
- **HV-scheme:** Provides stable high voltage for SiPM Geiger mode, with compensation (temperature, aging). In muon tomography AFE critical for positional resolution (~ 1 mm) and energy (ToT proxy for dE/dx).
- **Applications:** particle detectors (CMS/LHC) for SiPM readout; portable muon arrays for imaging (low-power <5 mW/channel); TOF-ion mass spectrometers (2025 ASIC: <1 ns timing).

3. How does it work?

AFE processes pulse from SiPM (width 10–50 ns, amplitude 1–100 mV).

- Amplifiers (TIA/Preamps):
- TIA converts current I_{in} to voltage V_{out} :

$$V_{out} = -I_{in} \cdot R_f,$$

where R_f — feedback resistor (1–10 kOhm).

- Bandwidth:

$$BW = \frac{1}{2\pi R_f C_f},$$

where C_f — capacitance for stability (pole-zero compensation).

- Noise:

$$i_n = \frac{\sqrt{\frac{4kTR_f + 2qI_{dark}}{BW}}}{\sqrt{BW}} + \frac{v_n}{R_f},$$

where v_n — opamp voltage noise (~1–5 nV/√Hz).

- Preamp/shaper: CR-RC filter (
 $\tau_{shape} \sim 10\text{--}100$ ns):

$$V(t) = \frac{Q}{C} \cdot \frac{t}{\tau} \exp\left(-\frac{t}{\tau}\right),$$

where Q — charge from SiPM.

SNR:

$$SNR = \sqrt{\frac{Q^2 2qI_{leak}\tau + 4kTR_s\tau + e_n^2 C_{in}^2 / \tau}{C^2}}$$

- Comparators/discriminators:
- Compare V_{in} with V_{th} (threshold from DAC, 1–100 mV). Output: TTL/LVDS if $V_{in} > V_{th}$.
- Response time: $t_{prop} < 10$ ns.
- CFD: $V_{th} = f \cdot V_{peak}$, $f = 0.2\text{--}0.5$.
- HV-scheme (High Voltage):
- Generates $V_{bias} = 30\text{--}70$ V from 5–12 V.
- DC-DC boost:

$$V_{textout} = \frac{V_{textin}}{1 - D},$$

where D — duty cycle PWM.

- Compensation:

$$\Delta V_{textbias} = k$$

\cdot

$$\Delta T, k = -20 - 60$$

$\text{mV}/^\circ\text{C}$

C .

4. What technologies are used?

- Amplifiers: low-noise opamps (OPA211: $v_n = 1.1$ $\text{nV}/\sqrt{\text{Hz}}$, $I_n = 0.1$ $\text{pA}/\sqrt{\text{Hz}}$).
- Comparators: high-speed CMOS (TLV3501: slew rate >100 V/μs).

- HV: DC-DC converters (LM2623, freq 1.3 MHz); DAC for fine-tune (resolution 10–50 mV).

5. How many components are needed?

- Amplifiers: 1–2/channel (TIA + preamp). For 4–16 SiPM/detector — 4–32.
- Comparators: 1/channel (for trigger). For multi-channel: quad-chips (LM339 ×4).
- HV: 1/group of SiPM (or per-channel for compensation).

6. What do the components consist of?

- Amplifiers: opamps (BJT/MOSFET), R/C for feedback, filters (LC for EMI).
- Comparators: Differential transistor pairs, hysteresis-resistors.
- HV-scheme: MOSFET/IGBT, inductors/transformers, Schottky-diodes, LDO.

7. Cost of components

- Amplifiers: OPA211 ~5–10 USD.
- Comparators: LM339 ~1–3 USD; high-speed ~5–15 USD.
- HV: LM2623 ~5–15 USD; HV9910 ~3–10 USD.
- Full AFE ASIC: ~20–50 USD/channel (2025: reduction for mass-production).

8. Additional calculations and features

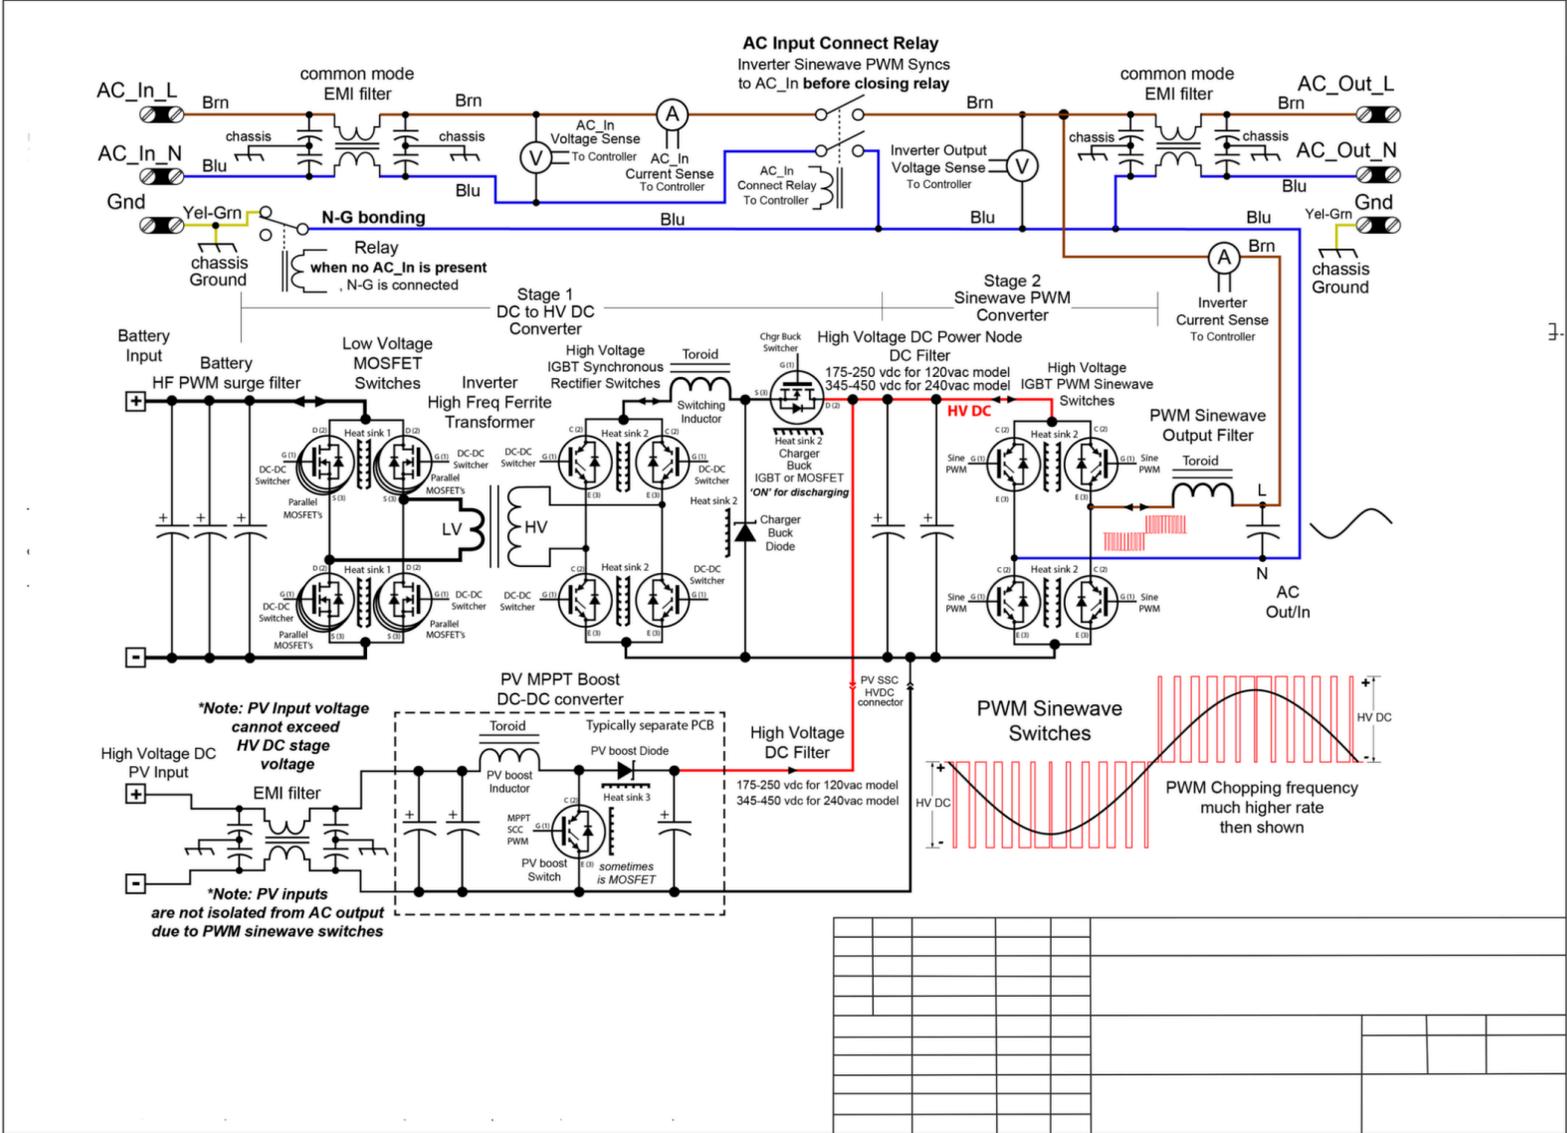
- Temperature stabilization: $dV_{os}/dT < 1$
 $\mu\text{V}/\text{C}$.
- Response time: <50 ns for comparators; ENC <100 e⁻ for TIA.

- Calculation: at signal $Q = 1$

tpC :

$$V_{\text{out}} = \frac{Q C_f}{\text{approx } 1} \text{V},$$

$$\text{tpC} = \frac{1}{\text{tpF}}.$$



Digital logic: Detailed technical description

Digital logic in the project — this is a block responsible for processing digitized signals from the analog tract, setting timestamps and forming events for further reconstruction.

Key components: TDC (Time-to-Digital Converter, time-to-digital converter), FPGA (Field-Programmable Gate Array, programmable logic integrated circuit) and MCU (Microcontroller Unit, microcontroller).

This provides high-precision time stamping of hits (resolution <5 ns), coincidences (matching events in layers) and data packing for transmission. In muon tomography, digital logic is critical for tracking: angular resolution $<5^\circ$ depends on time accuracy (

Δt

θ

$c \cdot t$

Δt

$\Delta t / \text{baseline}$, where baseline — distance between layers ≈ 50 cm).

As of September 2025, developments focus on low-power FPGA with built-in TDC (for example, for portable muon detectors in geology and nuclear safety) with resolution <1 ns and AI integration for on-chip filtering. Power consumption reduced to <50 mW/channel in new chips for TOF (time-of-flight) systems.

1. What is it?

Digital logic — combination of hardware (FPGA/TDC) and software (MCU firmware) elements for converting analog triggers into digital events.

- TDC measures signal arrival time (leading/trailing edges for ToT).
- FPGA handles logic (coincidences, filters).
- MCU manages the system (configuration, transmission).

General event format: {timestamp (64-bit, ns-resolution), module_id, layer_mask, strip_id, ToT (energy proxy), temp}.

In 2025: digitization at ASIC level with AI-accelerators for real-time track validation.

2. What is it used for?

- TDC: accurate measurement of Δt between start/stop (for example, PPS and hit). Resolution 1–5 ns for angular accuracy $\sim 1\text{--}5^\circ$ ($\theta = \arctan(c \cdot \Delta t / d)$, where $d = \text{baseline}$).
- FPGA: parallel processing (coincidences <100 ns), buffering, simple ML (for example, noise rejection). For multi-channel systems: up to 1000 channels/chip.
- MCU: control (thresholds, calibration), interfaces (I²C/SPI for HV-DAC), data transmission. In muon tomography: event filtering (>200 /min/stack), synchronization with GPS/IRIG-B.

Applications:

- particle physics (CERN ATLAS/CMS: TDC in muon triggers),
- portable imaging (geology/mining: low-power FPGA for field deployment),
- TOF-PET (medicine: <50 ps timing).

3. How does it work?

Digital logic converts LVDS/TTL from comparator into digital marks.

Pipeline:

- TDC
- Principle: measures Δt as code (for example, delay-line or oscillator-based).
- Delay-line: chain of inverters (resolution $\sim 10\text{--}50$ ps/element), $\text{code} = N_{\text{delay}} \cdot \tau_{\text{delay}}$.
- Oscillator: counter of ticks ($f_{\text{ref}} = 1\text{--}10$ GHz), $t_s = N \cdot (1/f_{\text{ref}}) + t_{\text{ext fine}}$ (DLL/PLL for sub-ps).
- Composition: start/stop inputs, coarse counter (64-bit, 40 MHz for ~ 25 ns LSB), fine interpolator (Vernier delay line: $\Delta t_{\text{ext fine}} = \tau_1 - \tau_2$ approx 10 ps).
- Resolution: $R = \frac{\text{text LSB}}{\sqrt{12}}$ approx 0.3 $\cdot \text{code}$ text LSB for Gaussian noise.
- Types:
 - TDC7200 (TI: 55 ps LSB, range 12 ns–8 ms, power 0.5 mW),
 - CAEN V1720 (12-bit ADC+TDC, 3.2 ps),
 - RP2040 PIO (quasi-TDC, 4–8 ns, low-cost).
 - Nonlinearity: INL/DNL <0.5 LSB (LUT calibration).
- FPGA
- Principle: programmable logic (LUTs, flip-flops) for parallel processing.
- Coincidence: AND-gate on hits in layers (window $\pm 30\text{--}100$ ns).
- Event packet: FIFO-buffer (ring buffer, 1–10k events).
- HDL (Verilog/VHDL):

```
always @(posedge clk)
```

```
if (hit_mask == 3'b111) event_pack();
```

- Processing: Bayesian filters for noise, ML-classifier (lightweight NN on LUTs for valid track vs noise).
- Clock: 100–500 MHz, latency <10 μ s.
- Types: iCE40UP5K (Lattice: 5k LUTs, 120 kbit RAM, power <50 mWt, for low-cost muon); Xilinx Spartan-6 (50k LUTs, for complex logic).
- MCU
- Principle: CPU for control (ISR for hits, DMA for buffers).
- Setup: I²C for DAC/thresholds.
- Transmission: UART/SPI to LoRa/SD.
- Types: RP2040 (dual Cortex-M0+, 133 MHz, PIO for custom TDC); STM32H7 (Cortex-M7, 480 MHz, FPU for calc).
- Firmware: C/C++, ring buffer for ISR < DMA.

4. What technologies are used?

- TDC: delay-line (Vernier for <10 ps), time-amplifier (TA for stretching Δt), ASIC (Nutt method: coarse+fine). In 2025 — hybrid TDC+ADC in one-chip for energy+time.
- FPGA: CMOS-based gates (7–28 nm nodes), DSP blocks for ML. Tools: IceStorm (open-source for iCE40), Vivado for Xilinx.
- MCU: ARM Cortex (M0/M4/M7), peripherals (TIM for counters, DMA for high-rate). RTOS (FreeRTOS) for multitasking.

5. How many components are needed?

- TDC: 1/channel (or multi-channel chip, like TDC7200 for 8 channels). In stacks: 3–6.
- FPGA: 1/stack (for 16–32 channels). For system: 1–2.
- MCU: 1/stack (or central for multi-stack).

6. What do the components consist of?

- TDC: counters, delay lines (inverters/chains), PLL/DLL, registers. Integrated: ADC for calib.
- FPGA: LUTs (4–6 inputs), FFs, BRAM (block RAM), I/O pins (LVDS for PPS).
- MCU: CPU core, peripherals (ADC/DAC/TIM), flash/RAM (256 kB–1 MB), interfaces (USB/Ethernet optional).

7. Cost of components

- TDC: TDC7200 ~10–30 USD; CAEN V1720 ~200–1000 USD (multi-channel).
- FPGA: iCE40UP5K ~15–50 USD; Spartan-6 ~50–200 USD.
- MCU: RP2040 ~5–20 USD; STM32H7 ~10–50 USD.

8. Additional calculations and features

- Synchronization:

$$t_s = PPS + ticks,$$

(64-bit counter, overflow ~584 years). Drift <1 ppm with OCXO (oven-controlled oscillator).

- Error handling: CRC32 on events, retransmit. Dead-time <1 % at 1 MHz rate.

- Calculations: for baseline 50 cm and

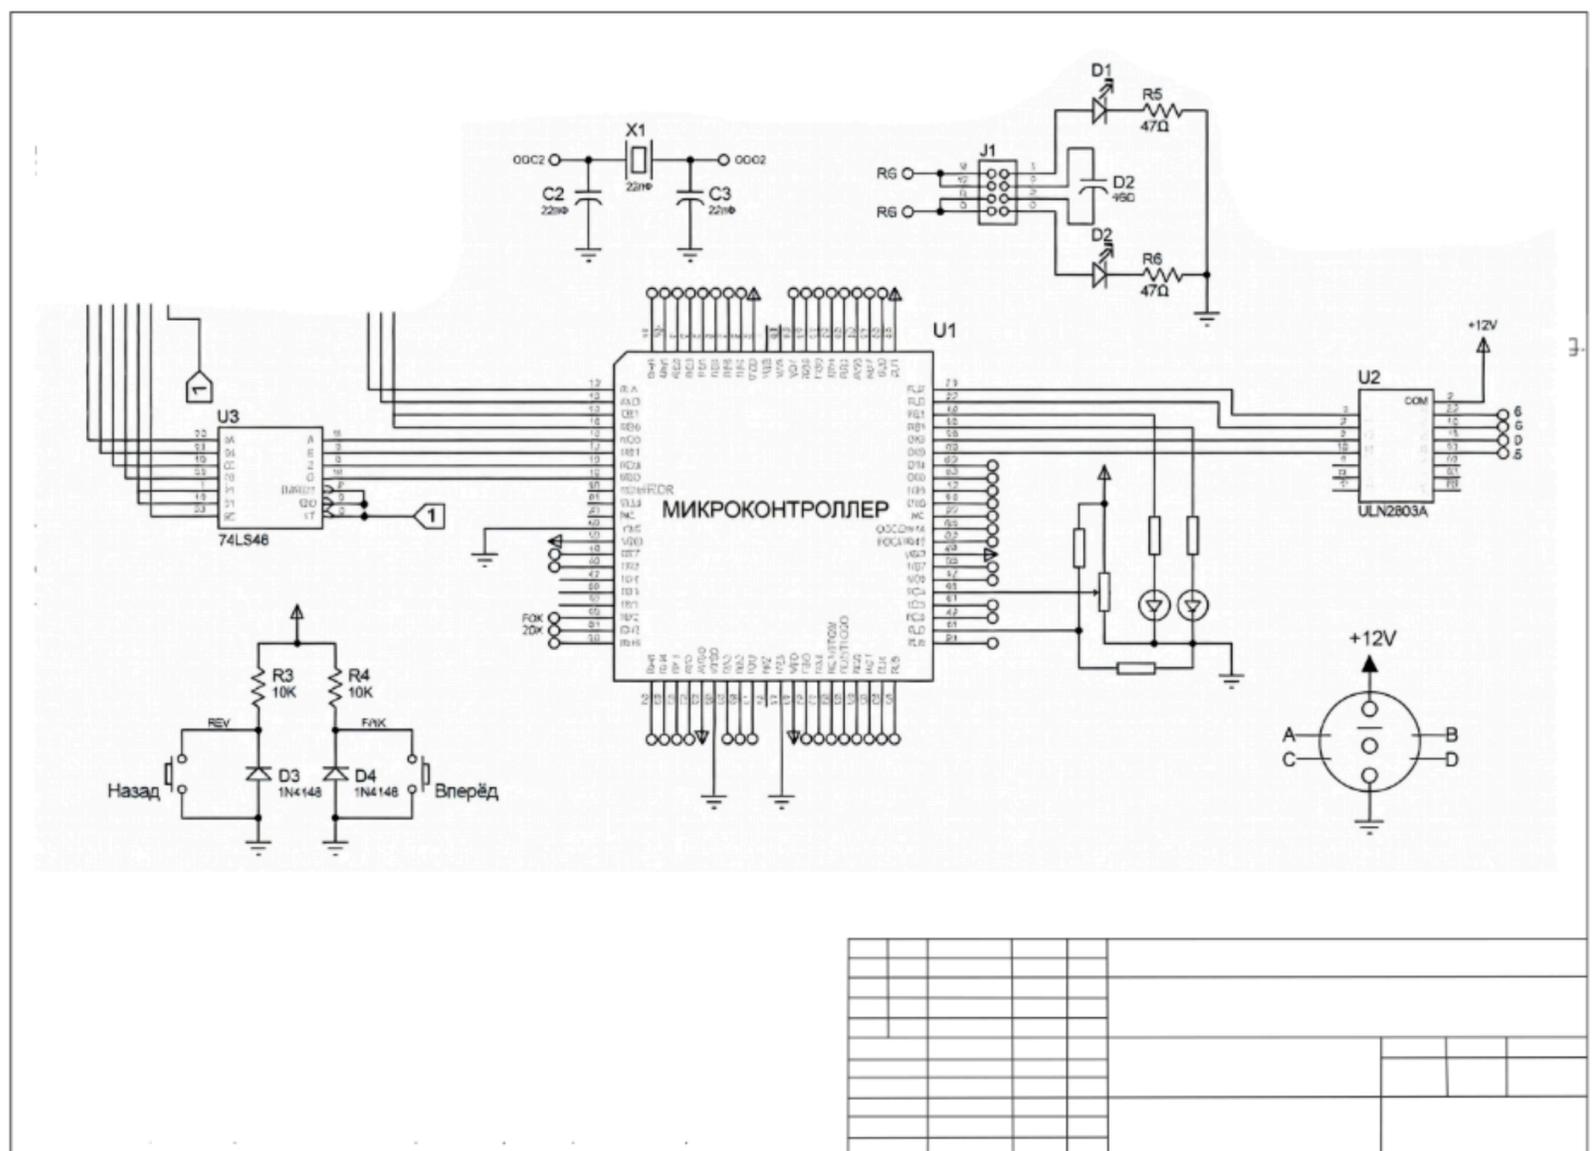
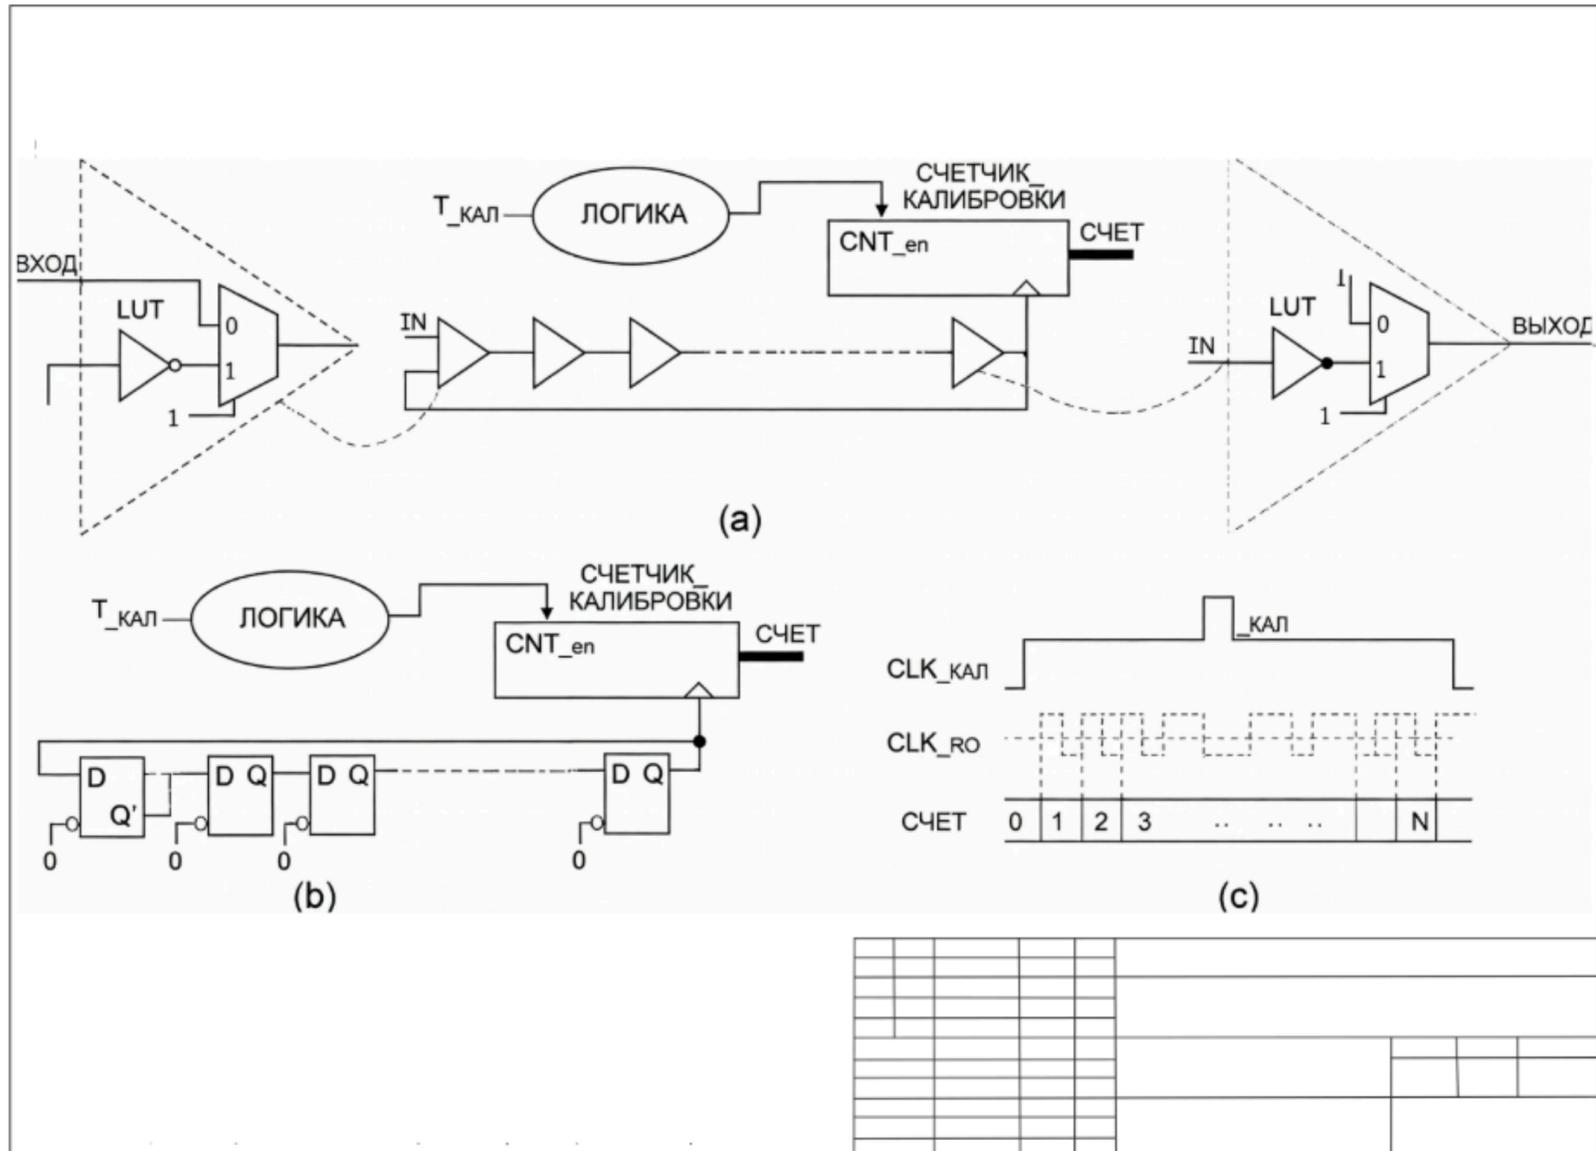
$\Delta t = 5$ ns:

$\Delta \theta \approx \frac{3 \cdot 10^8 \text{ m/s} \cdot 5 \cdot 10^{-9} \text{ s}}{0.5 \text{ m}} \approx 3^\circ$.

Throughput: 200 events/min at clock 100 MHz.

- Features 2025: AI on FPGA (TensorFlow Lite for track classification), radiation-hard (SEU mitigation with TMR — triple modular redundancy). In project: PIO-TDC as fallback, HDL for custom logic.

Example parameters: TDC resolution 1 ns, FPGA 16–32 channels @200 MHz, MCU 168 MHz with 1 MB RAM.



Communication and synchronization: Detailed technical description

Communication and synchronization — this is a subsystem ensuring data transmission between detector modules and precise time coordination of events (synchronization) in the muon tomograph system. It includes both wired channels (cables, for example CAT5e for PPS-signals), and wireless (LoRa, Wi-Fi, Zigbee), as well as synchronization sources (GPS for outdoor, IRIG-B or atomic clocks for indoor). Goal — provide low-latency transmission of compact events (<100 byte/event) with packet loss <0.1 % and time accuracy <10 ns for muon tracking (angular resolution <5°). In extreme conditions (caves, tropics) emphasis on durability: offline-buffer (SD) and retransmission. As of September 2025, IoT-network development emphasizes hybrid approaches: LoRa for long-range low-power communication, with AI integration for optimization (for example, adaptive modulation), and GPS/IRIG-B for sub-microsecond synchronization in scientific instruments. Power consumption minimized (<10 mW for LoRa), taking into account electromagnetic compatibility (EMC) for field applications.

1. What is it?

Communication and synchronization subsystem combines data transport (wired and wireless) and time references.

- Communication: transmission of binary frames (for example, Protobuf/CBOR with CRC32) from MCU to gateway/cloud.
- Synchronization: binding local timers to global time (UTC) for event correlation between modules (coincidences in window $\pm 30\text{--}100$ ns).

Two-mode operation: outdoor (GPS 1PPS + 10/20 MHz) and indoor (master-PPS via LVDS/CAT5e or IRIG-B/atomic clocks). Event format: {ts64, id, data} with confirmation (ACK) and retransmission. In 2025, integration with 5G/LTE for hybrid solutions is observed, but focus remains on low-power for remote systems (for example, LoRa-mesh).

2. What is it used for?

- Cables (wired communication): reliable transmission over medium distances (<100 m) with low latency (<1 ms), for indoor/labs (for example, PPS via LVDS for synchronization <1 ns).
- Wireless channels: for field deployments (remote areas, without cables), with low power consumption (for example, LoRa for 5–10 km in rocks), with buffering on connection loss.
- GPS/synchronization: absolute time binding (accuracy <1 μ s) for correlation in multi-module systems; indoor-fallback for caves (atomic clocks/IRIG-B).

Applications:

- muon tomography — track synchronization ($\Delta t < 5$ ns for $\theta < 5^\circ$), transmission ~200 events/min/stack;
- nuclear safety — realtime-alerts;
- geology — remote monitoring of volcanoes and mining.

3. How does it work?

Subsystem has two modes: high-speed (Wi-Fi for demonstrations) and reliable (LoRa for field conditions). Synchronization provides global reference for local counters.

Wired communication channels:

- Types:
- Coaxial (RG-58/RG-6 for RF-signals, attenuation <0.5 dB/m at 1 GHz);
- Fiber optic (LC/SC, <0.2 dB/km, insensitive to EMP);
- Ethernet (CAT5e for data and PPS).

CAT5e: twisted pair (4 pairs, 100 Ohm impedance), BW up to 100 MHz, distance <100 m for Gigabit (TIA-568). For LVDS-PPS: differential transmission ($V_{diff} = 250 - 400$ mV, common-mode rejection >60 dB) over pair (for example, pins 1-2) with 100 Ohm terminator.

Characteristics: skew <1 ns/100 m, jitter <50 ps (shielded version for EMI).

- Principle: electrical pulses (LVDS for PPS: rise time <1 ns, frequency up to 100 MHz) or optical pulses ($\lambda = 1310/1550$ nm, speed 10-100 Gbps). For multi-stack: daisy-chain or star.

Wireless communication channels:

- Wi-Fi (802.11n/ac/ax): speed up to 1 Gbps, radius <100 m, power ~100-500 mW. For demonstrations: 2.4/5/6 GHz, MCS up to 1024-QAM. Latency <10 ms, but unstable in remote-zones.
- LoRa (Long Range): CSS-modulation, frequencies 433/868/915/2400 MHz (ISM), BW 7.8-500 kHz, spreading factor 7-12. Speed 0.3-50 kbps, range 5-15 km LOS, <5 km in rocks. Path loss:

$$L = 32,4 + 20 \log d + 20 \log f,$$

where d in km, f in MHz.

Power <100 mW (TX 14-27 dBm), duty cycle <1 % (EU regulations). Errors <1 % with FEC. In 2025: LoRaWAN with ML for massive IoT (up to 1000 nodes/km²).

- Zigbee (802.15.4): radius 10-100 m, mesh-topology, speed 250 kbps at 2.4 GHz, power <50 mW. Used as fallback in dense networks indoor.
- Principle: radio waves (CSS for LoRa, O-QPSK for Zigbee). Protocol: CBOR/Protobuf (compression > 50 %), ACK+retransmit (timeout 1-10 s).

GPS-synchronization and fallback:

- GPS: 1PPS (TTL/LVDS-pulse width 10-100 ms, every second), accuracy < 1 μ s. NMEA protocol (\$GPRMC for UTC). For outdoor: 1PPS resets local counter, reference frequency 10/20 MHz ensures stability (phase noise <-100 dBc/Hz @ 10 Hz).
- IRIG-B: time code (100 bit/s, frame 1 s), PWM (2/5/8 ms for 0/1/index). Accuracy < 1 μ s. Interface RS-422/BNC, distance < 1 km.
- Atomic clocks: Rb/Cs, drift <10⁻¹¹/s (for example, mRO-50 Rb: accuracy < 5·10⁻¹¹, power < 5 W). Allan deviation $\sigma_y(\tau) < 10^{-12}$ at 1 s.
- Principle: 1PPS/IRIG-B sets the start of counting for local timer:

$$ts = PPS_{count} + fine_{ticks}.$$

4. What technologies are used?

- Cables: twisted pair (CAT5e: 24 AWG, crosstalk < -40 dB @100 MHz); fiber optic (SMF, losses < 0.2 dB/km).
- Wireless: LoRa (SX1276, CSS with BW/SF); Zigbee (CC2530, DSSS); Wi-Fi (ESP32, OFDM).
- Synchronization: GPS (u-blox NEO-M8, 1PPS jitter < 50 ns); IRIG-B (modules B12x); Rb-clocks (mRO-50, ADEV $< 5 \cdot 10^{-11}$).
- 2025: LoRaWAN with ML-routing; hybrid GPS+5G for urban.

5. How many components are needed?

- Cables: 1–2 per stack (CAT5e for PPS and data), 3–6 for multi-stack.
- Wireless: 1 module/stack (LoRa + fallback).
- Synchronization: 1 GPS/IRIG-B/atomic module or master.

6. What do the components consist of?

- Cables: copper pairs (CAT5e: 4 pairs, insulation PVC/PE), connectors (RJ45/M12).
- Wireless: RF-transceiver (SX1276: PA+LNA, quartz 32 MHz), antenna (whip/PCB).
- Synchronization: GPS-receiver (LNA+antenna), IRIG-B decoder (FPGA-logic), Rb-clocks (vapor cell + laser).

7. Cost of components

- CAT5e: ~ 0.5 – 1 USD/m.
- LoRa (SX1276): ~ 5 – 10 USD.
- GPS (u-blox): ~ 20 – 50 USD.
- IRIG-B module: ~ 50 – 200 USD.
- Rb-parts: ~ 500 – 1000 USD (miniature).

8. Additional calculations and features

- Latency: LoRa ~ 1 – 10 s/packet ($T_{air} = \frac{8}{DR} + PL$, where PL — payload, DR — data rate); CAT5e < 1 μ s/100 m.
- Synchronization accuracy: GPS < 1 μ s; IRIG-B < 1 μ s; Rb-holdover < 1 ns/h.
- RX-sensitivity for LoRa:

$$RX_{sens} = -174 + 10 \log(BW) + NF + SNR_{req},$$

where $SNR_{req} = -20$ dB for SF12.

- Power budget:

$$TX + G_{ant} - L_{path} > RX_{sens}.$$

- Features 2025: LoRa with AI for noise suppression; quantum clocks for stability $< 10^{-15}$. In project: dual-radio for durability, SD-buffer > 1 GB for offline-work.

Reconstruction and visualization: Detailed technical description

Reconstruction and visualization — this is the algorithmic and software level of the muon tomograph system, where raw data on muon tracks (timestamps, positions, ToT) are transformed into a 3D-map of object densities/anomalies, and then displayed in an interactive interface. Reconstruction uses tomographic algorithms for inverting the scattering/absorption problem of muons, and visualization — graphics for presenting results (e.g., voxel models, slices). This software runs on laptop/server/tablet (locally or in cloud), with focus on real-time processing for field applications. As of September 2025, key advancements include ML integration for denoising and superresolution (e.g., U-Net/GANs on real datasets), hybrid algorithms with geomatics data (GPR/LiDAR) and adaptive reconstruction for limited-view CT, which improves image quality by 30–50 % in archaeology and geology. Efficiency enhanced by GPU/edge computing, and momentum-dependent reconstruction provides more accurate material ID in nuclear safety.

1. What is it?

Reconstruction — mathematical inversion: from muon trajectories (angles, energies) the density distribution

$\rho(x, y, z)$ of the object is restored.

Visualization — rendering this map in 2D/3D (slices, volume rendering).

Pipeline: fast live (backprojection) → iterative (SIRT/ART) → ML-post (denoising).

Input format: events

$ts, pos, theta, varphi, ToT$;

output: voxel grid (for example, $100 \times 100 \times 100$ voxels, 3–10 cm).

In 2025: adaptive voxelization and generative AI for scenario modeling (for example, "risk of collapse").

2. What is it used for?

- **Reconstruction:** Building 3D-image of voids/anomalies (for example, chambers in pyramids, ores in mines) without destruction. Improvement: ML for superresolution (resolution < 1 cm from noisy data).
- **Visualization:** Interactive viewing (XY/XZ/YZ slices, heat maps, threshold sliders) for analysis by archaeologists/geologists. In nuclear safety — material detection by deviation (MST).
- **Applications:** Archaeology (hidden chambers); geology (voids/faults); nuclear (fuel debris).

3. How does it work?

Pipeline of transformation from raw tracks to 3D-model:

input — N tracks ($\theta, \varphi, pos, energy$);

output — density distribution

$$\rho(r) = \int \mu(r) dr,$$

where μ — absorption/scattering coefficient.

- **Fast live (ray-driven backprojection):**

For each muon ray-sum updates voxels:

$$\rho_v = \sum \frac{w_i L_i}{|A|},$$

where w_i — weight (1 for simple, ToT for energy-weighted), L_i — path in voxel.

Fast display (seconds for demo), but artifacts. Implementation: NumPy/Numba.

- **Iterative (ART/SIRT):**

ART:

$$\rho^{k+1} = \rho^k + \frac{\lambda}{|A|^2} A^T (b - A\rho^k),$$

where A — projection matrix (ray paths), b — measurements (angles/ToT), λ — relaxation (0.1–1).

Convergence in 10–50 iterations, but noise.

SIRT:

$$\rho^{k+1} = \rho^k + \frac{\lambda}{|A|^T} A^T (b - A\rho^k).$$

Better for limited angles.

Regularization: TV (total variation) or Tikhonov.

In 2025: momentum-integrated PoCA for accounting scattering:

$$\theta_{\text{scatter}} \sim \frac{1}{\sqrt{X}} \frac{1}{X_0 p},$$

where p — muon momentum, X_0 — radiation length.

ML-postprocessing: U-Net/TFLite for denoising and superresolution; GANs for simulation generation; clustering for detection masks.

- **Visualization:** Web UI (Three.js, shaders, heatmaps), slices (XY/XZ/YZ), export (PNG/PLY), AR-overlay via Unity.

4. What technologies are used?

- **Reconstruction:** Python (NumPy/Scipy, Numba, TomoPy), C++/CUDA for GPU.
- **ML:** TensorFlow Lite, PyTorch, Geant4 datasets.
- **Visualization:** Three.js/WebGL for 3D, PyQtGraph for desktop.

- In 2025: adaptive ART for limited-view, fusion with geodata (GPR/LiDAR).

5. How many components are needed?

- **Reconstruction:** 3–4 modules (backproj, iterative, ML, fusion).
- **Visualization:** 1 UI + 2–3 exports.
- **For system:** 1 server (RPi/laptop), multithreading for multi-stack.

6. What do the components consist of?

- **Reconstruction:** algorithms (matrix ops, optimizers), voxel grids (ndarray).
- **Visualization:** JS-lib (Three.js: Scene/Camera/Renderer), HTML/CSS for UI.

7. Cost of components

- Software: open-source (NumPy, SciPy, Three.js free).
- ML-training: cloud GPU ~ 100–500 USD.
- Hardware: laptop/server 500–2000 USD.

8. Additional calculations and features

- Parameters: voxel size 3–10 cm, iterations 10–50 to error $< 10^{-3}$.
- Calculations: for $N = 10^5$ tracks, $M = 10^6$ voxels: ART complexity $O(NM)$, ~minutes on CPU, seconds on GPU.
- TV-regularization:

$$\begin{aligned} & \text{norm} \\ & \rho = \\ & \sum \\ & \sqrt{(\rho_x - \rho_{x+1})^2 + \text{var}\epsilon}. \end{aligned}$$

- Features 2025: multi-orientation tomography (several views \rightarrow +accuracy), momentum-dependent recon for material discrimination, 3D reactor imaging.

Thus, reconstruction and visualization form a connected module that combines mathematical algorithms, ML-models and graphical interfaces for transforming raw events of the muon tracker into understandable and practically useful 3D-maps of the object's structure.

Appendices

1. Comparison with analogous methods

Criterion	MuoVis	X-ray / Gamma-scanning	Drilling / Sampling	US-/Seismo-methods
Principle	Natural muons	Ionizing radiation	Physical penetration	Wave methods
Non-destructive	✓	✓	×	✓
Requires radiation source	×	✓	×	×
Penetration depth	>10 m	1-3 m	no limits	1-5 m
Resolution (density)	High on large objects	Good on small	Very high	Medium
Risk for operator	Low	Medium-High	Medium	Low
Mobility	High	Limited	Low	Medium
Environmental impact	Insignificant	Medium	Significant	High

2. Cost of building the system

a) Basic MuoVis configuration (1 block)

Component	Qty	Price per unit, EUR	Sum, EUR
Detector stack (2 scintillators + SiPM)	1	450	450
FPGA + TDC module	1	350	350
Case and mounting	1	150	150
Power (LiFePO4 50 Ah)	1	200	200
Communication (LoRa/Wi-Fi)	1	70	70
Cables, brackets, small elements	-	80	80
Total (block)	-	-	1 300

b) System of 4 blocks (field kit)

Component	Qty	Price per unit, EUR	Sum, EUR
4 × Muoscope 3.0 block	4	1 300	5 200
Centralized processing	1	400	400
Protective cases	4	100	400
Documentation and licenses	-	-	200
Total (kit)			6 200

3. Cost of annual maintenance

Indicator	Value (EUR)	Justification
Regular calibration	300	Once a year
Replacement of consumables	150	SiPM, cables
Software support	250	Updates, bug-fixes
Field maintenance	300	Logistics, setup
Total/year	1 000	

4. Economic benefit (ROI)

Indicator	X-ray/Drilling	MuoVis
Survey cost	20 000 EUR	6 200 EUR
Survey time	5-10 days	1-3 days
Risk of anomalies after	Medium	Low
Downtime cost	2 000 EUR/day	500 EUR/day

Indicator	Traditional methods	MuoVis
Drilling samples	15 000 EUR	0 EUR
Time	>10 days	2-4 days
Object damage	Likely	None
Reputational risk	Medium	Low

5. Service life of key components

Component	Service life	Comment
Scintillators	10+ years	Low wear
SiPM photodetectors	7-10 years	Require protection from temperatures
FPGA + TDC electronics	8-12 years	Reprogramming, updates
LiFePO4 accumulator	5-8 years	300-500 cycles
Case / fasteners	15+ years	UV-stable materials