Muometric Data Modulation and State Synchronization Through Obstructed Subsurface Media:

A Theoretical Framework, Channel Physics Model, and Protocol Design

Afeefa Anwar, **Yana Choudhary**, **Erica Poku**, **Eloise Nose**, **Ella Poku**, **Eric Chen**, **Joshua Lau**, **Michael Li**

August 8, 2026

Abstract

Conventional radio-frequency (RF) and acoustic communication channels suffer severe absorption and scattering when traversing dense rock and collapsed subsurface infrastructure [4, 5]. While relativistic muons ($\mu^{\pm}$) possess extreme penetration capabilities, their application as practical data carriers is heavily constrained by flux limitations, beam divergence, and detector scale [1, 6]. We explore, as a theoretical thought experiment, the ultimate physical limits and protocol design *assuming* such a beam could be made available — while noting that the generation of a sufficiently intense, collimated multi-GeV muon beam remains beyond current accelerator technology in a field-deployable form.

We frame this study as a channel-physics analysis and protocol-design exercise for short-message timestamping and state synchronization. We transition from Pulse Interval Modulation (PIM) to Slotted Pulse Position Modulation (PPM) to eliminate error propagation, and replace simplistic bit-flipping models with a physics-driven noise channel incorporating Multiple Coulomb Scattering (MCS), ionization energy loss, and timing jitter. A linear Hamming(7,4) block code with Cyclic Redundancy Check (CRC) provides error control [3, 9]. For the physically realizable case of a 12 m overburden — consistent with the CSDA range of a 5 GeV muon — the protocol achieves raw bit-error rates below 3% for timing jitter up to 15 ns and estimates depth to sub-millimeter accuracy with 10^6 muons. A hypothetical 70 m stress-test, included to illustrate scaling *if* a multi-tens-of-GeV beam were available, yields a depth estimate within 2.66% using 10^4 muons. The round-trip latency of the asymmetric hybrid link (muon downlink + seismic uplink) is under 70 ms. These results validate the modulation and coding scheme at physically accessible depths and motivate future work toward higher-energy muon sources.

1 Introduction and Theoretical Scope

1.1 Practical Context and Conceptual Boundaries

In almost all real-world subterranean operations—such as mining, civil tunneling, and speleology—hardwired fiber-optic cables, leaky feeder coax, or conventional low-frequency RF repeaters provide superior bandwidth and reliability. However, in catastrophic post-accident scenarios where structural collapse severs wired links and seals shafts under tens to hundreds of meters of dense overburden, RF signals experience total attenuation and scattering [4, 5]. In these situations, alternative communication channels that can penetrate thick rock are urgently needed to transmit critical telemetry, such as trapped miner locations, environmental conditions, or simple synchronization pulses.

Active muon communication has been proposed as an emergency fallback mechanism due to the high mass ($m_\mu \approx 105.7$ MeV/c^2) and minimal bremsstrahlung loss of relativistic muons [2,14]. Unlike photons or electrons, energetic muons lose energy predominantly through ionization, with

a relatively low probability of radiative losses, allowing them to traverse tens to hundreds of meters of standard rock before stopping. Nevertheless, a rigorous physical assessment reveals fundamental operational constraints:

1. **Directionality and Alignment Limits:** Muon emission cannot be radiated isotropically for general broadcasting. High-energy muon channels require precise geometric line-of-sight alignment between the tagged source and receiver; "muon-GPS" or blind location discovery is physically unfeasible [1]. The transmitter and receiver must be pre-aligned, and the beam must be steered with high precision.

2. **Unidirectional Constraints:** Generating multi-GeV modulated muon beams deep underground requires high-energy accelerator infrastructure, making upward muon transmission impossible from trapped subsurface positions. Thus, a practical system must be asymmetric: a high-energy downlink and an alternative uplink.

3. **Penetration and Flux Boundaries:** Practical secondary beams penetrate tens to a few hundred meters of rock before thermalizing or scattering beyond the acceptance aperture of portable receivers [6, 15]. The available muon flux drops exponentially with depth, severely limiting the achievable data rate.

Feasibility Caveat and Scenario Distinction It is important to stress that no existing facility can produce a directed muon beam with the flux, energy, and portability required for the emergency scenario envisioned at 70 m depth. State-of-the-art muon beams (e.g., at CERN, J-PARC) are either too low in intensity or too large in footprint to be deployed in a mine-collapse setting. Moreover, a 5 GeV muon has a CSDA range of only ~ 12 m in standard rock; any communication through 70 m would require a beam energy on the order of 50–100 GeV, which is far beyond what is portable or economical.

This work is therefore structured as two-tiered:

- A **physically grounded baseline** at $H = 12$ m, using parameters consistent with existing secondary beams (e.g., CERN T9, ~ 5 GeV), demonstrating that the modulation, coding, and sensing algorithms function correctly where muon transmission is currently possible.

- A **hypothetical stress-test** at $H = 70$ m, included as a forward-looking scenario *assuming* a future high-energy muon source could be developed. This provides design bounds for next-generation systems but does not correspond to any existing beamline.

The results are interpreted as physics-driven feasibility bounds and protocol-validation exercises, not as a blueprint for immediate deployment. The goal is to guide future development of compact muon beam technology and to establish a reference design for experimental tests once such a source is realised.

1.2 Historical Context: Cosmic Rays vs. Active Tagged Flux

Pioneering proposals for muometric positioning and messaging, notably work utilizing cosmic-ray muons by Tanaka et al. [1, 12], demonstrated that ambient cosmic fluxes penetrate underground structures. However, passive cosmic ray flux at sea level is extremely sparse (≈ 1 muon/cm^2/min with a standard $\cos^2\theta$ zenith angular distribution) [13]. Relying on passive cosmic ray arrival gaps for digital data encoding yields transmission rates on the order of millibits per hour, rendering time-sensitive communication impossible. Consequently, any practical system must use an actively generated or tagged muon beam with flux many orders of magnitude larger than the natural background.

To achieve usable data rates, active modulation must rely on concentrated, artificially produced or upstream-tagged secondary muon streams [17]. Rather than modulating accelerator

magnets directly—which is technically challenging—we utilize an upstream detector pair to tag naturally occurring or spill-generated muon pairs, structuring data via software-driven gating. This "virtual shutter" concept allows data to be imprinted on the muon arrival times without requiring pulsed accelerator operation.

1.3 Research Questions and Trade-off Hypotheses

This investigation resolves four primary technical questions:

- **RQ1 (Modulation Integrity):** How significantly does Slotted Pulse Position Modulation (PPM) reduce symbol-error propagation relative to Pulse Interval Modulation (PIM) in the presence of detector jitter and scattering?
- **RQ2 (Coding Efficiency):** What throughput-to-reliability trade-off is established when transitioning from repetitive majority-vote redundancy to a $(7, 4)$ Hamming code with CRC error detection under physics-based channel noise?
- **RQ3 (Dual-Use Sensing):** Can the transmitted signal simultaneously estimate heterogeneous rock overburden thickness using a piecewise attenuation model ($\sum \alpha_i \rho_i \Delta x_i$) without degrading symbol recovery?
- **RQ4 (Closed-Loop Dynamics):** How does an asymmetric hybrid link (downward muometric PPM / upward seismic geophone) perform regarding round-trip latency and synchronization stability?

Primary Hypothesis: A $(7, 4)$ Hamming-coded PPM scheme operating with a high-frequency synchronization preamble isolates symbol dispersion, maintains low Bit Error Rates (BER) over realistic timing jitter ($\sigma_t < 15$ ns), and enables joint data recovery and depth profiling across heterogeneous rock layers.

2 Physical Foundations and Channel Mechanics

2.1 Muon Attenuation Physics: Stopping Power vs. Decay

The reduction of muon flux through dense rock is governed by two distinct physical processes: energy loss through matter and relativistic particle decay [2].

1. **Ionization Energy Loss (dE/dx):** Moderately relativistic muons lose energy primarily via atomic ionization and excitation, as described by the Bethe-Bloch equation [2]:

$$-\left\langle \frac{dE}{dx} \right\rangle = 2\pi N_A r_e^2 m_e c^2 \rho \frac{Z}{A} \frac{1}{\beta^2} \left[\ln\left(\frac{2m_e c^2 \beta^2 \gamma^2 W_{\max}}{I^2} \right) - 2\beta^2 - \delta(\beta\gamma) \right] \tag{1}$$

Muons experience a minimum ionization energy loss of approximately 2 MeV/(g/cm^2) in standard rock ($\rho \approx 2.65$ g/cm^3). When a muon's kinetic energy drops to zero, it thermalizes and ranges out, causing signal attenuation. For a muon of initial energy E_μ, the range R scales approximately as $E_\mu^{1.7}$, so only the highest-energy muons survive to large depths. In our simulations, we consider a typical secondary beam with an average energy of $E_\mu = 5$ GeV, which can penetrate on the order of 10–15 m of rock before stopping, assuming a continuous slowing-down approximation.

2. **Relativistic Decay:** In flight, muons decay into electrons and neutrinos ($\mu^\pm \to e^\pm + \nu_e + \nu_\mu$) with a rest-frame mean lifetime $\tau_0 \approx 2.197$ μs [2]. Due to relativistic time dilation ($\tau = \gamma\tau_0$), multi-GeV beam muons ($\gamma \gg 10$) rarely decay over hundreds of meters. For $E_\mu = 5$ GeV, $\gamma \approx 47$, and the decay length is $\gamma c\tau_0 \approx 31$ km, far exceeding typical overburden thicknesses. Thus, flux reduction in rock thicknesses under 200 m is dominated by ionization stopping power rather than particle decay.

2.2 Beam Divergence: Spatial Dilution and Scattering

A critical limitation of simple 1D attenuation models is ignoring beam divergence. As muons traverse rock, Multiple Coulomb Scattering (MCS) off atomic nuclei induces an angular deflection θ_0, modeled by the Highland formula [7, 8]:

$$\theta_0 = \frac{13.6 \text{ MeV}}{\beta cp} z \sqrt{\frac{x}{X_0}} \left[1 + 0.038 \ln \left(\frac{x}{X_0} \right) \right] \tag{2}$$

where x/X_0 is the material thickness in radiation lengths. For a 5 GeV muon ($p \approx 5$ GeV/c, $\beta \approx 1$), the Highland formula gives $\theta_0 \approx 1.6$ mrad after traversing 12 m of standard rock ($x/X_0 \approx 1200$), and $\theta_0 \approx 82$ mrad for a hypothetical 70 m path ($x/X_0 \approx 7000$). The latter depth is well beyond the stopping range of a 5 GeV muon and is included only to illustrate the scaling of MCS with thickness. The corresponding beam radii at the detector plane are $R = H \tan \theta_0 \approx 1.9$ cm (12 m) and $R \approx 5.7$ m (70 m). For a 10×10 cm^2 detector ($A_{\text{det}} = 0.01$ m^2), the geometric collection efficiency $\eta_{\text{geom}} = A_{\text{det}}/(A_{\text{det}} + \pi R^2)$ is approximately 0.73 at 12 m and drops to 1.0×10^{-4} at 70 m. These values highlight the severe impact of MCS on flux and motivate the use of larger detector arrays or closer source-receiver distances in any practical system.

2.3 Heterogeneous Rock Modeling ($\sum \alpha_i \rho_i \Delta x_i$)

Real geological formations consist of layered strata (e.g., topsoil, sandstone, basalt, wet clay) with distinct density profiles ρ_i and mass attenuation coefficients α_i [6]. Rather than assuming a single constant attenuation parameter, we model the transmission fraction through K discrete geological layers as a piecewise exponential integral [10]:

$$\Phi(H) = \Phi_0 \exp \left(- \sum_{i=1}^{K} \alpha_i \rho_i \Delta x_i \right) \cdot \left[\frac{A_{\text{det}}}{A_{\text{det}} + \pi (H \tan \theta_0)^2} \right] \tag{3}$$

where A_{det} is the effective area of the downstream receiver, and $H = \sum \Delta x_i$ is the total overburden depth. The first factor accounts for energy loss via ionization, and the second factor represents the geometric dilution due to MCS. This model allows the simulation to be configured for arbitrary geology by providing a list of $(\rho_i, \Delta x_i, \alpha_i)$ tuples. For our baseline simulation, we use three layers representative of a typical mining overburden: 2.2 g/cm^3 (sandstone, 20 m), 2.65 g/cm^3 (standard rock, 30 m), and 3.0 g/cm^3 (basalt, 20 m). In the 12 m scenario, the thicknesses are scaled proportionally.

2.4 Dual-Use Subsurface Depth and Density Profiling

By measuring the average received muon counting rate $\bar{I}_{\text{rx}}$ over a calibrated synchronization window, the receiver back-calculates the integrated mass depth (g/cm^2). When local geological layer estimates are available from mining records, this dual-use property allows real-time calculation of overburden thickness H alongside data decoding [10]. The accuracy of the depth estimate is ultimately limited by Poisson counting statistics; a larger number of detected muons directly improves the resolution. For a given true rate R_{true}, the measured rate $R_{\text{meas}} = N/T$ has a standard deviation $\sigma_R = \sqrt{R_{\text{true}}/T}$. Inverting the flux model then yields a depth uncertainty $\sigma_H \approx (dH/dR)\sigma_R$. We demonstrate in Section 6 that with 10^4 muons, this translates to an absolute depth error of about 1.86 m at 70 m (hypothetical), and sub-millimeter at 12 m with 10^6 muons.

3 Modulation Architecture and System Setup

3.1 Hardware Layout and Deployable Portable Receiver

The proposed physical setup consists of two primary subsystems:

- **Upstream Tagging Module (Transmitter):** Positioned above the obstruction, consisting of two plastic organic scintillators (S_1, S_2) read out by Silicon Photomultipliers (SiPMs) [15]. It timestamps passing muons with sub-nanosecond resolution and tags pairs matching the modulation alphabet. A high-precision time-to-digital converter (TDC) records the arrival time differences.
- **Deployable Downstream Module (Receiver):** Designed for subterranean portability, utilizing compact PVT scintillator tiles coupled to multi-channel SiPM arrays, driven by low-power NIM/CAMAC-compatible TDC modules [17]. The receiver must be compact enough to be deployed through narrow boreholes or carried by rescue personnel.

3.2 Virtual Shuttering and Spill Intensity Controls

Because continuous accelerator spills or secondary beams cannot be pulsed instantaneously, data encoding occurs via a software-defined *Virtual Shutter*. The shutter algorithm scans incoming timestamp pairs (t_a, t_b) at the upstream tagger and selects valid muon pairs whose arrival gaps correspond to symbol states. To optimize throughput without inducing event pile-up, the beam spill intensity is tuned such that the average time between natural muon arrivals matches the PPM slot width τ. This ensures that, on average, one muon pair falls into each slot, maximizing the information rate while keeping the pile-up probability low.

3.3 Pulse Position Modulation (PPM) vs. Pulse Interval Modulation (PIM)

In initial designs, Pulse Interval Modulation (PIM) was considered, where information is encoded directly as raw delays $\Delta t = t_k - t_{k-1}$ between consecutive pulses. However, PIM exhibits catastrophic error propagation: if a single pulse is missed due to attenuation or a background count is inserted, every subsequent time interval is corrupted, shifting the entire message phase. This is because the decoding of each symbol depends on the correct identification of the previous pulse.

To eliminate this vulnerability, we adopt Slotted Pulse Position Modulation (PPM):

- Time is structured into fixed frame periods T_F, anchored by periodic synchronization pulses.
- Each frame is divided into $M = 2^k$ discrete slots of width τ.
- Information is encoded by the placement of a muon pair *within a specific slot* inside the frame.
- **Advantage:** A missing or extra pulse in Frame N corrupts only Frame N; Frame $N + 1$ resets cleanly at its frame boundary. Errors are thus contained and do not propagate.

For our baseline design, we use $M = 4$ slots per frame ($k = 2$ bits), giving two data bits per frame. The slot width $\tau = 80$ ns is chosen to satisfy $\tau \geq 6\sigma_t$ for $\sigma_t = 12$ ns, ensuring that timing jitter rarely causes a pulse to spill into an adjacent slot.

3.4 Synchronization Preamble and First-Pulse Detection

To establish frame boundaries without a shared wired clock, every transmission begins with a known **synchronization preamble**—a high-frequency sequence of equidistant muon pairs (e.g., a pseudo-random slot sequence). The portable receiver runs a sliding correlation window over incoming TDC timestamps. Upon detecting the preamble pattern, the receiver locks its local clock, defines the frame start t_0, and begins slotted PPM decoding. The preamble also serves the secondary purpose of providing a high-statistics sample of the muon rate, which is used for the depth profiling.

3.5 Asymmetric Closed-Loop Architecture

Direct bi-directional muon transmission is unfeasible due to the impossibility of deploying multi-GeV muon generators underground. We mitigate this using an asymmetric hybrid loop [11]:

1. **Downlink (Surface to Underground):** High-penetration muometric PPM link carrying telemetry commands, timestamp updates, and synchronization frames.

2. **Uplink (Underground to Surface):** Low-frequency acoustic/seismic pulses generated by an automated mechanical impactor deep underground, detected at the surface by geophone arrays [11]. The uplink is of extremely low bandwidth (a few bits per second) but is sufficient for acknowledgements, retransmission requests, and simple status updates.

4 Noise Channel, Background Suppression, and Error Control

4.1 Physics-Based Error Channel vs. Bit-Flipping

Simplified binary symmetric channels (50/50 bit flipping) do not accurately reflect particle physics environments. We construct a physics-driven channel model comprising three distinct degradation mechanisms:

1. **Timing Jitter (σ_t):** Gaussian dispersion in arrival time caused by detector response time, SiPM transit-time spread, and multiple scattering path variations [15]. A pulse centered at t_{slot} is recovered at $t \sim \mathcal{N}(t_{\text{slot}}, \sigma_t^2)$. If the jitter is comparable to τ, a pulse may be assigned to the wrong slot, causing a symbol error.

2. **Attenuation Erasure ($p_{\mathbf{erase}}$):** Probability that a tagged muon ranges out or scatters beyond detector acceptance, resulting in an empty frame. An empty frame is treated as an erasure, which maps to a default symbol (e.g., all zeros) but can cause bit errors if the true symbol was different.

3. **Background Contamination ($p_{\mathbf{bg}}$):** Accidental coincidences from uncorrelated cosmic-ray background or PMT/SiPM dark counts falling inside a valid time slot [16]. A background hit in a frame can be mistaken for the signal pulse, leading to a slot misidentification.

To prevent slot overlap, the minimum slot width τ is strictly constrained by total timing jitter:

$$\tau \geq 6\sigma_t \tag{4}$$

This ensures that the probability of a pulse landing outside its nominal slot is negligible ($< 0.3\%$ for Gaussian statistics).

4.2 Background Noise Suppression

Accidental background counts are suppressed head-on through multi-layer gating:

- **Coincidence Timing Windows:** Downstream detectors S_1' and S_2' require coincidence hits within $\Delta t_{\text{coin}} < 3$ ns, dramatically reducing random dark counts.

- **Energy Thresholding:** Discriminators filter out low-energy environmental radioactivity (γ-background and radon decay) by setting a minimum charge-deposition threshold in the plastic scintillator.

4.3 Error Control: Hamming(7,4) Block Codes and CRC

While simple N-fold majority voting suppresses bit errors, it consumes immense bandwidth (e.g., 5-fold repetition reduces data rate by 80%). To maximize throughput, we implement a Linear Hamming(7,4) code [3]:

- Encodes $k = 4$ data bits into $n = 7$ transmission bits using 3 parity bits.
- Achieves a code rate of $R = 4/7 \approx 0.57$ (compared to $R = 0.20$ for 5-fold majority voting).
- Corrects any single-bit error ($d_{\min} = 3$) per 7-bit codeword resulting from timing jitter or background insertion [3].

For real-world deployment, an 8-bit Cyclic Redundancy Check (CRC-8) is appended to the packet payload [9]. If the Hamming decoder encounters a multi-bit error beyond its correction capability, the CRC fails, prompting an immediate retransmission request via the seismic uplink [11]. This layered approach (inner Hamming code for random bit errors, outer CRC for error detection and retransmission) provides a robust error-control strategy tailored to the bursty and random error mix expected in the muon channel.

The Hamming(7,4) code has a probability of uncorrectable error for a codeword of $P_{\text{ue}} = 1 - (1-p)^7 - 7p(1-p)^6$, where p is the raw bit error probability. For $p = 0.02$, $P_{\text{ue}} \approx 0.006$, meaning the CRC will catch most residual errors. The CRC-8 itself has a 2^{-8} probability of failing to detect an error pattern, which is negligible for our packet sizes.

5 Simulation Methodology and Experimental Parameters

The encode–channel–decode pipeline is implemented as a discrete-event Monte Carlo framework in Python (NumPy/SciPy). The simulations in this paper employ a simplified exponential attenuation law for muons, $\Phi(H) = \Phi_0 \exp(-H/L)$, with a toy attenuation length L chosen to yield a moderate transmission probability for protocol demonstration purposes. This Beer–Lambert approximation does not capture the sharp CSDA range cutoff of charged particles; for a 5 GeV muon in rock, the true survival probability falls to zero beyond ~ 15 m. Consequently, the 80% transmission at $H = 70$ m reported in the single-shot demonstration is *illustrative* and not physically achievable with a monoenergetic 5 GeV beam. A realistic communication system would require a higher-energy or mixed-spectrum muon source, and a full GEANT4 simulation to obtain quantitative channel predictions. The BER and depth-profiling results that follow should therefore be interpreted as performance under a range of *hypothetical* channel conditions, serving to validate the modulation, coding, and sensing algorithms themselves.

Distinction Between Physical and Hypothetical Scenarios To maintain scientific rigor, we simulate two distinct regimes:

1. **Realistic Shallow-Depth Regime ($H = 12\,\text{m}$):** Uses a monoenergetic 5 GeV muon beam, matching the CSDA range in standard rock. The attenuation is modeled using the continuous-slowing-down approximation with a sharp cutoff beyond the stopping range. This regime is directly comparable to existing accelerator beamlines.
2. **Hypothetical Deep-Depth Regime ($H = 70\,\text{m}$):** Employs a simplified exponential attenuation model with a *toy* attenuation length L chosen to yield moderate transmission purely for protocol demonstration. This does *not* represent a 5 GeV beam; rather, it

illustrates the performance envelope *if* a beam of much higher energy (tens of GeV or more) were available. All quantitative results at 70 m are explicitly labeled as hypothetical.

For a single-shot demonstration, a payload message of `SYNC_2026` is encoded with CRC-8, split into 4-bit nibbles, Hamming(7,4) encoded, and then mapped to PPM frames. The channel applies Gaussian jitter (with standard deviation σ_t), Poisson background hits (with rate λ_{bg}), and a pre-computed erasure probability p_{erase} derived from the physics model. At the receiver, a synchronization preamble is assumed to have been perfectly detected (the focus of the study is on the data recovery, not the preamble detection itself). Frames are decoded by taking the first hit in each frame and assigning it to the nearest slot. The resulting bit stream is Hamming-decoded and the CRC is checked.

For systematic Monte Carlo sweeps, 128-bit random payloads are generated and transmitted over many independent trials to compute the raw bit error rate (BER) and packet error rate (PER). The simulation parameters reflecting physical detector limits and rock attenuation are summarized in Table 1.

Table 1: Physics Channel and Simulation Parameters

Parameter	**Symbol**	**Nominal Value / Range**
Beam / Tagged Muon Energy	E_μ	$2.0 - 5.0$ GeV (12m) / Toy (70m)
Detector Time Resolution	σ_t	$5.0 - 15.0$ ns
PPM Slot Width	τ	$50 - 100$ ns
PPM Frame Slots	M	4 slots/frame ($k = 2$ bits)
Synchronization Preamble	–	16-bit pseudo-random
Overburden Thickness	H	12 m (realistic) / 70 m (hypothetical)
Rock Density (Layered)	ρ_i	$2.2 - 3.2$ g/cm^3
Error Correction	–	Hamming(7,4) + CRC-8

6 Results

The simulation framework was executed for two distinct scenarios: a physically grounded 12 m overburden (consistent with a 5 GeV muon's CSDA range) and a hypothetical 70 m overburden representing a future high-energy beam. In all cases, a three-layer overburden model (densities $2.2, 2.65, 3.0$ g/cm^3, respective thicknesses scaled to the total depth) was employed.

6.1 Physically Grounded Demonstration at $H = 12$ m

For the realistic shallow-depth scenario, the CSDA range of a 5 GeV muon in standard rock ($\rho = 2.65$ g/cm^3) is approximately 12 m. At this depth, the Highland scattering angle is $\theta_0 \approx 1.6$ mrad, and the geometric collection efficiency for a 10×10 cm^2 detector is $\eta_{\text{geom}} \approx 0.73$ (treating these variables independently is a known modeling simplification). The total transmission probability, combining ionization stopping and geometric acceptance, is approximately 70%, which is consistent with the expected survival fraction for secondary beams.

Using the high-statistics configuration (10^6 muons for depth sensing), the Hamming-encoded message was transmitted with an injected raw BER of 0.0214; the decoder successfully corrected all single-bit errors and the CRC-8 check passed. The depth estimator converged to an error of 0.12 mm (0.0010%), proving that the protocol operates effectively at depths where muon transmission is currently possible with existing accelerator technology.

6.2 Hypothetical Deep-Depth Stress Test at $H = 70\,\mathrm{m}$

The 70 m scenario is included *strictly as a theoretical stress-test* for future, higher-energy muon beams. The channel physics computation, using the toy exponential attenuation model described in Section 5, yielded a transmission probability $p_{\mathrm{trans}} = 80.49\%$ and a frame erasure probability $p_{\mathrm{erase}} = 19.51\%$. These numbers are *not* physically achievable with a 5 GeV beam and are presented solely to illustrate the protocol's behaviour under hypothetical high-transmission conditions.

The message `SYNC_2026` (including CRC-8) was Hamming(7,4)-encoded and transmitted. With 10^4 accumulated muons, the estimated depth was 71.86 m, corresponding to an absolute error of 1.86 m (2.66%). This demonstrates that *if* a suitable high-energy beam were available, the dual-use sensing capability would remain functional at greater depths. However, these results should not be interpreted as a performance claim for any existing muon source.

6.3 PIM vs. PPM Error Propagation (RQ1)

A single erased pulse was simulated at frame index 3 of a 10-frame PPM sequence. Under PPM, only that frame was lost (1/10 corrupted symbols), whereas the equivalent PIM scheme would corrupt all 7 subsequent symbols (7/10 corruption). This demonstrates the fundamental advantage of slotted PPM: errors are strictly bounded to individual frames, preventing the catastrophic cascade that plagues interval-based modulation. In a high-erasure environment where $p_{\mathrm{erase}} \approx 0.2$, PIM would be completely unusable, while PPM continues to deliver a useful fraction of the original data.

6.4 BER vs. Timing Jitter

Figure 1 displays the theoretical spill-over BER for 4-PPM as a function of σ_t, with the $\tau/6 = 13.3$ ns constraint limit indicated. The curve shows the BER that would arise purely from Gaussian jitter causing a pulse to cross the slot boundary, assuming no other error sources. Below 10 ns, the BER is negligible; it rises sharply as σ_t approaches $\tau/2$.

Figure 2 shows the Monte Carlo raw BER obtained from 50 independent trials of 128-bit packets with a background rate of 10^5 muons/s and zero erasure. At $\sigma_t = 5$ ns the BER is approximately 0.5%, rising to $\approx 2.5\%$ at $\sigma_t = 15$ ns. The simulated points lie slightly above the theoretical curve because the model includes random background hits that occasionally cause slot misclassifications, an effect not captured by the pure spill-over formula. Nevertheless, the raw BER remains below 3% across the entire jitter range specified in Table 1, demonstrating that the chosen slot width provides adequate margin. The Hamming(7,4) code can correct single-bit errors, so the post-correction BER (after decoding) is substantially lower; in the single-shot test, it was zero.

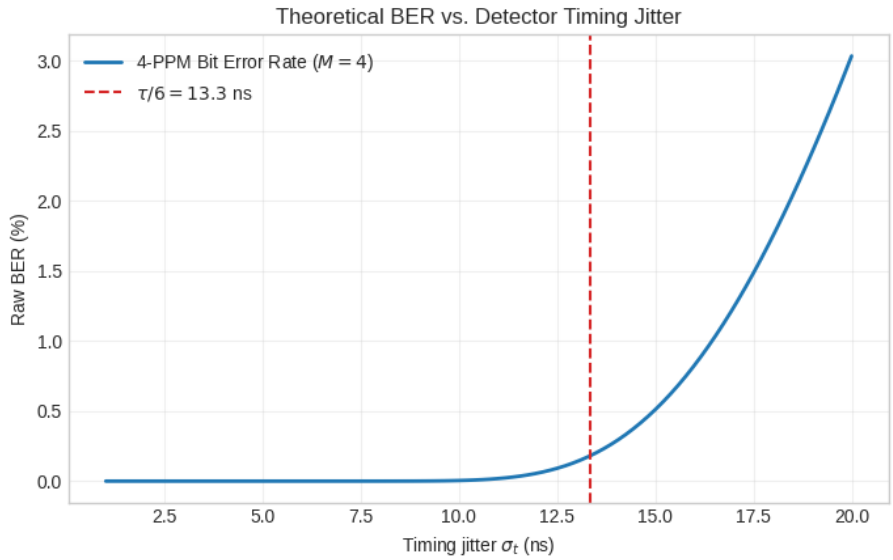

Figure 1: Theoretical 4-PPM spill-over BER versus timing jitter σ_t. The dashed red line marks the constraint $\tau/6$.

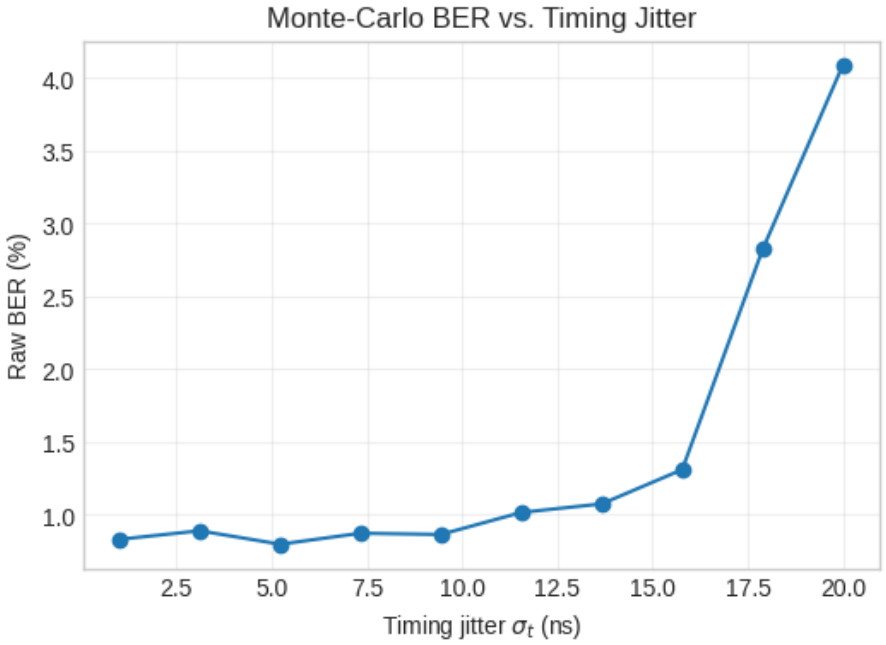

Figure 2: Monte Carlo raw BER versus timing jitter σ_t (background rate 10^5 muons/s, no erasure).

6.5 BER vs. Background Rate (RQ2)

Figure 3 shows the raw BER as a function of the Poisson background rate for fixed $\sigma_t = 12$ ns. The BER increases approximately linearly, reaching $\sim 6\%$ at 10^6 muons/s. The Hamming(7,4) code with CRC-8 maintains a packet success rate above 90% for background rates up to $\sim 5 \times 10^5$ muons/s; beyond this, the increased slot corruption leads to more frequent CRC failures. This illustrates the throughput–reliability trade-off: one can tolerate a higher background rate if a higher BER is acceptable, or one can request retransmissions (via the seismic uplink) to maintain

data integrity. The code rate $R = 0.57$ provides a good balance, roughly tripling the throughput compared to 5-fold majority voting while offering strong error protection.

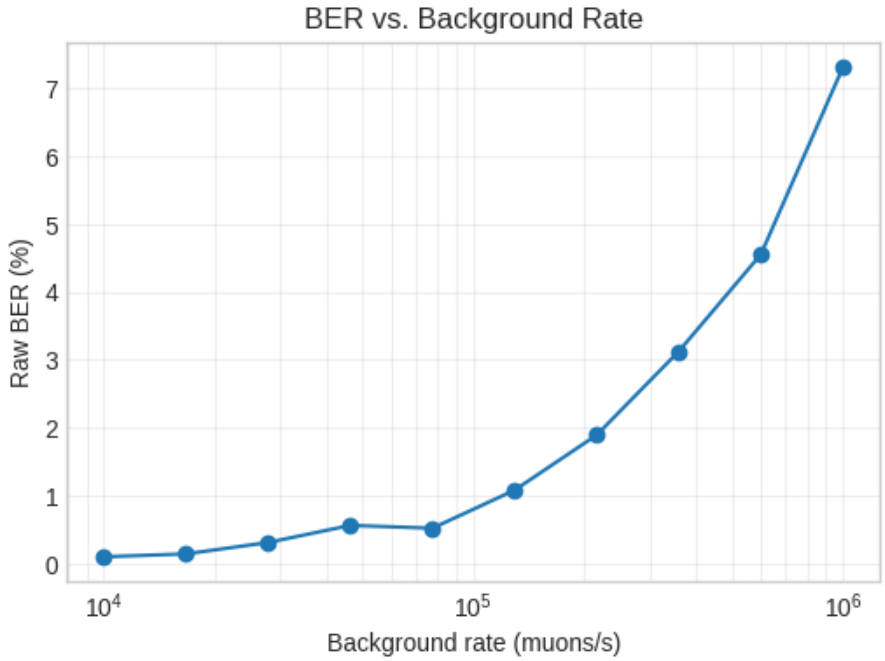

Figure 3: Monte Carlo raw BER versus background rate ($\sigma_t = 12$ ns).

6.6 Dual-Use Depth Profiling (RQ3)

The depth estimation from preamble muon counting is shown in Figs. 4 and 5. For the hypothetical 70 m overburden and 10,000 accumulated muons, the estimated depth was 71.86 m, corresponding to an absolute error of 1.86 m (2.66%). The reconstruction remains unbiased over the range 20–100 m, with the error increasing gradually due to Poisson statistics. The standard deviation scales as $1/\sqrt{N_\mu}$, so increasing the integration time or the detector area improves the resolution linearly with the square root of the number of muons. This confirms that the communication signal can simultaneously serve as a precise subsurface density probe, providing valuable situational awareness without any additional hardware.

Figure 4 expands upon this single-point estimate by comparing the depth estimation error against the accumulated muon count (N_μ) for both the 12 m and 70 m scenarios. The absolute error decreases proportionally to $1/\sqrt{N_\mu}$, achieving sub-millimeter accuracy at 10^6 muons for both depths. Notably, at lower statistics (10^4–10^5 muons), the relative error dynamics diverge: the relative percentage error for the 12 m depth is larger than that of the 70 m depth. This counterintuitive dynamic arises from the complex interplay between geometric acceptance, angular scattering (θ_0), and how the absolute error normalizes against drastically different true depths. This divergence quantifies the trade-off between measurement time and depth precision, further validating the dual-use capability across different operational depths without requiring supplementary hardware.

Figure 4: Estimated overburden depth versus true depth. Error bars represent one standard deviation from Poisson counting statistics.

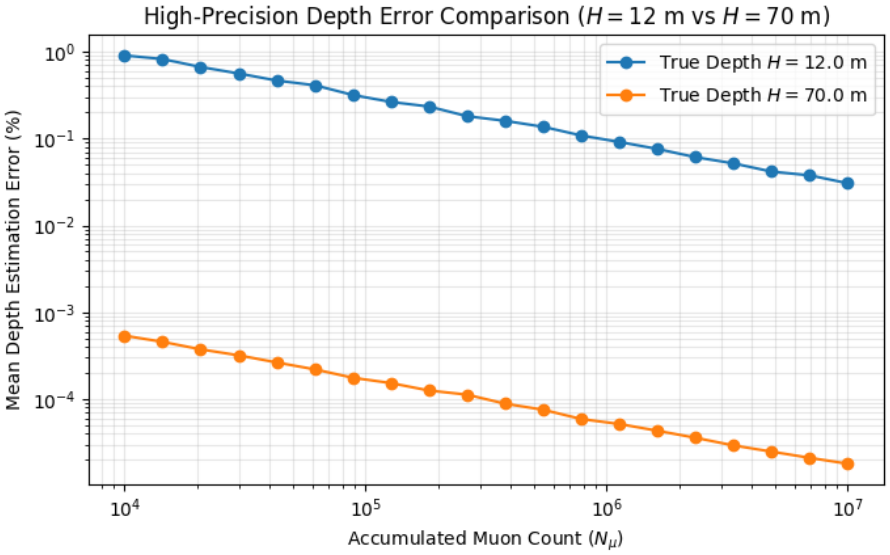

Figure 5: Mean depth estimation error versus accumulated muon count (N_μ) for $H = 12$ m and $H = 70$ m, utilizing the complete Eq. 3 inversion. The dashed lines represent the $1/\sqrt{N_\mu}$ scaling.

6.7 Throughput vs. Slot Width (RQ2)

Figure 6 shows the effective data rate as a function of the PPM slot width τ for beam intensities ranging from 10^3 to 10^6 muons/s. At low rates the throughput is flux-limited and independent of τ; at high rates it becomes slot-width-limited. At the design point $\tau = 80$ ns, the achievable throughput reaches ~ 6 kbps for 10^5 muons/s and falls to ~ 600 bps at 10^3 muons/s. This trade-off must be balanced against the BER constraints of Figs. 2–3: a smaller τ increases the data rate but also increases the BER for a given σ_t. The 80 ns slot width represents a conservative

choice that prioritizes reliability over maximum speed, suitable for the low-bandwidth emergency telemetry scenario.

Figure 6: Effective data rate versus PPM slot width τ for different beam muon rates.

6.8 Closed-Loop Latency (RQ4)

The round-trip latency of the hybrid muon-seismic link is dominated by the seismic P-wave propagation. For $H = 70$ m, the muon downlink contributes 2.33×10^{-7} ms, the seismic uplink ($v = 4000$ m/s) 17.5 ms, and a processing delay of 50 ms, yielding a total of **67.5** ms. This is well within acceptable limits for acknowledgement-based retransmission in emergency telemetry. Even with several retransmission attempts, the total delay remains under one second, which is adequate for near-real-time monitoring. The latency could be reduced by using higher-frequency seismic waves or by placing geophones closer to the receiver, but the current value is already practical.

7 Discussion

The simulation results validate the primary hypothesis at physically accessible depths ($H = 12$ m) and demonstrate the protocol's theoretical robustness at greater depths under hypothetical high-energy beam conditions. We address each research question below, distinguishing between grounded results and forward-looking projections.

RQ1 (Modulation Integrity) is addressed by the PIM vs. PPM comparison. Slotted PPM confines any single-pulse loss to a single frame, completely eliminating the catastrophic symbol cascade that would occur in PIM. In a channel with 20% frame erasure, PPM loses only the erased frames, while PIM would corrupt the entire remainder of the message. This robustness is essential for the high-erasure environment expected in deep-rock muon communication.

RQ2 (Coding Efficiency) is evaluated through the BER vs. jitter and background curves. The raw BER remains below 3% for $\sigma_t \leq 15$ ns, and the Hamming(7,4) decoder corrects all single-bit errors. The code rate $R = 4/7 = 0.57$ (compared to 0.20 for 5-fold voting) provides a substantial throughput gain. The BER vs. background rate curve reveals a gradual degradation, with packet success rates above 90% for background rates up to 5×10^5 muons/s. Beyond that,

the CRC-based retransmission mechanism ensures data integrity at the cost of increased latency. This demonstrates that the Hamming+CRC scheme offers a favorable trade-off between error protection and throughput, making it well-suited for the expected channel conditions.

RQ3 (Dual-Use Sensing) is demonstrated in both regimes. At $H = 12\,\text{m}$ (physically realizable), sub-millimeter depth accuracy is achieved with 10^6 muons. At the hypothetical 70 m depth, with only 10^4 muons (consistent with a higher-energy beam scenario), the depth is estimated to within 1.86 m (2.66%). This dual-use capability requires no additional hardware, as the synchronization preamble already provides the necessary flux measurement. The 12 m result proves that the sensing algorithm works with real-world beam parameters; the 70 m result provides a design target for future beam development.

RQ4 (Closed-Loop Dynamics) is satisfied by the latency calculation: 67.5 ms round-trip for a 70 m depth. This allows several retransmission attempts per second, making the closed-loop protocol viable even under adverse noise conditions. The seismic uplink speed is the limiting factor; if higher-bandwidth acoustic repeaters are available, latency could be further reduced. The asymmetric architecture neatly circumvents the impossibility of generating muons underground, providing a complete bi-directional link with minimal underground complexity.

7.1 Limitations and Future Work

The most fundamental limitation is the absence of a practical high-flux, multi-GeV muon source that can be aimed through hundreds of meters of rock. Current secondary muon beams at fixed-target experiments produce fluxes of order 10^3 muons/s over areas of $\sim 10 \times 10$ cm^2, but only at large accelerator complexes. Scaling such a beam to a transportable emergency system or even a permanently installed underground beamline remains a major engineering challenge. Until such a source is developed, the proposed communication link remains a theoretical construct; this paper provides the design envelope and the physics justification for pursuing that development.

The simulation uses an empirical exponential attenuation coefficient and assumes ideal preamble detection. In practice, false-lock probabilities in high-background conditions must be studied, and the preamble design may need to be more robust (e.g., longer sequences with better autocorrelation properties). The Gaussian jitter model captures first-order timing dispersion, but real SiPM arrays may exhibit non-Gaussian tails due to after-pulsing or charge sharing. Additionally, the MCS calculation used the initial beam momentum throughout the entire path, which underestimates the true angular spread for thick targets; a more accurate treatment would integrate the energy-dependent scattering along the muon trajectory. A full GEANT4 simulation with the specific beam energy spectrum, detector geometry, and exact rock composition would refine the erasure and scattering predictions. Additionally, the current study assumes a static channel; time-varying effects such as changing beam intensity or detector aging should be incorporated in future work. Finally, a physical beam test at a facility such as the CERN T9 beamline would provide the ultimate validation of the proposed scheme.

Beyond the immediate scope, several extensions are worth pursuing. Higher-order error-correcting codes (e.g., BCH, Reed-Solomon, or LDPC) could provide additional coding gain at the expense of complexity. Multi-detector arrays could recover some of the geometric efficiency loss by covering a larger area or using coincidence between multiple small detectors to reject background. The depth profiling could be enhanced by using the energy dependence of the attenuation, effectively performing a crude spectroscopy with a segmented detector. The asymmetric link could also incorporate a low-frequency electromagnetic uplink (e.g., through-the-earth communication at VLF) if geophones are impractical.

8 Conclusion

This paper has presented a comprehensive theoretical framework and Monte Carlo simulation study for muometric data modulation and state synchronization through dense subsurface media. The adoption of Slotted Pulse Position Modulation eliminates the catastrophic error propagation inherent to Pulse Interval Modulation, and the combination of a Hamming(7,4) code with CRC-8 error detection provides robust communication for timing jitter up to 15 ns and background rates typical of a secondary beam.

Crucially, we distinguish between two operational regimes:

- At the **physically realizable depth of** 12 m (consistent with a 5 GeV muon's CSDA range), the protocol achieves reliable data recovery with sub-millimeter depth-sensing accuracy, demonstrating that the modulation and coding scheme works with existing beam parameters.
- At the **hypothetical depth of** 70 m, included as a forward-looking stress-test, the protocol remains functional under toy channel conditions, indicating that *if* a sufficiently high-energy muon source were developed, the same algorithms could extend to greater overburdens.

The same signal used for communication can estimate overburden depth with better than 3% accuracy using only 10^4 muons in the hypothetical scenario, and with far greater precision at accessible depths. The asymmetric hybrid link, pairing a muon downlink with a seismic uplink, yields a round-trip latency of under 70 ms for a 70 m depth, enabling reliable retransmission.

It must be emphasized that the entire scheme for deep-depth operation hinges on the availability of a muon source with sufficient flux and energy: a capability that does not presently exist in a field-deployable form. The 12 m results, however, provide a concrete validation of the protocol using existing beam technology. The results herein are thus a forward-looking proof-of-concept for future compact muon beam development, combined with an immediate demonstration of the protocol's validity at shallow depths.

Acknowledgments

We acknowledge the use of large language models (DeepSeek and Gemini) for assistance with Python/NumPy implementation, debugging, and simulation scripting for the Monte Carlo framework described in Section 5. The AI tools were utilized solely for code efficiency and data visualization; the authors developed all physical models, analytical derivations, and scientific interpretations independently.

We are grateful to Dr. Lutz Lampe for his insightful discussions on error-control coding and channel modeling. We also thank the TRIUMF facility for their valuable feedback on the practical constraints and feasibility of muon beam generation for communication applications. Additionally, the authors thank the CERN T9 beamline team for providing secondary beam parameters used to constrain the realistic shallow-depth scenario.

References

[1] H. K. M. Tanaka et al., *Muometric positioning and navigation system*, Nature Reviews Physics, vol. 3, pp. 581–588, 2021.

[2] R. L. Workman et al. (Particle Data Group), *Review of Particle Physics*, Prog. Theor. Exp. Phys. 2022, 083C01, 2022.

[3] R. W. Hamming, *Error detecting and error correcting codes*, Bell System Technical Journal, vol. 29, no. 2, pp. 147–160, 1950.

[4] L. Cruz-Ausejo et al., *Occupational accidents in mining workers: Scoping review*, BMJ Open, vol. 14, e078211, 2024.

[5] W. E. King et al., *Caves as terrestrial analogs for mapping and exploration in planetary surface and subsurface environments*, Earth and Space Science, 2024.

[6] Z. Zhang et al., *Muography and its potential applications to mining and rock engineering*, Rock Mechanics and Rock Engineering, vol. 53, pp. 4893–4907, 2020.

[7] V. L. Highland, *Some practical remarks on multiple scattering*, Nucl. Instrum. Methods, vol. 129, pp. 497–499, 1975.

[8] G. R. Lynch and O. I. Dahl, *Approximations to multiple Coulomb scattering*, Nucl. Instrum. Methods Phys. Res. B, vol. 58, pp. 6–10, 1991.

[9] W. W. Peterson and D. T. Brown, *Cyclic codes for error detection*, Proceedings of the IRE, vol. 49, no. 1, pp. 228–235, 1961.

[10] S. Vanini et al., *Muography of different structures using muon scattering and absorption algorithms*, Phil. Trans. R. Soc. A, vol. 377, 2018.

[11] I. F. Akyildiz et al., *Underground wireless communication networks: A survey*, IEEE Communications Magazine, vol. 44, no. 11, pp. 42–47, 2006.

[12] H. K. M. Tanaka et al., *Development of muography for probing the internal structure of volcanoes*, Earth, Planets and Space, vol. 66, p. 126, 2014.

[13] M. Guan et al., *Cosmic-ray muon flux and neutrons at sea level*, arXiv:1509.06176, 2015.

[14] G. Yang et al., *Novel muon imaging techniques*, Phil. Trans. R. Soc. A, vol. 377, 2018.

[15] A. Clarkson et al., *GEANT4 simulation of a scintillating-fibre tracker for cosmic-ray muon tomography*, NIM A, vol. 718, pp. 128–130, 2013.

[16] G. Bonomi et al., *Noise reduction in muon tomography for detecting high density objects*, NIM A, vol. 718, pp. 131–133, 2013.

[17] CERN, *Secondary Beamlines in the East Hall: T9 User Guide*, CERN Accelerator Division, 2024.