Paper review progress

George Scriven   Xenofon Chiotopoulos

Communications Physics resubmission — detailed responses live in the Overleaf answer file.

29 April 2026

Outline

1. Recap of the paper
2. Reviewer overview
3. Major comments and our planned responses
4. Minor comments
5. New benchmarks supporting the response
    - Toy — classical (George): **ready**; quantum (George): **preliminary results ready**
    - Realistic LHCb MC (Xeno): **in progress**
    - Hardware/emulator (Xeno): **ready**
6. Open items & schedule
7. Extras — track-level metrics and tracker A/B

What the paper proposes

1-Bit Quantum Filter for particle-track reconstruction:

- Reformulates HHL: from *matrix inversion* to *spectral filtering*.
- The track-finding Hamiltonian collapses combinatorial-noise eigenstates onto a known eigenvalue λ_c.
- One ancilla qubit + a single phase-rotation step projects them out; signal eigenstates pass through with a $\cos^2(\lambda_j t/2)$ response.

Circuit complexity (N = matrix size = segment-pair count)

- **Direct Structural Synthesis** replaces Trotterisation $\Rightarrow \mathcal{O}(\sqrt{N} \log N)$ gates.
- Demonstrations on emulators of three QPU architectures (Quantinuum H2, IBM-Fez, IBM-Pittsburgh).

Three reviewers — big picture

- **Reviewer 1** — co-review notice only, no scientific content.
- **Reviewer 2** — 3 major + 4 minor comments. Critical:
  - Classical alternatives + $\mathcal{O}(N)$ load complexity
  - Realism / model limitations (claim that the method is "limited to ≤ 4 particles, ≤ 3 layers")
  - Generality of the conceptual advance
- **Reviewer 3** — positive overall:
  - Wants HL-LHC context and a fault-tolerant scaling discussion.
  - Detailed line-by-line edits to Methods + bibliography.

R2 Major 1 — Classical alternatives & $\mathcal{O}(N)$ load

Concern. Classical limits are not quantified. $\mathcal{O}(N)$ data loading may erase the quantum speed-up (Phys. Rev. D **101**, 094015).

Our response.

- No QRAM / external state preparation: $|b\rangle$ is loaded by Hadamards, $\mathbf{A}$ is embedded gate-by-gate by DSS.
- Both costs are already inside the $\mathcal{O}(\sqrt{N}\log N)$ count — no hidden $\mathcal{O}(N)$ load.
- An explicit step-by-step paragraph is being added to the Gate Complexity section. **[done]**

[open: HL-LHC throughput / classical-limit prose; "outperform classical" sentence — not yet drafted.]

Concern. Method only shown on straight tracks, ≤ 3 layers, ≤ 4 particles; will it survive realistic detector effects?

Clarifications already in the manuscript:

- The *toy* model already includes multiple scattering and detector resolution.
- The 4-particle / 3-layer ceiling refers to the *quantum-emulation reach* of the manuscript demos, not the algorithm. Statevector simulations already scale to 1000 particles. **[open: verify exact limit in the paper]**

New benchmarks:

- **Toy** (George): segment-level classical benchmark up to $n_{\rm trk} = 1000$, with hit-drop test — ready. Quantum equivalent — preliminary results up to $n_{\rm trk} = 60$.
- **Realistic MC** (Xeno): 1BQF on the Nicotra et al. events (arXiv:2308.00619), benchmarked vs. Search-by-Triplet — in progress.

Concern. The spectral-filter trick may be too tailored to this Hamiltonian.

Planned response:

- Within tracking: extends naturally to **curved tracks** via the original Hamiltonian (Debney/Pearson).
- More broadly: any problem where the noise eigenstates collapse onto a single known eigenvalue admits the same one-bit filter — useful when filtering a specific eigenvalue is enough and full inversion is not required.

[open: Prose not yet drafted; cross-domain example problems still being identified.]

R2 Minor comments

- **Hardware run on a real device.** Planned via Quantinuum + IBM. **[open: tight in 2 weeks]**
- **Eq. 6** ($|b\rangle$ as $|+\rangle^{\otimes n_s}$). Defined $|b\rangle = H^{\otimes n_s}|0\rangle = |+\rangle^{\otimes n_s}$; clarified the eigenbasis form is needed for the next step. **[done]**
- **Eq. 10 intermediate uncomputation step.** Drafted; merge into manuscript pending.
- **Error mitigation.** None used; will state explicitly and flag as future work.

R3 General comments

- **At what N does the asymptotic advantage matter for HL-LHC?** Will add an HL-LHC event-size reference line on the complexity figure.
- **HL-LHC timeline in the introduction.** Framed as scaling and dependency on fault-tolerant hardware, not as a single-paper deliverable.
- **Fault-tolerant scaling discussion.** **[open: short paragraph TBD]**
- **Sparsity at HL-LHC.** Refer to Nicotra et al.: matrix sparsity grows at most as $\sqrt{N}$ with event size, which keeps the gate count in the regime claimed. **[done]**
- **Why PVs / dense observables matter for non-LHCb readers.** Discussion expanded; PV used as a concrete example of an $\mathcal{O}(1)$-information observable that avoids $\mathcal{O}(N)$ tomography. **[done]**

R3 Line-by-line edits — summary

∼13 textual edits to Methods + a full bibliography rebuild.

Already done:

- θ defined before Eq. 2; "positive semi-definite" remark moved after Eq. 9.
- $\beta_j \to c_j$ to remove the notation clash with the Hamiltonian's β.
- Repetition removals; Aer wording fix; BW-friendly linestyles in Fig. 5.
- Fig. 4 caption clarifies the "#p" point labels; uncertainty source explained.
- Bibliography: DOIs uniform, missing DOIs added, capitalisation fixed.

Still to do:

- n_s definition; reference for "classical exhaustive search methods".
- Expansion of the fake-rate discussion (ties into R2 Major 2).
- Notation consistency for r_1, r_2, r_3.

Toy benchmarks — what we have and what is coming

Goal: address R2 Major 2 ("realism / scaling") by pushing the toy to LHCb-realistic multiplicity, in both the classical and quantum solvers, with realistic detector effects.

Classical toy (George) — ready

- Up to $n_{\mathrm{trk}} = 1000$
- Clean and 1% hit-drop
- Multiple scattering + detector resolution
- Segment-level efficiency, false-rate, spectrum

Quantum toy (George) — preliminary

- Same setup, run with the 1-bit quantum filter
- First direct quantum-vs-classical comparison shown later
- Reach so far: $n_{\mathrm{trk}} \leq 60$, single noise level

Following slides show the classical results, then the preliminary quantum comparison.

Headline — segment efficiency & false rate vs n_{trk}

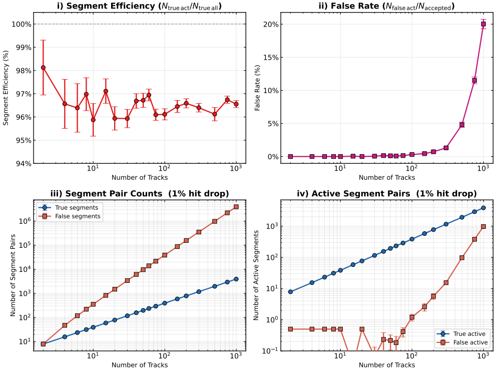

Classical solver, 1 % hit-drop. Segment efficiency stays at 96–98% across $n_{\mathrm{trk}} = 2$–1000. False-positive segment rate is $< 2\%$ up to $n = 300$, climbing to $\sim 20\%$ at $n = 1000$. Metrics use the segments **accepted from the solution vector**, not the geometrically-acceptable set.

Solver-level segment efficiency overlay (clean vs drop=1 %)

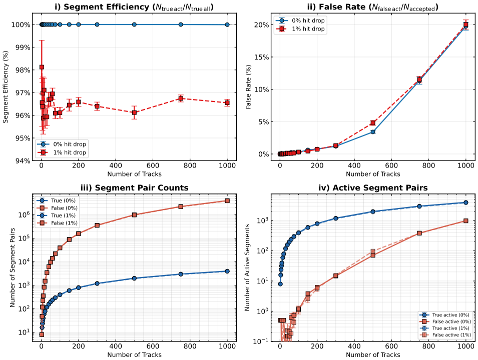

The drop=1% curve tracks the clean curve closely — the algorithm is robust to small hit inefficiencies.

1-bit quantum filter vs classical solver (preliminary)

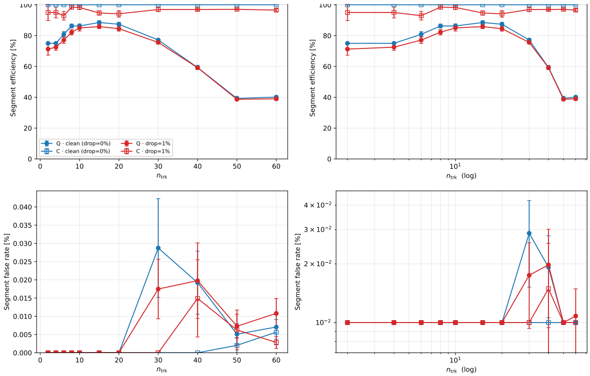

Segment efficiency (top) and false rate (bottom) for the 1-bit quantum filter (**Q**, filled circles) and the classical solver (**C**, open squares) at threshold τ=0.35. Clean (blue) vs 1 % hit-drop (red), $n_{\mathrm{trk}} = 2-60$. **Caveats:** emulator-only; single noise level; Q efficiency dips at very low n and degrades past $n \sim 30$; Q false rate stays $< 0.05\%$ throughout.

R2 Major 2 — realism / scaling:

- Algorithm is **not limited to 4 particles**; the toy goes to 1000 with realistic scattering and resolution.
- Survives a 1% hit-drop without collapse.
- The false-positive tail is small, well-characterised, and threshold-tunable.

R3 — conditioning & sparsity:

- Condition number κ grows only $\sim 3.4\times$ over two decades of n_{trk} *(preliminary — still being tested; correct up to Hamiltonian-diagonal corrections)*.
- The true/false separation gap stays clean at every n — the Hamiltonian itself is healthy.

Direct quantum-vs-classical comparison shown on the previous slide; full noise-level scan and larger-n runs are still open.

Realistic LHCb MC vs Search-by-Triplet — Xeno (in progress)

Plan (R2 Major 2): re-run the Nicotra et al. MC events (arXiv:2308.00619) through the 1-bit filter and compare to a Python implementation of Search-by-Triplet at the segment level.

Draft numbers from the response document (to be confirmed when results land):

- Mean segment-finding efficiency: **1BQF** 94.2 % vs. **SBT** 94.8 %
- Raw segment fake rate: **1BQF** 5.1–13.7 % (rises with n_{hits})
- Post track-building track-level fake rate: 4.3 % [Nicotra et al.]

[open: Final numbers not yet ready — placeholder slide.]

Hardware / emulator benchmarks — Xeno

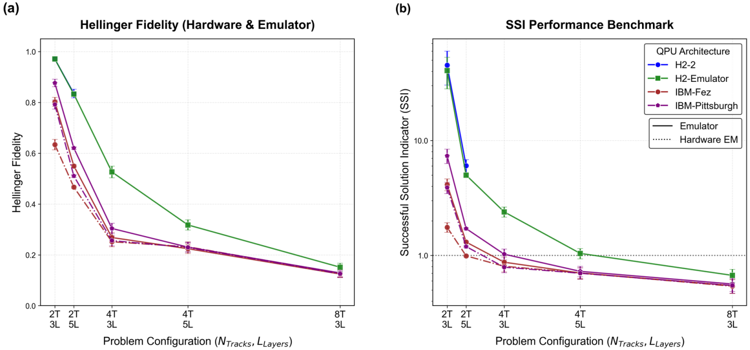

Hellinger fidelity (left) and SSI (right) for 5 problem configurations across H2 (Quantinuum), IBM-Fez, IBM-Pittsburgh. Solid lines = ideal emulator (no noise). Dotted lines = noisy results with error mitigation (*"EM"*): IBM lines are noise-model emulator + EM; the **H2-2 dotted point at 2T-3L is real Quantinuum hardware** — the only physical-device data point. H2 (real + emulator) keeps SSI > 1 except at 8T-3L; the IBM stack falls below SSI = 1 from 4T-3L onward.

Open items — 2 weeks to resubmission

Item	Status	Owner
$\mathcal{O}(N)$ load-complexity paragraph	drafted	Xeno
HL-LHC throughput / classical-limit prose	TODO	TBD
"Outperform classical" sentence	TODO	TBD
Realistic MC + SBT benchmark	running	Xeno
Quantum toy at high n_{trk}	running	George
Generality of the method (R2 Major 3)	TODO	TBD
Hardware run on a real device	uncertain	Xeno
Eq. 10 intermediate step	merge pending	Xeno
Error-mitigation paragraph	TODO	Xeno
Fake-rate expansion in Discussion	TODO	TBD
n_s def, exhaustive-search ref, r_i consistency	TODO	Xeno
Fault-tolerant scaling paragraph (R3)	TODO — no expert	TBD

Schedule & discussion

- Today: 29 April 2026. Resubmission window: $\sim$2 weeks.
- **Critical path:**
  - Realistic MC benchmark (Xeno)
  - Quantum toy at high n_{trk} (George)
  - Classical-limits prose (R2 Major 1)
  - Generality prose (R2 Major 3)
- **Stretch:** fault-tolerant paragraph, error-mitigation prose, hardware on a real device.
- Discussion:
  - Owners for the TBD rows
  - Whether to push back the deadline if MC + hardware are not ready
  - How to frame R2 Major 3 (full worked example or future work?)

Extra — tracker A/B at the track level

Track-level metrics are not what the reviewers asked for, but they are useful context. With naive connected-components grouping, track efficiency collapses 99% → 44% between $n = 300$ and $n = 1000$. With a tracker that adds an angle check and module exclusivity (*same* Hamiltonian, solver, threshold) it stays $\geq 99.96\%$ at every n.

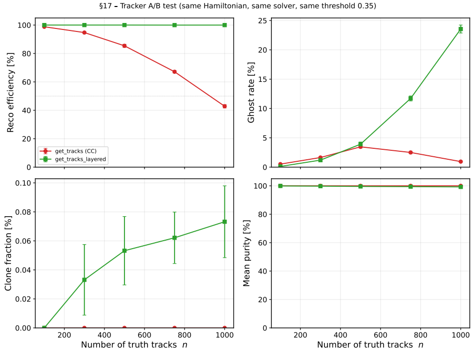

The previous low track-level numbers were a tracker artefact, not a Hamiltonian limit.