

# A 1-bit quantum filter for particle trajectory reconstruction

Received: 29 December 2025

Accepted: 15 July 2026

Published online: 04 August 2026

**Cite this article as:** Chiotopoulos, X., Nicotra, D., Scriven, G. *et al.* A 1-bit quantum filter for particle trajectory reconstruction. *Commun Phys* (2026). <https://doi.org/10.1038/s42005-026-02780-8>

**Xenofon Chiotopoulos, Davide Nicotra, George Scriven, Kurt Driessens, Marcel Merk, Jochen Schütz, Jacco de Vries & Mark H. M. Winands**

We are providing an unedited version of this manuscript to give early access to its findings. Before final publication, the manuscript will undergo further editing. Please note there may be errors present which affect the content, and all legal disclaimers apply.

If this paper is publishing under a Transparent Peer Review model then Peer Review reports will publish with the final article.

# A 1-bit quantum filter for particle trajectory reconstruction

**Xenofon Chiotopoulos<sup>1,2,4,\*</sup>, Davide Nicotra<sup>2</sup>, George Scriven<sup>2,3</sup>, Kurt Driessens<sup>4</sup>, Marcel Merk<sup>1,2</sup>, Jochen Schütz<sup>3</sup>, Jacco de Vries<sup>2</sup>, and Mark H.M. Winands<sup>4</sup>**

<sup>1</sup>Nikhef National Institute for Subatomic Physics, Science Park 105, 1098 XG Amsterdam, The Netherlands

<sup>2</sup>Maastricht University, Faculty of Science and Engineering, Gravitational Waves and Fundamental Physics department, Duboisdomain 30, 6229 GT Maastricht, The Netherlands

<sup>3</sup>Hasselt University, Faculty of Sciences and Data Science Institute, Agoralaan gebouw D, 3590 Diepenbeek, Belgium

<sup>4</sup>Maastricht University, Faculty of Science and Engineering, Department of Advanced Computing Sciences, Paul-Henri Spaaklaan 1, 6229 EN Maastricht, The Netherlands

\*xenofon.chiotopoulos@maastrichtuniversity.nl

## ABSTRACT

The transition to the High-Luminosity Large Hadron Collider (HL-LHC) presents a computational challenge where particle reconstruction complexity may outpace classical computing resources. While quantum computing offers potential speedups, standard algorithms like Harrow-Hassidim-Lloyd (HHL) require prohibitive circuit depths for near-term hardware. Here, we introduce a 1-Bit Quantum Filter, a domain-specific adaptation of HHL that reformulates tracking from matrix inversion to binary ground-state filtering. By replacing high-precision phase estimation with a single-ancilla spectral threshold and exploiting the Hamiltonian's sparsity, we achieve an asymptotic gate complexity of  $\mathcal{O}(\sqrt{N} \log N)$ , given Hamiltonian dimension  $N$ . We validate this approach on LHCb Monte Carlo events, demonstrating segment finding efficiency highly competitive with the classical state-of-the-art methods. Furthermore, we benchmark performance using the Quantinuum System Model H2 trapped-ion processor and IBM Heron R3 superconducting processor. This work establishes a quantum track reconstruction method capable of solving realistic event topologies on noise-free simulators and smaller tracking scenarios within the current constraints of the Noisy Intermediate Scale Quantum (NISQ) era. Remaining challenges toward a full end-to-end tracking solution include an efficient readout and Hamiltonian construction.

## Introduction

Reconstructing particle trajectories from detector hits is a central aspect of experimental high energy physics. The work in this paper anticipates on a hybrid classical-quantum implementation of particle track-finding for the LHCb experiment in LHC Run-5<sup>1</sup>, currently planned to start in 2036. In the High-Luminosity phase of the Large Hadron Collider (HL-LHC) at CERN, high-energy collisions will result in events where thousands of particles are simultaneously produced. These particles traverse multiple sensitive detection layers where they deposit small amounts of energy, generating discrete signals known as detector hits, from which the particle trajectories are reconstructed for physics analysis. The reconstruction process involves a significant combinatorial pattern recognition task, which must be solved in real-time at the high beam crossing rate of 40 MHz and requires expensive computational resources.

The transition to the HL-LHC presents a computational bottleneck that challenges the scalability of traditional track reconstruction algorithms. As the number of simultaneous collisions increases by an order of magnitude, the combinatorial complexity of track reconstruction grows factorially<sup>2</sup>, potentially outpacing the growth of the classical computing budgets. To address this, the high-energy physics community is increasing its exploration of disruptive computing paradigms, including machine learning and quantum computing. For instance, Graph Neural Networks (GNNs) have shown potential but require extensive training and their inference latency goes beyond the strict processing constraints of the real-time LHCb trigger<sup>3,4</sup>. Quantum computers may offer advantages including potential better energy consumption efficiency<sup>5,6</sup> and the ability to leverage high-dimensional Hilbert spaces to address combinatorial problems. Various approaches for particle tracking have been explored, ranging from Quantum Annealing<sup>7,8</sup> to quantum GNNs<sup>9</sup>. In previous work we explored using the Harrow-Hassidim-Lloyd (HHL<sup>10</sup>) algorithm as a global pattern recognition algorithm, TrackHHL<sup>11,12</sup>, achieving good performance but requiring circuit depths that far exceed the coherence times available on current Noisy Intermediate-Scale Quantum (NISQ) devices.

In this paper we present an extension to TrackHHL<sup>12</sup>, which introduces a specialized 1-Bit Quantum Filter to mitigate the

limitations of HHL. Inspired by the Denby-Peterson model<sup>13,14</sup>, track reconstruction is formulated as the ground-state search of an Ising-like Hamiltonian, which we approach as a spectral filtering task rather than a resource-intensive, high-precision matrix inversion. By restricting Quantum Phase Estimation (QPE) to a single-bit clock register and exploiting phase aliasing through a problem-informed evolution parameter, the algorithm deterministically filters out combinatorial noise while constructing the tracking solution. We leverage the Hamiltonian's sparsity and structure by implementing a custom time evolution using the proposed Direct Structural Synthesis (DSS) method, which synthesizes exact unitary operators to bypass Trotterization and reduce circuit depth. This approach achieves an analytical two-qubit gate complexity of  $\mathcal{O}(\sqrt{N} \log N)$ , providing a polynomial speedup over classical matrix inversion and an overall sampling complexity of  $\mathcal{O}(N \log N)$ , where  $N$  is the number of candidate track segments representing the dimension of the linear system. We validate our 1-Bit Quantum Filter using noiseless statevector simulations on both fast MC-toy events and on realistic LHCb Monte Carlo (MC) simulated events containing up to 1,000 hits, demonstrating a segment-finding efficiency of 94.2%, which is comparable with the classical state-of-the-art "Search by Triplet" algorithm<sup>2</sup>. Finally, we benchmark the algorithm's noise resilience through physical executions on IBM's Heron R3 superconducting processors<sup>15</sup> and Quantinuum's System Model H2 trapped-ion system<sup>16</sup>. We demonstrate that for our approach the all-to-all connectivity and lower error rates of the trapped-ion architecture provide improved solution fidelity for small systems with no error correction, while both systems currently suffer from intractable gate error accumulation beyond a problem size of 4 particle tracks and 3 layers. These findings establish a viable pathway to further study quantum particle track reconstruction at the scale of future HL-LHC events.

## Methods

### Hamiltonian Formulation for Trackfinding

The LHCb Vertex Locator (VELO)<sup>17</sup> serves as the innermost tracking detector in the spectrometer, positioned 5 mm from the proton-proton interaction point. It comprises 52 silicon pixel detector modules, divided equally between two retractable halves surrounding the beam line. At the High-Luminosity LHC each beam crossing causes  $\sim 20$  simultaneous collisions, which lead to  $\sim 1000$  produced particles along straight line trajectories due to the absence of a magnetic field. When charged particles traverse the detector modules they deposit energy, allowing the VELO to record precise two-dimensional hit coordinates  $(x, y)$  and derive a  $z$  coordinate from the detector module position. The subsequent reconstruction process aims to group these hits into distinct particle trajectories (tracks). One such simulated beam crossing event in the LHCb VELO is shown in the left panel of Fig. 1. To convert the task of track reconstruction into a format suitable for quantum processing, we formulate the problem within a Hamiltonian framework, adapting the approach initially proposed by Denby and Peterson<sup>13,14</sup> and subsequently refined for the LHCb VELO geometry<sup>11,12</sup>, briefly summarized below.




Figure 1

The system state is parameterized by binary variables  $\mathbf{S} = \{S_i, 1 \leq i \leq N, S_i \in \{0, 1\}\}$ , representing  $N$  potential track segments (doublets connecting hits on adjacent detector layers), where  $S_i = 1$  indicates an active segment contributing to a track, and  $S_i = 0$  indicates inactivity. This graph representation is illustrated in the right panel of Fig. 1. The task of identifying valid tracks is then mapped to finding the ground state of the Hamiltonian  $\mathcal{H}(\mathbf{S})$  with angular, spectral and gap terms, given by:

$$\mathcal{H}(\mathbf{S}) = \mathcal{H}_{\text{ang}}(\mathbf{S}, \epsilon) + \alpha \mathcal{H}_{\text{spec}}(\mathbf{S}) + \beta \mathcal{H}_{\text{gap}}(\mathbf{S}). \quad (1)$$

Here,  $\mathcal{H}_{\text{ang}}$  rewards straight connections between segments  $S_i$  and  $S_j$  where the angle they construct  $\theta_{i,j}$  is smaller than some

angular tolerance  $\varepsilon$ :

$$\mathcal{H}_{\text{ang}}(\mathbf{S}, \varepsilon) = -\frac{1}{2} \sum_{i,j} f(\theta_{i,j}, \varepsilon) S_i S_j, \quad f(\theta, \varepsilon) = \begin{cases} 1 & \text{if } \cos \theta \geq 1 - \varepsilon, \\ 0 & \text{otherwise.} \end{cases} \quad (2)$$

The terms  $\mathcal{H}_{\text{spec}}(\mathbf{S}) = \sum_i S_i^2$  and  $\mathcal{H}_{\text{gap}}(\mathbf{S}) = \sum_i (1 - S_i)^2$  serve as regularization constraints, weighted by parameters  $\alpha = 2.0$  and  $\beta = 1.0$  following the previous work<sup>11</sup>. To leverage quantum linear solvers, this discrete optimization problem is transformed into a linear system, following the mathematical framework established for quantum Hopfield neural networks<sup>18</sup>. By relaxing the discrete segment variables  $S_i$  to continuous variables forming a real-valued vector  $\mathbf{x} \in \mathbb{R}^N$ , the minimum of  $\mathcal{H}(\mathbf{x})$  can be found by setting its gradient to zero:

$$\nabla_{\mathbf{x}} \mathcal{H} = \mathbf{A}\mathbf{x} - \mathbf{b} = 0 \quad \implies \quad \mathbf{A}\mathbf{x} = \mathbf{b} \quad \implies \quad \mathbf{x} = \mathbf{A}^{-1}\mathbf{b}. \quad (3)$$

The resulting matrix  $\mathbf{A}$  encapsulates the quadratic couplings defined in  $\mathcal{H}$ . The vector  $\mathbf{b}$  originates from the linear terms in  $\mathbf{S}$  (arising from  $\mathcal{H}_{\text{gap}}$ ) and takes the form  $\mathbf{b} = (1, 1, \dots, 1)^T$  due to the choice of  $\beta = 1$ . The parameters  $\alpha$  and  $\beta$  are chosen such that the resulting matrix is positive semi-definite. The solution  $\mathbf{x}$  to this linear system represents the relaxed segment amplitudes, which must then be discretized (applying a threshold  $T$  via sampling high-amplitude segments) to identify the final set of track segments. This formulation maps the combinatorial track-finding problem onto a linear algebraic structure compatible with algorithms like HHL.

To quantify the computational scale of this approach, we define the system parameters in terms of the detector geometry and event size. The dimension of the linear system,  $N$ , corresponds to the total number of candidate segments in the set  $\mathbf{S}$ , where  $N \times N$  defines the size of the  $\mathbf{A}$  matrix<sup>11</sup>. For a detector with  $l$  layers and  $m$  reconstructible particle tracks, the system size scales as  $N = m^2(l - 1)$ . The total number of physical interaction terms, denoted by  $k$ , represents the number of angular constraints active in the system, and defines the number of non-zero off-diagonal elements in the Hamiltonian. For full tracks spanning the detector, this scales linearly with the number of particles as  $k = m(l - 2)$ . These parameters will define the gate complexity and sampling cost of our proposed quantum algorithm.

## Simulation Tool

To evaluate the performance of the reconstruction algorithm and in particular to explore the spectral and structural properties of the Hamiltonian formulation in Eq. (1), we developed a dedicated simulation tool. This tool captures the essential geometric and detection features of the LHCb VELO, such as particle multiple scattering, detector resolution as well as hit (in-)efficiency, while remaining computationally tractable for hybrid quantum-classical studies.

The tool implements a given number of sequential measurement planes in  $z$  that register hits of traversing particle tracks produced in events. The events consist of a configurable number of linear particle tracks originating from one or more collision vertices. Given the initial vertex position ( $x = 0, y = 0, z = z_0$ ), each track is parametrized by a state vector  $(x, y, t_x, t_y, p/q)_z$  at position  $z$ , where  $t_x = dx/dz$  and  $t_y = dy/dz$  denote the slopes of the trajectory in the transverse plane, and  $p/q$  is the charge-momentum ratio, included for completeness as in the absence of a magnetic field the particles propagate linearly between the detector layers. At each detector layer  $z = z_i$  the track deposits a *hit*, which is registered with an experimental realistic Gaussian measurement resolution. Subsequently, the track undergoes multiple scattering slightly changing its slopes  $t_x$  and  $t_y$  according to Gaussian multiple scattering  $\sigma_{\text{scatt}}$  of the particle in the material. The average scattering angle implemented is given by the material thickness of the detection planes and is inversely proportional to the momentum of the traversing particle, and therefore allows to study the performance as function of particle track momentum. Fig. 2a illustrates an event as generated from a single collision point.




Figure 2

A classical algorithm using classical matrix inversion techniques is used to test the potential of the Hamiltonian approach on the toy. Its performance is evaluated as a function of the number of collisions with the segment-finding *efficiency*: the fraction of correctly found track segments, and the *fake rate*: the fraction of fake segments in the total sample. Figs. 2b, 2c, 2d, 2e summarize the performance of the method in terms of segment finding efficiency  $\epsilon_{\text{segment}}$  and fake rate  $f_{\text{segment}}$  with:

$$\epsilon_{\text{segment}} = \frac{N_{\text{true accepted}}}{N_{\text{true generated}}} \quad \text{and} \quad f_{\text{segment}} = \frac{N_{\text{false accepted}}}{N_{\text{all accepted}}} \quad (4)$$

With the current Hamiltonian formulation, the efficiency remains close to 100%, independent of problem size. For high track densities the algorithm shows an increasing fake rate. Future work is envisaged to include additional bifurcation and occupancy terms from the Denby-Peterson model, as well as an improved angular term modeling competition between segment candidates.

### Quantum Algorithm

As described in the Methods section, this work builds upon the quantum tracking framework established in Ref.<sup>11,12</sup>, which uses the HHL<sup>10</sup> algorithm to solve linear systems of equations described by the Hamiltonian track finding formulation in Eq. (1). HHL's core advantage comes from the Quantum Phase Estimation (QPE<sup>19,20</sup>) subroutine to resolve these eigenvalues into a multi-qubit clock register. QPE consists of a Hamiltonian evolution followed by an inverse Quantum Fourier Transform (QFT<sup>21</sup>). The QPE's required precision determines the number of applications of the controlled unitary that implements the time evolution. A controlled rotation is then applied to an ancilla qubit to encode the inverse eigenvalues ( $1/\lambda$ ) onto the state amplitudes, followed by the uncomputation of the clock register. HHL is theoretically powerful and promising a logarithmic speedup, though its reliance on a high-precision estimation of the spectrum coupled with an expensive tomography creates significant bottlenecks<sup>22,23</sup>. The algorithm requires prohibitive circuit depths due to iterative application of the Hamiltonian evolution  $e^{-iAt}$  and QPE requires many clock register qubits, rendering it unfeasible for particle tracking on near-term hardware.


To address these limitations we present a 1-Bit Quantum Filter, a formalization and extension of our previous work<sup>12</sup>. We observe that for track reconstruction, exact matrix inversion is not necessary. The tracking problem is inherently binary, the task is to distinguish valid 'signal' segments from 'noise' segments. Consequently, we replace the resource-intensive high-precision QPE and controlled rotation using a 1-bit QPE with a deterministic eigenvalue filter. This reformulates the task from exact inversion to ground-state spectral filtering, drastically reducing the gate complexity and sampling requirements<sup>12</sup>. Additionally, we replace the general purpose approximate Suzuki-Trotter evolution<sup>24</sup> with a Direct Structural Synthesis (DSS) method, which synthesizes controlled unitary operators without the need for an expensive classical decomposition.

#### A 1-Bit Quantum Filter

This approach replaces resource-intensive eigenvalue inversion with a binary filtration mechanism that takes advantage of the specific spectral properties of the track-finding Hamiltonian. The algorithm begins with determining the evolution time

## Circuit Architecture and Decomposition

## (a) 1-Bit Quantum Filter Circuit



## (b) Direct Structural Synthesis (DSS) Decomposition



(c) Generalized Two-Level Unitary Structure ( $G_k$ )

Figure 3

$t$ , which is used in our 1-bit phase estimation subroutine. In general, the spectrum bounds  $(\lambda_{\min}, \lambda_{\max})$  can be estimated via the Gershgorin Circle Theorem<sup>25</sup>. However, for the particle tracking case the bounds are known as they converge to  $(\lambda_{\min} \approx 1, \lambda_{\max} \approx 5)$ <sup>11</sup>. Since the physical length of a track is bounded by the number of detector layers ( $l$ ), the eigenvalues of the track sub-matrices are bounded by the track length rather than the total number of candidate segments  $N$ . Due to this convergence, we can target the central eigenvalue  $\lambda_c$ . All the noise states exist in the eigenspace associated with the central eigenvalue, the numerical value of this is determined by  $\alpha + \beta$  from Eq. (1):

$$\lambda_c = \alpha + \beta \quad \text{and} \quad t = \frac{\pi}{\lambda_c}. \quad (5)$$

This choice of  $t$  is a deliberate deviation from standard QPE, where the evolution time is typically scaled as a function of  $\lambda_{\max}$ , to prevent phase aliasing<sup>10,26</sup>. We instead reduce the precision of the QPE to a single time register qubit and take advantage of the aliasing due to our choice of  $t$  in Eq. (5), which allows us to project the noise to a specific phase. We prepare three registers: the system register  $|0\rangle_S$  for encoding the  $\mathbf{b}$  vector, the time register used for the QPE  $|0\rangle_T$  and the ancilla register used to determine if the algorithm has run successfully  $|0\rangle_A$ . The system register (comprising  $n_s = \log_2 N$  qubits) is initialized in the uniform superposition state  $|b\rangle = H^{\otimes n_s} |0\rangle = |+\rangle^{\otimes n_s}$ , as  $\mathbf{b}$  is defined as a uniform vector. To analyze the spectral filtering mechanism, we rewrite the state in the eigenbasis of the Hamiltonian matrix  $\mathbf{A}$ :

$$|b\rangle = \sum_j \gamma_j |u_j\rangle, \quad (6)$$

where  $|u_j\rangle$  are the orthonormal eigenvectors with eigenvalues  $\lambda_j$ , and  $\gamma_j = \langle u_j | b \rangle$  represents the overlap amplitude. The proposed 1-Bit Quantum Filter proceeds in five steps, tracking the evolution of the state  $|\Psi\rangle$  across the ancilla  $|A\rangle$ , time  $|T\rangle$ , and system  $|S\rangle$  registers. The following steps match exactly the circuit shown in Fig. 3a.

**Step 1: State Preparation** We initialize the system register  $|0\rangle_S$  into a uniform superposition  $|b\rangle = |+\rangle^{\otimes n_s}$ , representing the vector  $\mathbf{b} = (1, 1, \dots, 1)^T$ . To track the phase evolution in subsequent steps, we express the state  $|+\rangle^{\otimes n_s}$  in the eigenbasis of the Hamiltonian matrix  $\mathbf{A}$  as in Eq. 6. The initial state is:

$$|\Psi_1\rangle = |0\rangle_A \otimes \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)_T \otimes \sum_j \gamma_j |u_j\rangle_S. \quad (7)$$

The time register is also prepared in superposition via a Hadamard gate.

**Step 2: 1-Bit Phase Estimation** We apply the controlled time evolution  $U = e^{iAt}$  conditioned on the  $|1\rangle_T$  state. This unitary is implemented by combining the Direct Structural Synthesis (DSS) method applied to non-diagonal terms, with a Phase Gate  $P(\phi)$  (where  $\phi = -(\alpha + \beta)t$ ) that accounts for the diagonal contributions of the Hamiltonian. Together, these components realize the time evolution of  $A$ , resulting in the phase  $e^{i\lambda_j t}$  being kicked back onto the  $|1\rangle_T$  state for each eigen-component:

$$|\Psi_{2a}\rangle_j = |0\rangle_A \otimes \frac{1}{\sqrt{2}} \sum_j \gamma_j \left( |0\rangle_T + e^{i\lambda_j t} |1\rangle_T \right) \otimes |u_j\rangle_S. \quad (8)$$

This creates a relative phase between the  $|0\rangle_T$  and  $|1\rangle_T$  states. Next, the inverse QFT is applied to the time register  $|0\rangle_T$ . In the 1-Bit case the inverse QFT is implemented using a single Hadamard on the time register to interfere the phases. Using the identity  $H(|0\rangle + e^{i\phi}|1\rangle) \propto \cos(\phi/2)|0\rangle - i\sin(\phi/2)|1\rangle$ <sup>26</sup>, it follows that:

$$|\Psi_{2b}\rangle_j = |0\rangle_A \otimes \sum_j \gamma_j \left[ \cos\left(\frac{\lambda_j t}{2}\right) |0\rangle_T - i\sin\left(\frac{\lambda_j t}{2}\right) |1\rangle_T \right] \otimes |u_j\rangle_S. \quad (9)$$

The phase information  $\lambda_j t$  has been converted into amplitude information in the  $|0\rangle_T$  and  $|1\rangle_T$  basis states.

**Step 3: Conditional Logic** We apply a Zero-Controlled NOT gate on the ancilla, flipping  $|0\rangle_A \rightarrow |1\rangle_A$  only when the time register is  $|0\rangle_T$ . The relative phase relationship is preserved, but now entangled with the ancilla register:

$$|\Psi_3\rangle = \sum_j \gamma_j \left[ \underbrace{\cos\left(\frac{\lambda_j t}{2}\right) |1\rangle_A |0\rangle_T}_{\text{Signal}} - \underbrace{i\sin\left(\frac{\lambda_j t}{2}\right) |0\rangle_A |1\rangle_T}_{\text{Noise}} \right] \otimes |u_j\rangle_S. \quad (10)$$

**Step 4: Uncomputation** To disentangle the time register, we apply the inverse of the phase estimation circuit (Inverse QPE). This sequence consists of applying a Hadamard gate, the inverse controlled time evolution  $C(e^{-iAt})$ , and a final Hadamard gate to the time register. Applying this sequence to  $|\Psi_3\rangle$  yields the exact full intermediate state:

$$|\Psi_{4a}\rangle = \sum_j \gamma_j \left[ \left( \cos^2 \frac{\lambda_j t}{2} |1\rangle_A + \sin^2 \frac{\lambda_j t}{2} |0\rangle_A \right) |0\rangle_T + \frac{i}{2} \sin(\lambda_j t) (|1\rangle_A - |0\rangle_A) |1\rangle_T \right] \otimes |u_j\rangle_S. \quad (11)$$

This explicit expansion demonstrates how the time register is restored. For the combinatorial noise subspace ( $j \in \text{Noise}$ ), our specific choice of evolution time  $t = \pi/\lambda_c$  ensures the time register perfectly uncomputes to  $|0\rangle_T$  and is entirely isolated into the rejected  $|0\rangle_A$  state. For the signal subspace ( $\lambda_j \neq \lambda_c$ ), the uncomputation splits across the time register with minor phase leakage. The idealized target state projected onto the successfully uncomputed  $|0\rangle_T$  subspace is grouped as:

$$|\Psi_4\rangle = \left[ \left( |1\rangle_A \otimes \sum_{j \in \text{Signal}} \gamma_j |u_j\rangle_S \right) + \left( |0\rangle_A \otimes \sum_{j \in \text{Noise}} \gamma_j |u_j\rangle_S \right) \right] \otimes |0\rangle_T. \quad (12)$$

**Step 5: Measurement & Filtration** Finally, we measure the ancilla. Our specific choice of  $t = \pi/\lambda_c$  ensures that for noise eigenvalues ( $\lambda = \lambda_c$ ), the probability of triggering the signal flag vanishes:

$$P(1|\lambda_c) = \cos^4 \left( \frac{\lambda_c t}{2} \right) + \frac{1}{4} \sin^2(\lambda_c t) = \cos^2 \left( \frac{\lambda_c t}{2} \right) = \cos^2 \left( \frac{\pi}{2} \right) = 0. \quad (13)$$

As a result, post-selecting on the measurement projecting onto the signal subspace  $|1\rangle_A$  deterministically filters out the combinatorial noise. The collapsed state must then be renormalized by the probability of success  $P_{\text{succ}}$ , yielding the solution state:

$$|\Psi_5\rangle = \frac{1}{\sqrt{P_{\text{succ}}}} \sum_{j \in \text{Signal}} \gamma_j \cos^2 \left( \frac{\lambda_j t}{2} \right) |u_j\rangle_S. \quad (14)$$

This spectral separation is guaranteed by the Hamiltonian structure defined in Eq. (1). Since noise segments correspond to isolated nodes in the graph, their eigenvalues are fixed exactly at  $\lambda_c$ , while valid track segments are spectrally shifted away. To exploit this, we replace the standard HHL rotation with a Zero-Controlled NOT gate (Step 3, Eq. 10). This gate targets the ancilla conditioned on the  $|0\rangle_T$  state, flipping it to  $|1\rangle_A$  only for signal components. Consequently, noise components (which map to  $|1\rangle_T$ ) fail to trigger the gate and remain in  $|0\rangle_A$ . Following the uncomputation step (Step 4, Eq. 12), we measure the ancilla and post-select on the  $|1\rangle_A$  outcome. This deterministically filters the noise contribution ( $P(1|\lambda_c) = 0$ ), leaving a final state composed strictly of valid track segments.

#### Time Evolution through Direct Structural Synthesis

A central component of our 1-bit Quantum Filter is the implementation of the controlled time evolution  $C(U) = C(e^{-iAt})$ . A general-purpose Suzuki-Trotter time evolution of  $A$  is resource intensive<sup>12</sup>, as it requires the  $A$  matrix to be represented as a sum of Pauli strings. An alternative approach is the Direct Structural Synthesis (DSS) method that exploits the specific structure of our Hamiltonian, placing it within the class of exact gate synthesis algorithms<sup>27</sup>. The DSS method is based on the fact that we can express  $A = cI - B$ . Since the diagonal matrix term  $cI$  commutes with the off-diagonal interaction matrix  $B$ , the unitary evolution operator  $U(t) = e^{-iAt}$  allows for an exact factorization:

$$U(t) = e^{-i(cI-B)t} = e^{-icIt} \cdot e^{iBt}. \quad (15)$$

Implementing the controlled version of this evolution,  $C(U)$ , separates into a controlled global phase rotation ( $C(e^{-icIt})$ ) and the controlled interaction terms ( $C(e^{iBt})$ ).

The first term  $C(e^{-icIt})$  can be efficiently implemented with a single  $P(\phi)$  gate on the control qubit, with  $\phi = -ct$ . The phase rotation can be implemented with a single gate as  $cI$  is a uniform identity matrix, its phase is kicked back deterministically from the target qubit to the control<sup>26</sup>. The matrix  $B = \sum_k B_k$  is a sum of its individual interaction components, where the index  $k$  represents a unique interaction term  $\{i, j\}$ , and the total number of interaction terms that pass the  $\mathcal{H}_{\text{ang}}(\mathbf{S}, \epsilon)$  angular function

can be written as  $k = m(l - 2)$  for  $m$  tracks. Each  $B_k$  represents a  $\sigma_x$ -like interaction between two basis states. The algorithm implements  $C(e^{iBt})$  by iterating through  $k$  components and applying the gate for each one sequentially:

$$G_{\text{DSS}} = \prod_k G_k \quad \text{where} \quad G_k = C(e^{iB_k t}). \quad (16)$$

A circuit based illustration of Eq. (16) is visualized in Fig. 3b. To strictly enforce the graph topology of  $B$ , each gate  $G_k$  is implemented as an exact Two-Level Unitary, often referred to as a Givens rotation<sup>26</sup>. For an interaction between basis states  $|i\rangle$  and  $|j\rangle$ , the operator  $G_k$  performs a rotation strictly within the subspace spanned by  $\{|i\rangle, |j\rangle\}$  and acts as the identity on all other states. This is synthesized using a ladder of CNOT gates to disentangle the basis states<sup>28</sup>, followed by a multi-controlled rotation  $C^{n-1}(R_X(\theta))$ , and a subsequent uncomputation step.

The exactness of DSS time evolution hinges on whether the individual interaction terms  $B_k$  commute. While this holds for terms acting on disjoint subspaces, the general case for a  $d$ -sparse Hamiltonian involves interaction terms sharing basis states. For these non-commuting terms, the factorization  $e^{iBt} \neq \prod_k e^{iB_k t}$  is not an equality, but rather a first-order Suzuki-Trotter product formula<sup>24</sup>:

$$e^{iBt} = \prod_k e^{iB_k t} + \mathcal{O}(t^2). \quad (17)$$

While this approximation introduces an error  $\mathcal{O}(t^2)$  scaling quadratically with time, it is crucial to note that this error is subspace dependent. Specifically, the primary function of our filter remains robust because the noise eigenvectors  $|\psi_c\rangle$  lie in the null space of the interaction matrix  $B$ . Since noise segments correspond to isolated vertices in the graph, they have no off-diagonal couplings, meaning  $B|\psi_c\rangle = 0$  exactly. Consequently, the time evolution is exact for the noise subspace, ensuring that combinatorial background suppression is unaffected by time-evolution errors. The Trotter error  $\mathcal{O}(t^2)$  acts non-trivially only on the signal eigenspace, introducing small phase perturbations that may slightly affect signal acceptance but not noise rejection.

While the noise subspace is immune to this error, the Trotter error acts non-trivially on the solution eigenspace ( $\lambda_j$ ). The first-order simulation  $U_{\text{sim}}(t)$  returns a perturbed phase  $\phi_{\text{sim}} = \phi_{\text{exact}} + \delta\phi$ , where  $\phi_{\text{exact}} = \lambda_j t$  and  $\delta\phi$  represents the  $\mathcal{O}(t^2)$  phase error. Consequently, the measured probability for the solution state is shifted:

$$P(0|\lambda_j) = \cos^2\left(\frac{(\lambda_j t + \delta\phi)}{2}\right) \quad \text{and} \quad P(1|\lambda_j) = \sin^2\left(\frac{(\lambda_j t + \delta\phi)}{2}\right). \quad (18)$$

The phase shift  $\delta\phi$  perturbs the interference condition. While  $\delta\phi$  does not cause spectral leakage for the critical  $\lambda_c$  state, a large enough  $\delta\phi$  can blur the signal phases  $\lambda_j$ . This can cause a portion of the solution states to be projected onto the reject bin ( $|1\rangle_T$ ) by mistake, or conversely, shift the acceptance probability  $P(0)$  closer to 1 or 0 depending on the sign of the error term. For the precision required in this application, the first-order approximation is sufficient; however, if higher precision QPE is critical, a higher-order product formula (e.g., 2nd-order Suzuki-Trotter with  $\mathcal{O}(t^3)$  scaling) could be employed to suppress the error scaling.

## Complexity Analysis

To quantify the impact of the changes to HHL we study the computational complexity of the 1-Bit Quantum Filter and derive the asymptotic complexity for both two-qubit gate count and sampling cost. We demonstrate that the total gate and sampling complexity scales as  $\mathcal{O}(\sqrt{N} \log N)$  and  $\mathcal{O}(N \log N)$ , respectively.

### Gate Complexity

We derive the asymptotic gate complexity for a single execution of the circuit. Based on the system parameters defined in Hamiltonian construction, the number of interaction terms scales as  $k \approx \sqrt{N}$ . Consequently, the total complexity scales as  $\mathcal{O}(\sqrt{N} \log N)$ , dominated by the sequential implementation of these off-diagonal terms in the time evolution. The total gate complexity  $C_{\text{gate}}$  is the sum of the complexity in steps 1-4 in Fig. 3a:

1. **Step 1, State Preparation:** This step loads the uniform vector  $|\mathbf{b}\rangle$  into the quantum circuit. To prepare this perfect superposition, we apply a single Hadamard gate to all  $n_s$  system qubits. While this can be executed in  $\mathcal{O}(1)$  depth, the total gate count is  $\mathcal{O}(n_s) = \mathcal{O}(\log N)$ .
2. **Step 2, 1-Bit Quantum Phase Estimation:** The QPE is implemented with the DSS method and a single phase rotation gate. The diagonal and off-diagonal components of  $\mathbf{A}$  are dynamically loaded into the quantum circuit using the single phase rotation and DSS methods respectively. The DSS method iterates  $k$  times and each iteration synthesizes an exact two-level unitary  $G_k$  using three sub-routines:

- (a) **Basis Transformation:** A ladder of CNOT gates is applied to map the interacting states  $|i\rangle$  and  $|j\rangle$  such that they differ at only a single qubit index (the pivot), seen in Fig 3c. The number of two-qubit gates scales with the Hamming distance  $d_H$  as  $\mathcal{O}(d_H)$ . Since  $d_H \leq n_s$ , this is upper-bounded by  $\mathcal{O}(n_s)$ .
- (b) **Control Logic:** To trigger the multi-controlled rotation,  $X$  gates are applied to any  $n_s - 1$  non-pivot qubits that are currently in the  $|0\rangle$  state. This ensures that the multi-controlled rotation will target the subspace defined by  $\{|i\rangle, |j\rangle\}$  correctly. This step contributes  $\mathcal{O}(n_s)$  single-qubit gates.
- (c) **Rotation:** A multi-controlled  $R_X(\theta)$  rotation is applied to the pivot qubit, controlled by the  $n_s - 1$  other qubits. Using the decomposition for multi-controlled rotations without ancilla<sup>29</sup>, this contributes  $\mathcal{O}(n_s)$  gates.
- (d) **Uncomputation:** The  $X$  gates and CNOT ladder are applied in reverse.

Summing these components, a single DSS iteration scales linearly with the register size  $\mathcal{O}(n_s)$ . For  $k$  iterations, the total complexity for Step 2 is  $\mathcal{O}(k \cdot n_s)$ .

- 3. **Step 3, Controlled Inversion:** The controlled inversion consists of two  $X$  gates and a single CNOT gate. This is constant  $\mathcal{O}(1)$  complexity.
- 4. **Step 4, Uncomputation:** This is identical to Step 2, thus doubling the complexity contribution of the phase estimation.

The complexity is primarily determined by the  $k$  repetitions required to encode every valid interaction term. Substituting  $n_s = \log N$  and the sparsity factor  $k \approx \sqrt{N}$ , the total gate complexity is:

$$C_{total} \approx k \cdot n_s \approx \mathcal{O}(\sqrt{N} \log N). \quad (19)$$

This scaling represents a polynomial speedup in  $N$  compared to classical exhaustive search methods<sup>30</sup>, providing a computational advantage analogous to Grover's algorithm<sup>31</sup>.

### Sampling Complexity

Unlike the standard HHL algorithm, where the probability of measuring the solution state is suppressed by the condition number  $\kappa$  as  $\mathcal{O}(1/\kappa^2)$ <sup>10</sup>, the performance of this 1-Bit Quantum Filter is governed by the geometric overlap of the initial state with the solution subspace. The algorithm succeeds when the ancilla is measured in the  $|1\rangle_A$  state, which requires the time register to collapse to  $|0\rangle_T$ . The total success probability is the sum of the conditional probabilities for each eigenvector  $|u_j\rangle$  weighted by its overlap with the input vector  $|\mathbf{b}\rangle$ :

$$P_{succ} = P(T = 0) = \sum_j P(0|\lambda_j) \cdot |\gamma_j|^2. \quad (20)$$

The Hamiltonian spectrum is partitioned into two orthogonal subspaces: the Signal subspace (valid tracks,  $\lambda \neq \lambda_c$ ) and the Noise subspace (inactive segments,  $\lambda = \lambda_c$ ). Substituting the filter response from Eq. (13), the success probability becomes:

$$P_{succ} \approx \underbrace{\frac{1}{2} \cos^2\left(\frac{\lambda_j \pi}{2\lambda_c}\right)}_{\text{Signal Pass Rate}} \sum_{j \in \text{Signal}} |\gamma_j|^2 + \underbrace{\frac{1}{2} \sin^2\left(\frac{\lambda_c \pi}{2\lambda_c}\right)}_{\text{Noise Nulling}} \sum_{c \in \text{Noise}} |\gamma_c|^2. \quad (21)$$

The filter deterministically eliminates the noise subspace (as the second term in Eq. (21) vanishes), while passing the signal component with an effective probability  $c_{filter} = \cos^2\left(\frac{\lambda_j \pi}{2\lambda_c}\right)$ . The summation term  $\sum_{j \in \text{Signal}} |\gamma_j|^2$  represents the overlap of the uniform superposition of  $|b\rangle$  with the valid track configurations. For a system with  $N$  total segments and  $k$  valid interaction terms, this overlap is the ratio of the subspace dimensions. Thus, the single-shot success probability scales linearly with the density of valid tracks, given  $M = m(l-1)$  valid track segments:

$$P_{succ} \approx c_{filter} \cdot \sum_{j \in \text{Signal}} |\gamma_j|^2 \approx \frac{M \cdot c_{filter}}{N}. \quad (22)$$

To derive the full sampling complexity we account for the reconstruction of the entire event. Since the output is probabilistic, this process can be modeled by the Coupon Collector's Problem<sup>32</sup>, where we sample with replacement from  $M$  distinct outcomes. To observe all distinct segments with high probability, the required number of successful measurements scales as

$S_{\text{dist}} \approx M \log M$ <sup>32</sup>. The total computational cost, defined as the total number of circuit executions  $S_{\text{total}}$ , is the required number of successes normalized by the single-shot probability  $P_{\text{succ}}$ :

$$S_{\text{total}} = \frac{S_{\text{dist}}}{P_{\text{succ}}} \approx \frac{M \log M}{c_{\text{filter}}(M/N)} \approx N \log M \approx \mathcal{O}(N \log \sqrt{N}) \approx \mathcal{O}(N \log N). \quad (23)$$

By expressing the track density in terms of the total system size  $N$ , we observe that the complexity scales as  $\mathcal{O}(N \log N)$ . This represents an advantage over HHL, as we eliminate the polynomial dependence on the condition number  $\kappa$ <sup>10</sup> and the precision overhead of the QPE. However, this does not overcome the fundamental challenge of quantum state tomography. To address this we would need to project onto more global event observables such as the particle's origin point (the "vertex") or employ some form of amplitude amplification on the  $|1\rangle_A$  subspace.

## Results

### Performance Evaluation

To strictly evaluate the efficacy of the proposed 1-Bit Quantum Filter, we applied the algorithm to simulated event data generated with the simulation tool. This study centers on two critical performance metrics: (1) the ability to find the correct segments under realistic quantum noise conditions, with low fake rates, and (2) the empirical validation of the theoretically derived sampling and two-qubit gate complexities.

For these benchmarks, we employed a controlled version of the simulation tool. To study the scaling of the quantum algorithm we restricted the toy model to a single primary vertex with no detector inefficiencies or multiple scattering. We then varied the detector geometry between 3 and 5 layers to study the behavior and scaling of the  $c_{\text{filter}}$  coefficient.

To run our 1-Bit Quantum Filter, we utilized Quantinuum's TKET compilation stack<sup>33</sup> and the Quantinuum Nexus platform<sup>34</sup> to perform noiseless simulations, noise-model emulations, and hardware executions targeting the Quantinuum System Model H2 Quantum Processing Unit (QPU), which utilizes high-connectivity trapped-ion hardware. These results are compared against IBM's Qiskit framework<sup>35</sup>, utilizing both the IBM Pittsburgh device and its associated noise model. This system represents the Heron R3 superconducting architecture<sup>15</sup>.




Figure 4

### Validation of Derived Complexities

To verify the theoretical two-qubit gate complexity of  $\mathcal{O}(\sqrt{N} \log N)$  (Eq. 19), we analyzed the post-compilation two-qubit gate counts for circuits targeting both IBM and Quantinuum architectures. Using the hardware compilation pipelines established

above, we generated the necessary circuits using noiseless statevector simulations with Qiskit's Aer simulator<sup>35</sup>, accelerated via an Nvidia L40S GPU. For the compilation step, we specifically applied Qiskit's optimization level 3 for the IBM Pittsburgh backend, and 'full peephole' optimization<sup>33</sup> within the TKET framework. The gate count analysis is restricted to the range  $N \in [2^3, 2^{17}]$ , as the classical optimization of circuits exceeding  $10^6$  two-qubit gates becomes computationally intractable beyond this point.

The compiled gate counts are fitted to the theoretical ansatz  $C_{total} = aN^b \log(N) + c$ . The logarithmic term is fixed to reflect the dependence on register size ( $n_s = \log N$ ), allowing us to isolate the polynomial scaling exponent  $b$ . This constraint minimizes the degrees of freedom relative to the dataset size preventing overfitting, in the asymptotic regime.

Figures 4a and 4b display the results with  $1\sigma$  confidence bands. We observe excellent agreement between the fitted coefficient  $b$  and the theoretical prediction of 0.5 in the asymptotic regime ( $N > 2^{10}$ ), where the confidence bands narrow significantly. The divergence observed at smaller system sizes is attributed to the constant-depth overhead not captured by the asymptotic model.

Comparing the two compilation pipelines reveals a distinct scaling behavior driven by hardware topology. The TKET results (Fig. 4b), which target the all-to-all connectivity of trapped-ion systems, successfully minimize the overall gate count as the hardware supports direct interaction between arbitrary qubit pairs removing the necessity for SWAP gates. This topological advantage is explicitly reflected in the fitted parameters, yielding tight confidence intervals, lower scaling pre-factors (e.g.,  $a = 23.27 \pm 0.51$  for 3 layers), and minimal constant overhead ( $d$ ). In contrast, Qiskit results (Fig. 4a) exhibit larger overhead and higher variance. This discrepancy arises from the limited connectivity of the IBM superconducting architecture, which requires the insertion of SWAP chains during the routing process to implement non-local Hamiltonian terms. This complex routing inflates both the scaling coefficient  $a$  rising to  $a = 37.60 \pm 6.27$  for 3 layers and  $a = 127.78 \pm 19.68$  for 5 layers.

Having established the theoretical total sampling complexity in Eq. (22), we see that the validation of our model is entirely dependent on the behavior of the success probability  $P_{succ}$  with system size  $N$ , as  $S_{dist}$  follows the standard coupon collector scaling. Rewriting Eq. (21) in terms of the system size  $N$ , the theoretical success probability is expected to scale as  $P_{succ} \approx (c_{filter} \sqrt{l-1})/\sqrt{N}$ . As such, we fit the simulation data to a power-law model of the form  $P_{succ} = a \cdot N^b$ . Validation of the theoretical derivation requires that the fitted exponent  $b$  converges to  $-0.5$  and the coefficient  $a$  aligns with the derived factor  $c_{filter} \sqrt{l-1}$ .

The dataset for this scaling analysis was generated using statevector simulations. Since determining  $P_{succ}$  does not require the computationally expensive hardware specific compilation step used in the gate analysis, we extended the range from the minimal configuration of 3 layers and 2 tracks ( $N = 2^3$ ) up to a denser environment of 5 layers and 1024 tracks ( $N = 2^{22}$ ). Experimentally,  $P_{succ}$  is quantified as the probability of the ancilla flag qubit projecting onto the  $|1\rangle$  subspace. Figure 4c presents the power-law fits for 3-layer and 5-layer detector configurations. The fitted coefficients  $a$  and exponents  $b$  exhibit excellent agreement with theoretical predictions for both configurations ( $R^2 > 0.999$ ). Specifically, the fitted exponents match the expected  $N^{-0.5}$  scaling ( $b = -0.4998 \pm 0.0001$  for 3 layers, and  $b = -0.4999 \pm 0.0001$  for 5 layers). Furthermore, the empirical pre-factors yield  $c_{filter}$  values of  $0.501 \pm 0.001$  and  $0.813 \pm 0.001$  for 3 and 5 layers respectively, tightly aligning with our theoretical  $c_{filter}$  derivations of 0.500 and 0.823. This confirms the validity of the derived scaling laws.

A deviation is observed for large system sizes ( $N > 2^{18}$ ), where the measured  $P_{succ}$  exceeds the theoretical prediction. This enhancement is not a computational artifact but a consequence of the combinatorial density. In regions of high track occupancy, the density of hits increases the probability of fake segments satisfying the angular constraints of the Hamiltonian Eq. (2). Because the current implementation relies on a simple binary threshold defined in Eq. (2), random geometric alignments can satisfy the angular constraints. These false positives satisfy the Hamiltonian constraints and contribute to the final solution state, increasing the  $P_{succ}$  even though they do not correspond to any real tracks. This was also observed using classical matrix inversion in the simulation tool, and could be resolved by adding additional Denby-Peterson terms to the Hamiltonian.

## Benchmarking Realistic Event Topologies

While the algorithm demonstrates theoretical scaling advantages, evaluating its robustness against realistic detector conditions is critical to assessing its practical relevance. To this end, we utilized quantum statevector simulations to process realistic LHCb Monte Carlo (MC) events. Specifically, we utilized the same  $B_s \rightarrow \phi\phi$  selection that has been previously used to obtain tracking efficiencies using classical inversion<sup>11</sup>. The  $B_s \rightarrow \phi\phi$  events were generated from  $pp$  collisions at a center-of-mass energy  $\sqrt{s} = 13.6$  TeV, corresponding to LHC run 3 and 4. The event generation and detector response were simulated using the standard combination of Pythia<sup>36</sup> with a specific LHCb configuration<sup>37</sup>, EvtGen<sup>38</sup> for hadronic decays, and Geant4<sup>39</sup> to simulate the material interactions<sup>40</sup>. We select events with up to 1000 hits in the right half of the VELO. The event is then processed using a quantum simulation of the 1-Bit Quantum Filter algorithm, which allows us to probe the algorithm's response in tracking environments with realistic detector inefficiencies.

We benchmarked the performance of our 1-Bit Quantum Filter simulation directly against the algorithm that is currently used in LHCb reconstruction, the Search by Triplet (SBT) algorithm<sup>2</sup>, utilizing its open-source Python implementation<sup>41</sup>. To




Figure 5

ensure a direct and fair comparison with the distribution sampled from the quantum simulation, performance was evaluated strictly at the segment level, using the segment finding efficiency  $\epsilon_{\text{segments}}$  and the fake rate  $f_{\text{segments}}$  as defined in Eq. 4.

As seen in Fig. 5a, the 1-Bit Quantum Filter remains effective in realistic events with up to 1,000 hits. The quantum algorithm achieved a mean  $\langle \epsilon_{\text{segments}} \rangle = 94.2\% \pm 2.2\%$  across the tested event sizes, which is close to the SBT's efficiency of  $\langle \epsilon_{\text{segments}} \rangle = 94.8\% \pm 1.7\%$ . However, examining the segment fake rate in Fig. 5b reveals a different trend. The SBT algorithm's  $f_{\text{segments}}$  exhibits a minimal increase with event size, ranging from 2.9% – 3.4%, whereas the quantum algorithm yields a range of 5.1% – 13.7%.

Fundamentally, these false positive segments originate from the construction of our Hamiltonian, which currently relies solely on an angular step function. As discussed previously, suppressing these with Denby-Peterson penalty terms<sup>42,43</sup> is possible, but remains a subject for further study due to the necessary trade-off with increased quantum circuit depth.

### Solution Fidelity on noisy hardware




Figure 6

The noise resilience of the proposed 1-bit quantum filter is assessed by executing the circuits on both physical quantum hardware and noise-model emulators for the IBM Pittsburgh (Heron R3) and Quantinuum H2-2 processor. The noise models

are derived from regular device calibration data, accurately reflecting the coherence times, gate fidelities, and readout error rates of the physical hardware at the time of execution. We employ the same compilation pipeline established for the validation of our theoretical complexities. For the superconducting executions, the Qiskit compilation step was modified to target the specific IBM hardware backend, while for the trapped-ion executions, the optimized TKET circuits were compiled and executed via the Quantinuum Nexus platform. For the IBM hardware executions, baseline error mitigation was applied using Twirled Readout Error eXtinction (TREX)<sup>44</sup> via Qiskit's Resilience Level 1. The analysis is restricted to  $N \in [2^3, 2^7]$  to ensure the accumulated gate errors remain within a regime where valid tracks can be resolved from the background. For statistical validation, executions on the physical IBM Pittsburgh processor, as well as all noise-model emulators, were performed using 10,000 measurement shots across 5 independent repetitions. However, physical executions on the Quantinuum H2-2 hardware were constrained by quantum resource availability, necessitating a reduced sampling scheme of 3 independent repetitions. Within these H2-2 hardware runs, the first three tracking configurations were evaluated using 500 shots, while the fourth configuration utilized 800 shots. Despite this reduction in sampling volume, the allocated shot counts proved statistically sufficient to reconstruct the final probability distributions and extract performance metrics.

To quantify the impact of noise on track reconstruction performance we construct the following two metrics:

1. **Hellinger Fidelity ( $F_H$ )**: Measures the geometric overlap between the noisy experimental distribution  $P_{exp}$  and the ideal noiseless baseline  $P_{ideal}$ :

$$F_H(P_{exp}, P_{ideal}) = \left( \sum_i \sqrt{P_{exp}(x_i) \cdot P_{ideal}(x_i)} \right)^2 \quad (24)$$

2. **Signal Separation Index (SSI)**: This metric quantifies the distinguishability of valid tracks from the background. We define it as the ratio of the total probability mass of the  $M$  valid track segments ( $S$ ) to that of the  $M$  highest-amplitude noise segments ( $N^*$ ):

$$SSI = \frac{\sum_{i \in S} P_i}{\sum_{j \in N^*} P_j} \quad (25)$$

The Hellinger Fidelity ( $F_H$ ) results, presented in Fig. 6a, demonstrate a distinct performance hierarchy among the evaluated architectures, a trend that is consistent across both the physical hardware executions and the noise emulators. The Quantinuum H2 device consistently achieves the highest fidelity across all problem configurations when compared to the IBM Pittsburgh device. Furthermore, distinct deviation trends emerge between the physical executions and their emulators. For the IBM architecture, the emulator consistently predicts a slightly higher fidelity than the physical hardware achieves. Conversely, the Quantinuum H2-2 device demonstrates tight agreement between hardware and emulator for small configurations, while the physical hardware actually outperforms the emulator at larger problem sizes. The observed difference in performance between hardware devices may be caused by the topological differences between the architectures. The all-to-all connectivity of the trapped-ion H2 device allows the compiler (TKET) to synthesize the unitary with approximately half the circuit depth required for the fixed-topology superconducting chips. Since error accumulation is exponential with depth, this topological advantage and higher qubit fidelity<sup>16</sup> enhances the solution quality.

The Signal Separation Index (SSI), shown in Fig. 6b, corroborates these findings by mapping the distinguishability of valid tracks from fake segment combinations. The dashed horizontal line at  $SSI = 1$  defines the critical threshold below which the quantum filter fails to separate signal from noise. When comparing the physical executions to their respective emulators, the SSI trends mirror those observed in the Hellinger fidelity. The IBM Pittsburgh hardware performance is below its emulator predictions whereas the Quantinuum H2 hardware and emulator demonstrate agreement, although with larger statistical uncertainties in the hardware data due to a limited number of measurement shots. Consistent with the fidelity data, both the hardware and emulator results for the H2 platform maintain a resolvable signal ( $SSI > 1$ ) up to the denser 4-track, 5-layer configuration. In contrast, executions on hardware and emulators using the superconducting architecture approach the noise floor more rapidly. These results indicate that for current NISQ hardware, the effective connectivity and its direct reduction of compilation overhead may be a dominant factor in extending the reach of quantum track reconstruction to larger detector geometries.

## Discussion

In this paper we presented a 1-Bit Quantum Filter, a domain-specific adaptation of the HHL algorithm tailored for the combinatorial challenge of particle track reconstruction. Our results demonstrate that by replacing standard quantum linear algebra routines in favour of problem-informed filtering, we can achieve a polynomial gate complexity  $\mathcal{O}(\sqrt{N} \log N)$ .

A key finding of this study is the computational advantage gained by relaxing the requirements of the HHL algorithm, originally designed to perform precise matrix inversion  $A^{-1}b$ . This precision preserves the relative amplitudes of the solution vector for any well conditioned matrix, however, for the particle tracking task this is excessive. The physics problem is inherently binary: we aim to distinguish the "signal" subspace (valid track segments) from the "noise" subspace (combinatorial background segments), rather than computing a continuous amplitude distribution. This 1-Bit Quantum Filter exploits this relaxation by replacing the resource intensive multi-qubit QPE with a single-ancilla projection. The modification effectively replaces the expensive arithmetic inversion of the eigenvalues ( $1/\lambda$ ) with a simpler spectral step function, significantly reducing the circuit depth and the number of required ancilla qubits. The shift from calculating the exact amplitudes of the solution spectrum to filtering only track segments is what renders the algorithm feasible for the problem sizes required by the HL-LHC, where the overhead of full QPE would be prohibitive. This scalability is further supported by previous studies demonstrating that the Hamiltonian matrix becomes increasingly sparse with increasing event sizes<sup>11</sup>. To contextualize this scalability, HL-LHC luminosities during LHCb Run-5 are expected to yield 1,000 to 2,000 charged particles per event. Due to the combinatorial density of hits, this multiplicity pushes the matrix size to  $N \approx 10^6 - 10^7$  candidate segments. In contrast to Search-by-Triplet<sup>2</sup>, which is a local tracking algorithm, our Hamiltonian formulation operates as a global pattern recognition method. Integrating this global filtering into future heterogeneous, hybrid tracking pipelines could resolve complex combinatorial problems which will be seen at the HL-LHC.

While the gate complexity of our approach scales favorably, the total runtime remains constrained by the sampling complexity, which scales linearly with the system size  $\mathcal{O}(N)$ . To reconstruct the full event, the quantum state must be repeatedly prepared and measured until all  $M$  valid track segments are found. For large events on the order of the HL-LHC the linear readout cost is expensive. Future research must therefore pivot from full state tomography to the direct estimation of global properties. Instead of reading out the full distribution, one could construct quantum observables that can extract higher level physics features that are easy to implement on hardware. One such feature is the  $z$ -coordinate location of the Primary Vertices (PVs). In high-energy physics, the PV represents the primary point of particle collision on the beamline and at the HL-LHC 100 PVs may be present in an event. Their extraction at a low level reconstruction stage would greatly benefit complete event reconstruction and selection. By extracting this compressed information directly from the quantum state, it may be possible to bypass the  $\mathcal{O}(N)$  track readout bottleneck entirely and achieve sub-linear readout costs.

The results obtained by running our 1-bit HHL algorithm in Quantinuum and IBM hardware show promising outcomes for small systems of about four particle tracks across three detector layers. Larger systems scale badly with size as separation of good and back track segments disappears due to buildup of gate errors and the noise floor of current quantum devices. The results in Fig. 6 on fidelity and Signal Separation Index show a more favorable result for Quantinuum, which indicates that the all-to-all connectivity is better suited to our use case than a fixed topology requiring multiple SWAP operations. While our conclusions highlight current hardware constraints, future fault-tolerant quantum computers will use Quantum Error Correction (QEC) which may mitigate physical connectivity limits. How these codes will be implemented across architectures remains an open question. Ultimately, the algorithm's scalability will depend on balancing QEC overhead against the challenges of scaling physical qubits.

The benchmarks presented here evaluate our 1-Bit Quantum Filter on both physical quantum hardware and corresponding noise models. While the noise models capture general machine errors, our dual approach successfully accounts for the nuanced differences between simulated and physical noise across different hardware platforms. While current NISQ-era hardware restricts the solvable problem size, these physical executions experimentally validate our proof-of-principle. Future work will focus on benchmarking the algorithm within logical error-correction layers, optimizing both the Hamiltonian formulation and this 1-Bit Quantum Filter, the direct extraction of global event observables such as Primary Vertices, and the efficient calculation and construction of the Hamiltonian matrix.

## Conclusion

Our 1-Bit Quantum Filter represents a development towards a potential quantum-classical hybrid particle reconstruction algorithm, focusing towards specialized, problem-specific spectral filtering. We have validated the theoretical two-qubit gate complexity  $\mathcal{O}(\sqrt{N} \log N)$  and simulated our 1-Bit Quantum Filter on realistic LHCb Monte Carlo events, a task previously inaccessible by quantum tracking algorithms. We demonstrate that for small systems the algorithm is robust against noise profiles of current generation hardware, particularly on architectures with high connectivity. As quantum hardware matures from noisy physical qubits to logically coherent systems, this 1-Bit Quantum Filter deserves further study to reconstruct events at the scale of the future HL-LHC. More broadly, this 1-Bit Quantum Filter technique can be adapted to similar optimization problems such as the ones arising in the context of graph optimization, associative memories and pattern denoising, reducing the practical requirements down to sizes that are accessible by quantum hardware of today.

## Data availability

The source data used in this study for Fig. 1, Fig. 2, Fig. 4, Fig. 5, Fig. 6 can be found in the data repository available at <https://github.com/Xenofon-Chiotopoulos/OneBQF>.

## Code availability

The code used in this work is open source and available at <https://github.com/Xenofon-Chiotopoulos/OneBQF>

## References

1. Albrecht, J. *et al.* A Roadmap for HEP Software and Computing R&D for the 2020s. *Comput. Softw. for Big Sci.* **3**, DOI: [10.1007/s41781-018-0018-8](https://doi.org/10.1007/s41781-018-0018-8) (2019).
2. Cámpora Pérez, D. H., Neufeld, N. & Riscos Núñez, A. Search by triplet: An efficient local track reconstruction algorithm for parallel architectures. *J. Comput. Sci.* **54**, 101422, DOI: [10.1016/j.jocs.2021.101422](https://doi.org/10.1016/j.jocs.2021.101422) (2021).
3. Ju, X. *et al.* Performance of a geometric deep learning pipeline for HL-LHC particle tracking. *The Eur. Phys. J. C* **81**, DOI: [10.1140/epjc/s10052-021-09675-8](https://doi.org/10.1140/epjc/s10052-021-09675-8) (2021).
4. Correia, A., Giasemis, F. I., Garroum, N., Gligorov, V. V. & Granado, B. Graph Neural Network-based track finding in the LHCb vertex detector. *J. Instrumentation* **19**, P12022, DOI: [10.1088/1748-0221/19/12/p12022](https://doi.org/10.1088/1748-0221/19/12/p12022) (2024).
5. Jaschke, D. & Montangero, S. Is quantum computing green? An estimate for an energy-efficiency quantum advantage. *Quantum Sci. Technol.* **8**, 025001, DOI: [10.1088/2058-9565/acae3e](https://doi.org/10.1088/2058-9565/acae3e) (2023).
6. Auffèves, A. Quantum Technologies Need a Quantum Energy Initiative. *PRX Quantum* **3**, 020101, DOI: [10.1103/PRXQuantum.3.020101](https://doi.org/10.1103/PRXQuantum.3.020101) (2022).
7. Zlokapa, A. *et al.* Charged particle tracking with quantum annealing optimization. *Quantum Mach. Intell.* **3**, DOI: [10.1007/s42484-021-00054-w](https://doi.org/10.1007/s42484-021-00054-w) (2021).
8. Saito, Masahiko *et al.* Quantum annealing algorithms for track pattern recognition. *EPJ Web Conf.* **245**, 10006, DOI: [10.1051/epjconf/202024510006](https://doi.org/10.1051/epjconf/202024510006) (2020).
9. Tüysüz, C. *et al.* Hybrid quantum classical graph neural networks for particle track reconstruction. *Quantum Mach. Intell.* **3**, DOI: [10.1007/s42484-021-00055-9](https://doi.org/10.1007/s42484-021-00055-9) (2021).
10. Harrow, A. W., Hassidim, A. & Lloyd, S. Quantum Algorithm for Linear Systems of Equations. *Phys. Rev. Lett.* **103**, DOI: [10.1103/physrevlett.103.150502](https://doi.org/10.1103/physrevlett.103.150502) (2009).
11. Nicotra, D. *et al.* A quantum algorithm for track reconstruction in the LHCb vertex detector. *J. Instrumentation* **18**, P11028, DOI: [10.1088/1748-0221/18/11/P11028](https://doi.org/10.1088/1748-0221/18/11/P11028) (2023).
12. Chiotopoulos, Xenofon *et al.* TrackHHL: A Quantum Computing Algorithm for Track Reconstruction at the LHCb. *EPJ Web Conf.* **337**, 01181, DOI: [10.1051/epjconf/202533701181](https://doi.org/10.1051/epjconf/202533701181) (2025).
13. Denby, B. Neural networks and cellular automata in experimental high energy physics. *Comput. Phys. Commun.* **49**, 429–448, DOI: [https://doi.org/10.1016/0010-4655\(88\)90004-5](https://doi.org/10.1016/0010-4655(88)90004-5) (1988).
14. Peterson, C. Track finding with neural networks. *Nucl. Instruments Methods Phys. Res. Sect. A: Accel. Spectrometers, Detect. Assoc. Equip.* **279**, 537–545, DOI: [https://doi.org/10.1016/0168-9002\(89\)91300-4](https://doi.org/10.1016/0168-9002(89)91300-4) (1989).
15. AbuGhanem, M. IBM quantum computers: evolution, performance, and future directions. *The J. Supercomput.* **81**, DOI: [10.1007/s11227-025-07047-7](https://doi.org/10.1007/s11227-025-07047-7) (2025).
16. Moses, S. *et al.* A Race-Track Trapped-Ion Quantum Processor. *Phys. Rev. X* **13**, DOI: [10.1103/physrevx.13.041052](https://doi.org/10.1103/physrevx.13.041052) (2023).
17. Aaij, R. *et al.* The LHCb Upgrade I. *J. Instrumentation* **19**, P05065, DOI: [10.1088/1748-0221/19/05/P05065](https://doi.org/10.1088/1748-0221/19/05/P05065) (2024).
18. Rebentrost, P., Bromley, T. R., Weedbrook, C. & Lloyd, S. Quantum hopfield neural network. *Phys. Rev. A* **98**, 042308, DOI: [10.1103/PhysRevA.98.042308](https://doi.org/10.1103/PhysRevA.98.042308) (2018).
19. Cleve, R., Ekert, A., Macchiavello, C. & Mosca, M. Quantum algorithms revisited. *Proc. Royal Soc. London. Ser. A: Math. Phys. Eng. Sci.* **454**, 339–354, DOI: [10.1098/rspa.1998.0164](https://doi.org/10.1098/rspa.1998.0164) (1998).
20. Kitaev, A. Y. Quantum measurements and the Abelian Stabilizer Problem. *Electron. Colloquium Comput. Complex.* **TR96**, DOI: [10.48550/arXiv.quant-ph/9511026](https://doi.org/10.48550/arXiv.quant-ph/9511026) (1995).

21. Coppersmith, D. An approximate Fourier transform useful in quantum factoring. *arXiv: Quantum Phys.* DOI: [10.48550/arXiv.quant-ph/0201067](https://doi.org/10.48550/arXiv.quant-ph/0201067) (2002).
22. Aaronson, S. Read the fine print. *Nat. Phys.* **11**, 291–293, DOI: [10.1038/nphys3272](https://doi.org/10.1038/nphys3272) (2015).
23. Dervovic, D. *et al.* Quantum linear systems algorithms: a primer. *ArXiv* **abs/1802.08227**, DOI: [10.48550/arXiv.1802.08227](https://doi.org/10.48550/arXiv.1802.08227) (2018).
24. Suzuki, M. Generalized Trotter's formula and systematic approximants of exponential operators and inner derivations with applications to many-body problems. *Commun. Math. Phys.* **51**, 183–190, DOI: [10.1007/BF01609348](https://doi.org/10.1007/BF01609348) (1976).
25. Horn, R. A. & Johnson, C. R. Matrix Analysis. DOI: [10.1017/CBO9780511810817](https://doi.org/10.1017/CBO9780511810817) (1985).
26. Nielsen, M. A. & Chuang, I. L. *Quantum Computation and Quantum Information: 10th Anniversary Edition* (Cambridge University Press, 2010).
27. Cowtan, A., Dilkes, S., Duncan, R., Simmons, W. & Sivarajah, S. Phase Gadget Synthesis for Shallow Circuits. *Electron. Proc. Theor. Comput. Sci.* **318**, 213–228, DOI: [10.4204/eptcs.318.13](https://doi.org/10.4204/eptcs.318.13) (2020).
28. Di Matteo, O. *et al.* Improving Hamiltonian encodings with the Gray code. *Phys. Rev. A* **103**, 042405, DOI: [10.1103/PhysRevA.103.042405](https://doi.org/10.1103/PhysRevA.103.042405) (2021).
29. Barenco, A. *et al.* Elementary gates for quantum computation. *Phys. Rev. A* **52**, 3457–3467, DOI: [10.1103/physreva.52.3457](https://doi.org/10.1103/physreva.52.3457) (1995).
30. Strandlie, A. & Frühwirth, R. Track and vertex reconstruction: From classical to adaptive methods. *Rev. Mod. Phys.* **82**, 1419–1458, DOI: [10.1103/RevModPhys.82.1419](https://doi.org/10.1103/RevModPhys.82.1419) (2010).
31. Grover, L. K. A fast quantum mechanical algorithm for database search. In *Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing*, STOC '96, 212–219, DOI: [10.1145/237814.237866](https://doi.org/10.1145/237814.237866) (Association for Computing Machinery, New York, NY, USA, 1996).
32. Motwani, R. & Raghavan, P. Randomized algorithms. In *CSUR*, DOI: [10.1145/234313.234327](https://doi.org/10.1145/234313.234327) (1995).
33. Sivarajah, S. *et al.* tket): a retargetable compiler for NISQ devices. *Quantum Sci. Technol.* **6**, 014003, DOI: [10.1088/2058-9565/ab8e92](https://doi.org/10.1088/2058-9565/ab8e92) (2020).
34. Quantinuum. Quantinuum Nexus: A Cloud-Based Quantum Computing Platform. <https://docs.quantinuum.com/nexus> (2024). Accessed: 2024-12-16.
35. Javadi-Abhari, A. *et al.* Quantum computing with Qiskit. *ArXiv* **abs/2405.08810**, DOI: [10.48550/arXiv.2405.08810](https://doi.org/10.48550/arXiv.2405.08810) (2024).
36. Sjöstrand, T., Mrenna, S. & Skands, P. Pythia 6.4 physics and manual. *J. High Energy Phys.* **2006**, 026, DOI: [10.1088/1126-6708/2006/05/026](https://doi.org/10.1088/1126-6708/2006/05/026) (2006).
37. Belyaev, I. *et al.* Handling of the generation of primary events in gauss, the lhcb simulation framework. *J. Physics: Conf. Ser.* **331**, 032047, DOI: [10.1088/1742-6596/331/3/032047](https://doi.org/10.1088/1742-6596/331/3/032047) (2011).
38. Lange, D. J. The evtgen particle decay simulation package. *Nucl. Instruments Methods Phys. Res. Sect. A: Accel. Spectrometers, Detect. Assoc. Equip.* **462**, 152–155, DOI: [https://doi.org/10.1016/S0168-9002\(01\)00089-4](https://doi.org/10.1016/S0168-9002(01)00089-4) (2001). BEAUTY2000, Proceedings of the 7th Int. Conf. on B-Physics at Hadron Machines.
39. Agostinelli, S. *et al.* Geant4—a simulation toolkit. *Nucl. Instruments Methods Phys. Res. Sect. A: Accel. Spectrometers, Detect. Assoc. Equip.* **506**, 250–303, DOI: [https://doi.org/10.1016/S0168-9002\(03\)01368-8](https://doi.org/10.1016/S0168-9002(03)01368-8) (2003).
40. Clemencic, M. *et al.* The lhcb simulation application, gauss: Design, evolution and experience. *J. Physics: Conf. Ser.* **331**, 032023, DOI: [10.1088/1742-6596/331/3/032023](https://doi.org/10.1088/1742-6596/331/3/032023) (2011).
41. Cámpora Pérez, D. H. velopix\_tracking: Python implementation of the Search by Triplet algorithm. [https://github.com/dcampa/velopix\\_tracking](https://github.com/dcampa/velopix_tracking) (2022).
42. Peterson, C. Track finding with neural networks. *Nucl. Instruments Methods Phys. Res. Sect. A: Accel. Spectrometers, Detect. Assoc. Equip.* **279**, 537–545, DOI: [https://doi.org/10.1016/0168-9002\(89\)91300-4](https://doi.org/10.1016/0168-9002(89)91300-4) (1989).
43. Denby, B. Neural networks and cellular automata in experimental high energy physics. *Comput. Phys. Commun.* **49**, 429–448, DOI: [https://doi.org/10.1016/0010-4655\(88\)90004-5](https://doi.org/10.1016/0010-4655(88)90004-5) (1988).
44. van den Berg, E., Mineev, Z. K. & Temme, K. Model-free readout-error mitigation for quantum expectation values. *Phys. Rev. A* **105**, 032620, DOI: [10.1103/PhysRevA.105.032620](https://doi.org/10.1103/PhysRevA.105.032620) (2022).

## Acknowledgements

We are thankful to the CERN Quantum Technology Initiative (QTI) for providing access to IBM Quantum computers and Quantinuum for providing access to their quantum platforms, resulting in the studies reported in this paper. We also thank the Coherence and Quantum Technology group of Eindhoven University for initial tests on their simulation platform. Finally, we are grateful for the support from the LHCb Data Processing and Analysis (DPA) project, and acknowledge the LHCb simulation and computing projects for the production of simulated samples.

## Funding Statement

X. Chiotopoulos discloses support for the research of this work from the NWO project *Fast Sensors and Algorithms for Space-time Tracking and Event Reconstruction (FASTER)* [grant number OCENW.XL21.XL21.076], D. Nicotra discloses support for part of the research from SURF innovation fund: *From Quark to Quantum* [SOIL.QQLHCb.01], and G. Scriven discloses support for the research of his work by the Special Research Fund of Hasselt University [grant number BOF24DOCUM01]. K. Driessens, M. Merk, J. Schütz, J. de Vries and M. Winands declare no relevant funding.

## Author contribution

### Authors and Affiliations

**Nikhef National Institute for Subatomic Physics, Science Park 105, 1098 XG Amsterdam, The Netherlands**

Xenofon Chiotopoulos, Marcel Merk

**Maastricht University, Faculty of Science and Engineering, Gravitational Waves and Fundamental Physics department, Duboisdomain 30, 6229 GT Maastricht, The Netherlands**

Xenofon Chiotopoulos, George Scriven, Davide Nicotra, Jacco de Vries, Marcel Merk

**Hasselt University, Faculty of Sciences and Data Science Institute, Agoralaan gebouw D, 3590 Diepenbeek, Belgium**

George Scriven, Jochen Schütz

**Maastricht University, Faculty of Science and Engineering, Department of Advanced Computing Sciences, Paul-Henri Spaaklaan 1, 6229 EN Maastricht, The Netherlands**

Xenofon Chiotopoulos, Kurt Driessens, Mark H.M. Winands

### Contributions

X.C. derived and implemented the 1-Bit Quantum Filter and the Direct Structural Synthesis (DSS) method, performed the theoretical and empirical data fitting complexity analysis with cross checks from D.N., performed the hardware and emulator executions on both platforms, and carried out the tracking benchmarking on LHCb Monte Carlo events. X.C. M.M., G.S., D.N., K.D., M.H.M.W. and J.d.V. contributed to the interpretation of these benchmarks. G.S. developed the simulation framework based on original work by D.N., updated the HHL implementation to the latest Qiskit version, and together with M.M. developed the toy simulation model. J.d.V., K.D., M.M. and D.N. involved in the discussions of the Hamiltonian formulation. J.S. contributed to the mathematical consistency of the work. All authors participated in the theoretical and algorithm discussions, contributed to the interpretation of the results and edited the manuscript.

### Corresponding author

Correspondence to Xenofon Chiotopoulos (xenofon.chiotopoulos@maastrichtuniversity.nl)

## Competing interests

The authors declare no competing interests.

## Figure Captions

### Fig. 1: Event simulation mapping to the tracking graph.

(a) A simulated event in the LHCb detector. The larger grey dots represent the primary vertices of multiple simultaneous proton-proton collisions, while the small black dots indicate the detector hits generated by the produced particles. The faint grey lines depict the reconstructed tracks of the produced particles. The red dots and lines indicate the hits and particles produced in a single given collision, represented by the red circle. (b) Illustration of the definition of the Hamiltonian graph construction<sup>11</sup>. Coloured segments represent active variables  $S_i = 1$  forming valid tracks, while grey segments correspond to inactive variables  $S_j = 0$ .

### Fig. 2: Tracking performance metrics of the classical matrix solver.

(a) Visual event display of a simulated collision with a single primary vertex, 5 detection layers, and 64 particle tracks originating from a single primary vertex. The valid track segments are shown as green lines, and the inactive combinatorial background segments are shown as black lines. (b,c,d,e) Classical-solver segment performance versus track multiplicity  $T$  at fixed angular tolerance  $\epsilon = 2$  mrad with a 1% random hit inefficiency. Markers have vertical error bars representing the standard error of the mean ( $\text{SEM} = \sigma/\sqrt{n}$ ) across independent toy events at each  $T$  ( $n = 60$  events for  $T \leq 100$ ,  $n = 30$  for  $150 \leq T \leq 300$ ,  $n = 15$  for  $T \geq 500$ ). (b) Segment efficiency, the fraction of true segments activated by the solver; red circles show the average as function of the number of tracks in the event. (c) Segment false rate, the fraction of incorrectly activated segments, shown as magenta squares. (d) The number of candidate segments prior to solving: blue circles are true segments, red-orange squares are false (ghost) segments. (e) The segments activated by the solver, for true active (blue-circle) and false active segments (red-orange-square) respectively.

### Fig. 3: Circuit Architecture.

(a) Circuit diagram of the proposed 1-Bit Quantum Filter, with registers that include the single ancilla flag  $|0\rangle_A$ , the single time register  $|0\rangle_T$  for the 1-bit QPE, and the system register  $|0\rangle_{S_1} \dots |0\rangle_{S_{n_s}}$ , where  $n_s = \log_2 N$  represents the total number of system qubits. The algorithm is composed of state preparation, 1-bit phase estimation with flagging logic, and uncomputation. The additional phase gate  $P(\phi)$  on the time register implements the diagonal term of the Hamiltonian. (b) The Direct Structural Synthesis (DSS) decomposition, explicitly showing that the controlled time evolution  $U_{\text{DSS}}$  is implemented as a sequential application of controlled interaction gates  $\prod_k G_k = \prod_k e^{iB_k t}$  corresponding to the non-zero off-diagonal terms of the Hamiltonian matrix. (c) The generalized structure of a single interaction gate  $G_k = e^{iB_k t}$ . For states with Hamming distance  $d_H > 1$ , a CNOT ladder ( $S_{\text{lad}}$ ) disentangles the states onto a pivot ( $S_{\text{piv}}$ ), while passive qubits ( $S_{\text{pass}}$ ) act as controls to strictly enforce subspace selection.

### Fig. 4: Asymptotic complexity scaling.

(a) Gate count complexity scaling for one-bit quantum filter implementations under IBM Pittsburgh hardware compilation. (b) Gate count complexity scaling under Quantinuum system model H2. Power-law fits with constrained logarithmic scaling ( $c = 1.0$ ) yield scaling exponents of  $b \approx 0.53$  for IBM and  $b \approx 0.52$  for Quantinuum, consistent with the  $\mathcal{O}(\sqrt{N} \log N)$  theoretical gate complexity ( $R^2 > 0.999$ ). The three-layer tracking configurations are represented by blue dotted lines with circle markers, and the five-layer configurations are represented by green dotted lines with triangle markers. In all panels, the text labels annotated next to the data points (e.g., 4p, 128p) indicate the number of particles simulated to generate the corresponding matrix size. (c) Success probability  $P_{\text{succ}}$  as a function of matrix size  $N$  for three-layer (blue solid line with circle markers) and five-layer (green solid line with triangle markers) tracking scenarios. The fitted exponents exactly match the theoretical  $N^{-1/2}$  scaling up to  $N = 10^6$ , where an upward deviation is observed in the residuals. The error bars represent one standard deviation of the mean, calculated across a sample size of  $n = 10$  independent statevector simulations, each utilizing  $10^7$  measurement shots. The fitted scaling coefficient  $c_{\text{filter}}$  matches the theoretically predicted value for both the three-layer and five-layer configurations.

### Fig. 5: Monte Carlo Performance comparisons.

The performance of our 1-Bit Quantum Filter simulation (green triangles) on realistic LHCb simulated (MC) events, is benchmarked against the classical state-of-the-art Search by Triplet<sup>2</sup> algorithm (blue circles) across varying event sizes. The horizontal error bars represent equal density bins and the vertical error bars represent a 1 standard deviation of the mean in that bin for which 79 events were used per bin. (a) Segment Finding Efficiency ( $\epsilon_{\text{segment}}$  Eq. 4), representing the fraction of true track segments successfully recovered as a function of the total number of detector hits. (b) Segment fake rate ( $f_{\text{segment}}$  Eq. 4), representing the fraction of incorrectly identified segments, as a function of the total number of detector hits.

### Fig. 6: Hardware performance comparison.

Both panels present metrics evaluating tracking results for problem sizes up to 8 tracks across 3 layers, utilizing both quantum

hardware and noise emulators. All results were obtained across the 5 tracking scenarios, except for the physical Quantinuum H2 hardware executions, which are limited to the first 4 scenarios due to deep circuits and constrained quantum resources. For IBM Pittsburgh all points were generated across 5 runs with 10,000 shots each. For system model H2 the points were generated across 5 runs with the first 3 using 500 shots and the 4th with 800 shots. Error bars denote  $1\sigma$  confidence intervals calculated directly across independent algorithm executions. **(a)** Hellinger fidelity ( $F_H$ ) relative to the noiseless baseline indicating the overlap of the probability distribution. **(b)** Signal Separation Index (SSI) quantifying the contrast between valid tracks and the noise floor. The dashed line at  $SSI = 1$  represents the distinguishability limit where signal amplitudes drop below the background noise.

ARTICLE IN PRESS