

UC Berkeley

UC Berkeley Previously Published Works

Title

Improving quantum simulation efficiency of final state radiation with dynamic quantum circuits

Permalink

<https://escholarship.org/uc/item/1b65763q>

Journal

Physical Review D, 106(3)

ISSN

2470-0010

Authors

Deliyannis, Plato

Sud, James

Chamaki, Diana

et al.

Publication Date

2022-08-01

DOI

10.1103/physrevd.106.036007

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed

Improving Quantum Simulation Efficiency of Final State Radiation with Dynamic Quantum Circuits

Plato Deliyannis,^{1,*} James Sud,^{2,†} Diana Chamaki,^{2,‡} Zoë Webb-Mack,^{1,§} Christian W. Bauer,^{1,¶} and Benjamin Nachman^{1,**}

¹*Physics Division, Lawrence Berkeley National Laboratory, Berkeley, CA 94720, USA*

²*Department of Physics, UC Berkeley, Berkeley, CA 94720, USA*

(Dated: March 21, 2022)

Reference [1] recently introduced an algorithm (QPS) for simulating parton showers with intermediate flavor states using polynomial resources on a digital quantum computer. We make use of a new quantum hardware capability called *dynamical quantum computing* to improve the scaling of this algorithm to significantly improve the method precision. In particular, we modify the quantum parton shower circuit to incorporate mid-circuit qubit measurements, resets, and quantum operations conditioned on classical information. This reduces the computational depth from $\mathcal{O}(N^5 \log_2(N)^2)$ to $\mathcal{O}(N^3 \log_2(N)^2)$ and the qubit requirements are reduced from $\mathcal{O}(N \log_2(N))$ to $\mathcal{O}(N)$. Using “matrix product state” statevector simulators, we demonstrate that the improved algorithm yields expected results for 2, 3, 4, and 5-steps of the algorithm. We compare absolute costs with the original QPS algorithm, and show that dynamical quantum computing can significantly reduce costs in the class of digital quantum algorithms representing quantum walks (which includes the QPS).

I. INTRODUCTION

High energy physics (HEP) simulations are one of the most natural and exciting applications of quantum computers, given the complex many-body quantum nature of HEP processes. Foundational work establishing the existence of polynomial scaling digital quantum algorithms for scattering calculations [2] has been followed by a variety of particle and nuclear physics investigations into simulations on quantum computers¹.

While most of these studies propose quantum algorithms for the full scattering process on a quantum computer, a complementary approach has been proposed to exploit factorization [4]. In particular, scattering cross sections approximately factor into pieces governed by physical processes occurring at different energy scales. One way to factorize a full calculation at a collider like the Large Hadron Collider (LHC) involves parton shower (PS) Monte Carlo. Parton showers govern the collinear radiation from high energy charged particles. Classical PSs approximate this radiation as a Markov Chain. This is an excellent approximation in some cases, but ignores certain interference effects. Recently, Ref. [1] introduced a quantum algorithm for parton showers (QPS) that models interference effects from intermediate flavor states. This algorithm requires only polynomial resources compared to existing exponentially scaling algorithms to model the

same physics. While the QPS does not describe the full properties of PSs in the Standard Model (dominated mostly by the strong force), it represents an important benchmark for developing and testing HEP algorithms on quantum computers.

Even though the QPS requires only polynomial quantum resources, it is still challenging to run on existing devices. This is because we are in the Noisy Intermediate Scale Quantum (NISQ) [5] computing era where qubit counts, connectivities, and coherence times are limited and quantum gate and readout operations have significant noise. Therefore, there is a strong motivation to improve the polynomial scaling of the current quantum algorithms like the QPS.

In this paper, we improve the original QPS algorithm [1] by using *dynamical quantum circuits*: we incorporate mid-circuit qubit measurements² and quantum gates applied dynamically based on results of classical processing on these measurements. The resulting quantum state prepared by the modified protocol is equivalent to the original. By re-setting the measured qubits to the ground state and re-using them for subsequent iterations the computational complexity in both qubits and gates can be reduced. Compared to the original algorithm, the new version uses a factor of $\mathcal{O}(N^2)$ fewer standard entangling gates, where N is the number of points used to discretize the PS.

This paper is organized as follows. Section II introduces dynamic quantum computing, in which a quantum processing unit (QPU) interacts with a classical processing unit (CPU) *during* computation. This is in contrast with nearly all current digital quantum algorithms, where a preset sequence of quantum gates are applied, and the

* pdeliyannis@lbl.gov

† jamessud@berkeley.edu

‡ dchamaki@berkeley.edu

§ zwebbmack@lbl.gov

¶ cwbauer@lbl.gov

** bpnachman@lbl.gov

¹ There are now over a hundred papers in this area; see Ref. [3] for an overview to this topic and the papers that cite Ref. [2] for specific studies.

² Throughout the paper, we use the terms ‘mid-circuit measurements’ and ‘remeasurement’ interchangeably.

system is measured only as the final step of the computation. Next, Sec. III begins by briefly reviewing the QPS algorithm [1], including its qubit and quantum gate requirements. Section III then continues by introducing a modified QPS that incorporates mid-circuit measurement and quantum-classical feedback. We have implemented this algorithm in `Qiskit` [6] and provide numerical results in Sec. IV. The simulations agree with those in Ref. [1] and we are able to make more precise predictions than were possible before, using the matrix product state simulator. The paper ends with conclusions in Sec. V.

II. DYNAMIC QUANTUM COMPUTING

Dynamic quantum computing involves dividing a program into (1) steps that can only be implemented in a quantum computer (QPU) and (2) steps that can be implemented more efficiently on a classical computer (CPU), and interfacing between the QPU and CPU wherever necessary. This scheme manifests in two categorically distinct ways.

First, one could execute an algorithm consisting of a sequence of alternating quantum and classical steps, where the result of each step is fed serially to the next. Variational algorithms such as the variational quantum eigensolver (VQE) [7] are primary examples of this scheme. To compute the smallest eigenvalue of a Hermitian operator \mathcal{H} , VQE alternates between a quantum step that computes the expectation of \mathcal{H} on some vector $|\psi(\vec{\theta})\rangle$, and a classical step that minimizes the expectation $\langle\psi(\vec{\theta})|\mathcal{H}|\psi(\vec{\theta})\rangle$ over parameters $\vec{\theta}$. In this scheme, each quantum step is independent from the previous, so the same QPU is reset and reused for all quantum steps. This means that the QPU's coherence time must be long enough only to accommodate the time to execute a single quantum step.

On the other hand, one could also construct an algorithm that requires rapid interfacing between QPU and CPU. For example, standard quantum error correction (QEC) procedures entail measuring a *syndrome* operator, followed by a *recovery operation* in which pre-defined quantum gates are applied conditional on measurement results [8]. Therefore, incorporating QEC to some general quantum program requires frequent interfacing between QPUs and CPUs in the form of measurements (QPU to CPU) and feedback (CPU to QPU). In addition to direct feedback based on measurement results, one could also perform classical computations on them before applying quantum gates conditioned on the results of those computations. This procedure describes a fully dynamic quantum-classical computer and is the basis for the algorithm presented in this paper. Figure 1 illustrates the operation of a dynamic quantum computer.

In contrast with VQE, where quantum and classical steps alternate serially, QEC and other algorithms that use a similar rapid interfacing scheme require that quantum resources (qubits) maintain coherence during

quantum-classical interfacing. Limited qubit coherence times are a major bottleneck for implementing QEC and other dynamic algorithms. However, coherence times continue to improve as hardware is developed and refined, and because of the importance of QEC for developing fault-tolerant quantum computers, dynamic hardware will be a major focus in the long term. Until now, demonstrations of dynamic computing with rapid digital interfacing include active qubit resets [9–11], quantum teleportation [12–16], and error-correction [17–26]. A more complex example was recently demonstrated by IBM [27], employing a hybrid quantum-classical version of phase estimation on two qubits, where an m -bit representation of the phase is computed using several shots of a hybrid circuit that contains m measurement-feed-forward cycles. For each cycle, an $R_z(\theta)$ gate is selected and applied conditionally on previous measurement results. Different measurement results select different θ values.

The hybrid phase estimation algorithm demonstrated by IBM [27] closely resembles the hybrid QPS we will introduce below in that it consists of single-qubit rotation gates with rotation angles that depend on classical information, i.e. mid-circuit measurement results. Therefore, IBM's demonstration of executing a hybrid quantum algorithm shows that our QPS algorithm can theoretically be implemented on real devices in the future. In Sec. III, we summarize the original QPS algorithm [1] before laying out the dynamical version.

III. QUANTUM PARTON SHOWER ALGORITHM

First, we summarize the quantum parton shower (QPS) algorithm originally presented in Ref. [1].

A. Physical Background

Parton shower algorithms are perturbative approaches to efficiently describe high-multiplicity final states by focusing on the soft and collinear regions of phase space. The QPS algorithm in Ref. [1] was developed for a simple quantum field theory involving two types of fermion fields, f_1 and f_2 , interacting with one scalar boson ϕ governed by the following Lagrangian:

$$\mathcal{L} = \bar{f}_1(i\cancel{\partial} + m_1)f_1 + \bar{f}_2(i\cancel{\partial} + m_2)f_2 + (\partial_\mu\phi)^2 + g_1\bar{f}_1f_1\phi + g_2\bar{f}_2f_2\phi + g_{12}[\bar{f}_1f_2 + \bar{f}_2f_1]\phi. \quad (1)$$

The first three terms in Eq. (1) describe the kinematic properties of the fermions and scalar while the latter three terms govern their interactions. The goal of a PS algorithm is to describe the collinear dynamics of the theory, which in this case correspond to the fermions radiating scalars ($f_i \rightarrow f_j\phi$) and scalars splitting into fermion pairs ($\phi \rightarrow f_i\bar{f}_j$).




FIG. 1: Dynamic computing workflow. The essential procedure consists of four steps: (1) Measure a subset of quantum resources in the QPU, represented here by the k -qubit register, (2) On the CPU perform some processing on the measured data, (3) Reset the measured qubits to the $|0\rangle$ state so they can be re-used, and (4) Based on CPU outputs, apply additional quantum operations on the QPU. Note that the QPU must maintain coherence throughout this procedure.

In classical PSs, the rates of these processes are described by splitting functions:

$$P_{i \rightarrow j\phi}(\theta) = g_{ij}^2 P_f(\theta), \quad i, j \in \{1, 2\} \quad (2)$$

$$P_{\phi \rightarrow i\bar{j}}(\theta) = g_{ij}^2 \hat{P}_\phi(\theta), \quad i, j \in \{1, 2\}, \quad (3)$$

where $g_i \equiv g_{ii}$. The splitting functions describe the probability for a particular particle at a given step (parameterized by the scale θ) in the parton shower evolution to undergo a transformation. There are many formally equivalent definitions of the scale; here we use a common choice of the opening angle of the emission with respect to the emitter (angularly ordered shower).

In addition to the splitting functions, another important quantity is the no-branching probability (Sudakov factor):

$$\Delta_{i,k}(\theta_1, \theta_2) = \exp \left[-g_i^2 \int_{\theta_1}^{\theta_2} d\theta' \hat{P}_k(\theta') \right], \quad (4)$$

which describes the probability that no emission occurs between scales θ_1 and θ_2 . With the splitting functions and Sudakov factors, we can sample from the cross-section using a Markov Chain algorithm that generates one emission at a time, conditioned on the previous emissions. In particular, at a given step in the algorithm with a fixed number of particles, the probability that none of them radiate or split is simply a product over Sudakov factors. If something does happen at a given step, the probabilities are proportional to the appropriate splitting function. In the limit $g_{12} \rightarrow 0$, the Markov Chain algorithm can be implemented in terms of emission probabilities, and is therefore classically efficient. However, if $g_{12} > 0$, there are now multiple histories with unmeasured intermediate fermion types which contribute to the same final state. To account for these interferences, we must implement the Markov Chain on at the quantum amplitude level, which necessitates keeping track of $\mathcal{O}(2^N)$ different histories, where the range of opening angles is discretized into N parts. This motivates the QPS algorithm, which computes the final state radiation with $g_{12} > 0$ using only polynomially many qubits and gates on a digital quantum computer.

B. Basis for the Quantum Algorithm

The interaction terms (Eq. (1)) of the Lagrangian can be written as a matrix equation:

$$\mathcal{L}_{\text{interaction}} = (\bar{f}_1 \quad \bar{f}_2) \begin{pmatrix} g_1 & g_{12} \\ g_{12} & g_2 \end{pmatrix} \begin{pmatrix} f_1 \\ f_2 \end{pmatrix} \phi. \quad (5)$$

Furthermore, the “interaction matrix” is real and symmetric, and can thus be diagonalized as

$$G \equiv \begin{pmatrix} g_1 & g_{12} \\ g_{12} & g_2 \end{pmatrix} = U^\dagger \begin{pmatrix} g_a & 0 \\ 0 & g_b \end{pmatrix} U. \quad (6)$$

By defining a change of basis

$$\begin{pmatrix} f_a \\ f_b \end{pmatrix} \equiv U \begin{pmatrix} f_1 \\ f_2 \end{pmatrix}, \quad (7)$$

the interactions (Eq. (1)) become diagonal:

$$\mathcal{L}_{\text{interaction}} = (\bar{f}_a \quad \bar{f}_b) \begin{pmatrix} g_a & 0 \\ 0 & g_b \end{pmatrix} \begin{pmatrix} f_a \\ f_b \end{pmatrix} \phi. \quad (8)$$

In this “diagonal basis”, splitting do not create interference between fermion types. In other words,

$$P_{i \rightarrow j\phi}(\theta) = \delta_{ij} g_{ij}^2 P_f(\theta), \quad i, j \in \{a, b\} \quad (9)$$

$$P_{\phi \rightarrow i\bar{j}}(\theta) = \delta_{ij} g_{ij}^2 P_\phi(\theta), \quad i, j \in \{a, b\}, \quad (10)$$

where δ_{ij} is Kronecker delta. This is also the case in the original basis if $g_{12} = 0$.

Therefore, to simulate interference between fermion types, we first rotate particle registers $|p\rangle$ encoding fermion/boson fields into the diagonal basis according to Eq. (7), proceed with a quantum analogue of the classical Markov Chain algorithm (generating a *history* of angles and particle types), and lastly rotate the final particle states back to the original basis. If $g_{12} > 0$, then the initial rotation to the diagonal basis creates a superposition of f_a and f_b fermions. Subsequent operations act on this superposition, and all intermediate amplitudes/histories are preserved throughout the quantum Markov Chain. This contrasts the classical MCMC parton shower, where there superpositions of multiple fermion types are not included. Note that in this model, there is no interference between histories where emissions occurred at different angles.

C. Original Quantum Algorithm

The QPS circuit uses three qubit registers to encode the particle state and history, three registers to store derived quantities about the number of particles, and a number of ancillary qubits. For N steps of the algorithm (i.e, discretizing the range of angles into N parts) and with n_I initial particles, the qubit counts are provided in Table I.

Register	Purpose	# qubits
$ p\rangle$	Particle state	$3(N + n_I)$
$ h\rangle$	Emission history	$\sum_{m=0}^{N-1} \lceil \log_2(m + n_I + 1) \rceil$
$ e\rangle$	Did emission happen?	1
$ n_\phi\rangle$	Number of ϕ	$\lceil \log_2(N + n_I) \rceil$
$ n_a\rangle$	Number of f_a	$\lceil \log_2(N + n_I) \rceil$
$ n_b\rangle$	Number of f_b	$\lceil \log_2(N + n_I) \rceil$
	Ancillas	$4\lceil \log_2(N + n_I) \rceil + 5$

TABLE I: Registers in the QPS quantum circuit [1] along with the number of qubits required for N steps and n_I initial particles.

The QPS algorithm then proceeds by iteratively applying a series of steps. As described schematically above, each step has six components: (1) basis rotation $R^{(m)}$, (2) count particles U_{count} , (3) determine emission $U_e^{(m)}$, (4) update history U_h , (5) update particles $U_p^{(m)}$, and (6) inverse basis rotation $R^{(m)\dagger}$. Using the registers described in Table I, Alg. 1 gives a high-level description of the QPS algorithm in terms of these six meta gates. An alternative description in the form of a computational protocol is provided in Alg. 1. A detailed description of each register and each quantum gate is provided in App. A. Additionally, Fig. 2 illustrates two steps of the circuit diagram for Alg. 1. Note that emission history register $|h\rangle_m$ is updated and controlled on for U_p once at step m , but left alone during previous and subsequent steps. This property of the history register is a consequence of the lack of interference between different histories with the same emission angles. This motivates the idea of measuring $|h\rangle_m$ after each step, and then re-using those qubits during subsequent steps, which would reduce qubit requirements for the QPS.

D. Modified Quantum Parton Shower

Using mid-circuit operations conditioned on measurement results, we can significantly simplify the quantum circuit of the previous section. A qubit that is untouched for the remainder of a circuit can be measured at any time without changing (in principle) the properties of the encoded quantum state. In fact, a qubit used only to control operations on other qubits can be measured – and all subsequent quantum controls replaced by classical controls – without affecting the distribution of measured states. By invoking this deferred measurement principle on Fig. 2, the history sub-register $|h\rangle_m$ at step m can be

Algorithm 1: Original QPS algorithm [1]

Data: Splitting functions $P_{i \rightarrow j\phi}$, $P_{\phi \rightarrow ij}$, couplings g_1 , g_2 , g_{12} , step parameter ϵ , number of steps N , and n_I initial particles.

Result: Full amplitude description of final state radiation.

begin

Initialize all qubit registers in the $|0\rangle$ state.

Encode initial particles $|p\rangle_1 \dots |p\rangle_{n_I}$.

for $j \leftarrow 0$ **to** $N - 1$ **do**

(1) **Basis rotation:** Rotate all particles in $|p\rangle$ to the diagonal basis (Eq. (6)) using Eq. (A5).

(2) **Count particles:** Using U_{count} (App. A0f), count the number of each particle type, storing the results in $|n_a\rangle$, $|n_b\rangle$, $|n_\phi\rangle$.

(3) **Determine emission:** Using U_e (App. A0g), encode whether an emission occurred this step on $|e\rangle$, where the probability of emission is controlled on $|n_a\rangle$, $|n_b\rangle$, $|n_\phi\rangle$.

(4) **Update history:** Using U_h (App. A0h), update the history register $|h\rangle_m$, which encodes *which* particle (if any) emitted this step. The relative amplitudes for particular emissions are controlled by $|n_a\rangle$, $|n_b\rangle$, $|n_\phi\rangle$, $|p\rangle$, $|e\rangle$. (Note that $|e\rangle$ is put back into the $|0\rangle$ state implicitly in U_h .)

(5) **Update particles:** Using U_p (App. A0i), update the particle state, controlled on which particle emitted (encoded in $|h\rangle_m$).

(6) **Inverse basis rotation:** Rotate all particles in $|p\rangle$ back to the original basis, using the inverse of Eq. (A5).

Measure all qubits.

measured directly after applying $U_h^{(m)}$, and the results used to classically control $U_p^{(m)}$. Therefore, in principle, the qubits of $|h\rangle_m$ can be reset to the ground state $|0\rangle$ and re-used for $|h\rangle_{m+1}$. Because $|h\rangle_m$ encodes a superposition of *which* particle $|p_1\rangle, \dots, |p_{n_I+m-1}\rangle$, or None emitted at step m , measuring $|h\rangle_m$ projects the wavefunction onto a definite emission history. In other words, the sequence of emission locations (particles) is stored classically during the circuit execution. This does not affect the dynamics of the simulation, as different emission histories do not interfere with each other – fermion superpositions within the particle register $|p\rangle$ remain intact and only affect operations within the same history. Therefore, by running the simulation a polynomial number of times (given some constant statistical tolerance), we still construct an accurate probability distribution of final states. In this section, we describe at a high level how measuring and resetting $|h\rangle$ after each step simplifies the QPS algorithm. A detailed treatment of the improvements is provided in App. A.

Measuring the history register at a given step tells us *which* particle emits at that step, so the entire emission history is available as classical information during the simulation. In fact, given that the simulation begins with a definite type (fermion or boson) of initial particle, we can

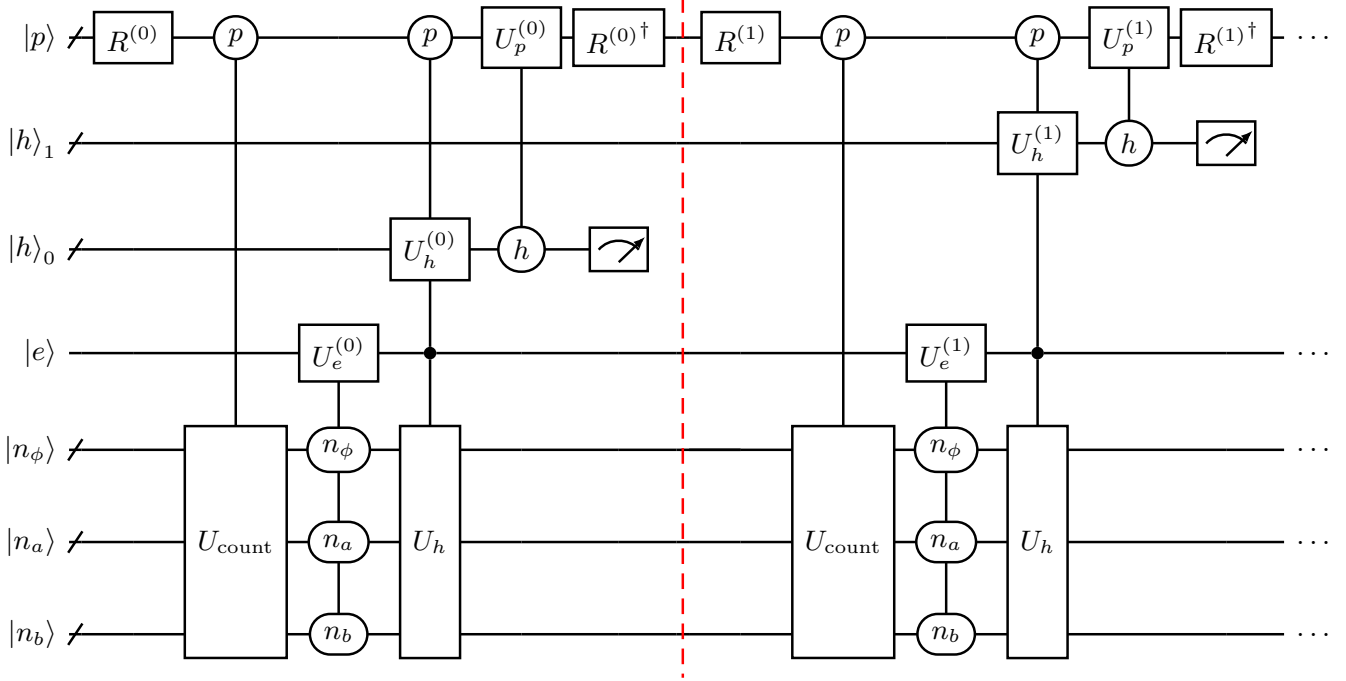


FIG. 2: High-level circuit diagram of the first two steps of the QPS algorithm. Round gates indicate control qubits.




FIG. 3: High-level circuit diagram of the first two steps of the improved QPS algorithm. Round gates indicate control qubits. Double wires indicate classical information stored on a CPU, and measured from $|h\rangle$. Double wires attached to U_p indicate that sub-gates in U_p are selectively applied based on measured values of h_0, h_1, \dots . The $|0\rangle$ attached to the same gate indicates a reset to the $|0\rangle$ state.

inductively infer from the emission history which *type* of particle emitted at each step. Note that while individual fermions may be in a superposition of different “flavors” a/b , for the two types of emissions considered – knowledge of the emitting particle type (f or ϕ) always implies the emitted particle type, as we consider the splittings $f \rightarrow f\phi$ and $\phi \rightarrow f\phi$. Therefore, the total number of particles (n_{tot}), number of bosons (n_ϕ), and number of fermions ($n_f = n_{f_a} + n_{f_b}$) are classical information. This lets us

remove two of the three counting registers $|n_\phi\rangle, |n_a\rangle, |n_b\rangle$, as follows.

In the original QPS, the U_e and U_h gates apply rotations with different angles, where each rotation (angle) is controlled on the corresponding possible value stored in $|n_a\rangle |n_b\rangle |n_\phi\rangle$. With mid-circuit measurements of $|h\rangle$, n_ϕ is classical information, so there is no need store and control on $|n_\phi\rangle$ at all. Now suppose we count the number of a -fermions – without loss of generality as we could instead

count the number of b -fermions – and apply rotations controlled on each possible value stored in $|n_a\rangle$. There is a one-to-one mapping of the possible values stored in $|n_a\rangle$ to the possible values of $|n_b\rangle$, given by $n_{\text{tot}} = n_a + n_b + n_\phi$, as n_{tot} and n_ϕ are stored on a CPU. Therefore, there is no need to store and control on $|n_b\rangle$, as the superposition of possible n_b 's is already implicitly encoded in $|n_a\rangle$. Thus, by measuring $|h\rangle$ at each step, $|n_b\rangle$ and $|n_\phi\rangle$ become redundant and can be removed.

Suppose we start with

$$\begin{aligned} n_{f,0} &= \text{initial number of fermions} = n_{f_a} + n_{f_b} \\ n_{\phi,0} &= \text{initial number of bosons} \\ n_{\text{tot},0} &= n_{f,0} + n_{\phi,0} \end{aligned} \quad (11)$$

Then, after measuring the history register, the CPU stores which of the initial particles emitted, so we adjust

$$\begin{aligned} n_{f,0} &\mapsto n_{f,1} \\ n_{\phi,0} &\mapsto n_{\phi,1} \\ n_{\text{tot},0} &\mapsto n_{\text{tot},1} \end{aligned} \quad (12)$$

in the CPU accordingly, and the emitted particle is encoded in $|p_1\rangle$. This process is repeated for each subsequent simulation step, so inductively $n_{f,i}, n_{\phi,i}, n_{\text{tot},i}$ are stored in the CPU throughout the simulation. Under the assumption that we can implement the workflow from Fig. 1, we can use this information to reduce both the computational complexity and absolute qubit/gate counts of QPS. Alg. 2 gives a high-level description of the improved QPS algorithm with mid-circuit measurements.

The second simplification is as follows. Storing ϕ or None in particle register is redundant, as given the emission history, the location of all ϕ and None particles in the simulation is stored in the CPU. Therefore, we encode particle m into a qubit sub-register $|p_m\rangle$ only if it is a fermion. To see that this can be done without significantly changing the QPS algorithm, we briefly comment on each gate that acts on $|p\rangle$: R , U_{count} , U_h , and U_p (see Fig. 3). The basis rotation R (Eq. (A5)) acts as the identity on ϕ 's, so simply not encoding ϕ 's does not affect the action of R . For each particle $|p_m\rangle$ in $|p\rangle$, the counting gate U_{count} applies an incrementer on $|n_a\rangle$ controlled on $|p_m = f_a\rangle$ (see App. B1). Therefore, not encoding ϕ 's on qubit does not affect the action of U_{count} . The U_h gate applies a two-level rotation to $|h\rangle$ for each particle in $|p\rangle$, where the rotation angle is controlled by particle type (see App. B3). Recall that by measuring the history register, the CPU stores whether each particle in the emission history is a f , ϕ or None. Therefore, for fermions, control on $|p\rangle$ as before, but for ϕ 's the corresponding rotation can simply be applied without controlling on $|p\rangle$. In the circuit diagram of U_h – Fig. 13 – this can be visualized by replacing quantum register $|p_i\rangle$ with classical wires if $|p_i\rangle = |\phi\rangle$. The gate sequence of Fig. 13 is unchanged, except that some rotations $U_h^{(m,i)}$ may no longer be controlled on $|p_i\rangle$, but applied directly to $|h\rangle$.

Algorithm 2: QPS with mid-circuit measurements

Data: Splitting functions $P_{i \rightarrow j\phi}$, $P_{\phi \rightarrow ij}$, couplings g_1, g_2, g_{12} , step parameter ϵ , number of steps N , and n_I initial particles.

Result: Full amplitude description of final state radiation.

begin

Initialize all qubit registers in the $|0\rangle$ state.

Encode initial particles $|p\rangle_1 \dots |p\rangle_{n_I}$.

for $j \leftarrow 0$ **to** $N - 1$ **do**

(1) Rotate all particles in $|p\rangle$ to the diagonal basis (Eq. (6)) using Eq. (A5).

(2) Using U_{count} (App. A0f), count the number of a -fermions, storing the result in $|n_a\rangle$.

(3) Using U_e (App. A0g), encode whether an emission occurred this step on $|e\rangle$, where the probability of emission is controlled on $|n_a\rangle$, and n_ϕ, n_{tot} from the CPU.

(4) Using U_h (App. A0h), update the history register $|h\rangle_m$, which encodes *which* particle (if any) emitted this step. The relative amplitudes for particular emissions are controlled by $|n_a\rangle$, $|e\rangle$, and n_ϕ, n_{tot} from the CPU. (Note that $|e\rangle$ is put back into the $|0\rangle$ state implicitly in U_h .)

(5) Measure the history register $|h\rangle$, storing the result on the CPU.

(6) If the measurement result indicates that a $\phi \rightarrow f\bar{f}$ emission occurred, apply U_p (Fig. 14) to update the particle state.

(7) Reset $|h\rangle$ to the $|0\rangle$ state.

(8) Rotate all particles in $|p\rangle$ back to the original basis, using the inverse of Eq. (A5).

Measure $|p\rangle$.

Finally, the particle update U_p must be applied only if a $\phi \rightarrow f\bar{f}$ emission occurs. If that is the case, then the circuit in Fig. 14 is applied to $|p\rangle$.

1. Qubit Costs

Figure. 3 illustrates the high-level circuit diagram for two steps of the improved QPS algorithm. Note the reduction in registers compared to the original QPS. The improved QPS circuit calls for four qubit registers, detailed in Table II below.

Register	Purpose	# qubits
$ p\rangle$	Particle state	$2(N + n_I)$
$ h\rangle$	Emission history	$\lceil \log_2(N + n_I) \rceil$
$ e\rangle$	Did emission happen?	1
$ n_a\rangle$	Number of f_a	$\lceil \log_2(N + n_I) \rceil$
	Ancillas	$2\lceil \log_2(N + n_I) \rceil + 1$

TABLE II: Registers in the improved QPS algorithm along with the number of qubits required for N steps and n_I initial particles.

Note that the particle state only calls for $2(N + n_I)$ qubits, compared to $3(N + n_I)$ originally (Table. I). By only encoding fermions in quantum registers, two qubits are sufficient to encode $f_a, f_b, \bar{f}_a, \bar{f}_b$. Also note that the number of required qubits varies between different circuit executions, as a simulation where more fermions are produced requires more quantum resources. In the worst case, $|p\rangle$ will still consist of $N + n_I$ sub-registers, while the actual number is n_f . Depending on parameters $g_1, g_2, g_{12}, \epsilon$, and the splitting functions $P_{i \rightarrow j\phi}, P_{\phi \rightarrow ij}$, the expectation for $\langle n_f \rangle$ may be significantly smaller than $N + n_I$. Therefore, the maximum number of qubits required for an N -step simulation is

$$\# \text{Qubits}_{\text{max}} = 2(n_I + N) + 4\lceil \log_2(n_I + N) \rceil + 2, \quad (13)$$

while the actual number is

$$\# \text{Qubits}_{\text{actual}} = 2\left[n_f + \lceil \log_2(n_I + N) \rceil + \lceil \log_2(n_f) \rceil + 1\right], \quad (14)$$

where the latter could be used if a variable number of qubits can be used for the circuit execution. The asymptotic qubit scaling is just $\mathcal{O}(N + n_I)$, compared to $\mathcal{O}(N \log_2(N + n_I))$ for the original QPS algorithm (Table I). This original qubit scaling is from storing the emission history at each step, which means using N subregisters $|h\rangle_m$ each with $\lceil \log_2(n_I + m + 1) \rceil \sim \mathcal{O}(\log_2(N + n_I))$ qubits. By measuring, resetting, and re-using $|h\rangle$ after each step, $|h\rangle$ just consists of $\lceil \log_2(N + n_I + 1) \rceil$ qubits, and the dominant scaling becomes the $\mathcal{O}(N + n_I)$ qubits of $|p\rangle$. The qubit scaling difference is illustrated in Fig. 4, which plots qubit count against N , with one starting particle, $n_I = 1$.

2. Gate Costs

We measure gate costs by writing operations in terms of the universal standard gate set consisting of two-qubit controlled not gates (CNOTs) and arbitrary single qubit gates $U(\theta, \phi, \lambda)$. Multi-controlled gates are decomposed into sequences of Toffoli (CCX) gates using a standard procedure that requires ancillary qubits equal to the number of controls minus one [8]. Then, Toffoli gates are decomposed into six two-qubit entangling gates [8]. As two-qubit entangling gates are far costlier to implement in real devices than single qubit gates, we quote gate counts in terms of CNOTs. Note that while we illustrate ‘‘classical controls’’ in our circuit diagrams (Fig. 14, Fig. 13) for quantum gates selected dynamically by the CPU, only the attached quantum gates are included in the gate count.

Table III summarizes the gate costs of each component of the improved QPS algorithm. The overall asymptotic scaling is

$$\mathcal{O}(N \cdot (N + n_I)^2 \cdot \log_2(N + n_I)^2), \quad (15)$$




FIG. 4: Qubit cost comparison between the original QPS and the improved version with mid-circuit measurements. This plot is simply an illustration of Eq. (13) and the sum over Table I with $n_I = 1$. For $n_I > 1$, the qubit count curves simply shift to the left, as n_I and N only appear together in the qubit scaling, as $N + n_I$.

Element	Step m cost	Total Scaling (N steps)
R	$2(n_I + m)$	$\mathcal{O}(N^2 + Nn_I)$
U_{count}	$13\ell(n_I + m)$	$\mathcal{O}(N(N + n_I) \cdot \log_2(n_I + N))$
U_e	$(n_I + m)(12\ell - 10)$	$\mathcal{O}(N(N + n_I) \cdot \log_2(n_I + N))$
U_h		$\mathcal{O}(N(N + n_I)^2 \cdot \log_2(n_I + N)^2)$
U_p	2	$2N$
Total		$\mathcal{O}(N(N + n_I)^2 \cdot \log_2(n_I + N)^2)$

TABLE III: Gate costs of the different circuit elements using re-measurement. $\ell \equiv \lceil \log_2(n_I + m + 1) \rceil$

which is a factor of $(N + n_I)^2$ more efficient than the original QPS gate scaling:

$$\mathcal{O}(N \cdot (N + n_I)^4 \cdot \log_2(N + n_I)^2). \quad (16)$$

This scaling improvement is due to the fact that at step m , $|n_a\rangle$ is a superposition of $n_I + m$ possible basis states, while $|n_a\rangle |n_b\rangle |n_\phi\rangle$ is a superposition of $(n_I + m)^3$ possible basis states. To implement U_h (see App. A), rotation gates controlled on $|n_a\rangle$ is applied to $|h\rangle$ for each possible value stored in $|n_a\rangle$. Therefore, in the original algorithm, U_h consisted of $\mathcal{O}((m + n_I)^3)$ controlled-rotations, while in the improved algorithm, only $\mathcal{O}(m + n_I)$ controlled rotations are applied.

Fig. 5 compares the actual gate counts of our improved QPS circuits with those of the original QPS circuits. The dashed line is the contribution from just the U_h gate, and Fig. 5 illustrates that the gate cost of U_h is dominant.




FIG. 5: Gate cost comparison: The dashed line represents the dominant contribution from U_h to the total gate count for QPS with mid-circuit measurements.

IV. NUMERICAL RESULTS

Using `Qiskit`'s matrix product state simulator [28, 29], we are able to simulate QPS with one initial particle ($n_I = 1$) up to several steps. For each simulation, we use a scale parameter $\theta = \theta_m$, defined by

$$\theta_m \equiv \epsilon^{m/N}, \quad \epsilon = 0.001, \quad (17)$$

splitting functions

$$P_{i \rightarrow i\phi}(\theta) = g_i^2 \hat{P}_f(\theta) = \frac{g_i^2 \log(\theta)}{4\pi} \quad (18)$$

$$P_{\phi \rightarrow i\bar{i}}(\theta) = g_i^2 \hat{P}_\phi(\theta) = \frac{g_i^2 \log(\theta)}{4\pi}, \quad (19)$$

couplings

$$g_1 = 2 \quad g_2 = 1 \quad g_{12} = 1, \quad (20)$$

and one initial f_1 (see Eq. (A2)),

$$|p\rangle_{\text{initial}} = |p_1\rangle = |100\rangle. \quad (21)$$

The couplings $g_1 = 2, g_2 = 1, g_{12} = 1$ are arbitrary, but chosen such that $g_a, g_b \neq 0$ and $g_a \neq g_b$ (Eq. (6)), in order to capture the full problem complexity. For simplicity, the couplings are also kept independent of step for simplicity (in reality, they would run with the scale). This means that the rotations are the same at each step. The numerical values of the diagonalized couplings are

$$g_a = \frac{3 + \sqrt{5}}{2} \approx 2.618 \quad (22)$$

$$g_b = \frac{3 - \sqrt{5}}{2} \approx 0.382, \quad (23)$$

and the rotation angle u is

$$u = \frac{\sqrt{5} - 1}{2\sqrt{5}} \approx 0.28 \text{ rad}. \quad (24)$$

With these parameters, the Sudakov factors, which give probabilities to have no emission from a given particle at step m , can be written as

$$\Delta_a(\theta_m) = \epsilon^{g_a^2/4\pi N} \quad (25)$$

$$\Delta_b(\theta_m) = \epsilon^{g_b^2/4\pi N} \quad (26)$$

$$\Delta_\phi(\theta_m) = \epsilon^{(g_a^2 + g_b^2)/4\pi N}. \quad (27)$$

Because the couplings are kept constant, these probabilities also remain constant at each step.

We run simulations with $g_{12} = 0$ in addition to $g_{12} = 1$. As explained in Sec. III A, when $g_{12} = 0$ the parton shower can be solved using a classical Markov Chain algorithm. Therefore, as a sanity check, we overlay analytical Markov Chain calculations, each with 10^9 shots, over simulation results with $g_{12} = 0$ in our plots.

Figures 6 and 7 present simulation results for $N = 2, 3, 4, 5$ steps and compare the outputs between the original QPS and QPS with mid-circuit measurements³. We have chosen two different observables for illustration.

First, Fig. 6 shows histograms of the total number of emissions (E). The main subplot illustrates the probability distributions of E for classical MCMC (black), original QPS (filled bars), and QPS with mid-circuit measurements (solid edges) simulations, with both $g_{12} = 0$ (blue) and $g_{12} = 1$ (red). The second subplot magnifies differences between the MCMC and $g_{12} = 0$ simulation distributions, which are due to statistical noise and exhibited the expected deviations. The third subplot magnifies differences between distributions obtained from original QPS and QPS with mid-circuit measurements, which are also within the expected statistical variations. With 10^5 shots per simulation, typical statistical errors are on the order of $\sigma \sim \sqrt{\frac{\text{Pr}(E)}{10^5}} \lesssim \sqrt{\frac{1}{10^5}} \approx 0.0032$. Error bars shown in the second and third subplots Fig. 6 are 1σ ranges for the difference distributions, and the simulation results exhibit deviations on the expected scale. In other words, the second subplot shows that quantum simulations with coupling turned off ($g_{12} = 0$) agree with the classical MCMC algorithm, as expected. Additionally, the two different versions of the quantum algorithm – original

³ We have stopped at 5 steps due to the simulation time. The present criteria for determining how many steps to use is that simulations with 10^5 shots have to take fewer than 3 hours running naively without any parallelization on a 8 GB RAM Mac. It would be possible to go a bit further with larger computing resources. For the remeasurement circuits, it took ~ 2.5 hours to achieve 10^5 shots for $g_{12} = 0$ and $g_{12} = 1$. We note that classical conditioning is not fully implemented in `Qiskit` (it is not possible to do arbitrary classical calculations), so we have to apply an exponential number of classically-conditioned gates.

and with mid-circuit measurements – agree with one another. Nevertheless, the classical ($g_{12} = 0$) and quantum ($g_{12} = 1$) algorithms yield fundamentally different results.

We briefly describe the qualitative features of Fig. 6. First, with the chosen parameters (Eqs. (17), (19) and (20)), it turns out that

$$\Pr(E = 0, g_{12} = 0) > \Pr(E = 0, g_{12} = 1). \quad (28)$$

Additionally, for $g_{12} = 1$, the probability of ϕ -emission is $1 - \epsilon^{7/4\pi N}$, compared to $1 - \epsilon^{5/4\pi N}$ for $g_{12} = 0$. Therefore, conditional on a ϕ particle being present in the system, ϕ -emissions occur more frequently when $g_{12} = 1$. This explains why the probability of having just one emission for $g_{12} = 1$ is so low compared to $g_{12} = 0$. Finally, the exact shape of E distributions depends on numerical parameter values, but the general shape exhibited in Fig. 6 – increasing density with increasing E up to a peak (which could be $E = N$), followed by a tail where density decreases as $E \rightarrow N$ – is expected to hold for all N and all parameter values. The probability of emission at a given step is higher when there are more particles in the system, which explains why having just one or two emissions is less probable compared to having several emissions. However at the tail end ($E \rightarrow N$) of this trend, the distribution decreases slightly, because combinatorially there are more histories with $E = N - 1$ than with $E = N$.

The second observable, Fig. 7, is the distribution of the “hardest” emission angle, which algorithmically corresponds to the first emission that occurred during the shower evolution. The emission probability decreases exponentially with $\log(\theta)$, or linearly with opening angle θ . Algorithmically, this is because the probability of first $f_i \rightarrow f_i\phi$ emission occurring at smaller angles (later steps) is just an exponential of the Sudakov factor (Eq. (27)). For $g_{12} = 0$,

$$\begin{aligned} \Pr(\text{First emission at step } m) &= (1 - \Delta_a(\theta_m)) \prod_{i=1}^{m-1} \Delta_a(\theta_m) \quad (29) \\ &= (\Delta_a)^{m-1} (1 - \Delta_a), \end{aligned}$$

and for $g_{12} = 1$,

$$\begin{aligned} \Pr(\text{First emission at step } m) &= u (\Delta_a)^{m-1} (1 - \Delta_a) \quad (30) \\ &+ (1 - u) (\Delta_b)^{m-1} (1 - \Delta_b). \end{aligned}$$

Fig. 7 shows the distribution of $\log_e(\theta_{\max})$, with probabilities displayed in the main subplot and differences displayed in the secondary subplots. For this observable there is again a demonstrated difference in results between the classical ($g_{12} = 0$) and quantum ($g_{12} = 1$) algorithms. Nevertheless, the third subplot illustrates that the two versions of QPS – original and with mid-circuit measurements – agree within expected statistical variations.

V. CONCLUSIONS AND OUTLOOK

In this paper, we have simplified the digital Quantum Parton Shower (QPS) algorithm presented in [1] by considering mid-circuit measurements and quantum gates that are dynamically selected based on these measurement results. The QPS is an iterative “quantum Markov Chain” algorithm with N steps, and by making a mid-circuit measurement on a subset of qubits at each step, subsequent multi-qubit controlled R_y rotations in the original QPS can be replaced by dynamically selected single-qubit R_y rotations. In this case, the number of required R_y rotation at each step is reduced by a factor of N^2 . Additionally, qubits measured mid-circuit can be reset to the initial $|0\rangle$ state and re-used during subsequent steps, which reduces the qubit costs significantly. The resulting Alg. 2 improves the quantum gate complexity from $\mathcal{O}(N \cdot (N + n_I)^4 \cdot \log(N + n_I)^2)$ to $\mathcal{O}(N \cdot (N + n_I)^2 \cdot \log(N + n_I)^2)$ (Fig. 5), and the qubit complexity from $\mathcal{O}((N + n_I) \log(N + n_I))$ to $\mathcal{O}(N + n_I)$ (Fig. 4), compared to the original algorithm (Alg. 1).

We implement our quantum circuits using `Qiskit` (where dynamical quantum operations is a relatively new feature), and present results for $N = 2, 3, 4, 5$ steps. We illustrate agreement between the original and improved versions of QPS, as well as agreement with classical MCMC simulations in the limit $g_{12} = 0$, where QPS can be efficiently computed classically. Errors are shown to be consistent with the expected statistical uncertainties in all cases.

More generally, we showed how adopting a hybrid quantum-classical computing platform can be used to make an originally quantum algorithm more efficient. Recent studies [27, 30] have demonstrated that dynamic/hybrid quantum computing is feasible, and even implemented shallow algorithm on current hardware. As qubit design continues to improve, we expect to be able to execute more complicated hybrid algorithms such as QPS on real devices, and eventually be able to compute classically inaccessible physical observables. Moreover, as dynamic quantum computing is an intrinsic component of quantum error correction, developing dynamic/hybrid computing platforms is likely necessary in order to realize fault-tolerant quantum computers. The improved QPS algorithm serves as an additional case for prioritizing development of dynamic computing platforms, as reduced qubit and gate complexities/costs raise the potential for realizing QPS to compute classically inaccessible physical quantities much sooner. It is likely that other digital quantum algorithms with similar features – Markov Chain, or iterative algorithms where interferences exist within but not between different histories – can benefit from employing a dynamic/hybrid structure, and we encourage algorithm developers to consider this approach.




FIG. 6: Probability vs. Number of emissions (E) for 2, 3, 4, and 5 step simulations. Error bars represent 1σ ranges, e.g., in each third subplot, the red error bars corresponds to the standard deviation of the difference distribution between simulation results obtained from original QPS and QPS with remeasurement. Classical MCMC data were obtained using 10^7 shots, so the statistical errors are suppressed by a factor of $1/10$, and are thus negligible. Error bars in each subplot represent statistical errors from $g_{12} = 0$ simulations.

VI. ACKNOWLEDGEMENTS

This work is supported by the U.S. Department of Energy, Office of Science under contract DE-AC02-

05CH11231. In particular, support comes from Quantum Information Science Enabled Discovery (QuantISED) for High Energy Physics (KA2401032) and the Office of Advanced Scientific Computing Research (ASCR) through




FIG. 7: Probability vs. $\log_e(\theta_{\max})$ for 2, 3, 4, and 5 step simulations. As in Fig. 6, error bars represent 1σ ranges for the difference distributions. Simulation data (red, blue) is normalized such that probabilities are equal to bar height times bar width. For $g_{12} = 0$ simulations (blue), the area of each bar is equal – up to statistical deviations – to the integral of the analytical curve (black) over the respective bin.

the Accelerated Research for Quantum Computing Program.

-
- [1] B. Nachman, D. Provasoli, W. A. de Jong, and C. W. Bauer, *Physical Review Letters* **126**, 10.1103/physrevlett.126.062001 (2021).
- [2] S. P. Jordan, K. S. M. Lee, and J. Preskill, *Science* **336**, 1130 (2012), arXiv:1111.3633 [quant-ph].
- [3] J. Preskill, *PoS LATTICE2018*, 024 (2018), arXiv:1811.10085 [hep-lat].
- [4] C. W. Bauer, M. Freytsis, and B. Nachman, *Phys. Rev. Lett.* **127**, 212001 (2021), arXiv:2102.05044 [hep-ph].
- [5] J. Preskill, *Quantum* **2**, 79 (2018).

- [6] M. S. ANIS *et al.*, [Qiskit: An open-source framework for quantum computing](#) (2021).
- [7] A. Peruzzo, J. McClean, P. Shadbolt, M.-H. Yung, X.-Q. Zhou, P. J. Love, A. Aspuru-Guzik, and J. L. O’Brien, *Nature Communications* **5**, [10.1038/ncomms5213](#) (2014).
- [8] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information: 10th Anniversary Edition* (Cambridge University Press, 2010).
- [9] D. Ristè, C. C. Bultink, K. W. Lehnert, and L. DiCarlo, *Physical Review Letters* **109**, [10.1103/physrevlett.109.240502](#) (2012).
- [10] Y. Salathé, P. Kurpiers, T. Karg, C. Lang, C. K. Andersen, A. Akin, S. Krinner, C. Eichler, and A. Wallraff, *Phys. Rev. Applied* **9**, [034011](#) (2018).
- [11] L. Xiang, Z. Zong, Z. Sun, Z. Zhan, Y. Fei, Z. Dong, C. Run, Z. Jia, P. Duan, J. Wu, Y. Yin, and G. Guo, *Phys. Rev. Applied* **14**, [014099](#) (2020).
- [12] L. Steffen, Y. Salathe, M. Oppliger, P. Kurpiers, M. Baur, C. Lang, C. Eichler, G. Puebla-Hellmann, A. Fedorov, and A. Wallraff, *Nature* **500**, [319](#) (2013).
- [13] M. D. Barrett, J. Chiaverini, T. Schaetz, J. Britton, W. M. Itano, J. D. Jost, E. Knill, C. Langer, D. Leibfried, R. Ozeri, and D. J. Wineland, *Nature* **429**, [737](#) (2004).
- [14] M. Riebe, H. Häffner, C. F. Roos, W. Hänsel, J. Benhelm, G. P. T. Lancaster, T. W. Körber, C. Becher, F. Schmidt-Kaler, D. F. V. James, and R. Blatt, *Nature* **429**, [734](#) (2004).
- [15] K. S. Chou, J. Z. Blumoff, C. S. Wang, P. C. Reinhold, C. J. Axline, Y. Y. Gao, L. Frunzio, M. H. Devoret, L. Jiang, and R. J. Schoelkopf, *Nature* **561**, [368](#) (2018).
- [16] A. Furusawa, J. L. Sørensen, S. L. Braunstein, C. A. Fuchs, H. J. Kimble, and E. S. Polzik, *Science* **282**, [706](#) (1998), <https://www.science.org/doi/pdf/10.1126/science.282.5389.706>.
- [17] D. Ristè, L. C. G. Govia, B. Donovan, S. D. Fallek, W. D. Kalfus, M. Brink, N. T. Bronn, and T. A. Ohki, *npj Quantum Information* **6**, [71](#) (2020).
- [18] N. Ofek, A. Petrenko, R. Heeres, P. Reinhold, Z. Leghtas, B. Vlastakis, Y. Liu, L. Frunzio, S. M. Girvin, L. Jiang, M. Mirrahimi, M. H. Devoret, and R. J. Schoelkopf, *Nature* **536**, [441](#) (2016).
- [19] P. Campagne-Ibarcq, E. Flurin, N. Roch, D. Darson, P. Morfin, M. Mirrahimi, M. H. Devoret, F. Mallet, and B. Huard, *Phys. Rev. X* **3**, [021008](#) (2013).
- [20] Z. K. Mineev, S. O. Mundhada, S. Shankar, P. Reinhold, R. Gutiérrez-Jáuregui, R. J. Schoelkopf, M. Mirrahimi, H. J. Carmichael, and M. H. Devoret, *Nature* **570**, [200](#) (2019).
- [21] T. B. Pittman, B. C. Jacobs, and J. D. Franson, *Phys. Rev. A* **71**, [052332](#) (2005).
- [22] J. Chiaverini, D. Leibfried, T. Schaetz, M. D. Barrett, R. B. Blakestad, J. Britton, W. M. Itano, J. D. Jost, E. Knill, C. Langer, R. Ozeri, and D. J. Wineland, *Nature* **432**, [602](#) (2004).
- [23] R. Vijay, C. Macklin, D. H. Slichter, S. J. Weber, K. W. Murch, R. Naik, A. N. Korotkov, and I. Siddiqi, *Nature* **490**, [77](#) (2012).
- [24] L. Hu, Y. Ma, W. Cai, X. Mu, Y. Xu, W. Wang, Y. Wu, H. Wang, Y. P. Song, C.-L. Zou, S. M. Girvin, L.-M. Duan, and L. Sun, *Nature Physics* **15**, [503](#) (2019).
- [25] C. K. Andersen, A. Remm, S. Lazar, S. Krinner, J. Heinssoo, J.-C. Besse, M. Gabureac, A. Wallraff, and C. Eichler, *npj Quantum Information* **5**, [69](#) (2019).
- [26] C. C. Bultink, T. E. O’Brien, R. Vollmer, N. Muthusubramanian, M. W. Beekman, M. A. Rol, X. Fu, B. Tarasinski, V. Ostroukh, B. Varbanov, A. Bruno, and L. DiCarlo, *Science Advances* **6**, [eaay3050](#) (2020), <https://www.science.org/doi/pdf/10.1126/sciadv.aay3050>.
- [27] A. D. Córcoles, M. Takita, K. Inoue, S. Lekuch, Z. K. Mineev, J. M. Chow, and J. M. Gambetta, *Physical Review Letters* **127**, [10.1103/physrevlett.127.100501](#) (2021).
- [28] U. Schollwöck, *Annals of Physics* **326**, [96–192](#) (2011).
- [29] G. Vidal, *Physical Review Letters* **91**, [10.1103/physrevlett.91.147902](#) (2003).
- [30] C. A. Ryan, B. R. Johnson, D. Ristè, B. Donovan, and T. A. Ohki, *Review of Scientific Instruments* **88**, [104703](#) (2017).

Appendix A: Details of the original Quantum Parton Shower (QPS) algorithm

This appendix provides details on the original QPS algorithm presented in [1].

a. Particle state $|p\rangle$

This register consists of $N + n_I$ 3-qubit sub-registers, one for each initial particle, and one for each emission step,

$$|p\rangle = |p_1\rangle^{\otimes 3} \otimes \cdots \otimes |p_{n_I+N}\rangle^{\otimes 3}. \quad (\text{A1})$$

We use three qubits to encode each particle, as there are six different types of particles ($f_1, f_2, \bar{f}_1, \bar{f}_2, \phi$, and None) in our model. The exact encoding is relevant, and we use the following,

$$|p_i\rangle = \begin{pmatrix} 000 \\ 001 \\ 010 \\ 011 \\ 100 \\ 101 \\ 110 \\ 111 \end{pmatrix} = \begin{pmatrix} \text{None} \\ \phi \\ f_1/f_a \\ f_2/f_b \\ \bar{f}_1/\bar{f}_a \\ \bar{f}_2/\bar{f}_b \end{pmatrix}. \quad (\text{A2})$$

Therefore, operations conditioned on whether a particle is a fermion are controlled by just the first qubit, and operations conditioned on whether a fermion is type- a or type- b are controlled by just the first and third qubits. Note that two computational basis states are extraneous.

b. Emission history $|h\rangle$

In original QPS, $|h\rangle$ encodes the location of emission at each step. In particular, at step m , $|h_m\rangle$ stores a binary number between 0 and $n_I + m$ that specifies which particle emitted at that step. The $|0\rangle$ state encodes that no emission occurred. Therefore, each sub-register $|h_m\rangle$ consists of $\lceil \log_2(n_I + m + 1) \rceil$ qubits, and in total $|h\rangle$ consists of

$$\sum_{m=0}^{N-1} (\lceil \log_2(n_I + m + 1) \rceil) \sim \mathcal{O}(N \log_2(n_I + N)) \quad (\text{A3})$$

qubits.

c. Emission $|e\rangle$

The emission register $|e\rangle$ stores a boolean that specifies whether an emission occurred at a given step. It is straightforward to uncompute $|e\rangle$ after each step, so just one qubit is sufficient to represent $|e\rangle$.

d. Number registers $|n_\phi\rangle, |n_a\rangle, |n_b\rangle$

The number registers are used to count the number of each particle type at each step. In particular, at step m each of $|n_\phi\rangle, |n_a\rangle, |n_b\rangle$ stores a binary number between 0 and $n_I + m$, the maximum possible number of particles at step m . Also note that the total

$$0 \leq n_\phi + n_a + n_b \leq n_I + m. \quad (\text{A4})$$

Each number register is uncomputed during each step and can be reused for subsequent steps. Therefore, for an N -step algorithm, each number register consists of $\lceil \log_2(n_I + N) \rceil$ qubits, for a total of $3\lceil \log_2(n_I + N) \rceil$ qubits between $|n_\phi\rangle, |n_a\rangle$, and $|n_b\rangle$.

Having set up the six quantum registers shown in Fig. 2, we now describe each gate in the circuit.

e. $R^{(m)}$ basis rotation

As described in Sec. III B, we rotate fermion states from the $1/2$ basis into the a/b basis by applying unitary U (Eq. (A5)). Given our particle state representation A2, rotating a single particle represented by three qubits entails applying the following unitary gate:

$$R = \begin{pmatrix} I & 0 & 0 & 0 \\ 0 & I & 0 & 0 \\ 0 & 0 & U & 0 \\ 0 & 0 & 0 & U \end{pmatrix}, \quad (\text{A5})$$

where I and U are 2×2 matrices. Therefore, to rotate to complete particle register $|p\rangle$ at step m , apply the product gate

$$R^{\otimes (n_I+m)}. \quad (\text{A6})$$

The gate in Eq. (A5) is just a controlled- U gate, where U is applied to the rightmost qubit controlled on the leftmost qubit, in the particle encoding (Eq. (A2)). Therefore, applying Eq. (A6) at the beginning and end of step m involves applying $(n_I + m)$ controlled- U gates, each of which can be decomposed into two CNOTs.

f. U_{count} particle counting

The counting gate maps the particle state $|p\rangle$ at step m to the number of each particle, which is stored in $|n_\phi\rangle, |n_a\rangle, |n_b\rangle$. Note that we count fermions and anti-fermions of the same type together, as this distinction does not affect emission probabilities. For each particle $|p_i\rangle$ in the particle state $|p\rangle$, we apply an increment gate $\boxed{+}$ controlled on particle type to each of $|n_\phi\rangle, |n_a\rangle, |n_b\rangle$, as illustrated in Fig. 8 The increment gate controlled on ϕ, f_a , and f_b implements respective transformations

$$|p_i\rangle |n_\phi\rangle \rightarrow |p_i = \phi\rangle |n_\phi + 1\rangle + |p_i \neq \phi\rangle |n_\phi\rangle \quad (\text{A7})$$

$$|p_i\rangle |n_a\rangle \rightarrow |p_i = f_a\rangle |n_a + 1\rangle + |p_i \neq f_a\rangle |n_a\rangle \quad (\text{A8})$$

$$|p_i\rangle |n_b\rangle \rightarrow |p_i = f_b\rangle |n_b + 1\rangle + |p_i \neq f_b\rangle |n_b\rangle, \quad (\text{A9})$$




FIG. 8: Complete illustration of the U_{count} gate.

and is illustrated in Fig. 9. In total, U_{count} consists of




FIG. 9: Increment gate controlled on $|p = \phi\rangle$. First, a multi-control gate that encodes whether $|p_i\rangle$ is a ϕ onto an ancilla $|c\rangle$ is applied. Then, a Ripple Adder [?] with $|c\rangle$ as an input carry, and an additional ancillary register initialized to $|0\rangle$ is applied. Therefore, $|n_\phi\rangle$ is mapped to $|n_\phi + 1\rangle$ if $c = 1$, i.e. if $|p_i\rangle$ is a ϕ . Increment gates controlled on $|p = f_a\rangle$ $|p = f_b\rangle$ are identical other than the control setting (see A2).

$3(n_I + m)$ controlled-incrementers, each of which requires $\mathcal{O}(\lceil \log_2(n_I + m + 1) \rceil)$ gates and $\lceil \log_2(n_I + m + 1) \rceil$ re-usable ancillas. Thus, the total gate complexity of U_{count} is $\mathcal{O}(\{m, n_I\} \cdot \lceil \log_2(n_I + m + 1) \rceil)$.

g. U_e emission update

At each step, the total probability that an emission occurs depends on the scale θ_m , splitting functions $P_f(\theta_m)$, $P_\phi(\theta_m)$, and numbers of each type of particle $|n_a\rangle$, $|n_b\rangle$, $|n_\phi\rangle$. Let

$$\Delta_a(\theta_m), \Delta_b(\theta_m), \Delta_\phi(\theta_m) \quad (\text{A10})$$

denote the probability of *no emission* from a particular f_a , f_b , or ϕ particle, respectively. (These are the Sudakov

factors defined by Eq. (4).) Then, the total probability of having *no emission* at step m is

$$\Pr(\text{No emission}) \equiv \Delta^{(m)} = \Delta_a(\theta_m)^{n_a} \Delta_b(\theta_m)^{n_b} \Delta_\phi(\theta_m)^{n_\phi}, \quad (\text{A11})$$

so that the probability of emission at step m is

$$\Pr(\text{Emit at step } m) = 1 - \Delta^{(m)}, \quad (\text{A12})$$

noting that $|n_a\rangle$, $|n_b\rangle$, $|n_\phi\rangle$ are quantum superpositions while θ_m , $P_f(\theta_m)$, $P_\phi(\theta_m)$ are classical parameters. The U_e gate puts qubit $|e\rangle$ in the $|0\rangle$ state with probability $\Delta^{(m)}$ and the $|1\rangle$ state with probability $1 - \Delta^{(m)}$. With $|e\rangle$ initially in the $|0\rangle$ state, this entails applying a rotation

$$U_e^{(m)} = \begin{pmatrix} \sqrt{\Delta^{(m)}} & -\sqrt{1 - \Delta^{(m)}} \\ \sqrt{1 - \Delta^{(m)}} & \sqrt{\Delta^{(m)}} \end{pmatrix}. \quad (\text{A13})$$

As $\Delta^{(m)}$ depends on the $|n_a\rangle$, $|n_b\rangle$, $|n_\phi\rangle$, this rotation must be controlled on all three count registers. As $\Delta^{(m)}$ takes a different value for each different combination of counts, we control on every possible combination, of which there are $\mathcal{O}(\{m, n_I\}^3)$. Each rotation is controlled on $3\lceil \log_2(n_I + m + 1) \rceil$ qubits. A standard method [8] decomposes multi-control gates into a sequence of Toffoli gates, which can each be implemented using six CNOTs. The total costs to implement a single $3\lceil \log_2(n_I + m + 1) \rceil$ -controlled rotation is

$$12 \cdot \left(3\lceil \log_2(n_I + m + 1) \rceil - 1 \right) \quad (\text{A14})$$

CNOTs and

$$3\lceil \log_2(n_I + m + 1) \rceil - 1 \quad (\text{A15})$$

Re-usable ancillary qubits.

h. U_h history update

The U_h gate encodes the probability of particular particles emitting. If there is no emission, then $|h\rangle \rightarrow |0\rangle$ with amplitude $\sqrt{\Delta^{(m)}}$. But if there is an emission, U_h maps $|h\rangle$ to a superposition of basis states that correspond to particle numbers 1 to $n_I + m$. Each state $|1\rangle$ through $|n_I + m\rangle$ gains amplitude equal to the square root probability that the corresponding particle emits. In particular, the total probability for a particular particle emitting is equal to the total probability of emission $1 - \Delta^{(m)}$ times the conditional emission probability for that particle type. Denote the relative emission probabilities

$$P_a(\theta_m) = 1 - \Delta_a(\theta_m) \quad (\text{A16})$$

$$P_b(\theta_m) = 1 - \Delta_b(\theta_m) \quad (\text{A17})$$

$$P_\phi(\theta_m) = 1 - \Delta_\phi(\theta_m). \quad (\text{A18})$$

Given that an emission occurred, the conditional probability that particle p is the emitter is

$$\begin{aligned} \Pr(p \text{ emits} \mid |e\rangle = 1) &= \frac{P_p(\theta_m)}{\sum_{p \in \{a, b, \phi\}} n_p P_p(\theta_m)} \\ &= \frac{P_p(\theta_m)}{\sum_{j=1}^{n_I+m} P_j(\theta_m)}. \end{aligned} \quad (\text{A19})$$

Therefore, U_h prepares the computational basis states of $|h\rangle$ with the following amplitude distribution:

$$\begin{aligned} |h\rangle &\mapsto \sqrt{\Delta^{(m)}} |0\rangle + \sqrt{1 - \Delta^{(m)}} \\ &\quad \times \sum_{j=1}^{n_I+m} |j\rangle \cdot \Pr(p_j \text{ emits} \mid |e\rangle = 1). \end{aligned} \quad (\text{A20})$$

Starting with $|h\rangle$ in the $|0\rangle$ state, this distribution is prepared by applying a series of two-level rotations from $|0\rangle$ to the other computational basis states of $|h\rangle$. Each rotation is controlled on $|e\rangle$, and applies the conditional emission probability of emission for each particle $|p_1\rangle \dots |p_{n_I+m}\rangle$. The rotation controlled on $|p_j\rangle$ is

$$U_h^{(m,k)} = \begin{pmatrix} \sqrt{1 - \frac{P_j(\theta_m)}{\sum_{k \geq j} P_k(\theta_m)}} & -\sqrt{\frac{P_j(\theta_m)}{\sum_{k \geq j} P_k(\theta_m)}} \\ \sqrt{\frac{P_j(\theta_m)}{\sum_{k \geq j} P_k(\theta_m)}} & \sqrt{1 - \frac{P_j(\theta_m)}{\sum_{k \geq j} P_k(\theta_m)}} \end{pmatrix}. \quad (\text{A21})$$

As the rotation angle depends on counts $|n_a\rangle, |n_b\rangle, |n_\phi\rangle$, rotations are controlled on the count registers, in addition to $|e\rangle$ and $|p\rangle$. Iterating through each particle $p_1 \dots p_{n_I+m}$, each of these rotations is followed by a decrement to $|n_a\rangle, |n_b\rangle$, or $|n_\phi\rangle$ controlled on that particle type. This means the relative emission probabilities given by Eq. (A19) are updated after each rotation. (Note that the denominator in the entries of Eq. (A21) is different for each rotation.) Figure 10 illustrates a single rotation-decrement iterate, and Fig. 11 illustrates the entire U_h gate.

The gate complexity of U_h is $\mathcal{O}(\{m, n_I\}^4 \log_2(m+n_I)^2)$, and $3\lceil \log_2(n_I+m+1) \rceil + 2$ re-usable ancilla qubits are required.

i. U_p particle update

The U_p gate updates the particle state $|p\rangle$ based on which particle emitted at a given step. At step m , $|p_m\rangle$ stores the newly radiated particle if any. For example, if particle i emits a ϕ at step m , then U_p takes

$$U_p |p_i = f_j\rangle |p_m = 0\rangle \rightarrow |f_j\rangle |\phi\rangle. \quad (\text{A22})$$

Keeping in mind that $|h\rangle$ and $|p\rangle$, then U_p is implemented by applying the gate in Fig. 12 controlled on $|h_m - j\rangle$ for each j from 1 to $n_I + m$. The single-qubit rotation U_r in

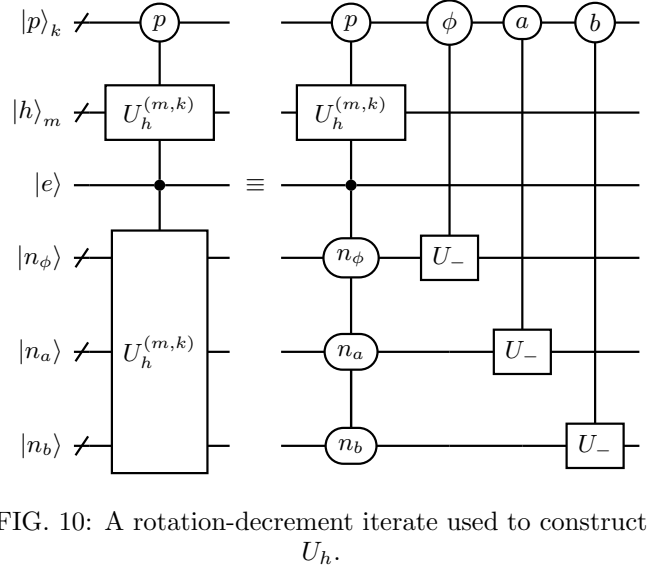


FIG. 10: A rotation-decrement iterate used to construct U_h .

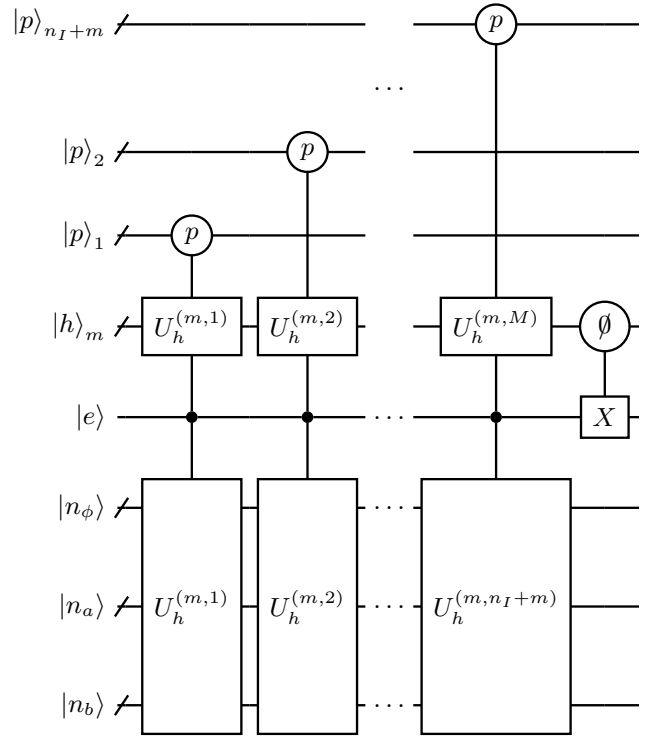


FIG. 11: Original U_h gate. Each sub-gate is defined by Fig. 10

Fig. 12 is the gate

$$U_r = \frac{1}{\sqrt{g_a^2 + g_b^2}} \begin{pmatrix} g_a & -g_b \\ g_b & g_a \end{pmatrix}. \quad (\text{A23})$$

In total the particle update step consists of $(n_I + m)$ applications of U_p controlled on $\lceil \log_2(n_I + m + 1) \rceil$ qubits. The U_p gate has a constant number of operations, so the overall gate complexity of the particle update is $\mathcal{O}(\{m, n_I\} \log_2(n_I + m + 1))$.




FIG. 12: U_p consists of applying this gate, controlled on $|h_m = j\rangle$, for each j from 1 to $n_I + m$.

Appendix B: Details of the QPS algorithm with mid-circuit measurements

1. U_{count} gate

We need only count the number of a single fermion type, so both gate and qubit counts are reduced by a factor of 3, compared to [1].

The total number of CNOTs required to implement U_{count} at step m is

$$\begin{aligned} & 13(n_{f,m})\lceil \log_2(n_{f,m} + 1) \rceil \\ & \leq 13(n_I + m)\lceil \log_2(n_I + m + 1) \rceil. \end{aligned} \quad (\text{B1})$$

2. U_e gate

Instead of conditioning on $|n_a\rangle$, $|n_b\rangle$, and $|n_\phi\rangle$, we only condition on $|n_a\rangle$. We compute Sudakov factors [1] for each possible value of $|n_a\rangle$, and apply the appropriate rotation matrices [1] conditioned on the value stored in $|n_a\rangle$. At the m th simulation step, $0 \leq n_a \leq n_I + m$, so there are at most $n_I + m$ conditional rotations, each conditioned on $\lceil \log_2(n_I + m + 1) \rceil$ qubits. In the original algorithm, there are $\mathcal{O}(m^3)$ rotations, which is the number of combinations of $|n_a\rangle$, $|n_b\rangle$, $|n_\phi\rangle$. Here n_ϕ and n_b are conditioned on classically, so the number of conditional rotations is reduced by a factor of m^2 . Thus, the computational complexity of the U_e gate is reduced to $\mathcal{O}(m \log_2^2[n_I + m + 1])$, compared to $\mathcal{O}(m^3 \log_2[n_I + m + 1])$ for the original [1].

3. U_h gate

Like the U_e gate, the rotations in this gate must be conditioned only the value stored in $|n_a\rangle$. Thus, the computational complexity of the U_h gate is reduced to $\mathcal{O}(m^2 \log_2^2[n_I + m + 1])$, compared to $\mathcal{O}(m^4 \log_2^2[n_I + m + 1])$ for the original [1]. The improved U_h gate is

illustrated in Fig. 13. The individual $U_h^{(m,k)}$ rotations are given by Eq. (A21).

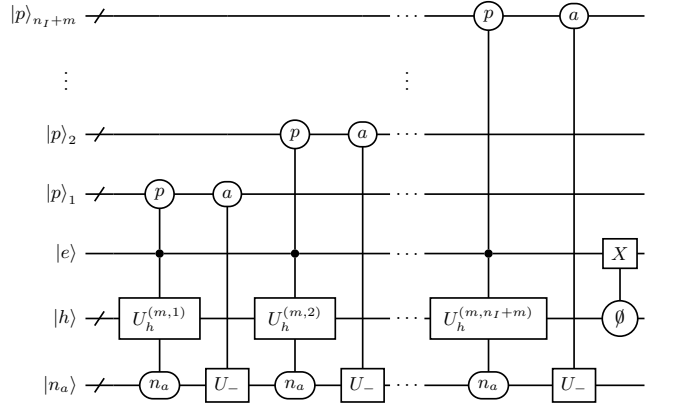


FIG. 13: The U_h gate at step m . Each $U_h^{(m,j)}$ gate denotes a sequence of two-level rotations (see App. A0h) controlled on the different states stored in $|n_a\rangle$ and $|p\rangle_j$. These controls are denoted by circular “gates” in the diagram. The relative probabilities for an emission from $|p\rangle_j$ depend on whether $|p\rangle_j$ is an f_a , f_b , or ϕ , as well as the number of each particle type in the system (which is fully encoded by $|n_a\rangle$ and previous measurements of $|h\rangle$). Therefore, we must apply $3\lceil \log_2(n_I + m + 1) \rceil$ different two-level rotations between $|h = 0\rangle$ and $|j\rangle$, each conditional on the values stored in $|n_a\rangle$ and $|p\rangle_j$.

Applying a controlled-decrement after each $U_h^{(m,j)}$ ensures that the correct relative emission probabilities are encoded into $|h\rangle$ (see Appendix A.5 in [1]), and also resets $|n_a\rangle$ to the $|0\rangle$ state. Finally, the last gate puts $|e\rangle$ back to the $|0\rangle$ state – after updating the history register, $|h\rangle \neq |0\rangle \iff |e\rangle = |1\rangle$, so we apply a NOT gate conditional on $|h\rangle \neq |0\rangle$.

4. U_p gate

Measuring the history register at step m collapses the wavefunction of the system such that a particular particle emitted or there was no emission. There is still quantum interference between the different fermion *types/flavors*, as measuring $|h\rangle$ does not affect any superpositions within particle states $|p\rangle$.

The selection of which U_p gate to apply is determined dynamically based on the measurement result. If the emitting particle is a ϕ , then apply the following particle update:

If the emitting particle is a fermion, then the particle update consists of entirely classical operations, as ϕ and None are not encoded on qubits. In our particle history table Table IV, we record that emitting particle j remains a fermion, and particle m is a ϕ . The computational cost of the U_p gate is reduced to a maximum constant of 2




FIG. 14: U_p gate to be applied if $p_j = \phi$. Note that prior to the emission, p_j is not encoded on qubits, so the particle update can be thought of as “initializing” new registers $|p_j\rangle$ and $|p_m\rangle$.

CNOTs, compared to $(n_I + m) \log_2 \lceil n_I + m + 1 \rceil$ for the original [1].

Step (m)	h_m	Particle $m + 1$
0	h_0	$\Rightarrow f/\phi/\text{None}$
1	h_1	$\Rightarrow f/\phi/\text{None}$
2	h_2	$\Rightarrow f/\phi/\text{None}$
3	h_3	$\Rightarrow f/\phi/\text{None}$
\vdots	\vdots	\vdots

TABLE IV: Information stored in a CPU. At each step (m) the history measurement determines whether particle $m + 1$ will be a fermion, boson, or None. This “particle history table” is filled out row by row as QPS is iterated through, and quantum gate are dynamically selected based on the values stored here.