

Quantum time-marching algorithms for solving linear transport problems including boundary conditions

Sergio Bengoechea^{1,*}, Paul Over¹ and Thomas Rung¹

¹*Institute for Fluid Dynamics and Ship Theory, Hamburg University of Technology, Hamburg D-21073, Germany.*

(Dated: November 7, 2025)

This article presents the first complete application of a quantum time-marching algorithm for simulating multidimensional linear transport phenomena with arbitrary boundaries, whereby the success probabilities are problem intrinsic. The method adapts the linear combination of unitaries algorithm to block encode the diffusive dynamics, while arbitrary boundary conditions are enforced by the method of images only at the cost of one additional qubit per spatial dimension. As an alternative to the non-periodic reflection, the direct encoding of Neumann conditions by the unitary decomposition of the discrete time-marching operator is proposed. All presented algorithms indicate optimal success probabilities while maintaining linear time complexity, thereby securing the practical applicability of the quantum algorithm on fault-tolerant quantum computers. The proposed time-marching method is demonstrated through state-vector simulations of the heat equation in combination with Neumann, Dirichlet, and mixed boundary conditions.

Keywords: Quantum Time Marching, Boundary Conditions, Block Encoding, Success Probabilities

I. INTRODUCTION

Thermofluid dynamic phenomena appear in a large range of applications in industry and academia, covering heat exchangers in power plants [1], aerodynamic design in aerospace engineering [2], cooling systems in electronics [3], or atmospheric and oceanic circulation modeling [4]. In general, such problems are solved numerically, since experimental investigations usually suffer from scaling effects and access to expensive testing facilities is often limited. However, an in-depth analysis requires extensive high-performance computing efforts, leading to greater demand for computational power as the spatio-temporal resolution increases. Alongside the challenges of further reducing the transistor size [5], classical Central Processor Unit (CPU) technology is approaching its capacity limits to provide the scale resolution required for these applications. Likewise, comparable scaling constraints apply to general-purpose computing on Graphics Processing Units (GPUs) [6]. With the objective of resolving smaller length scales [7], Quantum Computers (QCs) offer a promising alternative, comprising greater computational resources and more energy-efficient hardware in comparison to the classical counterparts [8, 9]. The improvements brought about by QCs are based on an exponentially large vector space for computation and the ability to apply operations simultaneously across the entire quantum register [10]. Nonetheless, the development of QCs confronts significant challenges to establish quantum advantage, including restricted gate depth, vulnerability to noise, state preparation, or computation readout, to name a few [11].

In the current era of Noisy-Intermediate-Scale Quantum (NISQ) devices [11], Variational Quantum Algorithms (VQAs) have emerged [12] as a viable hybrid classical-quantum strategy. VQAs mitigate the I/O problem of QCs by encoding the solution in the parameters of an ansatz and updating them by measuring a single ancilla qubit [13]. Despite these advantageous properties, VQA methods rely on solving non-convex optimization problems, limiting the scalability of such methods and preventing –for now– industrial applications for Computational Fluid Dynamics (CFD) [14, 15].

To extend the applicability of QCs to engineering problems, the advent of fault-tolerant machines opens the door for the adaptation of classical time-marching algorithms into the quantum framework [16, 17]. To this end, Hamiltonian simulations, used to solve the Schrödinger equation [10], are applied to linear [18–21] and nonlinear Partial Differential Equations (PDEs) [22]. The advantage of this approach is to recover the PDE time evolution using the Schrödinger equation, as it is the governing model on a QC. For example, linear transport problems described by the wave equation can be solved under internal and external Dirichlet and Neumann boundary conditions as proposed by Costa *et al.* [21].

The change in methodology increases the requirements for the probabilistic design of the quantum algorithm, making high success probabilities essential to guarantee its practical feasibility. This aspect has proven to be more critical than nonlinearity implementations, as these have already been addressed either directly [13, 22, 23] or via Carleman linearization [24]. The main challenge lies in modeling dissipation, which breaks the unitariness of quantum operations due to its irreversible nature and thereby penalizes the success probabilities. Consequently, any sequential application of non-unitary operations leads to an exponential decay in the cumulative success probability [25], making the direct implementation on a QC impractical without additional techniques [26], such as uniform singular value amplification [27] at

* Corresponding author: sergio.bengoechea@tuhh.de

each time step to bound the success probabilities. However, these measures introduce an unfavorable gate scaling, making the method, at best, quadratic in the simulation time. An example of an approach with nearly certain success probabilities was proposed by Brearley and Laizet [18], where the discrete time-marching operator of the advection equation was block encoded [28] in the Hamiltonian simulation to permit central and upwind-biased discretization schemes, as well as boundaries. The related work of Over *et al.* [19] suggests an extension including a diffusive term to recover the advection-diffusion equation while maintaining optimal success probabilities with linear time complexity. Another methodology is presented by An *et al.* [20], where a linear combination of Hamiltonian simulations is used to approximate the non-unitary dynamics described by general Ordinary Differential Equations (ODEs). The works of Jin *et al.* [29–32] and Lu and Yang [33] employ a different technique, called Schrödingerization, adapting a discrete ODE into a dilated system of Schrödinger equations, that is subsequently propagated in time via a Hamiltonian simulation on a QC.

The present publication introduces a novel alternative for treating non-unitary diffusive effects, maintaining optimal success probabilities, without the need for additional techniques [26, 27]. The work particularly focuses on the applicability and feasibility of the quantum time-marching algorithm proposed by Over *et al.* [19] to deal with engineering boundary conditions when solving time-dependent linear transport phenomena in multiple dimensions, such as heat conduction. The strategy applies the Linear Combination of Unitaries (LCU) algorithm for the operator of the governing second-order linear PDE, which is unitarily decomposed for diffusion problems. To incorporate arbitrary boundary conditions, the method of images associates the boundary type to the symmetry of the solution over the non-periodic dimension [34]. Using this domain decomposition technique, the boundaries can be reconstructed from adjacent interior points. An alternative algorithm is also provided for diffusion problems with Neumann conditions. In this algorithm, the bounded time-marching operator is decomposed into unitaries that are directly recombined using the LCU method, eliminating the need for spatial reflection. Both proposed methods to realize boundary conditions retain optimal success probabilities while maintaining linear time complexity. This property of the algorithms is of paramount importance for their applicability in CFD simulations, as satisfying the stability criteria for explicitly solving PDEs is associated with small time steps and, therefore, a large number of them.

The paper is organized as follows. The mathematical model is described in Sec. II. Section III provides a brief review of the quantum methods required to implement discrete operations, such as those presented in Sec. II on a fault-tolerant QC, along with the corresponding quantum circuits. Two-dimensional state-vector simulations of the heat equation with all types of boundary condi-

tions are benchmarked against classical Finite Differences (FD) solutions in Sec. V. Conclusions and future directions are summarized in Sec. VI. Additionally, the Appendix A illustrates a matrix-circuit analogy to facilitate the understanding of the quantum circuits, with particular emphasis on the LCU algorithm.

Throughout the publication, continuous vectors are indicated as $\vec{\square}$, discrete vectors in **bold**, quantum states in the Dirac (or bra-ket) notation, while matrices and gates are marked with capital letters. Furthermore, the *big-endian* convention applies to all quantum circuits and registers.

II. MATHEMATICAL MODEL

Consider a two-dimensional ($d = 2$) domain, with spatial coordinates $\vec{x} \in [0, 1] \times [0, 1]$ [m], the time $t \in [0, T]$ [s], a generic transport scalar ϕ [-], a (constant) dynamic viscosity Γ [m²/s] and the dynamics given by the unsteady diffusion equation, e.g., the heat conduction, viz.

$$\frac{\partial \phi}{\partial t} - \Gamma \nabla^2 \phi = 0. \quad (1)$$

This parabolic equation is discretized using an isotropic and equidistant spatial Δx [m] and temporal Δt [s] spacing with $N = N_{x_1}^d$ and N_t supporting points, respectively. For the sake of clarity, the discretization of Eqn. (1) is demonstrated for a one-dimensional ($d = 1$) problem with $N_{x_1} = \sqrt[d]{N}$ supporting points. The differential operators are approximated with a first-order accurate explicit forward Euler method in time and second-order central differences for the spatial derivative. Moreover, periodic boundary conditions are assumed for simplicity. Accordingly, the discrete form of Eqn. (1) is

$$\frac{\phi_j^{t+1} - \phi_j^t}{\Delta t} - \Gamma \frac{\phi_{j+1}^t - 2\phi_j^t + \phi_{j-1}^t}{\Delta x^2} = 0. \quad (2)$$

Here, the Toeplitz matrix $A \in \mathbb{R}^{N_{x_1} \times N_{x_1}}$, viz.

$$A = \begin{bmatrix} 1-2r_h & r_h & & 0 & r_h \\ r_h & 1-2r_h & \ddots & & 0 \\ & \ddots & \ddots & \ddots & \\ 0 & & \ddots & 1-2r_h & r_h \\ r_h & 0 & & r_h & 1-2r_h \end{bmatrix}, \quad (3)$$

is the time-marching operator in $\phi(t + \Delta t) = A\phi(t)$, where the non-dimensional parameter $r_h = \Gamma \Delta t / \Delta x^2$ [-] characterizes the spatio-temporal behavior of the diffusive transport. Following the time-marching idea, Eqn. (1) is solved sequentially in time, with the time step only being restricted by the stability criteria determined from the magnitude of r_h . This defines a convergence condition of the numerical method applied to the ODE,

requiring values below one in all supporting points for the duration of the entire solution process.


Prior to its application on a QC, the matrix A is subjected to modifications. First, the main diagonal is rescaled to one by $A/(1 - 2r_h)$. Second, this scaled matrix is decomposed w.r.t. its non-zero diagonals as

$$\frac{A}{1 - 2r_h} = I + \frac{r_h}{1 - 2r_h} (S_0 + S_0^\dagger), \quad (4)$$

using the unitaries identity I and the shift operator

$$S_0 = \begin{bmatrix} 0 & 1 & 0 & \cdots & 0 \\ & 0 & 1 & & \vdots \\ \vdots & & \ddots & \ddots & 0 \\ 0 & & & 0 & 1 \\ 1 & 0 & \cdots & & 0 \end{bmatrix}. \quad (5)$$

In Eqn. (4), $S_0^\dagger$ indicates the conjugate transpose of S_0 , and both of them can be efficiently implemented with the sequence of quantum gates depicted in Fig. 1.



(a) Quantum circuit for S_0 . (b) Quantum circuit for $S_0^\dagger$.

FIG. 1: Three-qubit example for the implementation of the shift operators S_0 and $S_0^\dagger$.

III. QUANTUM METHODS AND ALGORITHMS

This section introduces the fundamental methodologies used to implement the discrete operators of Sec. II on the QC. These are amplitude encoding (Sec. III A), Hamiltonian simulation (Sec. III B), block encoding (Sec. III C), and the LCU algorithm (Sec. III D). Subsequently, the algorithmic framework is described, with the focal point being the treatment of boundary conditions, cf. Sec. III E.

A. Amplitude encoding

Unlike classical computation, where information is represented by binary strings in the computational basis, arbitrary variable functions can be encoded in the amplitudes $\psi_j \in \mathbb{C}$ of a quantum state $|\psi\rangle \in \mathbb{C}^N$. Here, the Dirac (or bra-ket) notation, standard in quantum mechanics, is used. In this notation, a normalized column vector ψ is denoted as $|\psi\rangle$ and its conjugate transpose

reads $|\psi\rangle^\dagger = (|\psi\rangle^\top)^* = \langle\psi|$ [35]. Note that the bra-ket notation ($\langle\cdot|$ or $|\cdot\rangle$) always indicates a normalized quantum state. For a quantum register consisting of $n = \log_2(N)$ qubits, the state vector of size N is thus given as

$$|\psi\rangle = \sum_{j=0}^{N-1} \psi_j |e_j\rangle, \quad (6)$$

where $|e_j\rangle$ form an orthonormal basis of a Hilbert space and ψ_j^2 represents the probability of measuring the j -th basis state $|e_j\rangle$ conserving $\sum_j \psi_j^2 = 1$. Operations such as the (inner) scalar and (outer) dyadic product can be compactly expressed using the bra-ket notation as $\langle\cdot|\cdot\rangle$ and $|\cdot\rangle\langle\cdot|$, respectively. In the special case of concatenating qubit states represented as two-column vectors or two-row vectors, the Kronecker product $\otimes$ exemplarily given for an $n = 2$ qubit system is

$$\begin{aligned} \langle\rho\psi| &= \langle\rho| \otimes \langle\psi| \\ &= [\rho_1\psi_1, \rho_1\psi_2, \rho_2\psi_1, \rho_2\psi_2] \\ &= (|\rho\rangle \otimes |\psi\rangle)^\dagger = |\rho\psi\rangle^\dagger. \end{aligned} \quad (7)$$

In general, the state vector is prepared either by initializing the amplitudes ψ_j through a learning approach [14, 15], by employing divide-and-conquer strategies [22], by programming through Tensor Trains (TTs) [36–38], or in special cases by using analytical approaches [39].

B. Hamiltonian simulation

The Hamiltonian dynamics of a wave function $\psi(\vec{x}, t)$ are described by the time-dependent Schrödinger equation,

$$i \frac{d\psi(\vec{x}, t)}{dt} = \hat{H}\psi(\vec{x}, t), \quad \psi(\vec{x}, 0) = \psi^0. \quad (8)$$

This ODE represents a linear and dissipation-free (reversible) system, with the Hamiltonian $\hat{H}$ being Hermitian, $\hat{H} = \hat{H}^\dagger$, and $i = \sqrt{-1}$ denoting the imaginary unit. For a time-independent $\hat{H}$, the solution to Eqn. (8) is given by $\psi(\vec{x}, t) = \psi^0(\vec{x}) \exp(-i\hat{H}t)$ [10, 35]. Given the fact that the operator $\hat{U} = \exp(-i\hat{H}t)$ is unitary, a direct implementation of an equivalent quantum circuit exists, providing a deterministic implementation on a QC. The discretized Hamiltonian operator $H \in \mathbb{C}^{N \times N}$ is applied to the spatially discretized initial wave function $\psi^0 \in \mathbb{C}^N$, encoded by the initial quantum state $|\psi^0\rangle$, to approximate its temporal evolution by $|\psi^t\rangle = U|\psi^0\rangle$ [40], as it is depicted in Fig. 2.

There exist multiple algorithms based on this quantum state evolution [41], such as product formulas to expand a unitary evolution (e.g., Suzuki-Trotter [42]), Taylor series [43], quantum signal processing [44], or the Schrödingerization process [29–31] in which the underlying PDE is cast into a dilated system of Schrödinger equations (8) to be solved using Hamiltonian simulations. A




FIG. 2: Generic circuit for the temporal evolution of $|\psi^0\rangle$ given by the Hamiltonian simulation of Eqn. (8).

specific application of the Hamiltonian simulation is described in [18, 21, 45] where this technique is employed to block encode non-unitary operators, cf. Sec. III C.

C. Block encoding

Applying FD approximations to PDEs results in discrete operators M , e.g., matrix A in Sec. II. These matrices are generally non-unitary, so that $M^\dagger M \neq I$ and $MM^\dagger \neq I$, which poses a challenge when implementing them on a QC. Since quantum dynamics must be expressed through unitary operations to preserve norm and probability amplitudes, a common approach is to enlarge the underlying vector space and the operator matrix for a unitary embedding, cf. [28], according to

$$U_M = \begin{bmatrix} M/\alpha & * \\ * & * \end{bmatrix}, \quad (9)$$

where the sub-blocks indicated with $*$ are selected differently. Note that the unitary matrix U_M exists only if the singular values (σ_r) of any submatrix block are bounded by 1, i.e., its spectral norm $\|M\|_2 = \sigma_{\max}(M) \leq 1$ [28], where $\sigma_{\max}$ is the largest singular value. This implies that the l2-norm of each column of M must not exceed 1. To satisfy this condition, M is pre-scaled by a normalization factor $\alpha \in \mathbb{R}^+$ such that $\alpha \geq \|M\|_2$. Consequently, for the non-unitary matrix $M \in \mathbb{C}^{2^n \times 2^n}$ defined by n qubits, an accurate block encoding is achieved primarily by an optimal choice of the subnormalization factor, i.e., $\alpha = \|M\|_2$. For the sake of clarity, $\tilde{M}$ is assumed to be the optimal scaling of M , i.e., $\tilde{M} = M/\alpha$. In general, the scaling increases the condition number of $\tilde{M}$ [46], making block-encoding strategies for quantum linear system solvers not practical. In contrast, time-marching approaches do not encounter this problem due to the absence of matrix inversion.

To complete the block encoding, it remains to select the $*$ entries in Eqn. (9). This choice is not unique. A general method proposed by Camps *et al.* [28] uses non-parameterized sub-blocks for U_M as

$$U_M = \begin{bmatrix} \tilde{M} & \sqrt{I - \tilde{M}^\dagger \tilde{M}} \\ \sqrt{I - \tilde{M} \tilde{M}^\dagger} & -\tilde{M} \end{bmatrix}. \quad (10)$$

An alternative to compute the square root of $\tilde{M}^\dagger \tilde{M}$ in Eqn. (10) is also provided in Ref. [28], where the appearance of this term is avoided in the sub-blocks of U_M . This approach is tailor-made for certain dynamics that

recover Toeplitz matrices with constant non-zero diagonals upon discretization, and it requires more than one ancilla qubit.

A numerically motivated block encoding has been introduced by Lin *et al.* [47], which determines the sub-blocks of the structured matrix

$$U_M = \begin{bmatrix} sM & G \\ F & W \end{bmatrix}, \quad (11)$$

imposing $F^\dagger F = I - s^2 M^\dagger M$ and $FF^\dagger = I - s^2 MM^\dagger$ for U_M to be unitary with $s = 1/\sigma_{\max}(M)$. The remaining sub-blocks, G and W , are numerically constructed using QR decomposition on randomly generated matrices of the same size. Meanwhile, the first 2^N columns of U_M remain unaffected by the QR -based operation since these are orthogonal.

A different block encoding for time-marching operators derived from pure advective transport is based on the Hamiltonian simulation, cf. Sec. III B. This procedure follows Ref. [18] which encodes the non-unitary matrix $\tilde{M}$ as

$$\hat{U}_M = \begin{bmatrix} * & * \\ \tilde{M} & * \end{bmatrix} = \exp(-iH\theta), \quad (12)$$

with $H = \begin{bmatrix} 0 & -i\tilde{M}^\dagger \\ i\tilde{M} & 0 \end{bmatrix}$.

Brearley and Laizet [18] propose to transform Eq. (12) by applying trigonometric identities [45] as follows

$$\hat{U}_M = \begin{bmatrix} \cos(\sqrt{\tilde{M}^\dagger \tilde{M}} \theta) & -\tilde{M}^\dagger \frac{\sin(\sqrt{\tilde{M} \tilde{M}^\dagger} \theta)}{\sqrt{\tilde{M} \tilde{M}^\dagger}} \\ \tilde{M} \frac{\sin(\sqrt{\tilde{M}^\dagger \tilde{M}} \theta)}{\sqrt{\tilde{M}^\dagger \tilde{M}}} & \cos(\sqrt{\tilde{M} \tilde{M}^\dagger} \theta) \end{bmatrix}. \quad (13)$$

In Eqns. (12) and (13), the control parameter $\theta \in \mathbb{R}$ denotes a Hamiltonian evolution time, defining both the accuracy of the block encoding and its probability of success. Equation (13) is used to temporarily evolve the advection dynamics by applying the unitary operator $\hat{U}_M$ in [18]. The application of $\hat{U}_M$ to the prepared quantum state $|\psi^t\rangle$, indicated at the bottom of Fig. 3, implements

$$\hat{U}_M |0\rangle |\psi^t\rangle = \begin{bmatrix} \tilde{M} |\psi^t\rangle \\ I |\psi^t\rangle \end{bmatrix} \approx \begin{bmatrix} |\psi^{t+1}\rangle \\ |\psi^t\rangle \end{bmatrix}. \quad (14)$$

In the post-selection step, either the updated state $|\psi^{t+1}\rangle$ or the corresponding initial state $|\psi^t\rangle$ is measured in the computational basis $|0\rangle$ or $|1\rangle$, respectively. This choice is governed by the success probability per time step p^t , which for $\theta = 0$ results in zero and for $\theta = \pi/2$, the optimal value of $p^t \approx 1$ is reached.

Mind that U_M and $\hat{U}_M$ differ in structure with the block encoding of $\tilde{M}$ located in the upper-left and lower-left sub-blocks, respectively. This distinction must be considered for the circuits' implementation and the post-selection output as it is depicted in Fig. 3.

For all block encoding strategies, the choice of $*$ determines the success probabilities associated with it, as




FIG. 3: Exemplary quantum circuit implementations of the block encoding strategies for U_M (top) and $\hat{U}_M$ (bottom).

the method executes a non-deterministic operation on the QC, and introduces a penalty to the overall success probability. For example, the application of U_M to a quantum register initialized in the state $|0\rangle \otimes |\psi(\vec{x}, t)\rangle$ evolves the system to

$$|0\rangle (M/\alpha |\psi(\vec{x}, t)\rangle + * \mathbf{0}) + |1\rangle (* |0\rangle |\psi(\vec{x}, t)\rangle). \quad (15)$$

Here, the term $* \mathbf{0}$ indicates the action of the upper-right sub-block in Eqn. (9) on the lower half of the register, while $|1\rangle (* |0\rangle |\psi(\vec{x}, t)\rangle)$ corresponds to the lower row in Eqn. (9). This transformation is accurately implemented up to the subnormalization factor $\alpha = \|M\|_2$, conditioned on measuring $|0\rangle$ on the additional (ancilla) qubit. Thus, the block encoding prepares the target state $M/\alpha |\psi\rangle$ with a success probability of

$$p = (1/\alpha^2) \langle \psi | M^\dagger M | \psi \rangle = 1/\alpha^2 \|M |\psi\rangle\|_2^2 \leq 1/\alpha^2 \|M\|_2^2 \| |\psi\rangle \|_2^2 = 1/\alpha^2 \|M\|_2^2. \quad (16)$$

For $\alpha > 1$, the cumulative success probability of any time-marching approach successively decreases with each temporal evolution step. The success probabilities rapidly converge to zero, making the algorithm impractical except for $\alpha = 1$ [19]. In general, any matrix can be block encoded. However, for nearly unitary matrices when $\alpha \approx 1$, the success probability approaches one. This is the case for the block encoding strategy suggested by Brearley and Laizet [18], where the parameter θ can be selected to obtain a success probability near unity. In probabilistic terms, the strategy proposed in Ref. [18], cf. Eqns. (12) and (13), has clear advantages over the block encoding strategies in Eqns. (10) and (11).

Fang *et al.* [26] proposed using Quantum Singular Value Transformation (QSVT) to correct vanishing probabilities, thereby restoring practical feasibility. However, the proposed modification comes with an increase from linear to quadratic complexity in time. In the next sections, further alternatives to QSVT, maintaining $\alpha = 1$ are exploited while preserving the linear time complexity.

D. Linear combination of unitaries

The Linear Combination of Unitaries (LCU) algorithm combines a weighted sum of m unitary operators U_k on a QC [48]. These unitary operators are not necessarily obtained through the block encoding strategies applied to U_M in Sec. III C. With this approach, a decomposition into unitaries of a non-unitary matrix

$$M = \sum_{k=0}^{m-1} \kappa_k U_k, \quad (17)$$

with $\kappa_k \in \mathbb{R}^+ \cup \{0\}$, can be implemented with the circuits of Fig. 4, or Fig. 5 for the example of $m = 4$, and $m = 8$ unitaries, respectively. These quantum circuits prepare the final state $|\phi^{t+1}\rangle$, using $n_{\text{anc.}} = \log_2(m)$ ancilla qubits, conditioned to measure $|0\rangle^{\otimes n_{\text{anc.}}}$, viz.

$$\langle 0 | \underbrace{V^\dagger U_c V}_{U_M} | 0 \rangle |\phi^t\rangle = \frac{M}{\alpha} |\phi^t\rangle. \quad (18)$$

Comparing Eqn. (18) with Eqn. (9) from Sec. III C, the LCU can be interpreted as a block encoding strategy where α is the scaling factor and $1/\alpha^2$ is the penalty in the success probability, as indicated in Eqn. (16). The LCU method can also be seen as the implementation of a summation on a QC. Since this operation is not norm-preserving and, in general, cannot be performed deterministically, it is reflected in losses in the success probability.

In the application of the LCU algorithm, the first step is to decompose the arbitrary matrix into unitaries. To this end, the Pauli decomposition employs the complete basis formed by Pauli matrices (or their tensor products) to represent the operators acting on n qubits [49]. Once the unitary decomposition of M is achieved, each resulting unitary operator U_k is implemented in the circuits of Fig. 4 or Fig. 5, and it can be understood as a block matrix on the main diagonal of a larger matrix space. Their ordering on the diagonal matrix is defined by the binary selection of the controlled ancilla qubits, being $|00\rangle$ the first diagonal block matrix element, $|01\rangle$ the second diagonal block, etc. Appendix A describes the mapping between quantum gates and matrices for the quantum circuit of Fig. 4 in greater detail. The weights κ_k are contained in the first column of V_k as modified coefficients. Any excess unitaries in the power of two can easily be discarded by setting the corresponding κ_k coefficient in V_k to zero. Exemplified for $m = 4$, the unitary V_k reads

$$V_k = \begin{bmatrix} \sqrt{\frac{\kappa_0}{\kappa_0 + \kappa_1 + \kappa_2 + \kappa_3}} & & & \\ \sqrt{\frac{\kappa_1}{\kappa_0 + \kappa_1 + \kappa_2 + \kappa_3}} & \cdots & & \\ \sqrt{\frac{\kappa_2}{\kappa_0 + \kappa_1 + \kappa_2 + \kappa_3}} & & \cdots & \\ \sqrt{\frac{\kappa_3}{\kappa_0 + \kappa_1 + \kappa_2 + \kappa_3}} & & & \end{bmatrix}, \quad (19)$$

with the square roots appearing due to normalization, and the rest of the columns (...) forming an orthonormal

basis to make V_k unitary. This could be achieved by the application of the Gram-Schmidt process on randomly initialized columns.




FIG. 4: Quantum circuit to implement one explicit time step of the one-dimensional periodic heat equation, Eqn. (1).

In particular, the circuit of Fig. 4 prepares the initial state as

$$|00\rangle |\phi^t\rangle \rightarrow |00\rangle \tilde{A}_{00} |\phi^t\rangle + |01\rangle \tilde{A}_{01} |\phi^t\rangle + |10\rangle \tilde{A}_{10} |\phi^t\rangle + |11\rangle \tilde{A}_{11} |\phi^t\rangle, \quad (20)$$

with the matrices

$$\tilde{A}_{00} = \frac{\kappa_0 U_0 + \kappa_1 U_1 + \kappa_2 U_2 + \kappa_3 U_3}{\kappa_0 + \kappa_1 + \kappa_2 + \kappa_3} \quad (21)$$

$$\tilde{A}_{01} = \frac{1}{\kappa_0 + \kappa_1 + \kappa_2 + \kappa_3} \left(\sqrt{\kappa_1(\kappa_0 + \kappa_2 + \kappa_3)} U_1 - \sqrt{\kappa_1(\kappa_0 + \kappa_2 + \kappa_3)^{-1}} (\kappa_2 U_2 + \kappa_3 U_3 + \kappa_0 U_0) \right) \quad (22)$$

$$\tilde{A}_{10} = \sqrt{\frac{\kappa_2}{(\kappa_0 + \kappa_2 + \kappa_3)(\kappa_0 + \kappa_1 + \kappa_2 + \kappa_3)}} \left(\sqrt{(\kappa_0 + \kappa_3)} U_2 - \sqrt{(\kappa_0 + \kappa_3)^{-1}} (\kappa_3 U_3 + \kappa_0 U_0) \right) \quad (23)$$

$$\tilde{A}_{11} = \sqrt{\frac{\kappa_0 \kappa_3}{(\kappa_0 + \kappa_3)(\kappa_0 + \kappa_1 + \kappa_2 + \kappa_3)}} (U_3 - U_0). \quad (24)$$

From Eqn. (20), it is evident that the state $|00\rangle \tilde{A}_{00} |\phi^t\rangle$ needs to be post-selected out of the superposition of the four computational basis states to recover the action of $\tilde{A}_{00} = M/\alpha$ with $\alpha = \sum_k \kappa_k$. The probability of measuring that particular basis state $p_{|00\rangle}$ is expressed by the squared Euclidean norm of the quantum state in the Hilbert space, viz.

$$p_{|00\rangle} = \langle \phi | \tilde{A}_{00}^\dagger \tilde{A}_{00} | \phi \rangle = \|\tilde{A}_{00} | \phi \rangle\|_2^2 \leq \|\tilde{A}_{00}\|_2^2 \| | \phi \rangle \|_2^2 = \|\tilde{A}_{00}\|_2^2 = \sigma_{\max}(\tilde{A}_{00})^2 = \lambda_{\max}(\tilde{A}_{00}^\dagger \tilde{A}_{00}), \quad (25)$$

where $\sigma_{\max}$ and $\lambda_{\max}$ denote the largest singular value and eigenvalue, respectively. The same applies to the

measurement probabilities of the other computational basis $p_{|01\rangle} (\leq \sigma_{\max}(\tilde{A}_{01})^2)$, $p_{|10\rangle} (\leq \sigma_{\max}(\tilde{A}_{10})^2)$, and $p_{|11\rangle} (\leq \sigma_{\max}(\tilde{A}_{11})^2)$.

In principle, any unitary decomposition of M is possible. However, it is crucial to ensure an optimal probability of success. To this end, the unitary decomposition given in Eqn. (4) for the time-marching solution of the one-dimensional heat equation is considered. One step of the first-order temporal evolution $\phi(t + \Delta t) = A\phi(t)$ can be implemented exclusively with three unitaries $(I, S_0, S_0^\dagger)$ in the LCU circuit of Fig. 4 with $\kappa_0 = 1$, $\kappa_1 = \kappa_2 = r_h/(1 - 2r_h)$, and $\kappa_3 = 0$. Note that the additional constraint $r_h \leq 1/(2d)$ needs to be imposed to guarantee that all the LCU weights are $\kappa_k \geq 0$. Here, the LCU scaling factor $1/(\kappa_0 + \kappa_1 + \kappa_2 + \kappa_3) = 1 - 2r_h$ cancels the initial subnormalization factor $1/(1 - 2r_h)$ in Eqn. (4) completely and results in the exact recovering of A with $\alpha = 1$. Accordingly, the success probability of measuring the computational basis $|00\rangle$ for one time step is given by $p^t = \langle \phi^t | A^\dagger A | \phi^t \rangle = \|\phi^{t+1}\|_2^2 / \|\phi^t\|_2^2$ since the recovered operator A is not unitary [19]. For simulating one step of the two-dimensional heat equation, the quantum circuit in Fig. 5 can be applied in the same manner.

As a result, the time-marching operator A can be implemented in the LCU method to resolve the temporal evolution of the parabolic PDE in a time-marching fashion without vanishing cumulative probabilities p_c due to $\alpha = 1$. In particular, $p_c = \prod_0^T p^t = \|\phi^T\|_2^2 / \|\phi^0\|_2^2$ and converges for steady-state solutions, depending solely on the temporal distance between initial and final state. Therefore, it is an intrinsic property of the problem under investigation rather than of the applied numerical method.

E. Non-periodic boundary conditions

The method of images enforces boundary conditions by reflecting the bounded field and thereby exploiting the symmetric and antisymmetric properties of Neumann and Dirichlet boundaries, respectively [34]. This implies a two-fold increase in the degrees of freedom per non-periodic dimension, making the approach less appealing for classical computation. However, symmetry reflections on a QC can be implemented efficiently by introducing a single qubit for each non-periodic dimension in $\mathbb{R}^d$ [50], exploiting the exponential scaling in state space of the amplitude encoding strategy [10]. This permits boundary treatment through the even (sign-preserving) or odd (sign-changing) symmetry of neighboring points using discrete periodic operators, while conserving the advantage of non-vanishing success probabilities of the time-marching method, as described in Sec. III D.

The circuits depicted in Fig. 6, 7, and 8 illustrate the implementation of non-periodic boundary conditions for a single time step. In these circuits, the initial state $|0\phi^t\rangle$ is mirrored from the upper half of the register onto the lower half according to the desired boundary




FIG. 5: LCU circuit for the two-dimensional periodic heat equation.

conditions. Specifically, these circuits implement this reflection for one-dimensional (Fig. 6), two-dimensional (Fig. 7), and mixed two-dimensional boundaries (Fig. 8), respectively. After executing the unitary transformations V_k , U_c , and $V_k^\dagger$ which implement one explicit time step of the quantum time-marching algorithm, the conjugate transpose of the reflection circuit is applied to uncompute the initial reflection of the prepared state, returning the quantum register to $|0\rangle^{\otimes n_{\text{anc}}}$ conditioned to measure on the $|0\rangle^{\otimes n_{\text{anc}}}$ basis. In case of two-dimensional domains, the reflection circuits require the modifications shown in Fig. 7 and 8 for a lexicographic order of the two-dimensional field into a one-dimensional vector $\tilde{\phi}^t$ or $\tilde{\phi}^{t+1}$. Note that the even (sign-preserving) symmetry found for Neumann boundaries is implemented with Hadamard and controlled-NOT (CNOT) gates [50], while the odd (sign-changing) reflection in Dirichlet boundaries employs additional Z gates.

The boundary method has several limitations. First, boundaries in each non-periodic direction must be of the same type and value, since the reflection depends on the boundary type, and different boundary values would produce discontinuities in the reflected field. Second, a prescribed non-zero derivative for inhomogeneous Neumann boundaries cannot be implemented using central second-order stencils on the reflected field without applying a correction source term to the PDE. These limitations on the boundaries are summarized for the heat equation in Tab. I.

TABLE I: Boundary conditions for Eqn. (1) supported by the approaches described in Sec. III E.

	Homogeneous	Inhomogeneous
Dirichlet Reflection	✓	✓
Neumann Reflection	✓	×
Mixed Reflection	✓	×
Dirichlet no-reflection	×	×
Neumann no-reflection	✓	×

An alternative to the method of images for implementing homogeneous Neumann boundaries in the heat equation is presented next. In particular, the LCU method described in Sec. IIID is combined with the unitary de-

composition of the non-periodic time-marching operator from Eqn. (3), incorporating the corresponding boundary modifications given in Tab. II, as

$$\frac{A_{\text{Neumann}}}{1 - 2r_h} = I + \frac{r_h}{1 - 2r_h} (S_1 + S_2), \quad (26)$$

with

$$S_1 = \begin{bmatrix} 0 & 1 & \dots & 0 \\ 1 & 0 & & \\ \vdots & & 0 & 1 \\ & & 1 & 0 \\ 0 & \dots & & \ddots \end{bmatrix} = I \otimes X, \quad (27)$$

and

$$S_2 = \begin{bmatrix} 1 & & & 0 \\ & 0 & 1 & \\ & 1 & 0 & \\ & & 0 & 1 \\ & & 1 & 0 \\ & & & \ddots \\ 0 & & & & 1 \end{bmatrix}. \quad (28)$$

TABLE II: Boundary term modifications in Eqn. (3)

	$b_{1,1}$	$b_{1,N}$	$b_{N,1}$	$b_{N,N}$
Periodic	$1 - 2r_h$	r_h	r_h	$1 - 2r_h$
Dirichlet	$1 - 2r_h$	0	0	$1 - 2r_h$
Neumann	$1 - r_h$	0	0	$1 - r_h$

This direct embedding of the boundaries has the advantage of reducing the number of qubits required for the reflection, while preserving the optimal success probabilities per time step with linear time scaling [19]. Analogously to periodic cases, the LCU block encoding of A_{Neumann} results in $\alpha = 1$, which guarantees the convergence of the cumulative success probabilities for non-zero steady-state solutions, as shown in Sec. IIID.

The method can be extended to arbitrary dimensions. Here, one ($d = 1$) and two ($d = 2$) dimensions are described for the sake of clarity. For the one-dimensional




FIG. 6: Quantum circuit to implement one explicit time step of the one-dimensional non-periodic PDE defined by the U_c and V_k blocks. The total number of qubits is six with three qubits dedicated to $N = 8$ supporting points. For the Neumann boundaries, the even reflection of the register is prepared exclusively with Hadamard and CNOT gates, while the odd reflection of the Dirichlet condition also includes Z gates.




FIG. 7: Quantum circuit to implement one explicit time step of the two-dimensional non-periodic PDE defined by the U_c and V_k blocks. The total number of qubits is nine with four qubits dedicated to $N = 16$ supporting points. For the Neumann boundaries, the even reflection of the register is prepared exclusively with Hadamard and CNOT gates, while the odd reflection of the Dirichlet condition also includes Z gates.

case, the identity I and the circuits of Fig. 9 for S_1 and S_2 are applied to the LCU circuit of Fig. 4 with the coefficients $\kappa_0 = 1$, $\kappa_1 = \kappa_2 = r_h/(1 - 2dr_h)$, and $\kappa_3 = 0$. For two dimensions, the five unitaries I , $I \otimes S_1$, $I \otimes S_2$, $S_1 \otimes I$, and $S_2 \otimes I$ are used in the LCU circuit of Fig. 5 with the coefficients $\kappa_0 = 1$, $\kappa_1 = \kappa_2 = \kappa_3 = \kappa_4 = r_h/(1 - 2dr_h)$, and $\kappa_5 = \kappa_6 = \kappa_7 = 0$.

IV. COMPLEXITY

Given a problem in d dimensions with $N = N_{x_1}^d$ equidistant supporting points, the spatial information is encoded in the quantum register using $\log_2(N)$ qubits for the domain discretization, $\log_2(3+2(d-1))$ ancilla qubits for the LCU method, and one additional qubit per non-periodic dimension independent of the type of boundary conditions. As a result, the total number of qubits scales as $O(\log_2(N) + \log_2(d) + d)$. Furthermore, the number of attempts required for a successful execution of the quantum time-marching algorithm over the time horizon T is intrinsically determined by the inverse of p_c and it scales accordingly as $O(\|\phi(0)\|_2^2 / \|\phi(T)\|_2^2)$.

The temporal evolution is implemented by sequen-

tially applying the quantum circuits in Fig. 6, 7, or 8. These quantum circuits consist of a constant number of gates per time step, and due to the absence of amplitude amplification, the gate complexity grows linearly with the number of time steps $N_t = T/\Delta t$. The limiting stability condition r_h implies $\Delta t \propto \Delta x^2$, where $N_{x_1} \propto 1/\Delta x$, such that $N_t \propto TN_{x_1}^2$ and for d dimensions $N_t \propto TN_{x_1}^{2/d} = TN^{2/d}$. Given these contributions, the implementation has a gate complexity of $O((\log_2(N) + \log_2(d) + d)TN^{2/d})$. The composition of the complexity is identical to that in Ref. [19] suppressing logarithmic terms, except for the treatment of boundary conditions, which implies doubling N per dimension d .

Analog to the quantum algorithm, the classical implementation also scales in time with $N_t \propto TN_{x_1}^2$ in one dimension. For $N = N_{x_1}^d$ supporting points distributed over d dimensions, the classical complexity results instead in $O(NTN^{2/d}) = O(TN^{(2+d)/d})$. Comparing the exponent of the complexities, the factor of $2/(2+d)$ indicates a polynomial speed up of the quantum algorithm over classical algorithms. As the approximation schemes in the quantum and classical implementations are identical, the underlying error of the Euler method and the second-order discretization is not taken into account.




FIG. 8: Quantum circuit to implement one explicit time step of the two-dimensional non-periodic PDE defined by the U_c and V_k blocks. The total number of qubits is nine with four qubits dedicated to $N = 16$ supporting points. Mixed boundaries are implemented with Neumann in the first index and Dirichlet in the second.




FIG. 9: Three-qubit quantum circuits implementing S_1 and S_2 for homogeneous Neumann boundary conditions without reflection in the one-dimensional circuit of Fig. 4.

V. SIMULATIONS

The quantum methods for boundary implementation described above in Sec. III E are investigated in this section using a state-vector simulation for the two-dimensional non-periodic heat equation (Eqn. (1)). Homogeneous Neumann, Dirichlet, and mixed boundaries are incorporated into the quantum time-marching framework with the method of images, while the direct embedding of the unitary decomposition in the time-marching operator is used to implement homogeneous Neumann conditions. For the latter, the matrix A_{Neumann} is directly implemented with the LCU method to advance the linear transport by the quantum time-marching framework. Note that, for the evolution of N_t time steps, the sequential execution of the circuits shown in Figs. 6, 7, and 8 is equivalent to an N_t -sequence of U_c gates, as the other quantum gates simplify to identities.

The simulations refer to a spatial unit square $\vec{x} \in [0, 1] \times [0, 1]$ [m]. The initial condition ($t = 0$) for all cases studied is based on a two-dimensional Gaussian per-

turbation of the form $\phi(\vec{x}) = \exp(-200(x_1 - 0.25)^2 - 200(x_2 - 0.25)^2)$, placed in the lower left corner of the unit square. The temporal error measure ε is defined by the l2-norm of the difference between classical FD and Quantum March (QM) results, i.e., $\varepsilon^t = \|\phi_{\text{FD}}^t - \phi_{\text{QM}}^t\|_2$. The probability of success p^t and failure $1 - p^t$ per time step, together with the cumulative probability of success p_c per simulation time interval are also given in the results.

A. Neumann BC

This section reports the results obtained from two different implementations of homogeneous Neumann conditions along the route outlined by the methods described in Sec. III E.

First, the method of images is applied to enforce the boundary conditions using the reflection circuit illustrated in Fig. 7. The U_c gate of this circuit is defined by the controlled unitaries of the LCU circuit in Fig. 5 to implement the time-marching operator A of Eqn. (3). To this end, we define $U_0 = I$, $U_1 = I \otimes S_0$, $U_2 = I \otimes S_0^\dagger$, $U_3 = S_0 \otimes I$, and $U_4 = S_0^\dagger \otimes I$, with the LCU coefficients as $\kappa_0 = 1$, $\kappa_1 = \kappa_2 = \kappa_3 = \kappa_4 = r_h/(1 - 2dr_h)$, and $\kappa_5 = \kappa_6 = \kappa_7 = 0$. Furthermore, the quantum circuits given in Fig. 1 are used to implement the upper and lower diagonal shifts S_0 and $S_0^\dagger$. The results in Fig. 10 depict the QM solution of ϕ^t computed for $N_t = 12000$ explicit time steps with $r_h = 0.2$. The even (symmetric) reflection on the three additional quadrants II, III, and IV, without the register reset by the reflection part of the circuit in Fig. 7, is indicated in Fig. 10a, which corresponds to 1000 time steps. The full resolution contains 128×128 grid points, where $N = 64 \times 64$ points correspond to the (non-reflected) domain of interest in the first quadrant. A total of 17 qubits were used in these simulations, including 14 working qubits and three ancilla qubits.




FIG. 10: Comparison of QM and FD results of ϕ^t for the two-dimensional heat equation (Eqn. (1)) with homogeneous Neumann boundary conditions enforced by the method of images. Results refer to 128×128 grid points (including reflection quadrants) with 14 working qubits and three ancilla qubits, obtained from state-vector simulation for $N_t = 12000$ time steps with $r_h = 0.2$.


Second, the homogeneous Neumann conditions are implemented without reflection by combining the controlled unitaries of the LCU circuit (Fig. 5) with the circuits given in Fig. 9. To directly embed the two-dimensional matrix A_{Neumann} of Eqn. (26), the five unitaries in the LCU circuit are defined as $U_0 = I$, $U_1 = I \otimes S_1$, $U_2 = I \otimes S_2$, $U_3 = S_1 \otimes I$, and $U_4 = S_2 \otimes I$ with the same weights κ_k as in the boundary reflection method. Figure 11 displays the results of the QM simulations for an evolution of ϕ^t over $N_t = 12000$ time steps, again setting $r_h = 0.2$. Comparing the results of Fig. 10a and Fig. 11a, the different colorbars need to be taken into account. It should be noted that the absence of reflection quadrants spares two qubits, maintaining the resolution of 64×64 discrete points with 12 working qubits and three ancilla qubits.

The temporal evolution of the error ϵ^t reaches maximum values of approximately 10^{-11} in Fig. 10c and Fig. 11c. This validation against classical FD unequivocally demonstrates the accuracy of the proposed QM boundary methods. Furthermore, the probability of success per time step converges to $p^t = \|\phi^{t+1}\|_2^2 / \|\phi^t\|_2^2 = 1$ as the steady state is approached because $\alpha = 1$ in the block encoding of the time-marching operator,


cf. Sec. IIID. The application of a non-unitary operation to represent the irreversible diffusion dynamics of the heat equation explains the decay in the cumulative probability of success p_c for the transient solution, cf. Fig. 10b and Fig. 11b. However, the theoretically predicted p_c in Sec. IIID converges for both methods to $p_c = \|\phi^T\|_2^2 / \|\phi^0\|_2^2 = 0.123$, cf. Fig. 10b and Fig. 11b, which is intrinsic to the problem. Although this final converged value of $p_c = 0.123$ for the steady-state solution can only be increased at the cost of quadratic time complexity [26], the slope of its decay during the transient stage can be adjusted by modifying the time discretization, as Δt determines the difference between the consecutive solutions ϕ^t and ϕ^{t+1} . However, the selection of a smaller Δt with an increase in the number of time steps will not affect the final convergence to the problem's intrinsic probabilities.

B. Dirichlet BC


The homogeneous Dirichlet conditions are implemented by combining the quantum circuits of Fig. 5 and Fig. 7. Analogous to Sec. VA, the unitaries in U_c build-



(a) QM results of ϕ^t for $N_t = 0, 300, 1000$, and 12000 .



(b) Temporal evolution of the algorithm probabilities in the state-vector simulation.



(c) Temporal evolution of the error measure ϵ^t w.r.t. classical FD results.

FIG. 11: State-vector simulation of the two-dimensional heat equation (Eqn. (1)) with the homogeneous Neumann boundary conditions embedded directly in the time-marching operator. Results show 64×64 grid points for quadrant I (cf. Fig. 10) with 12 working qubits and three ancilla qubits. The stability parameter is $r_h = 0.2$ and $N_t = 12000$.


ing the time-marching operator A of Eqn. (3) are selected as, $U_0 = I$, $U_1 = I \otimes S_0$, $U_2 = I \otimes S_0^\dagger$, $U_3 = S_0 \otimes I$, and $U_4 = S_0^\dagger \otimes I$ with the circuits of Fig. 1 and the LCU weights $\kappa_0 = 1$, $\kappa_1 = \kappa_2 = \kappa_3 = \kappa_4 = r_h/(1 - 2dr_h)$, $\kappa_5 = \kappa_6 = \kappa_7 = 0$ with $d = 2$. The approach applied here differs from the one described in Sec. VA only in the type of reflection. The method of images maps the initial conditions by applying an odd (antisymmetric) reflection. This is achieved using the Z -gates in Fig. 7. The full resolution is 128×128 with $N = 64 \times 64$ grid points dedicated to the quadrant of interest (quadrant I). Again, a total of 17 qubits including three ancilla qubits are used in the state-vector simulation of $N_t = 12000$ time steps in combination with $r_h = 0.2$.

The results are depicted in Fig. 12. This figure shows ϕ^t for four exemplary time instants ($N_t = 0, 300, 1000, 12000$) of the QM temporal evolution in quadrant I (Fig. 12a), the algorithm probabilities (Fig. 12b), and the error measure ϵ^t over time (Fig. 12c). The validation of the method is clearly proven by the resulting low error values, with maximums of the order of 10^{-13} . As shown in Fig. 12b, the probability of success per time step p^t converges to one. However, the


cumulative probability p_c vanishes as the steady-state solution is reached. This is due to the homogeneous Dirichlet boundaries that force the steady-state solution to converge to zero, and consequently, the numerator in $p_c = \|\phi^T\|_2^2 / \|\phi^0\|_2^2$ also vanishes. Introducing inhomogeneous Dirichlet conditions would resolve this issue. Nevertheless, such an approach requires shifting the entire quantum register by a constant value, which is challenging to implement with universal quantum gates. To address this difficulty, TT-based techniques [51] such as Tensor Quantum Programming (TQP) [38, 52] can be employed to initialize the register, encoding the inhomogeneity into quantum gates.

C. Mixed BC


Finally, homogeneous mixed boundaries are simulated with the method described in Sec. III E. In the configuration under investigation, Neumann conditions are applied to the exterior boundaries whose normal points in the x_1 direction and Dirichlet conditions are applied in the x_2 direction, which implies an even (symmetric) and an odd (antisymmetric) reflection of the initial conditions,



(a) QM results of ϕ^t for $N_t = 0, 300, 1000$, and 12000 .



(b) Temporal evolution of the algorithm probabilities in the state-vector simulation.



(c) Temporal evolution of the error measure ε^t w.r.t. classical FD results. Convergence to zero due to a vanishing steady solution.

FIG. 12: Benchmarking of QM results for the two-dimensional heat equation (Eqn. (1)) with homogeneous Dirichlet boundary conditions enforced by the method of images. Results refer to 128×128 grid points (including reflection quadrants) with 14 working qubits and three ancilla qubits. State-vector simulation for $N_t = 12000$ time steps with $r_h = 0.2$.

respectively.

For the implementation, the reflection circuit of Fig. 8 is combined with the unitary block U_c of the LCU circuit in Fig. 5. The two-dimensional time-marching operator A of Eqn. (3) is constructed by means of $U_0 = I$, $U_1 = I \otimes S_0$, $U_2 = I \otimes S_0^\dagger$, $U_3 = S_0 \otimes I$, and $U_4 = S_0^\dagger \otimes I$. To this end, the circuits of Fig. 1 and the LCU coefficients $\kappa_0 = 1$, $\kappa_1 = \kappa_2 = \kappa_3 = \kappa_4 = r_h/(1 - 2dr_h)$, and $\kappa_5 = \kappa_6 = \kappa_7 = 0$ with $d = 2$ are employed. The simulation covers $N_t = 12000$ time steps with the stability parameter $r_h = 0.2$. The total spatial resolution is defined as 128×128 supporting points, including the reflection quadrants II, III, IV, where $N = 64 \times 64$ points are used for the quadrant of interest (I), cf. Fig. 10. The employed 17 qubits are divided into 14 working qubits and three ancilla qubits.


Figure 13a depicts the temporal evolution of ϕ^t in quadrant I for four exemplary time instants, i.e., $N_t = 0, 300, 1000, 12000$. There, the different effects of the boundaries on the spatial distribution of ϕ^t can clearly be observed. As in the previous experiments, the error values in Fig. 13c show excellent agreement with results obtained from classical FD approaches, with maximum

values in the order of 10^{-12} .


Paying attention to the probabilities of success, Fig. 13b illustrates the convergence of p^t to unity and p_c to zero. The latter is a consequence of the steady-state solution approaching zero due to the homogeneous Dirichlet condition along the x_2 boundaries. As discussed in Sec. VB, this limitation can be addressed by solving for non-homogeneous boundary conditions. To this end, the problem is substituted using a smooth extension, which requires adding a non-zero constant to the initial and the boundary conditions, to recover the homogeneous solution. The construction of the corresponding quantum gates to encode this adjustment can be implemented by TT techniques [38].

VI. CONCLUSION


This paper presents the complete implementation of time-marching algorithms to solve diffusion problems described by second-order PDEs with different types of boundary conditions on fault-tolerant QCs. The proposed approaches combine the LCU algorithm to capture



(a) QM results of ϕ^t for $N_t = 0, 300, 1000$, and 12000 .



(b) Temporal evolution of the algorithm probabilities in the state-vector simulation.



(c) Temporal evolution of the error measure ϵ^t w.r.t. classical FD results. Convergence to zero due to a vanishing steady solution.

FIG. 13: QM results for the two-dimensional heat equation (Eqn. (1)) with homogeneous mixed boundary conditions enforced by the method of images. The horizontal direction corresponds to Neumann while the vertical direction to Dirichlet. Results refer to 128×128 grid points (including reflection quadrants) with 14 working qubits and three ancilla qubits. State-vector simulation for $N_t = 12000$ time steps with $r_h = 0.2$.

the diffusive dynamics and the method of images to enforce boundaries. In particular, the implementation of diffusion while maintaining optimal success probabilities constitutes the strongest asset of the method. Moreover, it is noteworthy to remark that the logarithmic scaling in the number of spatial points makes the two-fold increase in the reflection for the boundary treatment particularly cheap and efficient on a QC. In the specific case of Neumann boundaries, the domain reflection in each non-periodic direction is replaced by directly block encoding the boundary into the unitary decomposition of the time-marching operator.

The presented results show unequivocally that all types of boundaries can be implemented on a fault-tolerant QC with very high accuracy compared to classical Euler-FD discretizations. Furthermore, optimal success probabilities with linear time scaling are demonstrated for two-dimensional simulations, while the proposed approach can straightforwardly be extended to arbitrary dimensions. The treatment of diffusive dynamics shows potential for implementing other non-unitary processes, while an application of the boundary treatment to other initial-boundary value problems in structural, fluid, or quantum

mechanics is an interesting direction of future research.

ACKNOWLEDGMENTS

This work and the research are supported by the European Union's Horizon Europe research and innovation program (HORIZON-CL4-2021-DIGITAL-EMERGING-02-10) under grant agreement No. 101080085 QCED. Furthermore, we thank Peter Brearley for the fruitful discussions.

AUTHOR CONTRIBUTIONS

Sergio Bengoechea: Conceptualization, Formal analysis, Investigation, Methodology, Software, Validation, Visualization, Writing - original draft, Writing - review & editing. **Paul Over:** Conceptualization, Formal analysis, Investigation, Methodology, Software, Validation, Visualization, Writing - original draft, Writing - review & editing. **Thomas Rung:** Funding acquisition, Project administration, Supervision, Resources, Writing

- original draft, Writing - review & editing.

DATA AVAILABILITY

The data is available via
doi.org/10.25592/uhhfdm.18048.

-
- [1] M. M. Aslam Bhutta, N. Hayat, M. H. Bashir, A. R. Khan, K. N. Ahmad, and S. Khan, *Applied Thermal Engineering* **32**, 1 (2012).
- [2] A. Jameson, *Journal of Scientific Computing* **3**, 233 (1988).
- [3] Z. Zhang, X. Wang, and Y. Yan, *e-Prime - Advances in Electrical Engineering, Electronics and Energy* **1**, 100009 (2021).
- [4] G. K. Vallis, *Atmospheric and Oceanic Fluid Dynamics: Fundamentals and Large-Scale Circulation*, 2nd ed. (Cambridge University Press, 2017).
- [5] H. N. Khan, D. A. Hounshell, and E. R. H. Fuchs, *Nature Electronics* **1**, 14 (2018).
- [6] H. Esmaeilzadeh, E. Blem, R. S. Amant, K. Sankaralingam, and D. Burger, in *2011 38th Annual International Symposium on Computer Architecture (ISCA)* (2011) pp. 365–376.
- [7] J. P. Slotnick, A. Khodadoust, J. Alonso, D. Darmofal, W. Gropp, E. Lurie, and D. J. Mavriplis, *CFD vision 2030 study: a path to revolutionary computational aerosciences*, Tech. Rep. (NASA, 2014).
- [8] D. Jaksch, P. Givi, A. J. Daley, and T. Rung, *AIAA Journal* **61**, 1885 (2023).
- [9] P. Givi, A. J. Daley, D. Mavriplis, and M. Malik, *AIAA Journal* **58**, 3715 (2020).
- [10] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information* (Cambridge University Press, 2010).
- [11] J. Preskill, *Quantum* **2**, 79 (2018).
- [12] M. Cerezo, A. Arrasmith, R. Babbush, S. C. Benjamin, S. Endo, K. Fujii, J. R. McClean, K. Mitarai, X. Yuan, L. Cincio, and P. J. Coles, *Nature Reviews Physics* **3**, 2522 (2021).
- [13] M. Lubasch, J. Joo, P. Moinier, M. Kiffner, and D. Jaksch, *Physical Review A* **101** (2020).
- [14] P. Over, S. Bengoechea, T. Rung, F. Clerici, L. Scandurra, E. de Villiers, and D. Jaksch, *Computers & Fluids* **288**, 106508 (2025).
- [15] S. Bengoechea, P. Over, D. Jaksch, and T. Rung, *AIAA Journal* (**in-press**) (2025).
- [16] Y. Sato, R. Kondo, I. Hamamura, T. Onodera, and N. Yamamoto, *Phys. Rev. Res.* **6**, 033246 (2024).
- [17] H. Alipanah, F. Zhang, Y. Yao, R. Thompson, N. Nguyen, J. Liu, P. Givi, B. J. McDermott, and J. J. Mendoza-Arenas, *Phys. Rev. Res.*, (2025).
- [18] P. Brearley and S. Laizet, *Phys. Rev. A* **110**, 012430 (2024).
- [19] P. Over, S. Bengoechea, P. Brearley, S. Laizet, and T. Rung, *Phys. Rev. A* **112**, L010401 (2025).
- [20] D. An, J.-P. Liu, and L. Lin, *Physical Review Letters* **131**, 150603 (2023).
- [21] P. C. S. Costa, S. Jordan, and A. Ostrander, *Phys. Rev. A* **99**, 012323 (2019).
- [22] E. Esmaeilifar, D. Ahn, and R. S. Myong, *Physics of Fluids* **36**, 106110 (2024).
- [23] A. J. Pool, A. D. Somoza, C. Mc Keever, M. Lubasch, and B. Horstmann, *Phys. Rev. Res.* **6**, 033257 (2024).
- [24] C. Sanavio, R. Scatamacchia, C. de Falco, and S. Succi, *Physics of Fluids* **36**, 057143 (2024).
- [25] S. S. Bharadwaj and K. R. Sreenivasan, *Phys. Rev. Res.* **7**, 023262 (2025).
- [26] D. Fang, L. Lin, and Y. Tong, *Quantum* **7**, 955 (2023).
- [27] A. Gilyén, Y. Su, G. H. Low, and N. Wiebe, in *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2019 (Association for Computing Machinery, New York, NY, USA, 2019) p. 193–204.
- [28] D. Camps, L. Lin, R. Van Beeumen, and C. Yang, *SIAM Journal on Matrix Analysis and Applications* **45**, 801 (2024).
- [29] S. Jin, N. Liu, and Y. Yu, *Phys. Rev. A* **108**, 032603 (2023).
- [30] S. Jin, X. Li, N. Liu, and Y. Yu, *Journal of Computational Physics* **498**, 112707 (2024).
- [31] S. Jin, X. Li, N. Liu, and Y. Yu, *SIAM Journal on Scientific Computing* **46**, B403 (2024).
- [32] S. Jin, N. Liu, and Y. Yu, *Phys. Rev. Lett.* **133**, 230602 (2024).
- [33] Z. Lu and Y. Yang, *Proceedings of the Combustion Institute* **40**, 105440 (2024).
- [34] A. M. Childs, J.-P. Liu, and A. Ostrander, *Quantum* **5**, 574 (2021).
- [35] D. J. Griffiths and D. F. Schroeter, *Introduction to Quantum Mechanics*, 3rd ed. (Cambridge University Press, 2018).
- [36] S.-J. Ran, *Phys. Rev. A* **101**, 032310 (2020).
- [37] P.-F. Zhou, R. Hong, and S.-J. Ran, *Phys. Rev. A* **104**, 042601 (2021).
- [38] A. Termanova, A. Melnikov, E. Mamenchikov, N. Belokonev, S. Dolgov, A. Berezutskii, R. Ellerbrock, C. Mansell, and M. R. Perelshtein, *New Journal of Physics* **26**, 123019 (2024).
- [39] I. V. Oseledets, *Constructive Approximation* **37**, 1 (2013).
- [40] D. W. Berry, A. M. Childs, and R. Kothari, in *2015 IEEE 56th Annual Symposium on Foundations of Computer Science* (2015) pp. 792–809.
- [41] A. M. Childs, D. Maslov, Y. Nam, N. J. Ross, and Y. Su, *Proceedings of the National Academy of Sciences* **115**, 9456 (2018).
- [42] N. Hatano and M. Suzuki, Finding exponential product formulas of higher orders, in *"Quantum Annealing and Other Optimization Methods* (Springer Berlin Hei-

- delberg, 2005) pp. 37–68.
- [43] D. W. Berry, A. M. Childs, R. Cleve, R. Kothari, and R. D. Somma, *Phys. Rev. Lett.* **114**, 090502 (2015).
 - [44] G. H. Low and I. L. Chuang, *Phys. Rev. Lett.* **118**, 010501 (2017).
 - [45] R. M. Gingrich and C. P. Williams, in *Proceedings of the Winter International Symposium on Information and Communication Technologies*, Vol. 58 (2004) pp. 1–6.
 - [46] L. Lapworth and C. Sünderhauf, *Quantum Science and Technology* **10**, 045064 (2025).
 - [47] S.-H. Lin, R. Dilip, A. G. Green, A. Smith, and F. Pollmann, *PRX Quantum* **2**, 010342 (2021).
 - [48] A. M. Childs and N. Wiebe, *Quantum Information & Computation* **12**, 901 (2012).
 - [49] L. Hantzko, L. Binkowski, and S. Gupta, *Physica Scripta* **99**, 085128 (2024).
 - [50] Y. Sano and I. Hamamura, *Phys. Rev. Appl.* **24**, 034062 (2025).
 - [51] N.-L. van Hülst, P. Siegl, P. Over, S. Bengoechea, T. Hashizume, M. G. Cecile, T. Rung, and D. Jaksch, Quantum-inspired tensor-network fractional-step method for incompressible flow in curvilinear coordinates (2025), arXiv:2507.05222.
 - [52] P. Siegl, G. S. Reese, T. Hashizume, N.-L. van Hülst, and D. Jaksch, Tensor-programmable quantum circuits for solving differential equations (2025), arXiv:2502.04425.

Appendix A: LCU circuit to matrix

As an example of state-vector simulation, the matrix-vector operations executed by the LCU circuit of Fig. 4 are illustrated here. Starting from the initial state

$$|00\phi^t\rangle = |0\rangle \otimes |0\rangle \otimes |\phi^t\rangle \in \mathbb{R}^{2^{n+n_{\text{anc}}}}, \quad (\text{A1})$$

the circuit sequentially applies the quantum gates $V_k, U_0, U_1, U_2, U_3, V_k^\dagger$, each represented by the corresponding matrix $\tilde{V}_k, \tilde{U}_0, \tilde{U}_1, \tilde{U}_2, \tilde{U}_3, \tilde{V}_k^\dagger \in \mathbb{R}^{2^{n+n_{\text{anc}}} \times 2^{n+n_{\text{anc}}}}$, evolving into the final state

$$|00\rangle |\phi^{t+1}\rangle + |01\rangle |*\rangle + |10\rangle |*\rangle + |11\rangle |*\rangle = \tilde{V}_k^\dagger \tilde{U}_3 \tilde{U}_2 \tilde{U}_1 \tilde{U}_0 \tilde{V}_k |00\phi^t\rangle. \quad (\text{A2})$$

The matrix $\tilde{V}_k$ is built with the matrices $V_k \in \mathbb{R}^{2^{n_{\text{anc}}} \times 2^{n_{\text{anc}}}}$ and $I \in \mathbb{R}^{2^n \times 2^n}$ as

$$\tilde{V}_k = V_k \otimes I = \begin{bmatrix} V_{k00} I & V_{k01} I & V_{k02} I & V_{k03} I \\ V_{k10} I & V_{k11} I & V_{k12} I & V_{k13} I \\ V_{k20} I & V_{k21} I & V_{k22} I & V_{k23} I \\ V_{k30} I & V_{k31} I & V_{k32} I & V_{k33} I \end{bmatrix}, \quad (\text{A3})$$

while the 00-controlled gate U_0 , the 01-controlled gate U_1 , the 10-controlled gate U_2 , and the 11-controlled gate U_3 result in the matrices

$$\tilde{U}_0 = \begin{bmatrix} U_0 & & & \\ & I & & \\ & & I & \\ & & & I \end{bmatrix}, \tilde{U}_1 = \begin{bmatrix} I & & & \\ & U_1 & & \\ & & I & \\ & & & I \end{bmatrix}, \quad (\text{A4})$$

$$\tilde{U}_2 = \begin{bmatrix} I & & & \\ & I & & \\ & & U_2 & \\ & & & I \end{bmatrix}, \text{ and } \tilde{U}_3 = \begin{bmatrix} I & & & \\ & I & & \\ & & I & \\ & & & U_3 \end{bmatrix}.$$