Efficient Quantum Circuits for Non-Unitary and Unitary Diagonal Operators with Space-Time-Accuracy trade-offs

Julien Zylberman, Ugo Nzongani, Andrea Simonetto, Fabrice Debbasch

Efficient Quantum Circuits for Non-Unitary and Unitary Diagonal Operators with Space-Time-Accuracy trade-offs

Julien Zylberman,[1] Ugo Nzongani,[2,3] Andrea Simonetto,[3] and Fabrice Debbasch[1]

[1] *Sorbonne Université, Observatoire de Paris, Université PSL, CNRS, LERMA, 75005 Paris, France*
[2] *CNRS, LIS, Aix-Marseille Université, Université de Toulon, Marseille, France*
[3] *Unité de Mathématiques Appliquées, ENSTA Paris, Institut Polytechnique de Paris, 91120 Palaiseau, France*

Unitary and non-unitary diagonal operators are fundamental building blocks in quantum algorithms with applications in the resolution of partial differential equations, Hamiltonian simulations, the loading of classical data on quantum computers (quantum state preparation) and many others. In this paper, we introduce a general approach to implement unitary and non-unitary diagonal operators with efficient-adjustable-depth quantum circuits. The depth, *i.e.*, the number of layers of quantum gates of the quantum circuit, is reducible with respect either to the width, *i.e.*, the number of ancilla qubits, or to the accuracy between the implemented operator and the target one. While exact methods have an optimal exponential scaling either in terms of size, *i.e.*, the total number of primitive quantum gates, or width, approximate methods prove to be efficient for the class of diagonal operators depending on smooth, at least differentiable, functions. Our approach is general enough to allow any method for diagonal operators to become adjustable-depth or approximate, decreasing the depth of the circuit by increasing its width or its approximation level. This feature offers flexibility and can match with the hardware limitations in coherence time or cumulative gate error. We illustrate these methods by performing quantum state preparation and non-unitary-real-space simulation of the diffusion equation: an initial Gaussian function is prepared on a set of qubits before being evolved through the non-unitary evolution operator of the diffusion process.

I. INTRODUCTION

Gate-based quantum computing leverages the principles of quantum mechanics to perform computations that would be prohibitively challenging on classical computers. The advantage of quantum algorithms lies in the possible decomposition of n-qubit operations into a reasonable number of hardware feasible single- and two-qubit operations. In this context, the development of new decompositions which improve on the current ones or widen the range of operations efficiently implementable on quantum computers, to include, *e.g*, non-unitary operations, is crucial to achieve novel applications and to lower the amount of resources.

In this paper, we focus on implementing one of the most fundamental operations in quantum computing, *i.e*, the diagonal operation, be it unitary or non-unitary. In fact, the computational cost of implementing diagonal operators is a limiting factor in many quantum algorithms with applications as varied as real-space simulation of unitary and non-unitary dynamics [1–6], quantum optimization [7, 8] or quantum state preparation [9–11].

The development of canonical decomposition of diagonal unitaries was first mentioned in the context of general unitary synthesis by A.Barenco et al. in 1996 [12]. Considering an n-qubit diagonal unitary, the authors propose to implement each eigenvalue sequentially with 2^n multi-controlled phase gates, each of these gates being in turn implementable with $\mathcal{O}(n^2)$ primitive quantum gates (or only $\mathcal{O}(n)$ if one uses one ancilla qubit). In 2004, S.Bullock et al. [13] were the first to present a quantum circuit related to the Walsh-Hadamard transform of the phases of the diagonal unitaries. Their quantum circuits, made only of controlled-NOT gate and one qubit z-axis rotations, implement any diagonal unitary with $2^{n+1}-3$ gates, and are asymptotically optimal in terms of size for exact implementation. It is only in 2014 that Welch et al. [14] developed the first approximate quantum circuits with no exponential scalings for diagonal unitaries depending on smooth, at least once differentiable, functions. The complexity of their quantum circuits, also based on Walsh-Hadamard transform, is directly related to the maximum value of the derivative of the function, making the synthesis efficient as long as the function has small variations. The same year, Welch et al. developed efficient quantum circuits in the Clifford+T basis for diagonal unitaries with $k \ll 2^n$ distinct eigenvalues [14]. In 2021, approximate quantum compiling have been developed to train diagonal ansatze [15]. In 2023, Sun et al. [16] developed a recursive scheme for exact implementation of diagonal unitaries based on a particular ordering of the Walsh operators. Their method has an optimal depth scaling as $\mathcal{O}(2^n/n)$ and can be implemented with ancilla qubits to reduce the depth of the quantum circuits. In 2023 also, Zylberman et al. [10] implemented diagonal unitaries using sparse Walsh Series: by considering the smallest number of Walsh coefficients approximating a smooth, at least differentiable, function, they were able to reduce the depth of the associated quantum circuits. Recently, Claudon et al. [17] were able to revisit the initial results of Barenco et al., by developing a poly-logarithmic depth quantum circuit for multi-controlled operations, improving the circuit complexity of implementing one eigenvalue of a diagonal operator from $\mathcal{O}(n)$ to $\mathcal{O}(\log(n)^3)$ using one ancilla qubit.

All above methods can be grouped into two families: (i) the methods that use a sequential decomposition and implement one eigenvalue at a time [12, 13, 17] (ii) the methods that use a Walsh (or Walsh-Hadamard) decomposition [10, 13, 14, 16]. All methods express diagonal unitaries as a product of simpler, efficiently implementable, diagonal unitaries. For generic diagonal unitary, exact decompositions involve an exponential amount of these simpler diagonal unitaries, leading to non-efficient quantum circuits. The hope is then that, for many applications, the relevant diagonal unitaries are actually non generic. For example, the relevant diagonal unitaries can be sparse, *i.e.*, only have a certain "small" number of eigenvalues different from unity, or the phases of the eigenvalues can be computed from functions with some smoothness properties, as for real-space simulations based on the implementation of evolution operators which are diagonal in a specific basis [14, 18].

In this work, we present an adjustable-depth framework to implement diagonal unitaries: from any decomposition into primitive diagonal operators, be it sparse or dense, exact or approximate, the depth of the associated quantum circuits can be reduced using an arbitrary number of ancilla qubits. The ancilla qubits are used to parallelize the implementation of the primitive diagonal operators, giving a first method to adjust the depth of the circuits. For the class of diagonal unitary depending on smooth, at least differentiable, functions, we prove that one can also adjust the depth of the circuits with respect to the approximation, giving a second tunable feature of the circuit. These contributions are summarized in different theorems:

- Full parallelization Theorem III.1 for any diagonal unitary decomposed into a product of p simpler, efficiently implementable, diagonal unitaries $\hat{U} = \prod_{j=0}^{p-1} \hat{U}_j$ given a sufficient number of ancilla qubits.
- Adjustable-depth Theorem III.2 for any diagonal unitary decomposed into a product of p simpler, efficiently implementable, diagonal unitaries $\hat{U} = \prod_{j=0}^{p-1} \hat{U}_j$ given a limited number of ancilla qubits.
- Approximate Theorem IV.1 to modify any exact methods into an approximate one for any diagonal unitary associated to a smooth, at least differentiable, function.

These first results (Theorems III.1, III.2) are similar in spirit to those in [19], but they were developed independently and they are here presented in full generality and with numerical simulations.

From now on we focus on diagonal non-unitaries. Implementing non-unitary operators with quantum computers has been often discussed in the last decades [20–23] and is crucial to broaden the range of applications accessible to quantum computing. Some algorithms of quantum state preparation (QSP) involve non-unitary operations [20, 24]. Linear combination of unitaries (LCU) achieves Hamiltonian simulation by implementing the Taylor expansion of the evolution operator into a series of non-unitary operators [25, 26]. Solving ordinary differential equations is also possible using non-unitary operations within the LCU framework [27, 28] and many quantum partial differential equations solvers are based on non-unitary operations [29, 30]. The quantum singular value transform (QSVT) makes it possible to implement polynomials of a block encoded operator and quantum algorithms such as Shor's algorithm [31], quantum search [32, 33] or quantum phase estimation [34] can be recast as a QSVT [29, 35]. Most of these algorithms are based on the block-encoding of a non-unitary operator into a larger unitary one by adding ancilla qubits and the computational resources are often given as the number of queries to a block-encoded operator. In particular, the development of quantum circuits for non-unitary diagonal operators has attracted attention.

In this article, we show how to block-encode a non-unitary diagonal operator from a unitary one, keeping the adjustable properties of the quantum circuits. We present the different strategies that one can use starting from a block-encoded diagonal operator: the block-encoding may be enough for the target application, but one may also need to measure an ancilla qubit to get a renormalized qubit state on which the diagonal non-unitary operator has been performed. The probability of success depends directly on the diagonal operator and on the qubit state. Two distincts approaches are presented: a non-destructive repeat-until-success scheme and an amplitude amplification scheme for which the depth of the quantum circuit is adjustable with respect to the probability of success. Finally, we illustrate these methods with two applications: a new quantum state preparation protocols of smooth, at least differentiable, functions is presented and the real-space simulation of the diffusion equation is performed by preparing an initial Gaussian distribution which evolves through the non-unitary evolution operator of the diffusion equation. These contributions can be summarize by the following points:

- Block-encoding Theorem V.1 of non-unitary diagonal operators $\hat{D}$ using adjustable-depth quantum circuits.
- Table of complexities I for unitary and non-unitary diagonal operators depending on smooth, at least differentiable, functions.
- Table of complexities II for sparse unitary and non-unitary diagonal operators.
- Non-destructive repeat-until-success protocol for non-unitary diagonal operators when amplitude amplification fails.

- Efficient quantum state preparation Theorem VIII.1 with space-time-accuracy trade-offs to prepare quantum states depending on differentiable functions.
- Resolution of the diffusion equation, see Fig. 9, with high fidelity, few qubits and low depth quantum circuits.

The remainder of the paper is organized as follows. Section II presents the sequential and Walsh-Hadamard decompositions and how to copy quantum registers. Section III introduces the general approach to adjust the depth of the quantum circuits with respect to the number of ancilla qubits and Section IV with respect to the accuracy. In Section V, the adjustable-depth quantum circuits for non-unitary diagonal operators are presented. Section VI summarizes the asymptotic scalings of the different methods for diagonal operators. Section VII details the amplitude amplification and the non-destructive repeat-until-success protocols. Finally, Section VIII presents two applications: quantum state preparation and the resolution of the diffusion equation. Additional technical details are left in the Appendices.

II. PRELIMINARIES

This section introduces preliminary notions used across the paper. We present the metrics chosen to estimate the computational resources of the quantum circuits, as well as the two decompositions of diagonal unitaries, *i.e.*, the sequential and the Walsh-Hadamard decomposition. We also introduce the copy unitary used to prepare the ancilla registers to parallelize the computations.

A. Depth, Size, Width, and Accuracy

Different metrics exist to evaluate the computational resources needed to implement a given quantum circuit. Size is the number of single-qubit and two-qubit gates. Depth is the number of layers of quantum gates, and it is related to the time needed to perform the desired operation. Width is the number of ancilla qubits. A metric for accuracy for us, is the error in spectral norm between the implemented operator and the target operator. The set of primitive single-qubit and two-qubit gates considered in this article is the set of all single-qubit gates complemented by the control NOT gate (CNOT). Other metrics such as the CNOT count or the T-count can also be estimated from the scalings provided in this paper. A unitary operator acting on n qubits is said to be efficiently implementable if there is a quantum circuit implementing it up to an error $\epsilon > 0$ in spectral norm with a size, depth and width scaling at worst as $\mathcal{O}(\text{poly}(n, 1/\epsilon))$[1]. By extension, the quantum circuit is said efficient. A non-unitary operator is said to be efficiently implementable on a set $\mathcal{H}$ of qubit states if the associated quantum circuit is efficient and the probability of successfully implementing the non-unitary operator on any n-qubit state $|\psi\rangle \in \mathcal{H}$ scales at worst as $\Omega(\text{poly}(1/n, \epsilon))$. Some of the quantum circuits presented in this article have a depth and/or a size independent of n, making them particularly relevant for large scale fault-tolerant algorithms.

B. Sequential decomposition

The sequential decomposition expresses the diagonal unitary $\hat{U} = \sum_{x=0}^{N-1} e^{i\theta_x} |x\rangle\langle x|$ acting on n qubits (with $N = 2^n$ eigenvalues) as a product of simple diagonal unitaries $\hat{U}_j = e^{i\theta_j} |j\rangle\langle j| + \sum_{x\neq j} |x\rangle\langle x|$ with $j \in \{0, ..., N-1\}$, $x = \sum_{i=0}^{n-1} x_i 2^i \in \{0, ..., N-1\}$, $x_i \in \{0, 1\}$ and $|x\rangle = |x_0\rangle \otimes \cdots \otimes |x_{n-1}\rangle$ being the computational basis of the Hilbert space $\mathcal{H}_2^{\otimes n}$. For each j, the operator $\hat{U}_j$ has by construction $2^n - 1$ eigenvalues equal to unity and the other eigenvalue is $\exp(i\theta_j)$, which is one of the eigenvalues of $\hat{U}$:

$$\hat{U} = \prod_{j=0}^{N-1} \hat{U}_j, \tag{1}$$

and with the notation $\hat{U} = \text{diag}(u_0, \ldots, u_{N-1})$:

$$\text{diag}(u_0, ..., u_{N-1}) = \text{diag}(u_0, 1, \ldots, 1) \times \text{diag}(1, u_1, \ldots, 1) \times \cdots \times \text{diag}(1, \ldots, 1, u_{N-1}). \tag{2}$$

[1] More formally, the "efficiency" is an asymptotic property that is rigorously defined on a sequence of unitaries $\{\hat{U}_n\}_{n=1}^{+\infty}$, where $\forall n$, $\hat{U}_n$ is acting on n qubits. The sequence $\{\hat{U}_n\}_{n=1}^{+\infty}$ is said to be efficiently implementable if $\forall \epsilon > 0, \forall n$, there exist a quantum circuit implementing $\hat{U}_n$ up to an error $\epsilon > 0$ in spectral norm with size, depth and width being $\mathcal{O}(\text{poly}(n, 1/\epsilon))$.

Thus, each $\hat{U}_j$ is a n-qubit unitary which changes the phase of one of the eigenvectors of the computational basis. By denoting $j = \sum_{i=0}^{n-1} j_i 2^i$ the binary decomposition of j (where $j_i \in \{0,1\}$ for all values of i), one can show that $\hat{U}_j$ corresponds to applying the phase $e^{i\theta_j}$ on one of the n qubits if all the qubits are in state $|j_i\rangle$. The operator $\hat{U}_j$ can therefore be written as a $(n-1)$-controlled phase gate $\hat{P}(\theta_j) = |0\rangle\langle 0| + e^{i\theta_j}|1\rangle\langle 1|$ if $j_{n-1} = 1$, or $e^{i\theta_j}|0\rangle\langle 0| + |1\rangle\langle 1| = \hat{X}\hat{P}(\theta_j)\hat{X}$ if $j_{n-1} = 0$, which is controlled by the qubits $i \in \{0, ..., n-2\}$ for which $j_i = 1$ and anti-controlled by the qubits i for which $j_i = 0$. The anti-controls can be rewritten as control operations using $\hat{X}$-Pauli gates:

$$\hat{U}_j = (\hat{X}_0^{j_0} \otimes ... \otimes \hat{X}_{n-1}^{j_{n-1}}) \Lambda_{\{0,...,n-2\}}(\hat{P}(\theta_j)) (\hat{X}_0^{j_0} \otimes ... \otimes \hat{X}_{n-1}^{j_{n-1}}). \tag{3}$$

with $\Lambda_{\{1,...,n-1\}}(\hat{P}(\theta_j))$ a $(n-1)$-controlled $\hat{P}(\theta_j)$ gates and $\hat{X}_i^{j_i}$ the $\hat{X}$-Pauli gate applied on qubit i if $j_i = 1$.

Let k_j be the number of 1's in the binary decomposition of j. Each $\hat{U}_j$ is implementable using $2k_j$ $\hat{X}$-Pauli gates and a $(n-1)$-multi-controlled gate. The recent work of Claudon et al. [17] introduces polylogarithmic-depth quantum circuits to implement $(n-1)$-multi-controlled gates by breaking it down into smaller multi-controlled gates which are implemented in parallel. Their approximate method implements the gate $\Lambda_{\{1,...,n-1\}}(\hat{P}(\theta_j))$ up to an error $\epsilon > 0$ with a quantum circuit of depth $\mathcal{O}(\log(n)^3 \log(1/\epsilon))$, size $\mathcal{O}(n \log(n)^4 \log(1/\epsilon))$ and without ancilla qubits (Proposition 2 in [17]). One can also implement exactly the n-control gate using one ancilla qubit (zeroed or borrowed[2]) with a depth $\mathcal{O}(\log(n)^3)$ and a size $\mathcal{O}(n \log(n)^4)$ (Corollary 1 in [17]) or without ancilla qubits with a linear depth $\mathcal{O}(n)$ and a linear size $\mathcal{O}(n)$ [36].

Being diagonal operators, the $\hat{U}_j$'s commute with one another. It thus makes sense to wonder if there is an optimal order in which to implement these operators. It is possible to use a Gray code order [37–39], *i.e.*, from two following operators $\hat{U}_j$ and $\hat{U}_{j'}$ only one bit differs from the binary decomposition of j and j', which cancels a maximum number of $\hat{X}$-Pauli gates from two consecutive operators. Thus, only one $\hat{X}$-Pauli gate remains between each multi-controlled gate, reducing the overall cost of the implementation (see Appendix A 1 for examples of quantum circuits with and without Gray code ordering).

C. Walsh-Hadamard decomposition

The exact synthesis of quantum circuits for diagonal unitaries has been proven asymptotically optimal in terms of size (up to a factor 2) using $2^{n+1} - 3$ CNOT and $\hat{R}_Z$ quantum gates [13]. As shown by Welch et al. [14], this construction is directly related to the Walsh-Hadamard representation of a set of $N = 2^n$ data points: a N-Walsh Series can represent exactly a set $\{\theta_k\}$ of N data points through the Walsh-Hadamard transform.

Walsh Series and Walsh functions were first introduced by Walsh in 1923 [40]. The Walsh function of order $j \in \{0, 1, ..., N-1\}$ is defined on $[0,1]$ by

$$w_j(x) = (-1)^{\sum_{i=0}^{n-1} j_i x_i}, \tag{4}$$

where j_i is the i-th bit in the binary expansion $j = \sum_{i=0}^{n-1} j_i 2^i$ and x_i is the i-th bit in the dyadic expansion $x = \sum_{i=0}^{\infty} x_i / 2^{i+1}$.

Walsh functions only take ± 1 values depending on their order j and are naturally implementable on quantum computers through their associated $\hat{Z}$-Pauli operators:

$$\hat{w}_j = \bigotimes_{i=0}^{n-1} (\hat{Z}_i)^{j_i}. \tag{5}$$

To implement the diagonal unitary $\hat{U} = \sum_{k=0}^{N-1} e^{i\theta_k} |k\rangle\langle k|$, which depends on the N real values $\{\theta_k\}$, one first defines the Walsh coefficients:

$$a_j = \frac{1}{N} \sum_{k=0}^{N-1} \theta_k w_j(k/N), \tag{6}$$

[2] Zeroed ancilla qubits are initially in state $|0\rangle$ and reset to $|0\rangle$ at the end of the computation. Borrowed ancilla qubits are in an arbitrary state $|\psi\rangle$ potentially entangled with other qubits and are restored to their initial state afterwards.

and write

$$\hat{U} = \prod_{j=0}^{N-1} \hat{W}_j, \tag{7}$$

where $\hat{W}_j = e^{ia_j \hat{w}_j}$.

Using the fact that a tensor product of $\hat{Z}$-Pauli gates can be rewritten using two CNOT stairs and one $\hat{Z}$-Pauli gate, one can show that each of the operators $\hat{W}_j$ can be implemented using one $\hat{R}_Z$ gate and $2k_j$ CNOT gates, with k_j the number of bit equals to unity in the binary decomposition of j. For $j = \sum_{i=0}^{p-1} j_i 2^i$ with $1 \leq p \leq n$ and $j_{p-1} = 1$, the operator $\hat{W}_j$ acts on the p-th qubits $0, ..., p-1$ as:

$$\hat{W}_j = \left(\prod_{i=0}^{p-2} \widehat{CNOT}_{i \to p-1}^{j_i} \right) \left(\hat{I}_2^{\otimes (p-1)} \otimes \hat{R}_Z(a_j) \right) \left(\prod_{i=0}^{p-2} \widehat{CNOT}_{i \to p-1}^{j_i} \right)^{-1}, \tag{8}$$

where $\widehat{CNOT}_{i \to p-1}^{j_i}$ is the CNOT gate controlled by the i-th qubit and applied on the $(p-1)$-th qubit if $j_i = 1$, the symbol $\prod_{i=0}^{p-2} \hat{A}_i = \hat{A}_0 ... \hat{A}_{p-2}$ is the product of operator $\hat{A}_i$ with indexes in increasing order, $\hat{I}_2^{\otimes (p-1)}$ is the identity operator on the qubits 0 to $p-2$ and $\hat{R}_Z(a_j) = e^{ia_j} |0\rangle \langle 0| + e^{-ia_j} |1\rangle \langle 1|$ acts on the qubit $p-1$. The operator $\hat{W}_0(a_0) = e^{ia_0} \hat{I}_2^{\otimes n}$ is a phase encoding the average value of the $\{\theta_k\}$.

Being diagonal operators, the $\hat{W}_j$ commute with one another. Similarly to the sequential decomposition, one can choose a Gray code ordering to cancel a maximum of CNOT gates from two consecutive $\hat{W}_j$ operator (see Appendix A 2 for examples of quantum circuit with and without Gray code ordering).

Note that sparse Walsh-Hadamard decomposition, *i.e.*, diagonal unitary expressed with only a number $s < N$ of $\hat{W}_j$ operators, are particularly efficient to approximate diagonal unitaries depending on smooth, at least once differentiable, functions [10]. In this case, the Gray code ordering may not be efficient to cancel a maximum of CNOT gates.

D. Preparation of the ancilla registers

The main results of this article use ancilla qubits to adjust the depth of the quantum circuits by parallelizing the implementation of the diagonal unitaries $\hat{U}_j$ or $\hat{W}_j$. This is possible by "copying"[3] the bit value of each of the n qubits into other ancilla qubits. Consider the i-th qubit in an arbitrary state $|\psi\rangle_i = \alpha |0\rangle_i + \beta |1\rangle_i$ and a register of m_i zeroed ancilla qubit in state $|0\rangle^{\otimes m_i}$. The copy unitary $\hat{C}_i$ copies acts as:

$$\hat{C}_i (\alpha |0\rangle_i + \beta |1\rangle_i) \otimes |0\rangle^{\otimes m_i} = \alpha |0\rangle^{\otimes (m_i+1)} + \beta |1\rangle^{\otimes (m_i+1)}. \tag{9}$$

This operation has already been introduced for other quantum circuit synthesis [41, 42]. We state and prove clearly the complexity associated to the copy operation with the following lemma.

Lemma II.1. *The "copy" unitary $\hat{C}_i$ acting on $m_i + 1$ qubits can be implemented using m_i CNOT gates with a depth $\lceil \log_2(m_i + 1) \rceil$.*

Proof. Note first that a set of q CNOT gates applied to different qubits can be implemented in parallel and therefore has a depth 1. Define $k = \lceil \log_2(m_i + 1) \rceil$. A direct recursion on $k' = 1, ..., k-1$ proves that the copy unitary doubles the number of "copied" qubit at each step, so each newly copied qubit can be used to copy another one at the next step as

$$\begin{aligned}
(\alpha |0\rangle + \beta |1\rangle) \otimes |0\rangle^{\otimes m_i} &\xrightarrow{\widehat{CNOT}_{0\to 1}} (\alpha |0\rangle^{\otimes 2} + \beta |1\rangle^{\otimes 2}) \otimes |0\rangle^{\otimes m_i - 1} \\
&\xrightarrow{\widehat{CNOT}_{0\to 2} \otimes \widehat{CNOT}_{1\to 3}} (\alpha |0\rangle^{\otimes 4} + \beta |1\rangle^{\otimes 4}) \otimes |0\rangle^{\otimes m_i - 3} \\
&\vdots \\
&\xrightarrow{\otimes_{j=0}^{2^{k'}-1} \widehat{CNOT}_{j\to(j+2^{k'})}} \alpha |0\rangle^{\otimes 2^{k'+1}} + \beta |1\rangle^{\otimes 2^{k'+1}} \otimes |0\rangle^{(m_i - (2^{k'+1}-1))},
\end{aligned} \tag{10}$$

[3] The name "copy" does not refer to the cloning operation $|\psi\rangle \otimes |0\rangle \to |\psi\rangle \otimes |\psi\rangle$ which is forbidden by the no-cloning theorem.

where $\widehat{CNOT}_{j\to(j+2^{k'})}$ is the CNOT gate controlled by the j-th qubit and applied on $(j+2^{k'})$-th zeroed qubit. After $k-1$ steps, the qubit state is $(\alpha\left|0\right\rangle^{\otimes 2^{k-1}}+\beta\left|1\right\rangle^{\otimes 2^{k-1}})\otimes\left|0\right\rangle^{\otimes(m_i-(2^{k-1}-1))}$. A last layer of CNOT gates suffices to produce the state $\alpha\left|0\right\rangle^{\otimes(m_i+1)}+\beta\left|1\right\rangle^{\otimes(m_i+1)}$. □

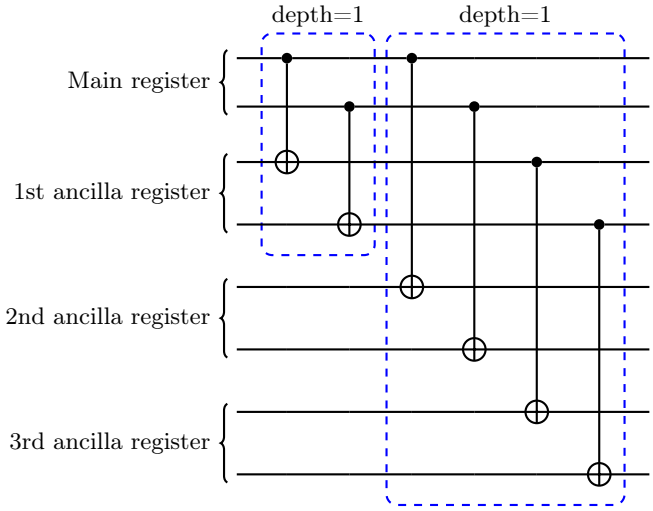

Figure 1: Quantum circuit to copy a register of 2 qubits into 3 ancilla registers. The qubits of the ancilla registers are in the $\left|0\right\rangle$ state. The operations framed together can be executed in parallel as they act on different qubits. The depth of the circuit scales logarithmically in the number of copies as the number of copying qubits doubles at each step.

A direct corollary of this lemma adresses the copy of a register of n qubits, where each of the n qubits is copied on m_i other qubits. This operation is performed by implementing each $\hat{C}_i$ in parallel through the operator $\widehat{\text{copy}}=\bigotimes_{i=0}^{n-1}\hat{C}_i$. Its size is given by $\sum_{i=0}^{n-1}m_i$ CNOT gates and its depth is $\lceil\log_2(m+1)\rceil$ where $m=\max_{i=0,\dots,n-1}m_i$. A small instance of the copy unitary is shown in Fig. 1 where a register of 2 qubits gets copied into 3 ancilla registers of 2 qubits each.

III. ADJUSTABLE-DEPTH FRAMEWORK USING ANCILLA QUBITS

In this section, we present an approach to adjust the depth of a quantum circuit implementing a diagonal unitary. Let $\hat{U}$ be a diagonal unitary acting on n qubits with a decomposition $\hat{U}=\prod_{j=0}^{p-1}\hat{U}_j$ where each $\hat{U}_j$ acts on $k_j\leq n$ qubits and is implementable with a quantum circuit of size s_j and depth d_j. The $\hat{U}_j$ can be sequential operators (3) as well as exponential of Walsh operators (8). Without ancilla qubits, one can implement $\hat{U}$ by performing each $\hat{U}_j$ one by one, leading to a quantum circuit of size $s=\sum_{j=0}^{p-1}s_j$ and depth $d=\sum_{j=0}^{p-1}d_j$. With ancilla qubits, one can parallelize the implementation of a maximum number of $\hat{U}_j$ operators. The parallelization can be partial or total, reducing the depth proportionally to the number of ancilla qubits available at the cost of preparing the ancilla registers with copy unitaries.

First, suppose one has enough available ancilla qubits to prepare $p-1$ registers:

$$\left|x\right\rangle\left|0\right\rangle^{\otimes k}\xrightarrow{\widehat{\text{copy}}}\left|x\right\rangle\left|\tilde{x}\right\rangle_1\dots\left|\tilde{x}\right\rangle_{p-1}, \tag{11}$$

where each register $\left|\tilde{x}\right\rangle_j$ contains only a copy of the k_j qubits on which $\hat{U}_j$ is acting non-trivially and $k=\sum_{j=1}^{p-1}k_j$. More formally, if one defines the support of $\hat{U}_j$ as the set of qubits on which $\hat{U}_j$ does not act like the identity, the state $\left|\tilde{x}\right\rangle_j$ contains only a copy of the qubits in the support of $\hat{U}_j$.

One can then implement each of the p diagonal unitaries $\hat{U}_j=\sum_{x=0}^{N-1}e^{i\theta_j(x)}\left|x\right\rangle\left\langle x\right|$ on a different register of qubits:

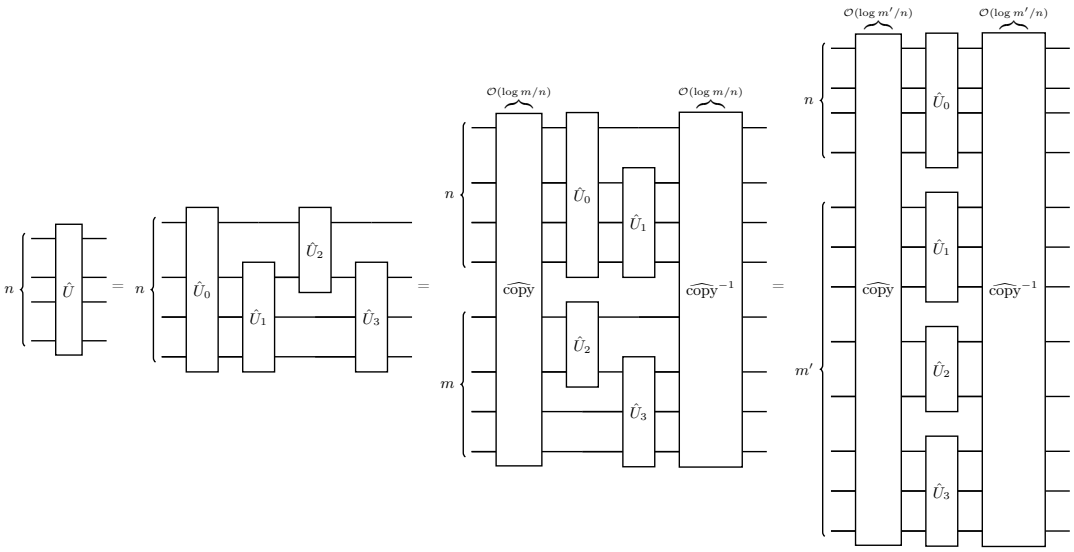

Figure 2: Adjustable-depth quantum circuit for a diagonal unitary $\hat{U} = \prod_{j=0}^{3} \hat{U}_j$ using ancilla qubits. The main register is composed of n qubits, the depth of the circuit is reduced by adding ancilla qubits. The "copy" operations are performed with a depth logarithmic in m/n.

$$\begin{aligned} |x\rangle |\tilde{x}\rangle_1 ... |\tilde{x}\rangle_{p-1} \xrightarrow{\bigotimes_{j=0}^{p-1} \hat{U}_j} & \hat{U}_0 |x\rangle \hat{U}_1 |\tilde{x}\rangle_1 ... \hat{U}_{p-1} |\tilde{x}\rangle_{p-1} \\ =& e^{i\theta_0(x)} e^{i\theta_1(x)} ... e^{i\theta_{p-1}(x)} |x\rangle |\tilde{x}\rangle_1 ... |\tilde{x}\rangle_{p-1} \\ =& e^{i\theta(x)} |x\rangle |\tilde{x}\rangle_1 ... |\tilde{x}\rangle_{p-1} \end{aligned} \tag{12}$$

with $e^{i\theta(x)} = e^{i\sum_{j=0}^{p-1} \theta_j(x)}$ being the eigenvalue of $\hat{U}$ associated to the $|x\rangle$ eigenvector. Finally, the inverse copy operation resets the ancilla qubits in the state $|0\rangle$:

$$e^{i\theta(x)} |x\rangle |\tilde{x}\rangle_1 ... |\tilde{x}\rangle_{p-1} \xrightarrow{\widehat{\text{copy}}^{-1}} e^{i\theta(x)} |x\rangle |0\rangle^{\otimes k} . \tag{13}$$

The full parallelization protocol has the following complexity:

Theorem III.1 (Full parallelization)**.** *The implementation of the diagonal operator* $\hat{U} = \prod_{j=0}^{p-1} \hat{U}_j$ *can be fully parallelized with a quantum circuit using* $k = \sum_{j=1}^{p-1} k_j$ *ancilla qubits, with a size* $s + 2k$ *and a depth bounded by* $\max_j(d_j) + 2\lceil \log_2(p) \rceil$.

Proof. The size of the quantum circuit is given by the sum of the size of each $\hat{U}_j$ and the size of the $\widehat{\text{copy}}$ and $\widehat{\text{copy}}^{-1}$ unitaries: one controlled-NOT gate is applied on each of the k ancilla qubits for the $\widehat{\text{copy}}$ and a second time for $\widehat{\text{copy}}^{-1}$, giving a final size of $s + 2k$. The depth corresponds to the sum of the depths of the two copy unitaries which, in the worst case of a qubit copied p times, is bounded by $2\lceil \log_2(p) \rceil$, and by the depth of the parallelized $\hat{U}_j$, which is given by the maximum of the depths of the $\hat{U}_j$'s. □

If the number m of available ancilla qubits is smaller than k, one can gather the $\hat{U}_j$'s into $p' < p$ groups and implement the groups in parallel to each other. Figure 2 gives an example of a diagonal unitary composed of $p = 4$ efficiently implementable diagonal operators which are implementable sequentially or in parallel depending on the number of ancilla qubits available. For $m = n$, one can implement half of the $\hat{U}_j$'s in parallel to the others, reducing at best by a factor of 2 the depth of the circuit (but doubling its width). The following Theorem III.2 gives an upper bound for the depth of the adjustable-depth protocol which decreases with the number of ancilla qubits, at the cost of performing the $\widehat{\text{copy}}$ and $\widehat{\text{copy}}^{-1}$ unitaries.

Theorem III.2 (Adjustable-depth). *Let $m \geq n$ be the number of ancilla qubits and $m' = \lceil m/n \rceil$. The diagonal unitary $\hat{U} = \prod_{j=0}^{p-1} \hat{U}_j$ can be implemented with a quantum circuit of size $s + 2m$ and depth d bounded as:*

$$d \leq \sum_{j=0}^{\lceil p/m' \rceil - 1} d_j + 2\lceil \log_2(m') \rceil \leq \lceil \frac{p}{m'} \rceil \max_j(d_j) + 2\lceil \log_2(m') \rceil, \tag{14}$$

where the $\hat{U}_j$ are ordered to have decreasing depth $d_0 \geq d_1 \geq ... \geq d_{p-1}$.

Proof. Similarly to the fully parallelized theorem, the size of the quantum circuit is given by the size s of the circuits implementing the $\hat{U}_j$'s added to the number $2m$ of controlled-NOT gates needed to perform the $\widehat{\text{copy}}$ and $\widehat{\text{copy}}^{-1}$ unitaries. The first inequality is given by the worst case scenario where the $\hat{U}_j$'s act non-trivially on n qubits, in which case one needs at least $m \geq n$ ancilla qubits to reduce the depth of the quantum circuit. In this situation, the number of groups that can be implemented in parallel is $m' = \lfloor m/n \rfloor + 1$ (the $+1$ term comes from the n initial qubits, *i.e.*, $m' = \lceil m/n \rceil$). It is always possible to create m' random groups of at most $\lceil p/m' \rceil$ operators. In the worst case, the maximum depth of the groups is the sum from $j = 0$ to $\lceil p/m' \rceil - 1$ of the maximum depths of the $\hat{U}_j$'s: $\sum_{j=0}^{\lceil p/m' \rceil - 1} d_j$. Since the depths of the two copy steps $2\lceil \log_2(m') \rceil$, one thus arrives at the first inequality. The second inequality trivially comes from the fact that $\forall j$, $d_j \leq \max_{j'}(d_{j'})$. □

This adjustable-depth protocol does not deal with finding an optimal strategy to group the $\hat{U}_j$'s. Such a strategy needs to consider the depth d_j of each operator and to group them in order to minimize the maximum depth of the groups. Finding this optimal solution or a "good" solution can take some classical time and it is in general a difficult combinatorial problem (even though some tailored algorithmic techniques may exist) [43]. We do not look into this here, however, the upper bound is reached for a diagonal unitary made only of sequential operators (3) acting non-trivially on n qubits with identical depth. For a diagonal unitary given by a Walsh-Hadamard decomposition, it is possible to group the unitaries with distinct support first, before parallelizing them.

Overall, the parallelization approach reduces the depth of the quantum circuit while increasing the width, making it possible to adjust the quantum circuit to the constraints of the hardware, *e.g.*, the decoherence time and the number of available qubits. Numerical results illustrating the trade-off between depth and width are given in the application section in Fig. 7 for the quantum state preparation of Gaussian states.

IV. ADJUSTABLE-DEPTH FRAMEWORK USING APPROXIMATIONS

This section presents two efficient methods to approximate diagonal unitaries $\hat{U}_f = e^{i\hat{f}}$ which depend on smooth or, at least, once differentiable functions f. The first method transforms any exact quantum circuit for diagonal unitaries into an efficient approximate one. The second method computes a good approximation of f using sparse Walsh Series. Both methods exploit the specific structure of $\hat{U}_f$ to get an efficient approximate quantum circuit, while exact methods suffer from exponential scalings either in terms of depth or width [16]. These two approximating methods have efficient scalings with a depth and size asymptotically independent of the number of qubits n. By comparison, approximate quantum compiling [15] does not provide any guarantees to reach a given accuracy in the large n limit.

Let f be a smooth, at least once differentiable, function defined on $[0, 1]$, $\hat{U}_{f,n} = e^{i\hat{f}_n} = \sum_{x=0}^{N-1} f(x/N) |x\rangle \langle x|$ be the target unitary acting on n qubits and $\hat{U}_{f,m} = \sum_{x=0}^{M-1} f(x/M) |x\rangle \langle x|$ be the restriction $\hat{U}_{f,n}$ to $m < n$ qubits (with $M = 2^m$). The following theorem states that implementing $\hat{U}_{f,m}$ using an exact method gives an efficient quantum circuit approximating the target unitary $\hat{U}_{f,n}$ up to an error $\epsilon > 0$ in spectral norm:

Theorem IV.1. *Given a quantum circuit implementing exactly any n-qubit diagonal unitary with size $s(n)$, depth $d(n)$ and width $w(n)$, one can ϵ-approximate $\hat{U}_{f,n}$ with a quantum circuit of size $s(m)$, depth $d(m)$ and width $w(m)$ with $m = \lceil \log_2(||f'||_{\infty,[0,1]}/\epsilon) \rceil$.*

Proof. The proof is based on the Walsh-Hadamard representation of the target unitary, independently of its method of implementation. First, we prove that implementing exactly the diagonal unitary $\hat{U}_{f,m}$ on $m = \lceil \log_2(||f'||_{\infty,[0,1]}/\epsilon) \rceil$ qubits gives an ϵ-approximation of $\hat{U}_{f,n}$ in spectral norm. Define $S_{f,M} = \sum_{j=0}^{M-1} a_j^f w_j$ as the M-Walsh Series of f, with $a_j^f = \frac{1}{M} \sum_{x=0}^{M-1} f(x/M) w_j(x/M)$ the j-th Walsh coefficients associated to the function f and $M = 2^m$. The M-Walsh Series $S_{f,M}$ of f is a piecewise function taking at most M different values and defined by $\forall x \in [k/M, (k+1)/M[$, $S_{f,M}(x) = S_{f,M}(k/2^m) = f(k/2^m)$ with $k \in \{0, ..., M-1\}$ (see Lemma 1.1 of [10]). Thus, the n-qubit operator $\hat{S}_{f,M} = \sum_{x=0}^{N-1} S_{f,M}(x/N) |x\rangle \langle x|$ has at most M distinct eigenvalues and can be rewritten as:

$$\begin{aligned}\hat{S}_{f,M} &= \sum_{x_0,...,x_{n-1}=0}^{1} S_{f,M}\left(\sum_{i=0}^{n-1} x_i/2^{i+1}\right)|x_0,...,x_{n-1}\rangle\langle x_0,...,x_{n-1}| \\ &= \sum_{x_0,...,x_{m-1}=0}^{1} S_{f,M}\left(\sum_{i=0}^{m-1} x_i/2^{i+1}\right)|x_0,...,x_{m-1}\rangle\langle x_0,...,x_{m-1}| \\ &\otimes \sum_{x_m,...,x_{n-1}=0}^{1} |x_m,...,x_{n-1}\rangle\langle x_m,...,x_{n-1}| \\ &= \hat{f}_m \otimes \hat{I}_2^{\otimes n-m},\end{aligned} \tag{15}$$

where $\hat{f}_m = \sum_{x=0}^{M-1} f(x/M)\,|x\rangle\langle x|$ is an m-qubit operator.

Thus, $\hat{U}_{f,m} \otimes \hat{I}_2^{\otimes n-m} = e^{i\hat{f}_m \otimes \hat{I}_2^{\otimes n-m}} = e^{i\hat{S}_{f,M}} = \hat{U}_{S_{f,M},n}$, which is the diagonal unitary with eigenvalues associated to the M-Walsh series of f. Furthermore, M-Walsh Series approximate the function itself up to an error depending on the maximum value of the derivative of f and decreasing linearly with M [14, 40]: $||f - S_{f,M}||_{\infty,[0,1]} \leq \frac{||f'||_{\infty,[0,1]}}{2^m}$. In terms of spectral norm $||.||_2$, this implies that $\hat{U}_{S_{M,f},n}$ is an approximation of the target diagonal unitary $\hat{U}_{f,n}$:

$$||\hat{U}_{f,n} - \hat{U}_{S_{f,M},n}||_2 \leq \frac{||f'||_{\infty,[0,1]}}{2^m}. \tag{16}$$

Finally, one can choose $m = \lceil \log_2(||f'||_{\infty,[0,1]}/\epsilon)\rceil$ of qubits to implement the target unitary $\hat{U}_{f,n}$ up to a given error $\epsilon > 0$ in spectral norm with a quantum circuit of size $s(m)$, depth $d(m)$ and width $w(m)$. □

In the particular case $m \geq n$, one implements exactly the diagonal unitary on n qubits and the error vanishes. On the contrary, for a given $\epsilon > 0$, this method gives a quantum circuit which approximates diagonal unitaries with complexities asymptotically independent of the number of qubits n.

A second way to approximate efficiently a diagonal unitary defined through a given function f is to consider a sparse Walsh series. For instance, by retaining only the largest coefficients in the M-Walsh Series, one can reach a given infidelity with quantum circuits of smaller size and depth [24]. More generally, the problem of finding the "best" Walsh series with the smallest number of terms approximating up to a given $\epsilon > 0$ a function f is called the minimax problem [44]. While the minimax problem has no known canonical solutions, sparse Walsh series appear numerically as the most efficient method to approximate smooth, at least differentiable, functions. Numerical results illustrating the trade-off between depth and accuracy are given in the application section, for the quantum state preparation of Gaussian states, see Fig. 8. The trade-off is presented for the approximate method using an M-Walsh series and for a sparse Walsh-Series.

V. IMPLEMENTATION OF NON-UNITARY DIAGONAL OPERATORS

In this section, we outline first the process of implementing a non-unitary diagonal operator from two diagonal unitaries using one ancilla qubit. Then, the methods introduced previously are adapted to deliver ways of implementing non-unitary diagonal operators with adjustable-depth quantum circuits.

Consider the non-unitary diagonal operator $\hat{D} = \mathrm{diag}(d_0,...,d_{N-1})$, $N = 2^n$ with real eigenvalues $\{d_x\}_{x=0}^{N-1}$. The block encoding approach consists in implementing $\hat{D}$ in a larger unitary operator $\hat{U}$ using additional ancilla qubits as:

$$\hat{U}_D = \begin{pmatrix} \hat{D} & * \\ * & * \end{pmatrix}, \tag{17}$$

where the other block "$*$" are chosen to ensure that $\hat{U}_b$ is unitary.

More formally, $\hat{U}_D$ is said to be a (α, m, ϵ)-block encoding of $\hat{D}$ with $\alpha > 0$ and $\epsilon > 0$ if

$$\|\frac{1}{\alpha}\hat{D} - (\langle 0|^{\otimes m} \otimes \hat{I}_N)\hat{U}_D(|0\rangle^{\otimes m} \otimes \hat{I}_N)\|_2 \leq \epsilon, \tag{18}$$

in spectral norm. Note that α is often the largest eigenvalue of $\hat{D}$ and m is the number of additional qubits used to block-encode $\hat{D}$ into a unitary operator. In the case where the block-encoding is exact, we note that $\hat{U}_D$ is a (α, m)-block encoding of $\hat{D}$.

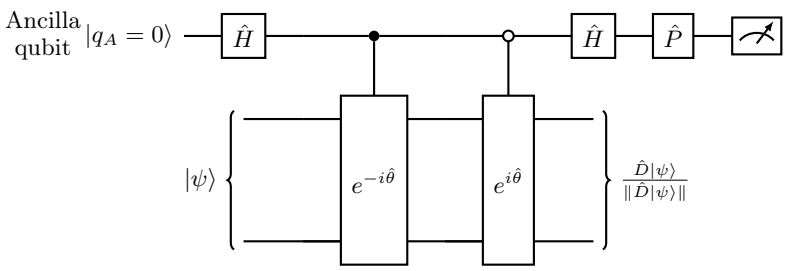

Figure 3: Quantum circuit to block-encode a non-unitary diagonal operator $\hat{D} = \mathrm{diag}(d_0, ..., d_{2^n-1})$ using $\hat{\theta} = \arcsin(\hat{D}/d_{\max})$, with $d_{\max} = \max_{i=0,...,2^n-1} |d_i|$.

Figure 3 depicts the sequence of unitary operations needed to block-encode a non-unitary diagonal operator using two controlled diagonal unitaries and one ancilla qubit without using parallelization protocols. From a qubit state $|\psi\rangle$ and an ancilla qubit $|0\rangle$, one obtains:

$$\begin{aligned} |\psi\rangle|0\rangle \xrightarrow{\hat{H}} |\psi\rangle \frac{|0\rangle + |1\rangle}{\sqrt{2}} \xrightarrow{e^{i\hat{\theta}\otimes\hat{Z}}} \frac{e^{i\hat{\theta}}|\psi\rangle|0\rangle + e^{-i\hat{\theta}}|\psi\rangle|1\rangle}{\sqrt{2}} \\ \xrightarrow{\hat{P}\hat{H}} -i\cos(\hat{\theta})|\psi\rangle|0\rangle + \sin(\hat{\theta})|\psi\rangle|1\rangle . \end{aligned} \tag{19}$$

The gate $\hat{P} = \begin{pmatrix} 1 & 0 \\ 0 & -i \end{pmatrix}$ is optional and only serves to suppress a factor i in the final state. The operator $\hat{\theta}$ encodes the eigenvalues of $\hat{D}$ as $\hat{\theta} = \arcsin(\hat{D}/d_{\max})$, where $d_{\max} = \max_{x=0,...,N-1} |d_x|$ is the largest eigenvalue of $\hat{D}$. Without measuring the ancilla, this quantum circuit is an exact $(d_{\max}, 1)$ block-encoding of $\hat{D}$. This block-encoding is sufficient for some applications but one may need to output only the state proportionate to $\hat{D}|\psi\rangle$ by measuring the ancilla qubit in state $|1\rangle$:

$$|\phi\rangle = -i\sqrt{\hat{I} - \left(\frac{\hat{D}}{d_{\max}}\right)^2}|\psi\rangle|0\rangle + \frac{\hat{D}}{d_{\max}}|\psi\rangle|1\rangle \xrightarrow{\text{if } |q_A\rangle=|1\rangle} \frac{\hat{D}|\psi\rangle}{\|\hat{D}|\psi\rangle\|}, \tag{20}$$

with a probability of success given by

$$\mathbb{P}(1) = \|\frac{\hat{D}}{d_{\max}}|\psi\rangle\|^2 = |\langle\psi|(\frac{\hat{D}}{d_{\max}})^2|\psi\rangle| = \sum_{x=0}^{2^n-1} d_x^2|\psi_x|^2/d_{\max}^2. \tag{21}$$

The specific structures of $\hat{D}$ and $|\psi\rangle$ directly affect the probability of success. When the diagonal operator $\hat{D}$ and the qubit state $|\psi\rangle$ depend on continuous real-valued functions f and g defined on $[0, 1]$, *i.e.*, $\hat{D} = \sum_{x=0}^{N-1} f(x/N)|x\rangle\langle x|$ and $|\psi\rangle = \frac{1}{||g||_{2,N}} \sum_{x=0}^{N-1} g(x/N)|x\rangle$, the probability of success tends to an n-independent limit, enabling an efficient implementation:

$$\mathbb{P}(1) = \frac{\|fg\|_{2,[0,1]}^2}{\|g\|_{2,[0,1]}^2 \|f\|_\infty^2} = \Theta(1), \tag{22}$$

with $\|fg\|_{2,[0,1]}^2 = \int_0^1 f(x)^2 g(x)^2 \mathrm{d}x$, $\|g\|_{2,[0,1]}^2 = \int_0^1 g(x)^2 \mathrm{d}x$ and $\|f\|_\infty = \sup_{x\in[0,1]} |f(x)|$. The notation $\mathbb{P}(1) = \Theta(1)$ means that the probability of success is bounded by two constants $0 < C_1 \leq \mathbb{P}(1) \leq C_2 \leq 1$ in the large n limit.

When the diagonal operator $\hat{D}$ is s-sparse, i.e., $\hat{D} = \sum_{x\in S} d_x|x\rangle\langle x|$, with $S \subseteq \{0, 1, ..., N-1\}$ and $|S| = s$, and the qubit state depends on a continuous function, *i.e.*, $|\psi\rangle = \frac{1}{||g||_{2,N}} \sum_{x=0}^{N-1} g(x/N)|x\rangle$, the probability of success diminishes exponentially with the number of qubits:

$$\mathbb{P}(1) = \frac{1}{||g||_{2,N}} \sum_{x\in S} \frac{d_x^2}{d_{\max}^2} g(x/N)^2 = \Theta(1/N), \tag{23}$$

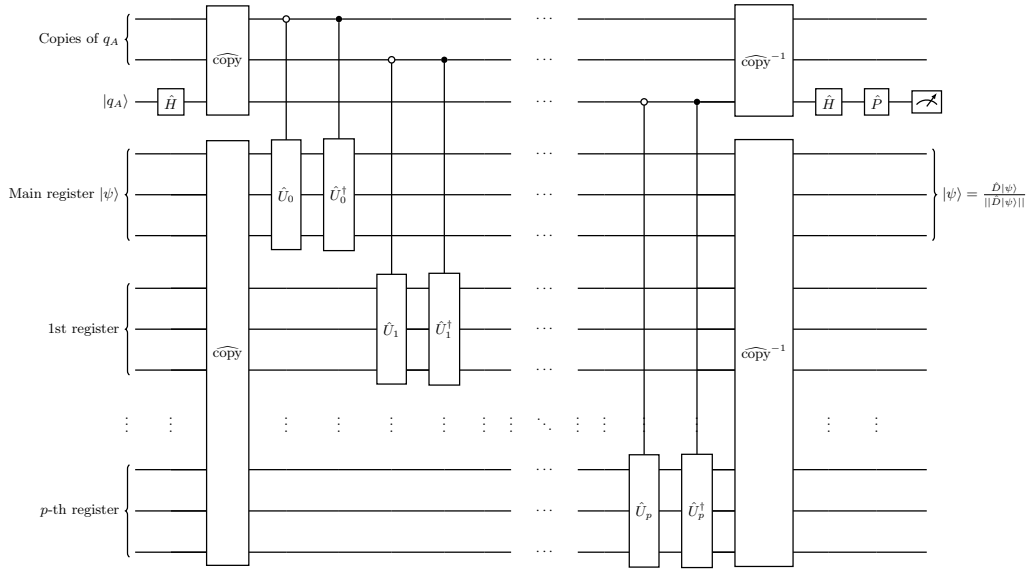

Figure 4: Adjustable-depth quantum circuit for the implementation of non-unitary diagonal operator $\hat{D}$. The main register $|\psi\rangle$ is composed of n qubits. The ancilla qubit is $|q_A\rangle$. The upper register is composed of ancilla qubits that store the bit value of $|q_A\rangle$, it is used to perform several controls simultaneously. The p ancilla registers below the main register are used to parallelize the implementation of $\hat{U}$. If the output of the measurement of $|q_A\rangle$ is 0, the target state has been successfuly implemented.

with $N = 2^n$. In the other cases, for instance for a sparse qubit state and a non-sparse diagonal operator, no general formula garantees a significant probability of success in the large n limit. The probability of success can also be improved using amplitude amplification as discussed in Section VII.

Now, consider the problem of implementing a non-unitary diagonal operator $\hat{D}$ using an adjustable-depth quantum circuit associated to a Walsh or a sequential decomposition of $\hat{U} = e^{i \arcsin(\hat{D}/d_{\max})} = \prod_{j=0}^{p-1} \hat{U}_j$. In Fig. 3, the controls of $\hat{U}$ and $\hat{U}^\dagger$ do not preserve the adjustable-depth feature of the methods presented in the previous sections since one ancilla qubit cannot control several blocks at the same time. In order to overcome this issue, one can use p copies of the ancilla qubit $|q_A\rangle$ to control p blocks in parallel at the same time. That is why, one needs to prepare a copy register of $|q_A\rangle$ before the control of the diagonal operators. The following theorem states and proves that Fig. 4 represents an adjustable-depth quantum circuit for a block-encoding of $\hat{D}$ up to the measurement of the ancilla qubit $|q_A\rangle$.

Theorem V.1 (Block-encoding of $\hat{D}$)**.** *Let $\hat{D}$ be a non-unitary diagonal operator and $\hat{U} = e^{i \arcsin(\hat{D}/(\alpha d_{\max}))}$ with $\alpha \geq 1$. Suppose there is a decomposition $\tilde{U} = \prod_{j=0}^{p-1} \hat{U}_j$ approximating $\hat{U}$ up to an error $\epsilon > 0$ in spectral norm, where each $\hat{U}_j$ is a diagonal unitary acting on $k_j \leq n$ qubits. Then, the quantum circuit defined Fig. 4[4] is an $(\alpha d_{\max}, k + p, \epsilon)$-block-encoding of $\hat{D}$, where $k = \sum_{j=1}^{p-1} k_j$.*

The proof is presented in Appendix B.

In terms of complexity, copying the ancilla qubit $|q_A\rangle$ always requires less controlled-not gates than copying the main register. Additionally, controlling an operator $\hat{U}_j$ has the same asymptotic cost than implementing $\hat{U}_j$: remark that for any operator of the form $\hat{V}^{-1}\hat{U}\hat{V}$, the controlled operation is given by $C(\hat{V}^{-1}\hat{U}\hat{V}) = \hat{V}^{-1}C(\hat{U})\hat{V}$. Therefore, if $\hat{U}_j$ is a sequential operator given by Eq. (3), one only needs to transform the $(n-1)$-controlled phase gate into a n-controlled phase gate and if $\hat{U}_j$ is a Walsh-Hadamard operator given by Eq. (8), one only needs to control the $\hat{R}_Z$ gates. As a consequence, the asymptotic cost in terms of size and depth of implementing a diagonal operator $\hat{D}$ is directly given by the cost of implementing the unitary $e^{i \arcsin(\hat{D}/d_{\max})}$ up to the same error.

[4] without the measurement.

VI. EFFICIENT-ADJUSTABLE-DEPTH METHODS

In this section, we combine the previous theorems to obtain the overall complexities for the efficient-adjustable-depth quantum circuits for unitary and non-unitary diagonal operators without exponential scalings.

A. Unitary and non-unitary diagonal operators defined in terms of smooth functions

Consider a diagonal unitary $\hat{U} = \sum_{x=0}^{N-1} e^{if(x/N)} |x\rangle \langle x|$ defined in terms of a smooth, at least differentiable, function f. For each decomposition of $\hat{U}$, sequential or Walsh-Hadamard, the combination of the adjustable-depth Theorem III.2 or the fully parallelized Theorem III.1 gives three possibilities: without ancilla, adjustable-depth with a given number $p > 0$ of ancilla and fully parallelized. Then, the efficiency is achieved by approximating the diagonal unitary up to an error $\epsilon > 0$ using Theorem IV.1 and the associated method.

For a non-unitary diagonal operator $\hat{D} = \sum_{x=0}^{N-1} f(x/N) |x\rangle \langle x|$ defined in terms of a smooth, at least differentiable, function f, the previous section guarantees that the asymptotic cost of implementing $\hat{D}$ is given by the cost of implementing $\hat{U} = e^{i \arcsin(\hat{D}/(\alpha d_{\max}))}$. The parameter $\alpha > 1$ has been introduced to ensure that the function $\arcsin(f/(\alpha \|f\|_\infty))$ is differentiable[5] and that the approximation Theorem IV.1 can be applied.

The complexities for each of these methods are presented in Table I and proven in Appendix C.

Efficient quantum circuit for diagonal operators depending on differentiable functions			
Method	Depth	Size	Ancilla
Sequential	$\tilde{\mathcal{O}}(1/\epsilon)$	$\tilde{\mathcal{O}}(1/\epsilon)$	0 or 1
Sequential adjustable-depth	$\tilde{\mathcal{O}}(\frac{1}{\epsilon m} + \log(m))$	$\tilde{\mathcal{O}}(1/\epsilon + m)$	$m = \Omega(\log(1/\epsilon))$
Sequential fully parallelized	$\mathcal{O}(\log(1/\epsilon))$	$\tilde{\mathcal{O}}(1/\epsilon)$	$\tilde{\mathcal{O}}(1/\epsilon)$
Walsh-Hadamard	$\mathcal{O}(1/\epsilon)$	$\mathcal{O}(1/\epsilon)$	0 or 1
Walsh-Hadamard adjustable-depth	$\tilde{\mathcal{O}}(\frac{1}{\epsilon m} + \log(m))$	$\tilde{\mathcal{O}}(1/\epsilon + m)$	$m = \Omega(\log(1/\epsilon))$
Walsh-Hadamard fully parallelized	$\mathcal{O}(\log(1/\epsilon))$	$\tilde{\mathcal{O}}(1/\epsilon)$	$\tilde{\mathcal{O}}(1/\epsilon)$
Walsh-recursive	$\mathcal{O}(\frac{1}{\epsilon \log(1/\epsilon)})$	$\mathcal{O}(1/\epsilon)$	0 or $\mathcal{O}(\log(1/\epsilon))$
Walsh-optimized adjustable-depth	$\mathcal{O}(\frac{1}{\epsilon m} + \log m)$	$\tilde{\mathcal{O}}(1/\epsilon + m)$	$m = \Omega(\log(1/\epsilon))$
Walsh-optimized fully parallelized	$\mathcal{O}(\log(1/\epsilon))$	$\mathcal{O}(1/\epsilon)$	$\mathcal{O}(\frac{1}{\epsilon \log(1/\epsilon)})$

Table I: Table of depth, size and m-ancilla qubits for different efficient quantum circuits to implement, up to an error $\epsilon > 0$ in spectral norm, n-qubit unitary and non-unitary diagonal operators associated to real-valued function f with bounded first derivative. All the scalings are proven in Appendix C 2 a for diagonal unitaries and in Appendix C 2 b for non-unitary diagonal operators. The notation $\tilde{\mathcal{O}}(\cdot)$ stands for $\mathcal{O}(\cdot)$ up to a polylogarithmic factor.

The last three lines of the Table I are obtained by using the approximation Theorem IV.1 and the block-encoding Theorem V.1 with the exact methods developed in [16]. When no ancilla qubits are available, they provide a recursive scheme for diagonal unitaries with a complexity which is optimal in terms of depth as $\mathcal{O}(2^n/n)$. However, the size is not optimal and is only bounded by $2^{n+3} + \mathcal{O}(\log(n^2/\log(n)))$ (Lemma 24 of [16]). A second method uses a number $2n \leq m \leq 2^n/n$ of ancilla qubits to parallelize the Walsh operators in an order which optimizes the number of CNOT gates, achieving a depth $\mathcal{O}(2^n/m + \log(m))$ and a size $\mathcal{O}(2^n)$ (Lemma 10 in [16]).

B. Sparse unitary and non-unitary diagonal operators

Sparsity makes diagonal operators efficiently implementable. Sparse diagonal operators either possess s eigenvalues different from unity, or they are made of a product of s Walsh operators $\hat{W}_j$ defined in Eq. (7). Using the full parallelization Theorem III.1 and the adjustable-depth Theorem III.2, we summarize all the complexities to implement sparse diagonal operators be it unitary or non-unitary in Table II.

Sparse Walsh-Hadamard decompositions are particularly efficient to implement diagonal unitaries depending on real-valued functions with bounded first derivative see Fig. 7 and Fig. 8 or for bounded piecewise continuous functions

[5] α can be taken constant equal to 1.1 with the only backward of reducing the probability of success by a factor α^2.

Efficient quantum circuit for sparse diagonal operators			
Method	Depth	Size	Ancilla
Sequential no ancilla exact	$\mathcal{O}(ns)$	$\mathcal{O}(ns)$	0 or 1
Sequential no ancilla approximated	$\mathcal{O}(s\log^3(n)\log(\frac{s}{\epsilon}))$	$\mathcal{O}(ns\log^4(n)\log(\frac{s}{\epsilon}))$	0 or 1
Sequential one ancilla exact	$\mathcal{O}(s\log^3(n))$	$\mathcal{O}(ns\log^4(n))$	1 or 2
Sequential Adjustable-depth	$\mathcal{O}(\frac{s\log^3(n)}{(m/n)}+\log(m/n))$	$\mathcal{O}(ns\log^4(n)+m)$	$m=\Omega(n)$
Sequential fully parallelized	$\mathcal{O}(\log(s)+\log^3(n))$	$\mathcal{O}(ns)$	$\mathcal{O}(ns)$
Walsh no ancilla	$\mathcal{O}(sk)$	$\mathcal{O}(sk)$	0 or 1
Walsh Adjustable-depth	$\mathcal{O}(\frac{sk}{(m/k)}+\log(m/k))$	$\mathcal{O}(sk+m)$	$m=\Omega(k)$
Walsh fully parallelized	$\mathcal{O}(k+\log(s))$	$\mathcal{O}(sk)$	$\mathcal{O}(sk)$

Table II: Table of depth, size and m-ancilla qubits for different quantum circuits implementing unitary or non-unitary n-qubit diagonal operators with sparse decompositions. s is the number of sequential or Walsh-Hadamard operators in the decomposition of the diagonal operator. All scalings are proven in Appendix C 3 a for diagonal unitaries and Appendix C 3 b for non-unitary diagonal operators.

(see Appendix A 3 in [10]). In particular, to reach a given accuracy, sparse Walsh-Hadamard decompositions are often more efficient than M-Walsh series (presented in Table I) when the most significant Walsh coefficients are implemented before the others.

VII. AMPLITUDE AMPLIFICATION AND NON-DESTRUCTIVE REPEAT-UNTIL-SUCCESS PROTOCOLS

Amplitude amplification enables to increase the probability of success of performing a non-unitary diagonal operator $\hat{D}$ on a given state $|\psi\rangle$. When the cost of preparing the state $|\psi\rangle$ from $|0\rangle^{\otimes n}$ is not prohibitively high, one can perform the usual amplitude amplification protocol [34], otherwise one can perform either the oblivious amplitude amplification protocol [45] or a repeat-until-success scheme where, at each failure, the "wrong" state is corrected to become the "good" state up to a new probability of success.

The amplitude amplification protocol increases the probability of success of measuring the ancilla qubit in state $|1\rangle$ by making repeated reflexions composed of the unitary $\hat{U}_\psi$ preparing the state $|\psi\rangle$ from $|0\rangle^{\otimes n}$ and of the unitary $\hat{U}_D$ which block-encodes $\hat{D}$. Considering the final state $|\phi\rangle = \hat{U}_D|\psi\rangle\otimes|0\rangle_{q_A} = \hat{U}_D(\hat{U}_\psi\otimes\hat{I}_2)|0\rangle^{\otimes n}\otimes|0\rangle_{q_A}$ defined in Eq. (20), the amplitude amplification protocol is composed of $\hat{U}_\phi = \hat{I}-2|\phi\rangle\langle\phi|$ and $\hat{U}_P = -(\hat{I}-2\hat{P})$ with $\hat{P}$ the projector on the good subspace: $\hat{P} = \hat{I}_2\otimes\cdots\otimes\hat{I}_2\otimes|1\rangle_{q_A}\langle 1|_{q_A}$. The unitary $\hat{U}_\phi$ can be rewritten as $\hat{U}_\phi = {\hat{U}_D}^{-1}({\hat{U}_\psi}^{-1}\otimes\hat{I}_2)\Lambda_{1,\ldots,n\to q_A}(-\hat{Z})(\hat{U}_\psi\otimes\hat{I}_2)\hat{U}_D$ with $\Lambda_{1,\ldots,n\to q_A}(-\hat{Z})$ an anti-controlled minus $\hat{Z}$-Pauli gate and $\hat{U}_P = \hat{I}_2\otimes\cdots\otimes\hat{I}_2\otimes(-\hat{Z})$. The scheme of the quantum circuit associated to a step of amplitude amplification is shown on Fig. 5. The number k of steps needed to perform the amplitude amplification is given by $k=\lfloor\pi/(4\beta)\rfloor$ with $\beta=\arcsin(\sqrt{\mathbb{P}(1)})$. The number k can be estimated from $\mathbb{P}(1)$ using one of the formula (21, 22, 23). In the case where the diagonal unitary and the qubit state are associated to continuous functions, the number of amplitude amplification steps needed to reach $\mathbb{P}(1)\simeq 1$ is asymptotically independent of the number of qubits n. The amplitude amplification protocol is illustrated Fig. 6 for the quantum state preparation of Gaussian states proving numerically that the probability of success reaches $\mathbb{P}(1)\simeq 1$ in a number $k=\lfloor\pi/(4\beta)\rfloor$ of steps.

Conversely, a repeat-until-success protocol is relevant when the cost of preparing the state $|\psi\rangle$ from $|0\rangle^{\otimes n}$ is prohibitively high. Here we propose a non-destructive repeat-until-success protocol: consider the diagonal unitary $e^{i\arcsin(\hat{D}/(\alpha d_{\max}))}$ where a constant parameter $\alpha>1$ ensures that the failure operator $\hat{D}'=\sqrt{\hat{I}-(\hat{D}/(\alpha d_{\max}))^2}$ is invertible. If a failure occurs, one can implement the non-unitary diagonal operator $\hat{D}({\hat{D}'}^{-1})$ using the unitary operator $e^{i\arcsin(\hat{D}({\hat{D}'}^{-1})/(\alpha d'_{\max}))}$ with $d'_{\max}=\max_{i=0,\ldots,N-1}|d_i|/(\sqrt{1-(d_i/(\alpha d_{\max}))^2})$. Applying this operator on the state $\frac{\hat{D}'|\psi\rangle}{\|\hat{D}'|\psi\rangle\|}$ leads to the target state $\frac{\hat{D}|\psi\rangle}{\|\hat{D}|\psi\rangle\|}$ with a probability of success $\mathbb{P}'(1)=\frac{\|\hat{D}|\psi\rangle\|^2}{\alpha^2 d'^2_{\max}\|\hat{D}'|\psi\rangle\|^2}$. This protocol can be implemented until a success is reached, avoiding destroying and preparing again the state $|\psi\rangle$. Figure 6 illustrates this protocol for the quantum state preparation of Gaussian states, showing that the probability of success after k failures reaches a constant value. Therefore, the average number of failures is finite and can be estimated numerically.

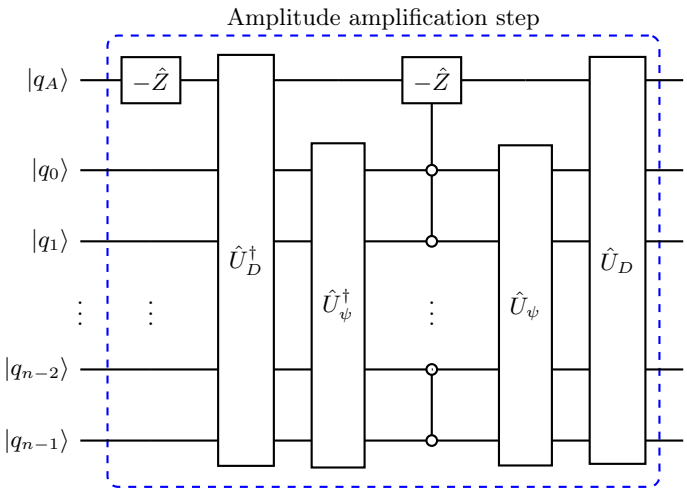

Figure 5: Quantum circuit for amplitude amplification. Note that for QSP, operator $\hat{U}_\psi$ is an Hadamard tower on the register $|q\rangle$, *i.e.*, $\hat{U}_\psi = \hat{H}^{\otimes n}$. Operator $\hat{U}_D$ is the block-encoding of the non-unitary diagonal operator $\hat{D}$.

VIII. APPLICATIONS

We apply the methods previously presented to two problems. The first one is the quantum state preparation problem for which we improve the current state of the art of preparing qubit-state depending on continuous, differentiable function f and we provide new methods with space-time-accuracy trade-offs. The second major application is the resolution of partial differential equation where the associated evolution operator is implementable using diagonal operators and the Quantum Fourier Transform. As an example, we solve the diffusion equation, also called heat equation, which possesses a non-unitary evolution.

A. Quantum State Preparation

The quantum state preparation problem consists in loading a set of classical data $\{y_x\}_{x=0}^{2^n-1}$ into a n-qubit state $|\psi\rangle = \frac{1}{\mathcal{N}}\sum_{x=0}^{2^n-1} y_x |x\rangle$ up to a normalization factor $\mathcal{N}$. In the following we provide an efficient quantum state preparation algorithm for n-qubit states depending on a differentiable function f is defined by:

$$|f\rangle = \frac{1}{||f||_{2,N}} \sum_{x=0}^{N-1} f(x/N)\,|x\rangle\,, \tag{24}$$

with $||f||_{2,N} = \sqrt{\sum_{x=0}^{N-1} |f(x/N)|^2}$ and $N = 2^n$. Now, consider the non-unitary operator $\hat{D} = \sum_{x=0}^{N-1} f(x/N)\,|x\rangle\langle x|$. The quantum state preparation can be performed by implementing $\hat{D}$ with one of the method previously presented Table I and applying $\hat{D}$ on the uniform superposition of all states $|s\rangle = \hat{H}^{\otimes n}|0\rangle^{\otimes n} = \frac{1}{\sqrt{N}}\sum_{x=0}^{N-1}|x\rangle$. The following Theorem summarizes this result:

Theorem VIII.1 (Quantum State Preparation)**.** *Let f be a real-valued differentiable function in $L^2([0,1])$, n the number of qubits and $\epsilon > 0$. There is a quantum circuit that efficiently prepare the state $|f\rangle$ defined in Eq. (24) up to an infidelity $\epsilon > 0$ using $p \geq 1$ ancilla qubits with a depth scaling at worst as $\mathcal{O}(1/(p\sqrt{\epsilon}) + \log(p))$, size $\mathcal{O}(n + 1/(p\sqrt{\epsilon}))$ and a probability of success $\mathbb{P}(1) = \|f\|^2_{2,[0,1]}/\|f\|^2_\infty = \Theta(1)$.*

Proof. This theorem is a direct consequence of the scaling given in Table I. The associated quantum circuits are the same as the one for the non-unitary diagonal operator. Only a one-depth layer of n Hadamard gates needs to be added to prepare the uniform superposition $|s\rangle = \frac{1}{\sqrt{N}}\sum_{x=0}^{N-1}|x\rangle$. The probability of success is given by Eq. (22) for the uniform function $g = 1$ and $\mathbb{P}(1) = \|f\|^2_{2,[0,1]}/\|f\|^2_\infty = \Theta(1)$. □

In the case of a small probability of success, one can still perform a constant number $k = \lfloor \pi/(4\arcsin(\sqrt{\mathbb{P}(1)})) \rfloor$ of amplitude amplification steps to reach $\mathbb{P}(1) \simeq 1$.

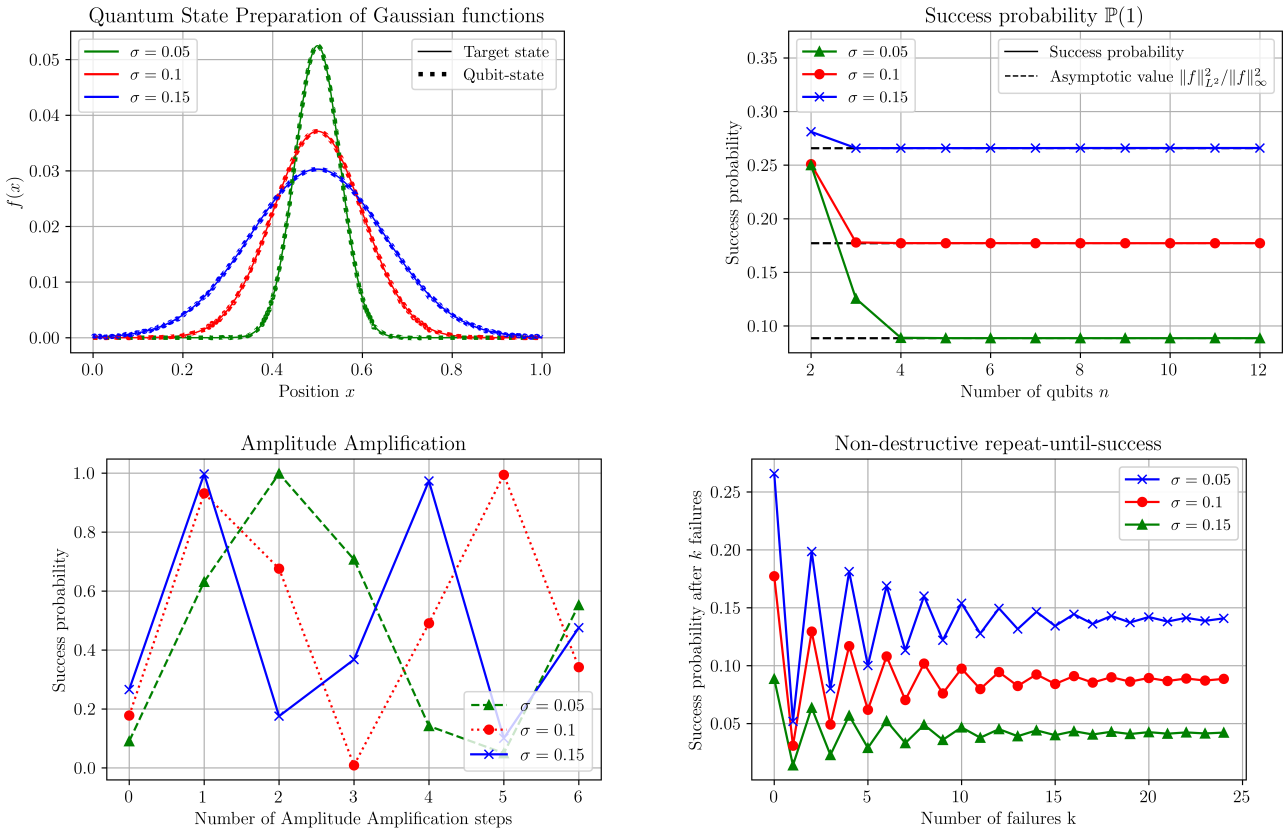

Figure 6: Quantum state preparation (upper-left) of Gaussian state $f(x) \propto e^{-0.5(x-0.5)^2/\sigma^2}$ for $\sigma = 0.05, 0.1$ and 0.15 on $n = 12$ qubits using respectively $s = 30, 45, 90$ of Walsh-Hadamard operators with 2-norm error between the qubit state and the target state of $0.0054, 0.0052, 0.0054$ respectively. Probability of success of the QSP protocol as a function of the number of qubits (upper right), as a function of the number of the number of amplitude amplification steps (lower left) and probability of success after k failures for the non-destructive repeat-until-success scheme (lower right). The numerical values of $\mathbb{P}(1)$ reach their theoretical asymptotic limit $\|f\|^2_{L^2}/\|f\|^2_\infty$ even for a small number of qubits and the number of amplitude amplification steps needed to reach the first pic of probability verify $k = \lfloor \pi/(4 \arcsin(\sqrt{\mathbb{P}(1)})) \rfloor = 1, 1, 2$ respectively. For the non-destructive repeat-until-success scheme, the average number of failures before reaching a success are $6, 10, 22$ respectively.

In Figure 6, the quantum state preparation of Gaussian states is performed and the probability of success is presented as a function of the number of qubits, showing an independent of n limit. The amplitude amplification protocol is also illustrated, proving numerically the efficiency of our method. The space-time trade-off is illustrated in Fig. 7 for an exact Walsh-Hadamard decomposition, an approximate one using a M-Walsh series and using a sparse Walsh series. The trade-off is particularly efficient for a small number of ancilla where one diminishes by several order the depth of the circuits. Note that increasing the number of ancilla qubits, increases the cost of the copy operations as well. As a consequence, for a large number of ancilla qubits, adding some ancilla qubits may not improve anymore the depth of the circuit. The trade-off between time and accuracy is illustrated in Fig. 8 for the quantum state preparation of Gaussian states. The infidelity of the prepared states decreases with the number of operators one implements as stated in the approximation Theorem IV.1. Furthermore, sparse Walsh-Hadamard decomposition allows to reach a given infidelity with smaller depth quantum circuits than M-Walsh Series.

In [10], n-qubit states are implemented using the Taylor expansion of the operator $\hat{I} - e^{i\hat{f}\epsilon_0} \simeq \hat{f}\epsilon_0$ with $\hat{f} = \sum_{x=0}^{2^n-1} f(x) |x\rangle \langle x|$. In comparison, the methods presented here implement directly $\hat{f}$ using $e^{\pm i \arcsin(\hat{f}/(\alpha f_{\max}))}$ avoiding an additional error coming from the Taylor expansion. Theorem VIII.1 also improves the probability of success which is scaling with ϵ in [10] while here it reaches a constant value.

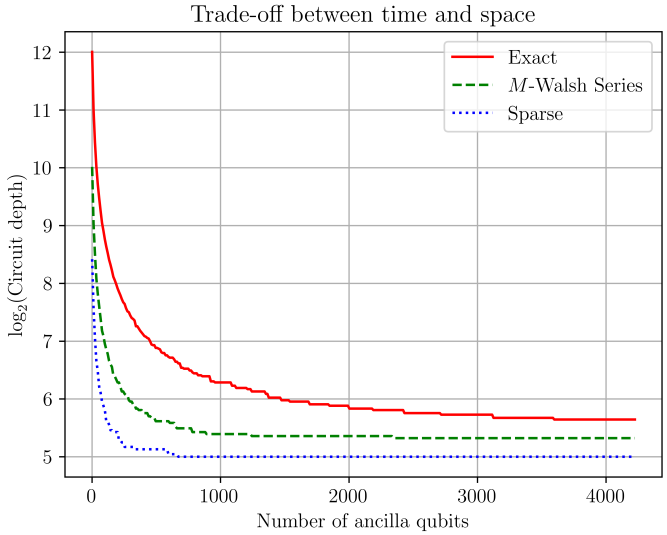

Figure 7: Depth as a function of the number of ancilla qubits for the quantum state preparation of a Gaussian state $f(x) \propto e^{-0.5(x-0.5)^2/\sigma^2}$ with $\sigma = 0.1$ on $n = 10$ qubits for three different Walsh-Hadamard decompositions: exact with 210 operators, approximate with 2^8 operators and an infidelity $1 - F = 5.96 \times 10^{-5}$ and sparse with 70 Walsh-Hadamard operators and an infidelity $1 - F = 6.06 \times 10^{-5}$.

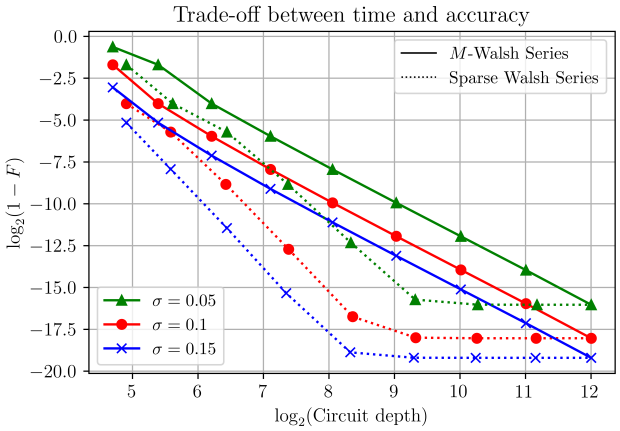

Figure 8: Infidelity as a function of the depth of the quantum circuit associated to the quantum state preparation of Gaussian states $f(x) \propto e^{-0.5(x-0.5)^2/\sigma^2}$ with $\sigma = 0.05, 0.1, 0.15$ on $n = 12$ qubits without ancilla qubits. The dots from upper left to lower right correspond to the M-Walsh Series with $M = 2^m$ Walsh-Hadamard operators (full lines) and to the sparse Walsh Series composed of the 2^m Walsh-Hadamard operators with the largest coefficients in absolute value with $m = 2, 3, \ldots, 10$.

B. Real-space simulation of the Heat equation

As an illustration of the methods presented in this article, we solve the following one-dimensional diffusion equation for a given initial condition f_0 defined on $[0, 1]$ and a constant diffusion coefficient κ:

$$\partial_t f = \kappa \partial_{xx} f, \tag{25}$$

where f is a function of the time $[0, T]$ and space $[0, 1]$. For simplicity, we consider periodic boundary conditions. The numerical resolution is based on a real-space simulation where the space $[0, 1]$ of the simulation is discretized into N equi-length intervals. The spatial derivatives are discretized using a central finite difference $\partial_x f \to \frac{f(x+\Delta x) - f(x-\Delta x)}{2\Delta x}$ with $\Delta x = 1/N$ the length of a space interval.

The real space encoding consists in approximating the solution f of Eq. (25) at the $N = 2^n$ spatial points $\mathcal{X}_n = \{0, 1/N, ..., (N-1)/N\}$ and at each time $t > 0$ by an n-qubit state $\left|\tilde{f}\right\rangle_t$ solution of the discretized diffusion equation:

$$\begin{aligned} \partial_t \left|\tilde{f}\right\rangle_t &= \kappa \frac{(\hat{S} - \hat{S}^\dagger)^2}{4\Delta x^2} \left|\tilde{f}\right\rangle_t, \\ \left|\tilde{f}\right\rangle_{t=0} &= |f_0\rangle \end{aligned} \tag{26}$$

where $\frac{(\hat{S}-\hat{S}^\dagger)^2}{4\Delta x^2}$ is the discretized Laplacian operator with $\hat{S} = \sum_{x\in\mathcal{X}_n} |x+1\rangle\langle x|$ the shift operator, also called increment operator[6]. The solution $\left|\tilde{f}\right\rangle_t$ is given by the evolution operator applied on the initial condition $|f_0\rangle$:

$$\left|\tilde{f}\right\rangle_t = e^{\frac{\kappa t}{4\Delta x^2}(\hat{S}-\hat{S}^\dagger)^2} |f_0\rangle . \tag{27}$$

The evolution operator of the diffusion equation can be diagonalized in Fourier space using the Quantum Fourier Transform [38, 46] thanks to the fact that the shift operators are circulant matrices which are diagonal in Fourier space [47]. The diagonal operator is non-unitary and depends on a smooth function $f(x) = e^{-\frac{\kappa}{\Delta x^2}\sin(2\pi x)^2 t}$, so the evolution is given by:

$$\left|\tilde{f}\right\rangle_t = \widehat{QFT}^{-1} e^{-\frac{\kappa t}{\Delta x^2}\sin(2\pi\hat{x})^2} \widehat{QFT} |f_0\rangle , \tag{28}$$

where $\hat{x} = \sum_{x\in\mathcal{X}_n} x|x\rangle\langle x|$ is the position operator.

The initial condition $|f_0\rangle$ is encoded in an n-qubit state thanks to the QSP protocol presented in the previous section (Theorem VIII.1) and the non-unitary diagonal operator is implemented using a sequential decomposition. Figure 9 presents an $n = 10$-qubit simulation where the initial Gaussian state $f_0(x) = e^{-0.5(x-0.5)^2/\sigma^2}$ with $\sigma = 0.1$ evolves through the heat equation. The amplitude of the numerical solution $\left|\tilde{f}\right\rangle_t$ is shown as a function of the position for different times t and compared with the analytical solution f of Eq. (25) computed as a Fourier Series:

$$f(x,t) = \sum_{q=0}^{+\infty} e^{-4\kappa q^2\pi^2 t}(\alpha_q \cos(2q\pi(x-0.5)) + \beta_q \sin(2q\pi(x-0.5))), \tag{29}$$

with the Fourier coefficients, for $q \geq 1$:

$$\begin{aligned} \alpha_q &= 2\int_0^1 f_0(x)\cos(2q\pi(x-0.5))\mathrm{d}x \\ \beta_q &= 2\int_0^1 f_0(x)\sin(2q\pi(x-0.5))\mathrm{d}x \end{aligned}, \tag{30}$$

and $\alpha_0 = \int_0^1 f_0(x)\mathrm{d}x$, $\beta_0 = 0$.

The number of sequential operators needed for the simulation is very small (14, 8, 6 and 1) due to the particular structure of the non-unitary evolution operator which is increasingly sparse with the time t (in Fourier space).

IX. CONCLUSION

In conclusion, efficient quantum circuits for unitary and non-unitary diagonal operator exist when the diagonal operators have some structure (depending on differentiable functions or sparse). The quantum circuits we introduced offer trade-offs between the depth, the number of ancilla qubits and the error between the implemented operator and the target one. These particular properties enable the user to design quantum circuits adjusted to the characteristic of the hardware, such as the coherence time of the qubits, the number of available qubits or the cumulative error per gate. Numerical evidence show that these results are generalizable to diagonal operators depending on bounded, continuous functions even when the function is not differentiable [10] but additional work remains to prove it. Also, current work investigates how to apply these quantum circuits to other partial differential equations.

[6] with $|n+1\rangle = |0\rangle$ due to the periodic boundary conditions.

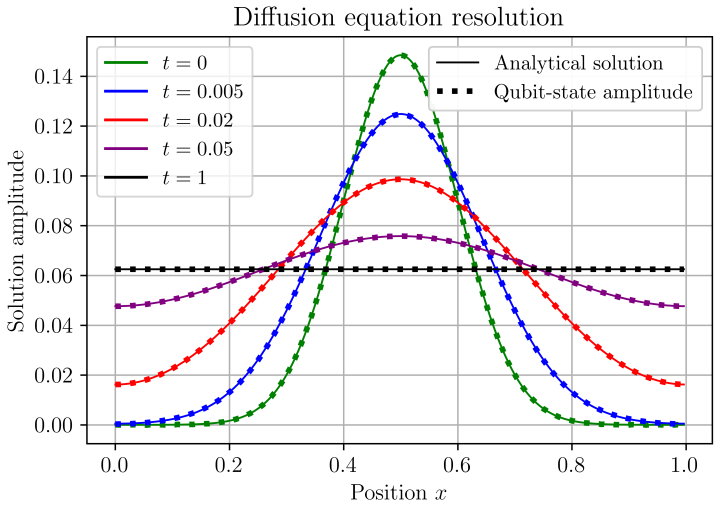

Figure 9: Resolution of the heat equation at given times t by computing the $n = 8$ qubit state given by Eq. (28) (dashed lines) and comparison with the analytical solution (full line) given by Eq. (29) for an initial condition $f_0(x) = e^{-0.5(x-0.5)^2/\sigma^2}$ for $\sigma = 0.1$. The number of sequential operators used in the implementation of the non-unitary diagonal operator is $s = 14, 8, 6, 1$ respectively for time $t = 0.005, 0.02, 0.05, 1$. The infidelity between the analytical solution and the implemented qubit states are $1 - F = 8.5 \times 10^{-4}, 1.8 \times 10^{-3}, 3.1 \times 10^{-4}, 3.5 \times 10^{-4}, 1.7 \times 10^{-15}$ respectively for time $t = 0.005, 0.02, 0.05, 1$.

ACKNOWLEDGEMENTS

The authors would like to thank G.Di Molfetta, B.Claudon and C.Feniou for their useful feedbacks on the form and content of this manuscript. The quantum circuit diagrams were generated with the quantikz package [48]. U.Nzongani acknowledges support from the PEPR EPiQ ANR-22-PETQ-0007, by the ANR JCJC DisQC ANR-22-CE47-0002-01.

[1] Christof Zalka. Simulating quantum systems on a quantum computer. *Proceedings of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences*, 454(1969):313–322, 1998.

[2] Stephen Wiesner. Simulations of many-body quantum systems by a quantum computer. *arXiv preprint quant-ph/9603028*, 1996.

[3] Ivan Kassal, Stephen P Jordan, Peter J Love, Masoud Mohseni, and Alán Aspuru-Guzik. Polynomial-time quantum algorithm for the simulation of chemical dynamics. *Proceedings of the National Academy of Sciences*, 105(48):18681–18686, 2008.

[4] Stephen P Jordan, Keith SM Lee, and John Preskill. Quantum algorithms for quantum field theories. *Science*, 336(6085):1130–1133, 2012.

[5] Ljubomir Budinski. Quantum algorithm for the advection–diffusion equation simulated with the lattice boltzmann method. *Quantum Information Processing*, 20(2):57, 2021.

[6] Noah Linden, Ashley Montanaro, and Changpeng Shao. Quantum vs. classical algorithms for solving the heat equation. *Communications in Mathematical Physics*, 395(2):601–641, 2022.

[7] Edward Farhi, Jeffrey Goldstone, Sam Gutmann, Joshua Lapan, Andrew Lundgren, and Daniel Preda. A quantum adiabatic evolution algorithm applied to random instances of an np-complete problem. *Science*, 292(5516):472–475, 2001.

[8] Nicola Mariella and Andrea Simonetto. A quantum algorithm for the sub-graph isomorphism problem. *ACM Transactions on Quantum Computing*, 4(2), feb 2023.

[9] Lov Grover and Terry Rudolph. Creating superpositions that correspond to efficiently integrable probability distributions. *arXiv preprint quant-ph/0208112*, 2002.

[10] Julien Zylberman and Fabrice Debbasch. Efficient quantum state preparation with walsh series. *Phys. Rev. A*, 109:042401, Apr 2024.

[11] Xiao-Ming Zhang, Tongyang Li, and Xiao Yuan. Quantum state preparation with optimal circuit depth: Implementations and applications. *Physical Review Letters*, 129(23):230504, 2022.

[12] Adriano Barenco, Charles H Bennett, Richard Cleve, David P DiVincenzo, Norman Margolus, Peter Shor, Tycho Sleator, John A Smolin, and Harald Weinfurter. Elementary gates for quantum computation. *Physical review A*, 52(5):3457, 1995.

[13] Stephen S Bullock and Igor L Markov. Asymptotically optimal circuits for arbitrary n-qubit diagonal comutations. *Quantum Information & Computation*, 4(1):27–47, 2004.

[14] Jonathan Welch, Daniel Greenbaum, Sarah Mostame, and Alán Aspuru-Guzik. Efficient quantum circuits for diagonal unitaries without ancillas. *New Journal of Physics*, 16(3):033040, 2014.

[15] Michael R Geller, Zoë Holmes, Patrick J Coles, and Andrew Sornborger. Experimental quantum learning of a spectral decomposition. *Physical Review Research*, 3(3):033200, 2021.

[16] Xiaoming Sun, Guojing Tian, Shuai Yang, Pei Yuan, and Shengyu Zhang. Asymptotically optimal circuit depth for quantum state preparation and general unitary synthesis. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 2023.

[17] Baptiste Claudon, Julien Zylberman, César Feniou, Fabrice Debbasch, Alberto Peruzzo, and Jean-Philip Piquemal. Polylogarithmic-depth controlled-not gates without ancilla qubits. *arXiv preprint arXiv:2312.13206*, 2023.

[18] Andrew M Childs, Jiaqi Leng, Tongyang Li, Jin-Peng Liu, and Chenyi Zhang. Quantum simulation of real-space dynamics. *Quantum*, 6:860, 2022.

[19] Cristopher Moore and Martin Nilsson. Parallel quantum computation and quantum codes. *SIAM journal on computing*, 31(3):799–815, 2001.

[20] Robert M Gingrich and Colin P Williams. Non-unitary probabilistic quantum computing. In *ACM International Conference Proceeding Series*, volume 58, pages 1–6. Citeseer, 2004.

[21] Hiroaki Terashima and Masahito Ueda. Nonunitary quantum circuit. *International Journal of Quantum Information*, 3(04):633–647, 2005.

[22] H. J. Briegel, D. E. Browne, W. Dür, R. Raussendorf, and M. Van den Nest. Measurement-based quantum computation. *Nature Physics*, 5(1):19–26, January 2009.

[23] Ammar Daskin and Sabre Kais. An ancilla-based quantum simulation framework for non-unitary matrices. *Quantum Information Processing*, 16:1–17, 2017.

[24] Anthony W. Schlimgen, Kade Head-Marsden, LeeAnn M. Sager-Smith, Prineha Narang, and David A. Mazziotti. Quantum state preparation and nonunitary evolution with diagonal operators. *Phys. Rev. A*, 106:022414, Aug 2022.

[25] Andrew M. Childs and Nathan Wiebe. Hamiltonian simulation using linear combinations of unitary operations. *Quantum Info. Comput.*, 12(11–12):901–924, nov 2012.

[26] Dominic W Berry, Andrew M Childs, Richard Cleve, Robin Kothari, and Rolando D Somma. Simulating hamiltonian dynamics with a truncated taylor series. *Physical review letters*, 114(9):090502, 2015.

[27] Andrew M Childs, Robin Kothari, and Rolando D Somma. Quantum algorithm for systems of linear equations with exponentially improved dependence on precision. *SIAM Journal on Computing*, 46(6):1920–1950, 2017.

[28] Dominic W Berry, Andrew M Childs, Aaron Ostrander, and Guoming Wang. Quantum algorithm for linear differential equations with exponentially improved dependence on precision. *Communications in Mathematical Physics*, 356:1057–1081, 2017.

[29] András Gilyén, Yuan Su, Guang Hao Low, and Nathan Wiebe. Quantum singular value transformation and beyond: exponential improvements for quantum matrix arithmetics. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, pages 193–204, 2019.

[30] Guang Hao Low and Isaac L Chuang. Hamiltonian simulation by qubitization. *Quantum*, 3:163, 2019.

[31] Peter W. Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM Journal on Computing*, 26(5):1484–1509, October 1997.

[32] Lov K. Grover. A fast quantum mechanical algorithm for database search, 1996.

[33] Pulak Ranjan Giri and Vladimir E Korepin. A review on quantum search algorithms. *Quantum Information Processing*, 16:1–36, 2017.

[34] Gilles Brassard, Peter Høyer, Michele Mosca, and Alain Tapp. Quantum amplitude amplification and estimation, 2002.

[35] John M Martyn, Zane M Rossi, Andrew K Tan, and Isaac L Chuang. Grand unification of quantum algorithms. *PRX quantum*, 2(4):040203, 2021.

[36] Craig Gidney. https://algassert.com/circuits/2015/06/22/using-quantum-gates-instead-of-ancilla- bits.html. 2015.

[37] F. Gray. Pulse code communication, 1953. US Patent 2,632,058.

[38] Michael A Nielsen and Isaac Chuang. Quantum computation and quantum information, 2002.

[39] Kenneth George Beauchamp. *Applications of Walsh and related functions, with an introduction to sequency theory*, volume 2. Academic press, 1984.

[40] Joseph L Walsh. A closed set of normal orthogonal functions. *American Journal of Mathematics*, 45(1):5–24, 1923.

[41] Diogo Cruz, Romain Fournier, Fabien Gremion, Alix Jeannerot, Kenichi Komagata, Tara Tosic, Jarla Thiesbrummel, Chun Lam Chan, Nicolas Macris, Marc-André Dupertuis, and Clément Javerzac-Galy. Efficient quantum algorithms for ghz and w states, and implementation on the ibm quantum computer. *Advanced Quantum Technologies*, 2(5–6), April 2019.

[42] Ugo Nzongani, Julien Zylberman, Carlo-Elia Doncecchi, Armando Pérez, Fabrice Debbasch, and Pablo Arnault. Quantum circuits for discrete-time quantum walks with position-dependent coin operator. *Quantum Information Processing*, 22(7):270, 2023.

[43] D. Bertsimas and R. Weismantel. *Optimization Over Integers*. Dynamic Ideas, 2005.

[44] Chung-Kwong Yuen. Function approximation by walsh series. *IEEE Transactions on Computers*, 100(6):590–598, 1975.

[45] Dominic W. Berry, Andrew M. Childs, Richard Cleve, Robin Kothari, and Rolando D. Somma. Exponential improvement in precision for simulating sparse hamiltonians. In *Proceedings of the Forty-Sixth Annual ACM Symposium on Theory of Computing*, STOC '14, page 283–292, New York, NY, USA, 2014. Association for Computing Machinery.

[46] D. Coppersmith. An approximate fourier transform useful in quantum factoring, 2002.
[47] Robert M. Gray. *Toeplitz And Circulant Matrices: A Review (Foundations and Trends(R) in Communications and Information Theory)*. Now Publishers Inc., Hanover, MA, USA, 2006.
[48] Alastair Kay. Tutorial on the quantikz package. *arXiv preprint arXiv:1809.03842*, 2018.

Appendix A: Examples of quantum circuits

In this section, illustrative examples of quantum circuits are presented for both the sequential and Walsh-Hadamard decompositions with $n = 3$ qubits. First, the non-optimized circuits are displayed, then the optimize one with a Gray code ordering to reduce the number of gates. Finally, the quantum circuit associated to the parallelization scheme, shown on Fig. 2, with $p = 3$ ancilla qubits is presented. Consider the 3-qubit diagonal unitary:

$$\hat{U}_\theta = e^{i\hat{\theta}} = \begin{pmatrix} e^{i\theta_0} & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & e^{i\theta_1} & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & e^{i\theta_2} & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & e^{i\theta_3} & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & e^{i\theta_4} & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & e^{i\theta_5} & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & e^{i\theta_6} & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & e^{i\theta_7} \end{pmatrix}, \tag{A1}$$

where each $e^{i\theta_x}$ corresponds to the eigenvalue of basis state eigenvector $|x\rangle$. The main register of the following quantum circuits is $|q = q_{n-1} \dots q_1 q_0\rangle$ where $q_i \in \{0, 1\}$ and $q = \sum_{j=0}^{n-1} q_j 2^j$.

1. Sequential decomposition

The sequential decomposition of diagonal unitaries leads to quantum circuits composed of multi-controlled phase and NOT gates. Figure 10 gives an example of an exact decomposition of an arbitrary 3-qubit diagonal unitary where the multi-controlled gate are implemented in decimal ordering.

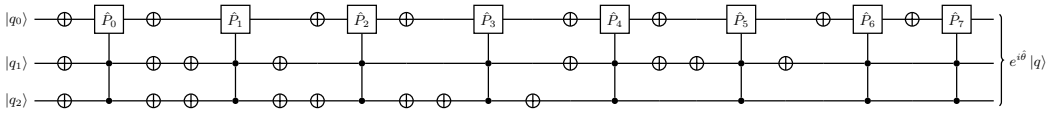

Figure 10: Sequential decomposition without Gray code for $n = 3$ qubits with the phase gates $\hat{P}_x = \hat{P}(\theta_x) = \begin{pmatrix} 1 & 0 \\ 0 & e^{i\theta_x} \end{pmatrix}$ with $x \in [0, 7]$.

A significant number of NOT gates cancel out. One can use a Gray code [37] to maximise the number of cancellation. The Gray code, also called reflected binary code, was first introduced by Frank Gray in 1953. It is a way to enumerate a set of binary element such that the Hamming distance of two neighboring element is equal to 1, *i.e.*, there is only a one bit difference between two neighboring bit strings. For instance, the 3-bits Gray code is shown on Table III.

The quantum circuit implementing an arbitrary 3-qubit diagonal unitary using a Gray code ordering is represented Fig. 11.

Lastly, we show the quantum circuit implementing the adjustable-depth sequential decomposition of $e^{i\hat{f}}$ on Fig. 12.

2. Walsh-Hadamard decomposition

Diagonal unitaries can be exactly implemented as a product of Walsh-Hadamard decomposition $\hat{U} = \prod \hat{W}_j$ is implemented with a quantum circuit composed of CNOT and $\hat{R}_Z$ gates. Its naïve form is shown on Fig. 13.

As with the sequential decomposition, the number of CNOTs can be reduced by changing the order of the operator using Gray code [14, 37]. The obtained circuit is shown on Fig. 14.

Decimal with Gray code ordering	*Binary representation*
0	000
1	001
3	011
2	010
6	110
7	111
5	101
4	100

Table III: 3-bits Gray code.

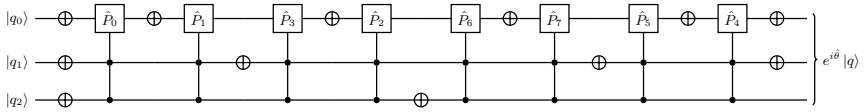

Figure 11: Sequential decomposition with Gray code for $n = 3$ qubits the phase gates $\hat{P}_x = \hat{P}(\theta_x) = \begin{pmatrix} 1 & 0 \\ 0 & e^{i\theta_x} \end{pmatrix}$ with $x \in [0, 7]$.

Finally, we show the quantum circuit implementing the adjustable-depth Walsh-Hadamard decomposition of $e^{i\hat{f}}$ on Fig. 15.

Appendix B: Proof of Theorem V.1

Proof. Consider a decomposition $\tilde{U} = \prod_{j=0}^{p-1} \hat{U}_j$ where each $\hat{U}_j$ acts on $k_j \leq n$ qubits and consider a number $k + p$ ancilla qubits, with $k = \sum_{j=1}^{p-1} k_j$. The registers of ancilla qubits are denoted by $R_1, \ldots, R_{p-1}$ and each R_j contains k_j qubits that will serve to copy only the qubits on which $\hat{U}_j$ is acting non-trivially. The register R_{q_A} contains $p-1$ qubits on which q_A will be copied. The proof of correctness of the quantum circuit Fig. 4 starts by considering a state $|x\rangle$ of the computational basis of the main register. For simplicity, tensor products between the qubit states and tensor products with identity operators are omitted:

$$
\begin{aligned}
&|x\rangle |0\rangle^{\otimes k} |0\rangle_{q_A} |0\rangle^{\otimes(p-1)} \\
&\xrightarrow{\hat{H}} |x\rangle |0\rangle^{\otimes k} \frac{1}{\sqrt{2}}(|0\rangle_{q_A} + |1\rangle_{q_A}) |0\rangle^{\otimes(p-1)} \\
&\xrightarrow{\widehat{\text{copy}}_{R \to R_1 \cup \ldots \cup R_{p-1}} \otimes \widehat{\text{copy}}_{q_A \to R_{q_A}}} |x\rangle |\tilde{x}_1\rangle \ldots |\tilde{x}_{p-1}\rangle \frac{1}{\sqrt{2}}(|0\rangle^{\otimes p} + |1\rangle^{\otimes p}) \\
&\xrightarrow{\bigotimes_{j=0}^{p-1} \hat{C}_{\tilde{q}_j \to R_j}(\hat{U}_j)} \frac{1}{\sqrt{2}}(\hat{U}_0 |x\rangle \hat{U}_1 |\tilde{x}_1\rangle \ldots \hat{U}_{p-1} |\tilde{x}_{p-1}\rangle |0\rangle^{\otimes p} \\
&+ |x\rangle |\tilde{x}_1\rangle \ldots |\tilde{x}_{p-1}\rangle |1\rangle^{\otimes p}) \\
&\xrightarrow{\bigotimes_{j=0}^{p-1} \hat{C}_{q_j \to R_j}(\hat{U}_j^\dagger)} \frac{1}{\sqrt{2}}(\hat{U}_0 |x\rangle \hat{U}_1 |\tilde{x}_1\rangle \ldots \hat{U}_{p-1} |\tilde{x}_{p-1}\rangle |0\rangle^{\otimes p} \\
&+ \hat{U}_0^\dagger |x\rangle \hat{U}_1^\dagger |\tilde{x}_1\rangle \ldots \hat{U}_{p-1}^\dagger |\tilde{x}_{p-1}\rangle |1\rangle^{\otimes p}) \\
&= \frac{1}{\sqrt{2}}(e^{i\tilde{\theta}_x} |x\rangle |\tilde{x}_1\rangle \ldots |\tilde{x}_{p-1}\rangle |0\rangle^{\otimes p} + e^{-i\tilde{\theta}_x} |x\rangle |\tilde{x}_1\rangle \ldots |\tilde{x}_{p-1}\rangle |1\rangle^{\otimes p}) \\
&\xrightarrow{\widehat{\text{copy}}^{-1}_{R \to R_1 \cup \ldots \cup R_{p-1}} \otimes \widehat{\text{copy}}^{-1}_{q_A \to R_{q_A}}} \frac{1}{\sqrt{2}}(e^{i\tilde{\theta}_x} |x\rangle |0\rangle^{\otimes k} |0\rangle_{q_A} + e^{-i\tilde{\theta}_x} |x\rangle |0\rangle^{\otimes k} |1\rangle_{q_A}) |0\rangle^{\otimes(p-1)} \\
&\xrightarrow{\hat{P}\hat{H}} (-i\cos(\theta_x) |x\rangle |0\rangle^{\otimes k} |0\rangle_{q_A} + \sin(\theta_x) |x\rangle |0\rangle^{\otimes k} |1\rangle_{q_A}) |0\rangle^{\otimes(p-1)}
\end{aligned}
\tag{B1}
$$

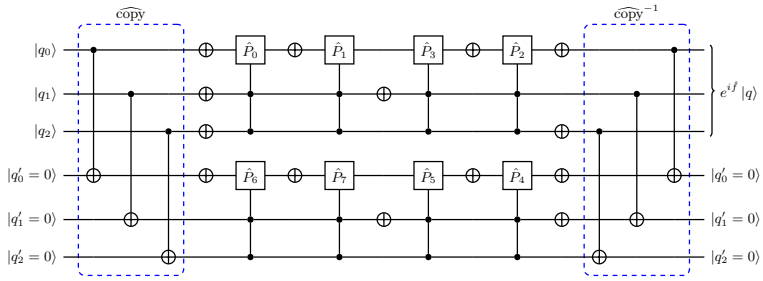

Figure 12: Adjustable-depth quantum circuit implementing the sequential decomposition with Gray code for $n = 3$ qubits with $p = 3$ ancilla qubits.

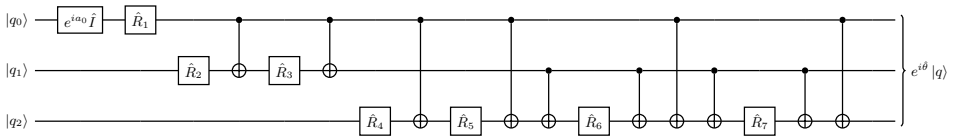

Figure 13: Walsh decomposition with no Gray code for $n = 3$ qubits with $\hat{R}_x \equiv \hat{R}_Z(-2a_x)$ with $x \in [1, 7]$. This circuit is shown on Fig. 5 of [14]. The first gate corresponds to the global phase gate $e^{ia_0}\hat{I} = \begin{pmatrix} e^{ia_0} & 0 \\ 0 & e^{ia_0} \end{pmatrix}$.

with $\tilde{\theta}_x = \sum_{j=0}^{p-1} \theta_j(x)$ and $\hat{U}_j = \sum_{x=0}^{N-1} e^{i\theta_j(x)} |x\rangle\langle x|$. The operator $\hat{C}_{\overline{q}_j \to R_j}(\hat{U}_j)$ is defined as the unitary $\hat{U}_j$ applied on the register R_j and anti-controlled by the j-th qubit of R_{q_A} (q_0 denotes q_A) and the operator $\hat{C}_{q_j \to R_j}(\hat{U}_j^\dagger)$ is defined as the unitary $\hat{U}_j^\dagger$ applied on the register R_j and controlled by the j-th qubit of R_{q_A}. The unitary $\widehat{\text{copy}}_{R \to R_1 \cup \ldots \cup R_{p-1}}$ corresponds to the copy of main register R into the copy registers $R_1, \ldots, R_{p-1}$ and $\widehat{\text{copy}}_{q_A \to R_{q_A}}$ corresponds to the copy of q_A into the register R_{q_A}.

The unitary $\hat{U}_D$ corresponding to the previous operations (*i.e.* Fig. 4) is a $(\alpha d_{\max}, k + p, \epsilon)$-block encoding of $\hat{D} = \sum_{x=0}^{N} d_x |x\rangle\langle x|$:

$$\begin{aligned} &\| \frac{1}{\alpha d_{\max}} \hat{D} - (\langle 0|^{\otimes k} \langle 1|_{q_A} \langle 0|^{\otimes (p-1)} \otimes \hat{I}_R) \hat{U}_D (|0\rangle^{\otimes k} |1\rangle_{q_A} |0\rangle^{\otimes (p-1)} \otimes \hat{I}_R) \|_2 \\ &= \| \sum_{x=0}^{N-1} \frac{d_x}{\alpha d_{\max}} |x\rangle\langle x| - \sum_{x=0}^{N-1} \sin(\tilde{\theta}_x) |x\rangle\langle x| \|_2 \\ &= \max_x \| \frac{d_x}{\alpha d_{\max}} - \sin(\tilde{\theta}_x) \| \leq \epsilon \end{aligned} \tag{B2}$$

The inequality comes from the hypothesis that $\tilde{U} = \sum_{x=0}^{N} e^{i\tilde{\theta}_x} |x\rangle\langle x|$ is an ϵ approximation of $\hat{U} = e^{i \arcsin(\hat{D}/(\alpha d_{\max}))}$ in spectral norm. □

Appendix C: Scalings of the different quantum circuits

1. Exact methods

a. *Diagonal unitaries*

The following Lemmas and corollaries summarize the complexity of implementing exactly any n-qubit diagonal unitary.

Lemma C.1 (Exact sequential decomposition without ancilla)**.** *Any n-qubit diagonal unitary is implementable through its sequential decomposition with a quantum circuit of depth $\mathcal{O}(n2^n)$ and size $\mathcal{O}(n2^n)$ without using ancilla qubits.*

Proof. The quantum circuit is composed of 2^n operator $\hat{U}_j$ defined in Eq. (3). Thanks to the fact that the unitaries $\hat{U}_j$ commute with each other, one can choose the order of implementation of the $\hat{U}_j$ which cancels a maximum number

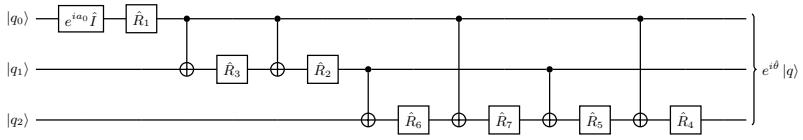

Figure 14: Walsh decomposition with Gray code for $n = 3$ qubits with $\hat{R}_x \equiv \hat{R}_Z(-2a_x)$ with $x \in [1, 7]$. This circuit is shown on Fig. 5 of [14]. The first gate corresponds to the global phase gate $e^{ia_0}\hat{I} = \begin{pmatrix} e^{ia_0} & 0 \\ 0 & e^{ia_0} \end{pmatrix}$.

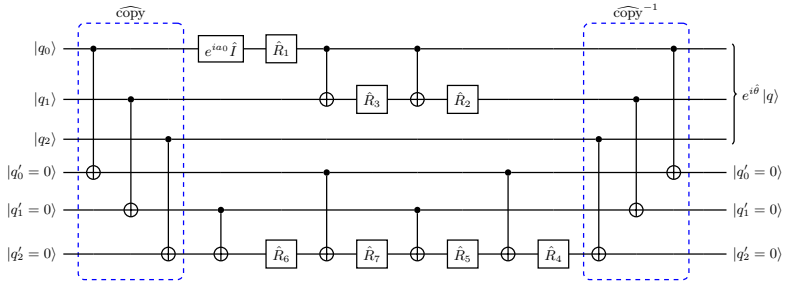

Figure 15: Adjustable-depth quantum circuit implementing the Walsh decomposition with Gray code for $n = 3$ qubits with $p = 3$ ancilla qubits.

of $\hat{X}$ gates. This is achieved using a Gray code [37–39]: between two following $\hat{U}_j$, $\hat{U}_{j'}$, only one bit changes in the binary decomposition of j and j'. This leads to an ordering where only one $\hat{X}$ gate remains between two consecutive $\Lambda_{\{0,\dots,n-2\}}(\hat{P}(\theta_j))$, $\Lambda_{\{0,\dots,n-2\}}(\hat{P}(\theta_{j'})$. Then, each of the $\Lambda_{\{0,\dots,n-2\}}(\hat{P}(\theta_j))$ is implemented using the scheme of [36] which is an exact scheme with size and depth linear in n without using ancilla qubits. □

Lemma C.2 (Exact sequential decomposition with one ancilla)**.** *Any n-qubit diagonal unitary is implementable through its sequential decomposition with a quantum circuit of depth $\mathcal{O}(\log(n)^3 2^n)$ and size $\mathcal{O}(n\log(n)^4 2^n)$ using one ancilla qubit.*

Proof. Similarly than previously, the 2^n sequential operator $\hat{U}_j$ are implemented in a gray-code order. Then, each of them is implemented with scheme presented in [17] which uses one ancilla qubit to achieve a polylogarithmic depth $\mathcal{O}(\log(n)^3)$ to implement any n-controlled operations at the cost of a size $\mathcal{O}(n\log(n)^4)$ (see Corollary 1 in [17]). Therefore, the overall depth for an exact implementation of a given arbitrary diagonal unitary is $\mathcal{O}(\log(n)^3 2^n)$ and size $\mathcal{O}(n\log(n)^4 2^n)$ using one ancilla qubit. □

Lemma C.3 (Exact Walsh-Hadamard decomposition without ancilla (Theorem 1.3 in [13]))**.** *Any n-qubit diagonal unitary is implementable through its Walsh-Hadamard decomposition with a quantum circuit of depth $\mathcal{O}(2^n)$ and size $\mathcal{O}(2^n)$ without using ancilla qubits.*

Corollary C.1 (Fully parallelized exact sequential decomposition)**.** *Any n-qubit diagonal unitary is implementable through its sequential decomposition with a fully parallelized quantum circuit of depth $\mathcal{O}(n)$ and size $\mathcal{O}(n2^n)$ using $\mathcal{O}(n2^n)$ ancilla qubits.*

Proof. Lemma C.1 or lemma C.2 associated to the full parallelization Theorem III.1. □

Corollary C.2 (Fully parallelized exact Walsh-Hadamard decomposition)**.** *Any n-qubit diagonal unitary is implementable through its Walsh-Hadamard decomposition with a fully parallelized quantum circuit of depth $\mathcal{O}(n)$ and size $\mathcal{O}(n2^n)$ without using $\mathcal{O}(n2^n)$ ancilla qubits.*

Proof. Lemma C.3 with the full parallelization Theorem III.1. □

Corollary C.3 (Partially parallelized exact sequential decomposition using linear depth decomposition)**.** *Any n-qubit diagonal unitary is implementable through its sequential decomposition with a partially parallelized quantum circuit of depth $\mathcal{O}(n^2 2^n/m + \log(m/n))$ and size $\mathcal{O}(n2^n + m)$ using $m \in [n, \mathcal{O}(n2^n)]$ ancilla qubits.*

Proof. Lemma C.1 with the partial parallelization Theorem III.2. □

Corollary C.4 (Partially parallelized exact sequential decomposition using polylogarithmic depth decomposition). *Any n-qubit diagonal unitary is implementable through its sequential decomposition with a partially parallelized quantum circuit of depth $\mathcal{O}(n\log(n)^3 2^n/m+\log(m/n))$ and size $\mathcal{O}(n\log(n)^4 2^n+m)$ using $m\in[n+2,\mathcal{O}(n2^n)]$ ancilla qubits.*

Proof. Lemma C.2 with the partial parallelization Theorem III.2. □

Corollary C.5 (Partially parallelized exact Walsh-Hadamard decomposition). *Any n-qubit diagonal unitary is implementable through its Walsh-Hadamard decomposition with a partially parallelized quantum circuit of depth $\mathcal{O}(n^2 2^n/m+\log(m/n))$ and size $\mathcal{O}(n2^n+m)$ using $m\in[n,\mathcal{O}(n2^n)]$ ancilla qubits.*

Proof. Lemma C.3 with the partial parallelization Theorem III.2. □

Lemma C.4 (Lemma 11 [16]). *Any n-qubit diagonal unitary can be implemented by a quantum circuit of depth $\mathcal{O}(2^n/n)$ and size $\mathcal{O}(2^n)$ without ancillary qubits.*

Lemma C.5 (Lemma 20 [16]). *For any $m\in[2n,2^n/n]$, any n-qubit diagonal unitary can be implemented by a quantum circuit of depth $\mathcal{O}(2^n/m+\log(m))$ and size $\mathcal{O}(2^n+nm)$ with m ancillary qubits.*

Lemma C.6 (Lemma 20 [16] with a maximum number of ancilla qubits). *Any n-qubit diagonal unitary can be implemented by a quantum circuit of depth $\mathcal{O}(n)$ and size $\mathcal{O}(2^n)$ with $m=2^n/n$ ancillary qubits.*

b. *Non-unitary diagonal operators*

The following corollaries summarize the complexity of implementing exactly a block-encoding $\hat{U}_D$ of any n-qubit non-unitary diagonal operator $\hat{D}$ through the controlled-diagonal unitaries $e^{\pm i\arcsin(\hat{D}/(\alpha d_{\max}))}$ with $\alpha\geq 1$.

Corollary C.6 (Exact block-encoding using a sequential decomposition and one ancilla qubit). *For any $\alpha\geq 1$, any n-qubit non-unitary diagonal operator $\hat{D}$ can be $(\alpha d_{\max},1,0)$-block-encoded with a quantum circuit of depth $\mathcal{O}(n2^n)$ and size $\mathcal{O}(n2^n)$ using one ancilla qubit.*

Proof. Lemma C.1 and the block-encoding Theorem V.1. □

Corollary C.7 (Exact sequential block-encoding with two ancilla qubits). *For any $\alpha\geq 1$, any n-qubit non-unitary diagonal operator $\hat{D}$ can be $(\alpha d_{\max},2,0)$-block-encoded with a quantum circuit of depth $\mathcal{O}(\log(n)^3 2^n)$ and size $\mathcal{O}(n\log(n)^4 2^n)$ using two ancilla qubits.*

Proof. Lemma C.2 and the block-encoding Theorem V.1: the $(n-1)$-controlled phase gates are controlled by one ancilla qubit, becoming n-controlled phase gates. The last ancilla qubit is used to implement the n-controlled phase gates exactly with Corollary 1 of [17]. □

Corollary C.8 (Exact Walsh-Hadamard block-encoding with one ancilla). *For any $\alpha\geq 1$, any n-qubit non-unitary diagonal operator $\hat{D}$ can be $(\alpha d_{\max},1,0)$-block-encoded with a quantum circuit of depth $\mathcal{O}(2^n)$ and size $\mathcal{O}(2^n)$ using one ancilla qubit.*

Proof. Lemma C.3 and the block-encoding Theorem V.1 □

Corollary C.9 (Fully parallelized block-encoding with exact sequential decomposition). *For any $\alpha\geq 1$, any n-qubit non-unitary diagonal operator $\hat{D}$ can be $(\alpha d_{\max},m,0)$-block-encoded with a quantum circuit of depth $\mathcal{O}(n)$ and size $\mathcal{O}(n2^n)$ using $m=\mathcal{O}(n2^n)$ ancilla qubits.*

Proof. Lemma C.1 or Lemma C.2 associated to the full parallelization Theorem III.1 and the block-encoding Theorem V.1 □

Corollary C.10 (Fully parallelized block-encoding with exact Walsh-Hadamard decomposition). *For any $\alpha\geq 1$, any n-qubit non-unitary diagonal operator $\hat{D}$ can be $(\alpha d_{\max},m,0)$-block-encoded with a quantum circuit of depth $\mathcal{O}(n)$ and size $\mathcal{O}(n2^n)$ using $m=\mathcal{O}(n2^n)$ ancilla qubits.*

Proof. Lemma C.3 with the full parallelization Theorem III.1 and the block-encoding Theorem V.1. □

Corollary C.11 (Partially parallelized block-encoding with exact sequential decomposition). *For any $\alpha\geq 1$, any n-qubit non-unitary diagonal operator $\hat{D}$ can be $(\alpha d_{\max},m,0)$-block-encoded with a quantum circuit of depth $\mathcal{O}(n^2 2^n/m+\log(m/n))$ and size $\mathcal{O}(n2^n+m)$ using $m\in[\Omega(n),\mathcal{O}(n2^n)]$ ancilla.*

Proof. Lemma C.1 with the partial parallelization Theorem III.2 and the block-encoding Theorem V.1 □

Corollary C.12 (Partially parallelized block-encoding with exact sequential decomposition using polylogarithmic depth decomposition)**.** *For any $\alpha \geq 1$, any n-qubit non-unitary diagonal operator $\hat{D}$ can be $(\alpha d_{\max}, m, 0)$-block-encoded with a quantum circuit of depth $\mathcal{O}(n\log(n)^3 2^n/m + \log(m/n))$ and size $\mathcal{O}(n\log(n)^4 2^n + m)$ using m ancilla qubits with $m = \Omega(n)$ and $m = \mathcal{O}(2^n/n)$.*

Proof. Lemma C.2 with the partial parallelization Theorem III.2 and the block-encoding Theorem V.1 □

Corollary C.13 (Partially parallelized block-encoding with exact Walsh-Hadamard decomposition)**.** *For any $\alpha \geq 1$, any n-qubit non-unitary diagonal operator $\hat{D}$ can be $(\alpha d_{\max}, m, 0)$-block-encoded with a quantum circuit of depth $\mathcal{O}(n^2 2^n/m + \log(m/n))$ and size $\mathcal{O}(n2^n + m)$ using m ancilla qubits with $m = \Omega(n)$ and $m = \mathcal{O}(2^n/n)$.*

Proof. Lemma C.3 with the partial parallelization Theorem III.2 and the block-encoding Theorem V.1. □

Corollary C.14 (Lemma 11 [16] + block-encoding)**.** *For any $\alpha \geq 1$, any n-qubit non-unitary diagonal operator $\hat{D}$ can be $(\alpha d_{\max}, m, 0)$-block-encoded with a quantum circuit of depth $\mathcal{O}(2^n/n)$ and size $\mathcal{O}(2^n)$ using n ancilla qubits.*

Proof. The Lemma 11 [16] associated with the block-encoding Theorem V.1 where each $\hat{U}_j$ is not a sequential or Walsh-Hadamard operator, but a primitive quantum gates of the decomposition given by the quantum circuits of [16]. Indeed, any n-qubit unitary operation with size $s(n)$ and depth $d(n)$ can be controlled by a qubit q_A with size $\mathcal{O}(s(n)+n)$ and depth $d(n) + \mathcal{O}(log(n))$ using $n-1$ copies of q_A to control in parallel the different gates. □

Corollary C.15 (Lemma 20 [16] + block-encoding)**.** *For any $\alpha \geq 1$, any n-qubit non-unitary diagonal operator $\hat{D}$ can be $(\alpha d_{\max}, m, 0)$-block-encoded with a quantum circuit of depth $\mathcal{O}(2^n/m+\log(m))$ and size $\mathcal{O}(n2^n)$ using m ancilla qubits with $m = \Omega(n)$ and $m = \mathcal{O}(2^n/n)$.*

Proof. Lemma 20 [16] associated to the block-encoding Theorem V.1. □

Corollary C.16 (Lemma 20 [16] with a maximum number of ancilla qubits + block-encoding)**.** *For any $\alpha \geq 1$, any n-qubit non-unitary diagonal operator $\hat{D}$ can be $(\alpha d_{\max}, m, 0)$-block-encoded with a quantum circuit of depth $\mathcal{O}(n)$ and size $\mathcal{O}(n2^n)$ using $m = \mathcal{O}(2^n/n)$ ancilla qubits.*

Proof. Lemma 20 [16] associated with the block-encoding Theorem V.1. □

2. Approximate methods

a. *For diagonal unitaries depending on differentiable functions*

In the following, we consider a n-qubit diagonal unitary $\hat{U} = \sum_{x=0}^{N-1} e^{if(x/N)} |x\rangle\langle x|$, with $N = 2^n$ depending on a real-valued function f defined on $[0,1]$ with bounded first derivative.

Corollary C.17 (Approximate sequential decomposition without ancilla)**.** *Any n-qubit diagonal unitary depending on a real-valued function f defined on $[0,1]$ with bounded first derivative is implementable up to an error $\epsilon > 0$ in spectral norm through its sequential decomposition with a quantum circuit of depth $\mathcal{O}(\log(1/\epsilon)/\epsilon)$ and size $\mathcal{O}(\log(1/\epsilon)/\epsilon)$ without ancilla qubits.*

Proof. This corollary is a direct consequence of Lemma C.1 and the approximate Theorem IV.1. □

Corollary C.18 (Approximate sequential decomposition with one ancilla)**.** *Any n-qubit diagonal unitary depending on a real-valued function f defined on $[0,1]$ with bounded first derivative is implementable up to an error $\epsilon > 0$ in spectral norm through its sequential decomposition with a quantum circuit of depth $\mathcal{O}(\log(\log(1/\epsilon))^3/\epsilon)$ and size $\mathcal{O}(\log(1/\epsilon)\log(\log(1/\epsilon))^4/\epsilon)$ using one ancilla qubit.*

Proof. This corollary is a direct consequence of Lemma C.2 and the approximate Theorem IV.1. □

Corollary C.19 (Approximate Walsh-Hadamard decomposition without ancilla [14])**.** *Any n-qubit diagonal unitary depending on a real-valued function f defined on $[0,1]$ with bounded first derivative is implementable up to an error $\epsilon > 0$ in spectral norm through its Walsh-Hadamard decomposition with a quantum circuit of depth $\mathcal{O}(1/\epsilon)$ and size $\mathcal{O}(1/\epsilon)$ without ancilla qubit.*

Proof. This corollary is a direct consequence of Lemma C.3 and the approximate Theorem IV.1. □

Corollary C.20 (Fully parallelized approximate sequential decomposition)**.** *Any n-qubit diagonal unitary depending on a real-valued function f defined on $[0,1]$ with bounded first derivative is implementable up to an error $\epsilon > 0$ in spectral norm through its sequential decomposition with a quantum circuit of depth $\mathcal{O}(\log(1/\epsilon))$ and size $\mathcal{O}(\log(1/\epsilon)/\epsilon)$ using $\mathcal{O}(\log(1/\epsilon)/\epsilon)$ ancilla qubits.*

Proof. This corollary is a direct consequence of corollary C.1 associated with the approximate Theorem IV.1. □

Corollary C.21 (Fully parallelized approximate Walsh-Hadamard decomposition)**.** *Any n-qubit diagonal unitary depending on a real-valued function f defined on $[0,1]$ with bounded first derivative is implementable up to an error $\epsilon > 0$ in spectral norm through its Walsh-Hadamard decomposition with a quantum circuit of depth $\mathcal{O}(\log(1/\epsilon))$ and size $\mathcal{O}(\log(1/\epsilon)/\epsilon)$ using $\mathcal{O}(\log(1/\epsilon)/\epsilon)$ ancilla qubits.*

This corollary is a direct consequence of corollary C.2 with the approximate Theorem IV.1.

Corollary C.22 (Partially parallelized approximate sequential decomposition using linear depth decomposition)**.** *Any n-qubit diagonal unitary depending on a real-valued function f defined on $[0,1]$ with bounded first derivative is implementable up to an error $\epsilon > 0$ in spectral norm through its sequential decomposition with a quantum circuit of depth $\mathcal{O}(\log(1/\epsilon)^2/(\epsilon m) + \log(m/\log(1/\epsilon)))$ and size $\mathcal{O}(\log(1/\epsilon)/\epsilon + m)$ using m ancilla qubits with $m = \Omega(\log(1/\epsilon))$ and $m = \mathcal{O}(\log(1/\epsilon)/\epsilon)$.*

Proof. This corollary is a direct consequence of corollary C.3 with the approximate Theorem IV.1. □

Corollary C.23 (Partially parallelized approximate sequential decomposition using polylogarithmic depth decomposition)**.** *Any n-qubit diagonal unitary depending on a real-valued function f defined on $[0,1]$ with bounded first derivative is implementable up to an error $\epsilon > 0$ in spectral norm through its sequential decomposition with a quantum circuit of depth $\mathcal{O}(\log(1/\epsilon)\log(\log(1/\epsilon))^3/(\epsilon m) + \log(m/\log(1/\epsilon)))$ and size $\mathcal{O}(\log(1/\epsilon)\log(\log(1/\epsilon))^4/\epsilon + m)$ using m ancilla qubits with $m = \Omega(\log(1/\epsilon))$ and $m = \mathcal{O}(\log(1/\epsilon)/\epsilon)$.*

Proof. This corollary is a direct consequence of corollary C.4 with the approximate Theorem IV.1. □

Corollary C.24 (Partially parallelized approximate Walsh-Hadamard decomposition)**.** *Any n-qubit diagonal unitary depending on a real-valued function f defined on $[0,1]$ with bounded first derivative is implementable up to an error $\epsilon > 0$ in spectral norm through its Walsh-Hadamard decomposition with a quantum circuit of depth $\mathcal{O}(\log(1/\epsilon)^2/(\epsilon m) + \log(m/\log(1/\epsilon)))$ and size $\mathcal{O}(\log(1/\epsilon)/\epsilon + m)$ using m ancilla qubits with $m = \Omega(\log(1/\epsilon))$ and $m = \mathcal{O}(\log(1/\epsilon)/\epsilon)$.*

Proof. This corollary is a direct consequence of corollary C.5 with the approximate Theorem IV.1. □

Corollary C.25 (Walsh-recursive)**.** *Any n-qubit diagonal unitary depending on a real-valued function f defined on $[0,1]$ with bounded first derivative is implementable up to an error $\epsilon > 0$ in spectral norm with a quantum circuit of depth $\mathcal{O}(1/(\epsilon\log(1/\epsilon)))$ and size $\mathcal{O}(1/\epsilon)$ without ancilla qubits.*

Proof. This corollary is a direct consequence of Lemma 11 [16] with the approximate Theorem IV.1. □

Corollary C.26 (Walsh-optimized adjustable-depth)**.** *For any $\epsilon > 0$, any n-qubit diagonal unitary depending on a real-valued function f defined on $[0,1]$ with bounded first derivative is implementable up to an error ϵ in spectral norm with a quantum circuit of depth $\mathcal{O}(1/(\epsilon m) + \log(m))$ and size $\mathcal{O}(1/\epsilon + \log(1/\epsilon)m)$ using m ancilla qubits with $m = \Omega(\log(1/\epsilon))$ and $m = \mathcal{O}(1/(\epsilon\log(1/\epsilon)))$.*

Proof. This corollary is a direct consequence of Lemma 20 [16] with the approximate Theorem IV.1. □

Corollary C.27 (Walsh-optimized fully parallelized)**.** *For any $\epsilon > 0$, any n-qubit diagonal unitary depending on a real-valued function f defined on $[0,1]$ with bounded first derivative is implementable up to an error ϵ in spectral norm with a quantum circuit of depth $\mathcal{O}(\log(1/\epsilon))$ and size $\mathcal{O}(1/\epsilon)$ with $m = \mathcal{O}(1/(\epsilon\log(1/\epsilon)))$ ancilla qubits.*

Proof. This corollary is a direct consequence of Lemma 20 [16] for a maximum number of ancilla qubits with the approximate Theorem IV.1. □

b. For non-unitary diagonal operators depending on differentiable functions

The following corollaries summarize the complexity of implementing approximately a block-encoding $\hat{U}_D$ of a n-qubit non-unitary diagonal operator $\hat{D} = \sum_{x=0}^{N-1} f(x/N) |x\rangle \langle x|$, depending on a real-valued function f defined on $[0,1]$ with bounded first derivative, through the controlled-diagonal unitaries $e^{\pm i \arcsin(\hat{D}/(\alpha d_{\max}))}$ with $\alpha > 1$.

Corollary C.28 (Approximate block-encoding using a sequential decomposition and one ancilla qubit)**.** *For any $\alpha > 1$, any n-qubit non-unitary diagonal operator $\hat{D}$ depending on a real-valued function f defined on $[0,1]$ with bounded first derivative can be $(\alpha d_{\max}, 1, \epsilon)$-block-encoded with a quantum circuit of depth $\mathcal{O}(\log(1/\epsilon)/\epsilon)$ and size $\mathcal{O}(\log(1/\epsilon)/\epsilon)$ using one ancilla qubit.*

Proof. This corollary is a direct consequence of corollary C.17 applied on $\hat{U} = \sum_{x=0}^{N-1} e^{ig(x/N)} |x\rangle \langle x|$ with $g(x) = \arcsin(f(x)/(\alpha \|f\|_\infty))$ such that $\|g'\|_\infty \leq \|f'\|_\infty / (\|f\|_\infty \sqrt{\alpha^2 - 1})$ and the block-encoding Theorem V.1. □

Corollary C.29 (Approximate sequential block-encoding with two ancilla qubits)**.** *For any $\alpha > 1$, any n-qubit non-unitary diagonal operator $\hat{D}$ depending on a real-valued function f defined on $[0,1]$ with bounded first derivative can be $(\alpha d_{\max}, 1, \epsilon)$-block-encoded with a quantum circuit of depth $\mathcal{O}(\log(\log(1/\epsilon))^3/\epsilon)$ and size $\mathcal{O}(\log(1/\epsilon)(\log(\log(1/\epsilon))^4/\epsilon)$ using two ancilla qubits.*

Proof. This corollary is a direct consequence of corollary C.18 and the block-encoding Theorem V.1. □

Corollary C.30 (Approximate Walsh-Hadamard block-encoding with one ancilla)**.** *For any $\alpha > 1$, any n-qubit non-unitary diagonal operator $\hat{D}$ depending on a real-valued function f defined on $[0,1]$ with bounded first derivative can be $(\alpha d_{\max}, 1, \epsilon)$-block-encoded with a quantum circuit of depth $\mathcal{O}(1/\epsilon)$ and size $\mathcal{O}(1/\epsilon)$ using one ancilla qubit.*

Proof. This corollary is a direct consequence of corollary C.19 and the block-encoding Theorem V.1. □

Corollary C.31 (Fully parallelized block-encoding with approximate sequential decomposition)**.** *For any $\alpha > 1$, any n-qubit non-unitary diagonal operator $\hat{D}$ depending on a real-valued function f defined on $[0,1]$ with bounded first derivative can be $(\alpha d_{\max}, 1, \epsilon)$-block-encoded with a quantum circuit of depth $\mathcal{O}(\log(1/\epsilon))$ and size $\mathcal{O}(\log(1/\epsilon)/\epsilon)$ using $m = \mathcal{O}(\log(1/\epsilon)/\epsilon)$ ancilla qubits.*

Proof. This corollary is a direct consequence of corollary C.20 and the block-encoding Theorem V.1. □

Corollary C.32 (Fully parallelized block-encoding with approximate Walsh-Hadamard decomposition)**.** *For any $\alpha > 1$, any n-qubit non-unitary diagonal operator $\hat{D}$ depending on a real-valued function f defined on $[0,1]$ with bounded first derivative can be $(\alpha d_{\max}, 1, \epsilon)$-block-encoded with a quantum circuit of depth $\mathcal{O}(\log(1/\epsilon))$ and size $\mathcal{O}(\log(1/\epsilon)/\epsilon)$ using $m = \mathcal{O}(\log(1/\epsilon)/\epsilon)$ ancilla qubits.*

Proof. This corollary is a direct consequence of corollary C.21 and the block-encoding Theorem V.1. □

Corollary C.33 (Partially parallelized block-encoding with approximate sequential decomposition)**.** *For any $\alpha > 1$, any n-qubit non-unitary diagonal operator $\hat{D}$ depending on a real-valued function f defined on $[0,1]$ with bounded first derivative can be $(\alpha d_{\max}, 1, \epsilon)$-block-encoded with a quantum circuit of depth $\mathcal{O}(\log(1/\epsilon)^2/(\epsilon m) + \log(m/\log(1/\epsilon)))$ and size $\mathcal{O}(\log(1/\epsilon)/\epsilon + m)$ using m ancilla qubits with $m = \Omega(\log(1/\epsilon))$ and $m = \mathcal{O}(\log(1/\epsilon)/\epsilon)$.*

Proof. This corollary is a direct consequence of corollary C.22 and the block-encoding Theorem V.1. □

Corollary C.34 (Partially parallelized block-encoding with approximate sequential decomposition using polylogarithmic depth decomposition)**.** *For any $\alpha > 1$, any n-qubit non-unitary diagonal operator $\hat{D}$ depending on a real-valued function f defined on $[0,1]$ with bounded first derivative can be $(\alpha d_{\max}, 1, \epsilon)$-block-encoded with a quantum circuit of depth $\mathcal{O}(\log(1/\epsilon)\log(\log(1/\epsilon))^3/(\epsilon m) + \log(m/\log(1/\epsilon)))$ and size $\mathcal{O}(\log(1/\epsilon)\log(\log(1/\epsilon))^4/\epsilon + m)$ using m ancilla qubits with $m = \Omega(\log(1/\epsilon))$ and $m = \mathcal{O}(\log(1/\epsilon)/\epsilon)$.*

Proof. This corollary is a direct consequence of corollary C.23 and the block-encoding Theorem V.1. □

Corollary C.35 (Partially parallelized block-encoding with approximate Walsh-Hadamard decomposition)**.** *For any $\alpha > 1$, any n-qubit non-unitary diagonal operator $\hat{D}$ depending on a real-valued function f defined on $[0,1]$ with bounded first derivative can be $(\alpha d_{\max}, 1, \epsilon)$-block-encoded with a quantum circuit of depth $\mathcal{O}(\log(1/\epsilon)^2/(\epsilon m) + \log(m/\log(1/\epsilon)))$ and size $\mathcal{O}(\log(1/\epsilon)/\epsilon + m)$ using m ancilla qubits with $m = \Omega(\log(1/\epsilon))$ and $m = \mathcal{O}(\log(1/\epsilon)/\epsilon)$.*

Proof. This corollary is a direct consequence of corollary C.24 and the block-encoding Theorem V.1. □

Corollary C.36 (Walsh-recursive). *For any $\alpha > 1$, any n-qubit non-unitary diagonal operator $\hat{D}$ depending on a real-valued function f defined on $[0,1]$ with bounded first derivative can be $(\alpha d_{\max}, 1, \epsilon)$-block-encoded with a quantum circuit of depth $\mathcal{O}(1/(\epsilon \log(1/\epsilon)))$ and size $\mathcal{O}(1/\epsilon)$ using $m = \mathcal{O}(\log(1/\epsilon))$ ancilla qubits.*

Proof. This corollary is a direct consequence of corollary C.25 and the block-encoding Theorem V.1. □

Corollary C.37 (Walsh-optimized adjustable-depth). *For any $\alpha > 1$, any n-qubit non-unitary diagonal operator $\hat{D}$ depending on a real-valued function f defined on $[0,1]$ with bounded first derivative can be $(\alpha d_{\max}, 1, \epsilon)$-block-encoded with a quantum circuit of depth $\mathcal{O}(1/(\epsilon m)+\log(m))$ and size $\mathcal{O}(\log(1/\epsilon)/\epsilon)$ using m ancilla qubits with $m = \Omega(\log(1/\epsilon))$ and $m = \mathcal{O}(1/(\epsilon \log(1/\epsilon)))$.*

Proof. This corollary is a direct consequence of corollary C.26 and the block-encoding Theorem V.1. □

Corollary C.38 (Walsh-optimized fully parallelized). *For any $\alpha > 1$, any n-qubit non-unitary diagonal operator $\hat{D}$ depending on a real-valued function f defined on $[0,1]$ with bounded first derivative can be $(\alpha d_{\max}, 1, \epsilon)$-block-encoded with a quantum circuit of depth $\mathcal{O}(\log(1/\epsilon))$ and size $\mathcal{O}(1/\epsilon)$ using $m = \mathcal{O}(1/(\epsilon \log(1/\epsilon)))$ ancilla qubits.*

Proof. This corollary is a direct consequence of corollary C.27 and the block-encoding Theorem V.1. □

3. Sparse diagonal operators

a. *Sparse diagonal unitaries*

The following Lemmas and corollaries summarize the complexity of implementing any n-qubit diagonal unitary with a decomposition containing only s sequential operators or s Walsh-Hadamard operators.

Lemma C.7 (Sparse sequential decomposition without ancilla). *Any n-qubit diagonal unitary with a s-sparse sequential decomposition is exactly implementable with a quantum circuit of depth $\mathcal{O}(ns)$ and size $\mathcal{O}(ns)$ without using ancilla qubits.*

Proof. The quantum circuit is composed of s operator $\hat{U}_j$ defined in Eq. (3). Each of the s sequential operator contains at worst $2n$ $\hat{X}$-Pauli gates and one $\Lambda_{\{0,\dots,n-2\}}(\hat{P}(\theta_j))$ which is implemented using the scheme of C.Gidney [36]. It is an exact method for multi-controlled gates with size and depth linear in n without using ancilla qubits. □

Lemma C.8 (Approximate sparse sequential decomposition without ancilla). *Any n-qubit diagonal unitary with a s-sparse sequential decomposition is implementable up to an error $\epsilon > 0$ with a quantum circuit of depth $\mathcal{O}(s \log(n)^3 \log(s/\epsilon))$ and size $\mathcal{O}(sn \log(n)^4 \log(s/\epsilon))$ without using ancilla qubits.*

Proof. The quantum circuit is composed of s operator $\hat{U}_j$ defined in Eq. (3). Each of the s sequential operator contains at worst $2n$ $\hat{X}$-Pauli gates and one $\Lambda_{\{0,\dots,n-2\}}(\hat{P}(\theta_j))$. Each of the s multi-controlled phase gate is implemented using the approximative method of Claudon et al. (Proposition 2 in [17]) with error $\epsilon' = \epsilon/s$, depth $\mathcal{O}(\log(n)^3 \log(s/\epsilon))$ and size $\mathcal{O}(n \log(n)^4 \log(s/\epsilon))$. □

Lemma C.9 (Sparse sequential decomposition with one ancilla). *Any n-qubit diagonal unitary with a s-sparse sequential decomposition is exactly implementable with a quantum circuit of depth $\mathcal{O}(s \log(n)^3)$ and size $\mathcal{O}(sn \log(n)^4)$ using one ancilla qubit.*

Proof. The quantum circuit is composed of s operator $\hat{U}_j$ defined in Eq. (3). Each of the s sequential operator contains at worst $2n$ $\hat{X}$-Pauli gates and one $\Lambda_{\{0,\dots,n-2\}}(\hat{P}(\theta_j))$. Each of the s multi-controlled phase gate is implemented using the exact method of Claudon et al. (Corollary 1 in [17]) with depth $\mathcal{O}(\log(n)^3)$, size $\mathcal{O}(n \log(n)^4)$ using one ancilla qubit. □

Corollary C.39 (Adjustable-depth sparse sequential decomposition). *Any n-qubit diagonal unitary with a s-sparse sequential decomposition is exactly implementable with a quantum circuit of depth $\mathcal{O}(s \log(n)^3/(m/n) + \log(m/n))$ and size $\mathcal{O}(sn \log(n)^4 + m)$ using $m \geq n+2$ ancilla qubits.*

Proof. This corollary is a direct consequence of Lemma C.9 and the adjustable-depth Theorem III.2. □

Corollary C.40 (Fully parallelized sparse sequential decomposition). *Any n-qubit diagonal unitary with a s-sparse sequential decomposition is exactly implementable with a quantum circuit of depth $\mathcal{O}(\log(n)^3 + \log(s))$ and size $\mathcal{O}(sn \log(n)^4)$ using $m = \mathcal{O}(ns)$ ancilla qubits.*

Proof. This corollary is a direct consequence of lemma C.9 and the fully parallelized Theorem III.1. □

Lemma C.10 (Sparse Walsh-Hadamard decomposition without ancilla)**.** *Any n-qubit diagonal unitary with a s-sparse Walsh-Hadamard decomposition is exactly implementable with a quantum circuit of depth $\mathcal{O}(sk)$ and size $\mathcal{O}(sk)$ without using ancilla qubits, where k is the maximum number of 1 in the binary decomposition of the indexes j of the Walsh-Hadamard operators.*

Proof. The quantum circuit is composed of s operator $\hat{W}_j$ defined in Eq. (8) for $j \in S \subseteq \{0, 1, ..., 2^n - 1\}$, with $\|S\| = s$. Each Walsh-Hadamard operator $\hat{W}_j$ for $j = \sum_{i=0}^{n-1} j_i 2^i \in S$ contains exactly $2k_j$ controlled-NOT gates and one $\hat{R}_Z$ gate where k_j is the number of 1 in the binary decomposition of j, a.k.a. $k_j = \sum_{i=0}^{n-1} j_i$. By defining $k = \max_{j \in S} k_j$, the size of the quantum circuit is bounded by $\mathcal{O}(sk)$ and the depth is also bounded by $\mathcal{O}(sk)$ where $k \leq n$ is independant of n. □

Corollary C.41 (Adjustable-depth sparse Walsh-Hadamard decomposition)**.** *Any n-qubit diagonal unitary with a s-sparse Walsh-Hadamard decomposition is exactly implementable with a quantum circuit of depth $\mathcal{O}(sk/(m/k) + \log(m/k))$ and size $\mathcal{O}(sk + m)$ without using $m \geq k$ ancilla qubits, where k is the maximum number of 1 in the binary decomposition of the indexes j of the Walsh-Hadamard operators.*

Proof. This corollary is a direct consequence of lemma C.10 and the adjustable-depth Theorem III.2. □

Corollary C.42 (Fully-parallelized sparse Walsh-Hadamard decomposition)**.** *Any n-qubit diagonal unitary with a s-sparse Walsh-Hadamard decomposition is exactly implementable with a quantum circuit of depth $\mathcal{O}(k + \log(s))$ and size $\mathcal{O}(sk)$ without using $m = \mathcal{O}(sk)$ ancilla qubits, where k is the maximum number of 1 in the binary decomposition of the indexes j of the Walsh-Hadamard operators.*

Proof. This corollary is a direct consequence of lemma C.10 and the fully parallelized Theorem III.1. □

b. Sparse non-unitary diagonal operators

The following corollaries summarize the complexity of implementing any n-qubit non-unitary diagonal operators $\hat{D}$ with an associated s-sparse operator $e^{\pm i \arcsin(\hat{D}/(\alpha d_{\max}))}$, with $\alpha > 1$.

Corollary C.43 (Block-encoding of sparse sequential decomposition with one ancilla)**.** *For any $\alpha \geq 1$, any n-qubit non-unitary diagonal operator with a s-sparse sequential decomposition can be $(\alpha d_{\max}, 1, 0)$-block-encoded with a quantum circuit of depth $\mathcal{O}(ns)$ and size $\mathcal{O}(ns)$ using one ancilla qubit.*

Proof. This corollary is a direct consequence of lemma C.7 and the block-encoding Theorem V.1. Remark that if $\hat{D}$ has s non-zero real eigenvalues, then $\hat{U} = e^{\pm i \arcsin(\hat{D}/(\alpha d_{\max}))}$ has at most s eigenvalues different than 1. □

Corollary C.44 (Block-encoding of approximate sparse sequential decomposition with one ancilla)**.** *For any $\alpha \geq 1$, any n-qubit non-unitary diagonal operator with a s-sparse sequential decomposition can be $(\alpha d_{\max}, 1, \epsilon)$-block-encoded with a quantum circuit of depth $\mathcal{O}(s \log(n)^3 \log(s/\epsilon))$, size $\mathcal{O}(sn \log(n)^4 \log(s/\epsilon))$ using one ancilla qubit.*

Proof. This corollary is a direct consequence of lemma C.8 and the block-encoding Theorem V.1. □

Corollary C.45 (Block-encoding of sparse sequential decomposition with two ancilla qubits)**.** *For any $\alpha \geq 1$, any n-qubit non-unitary diagonal operator with a s-sparse sequential decomposition can be $(\alpha d_{\max}, 2, 0)$-block-encoded with a quantum circuit of depth $\mathcal{O}(s \log(n)^3)$, size $\mathcal{O}(sn \log^4(n))$ using two ancilla qubit.*

Proof. This corollary is a direct consequence of lemma C.9 and the block-encoding Theorem V.1. □

Corollary C.46 (Adjustable-depth sparse sequential decomposition)**.** *For any $\alpha \geq 1$, any n-qubit non-unitary diagonal operator with a s-sparse sequential decomposition can be $(\alpha d_{\max}, m, 0)$-block-encoded with a quantum circuit of depth $\mathcal{O}(s \log(n)^3/(m/n) + \log(m/n))$, size $\mathcal{O}(sn \log^4(n) + m)$ using m ancilla qubits with $m = \Omega(n)$ and $m = \mathcal{O}(ns)$.*

Proof. This corollary is a direct consequence of corollary C.39 and the block-encoding Theorem V.1. □

Corollary C.47 (Fully parallelized sparse sequential decomposition)**.** *For any $\alpha \geq 1$, any n-qubit non-unitary diagonal operator with a s-sparse sequential decomposition can be $(\alpha d_{\max}, m, 0)$-block-encoded with a quantum circuit of depth $\mathcal{O}(\log(n)^3 + \log(s))$, size $\mathcal{O}(sn \log^4(n))$ using $m = \mathcal{O}(ns)$ ancilla qubits.*

Proof. This corollary is a direct consequence of Corollary C.40 and the block-encoding Theorem V.1. □

Corollary C.48 (Sparse Walsh-Hadamard decomposition with one ancilla)**.** *Any n-qubit non-unitary diagonal $\hat{D}$ with real eigenvalues such that $\hat{U}{=}e^{\pm i \arcsin(\hat{D}/(\alpha d_{\max}))}$ has an $\epsilon > 0$ approximative s-sparse Walsh-Hadamard decomposition can be $(\alpha d_{\max}, 1, \epsilon)$-block-encoded with a quantum circuit of depth $\mathcal{O}(sk)$, size $\mathcal{O}(sk)$ using one ancilla qubit.*

Proof. This corollary is a direct consequence of lemma C.10 and the block-encoding Theorem V.1. □

Corollary C.49 (Adjustable-depth sparse Walsh-Hadamard decomposition)**.** *Any n-qubit non-unitary diagonal $\hat{D}$ with real eigenvalues such that $\hat{U}{=}e^{\pm i \arcsin(\hat{D}/(\alpha d_{\max}))}$ has an $\epsilon > 0$ approximative s-sparse Walsh-Hadamard decomposition can be $(\alpha d_{\max}, m, \epsilon)$-block-encoded with a quantum circuit of depth $\mathcal{O}(sk/(m/k) + \log(m/k))$, size $\mathcal{O}(sk + m)$ using $m = \Omega(k)$ ancilla qubits.*

Proof. This corollary is a direct consequence of corollary C.41 and the block-encoding Theorem V.1. □

Corollary C.50 (Fully-parallelized sparse Walsh-Hadamard decomposition)**.** *Any n-qubit non-unitary diagonal $\hat{D}$ with real eigenvalues such that $\hat{U}{=}e^{\pm i \arcsin(\hat{D}/(\alpha d_{\max}))}$ has an $\epsilon > 0$ approximative s-sparse Walsh-Hadamard decomposition can be $(\alpha d_{\max}, m, \epsilon)$-block-encoded with a quantum circuit of depth $\mathcal{O}(k + \log(s))$, size $\mathcal{O}(sk)$ using $m = \mathcal{O}(sk)$ ancilla qubits.*

Proof. This corollary is a direct consequence of corollary C.50 and the block-encoding Theorem V.1. □