Reduction of the asymptotic prefactors for an optimally bounded quantum state preparation algorithm

Giacomo Belli[1,*], Michele Amoretti[1]

[1]Quantum Software Laboratory, Department of Engineering and Architecture, University of Parma, 43124 Parma, Italy

Abstract

The preparation of arbitrary n-qubit quantum states is a cross-cutting subroutine for many quantum algorithms, and the effort to reduce its circuit complexity is a significant challenge. In the literature, the quantum state preparation algorithm by Sun et al. is known to be optimally bounded, defining the asymptotically optimal width-depth trade-off bounds with and without ancillary qubits. In this work, a simpler algebraic decomposition is proposed to separate the preparation of the real part of the desired state from the complex one, resulting in a reduction in terms of circuit depth, total gates, and CNOT count when m ancillary qubits are available. The reduction in complexity is due to the use of a single operator Λ for each uniformly controlled gate, instead of the three in the original decomposition. Using the PennyLane library, this new algorithm for state preparation has been implemented and tested in a simulated environment for both dense and sparse quantum states, including those that are random and of physical interest. Its performance was compared with that of other QSP algorithms, including exact and approximate paradigms without ancillary qubits.

Keywords

Quantum state preparation, ancillary qubits, optimal bound, Fourier space, diagonal operators

1. Introduction

Quantum state preparation (QSP) is a fundamental subroutine in many quantum algorithms [1]. From quantum linear system solvers [2] to Hamiltonian simulation [3], from quantum chemistry [4] to quantum machine learning [5], the constant search for more efficient algorithms is still an open challenge [6, 7, 8, 9]. The goal of a QSP algorithm is to construct a unitary operator $U \in U(2^n)$ such that $|\psi\rangle = U|0\rangle^{\otimes n}$, where $|0\rangle^{\otimes n}$ is the conventional n-qubit initial state and $|\psi\rangle = \sum_{i=0}^{2^n-1} c_i|i\rangle$ is the desired quantum state in the computational basis, normalized with respect to the l_2-norm. The preparation can follow an exact or approximate paradigm; in the approximate version, one tries to prepare a state $|\phi\rangle$ that is ϵ-close to $|\psi\rangle$ with respect to some metric. Furthermore, the QSP algorithm can benefit from the use of ancillary registers in addition to the input registers to decrease some computational costs, such as depth, which corresponds to execution time and gate counts. Contextually, of particular interest is the study of the optimal space-time trade-off bounds for quantum state preparation circuits [10].

The study of QSP began in its exact version and without ancillary qubits using the method proposed by Nielsen and Chuang [11], characterized by a depth upper bound of $\mathcal{O}(2^n)$. In the same period, Grover and Rudolph [12] proposed a *hierarchical* QSP based on layers of conditional rotations for the special case of superpositions of integrable distributions, with a depth upper bound of $\mathcal{O}(n2^n)$. Their algorithm contained what we might define today as the *embryonic form* of Uniformly Controlled Rotations (UCR) [13], even though the UCR class had not yet been formalized. In 2005, Möttönen et al. [14] proposed a 3-stage QSP that properly formalizes the UCR class, with a depth of $\mathcal{O}(2^n)$, $2^{n+2} - 4n - 4$ CNOT, and $2^{n+2} - 5$ rotations. Each UCR layer can be implemented via Gray Code [15], and its parameters can be determined via Binary Search Tree (BST) diagrams, making QSP modular [14]. Immediately after, Bergholm et al. [16], still within the framework of a general QSP, introduced and systematically analyzed the Uniformly Controlled Gate (UCG) class as a generalization of the UCR class.

ICTCS 2026: Italian Conference on Theoretical Computer Science, September 7–9, 2026, Udine, Italy

*Corresponding author.

giacomo.belli@unipr.it (G. Belli); michele.amoretti@unipr.it (M. Amoretti)

https://github.com/qslab-unipr/qsp-sun (G. Belli); https://www.qslab.unipr.it/ (M. Amoretti)

0009-0008-3739-9525 (G. Belli); 0000-0002-6046-1904 (M. Amoretti)

Their algorithm had an upper bound of $2^{n+1} - 2n - 2$ for the number of CNOT gates, with a depth also of order $\mathcal{O}(2^n)$.

Following the line of research on QSP, where the UCG is the key building block and no ancillary qubits are used, from 2006 to 2016, some important results structured a bridge (not only theoretical) between Gate Synthesis and QSP, where efficient general decomposition techniques (QR/CS) are combined with the modularity of UCG/UCR [13]. While Shende et al. [17] proposed a framework for synthesis pipelines, with tighter asymptotic bounds than those in the literature at the time and a handling of the nearest-neighbor topology, Plesch and Brukner [18] focused on minimizing the CNOT count for a general QSP with arbitrary decomposition, achieving $\frac{23}{24}2^n - 2^{\frac{n}{2}+1} + \frac{5}{3}$ for even n and $\frac{115}{96}2^n$ for odd n. Finally, Iten et al. [19] provided the framework for the systematic synthesis of isometries, within which QSP represents a special case. These milestones in the development of QSP clarify the link with unitary synthesis and mark the transition from the preparation of specific quantum states to the design of unitary operators, where QSP can be seen as a sub-routine for preparing columns.

The use of ancillae in the history of QSP began with Zhang et al. [20], whose algorithm can generate the target state in $\mathcal{O}(n^2)$ depth with $\mathcal{O}(2^{2n})$ ancillary qubits, but only with a limited success probability. In the same year, Rosenthal [21] tackled the topic of space-time trade-off by obtaining an algorithm with depth $\mathcal{O}(n)$ using $\mathcal{O}(n2^n)$ ancillary qubits. The *asymptotically* optimal space-time trade-off bounds were found by Sun et al. [22] in their algorithm, hereinafter referred to as SUN, with depth $\mathcal{O}(\frac{2^n}{m+n} + n\log n)$ and width $\mathcal{O}(2^n)$, where m is the number of ancillary qubits. The QSP circuit by Zhang et al. [23] of depth $\mathcal{O}(n)$ and using $\Theta(2^n)$ ancillary qubits is a special case of [22]. This line of research culminated in the work of Yuan et al. [10], which completely solves the circuit complexity with or without ancillary qubits.

This paper provides the following contributions:

1. the optimized version of SUN, hereafter denoted as OSUN, in terms of depth and gate counts over the entire first range of [22, Theorem 1];
2. the implementation of each UCG-level with a single Λ-type constructor (instead of the 3 of the original version [22]), enabled by the introduction of the already known separation strategy between real and complex parts within the ancillary paradigm;
3. a reduction of the prefactors of the asymptotic dominant terms, and their explicit calculation, for the space-time trade-off in the first range of [22, Theorem 1];
4. the implementation of the new algorithm via the PennyLane library, available on the GitHub repository [24];
5. a complete verification of the algorithm's performance through an extensive battery of tests (sparse and dense states, with complex or real coefficients, and states of interest for research) executed in a simulated environment on the HPC system of the University of Parma;
6. an explicit comparison in terms of depth, total gates, CNOT, and gate type with other QSP algorithms and with the original version by Sun et al. (the standard for optimal complexity).

The remainder of the paper is organized as follows. In Section 2, the QSP algorithm by Sun et al. [22] is presented with its essential features, thus providing the theoretical framework and the main equations to understand its optimization. In Section 3, the decomposition that splits the QSP unitary into two parts, one for the real part of the desired state and one for the remaining complex part, is derived, and the role of the single Λ-type constructor is discussed. The complete circuit for the $n = 2$ case is shown as an example, and Section 4 presents the results of the simulation tests. Contextually, the comparison with other QSP algorithms and with the implementation of the original version of Sun et al.'s QSP [25, 24] are shown. Finally, Section 5 concludes the paper with a summary of the main results and a discussion of future work. An extended version of this work is available at [26].

2. SUN: traditional framework, diagonal operators and ancillae

In [22], Sun et al. tightly characterized the depth-size complexity of the QSP problem, except for a logarithmic gap in the small parametric range $m = \left[\omega\left(\frac{2^n}{n\log n}\right), o(2^n)\right]$. Their algorithms, which

also cover the case without ancillary qubits, allow for the construction of optimal quantum circuits where the upper and lower depth bounds coincide exactly in the two width ranges $m = \mathcal{O}\left(\frac{2^n}{n \log n}\right)$ and $m = \Omega(2^n)$. Improving on the work presented in [25], regarding the first implementation of SUN, this paper focuses on the first of the two ranges with optimal *asymptotic* complexity. In this specific domain, the traditional framework for QSP [14] can be used; it consists of a ladder structure of Uniformly Controlled Gates (UCGs) throughout the quantum register, as shown in Figure 1a. In matrix

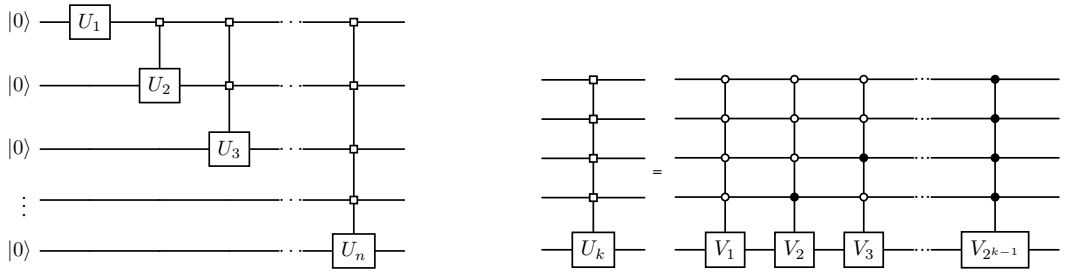

(a) Ladder structure of uniformly controlled gates (UCGs); each qubit k is handled by an UCG that applies a single-qubit unitary to it, conditioned on the basis state of the first $k-1$ qubits.

(b) The decomposition of a single UCG in the complete sequence of controlled gates, as stated by its own definition; by means of a cyclic Gray Code, it is possible to implement it easily [27]. However, without appropriate optimizations [19] this decomposition is suboptimal [22].

Figure 1: QSP traditional framework.

form, each UCG is a block-diagonal operator U_k belonging to $\mathbb{C}^{2^k \times 2^k}$, where each block $V_j \in U(2)$ is a single-qubit gate for any $1 \leqslant j \leqslant 2^{k-1}$, in turn decomposable as $V_j = e^{i\alpha_j} R_z(\beta_j) R_y(\gamma_j) R_z(\delta_j)$ for some parameters $(\alpha, \beta, \gamma, \delta) \in \mathbb{R}$.

Actually, it is convenient to rewrite each block V_j only in terms of diagonal operators thanks to the change of gate set $R_y(\gamma_j) \equiv SH \cdot R_z(\gamma_j) \cdot HS^\dagger$. In this way, the n-qubit extension for the arbitrary UCG U_n becomes:

$$U_n = [\text{diag}(e^{i\alpha_1}, ..., e^{i\alpha_{2^{n-1}}}) \otimes \mathbb{I}_1] \cdot F_n[R_z(\vec{\beta})] \cdot [\mathbb{I}_{n-1} \otimes (SH)] \cdot F_n[R_z(\vec{\gamma})] \cdot [\mathbb{I}_{n-1} \otimes (HS^\dagger)] \cdot F_n[R_z(\vec{\delta})] \quad (1)$$

where $F_n[R_z(\vec{\theta})] = \text{diag}(R_z(\theta_1), R_z(\theta_2), ..., R_z(\theta_{2^{n-1}}))$. It is important to note that the *central* block-diagonal operator $F_n(R_z(\vec{\gamma}))$ is solely responsible for preparing the real part of the desired quantum state, as the vector $\vec{\gamma}$ collects the parameters of the blocks R_y after the change of the gate set. Up to global phases[1], this is already the decomposition used in [22], where diagonal Λ-type operators,

$$\Lambda_n(\vec{\theta}) = \text{diag}(1, e^{i\theta_1}, e^{i\theta_2}, ..., e^{i\theta_{2^n-1}}) \ \in \mathbb{C}^{2^n \times 2^n} \quad \text{for} \quad \vec{\theta} = (\theta_1, \theta_2, ..., \theta_{2^n-1}) \in \mathbb{R}^{2^n-1} \quad (2)$$

are introduced to efficiently parallelize on an ancillary register [25]. Therefore, the arbitrary UCG U_n can be rewritten as:

$$U_n = \ e^{i(\alpha_1 - \beta_1)} \Lambda_n^{(1)}(\vec{\alpha^*}, \vec{\beta^*}) \cdot [\mathbb{I}_{n-1} \otimes (SH)] \cdot e^{-i\gamma_1} \Lambda_n^{(2)}(\vec{\gamma^*}) \cdot [\mathbb{I}_{n-1} \otimes (HS^\dagger)] \cdot e^{-i\delta_1} \Lambda_n^{(3)}(\vec{\delta^*}) \quad (3)$$

where $\vec{\alpha^*}, \vec{\beta^*}, \vec{\gamma^*}, \vec{\delta^*} \in \mathbb{R}^{2^n-1}$ are the vectors that group the updated Λ_n parameters after the collection of global phases. The corresponding circuit is shown in Figure 2. In the end, the traditional QSP circuit of Figure 1a reduces to the concatenation of Λ_n circuits, interspersed with the two unitary operators derived from the pair of gates S and H.

In [22, Section IV], the procedure to implement any operator $\Lambda_n \in \mathbb{C}^{2^n \times 2^n}$ is discussed and proven. The equivalent quantum circuit has depth $\mathcal{O}(\log_2 m + \frac{2^n}{m})$ and width $\mathcal{O}(2^n)$, within the ancillary range $m \in [2n, \frac{2^n}{n}]$. As shown in Equation 2, matrices Λ_n perform phase shifts on each state of the computational basis, a task that can be efficiently parallelized through sums in Fourier space [25]. The key algorithm to design the Λ_n circuit has 5 stages that define the formalism for the scalable implementation; they are depicted from a high-level perspective in Figure 3. The details of their implementation, as well as the management of all parameters, are discussed in [25]. This work does not modify the basic implementation of operators Λ_n in any way.

[1] A detailed mathematical proof is provided in [25].

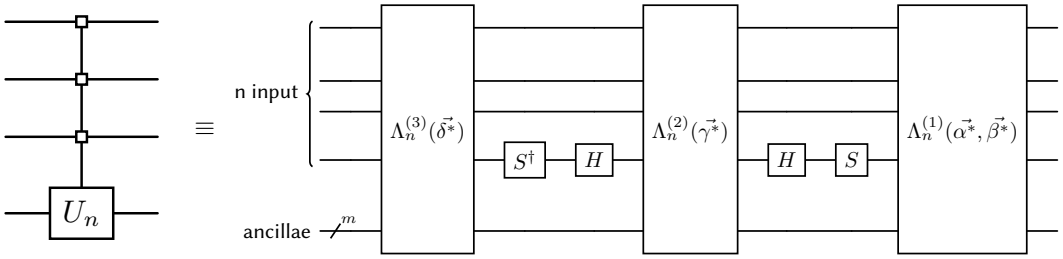

Figure 2: UCG in terms of Λ-type operators.

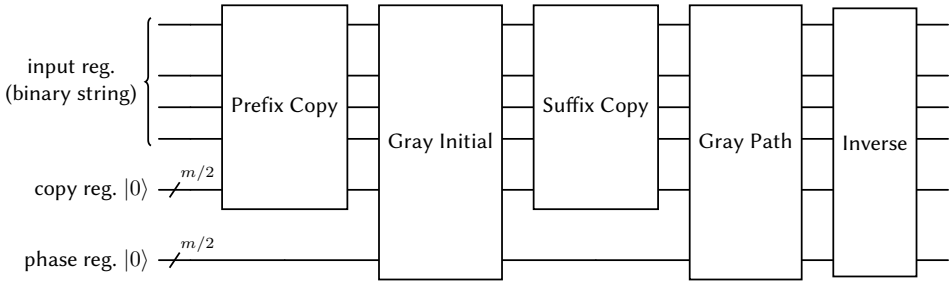

Figure 3: Λ_n quantum circuit divided into 5 sub-unitaries (stages).

3. The decomposition with a single Λ-operator for each UCG

As proved in [14] for the exact paradigm, and also in [28] for the variational one, the traditional structure for quantum state preparation can be split into two unitaries: one devoted to the preparation of the real part of the desired state and the other to the complex component. This separation is allowed by the conventional choice of initializing the state preparation circuits with $|0\rangle^{\otimes n}$, which in turn implies that only the first column of the total QSP unitary is relevant for preparation purposes. In this way, by concatenating the circuit block U_{mod} responsible for the preparation of the real part with a diagonal unitary D_{ph} of complex phases, a complete QSP circuit is obtained, as illustrated in Figure 4. The

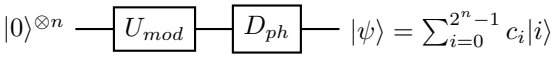

Figure 4: QSP structure with the separation of the real part from the complex one.

unitary U_{mod} can be implemented through the traditional QSP framework [14], which is composed of Uniformly Controlled Gates (UCGs) arranged in a ladder across the entire quantum register; but, since it only needs to prepare the real part of the desired state, each UCG is a R_y-block diagonal matrix $U_{n,y}$, where each block is $V_j = \cancel{e^{i\alpha_j} R_z(\beta_j)} R_y(\gamma_j) \cancel{R_z(\delta_j)} \in SU(2)$. In order to effectively parallelize according to SUN, gate set switching $R_y(\gamma_j) \equiv SH \cdot R_z(\gamma_j) \cdot HS^\dagger$ is mandatory, and now, contrary to the decomposition of Equation (1), each UCG $U_{n,y}$ is encoded in a single diagonal operator $F_n[R_z]$, as illustrated in Figure 5. In analogy to Figure 2, the circuit corresponding to the controlled operator $U_{n,y}$, which exploits the parallelization on the ancillary register with a single matrix $\Lambda_n^{(2)}$, is represented in

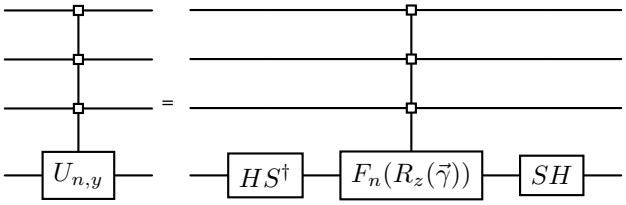

Figure 5: Reduced quantum circuit for each UCG-level inside U_{mod}.

Figure 6. For this reason, this algorithm can be called *single-Λ ancillae-based QSP*. In [25, Section 2] it

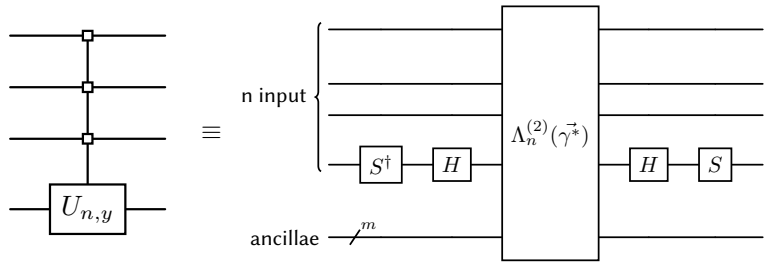

Figure 6: UCG $U_{n,y}$ in terms of a single Λ-type operators.

is shown that the ladder structure of UCG induces a simple recursive scheme on the first column of the total unitary, useful for deriving the parameters of the R_y gates according to the traditional Binary Search Tree (BST) method. On the other hand, the diagonal unitary D_{ph} can be implemented directly via a Λ-type operator after an appropriate global phase recollection. If $|\psi\rangle = \sum_{k=0}^{2^n-1} |c_k| e^{i\phi_k} |k\rangle$, the complex phases $e^{i(\phi_k - \phi_0)}$ will be the diagonal entries of D_{ph} for $0 \leqslant k \leqslant 2^n - 1$, where $e^{i\phi_0}$ is the global phase.

Ultimately, according to the decomposition summarized in the 2-unitary block structure of Figure 4, the overall reduced-depth quantum circuit for the preparation of a 2-qubit (4-ancilla) arbitrary state is depicted in Figure 7.

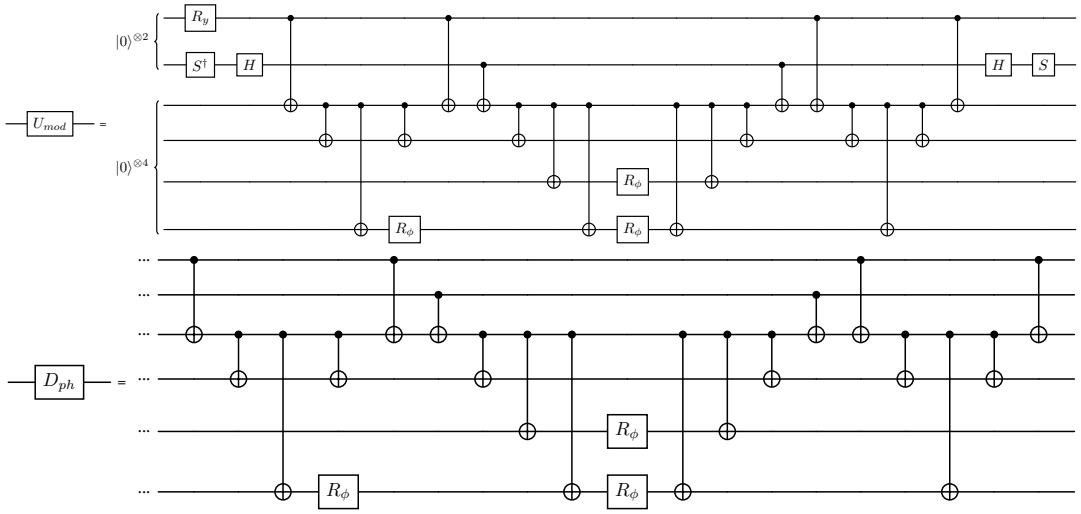

Figure 7: Single-Λ ancillae-based QSP circuit for $n = 2$, divided into the two unitary blocks U_{mod} and D_{ph}.

3.1. Complexity analysis for OSUN

The optimized version of SUN described in Section 3, denoted as OSUN, uses only one (instead of three) circuit block Λ_k for each UCG dedicated to the real part of the desired state, and a final block Λ_n for the imaginary part, in the parametric range $2n \leqslant m \leqslant \frac{2^n}{n}$. Therefore, the whole n-qubit QSP circuit has depth $D(n,m) = \sum_{k=1}^{n} [D_\Lambda(k,m) + 4] + 1 + D_\Lambda(n,m)$, which essentially leads to the same complexity: the improvement term is hidden in the constant factors of the two main asymptotic contributions. For this reason, reproducing the complexity analysis of SUN using a model with *explicit constants* will allow us to formally define the improvement.

SUN complexity model with explicit constant

In the parametric range $2n \leqslant m \leqslant \frac{2^n}{n}$, each operator Λ_j has depth $D_\Lambda^{anc}(j,m) \leqslant a\log_2 m + b\frac{2^j}{m}$ for any real numbers $a, b > 0$ and $j \in [1,n]$. In each level j, the UCG U_j also counts four single-qubit layers and one global phase, for a total cost indicated below as c. When no ancillary qubits are available, each operator Λ_j has depth $D_\Lambda^{no-anc}(j) \leqslant d\frac{2^j}{j}$, for any real $d > 0$. Therefore, the construction of a single smooth bound $m \geqslant 0$ for the depth of each UCG U_j leads one to consider:

$$D_j^{case}(m) \leqslant 3D_\Lambda^{case}(j,m) + c = \begin{cases} 3a\log_2 m + 3b\frac{2^j}{m} + c & \text{case} = anc \\ 3d\frac{2^j}{j} + c & \text{case} = no-anc \end{cases}$$

In the following, it is shown that there exists a real $k > 0$ such that $D_j^{case}(m) \leqslant k(n + \frac{2^n}{n+m})$ for $m \geqslant 0$, considering two cases:

a) in the range $2n \leqslant m \leqslant \frac{2^n}{n}$, $\log_2 m \leqslant \log_2(\frac{2^n}{n}) = n - \log_2 n \leqslant n$, thus $3a\log_2 m \leqslant 3an$. Instead, $\frac{2^j}{m} = \frac{m+n}{m}\cdot\frac{2^j}{m+n} \leqslant \frac{3}{2}\frac{2^j}{m+n}$, thus $3b\frac{2^j}{m} \leqslant \frac{9}{2}b\frac{2^j}{m+n}$. Therefore,

$$D_j^{anc}(m) \leqslant 3an + \frac{9b}{2}\frac{2^j}{m+n} + c \leqslant k_1\left(n + \frac{2^n}{m+n}\right)$$

for $k_1 \geqslant \max\{3a, \frac{9b}{2}, c\}$.

b) in the range $0 \leqslant m < 2n$, it is simple to see that $\frac{2^j}{j+m} \in (\frac{2^j}{2n+j}, \frac{2^j}{j}]$, since $j \leqslant j+m < 2n+j$. Then $\frac{2^j}{j} \leqslant (\frac{2n}{j}+1)\frac{2^j}{j+m}$ and

$$\begin{aligned} D_j^{no-anc}(m) &\leqslant 3d\frac{2^j}{j} + c \leqslant 3d\left(\frac{2n}{j}+1\right)\frac{2^j}{j+m} + c \leqslant 3d\left(\frac{2n}{j}+1\right)\left(n + \frac{2^n}{n+m}\right) + c \\ &\leqslant k_2\left(n + \frac{2^n}{n+m}\right) \end{aligned}$$

for $k_2 \geqslant \max\left\{3d\left(\frac{2n}{j}+1\right), c\right\}$.

As a consequence, taking $k = \max\{k_1, k_2\}$, $D_j^{case}(m) \in \mathcal{O}(n + \frac{2^n}{n+m})$.

The total depth of the circuit can be computed, respectively, in the range $2n \leqslant m \leqslant \frac{2^n}{n}$ by

$$\begin{aligned} D(n,m)^{anc} &= \sum_{j=1}^{n} D_j^{anc}(m) \leqslant 3an\log m + \frac{3b}{m}(2^{n+1}-2) + cn = \mathcal{O}(n^2) + \mathcal{O}\left(\frac{2^n}{m}\right) \\ &= \mathcal{O}\left(n^2 + \frac{2^n}{m}\right) \end{aligned} \tag{4}$$

and in the range $0 \leqslant m < 2n$ by

$$D(n)^{no-anc} = \sum_{j=1}^{n} D_j^{no-anc} \leqslant 3dS_n + cn$$

where $S_n := \sum_{j=1}^{n}\frac{2^j}{j} = \sum_{j=1}^{\lfloor n/2\rfloor}\frac{2^j}{j} + \sum_{\lfloor n/2\rfloor+1}^{n}\frac{2^j}{j}$. In the latter expression, $\sum_{j=1}^{\lfloor n/2\rfloor}\frac{2^j}{j} \leqslant \sum_{j=1}^{\lfloor n/2\rfloor}2^j \leqslant 2^{n/2+1}$ while $\sum_{\lfloor n/2\rfloor+1}^{n}\frac{2^j}{j} \leqslant \frac{2}{n}\sum_{\lfloor n/2\rfloor+1}^{n}2^j \leqslant \frac{2}{n}2^{n+1}$, so that $S_n \leqslant 2^{n/2+1} + \frac{4}{n}2^n \leqslant k_3\frac{2^n}{n}$ and

$$D(n)^{no-anc} \leqslant 3dk_3\frac{2^n}{n} + cn = \mathcal{O}\left(\frac{2^n}{n}\right) \tag{5}$$

Consistency at the edge $m = 2n$ is ensured by $D(n,2n)^{anc} = 3an\log(2n) + \frac{3b}{2n}(2^{n+1}-2) + cn = \mathcal{O}\left(\frac{2^n}{n}\right)$.

OSUN complexity model with explicit constant

Based on the algebraic reduction described in Section 3, valid in the parametric range $m \in [2n, \frac{2^n}{n}]$, the optimized version OSUN has a ladder structure of R_y-based UCGs starting from level $j = 2$. This is because level $j = 1$ only has one uncontrolled gate R_y (refer to Figure 7 as an example), which is possible due to the hypothesis of preparing the real part of the desired state separately from the complex part. This simplification has the following consequences:

- level $j = 1$ has a constant depth $c_0 = 2$;
- for each level $j \in [2, n]$, there are two single-gate layers, which we denote as cost $c_1 = 2$;
- for each level $j \in [2, n]$, there is one Λ-type constructor for the real part, plus another one for $j = n$ regarding the imaginary part of the desired state;
- c_3 is the cost of the global phase.

Then, based on the decomposition summarized by Figure 6, the total depth can be evaluated as

$$D_{opt}^{anc}(n, m) = c_0 + \sum_{j=2}^{n} [D_{\Lambda}^{anc}(j, m) + c_1] + D_{\Lambda}^{anc}(n, m) + c_3$$

where $D_{\Lambda}^{anc}(j, m) \leqslant a \log m + b\frac{2^j}{m}$ for real $a, b > 0$, as before. Considering that $\sum_{j=2}^{n} b\frac{2^j}{m} = \frac{b}{m}(2^{n+1} - 4)$,

$$D_{opt}^{anc}(n, m) \leqslant an \log m + \frac{b}{m}(3 \cdot 2^n - 4) + c_1 n + \mathcal{O}(1) \tag{6}$$

From a quick comparison with Equation 4, it is clear that term $n \log m$ shows a factor of 3 improvement, while term $\frac{2^n}{m}$ shows a factor of 2 improvement; on the contrary, the linear terms are similar and absorbed by $\mathcal{O}(n^2)$ in the global bound. It is also clear that the algebraic reduction has not changed the complexity class: on the range $2n \leqslant m \leqslant \frac{2^n}{n}$, $an \log m \leqslant an^2$, so that $D_{opt}^{anc}(n, m) \in \mathcal{O}(n^2 + \frac{2^n}{m})$. In conclusion, the new construction has the same asymptotic complexity but systematically lowers the prefactors of the two dominant contributions (linear-logarithmic and exponential).

For completeness, the analysis of OSUN in the range $0 \leqslant m < 2n$ is reported below. Applying the same assumptions as in the previous range, the total depth is given by

$$D_{opt}^{no-anc}(n) = c_0 + \sum_{j=2}^{n} [D_{\Lambda}^{no-anc}(j, m) + c_1] + D_{\Lambda}^{no-anc}(n, m) + c_3$$

where now $D_{\Lambda}^{no-anc}(j) \leqslant d\frac{2^j}{j}$, for real $d > 0$. Remembering that $S_n := \sum_{j=1}^{n} \frac{2^j}{j}$, it is worth that $\sum_{j=2}^{n} \frac{2^j}{j} \leqslant S_n \leqslant k\frac{2^n}{n}$, so that

$$D_{opt}^{no-anc}(n) \leqslant d(k + 1)\frac{2^n}{n} + cn + \mathcal{O}(1) = \mathcal{O}\left(\frac{2^n}{n}\right) \tag{7}$$

The comparison with Equation 5 also shows, in this case, the reduction of the prefactor of the dominant term.

4. Simulations and results

The *single-*Λ *ancillae-based* QSP algorithm (OSUN) has been tested using the PennyLane simulator [29] on a Linux machine equipped with an AMD EPYC 7282 CPU and 256 GB of RAM. First, the preparation of some quantum states of interest for research was verified: the four Bell states for 2 qubits, as well as the GHZ, W, and Dicke states for 3 and 4 qubits. The results are reported in Table 1, where the columns "Time Classical" and "Time Quantum" indicate, respectively, the time for determining the parameters (classical subroutine) and the time for the execution of the quantum circuit in a simulated

Table 1
Preparation of well-known quantum states. The results refer to a sample of 20 states.

Simulation results for specific quantum states								
n	Target State	Avg Fidelity $\sigma^2 \leqslant 10^{-32}$	Avg Trace Distance $\sigma^2 = 0.0$	Depth	Total gates	CNOT count	Avg Time Classical (s)	Avg Time Quantum (s)
2	Bell Φ_+	$1 - 2 \times 10^{-16}$	2.78×10^{-16}	20	26	18	1.41×10^{-3}	3.54×10^{-2}
	Bell Φ_-	$1 - 2 \times 10^{-16}$	2.78×10^{-16}	39	47	36	7.11×10^{-3}	4.20×10^{-2}
	Bell Ψ_+	$1 - 2 \times 10^{-16}$	2.78×10^{-16}	20	26	18	1.17×10^{-3}	3.77×10^{-2}
	Bell Ψ_-	$1 - 2 \times 10^{-16}$	2.78×10^{-16}	39	47	36	6.89×10^{-3}	2.75×10^{-2}
3	GHZ	$1 - 6 \times 10^{-16}$	4.16×10^{-16}	43	67	48	4.13×10^{-3}	4.26×10^{-2}
	W	$1 - 6 \times 10^{-16}$	3.05×10^{-16}	43	67	48	2.81×10^{-3}	4.84×10^{-2}
	Dicke	$1 - 6 \times 10^{-16}$	3.05×10^{-16}	43	67	48	2.89×10^{-3}	5.81×10^{-2}
4	GHZ	$1 - 9 \times 10^{-16}$	7.49×10^{-16}	74	142	104	8.65×10^{-3}	7.30×10^{-2}
	W	$1 - 9 \times 10^{-16}$	5.83×10^{-16}	74	142	104	6.70×10^{-3}	8.27×10^{-2}
	Dicke	$1 - 9 \times 10^{-16}$	5.27×10^{-16}	74	142	104	6.74×10^{-3}	7.60×10^{-2}

environment. Depending on their amplitudes, not all states require the general circuit of the complex case (see Figure 4 for the general structure and Figure 7 as an example for 2 qubits): states with positive real coefficients can be prepared with a reduced circuit, which corresponds to the only unitary U_{mod} of Figure 4. In a second step, the effectiveness of the algorithm has been verified in the preparation of random quantum states, both dense and sparse, with complex coefficients, from 2 to 5 qubits. The results are reported in Table 2.

Table 2
Preparation of random quantum states with complex coefficients. The results refer to a sample of 20 states.

Simulation results for random states with complex amplitudes							
n	Type of state	Avg Fidelity $\sigma^2 \in [10^{-32}, 10^{-31}]$	Avg Trace Distance $\sigma^2 \in [10^{-33}, 10^{-32}]$	Depth	CNOT count	Avg Time Classical (s)	Avg Time Quantum (s)
2	dense	$1 - 2 \times 10^{-16}$	2.43×10^{-16}	39	36	7.88×10^{-3}	6.11×10^{-2}
	sparse	$1 - 3 \times 10^{-16}$	2.58×10^{-16}			1.06×10^{-2}	3.87×10^{-2}
3	dense	$1 - 6 \times 10^{-16}$	3.79×10^{-16}	66	78	7.21×10^{-3}	8.90×10^{-2}
	sparse	$1 - 3 \times 10^{-16}$	3.07×10^{-16}			3.69×10^{-3}	5.92×10^{-2}
4	dense	$1 - 9 \times 10^{-16}$	5.71×10^{-16}	105	160	8.46×10^{-3}	8.91×10^{-2}
	sparse	$1 - 9 \times 10^{-16}$	5.48×10^{-16}			8.56×10^{-3}	8.64×10^{-2}
5	dense	$1 - 2 \times 10^{-15}$	7.48×10^{-16}	166	276	6.30×10^{-2}	1.55×10^{-1}
	sparse	$1 - 2 \times 10^{-15}$	7.47×10^{-16}			5.05×10^{-2}	1.39×10^{-1}

Finally, the novel QSP algorithm has been tested from 6 to 10 qubits for random states, both dense and sparse, with complex coefficients. The results are reported in Table 3. The *single-*Λ *ancillae-based* QSP algorithm has been compared with other QSP algorithms: i) the original SUN algorithm where each UCG is implemented with three Λ-operators: ii) the well-known algorithm by Möttönen et al. [14], here denoted as MOTT, which is considered a reference standard for exact QSP without ancillary qubits; iii) the variational algorithm based on a quantum neural network algebraically structured on the Standard Recursive Block Basis (SRBB), here denoted as VQSP. The comparison with MOTT and VQSP is meaningful because both algorithms are based on the traditional QSP structure defined by the synthesis of the UCG levels (Figure 1a). Some metrics are shown in Figures 8, 9, 10, and 11. A comparison with the developments that followed the algorithm of Möttönen et al., such as those proposed by Bergholm [16], by Plesch and Brukner [18], or with the generalizations introduced by Iten [19], has not been covered in this work and will be the heart of a more systematic future benchmarking[2]. The reduced decomposition proposed in this work shows a significant reduction in depth already at the lower end of the first range of [22, Theorem 1], where the auxiliary register grows only polynomially ($m = 2n$). Other points in Sun et al.'s parametric domain may show even greater reductions in light of the strategy here proposed.

[2]The authors are aware of the recent result published by Carvalho et al. [30] on the UCG class, which will also be compared.

Table 3
Preparation of random quantum states with complex coefficients. The results refer to a sample of 20 states.

Simulation results for random states with complex amplitudes							
n	Type of state	Avg Fidelity $\sigma^2 \in [10^{-31}, 10^{-30}]$	Avg Trace Distance $\sigma^2 \in [10^{-32}, 10^{-31}]$	Depth	CNOT count	Avg Time Classical (s)	Avg Time Quantum (s)
6	dense	$1 - 2 \times 10^{-15}$	8.22×10^{-16}	246	510	1.43×10^{-1}	4.23×10^{-1}
	sparse	$1 - 3 \times 10^{-15}$	8.99×10^{-16}			1.62×10^{-1}	4.30×10^{-1}
7	dense	$1 - 3 \times 10^{-15}$	1.11×10^{15}	427	930	3.40×10^{-1}	2.22
	sparse	$1 - 3 \times 10^{-15}$	1.21×10^{-15}			3.88×10^{-1}	2.23
8	dense	$1 - 3 \times 10^{-15}$	1.32×10^{-15}	647	1748	9.74×10^{-1}	4.23×10^{2}
	sparse	$1 - 3 \times 10^{-15}$	1.11×10^{-15}			8.98×10^{-1}	4.53×10^{2}
9	dense	$1 - 3 \times 10^{-15}$	1.66×10^{-15}	1046	3354	3.70	2.63×10^{3}
	sparse	$1 - 3 \times 10^{-15}$	1.30×10^{-15}			3.59	2.58×10^{3}
10	dense	$1 - 3 \times 10^{-15}$	1.55×10^{-15}	1871	6486	14.4	1.97×10^{4}
	sparse	$1 - 3 \times 10^{-15}$	1.52×10^{-15}			14.9	2.06×10^{4}

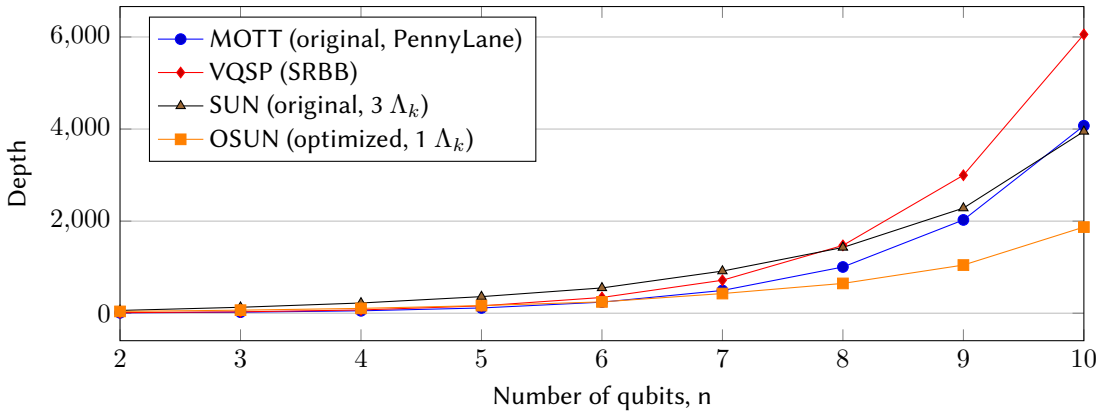

Figure 8: Depth vs Number of qubits.

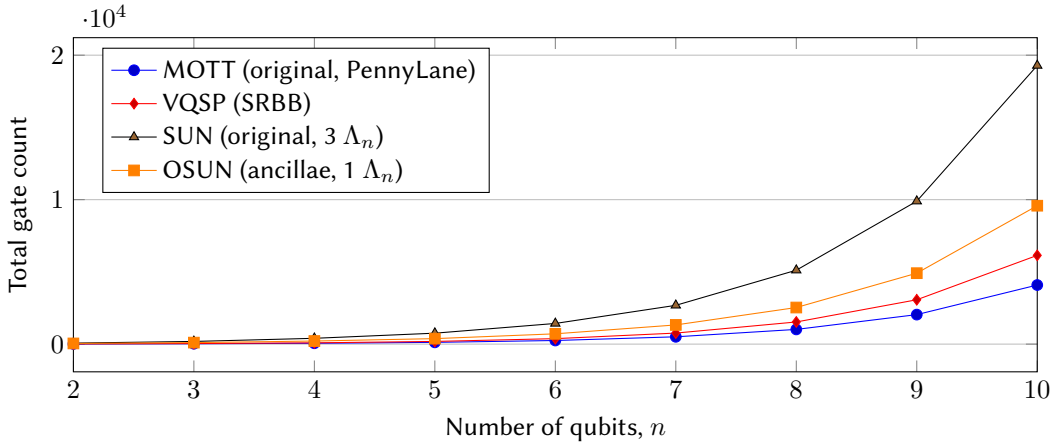

Figure 9: Total gates vs Number of qubits.

5. Conclusions

In the traditional QSP framework based on increasing UCG levels, the hypothesis of preparing the real part of the desired state separately from the complex one leads to an algebraic simplification. In fact, if the UCG ladder structure is entrusted with the task of preparing only the real moduli of each component of the target state, each UCG can be built with only R_y-blocks. This reduction has a precise

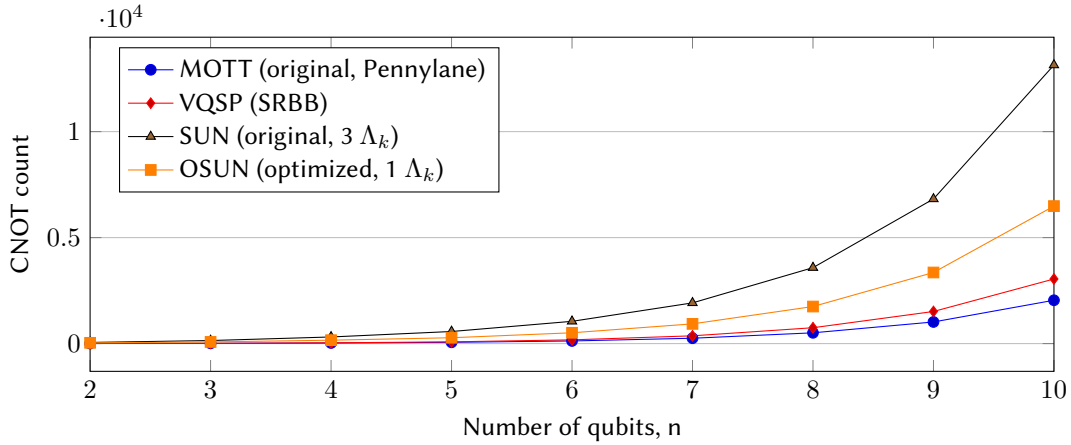

Figure 10: CNOT vs Number of qubits.

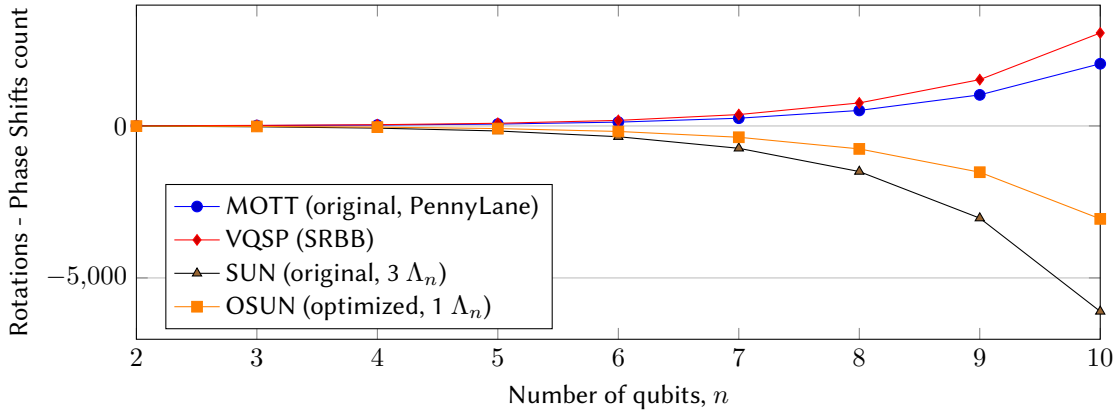

Figure 11: Difference in gate set vs Number of qubits.

consequence for the QSP algorithm SUN, known in the literature for defining the optimal asymptotic space-time trade-off bounds: for $2n \leqslant m \leqslant \frac{2^n}{n}$ ancillary qubits, each UCG can be implemented with one Λ operator (instead of 3), leading to systematic depth improvements. The complex part of the desired state is instead prepared with a single, maximum-dimensional Λ operator at the end of the UCG scheme. This new QSP algorithm, denoted as OSUN, does not change the complexity of SUN but improves the prefactors of the dominant terms, specifically by a factor of 3 in the linear term and by a factor of 2 in the exponential term. It is important to note that this algebraic reduction applies to the entire first parametric range of [22, Theorem 1], so future extensions of the algorithm to points in the domain where the ancillary register grows exponentially would benefit from the same optimization.

The OSUN algorithm has been implemented with the PennyLane library, and the code is available on GitHub [24]. The effectiveness of the algorithm has been verified in simulations involving up to 10 qubits for various types of target states. For each case, metrics such as depth, number of CNOTs, classical precomputation time, and quantum circuit simulation time have been reported. Furthermore, the OSUN algorithm has been compared with the original version SUN and with other QSP algorithms based on the same UCG scheme but without ancillary qubits. The results clearly show that OSUN has a shallower depth with a smoother growth curve compared to the original version, thanks to the optimization of the prefactors of the dominant terms.

Future developments could concern the extension of OSUN to other points in the parametric domain to better exploit the depth reductions afforded by the ancillary parallelization. Other topics to consider are qubit connectivity and how it affects performance on real hardware, or to what extent a translation in the Clifford+T gate set could affect the complexity.

Acknowledgments

Giacomo Belli and Michele Amoretti acknowledge financial support from the European Union - NextGenerationEU, PNRR MUR project PE0000023-NQSTI. This research benefits from the High Performance Computing facility of the University of Parma, Italy (HPC.unipr.it).

Declaration on Generative AI

The author(s) have not employed any Generative AI tools.

References

[1] J. Iaconis, S. Johri, E. Y. Zhu, Quantum state preparation of normal distributions using matrix product states, npj Quantum Information 10 (2024) 15.

[2] M. E. Morales, L. Pira, P. Schleich, K. Koor, P. Costa, D. An, A. Aspuru-Guzik, L. Lin, P. Rebentrost, D. W. Berry, Quantum linear system solvers: A survey of algorithms and applications, arXiv preprint arXiv:2411.02522 (2024).

[3] D. W. Berry, M. Kieferová, A. Scherer, Y. R. Sanders, G. H. Low, N. Wiebe, C. Gidney, R. Babbush, Improved techniques for preparing eigenstates of fermionic hamiltonians, npj Quantum Information 4 (2018) 22.

[4] D. W. Berry, Y. Tong, T. Khattar, A. White, T. I. Kim, G. H. Low, S. Boixo, Z. Ding, L. Lin, S. Lee, et al., Rapid initial-state preparation for the quantum simulation of strongly correlated molecules, PRX Quantum 6 (2025) 020327.

[5] M. Rath, H. Date, Quantum data encoding: A comparative analysis of classical-to-quantum mapping techniques and their impact on machine learning accuracy, EPJ Quantum Technology 11 (2024) 72.

[6] K. Gui, A. M. Dalzell, A. Achille, M. Suchara, F. T. Chong, Spacetime-efficient low-depth quantum state preparation with applications, Quantum 8 (2024) 1257.

[7] D. Malz, G. Styliaris, Z.-Y. Wei, J. I. Cirac, Preparation of matrix product states with log-depth quantum circuits, Physical Review Letters 132 (2024) 040404.

[8] R. Mao, G. Tian, X. Sun, Toward optimal circuit size for sparse quantum state preparation, Physical Review A 110 (2024) 032439.

[9] S. Lee, J. Lee, H. Zhai, Y. Tong, A. M. Dalzell, A. Kumar, P. Helms, J. Gray, Z.-H. Cui, W. Liu, et al., Evaluating the evidence for exponential quantum advantage in ground-state quantum chemistry, Nature communications 14 (2023) 1952.

[10] P. Yuan, S. Zhang, Optimal (controlled) quantum state preparation and improved unitary synthesis by quantum circuits with any number of ancillary qubits, Quantum 7 (2023) 956.

[11] M. A. Nielsen, I. L. Chuang, Quantum Computation and Quantum Information, 2000.

[12] L. Grover, T. Rudolph, Creating superpositions that correspond to efficiently integrable probability distributions, arXiv preprint quant-ph/0208112 (2002).

[13] M. Möttönen, J. J. Vartiainen, Decompositions of general quantum gates, Trends in quantum computing research (2006) 149.

[14] M. Mottonen, J. Vartiainen, V. Bergholm, M. Salomaa, Transformation of quantum states using uniformly controlled rotations, Quantum Information and Computation 5 (2005) 467–473.

[15] J. J. Vartiainen, M. Möttönen, M. M. Salomaa, Efficient decomposition of quantum gates, Physical review letters 92 (2004) 177902.

[16] V. Bergholm, J. J. Vartiainen, M. Möttönen, M. M. Salomaa, Quantum circuits with uniformly controlled one-qubit gates, Physical Review A—Atomic, Molecular, and Optical Physics 71 (2005) 052330.

[17] V. Shende, S. Bullock, I. Markov, Synthesis of quantum-logic circuits, IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems 25 (2006) 1000–1010.

[18] M. Plesch, Č. Brukner, Quantum-state preparation with universal gate decompositions, Physical Review A—Atomic, Molecular, and Optical Physics 83 (2011) 032302.
[19] R. Iten, R. Colbeck, I. Kukuljan, J. Home, M. Christandl, Quantum circuits for isometries, Physical Review A 93 (2016) 032318.
[20] X.-M. Zhang, M.-H. Yung, X. Yuan, Low-depth quantum state preparation, Physical Review Research 3 (2021) 043200.
[21] G. Rosenthal, Query and depth upper bounds for quantum unitaries via grover search, arXiv preprint arXiv:2111.07992 (2021).
[22] X. Sun, G. Tian, S. Yang, P. Yuan, S. Zhang, Asymptotically optimal circuit depth for quantum state preparation and general unitary synthesis, IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems 42 (2023) 3301–3314.
[23] X.-M. Zhang, T. Li, X. Yuan, Quantum state preparation with optimal circuit depth: Implementations and applications, Physical Review Letters 129 (2022) 230504.
[24] G. Belli, A. Bersellini, M. Amoretti, Quantum State Preparation with Λ-type Operators (qsp-sun), https://github.com/qslab-unipr/qsp-sun, 2024.
[25] G. Belli, A. Bersellini, M. Amoretti, Implementation of an optimally bounded algorithm for quantum state preparation, in: International Conference on Reversible Computation, Springer, 2025, pp. 37–53.
[26] G. Belli, M. Amoretti, Algebraic reduction to improve an optimally bounded quantum state preparation algorithm, arXiv preprint arXiv:2602.06535 (2026).
[27] M. Möttönen, J. J. Vartiainen, V. Bergholm, M. M. Salomaa, Quantum circuits for general multiqubit gates, Physical review letters 93 (2004) 130502.
[28] G. Belli, M. Mordacci, M. Amoretti, Srbb-based quantum state preparation, in: Proceedings of the 22nd ACM International Conference on Computing Frontiers, 2025, pp. 172–175.
[29] V. Bergholm, J. Izaac, M. Schuld, C. Gogolin, S. Ahmed, V. Ajith, M. S. Alam, G. Alonso-Linaje, B. AkashNarayanan, A. Asadi, et al., Pennylane: Automatic differentiation of hybrid quantum-classical computations, arXiv preprint arXiv:1811.04968 (2018).
[30] J. A. de Carvalho, C. Batista, T. de Veras, I. Araujo, A. J. da Silva, Quantum multiplexer simplification for state preparation, ACM Transactions on Quantum Computing 6 (2025) 1–12.