



UNIVERSIDADE FEDERAL DE PERNAMBUCO
CENTRO DE INFORMÁTICA
PROGRAMA DE PÓS-GRADUAÇÃO CIÊNCIA DA COMPUTAÇÃO

Rafaella Ferreira do Vale

Decomposition of multicontrolled single-qubit special unitary operators

Recife

2026

Rafaella Ferreira do Vale

Decomposition of multicontrolled single-qubit special unitary operators

Tese apresentada ao Programa de Pós-graduação em Ciência da Computação do Centro de Informática da Universidade Federal de Pernambuco, como requisito parcial para obtenção do grau de Doutor em Ciência da Computação.

Área de Concentração: Inteligência Computacional

Orientador: Adenilton José da Silva

Recife

2026

.Catalogação de Publicação na Fonte. UFPE - Biblioteca Central

Vale, Rafaela Ferreira do.

Decomposition of multicontrolled single-qubit special unitary operators / Rafaela Ferreira do Vale. - Recife, 2026.
75f.: il.

Tese (Doutorado) - Universidade Federal de Pernambuco, Centro de Informática, Programa de Pós-Graduação em Ciência da Computação, 2026.

Orientação: Adenilton José da Silva.

Inclui Referências.

1. Computação Quântica; 2. Circuitos Quânticos; 3. Otimização de Circuitos Quânticos; 4. Decomposição de Portas Quânticas. I. Silva, Adenilton José da. II. Título.

UFPE-Biblioteca Central

Rafaella Ferreira do Vale

“Decomposition of multicontrolled single-qubit special unitary operators”

Tese de Doutorado apresentada ao Programa de Pós-Graduação em Ciência da Computação da Universidade Federal de Pernambuco, como requisito parcial para a obtenção do título de Doutor em Ciência da Computação. Área de Concentração: Engenharia de Software e Linguagens de Programação

Aprovada em: 27/03/2026.

Orientador: Prof. Dr. Adenilton José da Silva

BANCA EXAMINADORA

Prof. Dr. Stefan Michael Blawid
Centro de Informática / UFPE

Profa. Dra. Regina Melo Silveira
Depto. de Engenharia de Computação e
Sistemas Digitais / USP

Prof. Dr. Israel Ferraz de Araújo
Data Cybernetics

Prof. Dr. Wilson Rosa de Oliveira Junior
Depto. de Estatística e Informática / UFRPE

Profa. Dra. Jerusa Marchi
Depto. de Estatística e Informática / UFSC

ACKNOWLEDGEMENTS

This thesis would not have been possible without the support of my family and my friends. Thank you for believing in me even when the path seemed uncertain. I would like to extend my sincere gratitude to my supervisor, Prof. Adenilton Silva, for your invaluable mentorship and your patience throughout this project. I am also grateful to the defense committee members, Israel Ferraz, Wilson Oliveira, Stefan Michael Blawid, Regina Silveira, and Jerusa Marchi, for their time, attention, and constructive evaluation of this thesis. Finally, I would like to acknowledge Instituto Nacional de Ciência e Tecnologia em Computação Quântica Aplicada.

STATEMENT ON THE USE OF ARTIFICIAL INTELLIGENCE

The author declares that generative artificial intelligence (AI) tools were used in a limited and supportive role during the preparation of this manuscript. Specifically, such tools were employed for language refinement and English editing, including improvements in grammar, clarity, and overall readability of the text. No AI tools were used to generate original scientific content, results, figures, or data presented in this work. All scientific ideas, methodologies, analyses, and conclusions remain the sole responsibility of the author. The author has carefully reviewed and validated the final manuscript to ensure its accuracy, originality, and compliance with ethical standards, including the avoidance of plagiarism. Generative AI tools did not contribute intellectually to the research.

RESUMO

Avanços recentes na computação quântica têm recebido atenção significativa da comunidade científica, impulsionando intensas iniciativas em pesquisa e desenvolvimento. Tais esforços visam aproveitar fenômenos quânticos na resolução de problemas computacionalmente desafiadores. Na busca por explorar esse potencial, a investigação de estratégias para superar as limitações de dispositivos quânticos ruidosos de escala intermediária (NISQ) emergiu como uma frente de pesquisa decisiva. Essa linha de pesquisa tem o potencial de viabilizar aplicações práticas de tecnologias quânticas no curto prazo. Este trabalho concentra-se na decomposição de operadores quânticos multicontrolados para auxiliar na mitigação de algumas das restrições de dispositivos NISQ. Em particular, são estudadas versões multicontroladas de operadores $SU(2)$ de um qubit, com n qubits de controle. Aprimoramos a síntese desses operadores em relação ao estado da arte à época da proposição deste trabalho, reduzindo o número de portas CNOT necessárias de $28n - 88$ (n par) ou $28n - 92$ (n ímpar) para $20n - 38$ (n ímpar) ou $20n - 42$ (n par)—ou para $16n$, caso a matriz do operador unitário especial apresente diagonal principal ou secundária com elementos reais. Esse resultado motiva investigações futuras em busca de abordagens ainda mais eficientes para a decomposição de operadores de um qubit com múltiplos controles.

Palavras-chave: Computação quântica. Circuitos quânticos. Otimização de circuitos quânticos. Decomposição de portas quânticas.

ABSTRACT

Recent advancements in quantum computing have received significant attention from the scientific community, driving intensive research and development initiatives. These efforts aim to leverage quantum phenomena to address computationally challenging problems. With the intention of exploring this potential, the investigation of strategies to overcome the limitations of noisy intermediate-scale quantum (NISQ) devices has emerged as a crucial research frontier. This line of research holds promise for paving the way for more practical applications of short-term quantum technologies. This work focuses on the decomposition of multicontrolled quantum operators to help mitigate some of these constraints of NISQ devices. In particular, multicontrolled single-qubit $SU(2)$ operators with n control qubits are of interest. We have improved the synthesis of these operators compared to the state-of-the-art at the time of proposal of this work, lowering the required number of CNOT gates from $28n - 88$ (for even n) or $28n - 92$ (for odd n) to $20n - 38$ (for odd n) or $20n - 42$ (for even n)—or to $16n - 40$ if the special unitary operator's matrix contains real-valued main diagonal or off-diagonal). This result motivates future investigations in the search for more efficient approaches to decompose any multicontrolled single-qubit operator.

Keywords: Quantum computing. Quantum circuits. Quantum circuit optimization. Quantum gate decomposition.

LIST OF FIGURES

Figure 1 – Example of quantum circuit diagram for the 3-qubit operator $E(C \otimes D)(B \otimes I_2)(A \otimes I_4)$. Source: The author (2026).	18
Figure 2 – Decomposition scheme for the Toffoli operator. Source: Adapted from Iten et al. (2016).	26
Figure 3 – Decomposition scheme for the relative phase Toffoli operator. Source: Adapted from Iten et al. (2016).	27
Figure 4 – The quantum circuit for a simplified version of Deutsch's algorithm (DEUTSCH, 1985). Source: Adapted from Nielsen and Chuang (2010).	29
Figure 5 – Barenco et al. (1995) circuit for k -controlled σ_x gates with $k - 2$ auxiliary qubits. Source: Adapted from Barenco et al. (1995) and (VALE et al., 2024) © 2024 IEEE.	35
Figure 6 – Barenco et al. (1995) circuit for multicontrolled σ_x gates with one auxiliary qubit. Source: Adapted from Barenco et al. (1995) and Vale et al. (2024) © 2024 IEEE.	36
Figure 7 – The “flip target” part of the circuit from Figure 5 following gate canceling optimization. Source: Adapted from Iten et al. (2016).	37
Figure 8 – He et al. (2017) decomposition scheme for multicontrolled σ_x gates with one auxiliary qubit. Source: Adapted from He et al. (2017) and Vale et al. (2024) © 2024 IEEE.	38
Figure 9 – Barenco et al. (1995) decomposition scheme for general multicontrolled single-qubit gates. Source: Adapted from Iten et al. (2016).	39
Figure 10 – Silva and Park (2022) decomposition scheme for general multicontrolled single-qubit gates. Source: Adapted from Silva and Park (2022).	41
Figure 11 – Saeedi and Pedram (2013) decomposition scheme for multicontrolled $R_x(\pi)$. Source: Adapted from Saeedi and Pedram (2013).	42
Figure 12 – Barenco et al. (1995) decomposition scheme for multicontrolled $SU(2)$ gates. Source: Adapted from Barenco et al. (1995).	43
Figure 13 – Region of the multicontrolled $SU(2)$ scheme by Iten et al. (2016) where the canceling of the resetting part of the $C_{k_2}(\sigma_x)$ gates happens. Source: Adapted from Iten et al. (2016).	44

Figure 14 – Initial decomposition scheme, adapted from He et al. (2017), proposed for multicontrolled $SU(2)$ gates whose corresponding matrices contain a real-valued off diagonal. Source: Adapted from Vale et al. (2024) © 2024 IEEE. 46

Figure 15 – Quantum circuit of the U_1 operator restricted to 4 qubits (ACASIETE et al., 2020). Source: Adapted from Acasiete et al. (2020). 54

Figure 16 – The decomposition of multicontrolled Q and a possible circuit optimization. Source: Adapted from Vale et al. (2024) © 2024 IEEE. 57

Figure 17 – Proposed generalized decomposition scheme for multicontrolled $SU(2)$ gates. Source: Adapted from Vale et al. (2024) © 2024 IEEE. 58

LIST OF TABLES

Table 1 – A non-exhaustive list of single-qubit and multiqubit gates and their respective representations as a matrix.	24
Table 2 – CNOT gate count and required number of auxiliary qubits of the different decomposition schemes for multicontrolled single-qubit gates in the related literature.	34
Table 3 – Comparison of CNOT gate count and required number of auxiliary qubits of this proposal and the different decomposition schemes for multicontrolled $U(2)$ and $SU(2)$ gates in the related literature.	63

CONTENTS

1	INTRODUCTION	13
1.1	MOTIVATION	15
1.2	OBJECTIVES	17
2	BACKGROUND	18
2.1	SINGLE-QUBIT GATES	20
2.2	MULTIQUBIT GATES	24
2.3	IMPORTANT PROPERTIES IN QUANTUM CIRCUITS	28
2.4	CHALLENGES IN NEAR-TERM QUANTUM COMPUTING	30
3	REVIEW OF LITERATURE	34
3.1	DECOMPOSITIONS FOR MULTICONTROLLED σ_x GATES	34
3.1.1	Barenco et al. (1995)	35
3.1.2	Iten et al. (2016)	36
3.1.3	He et al. (2017)	37
3.2	DECOMPOSITIONS FOR MULTICONTROLLED SINGLE-QUBIT GATES	39
3.2.1	Barenco et al. (1995)	39
3.2.2	Silva and Park (2022)	40
3.3	DECOMPOSITIONS FOR MULTICONTROLLED $SU(2)$ GATES	42
3.3.1	Barenco et al. (1995)	43
3.3.2	Iten et al. (2016)	43
4	PROPOSED DECOMPOSITION SCHEME FOR MULTICONTROLLED SPECIAL UNITARY SINGLE-QUBIT OPERATORS	45
4.1	DECOMPOSITION SCHEME FOR MULTICONTROLLED $SU(2)$ GATES WITH ONE REAL-VALUED DIAGONAL	45
4.1.1	Cost	52
4.1.2	Application to multicontrolled R_x, R_y, and R_z gates	53
4.1.3	Application example	53
4.2	GENERALIZED DECOMPOSITION SCHEME FOR ANY MULTICONTROLLED $SU(2)$ OPERATOR	55
4.2.1	Cost	58
5	DISCUSSION	60

6	CONCLUSION	64
	REFERENCES	67

1 INTRODUCTION

Efforts devoted to quantum computing research are proving fruitful and providing promising applications. Quantum algorithms suggest potential speedups for tasks such as factoring (SHOR, 1999), search in unstructured databases (GROVER, 1996), machine learning (BLAMONTE et al., 2017; FARHI; NEVEN, 2018; HAVLÍČEK et al., 2019), cryptography (GISIN et al., 2002; PIRANDOLA et al., 2020), chemistry (PERUZZO et al., 2014; KANDALA et al., 2017; CAO et al., 2019), and finance (ORÚS; MUGEL; LIZASO, 2019; HERMAN et al., 2022). The advantages discovered in quantum computing tap into pivotal phenomena in quantum mechanics. However, realizing useful speedups in practice requires that implementations make efficient use of resources and error rates are low enough to support large computations (BRAVYI et al., 2024; ACHARYA et al., 2025).

Despite the positive outlook, significant obstacles to quantum technology must be overcome before the capabilities of quantum computers can be leveraged to tackle these potential applications. There are significant challenges to building perfectly isolated quantum devices, and techniques to reduce the impact of these issues are, for now, also constrained by current technology (ALMUDEVER et al., 2017; PRESKILL, 2018; CÓRCOLES et al., 2019). To alleviate the effect of these technical obstacles, we need to join efforts in both hardware and software advancement as development continues. In parallel, interdisciplinary research should also advance while the necessary level of technology for practical applications is still being developed.

In recent years, there has been a shift in the field from noisy intermediate-scale quantum (NISQ) experiments towards demonstrating components needed for fault-tolerant quantum computing, especially quantum error correction (QEC) and reliable logical qubits. This transition is reflected in the changing objectives being pursued by researchers. Early milestones emphasized quantum supremacy, targeting demonstrations of outperformance by a large margin on tasks that are very hard for classical computing. A prime example was Google's experiment performed in 200 seconds on the 53-qubit Sycamore processor (ARUTE et al., 2019), a random quantum circuit sampling task estimated to take a classical supercomputer approximately 10,000 years to complete.

As the emphasis shifted to advantage, experiments increasingly sought to show that quantum processors can outperform classical methods under more realistic conditions and with more practical applicability. In this vein, Google's Quantum Echoes algorithm (Google Quantum AI and

Collaborators et al., 2025), implemented on the Willow superconducting chip (ACHARYA et al., 2025), measured out-of-time order correlators (OTOC) in a way that verifiably showed a path to practical quantum advantage. The significance of testing on the Willow architecture lies in the previous demonstration of a below-threshold surface code implementation (ACHARYA et al., 2025), where logical error rates decrease with increasing code distance. These results also exhibit repeatable performance without deterioration, which are required capabilities for full fault tolerance. Advantage was also reported with experiments on the Jiuzhang 4.0 photonic quantum processor (LIU et al., 2025), overcoming photon loss. In this work, the device was used to execute a Gaussian boson sampling task in 25.6 microseconds, which is estimated to require over 10^{42} years in a state-of-the-art supercomputer. Microsoft and Quantinuum's experiments (PAETZNICK et al., 2024) on a trapped-ion processor showed repeated fault-tolerant error correction while maintaining logical error rates below the physical error rates. These results are another step in the move from NISQ computing towards reliable quantum computing.

Contrasting with the advantage focus, IBM's contribution with its 120-qubit Nighthawk processor (KIM et al., 2023) was to measure accurate expectation values for circuit volumes, surpassing brute-force classical computation. The aim was to provide evidence for the utility of noisy quantum computing before fault-tolerant computing, without selecting problems that could take advantage from the demonstrated capability. Overall, the move from NISQ to fault-tolerant computing reflects a continuous evolution in objectives guided by advances in technical capabilities. As the field matures, improvements in qubit quality, gate fidelity, measurement, connectivity, control, and calibration mark quantum computing's convergence towards the goal of large-scale fault tolerance.

The method proposed here is only one of many possible ways to mitigate issues found in near-term devices without directly controlling the physical implementation. From the software side, we can take different approaches. Usually, they pertain to adjustments made to quantum circuits to produce an equivalent but more efficient quantum circuit. One of the problems that we can pursue involves converting multiqubit gates to efficient quantum circuits, as opposed to using universal techniques that are computationally more expensive. In tasks of this nature, aiming for a lower number of quantum gates is a common objective for mitigating errors on real hardware in the current era as physical quantum gates get more accurate (PRESKILL, 2018; ALMUDEVER et al., 2017; SAKI; ALAM; GHOSH, 2021).

We aim to study and propose new methods to decompose multicontrolled gates. Controlled gates are essential to express conditional operations in quantum circuits. The controlled NOT

or CNOT gate, for example, is one of the more frequently used operations in quantum algorithms. Assuming we are working with a universal gate set containing the expensive CNOT gate (BLATT; WINELAND, 2008), counting the number of occurrences of this gate in quantum circuits is a good measure of how costly a circuit is. Quantum gates have an associated gate error due to imperfections in their physical realization in quantum devices (BARENDS et al., 2014). In fact, two-qubit gates dominate the error rate of quantum circuits (TANNU; QURESHI, 2019), the CNOT gate being the most common in quantum gate sets. Other controlled gates, particularly with multiple controls, must be decomposed into elementary gates since most of them are not present in the gate sets of quantum devices. For this research, the focus is initially on special unitary single-qubit quantum gates, or $SU(2)$ gates, whose matrix representations form a group of 2×2 unitary matrices with determinant 1. Specifically, the controlled version of $SU(2)$ gates with multiple control qubits is the object of study of this work.

A brief overview of how this document is structured follows. In the remainder of this chapter, we provide a comprehensive contextualization to motivate the subject of this research and delineate our research objectives. We present the background information in Chapter 2, delving into the key theoretical concepts necessary to understand the foundation of all mentioned works and our proposal. Building upon this groundwork, Chapter 3 offers a comprehensive review of the existing literature, highlighting the most influential methods and how they compare. With this contextual backdrop, we detail our contributions in Chapter 4, outlining the relevant steps and the techniques we employ in our scheme. Chapter 5 discusses the improvements of our approach compared to the literature at the time, and also presents more recent methods that offer an advantage over the proposal. To conclude, Chapter 6 offers the closing remarks, where we not only summarize the main contributions, potential applications, and limitations of our work but also provide an outlook on the prospects of this project and this line of research.

1.1 MOTIVATION

In the noisy intermediate-scale quantum (NISQ) era (PRESKILL, 2018), due to limitations on the physical realization of quantum processors, limited coherence and imperfect control make the execution of deep circuits unreliable. Subsequent applications of faulty gates incur errors, which could be addressed using quantum error correction (SHOR, 1996; CAMPBELL; TERHAL; VUILLOT, 2017) as long as the qubit availability constraint is addressed (ALMUDEVER et al., 2017; PRESKILL, 2018).

In the meantime, as quantum hardware technology advances towards fault tolerance and quantum devices scale up, we must find new ways to mitigate these restrictions. We can approach mitigations from a software perspective (BIAMONTE et al., 2017; ALMUDEVER et al., 2017; PRESKILL, 2018; CÓRCOLES et al., 2019; SIVARAJAH et al., 2020; CONTRIBUTORS, 2023). Since error rates are typically dominated by two-qubit operations (TANNU; QURESHI, 2019), a practical goal is not only to mitigate errors, but also to compile and synthesize controlled operations into circuit forms that reduce two-qubit gate usage and the depth of quantum circuits. Reducing quantum circuit depth means reducing the number of steps of an algorithm, thereby increasing the likelihood of reliably executing the computations by the end of the circuit. We put forth that this synthesis focus remains relevant as the field moves towards fault-tolerant computing.

Another challenge, in addition to faulty quantum gates in general, is to efficiently translate specific operations that do not necessarily have a direct physical implementation. Specifically, this research is concerned with controlled gates. These gates form a fundamental class of quantum gates that allow the expression of conditional operations in quantum circuits. Quantum gates with multiple control qubits pose an essential challenge in the implementation of quantum circuits because they need to be decomposed into elementary quantum gates that are native to the quantum device or that can be synthesized with ease. As much as universal decomposition methods exist (BARENCO et al., 1995; VARTIAINEN; MÖTTÖNEN; SALOMAA, 2004; MÖTTÖNEN et al., 2004; BERGHOLM et al., 2005; SHENDE; BULLOCK; MARKOV, 2005; LI; ROBERTS; YIN, 2013; ITEN et al., 2016; BRUGIÈRE et al., 2020; KROL et al., 2022), generic approaches have an exponential cost (SHENDE; BULLOCK; MARKOV, 2005), whereas those designed for specific classes of gates tend to produce less expensive quantum circuits (BARENCO et al., 1995; BERGHOLM et al., 2005; SHENDE; BULLOCK; MARKOV, 2005; SAEEDI; PEDRAM, 2013; ITEN et al., 2016; LIU; WEI, 2020; MALVETTI; ITEN; COLBECK, 2021; SILVA; PARK, 2022).

Beyond the benefits provided to program runtime in hardware, decomposing particular types of quantum gates in predictable quantum circuits can aid in algorithm design. In fact, modularity in circuit design can benefit from efficient decomposition schemes for specific gates. If an efficient method to decompose a quantum gate is known, the details can be abstracted from the algorithm so that focus can be directed to other important parts of the circuit. Additionally, prior knowledge of the cost in terms of the number of gates resulting from decomposing specific gates using a known method facilitates the calculation of the total cost of the algorithm during design. However, ensuring that the decomposition methods used

in practice reduce the circuit size for real devices should remain a priority.

To sum up, efficient synthesis methods for quantum gates are a requirement of NISQ devices to compensate for the issues with quantum noise and error, both notably challenging phenomena in the field, which can overcome quantum advantage in practice. Furthermore, despite progress towards fault tolerance, realizing practical quantum speedups still depends on how efficiently quantum circuits can be synthesized from the hardware's native gate set. In this critical period of limited gate fidelity and the number of available qubits for computation still growing, when an expensive gate such as the CNOT dominates both gate count and accumulated error, improved decompositions directly translate into lower circuit costs and reduced error propagation. Doing this should lessen the penalty for executing computations on NISQ devices and regain some of the lost accuracy to allow more quantum algorithms to be run on real hardware.

1.2 OBJECTIVES

The objective of this work is to propose a decomposition method for multicontrolled $SU(2)$ gates that achieves lower cost than other available methods in the literature for the same quantum gate class without the need for auxiliary qubits.

The specific objectives are the following:

1. Reduce the required number of CNOTs compared to other approaches.
 - a) Partial objective: Reduce the depth of the circuit compared to other approaches.
2. Apply the proposed method to the multicontrolled versions of the commonly encountered rotation operator gates R_x , R_y , and R_z .

2 BACKGROUND

The quantum circuit model is a common way to describe how manipulating quantum bits transforms a quantum state. The transforming elements are called quantum gates, and, as such, they are the building blocks that make this model a convenient tool for depicting quantum algorithms. The diagram of a quantum circuit consists of horizontal wires representing the quantum bits and a sequence of quantum gates over them, each interacting with specific target quantum bits. Interpreting the transformation in a quantum circuit requires reading the gates sequentially from left to right, exemplified by the three-qubit quantum circuit in Figure 1. Additionally, quantum gates illustrated in the same horizontal position indicate simultaneous execution in one timestep.




Figure 1 – Example of quantum circuit diagram for the 3-qubit operator $E(C \otimes D)(B \otimes I_2)(A \otimes I_4)$. Source: The author (2026).

Also known as qubits, quantum bits are the quantum computing counterparts of classical bits, meaning they are the fundamental unit of quantum information. Unlike classical bits, however, their range of values is not restricted to either 0 or 1. The quantum state of a qubit can be in a superposition of the orthonormal basis states $\begin{pmatrix} 1 \\ 0 \end{pmatrix}$ and $\begin{pmatrix} 0 \\ 1 \end{pmatrix}$. These states form what is typically called the computational basis. We can use the Dirac notation to express the quantum superposition $\begin{pmatrix} \alpha \\ \beta \end{pmatrix}$ as $\alpha|0\rangle + \beta|1\rangle$, where $|0\rangle$ and $|1\rangle$ are the computational basis states. The coefficients α and β are complex numbers called the probability amplitudes of the quantum state. Figure 1 shows a circuit acting on three qubits $|q_0\rangle$, $|q_1\rangle$ and $|q_2\rangle$, where $|q_j\rangle = \alpha_j|0\rangle + \beta_j|1\rangle$.

With qubits being described by a 2-dimensional complex vector, the combined quantum state can be obtained via the tensor product of the qubits in the states $|q_0\rangle$, $|q_1\rangle$, and $|q_2\rangle$. In this 3-qubit state example, the general quantum state of the circuit is then a 2^3 -dimensional

superposition state which can be spanned by the basis states $\{|0\rangle \otimes |0\rangle \otimes |0\rangle, |0\rangle \otimes |0\rangle \otimes |1\rangle, \dots, |1\rangle \otimes |1\rangle \otimes |1\rangle\}$. A product state of the form $|q_0\rangle \otimes |q_1\rangle \otimes |q_2\rangle$ can be abbreviated as $|q_0, q_1, q_2\rangle$, or $|000\rangle, |001\rangle, \dots, |111\rangle$ for the basis states. The circuit transforms the input state $|q_0, q_1, q_2\rangle$ into another valid state $|q'_0, q'_1, q'_2\rangle$.

The probability of obtaining the outcome $|0\rangle$ upon measuring the qubit $\alpha|0\rangle + \beta|1\rangle$ is $|\alpha|^2$, and the probability of obtaining the outcome $|1\rangle$ upon measuring the qubit is $|\beta|^2$. The sum of the probabilities must add up to 1. The act of measuring qubits is how we acquire information about a quantum state, and we can interpret that as classical information. That said, the original quantum data is irrecoverably lost as soon as we obtain the classical output.

As for the case of qubits, classical logic gates have a counterpart in quantum logic gates. Quantum gates can be applied to only one qubit, like the gates A and C in Figure 1, or to multiple qubits simultaneously, like the gates B , D , and E . When we apply a quantum gate to a set of qubits, the quantum state of these qubits is modified accordingly. We can associate a quantum gate with a unitary matrix that performs an identical transformation. That is, all quantum gates possess the unitarity constraint. A unitary matrix U holds the property that $U^\dagger U = I$ (or $UU^\dagger = I$), where its conjugate transpose $U^\dagger$ is also its inverse. Therefore, the application of a quantum gate to a quantum state followed by its conjugate transpose is equivalent to not operating on the qubits since all quantum operations are reversible. Another consequence of this constraint is that all quantum gates preserve the magnitude of the vector on which they operate. In this case, these vectors are unit vectors representing a qubit system or quantum state. Because quantum gates can be represented by unitary matrices, they are equivalent to unitary operators acting on finite dimensional spaces. This work, therefore, will refer to quantum gates and quantum operators interchangeably.

It is also possible, if the system architecture permits, to apply quantum gates in parallel. In other words, each quantum gate affects distinct sets of target qubits in a single timestep. The dashed vertical lines in the exemplified circuit in Figure 1 mark different timesteps, each identifying the current operation on the qubits. Since matrices can represent these transformations, we can see the three-qubit operator $A \otimes I_4$ as the first layer of the circuit, where I_4 is the 4×4 identity matrix. Likewise, $B \otimes I_2$ corresponds to the second layer, $C \otimes D$ to the third layer, and E to the fourth layer. The third layer depicts two operators applied simultaneously in one timestep. The number of steps or layers applied in sequence, assuming quantum gates acting on disjoint sets of qubits can be executed in parallel, is known as the depth of a quantum circuit. This measures the longest path of gates from the beginning up until the end of

a circuit. Another important piece of information that may also influence circuit depth is the gate count, which is the total number of gates in the circuit. These are frequently used cost functions of quantum circuits (SELINGER, 2013) due to how the increase of these quantities may impact performance by, for instance, extending a circuit's execution beyond qubit lifetime or resulting in error accumulation (SAKI; ALAM; GHOSH, 2021). Both can be used to assess the complexity of a quantum algorithm, and both depend on the circuit design techniques and the hardware architecture of quantum machines.

It is important to note that simply combining a set of quantum gates in a quantum circuit to perform a desired computation does not guarantee the efficiency of the construction. That is especially true for NISQ devices, given the obstacles associated with their physical implementation mentioned in Chapter 1. One can also tailor quantum circuits to take advantage of specific hardware architectures (HÄNER et al., 2018; STEIGER; HÄNER; TROYER, 2018; MURALI et al., 2019b; MURALI et al., 2019a; KILLORAN et al., 2019; FÖSEL et al., 2021; HIETALA et al., 2021). However, even in the case of hardware-agnostic schemes, we can adopt procedures to leverage core attributes of quantum mechanics, such as quantum parallelism, or avoid common redundancies like consecutive applications of inverse gates. This work, however, does not follow a hardware-specific approach other than being motivated by NISQ devices.

Now we briefly present some commonly used single-qubit and multiqubit gates in quantum circuits, each relevant to the proposed methods in this work.

2.1 SINGLE-QUBIT GATES

As the name implies, single-qubit gates operate on one target qubit. They are the simplest elements of quantum circuits but do not form a universal set of gates. Therefore, circuits constructed only with single-qubit gates cannot exploit quantum mechanical properties that are significantly advantageous compared to classical computation.

We can conceptualize a qubit geometrically as a unit vector in a three-dimensional sphere known as the Bloch sphere. The vector starts at the center, its endpoint touching the surface. Typically, the north pole of the Bloch sphere represents the basis state $|0\rangle$, whereas the south pole represents $|1\rangle$. Considering operations on a qubit, one can interpret them as rotations around this sphere, which, in turn, can be defined by unitary matrices.

The simplest example of a one-qubit operator is the identity gate

$$I = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}, \quad (2.1)$$

which does not perform any action on its target. The Pauli operators σ_x , σ_y , and σ_z correspond to rotations by π radians about the x , y , and z axes, respectively.

$$\sigma_x = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \quad (2.2)$$

$$\sigma_y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix} \quad (2.3)$$

$$\sigma_z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} \quad (2.4)$$

They are also known as the X , Y , and Z gates. We can see Pauli operators as flip operators on different bases. The σ_x gate, in particular, is the quantum logic gate counterpart to the classical NOT logic gate by functioning as a qubit flip between the computational basis states.

The Hadamard gate has the following matrix representation.

$$H = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix} \quad (2.5)$$

It acts on the basis states $|0\rangle$ and $|1\rangle$ by putting them in uniform superposition as $(|0\rangle + |1\rangle)/\sqrt{2}$ and $(|0\rangle - |1\rangle)/\sqrt{2}$, respectively, a valuable property in many algorithms. That means there is an equal probability of measuring $|0\rangle$ or $|1\rangle$ in both cases. The difference lies in the negative sign of the probability amplitude of $|1\rangle$ in the state $H|1\rangle$, however that is not reflected in the measurement probabilities and is, therefore, not information that can be retrieved after computation. The Hadamard gate is a versatile quantum gate that can be used to manipulate a quantum state by offering a way of using quantum mechanical properties like quantum entanglement, quantum parallelism, and quantum interference, which are concepts explained further.

Manipulating local phases in quantum states is another crucial mechanism of quantum computing. In fact, there exists a set of gates called phase gates which affect the amplitudes

of quantum states. A single-qubit local phase gate can be expressed as

$$P(\phi) = \begin{pmatrix} 1 & 0 \\ 0 & e^{i\phi} \end{pmatrix}, \quad (2.6)$$

which maps $|1\rangle$ to $e^{i\phi}|1\rangle$ and leaves $|0\rangle$ alone. The effect of a local phase gate is to change the relative phase between basis states, $|0\rangle$ and $|1\rangle$ in this case, by altering the angle of the probability amplitude associated with a basis state without changing its magnitude. This angle is referred to as a local phase. Some phase gates with a particular value for the angle ϕ are recurring in algorithms. Two examples are the gates

$$S = \begin{pmatrix} 1 & 0 \\ 0 & i \end{pmatrix} \quad (2.7)$$

and

$$T = \begin{pmatrix} 1 & 0 \\ 0 & e^{i\pi/4} \end{pmatrix}, \quad (2.8)$$

also representable by $P(\pi/2)$ and $P(\pi/4)$. Lastly, the global phase gate $Ph(\phi)$ equally shifts the phases of both $|0\rangle$ and $|1\rangle$. Because of this, the action of this gate is to introduce a global phase to the quantum state of a qubit, which is identical to a scalar multiplication of the qubit. As such, $e^{i\phi}I$ can also denote a global phase gate. It is worth noting that, as qubits on the Bloch sphere are seen as representing unit vectors, operating on them with unitary operators does not change their norm from 1.

One more set of single-qubit operators that can be defined are the special unitary gates on one qubit, i.e., $SU(2)$ gates, whose controlled version is the main object of interest in this work. As the name suggests, these are gates sharing the property of having a 2×2 matrix representation with determinant 1, forming the special unitary subgroup of degree 2. They have the general matrix form

$$W = \begin{pmatrix} \alpha & -\beta^* \\ \beta & \alpha^* \end{pmatrix}, \quad (2.9)$$

where α and β are complex numbers and $|\alpha|^2 + |\beta|^2 = 1$. Three recurring examples of $SU(2)$ gates are the rotation operators, which have the following matricial definitions.

$$R_x(\theta) = \begin{pmatrix} \cos \theta/2 & -i \sin \theta/2 \\ -i \sin \theta/2 & \cos \theta/2 \end{pmatrix} \quad (2.10)$$

$$R_y(\theta) = \begin{pmatrix} \cos \theta/2 & -\sin \theta/2 \\ \sin \theta/2 & \cos \theta/2 \end{pmatrix} \quad (2.11)$$

$$R_z(\theta) = \begin{pmatrix} e^{-i\theta/2} & 0 \\ 0 & e^{i\theta/2} \end{pmatrix} \quad (2.12)$$

The operators R_x , R_y , and R_z describe general rotations by an angle θ about the corresponding axes. Conveniently, the conjugate transpose of $R_n(\theta)$ is denoted by $R_n^\dagger(\theta) = R_n(-\theta)$, where n is the axis of rotation, since rotating a vector by the same angle in the opposite direction would bring it back to its former position. To write any single-qubit operator, one only needs R_y and R_z gates and a global phase employing the ZYZ decomposition (BARENCO et al., 1995; NIELSEN; CHUANG, 2010) below:

$$U = e^{i\alpha} R_z(\beta) R_y(\gamma) R_z(\delta). \quad (2.13)$$

Excluding the global phase $e^{i\alpha}$, the rotation operators are sufficient to describe an $SU(2)$ operator. Additionally, if A , B and C are $SU(2)$ operators such that $ABC = I$, then

$$A\sigma_x B\sigma_x C = W \in SU(2), \quad (2.14)$$

allowing a general way to represent any special unitary single-qubit operator (BARENCO et al., 1995). This is a useful decomposition for $SU(2)$ operators that appears in some of the decomposition methods mentioned in Chapter 3. Table 1 lists the aforementioned operators and their corresponding matrix representations.

Table 1 – A non-exhaustive list of single-qubit and multiqubit gates and their respective representations as a matrix.

Gate name	Matrix representation
Identity	$I = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$
Hadamard	$H = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}$
General local phase	$P(\phi) = \begin{pmatrix} 1 & 0 \\ 0 & e^{i\phi} \end{pmatrix}$
Phase gate S ($\phi = \pi/2$)	$S = \begin{pmatrix} 1 & 0 \\ 0 & i \end{pmatrix}$
Phase gate T ($\phi = \pi/4$)	$T = \begin{pmatrix} 1 & 0 \\ 0 & e^{i\pi/4} \end{pmatrix}$
Pauli σ_x , X , or NOT	$\sigma_x = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$
Pauli σ_y or Y	$\sigma_y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}$
Pauli σ_z or Z	$\sigma_z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}$
X -rotation	$R_x(\theta) = \begin{pmatrix} \cos \theta/2 & -i \sin \theta/2 \\ -i \sin \theta/2 & \cos \theta/2 \end{pmatrix}$
Y -rotation	$R_y(\theta) = \begin{pmatrix} \cos \theta/2 & -\sin \theta/2 \\ \sin \theta/2 & \cos \theta/2 \end{pmatrix}$
Z -rotation	$R_z(\theta) = \begin{pmatrix} e^{-i\theta/2} & 0 \\ 0 & e^{i\theta/2} \end{pmatrix}$
$U \in U(2)$	$U = e^{i\alpha} R_z(\beta) R_y(\gamma) R_z(\delta)$
$W \in SU(2)$	$W = A\sigma_x B\sigma_x C = \begin{pmatrix} \alpha & -\beta^* \\ \beta & \alpha^* \end{pmatrix}, A, B, C \in SU(2) \text{ and } ABC = I$
CNOT	$\text{CNOT} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix} = \text{diag}(I, \sigma_x)$
Toffoli	$\text{CCNOT} = \text{diag}(I, I, I, \sigma_x) = \text{diag}(I_4, \text{CNOT})$

Source: The author (2026)

2.2 MULTIQUBIT GATES

Multiqubit gates target two or more qubits simultaneously, enabling a quantum algorithm to tap into quantum mechanics phenomena that set quantum computing apart from classical computing. Quantum superposition was briefly mentioned in Section 2.1. This is a property of quantum mechanics where a qubit can be one of a continuous range of possible values until measured. Measurement of the qubit results in a single outcome determined probabilistically, unlike how classical bits are deterministically read out as either 0 or 1. Classical computing


so far has not been able to perfectly replicate superposition without incurring an exponential cost in physical resources, but this property proves to be a beneficial resource in quantum computing. It permits simultaneous calculations to happen, which can in turn allow faster solutions of certain problems than a classical program could obtain. Aside from superposition, a phenomenon that distinguishes quantum computing from its classical counterpart is quantum entanglement, which arises when two or more interacting qubits become correlated, even after being separated. Entangled states cannot be expressed as a product state, but can still be expressed as a system in a quantum superposition.

Together with single-qubit gates, gates targeting multiple qubits can be used to express arbitrary circuits that capitalize on the potential of quantum mechanics. One can express a multiqubit gate as a $2^n \times 2^n$ matrix, where n is the number of affected qubits. The general matrix representation of a circuit can be obtained via matrix multiplication of successive gate applications in a qubit, combined through tensor product of the resulting matrices for each qubit. The simplest case, represented by the generalized I_{2^n} matrix, mirrors the n -qubit identity gate. Matrices of multiqubit operators also possess the unitarity property. This way, any unitary matrix describes a viable quantum operation. However, simulating these operations with a reasonable gate set or physically realizing them in a quantum processor may be challenging. In general, one needs to break down a multiqubit operator into basic operations executable on hardware, employing a set of quantum gates recognized by the quantum processor that will run it. For demonstrational purposes, one can design quantum circuits using any universal set of gates if hardware restrictions are ignored.

An important feature of some quantum gates is that they can be controlled by qubits. This characteristic increases the capability to convey conditional operations in quantum circuits. One or more controlling qubits can enable a quantum operator on one or more target qubits depending on the current state or value the controlling qubits possess. In theory, there exists a controlled version of any quantum gate. That said, one can only expect a restricted number of controlled gates to be part of gate sets in quantum devices. To access the remaining controlled gates, one must decompose them into available gates for execution. The CNOT or controlled σ_x gate is one of the most important examples of this type of gate that are common in real gate sets, and it affects the target qubit by flipping it when the control qubit is in state $|1\rangle$.

The CNOT gate is especially applicable when decomposing other, more complicated operations. Together with single-qubit gates, they form a universal gate set capable of expressing any quantum operator (BARENCO et al., 1995). In reality, they are not perfectly implemented on

Another recurring quantum operator, the Toffoli operator, is obtained by adding another control qubit to the CNOT gate. In fact, we can generalize the controlled σ_x for any number of control qubits. This generalization frequently features in the decomposition approaches mentioned throughout this work, including the proposed schemes. Like any multiqubit gate, the Toffoli gate can be decomposed as a sequence of CNOTs and single-qubit gates, requiring at minimum six CNOTs (MASLOV, 2016). The circuit in Figure 2 is an example seen in (ITEN et al., 2016), that presents one possible way to decompose the Toffoli gate using CNOTs and the single-qubit gates Hadamard, S and T from Equations 2.5, 2.7 and 2.8, with as few CNOTs as possible. The circuit contains a Toffoli operator on the left-hand side, shown as an σ_x gate depicted with the $\oplus$ symbol on the target qubit, controlled by the top two qubits, which are the two $\bullet$ symbols vertically connected to the $\oplus$ symbol. On the right-hand side, the decomposition scheme for the Toffoli gate is provided using CNOTs, shown as a control qubit ($\bullet$) connected vertically to a target qubit ($\oplus$), as well as the single-qubit gates H , T and its conjugate transpose $T^\dagger$, and S .



In some quantum circuits, a worthwhile trick is to use approximate versions of this gate (MASLOV, 2016). These are Toffoli gates up to some diagonal operator, the latter introducing relative phases to the different basis states. This can be a crucial optimization strategy in circuits designed such that the positioning of Toffoli gates permits gate canceling. Figure 3 illustrates a decomposition circuit for a relative phase Toffoli gate. The left-hand side again contains a Toffoli gate, this time followed by a diagonal gate Δ spanning all three qubits, responsible for including the relative phases. The diagonal gate can also be shown to the left of the Toffoli

gate since they commute. On the right-hand side, the synthesis of the relative phase Toffoli circuit needs three CNOT gates and four single-qubit gates, as long as the first two gates are the same and inverse to the last two. The single-qubit gates in this example are the R_y operator and its inverse.




Figure 3 – Decomposition scheme for the relative phase Toffoli operator. Source: Adapted from Iten et al. (2016).

In the decomposition approaches mentioned in this work, we assume that the angle of rotation used in relative phase Toffoli gates is $\pi/4$, which introduces an amplitude of -1 to the state $|010\rangle$. The usefulness of this gate, when gate canceling is made possible, is owed to it requiring fewer CNOTs than the exact Toffoli gate. Cancellation of gates depends on pairing adjacent relative phase Toffoli gates acting on the identical qubits. Since the circuit is its own inverse, this makes at least partial canceling of gates possible if there is no operation between them, either due to no operators present or other cancelable actions. After cancellations, the amplitudes that would differ between the exact Toffoli gate and its relative phase version are no longer present. As mentioned, this is a useful optimization employed in decomposition circuits used in this work.

Appending auxiliary qubits as working qubits to the qubit set intended to be used for the algorithms is another possible way to tackle computationally heavy tasks. Decompositions of gates that use auxiliary qubits can reduce the overall gate count of the circuit (BARENCO et al., 1995; ITEN et al., 2016; HE et al., 2017; KIM; CHOI, 2018; BAKER et al., 2019; LIU; WEI; KWEK, 2020; MALVETTI; ITEN; COLBECK, 2021; TOMESH; ALLEN; SALEEM, 2022). This exploitable resource allows the design of quantum circuits that compromise on qubit overhead to save on the number of operations, and it is a requirement if the gate implemented comes from an irreversible function (MILLER; WILLE; DRECHSLER, 2010), which is a situation we do not consider in this work. It is possible to select a subset of the available qubits to serve exclusively as auxiliary qubits. Some decompositions can also borrow qubits already in use, aptly called dirty auxiliary qubits, and return them to their original state at the end of the operation. Usually, the latter is more expensive, but as additional qubits are not always viable, especially when qubits are still scarce in current devices, this type of workaround may be more viable. For this

reason, in this research, the proposed schemes do not assume the availability of extra qubits to reserve for auxiliary space. Additionally, clean auxiliary qubits, which are those with a known initialization state, may offer smaller circuits than those using dirty auxiliary qubits. However, during computations we cannot guarantee the state of a qubit to be “clean”. Therefore we study decompositions that do not assume specific values of auxiliary qubits and can be used at any point in a larger computation.

2.3 IMPORTANT PROPERTIES IN QUANTUM CIRCUITS

A crucial characteristic that can be leveraged in the design of quantum algorithms is quantum parallelism. This is different to the parallelism alluded to previously when explaining the application of gates in parallel in Figure 1. In simple terms, quantum parallelism refers to a function evaluation for many different inputs simultaneously (NIELSEN; CHUANG, 2010), coming from the linearity of operators applied to a superposition of basis states. Deutsch’s algorithm (DEUTSCH, 1985) is an example of application of quantum parallelism which demonstrates an advantage that quantum computing offers over classical computing. In this algorithm, Deutsch uses this quantum property to obtain a binary outcome upon measuring the output of the quantum circuit, revealing something about the nature of a binary function.

To start, one can assume there is a function $f : \{0, 1\} \rightarrow \{0, 1\}$ and the goal is to determine whether f is constant, outputting either 0 or 1 for any input; or balanced, outputting different values for each input. To compute the function f , assume there is an operator U_f given by the transformation $|x, y\rangle \mapsto |x, y \oplus f(x)\rangle$. This type of quantum operator is called an oracle or a black box, which is a quantum circuit whose inner workings are not relevant to the description of the task (NIELSEN; CHUANG, 2010). In this case, the quantum oracle can compute $f(x)$ and output an answer for the original question. Since U_f , from its definition, is unitary, this operator is valid. As an oracle, it performs the evaluations for f and marks a qubit with the information it found. A simplified version of the algorithm presented by (NIELSEN; CHUANG, 2010) is summarized by the circuit shown in Figure 4. In this case, the answer will be found on the first qubit.

In this algorithm, this can be done by initializing two qubits in the state $|01\rangle$. To make it so the evaluations occur for all values of the binary function, a trick is to use the Hadamard operator to put in superposition all possible values of a qubit. The first step of the algorithm




Figure 4 – The quantum circuit for a simplified version of Deutsch's algorithm (DEUTSCH, 1985). Source: Adapted from Nielsen and Chuang (2010).

is to apply an Hadamard gate to each qubit, resulting in the state

$$|\psi_1\rangle = \left(\frac{|0\rangle + |1\rangle}{\sqrt{2}} \right) \left(\frac{|0\rangle - |1\rangle}{\sqrt{2}} \right). \quad (2.15)$$

The parentheses show the current state of each qubit in the circuit. The reason for an Hadamard gate on the second qubit is a key part of this algorithm, which will become explicit in the next steps. Below is the effect of applying the oracle gate U_f when the second qubit is in the state $(|0\rangle - |1\rangle)/\sqrt{2}$.

$$U_f |x\rangle \left(\frac{|0\rangle - |1\rangle}{\sqrt{2}} \right) = \begin{cases} |x\rangle \left(\frac{|0\rangle - |1\rangle}{\sqrt{2}} \right), & \text{if } f(x) = 0 \\ |x\rangle \left(\frac{|1\rangle - |0\rangle}{\sqrt{2}} \right), & \text{if } f(x) = 1 \end{cases} \quad (2.16)$$

This can be rewritten as $U_f |x\rangle (|0\rangle - |1\rangle)/\sqrt{2} = (-1)^{f(x)} |x\rangle (|0\rangle - |1\rangle)/\sqrt{2}$. Note that, if the second qubit had been initialized as $|0\rangle$, this distinction of outcomes would not exist. Since the first qubit is in the state $(|0\rangle + |1\rangle)/\sqrt{2}$, then the function will be evaluated for both 0 and 1. If f is constant, both terms $|0\rangle$ and $|1\rangle$ in the first qubit will retain the same probability amplitude with equal signs. However, if f is balanced, the signs will be opposite. Therefore, the following state $|\psi_2\rangle$ will depend on f being a constant or balanced function. The resulting state is then

$$|\psi_2\rangle = \begin{cases} \pm \left(\frac{|0\rangle + |1\rangle}{\sqrt{2}} \right) \left(\frac{|0\rangle - |1\rangle}{\sqrt{2}} \right), & \text{if } f(x) \text{ is constant} \\ \pm \left(\frac{|0\rangle - |1\rangle}{\sqrt{2}} \right) \left(\frac{|0\rangle - |1\rangle}{\sqrt{2}} \right), & \text{if } f(x) \text{ is balanced.} \end{cases} \quad (2.17)$$

When the final Hadamard gate operates on the first qubit, the conditional state is then

$$|\psi_3\rangle = \begin{cases} \pm |0\rangle \left(\frac{|0\rangle - |1\rangle}{\sqrt{2}} \right), & \text{if } f(x) \text{ is constant} \\ \pm |1\rangle \left(\frac{|0\rangle - |1\rangle}{\sqrt{2}} \right), & \text{if } f(x) \text{ is balanced,} \end{cases} \quad (2.18)$$

meaning that, at the end of the circuit, the first qubit will reveal that f is constant when measured as 0 or that f is balanced when measured as 1. Given that $f(0) \oplus f(1) = 0$ when $f(0) = f(1)$, and $f(0) \oplus f(1) = 1$ when $f(0) \neq f(1)$, then the final state $|\psi_3\rangle$ can also be

rewritten as

$$|\psi_3\rangle = \pm |f(0) \oplus f(1)\rangle \left(\frac{|0\rangle - |1\rangle}{\sqrt{2}} \right). \quad (2.19)$$

Deutsch's algorithm can be generalized for the case of an n -qubit x with a function $f : \{0, \dots, 2^n - 1\} \rightarrow \{0, 1\}$. This generalization is known as the Deutsch–Josza algorithm (DEUTSCH; JOZSA, 1992).

The way of applying the Hadamard gate in Deutsch's algorithm highlights a use for it that permits quantum mechanics phenomena, enabling a means to extract the solution to the problem. By initializing the state as $|01\rangle$ and performing the Hadamard gate applications the way they are shown in the circuit, important features of quantum mechanics are utilized. Quantum interference is used, together with quantum parallelism, to amplify or cancel out probability amplitudes, causing the measurement outcome of the first qubit to reveal this global property of the binary function. In this case, the negative local phase of the second qubit originated after the Hadamard gate created alternative cases that interfere with each other based on f being constant or balanced. This is something that cannot be done classically with the same efficiency, as the multiple function evaluations would be performed separately and the combined result would inform about this function property. The alternatives of having a constant or a balanced function are mutually exclusive. In the quantum setting, quantum superposition and local phases of quantum states can be combined to cause interference, making the resulting alternatives, such as this property of a function, measurable in the outcome. Beyond quantum superposition and quantum parallelism, which can in theory be modeled classically with enough resources, quantum interference is a key phenomenon of quantum mechanics that sets quantum computing apart from classical computing. In the following section, we present some of the obstacles that need to be overcome so that taking full advantage of these characteristics becomes possible in quantum computing.

2.4 CHALLENGES IN NEAR-TERM QUANTUM COMPUTING

Quantum noise presents an intrinsic barrier to the perfect execution of quantum computations. It can originate from multiple physical sources (RESCH; KARPUZCU, 2021), such as unwanted interactions that may happen between quantum systems and the environment (NIELSEN; CHUANG, 2010) or between qubits (WOOD; GAMBETTA, 2018). Other possible sources of noise can be the application of imprecise rotations of a quantum gate (BARNES et al., 2017) and, in

quantum systems that are not two-level systems, qubits leaking out from the suitable states that represent 0 or 1 to noncomputational states (FOWLER, 2013). Quantum decoherence is a critical and challenging phenomenon that presents an obstacle to reliable quantum computing. Yanofsky and Mannucci (2008, p. 308) define quantum decoherence as “the loss of purity of the state of a quantum system as the result of entanglement with the environment”. Since there is a desire for an isolated quantum system without the loss of the properties necessary to perform quantum computation, decoherence is something that must be avoided before a computational task is finished. These interactions referred to earlier between a quantum system and its environment result in disturbances that lead to the disruption of coherence and fidelity of quantum states. In other words, noise causes quantum systems to decohere and behave classically, while the loss of fidelity means that the quality of measurement outcomes is negatively impacted (NIELSEN; CHUANG, 2010).

This is an obstacle for long enough computations, making many possible applications of quantum computing impractical for near-term devices. Given the fragility of quantum phenomena like superposition and entanglement, and the inherent imperfections of quantum gate implementations resulting in error accumulation, quantum noise becomes a significant challenge in the realization of scalable quantum computing (PRESKILL, 2018; SAKI; ALAM; GHOSH, 2021). Consequently, mitigation strategies need to be employed in order to address the effects of quantum noise. The main strategies to handle noise currently require error correction or some other forms of mitigation. Ideally, fault-tolerant quantum computation is achieved by employing quantum error correction (QEC) techniques to recover from eventual computation errors. However, the process of encoding the quantum information needed to perform recovery operations, together with the process of recovering the information, implies an inevitable accumulation of errors (SHOR, 1996; PRESKILL, 1998) that must be managed reliably. In order for QEC to be reliable, the system gate error rates must be below a reasonable threshold. Additionally, due to the high qubit overhead required by QEC, fault tolerance in near-term quantum devices still proves to be infeasible.

At this stage of quantum computing hardware development, qubits remain inherently fragile, exhibiting limited lifetime due to decoherence and other sources of noise and error (LI; DING; XIE, 2019). Since gate errors can be avoided with a lower number of gates, and similarly shorter quantum circuits have a lower chance to reach decoherence time, reducing quantum circuit operations and depth can be a measure to mitigate both noise and error. As mentioned, QEC has an implementation cost that is difficult to be met in the current era. Consequently,

optimizations to quantum algorithms that lower overall circuit cost are easier to implement. The cost of a quantum circuit can be measured by counting the number of basic operations to simulate it, which can be translated into the total gate count (BARENCO et al., 1995), however the term may be applied more broadly. The quantum circuit cost can encompass the cost to run a quantum circuit in terms of the total number of gates required, the total number of qubits being operated on, or the minimal number of timesteps with gate operations, also known as circuit depth. With that said, each of these aspects can contribute to exacerbating the effects of noise:

- More gates mean a higher chance to execute faulty operations resulting in more error accumulation;
- More qubits are affected by the qubit connectivity of the quantum architecture, which can require more gates to be added for non-adjacent qubit interactions and result in more noise;
- Longer circuit depth impacts the duration of the quantum algorithm, which has an increased chance to hit decoherence time before the end of its runtime and continue execution despite the loss of quantum information.

Qubit connectivity, in particular, is affected by how qubits are physically realized on NISQ devices. Hardware constraints exist in physical qubit coupling, thereby hindering direct interaction between non-neighboring qubits without the incorporation of additional quantum gates (MASLOV; FALCONER; MOSCA, 2008; LI; DING; XIE, 2019). Due to the problem of mapping logical qubits to physical ones being NP-complete, this poses yet another obstacle to near-term quantum computing, particularly when it comes to the realization of multiqubit gates. In devising new algorithms for quantum computing, for simplicity, it is usual to assume full connectivity between qubits in the system. This is also the case in this work since, as was mentioned before, a hardware-agnostic architecture is assumed.

The Clifford group of quantum operators is a subset of quantum operators that can be efficiently simulated on a classical computer (GOTTESMAN, 1998; AARONSON; GOTTESMAN, 2004). The Hadamard, S and CNOT gates comprise this set, possessing quantum operators frequently used in many algorithms. However, this subset by itself is not sufficient for universal quantum computing (GOTTESMAN, 1998; SHOR, 1996). By also including the Toffoli operator, one can obtain a universal set of operators for quantum computation (SHOR, 1996).

Alternatively, as evidenced by the exact decomposition for the Toffoli operator in the quantum circuit illustrated in Figure 2, the Clifford group with the T operator forms a universal set, also known as Clifford+ T . This operator set underscores how cost can be estimated differently depending on the task. Efficient, fault tolerant implementation of the Clifford group has been demonstrated (SHOR, 1996), however the T operator is difficult to implement with efficiency without converting correctable errors into uncorrectable ones, which is essential for fault tolerant computing (GOSSET et al., 2013; CAMPBELL; TERHAL; VUILLOT, 2017; GIDNEY, 2018). Therefore, in QEC, more weight is given to the T -count and T -depth, which refer to the number of T operators and the number of timesteps where T operators occur. Minimizing both thus offers an avenue for reducing overhead in fault tolerant computing.

In general quantum circuit design where fault tolerance is not the focus, the CNOT operator is typically the emphasis when measuring cost. In NISQ hardware, two-qubit gate implementations have a significant error rate which cannot yet be overcome by QEC (PRESKILL, 2018). While this is true, the CNOT operator still proves to be more robust to noise than other two-qubit operators (HARROW; NIELSEN, 2003). As a useful quantum operator for permitting quantum entanglement and universal quantum computation with other single-qubit operators, despite being noisy, it is ever-present in many quantum algorithms. Owing to that, minimizing the CNOT count in quantum circuits is a crucial noise mitigation strategy for NISQ devices until fault tolerant quantum computing can effectively be achieved.

3 REVIEW OF LITERATURE

Here we explain the most relevant decomposition methods related to this work, particularly for multicontrolled single-qubit gates, including steps necessary for the proposed constructions. Although there are more general ways to decompose single-qubit gates with multiple controls, including universal methods that apply to any multiqubit operator, we introduce only specific ones since they produce more efficient circuits for the types of single-qubit gates mentioned in the following sections. For simplicity, from now on a $C_{k,n}(U)$ gate will refer to an n -qubit gate that performs the U operator with k control qubits. Table 2 summarizes the complexities in terms of CNOT counts for all decomposition schemes mentioned in this chapter, along with the minimum number of auxiliary qubits required by each method.

Table 2 – CNOT gate count and required number of auxiliary qubits of the different decomposition schemes for multicontrolled single-qubit gates in the related literature.

Gate	Reference	# of aux. qubits	# of CNOTs
$C_{k,n}(\sigma_x)$	Barenco et al. (1995)	$k - 2$	$24k - 48$
$C_{k,n}(\sigma_x)$	Iten et al. (2016)	$k - 2$	$8k - 6$
$C_{n-2,n}(\sigma_x)$	Barenco et al. (1995)	1	$24n - 84$
$C_{n-2,n}(\sigma_x)$	Iten et al. (2016)	1	$16n - 40$
$C_{n-2,n}(\sigma_x)$	He et al. (2017)	1	$24n - 96$
$C_{n-1,n}(U), U \in U(2)$	Barenco et al. (1995)	0	$16n^2 - 60n + 42$
$C_{n-1,n}(U), U \in U(2)$	Silva and Park (2022)	0	$4n^2 - 12n + 10$
$C_{n-1,n}(W), W \in SU(2)$	Barenco et al. (1995)	0	$48n - 162$
$C_{n-1,n}(W), W \in SU(2)$	Iten et al. (2016)	0	$28n - 88$ (n even) or $28n - 92$ (n odd)

Source: The author (2026).

3.1 DECOMPOSITIONS FOR MULTICONTROLLED σ_x GATES

Decomposing multicontrolled σ_x gates has several applications given this operator is widely used in many quantum algorithms. As implied before, the more qubits that are free to be used as working qubits for a multiqubit gate decomposition, the fewer gates a quantum circuit may need to implement the decomposition. The following decomposition schemes for multicontrolled σ_x have varying numbers of auxiliary qubits, so one can utilize them according to resource availability. In all cases, the auxiliary qubits do not have to be “clean” since they are reverted to their original states before the circuit operation completes. Table 2 contains

information about the number of auxiliary qubits and CNOTs in the circuits produced by the methods described below.

3.1.1 Barenco et al. (1995)

The construction of controlled σ_x gates with n qubits and k control qubits, shown by Barenco et al. (1995), needs at least $k - 2$ auxiliary qubits to operate with a linear number of CNOT gates. The circuit consists of a roughly V-shaped sequence of Toffoli gates. The need for $k - 2$ auxiliary qubits for each control qubit is due to the pairing of a control qubit with one auxiliary qubit in the chain of Toffoli gates. Two of the Toffoli gates in the sequence act on an auxiliary qubit using two control qubits of the original operator being decomposed, meaning k control qubits except two are matched with an auxiliary qubit.

The goal is to flip the target qubit employing the auxiliary qubits when all controls are $|1\rangle$. The first Toffoli flips the target if the corresponding auxiliary qubit is in state $|1\rangle$. Otherwise, the control qubits set that qubit to $|1\rangle$ so that the second Toffoli gate inverts the target. In Figure 5, the dashed vertical line identified by “flip target” marks the beginning of this process. After that, the remaining gates reset the auxiliary qubits to their initial values from the beginning of the circuit. The second dashed line identifies this step in Figure 5 as “revert auxiliary.” For each auxiliary qubit except one, four Toffoli gates are applied. For the remaining auxiliary qubits, two Toffoli gates are applied and two more Toffoli gates operate on the target qubit. This results in $4(k - 2)$ Toffoli gates needed in the decomposition, as shown in Lemma 7.2 of Barenco et al. (1995), for a total of $24k - 48$ CNOTs.




Figure 5 – Barenco et al. (1995) circuit for k -controlled σ_x gates with $k - 2$ auxiliary qubits. Source: Adapted from Barenco et al. (1995) and (VALE et al., 2024) © 2024 IEEE.

We can use this gate to construct a circuit for multicontrolled σ_x gates with only one

qubit free to be used as an auxiliary. The control qubits split into two groups comprising, for example, the top $k_1 = \lceil n/2 \rceil$ controls and the remaining $k_2 = n - k_1 - 1$ controls. The circuit from Figure 5 is applied four times, alternating between the k_1 controls acting on the auxiliary qubit and the k_2 control qubits (plus the auxiliary qubit as an extra control) on the actual target of the circuit. Figure 6 illustrates the resulting construction. Each gate borrows qubits from the other group of controls as dirty auxiliary qubits.




Figure 6 – Barenco et al. (1995) circuit for multicontrolled σ_x gates with one auxiliary qubit. Source: Adapted from Barenco et al. (1995) and Vale et al. (2024) © 2024 IEEE.

Direct implementation of the $C_{n-2,n}(\sigma_x)$ circuit shown in Figure 6 results in $8(n - 5)$ Toffoli gates, since the top gates require $4(k_1 - 2)$ Toffoli gates and the bottom ones require $4(k_2 - 2)$ Toffoli gates. The number of Toffoli gates before optimizations is therefore

$$\begin{aligned}
 2 \times 4(k_1 - 2) + 2 \times 4(k_2 - 2) &= 8(k_1 + k_2) - 32 \\
 &= 8(k_1 - n - k_1 - 1) - 32 \quad (3.1) \\
 &= 8n - 40,
 \end{aligned}$$

totaling $48n - 240$ CNOT gates. However, the $C_{k_1,n}(\sigma_x)$ gates targeting the auxiliary qubit are paired to reset each other. Therefore relative phase Toffoli gates, seen in Figure 3, can be the only version of the Toffoli gate used in their execution. In other words, four Toffoli gates each contribute with six CNOTs, and the remaining $8n - 36$ contribute with $24n - 108$ CNOTs, adding up to $24n - 84$ CNOT gates.

3.1.2 Iten et al. (2016)

Iten et al. (2016) demonstrated how to improve upon both schemes for multicontrolled σ_x gates with auxiliary qubits introduced by Barenco et al. (1995) by relying on relative phase Toffoli gates and gate canceling. The action of the circuit in Figure 6 would not change if all gates affecting the auxiliary qubits of the $C_{k_2,n}(\sigma_x)$ circuit were relative phase Toffoli gates. Figure 7 illustrates the modifications made to the construction of Figure 5 that reflect this change for the $C_{k_2,n}(\sigma_x)$ gates.




Figure 7 – The “flip target” part of the circuit from Figure 5 following gate canceling optimization. Source: Adapted from Iten et al. (2016).

These modifications reduce the number of CNOTs each gate contributes from six to three. Considering that the circuit pairs up all gates of the same type acting on the same target qubit, the phases they introduce cancel out due to partial canceling of the relative phase Toffoli circuit. Thus, except for three of the control qubits, eight CNOTs per qubit are produced. Including the two top relative phase Toffoli gates, where no canceling happens, and the exact Toffoli gates affecting the actual target qubit of the operation, this circuit contains $8(k-3) + 18 = 8k - 6$ CNOTs in total. These optimizations can then improve the multicontrolled σ_x with one auxiliary qubit. Using this modification to the circuit in Figure 6 and only relative phase Toffoli gates on the gates controlled by the k_1 top qubits, the latter already present in the construction by Barenco et al. (1995), one can implement this gate using $16n - 40$ CNOTs.

3.1.3 He et al. (2017)

He et al. (2017) proposed a circuit to decompose $C'_{k,n}(\sigma_x)$ gates with two or more control qubits and one auxiliary qubit that takes advantage of operations performed on the auxiliary qubit. In short, these operations utilize paired inverse gates to cancel unneeded actions to introduce the phase of -1 only when all controls and the target qubit are $|1\rangle$. When that happens, we can see it as propagating the phase to the target to simulate a controlled σ_z gate. This effect then becomes a qubit flip on the target. When the control state is 1, the action on the target qubit is equivalent to $H\sigma_z H = \sigma_x$. Otherwise, no action happens since the Hadamard operator is self-adjoint. Figure 8 depicts the complete circuit. The qubits $|t\rangle$ and $|a\rangle$ are the target and auxiliary qubits, respectively. The remaining top wires represent two groups of k_1 and k_2 qubits which are the control qubits of the gate.




Figure 8 – He et al. (2017) decomposition scheme for multicontrolled σ_x gates with one auxiliary qubit. Source: Adapted from He et al. (2017) and Vale et al. (2024) © 2024 IEEE.

Upon closer inspection of this circuit on different control states, four main situations decide which operators activate on the auxiliary qubit. The first is when both groups of qubits, the k_1 control qubits and the k_2 controls plus the target, have at least one qubit in the state $|0\rangle$, disabling all σ_x gates and yielding

$$S^\dagger S S^\dagger S = I. \quad (3.2)$$

When all of the qubits from the group of k_1 controls are $|1\rangle$, and at least one of the qubits from the group of k_2 controls or the target qubit is $|0\rangle$, then the sequence of operators

$$S^\dagger S \sigma_x S^\dagger S \sigma_x = I \quad (3.3)$$

applies to the auxiliary qubit. Conversely, when all k_2 controls and the target are $|1\rangle$, and at least one of the k_1 controls is $|0\rangle$, the resulting transformation is

$$S^\dagger \sigma_x S S^\dagger \sigma_x S = I. \quad (3.4)$$

Finally, when all control qubits and the target qubit are in state $|1\rangle$, the action on the auxiliary qubit is

$$S^\dagger \sigma_x S \sigma_x S^\dagger \sigma_x S \sigma_x = -I, \quad (3.5)$$

producing the effect of the controlled σ_z . With the Hadamard operators, the simulation of the controlled σ_x on the target qubit is satisfied.

Although He et al. (2017) present this decomposition scheme as consisting of only relative phase Toffoli gates (HE et al., 2017, Figures 2 and 3), this incorrectly produces a relative phase circuit. Exact Toffoli gates must target the auxiliary qubit to correct the relative phase. Nevertheless, the original estimate of the number of CNOTs in the decomposed quantum circuit is correct, totaling $24n - 96$ CNOTs, where the number of qubits n includes the auxiliary qubit.

3.2 DECOMPOSITIONS FOR MULTICONTROLLED SINGLE-QUBIT GATES

It is valuable to have ways to synthesize circuits for any gates in the $U(2)$ group and not just σ_x gates when multiple controls are involved. The decompositions shown in this section offer an improvement over general n -qubit gate decomposition methods, which have a CNOT gate count in the order of $O(4^n)$ (SHENDE; BULLOCK; MARKOV, 2005). Since these methods work for any controlled single-qubit gate, they are also appropriate for $SU(2)$ operators, which this work focuses on, although there is a penalty on the number of gates and depth of circuits compared to those produced by methods explained in Section 3.3.

3.2.1 Barenco et al. (1995)

An intuitive approach to decompose multicontrolled single-qubit operators with no auxiliary qubits is by employing roots of gates. In Figure 9, the quantum circuit illustrates a method presented by Barenco et al. (1995) to decompose $C_{n-1,n}(U)$ gates that requires recursively computing roots of U up to the 2^{n-2} th root. The disadvantage of exploring this property is the ever-decreasing power of U incurring numerical errors in computations and proving impractical for large numbers of qubits. The regular presence of free qubits to borrow as dirty auxiliary qubits in multicontrolled σ_x applications makes using the scheme from Figure 6 or its improved variant by Iten et al. (2016) possible.




Figure 9 – Barenco et al. (1995) decomposition scheme for general multicontrolled single-qubit gates. Source: Adapted from Iten et al. (2016).

The costs of the two controlled gates produced at each step of the recursion, plus the additional gate at the last computation step of the $\sqrt[n-2]{U}$ gate, do not depend on the number of qubits. Hence they only increase the CNOT count by a constant. Moreover, the multicontrolled σ_x gates come in pairs, starting with $n - 2$ control qubits but removing one control at each step. This results in a CNOT count in the order of $O(24n^2)$ produced by these gates. Including optimizations by Iten et al. (2016), the estimate is that this construction needs at most $16n^2 - 60n + 42$ CNOTs.

3.2.2 Silva and Park (2022)

Silva and Park (2022) introduced an improved quantum circuit to decompose any multi-controlled $U(2)$ gate based on an approach by Saeedi and Pedram (2013) specifically for R_x gates. The authors showed that any $C_{n,n+1}(U)$ gate can be decomposed as

$$C_{n,n+1}(U) = Q_n^\dagger P_n(U)^\dagger Q_n(q_1 \sqrt[2^{n-1}]{U} q_{n+1}) P_n(U), \quad (3.6)$$

where

$$P_n(U) = \prod_{k=2}^n q_k \sqrt[2^{n-k+1}]{U} q_{n+1} \quad (3.7)$$

and

$$Q_n = \prod_{k=1}^{n-1} C_{k,k+1} R_x(\pi). \quad (3.8)$$

The quantum circuit has four stages, as shown in Figure 10. The first stage applies the operators $\sqrt{U}, \dots, \sqrt[2^{n-1}]{U}$ plus an extra $\sqrt[2^{n-1}]{U}$ on the target qubit q_{n+1} , each controlled by one qubit from $q_n, \dots, q_1$, respectively. This sequence corresponds to the action of the operator U . Next, qubits $q_n, \dots, q_2$ are the targets of $R_x(\pi)$, each application controlled by all the qubits above. The following two stages are inverse to the first two, except the $\sqrt[2^{n-1}]{U}^\dagger$ gate controlled by qubit q_1 does not occur. All affected control qubits are then reset to their original values at the end, equating to the inverse of the second stage.

When all controls are in state $|1\rangle$, all operations of the first and second stages are activated, flipping each qubit $q_2, \dots, q_n$ to $-i|0\rangle$. Consequently, this skips the third stage, and the action of U transforms the target qubit. When q_1 is $|0\rangle$, any gates controlled by this qubit, including all operators in the second stage, are skipped, and the remaining gates cancel out. The target qubit remains unchanged. If q_1 is $|1\rangle$, and at least one control qubit is $|0\rangle$, the controls in state $|0\rangle$ disable the corresponding gates at the first stage. During the following sequence, any qubits below them remain unflipped. The first control qubit q_j in state $|0\rangle$, however, is flipped. That deactivates the inverse gates controlled by any qubits before q_j , activates the gate controlled by q_j , and causes the remaining gates with unaffected control qubits to cancel the corresponding gates from the first stage. The result is no action on the target qubit.

Saeedi and Pedram (2013) had proposed an alternative construction for the $C_{n-1,n}(R_x(\pi))$ circuit, depicted in Figure 11. When this construction and its inverse are incorporated into the quantum circuit from Figure 10, replacing the second and fourth stages, respectively, the number of two-qubit gates required in the decomposition scheme is $2n^2 - 6n + 5$. With each




Figure 10 – Silva and Park (2022) decomposition scheme for general multicontrolled single-qubit gates. Source: Adapted from Silva and Park (2022).

two-qubit gate requiring at most two CNOTs, this translates to up to $4n^2 - 12n + 10$ CNOT gates for the decomposition scheme proposed by Saeedi and Pedram (2013) and similarly for the proposal by Silva and Park (2022).

This gate decomposition improves upon the scheme introduced by Barenco et al. (1995), yet the quadratic cost in the number of gates remains. As we will see in Section 3.3, there is a subset of $U(2)$ operators, namely the special unitary or $SU(2)$ operators, whose controlled versions have linear decompositions. As a result, if we know beforehand that the class of quantum gates we are dealing with is a multicontrolled $SU(2)$ gate, we can pick a linear decomposition scheme among the available ones. In the proposed approach presented in Chapter 4, we improve upon the methods shown in the present chapter when dealing with multicontrolled $SU(2)$ operators.




Figure 11 – Saeedi and Pedram (2013) decomposition scheme for multicontrolled $R_x(\pi)$. Source: Adapted from Saeedi and Pedram (2013).

3.3 DECOMPOSITIONS FOR MULTICONTROLLED $SU(2)$ GATES

Now that we have described the decomposition methods for multicontrolled σ_x and general single-qubit gates, the focus can shift to decomposing $SU(2)$ gates with multiple control qubits. This set of quantum operators has fewer known decomposition methods designed for it than σ_x gates or general single-qubit gates. Yet, without relying on a generalized scheme, the following approaches exhibit linear complexity in basic quantum gates and circuit depth. The proposal in this work has a further advantage over the strategies presented in the literature. Examples of algorithms using $SU(2)$ gates are rotation-based state preparation using multicontrolled R_z or R_y rotations (MOTTONEN et al., 2004; BERGHOLM et al., 2005; SHENDE; BULLOCK; MARKOV, 2005), probability-loading state preparation (GROVER; RUDOLPH, 2002; MARIN-SANCHEZ; GONZALEZ-CONDE; SANZ, 2023), and circuit-based approaches for simulating lattice gauge Hamiltonians using multicontrolled R_x rotations (BALAJI et al., 2025).

3.3.1 Barenco et al. (1995)

To construct a quantum circuit that synthesizes this type of gate, Barenco et al. (1995) use the fact that Equation (2.14) can form $SU(2)$ operators. From this, one can note that adding one control qubit to an $SU(2)$ gate means adding one control to both σ_x gates. When the control state is $|0\rangle$, the target qubit is unaffected since $ABC = I$. The bottom wire simulates Equation (2.14) when the control state is $|1\rangle$. That is enough to realize a circuit for any number of control qubits. Figure 12 shows a circuit that uses two multicontrolled σ_x operators borrowing one of the controls as a dirty auxiliary qubit. Using single-control $SU(2)$ gates in this construction frees up one qubit that the multicontrolled σ_x gates can borrow. Then, the circuit can integrate one of the decomposition schemes from Section 3.1 here to lower its complexity.




Figure 12 – Barenco et al. (1995) decomposition scheme for multicontrolled $SU(2)$ gates. Source: Adapted from Barenco et al. (1995).

As initially presented, this scheme uses the decomposition methods described in Section 3.1.1. Consequently, the number of CNOTs required by two multicontrolled σ_x gates is $48n - 168$. As for the remaining three gates with one control qubit, each requires up to two CNOTs. This decomposition can be realized with at most $48n - 162$ CNOT gates.

3.3.2 Iten et al. (2016)

The decomposition scheme for $SU(2)$ gates by Iten et al. (2016) uses the same idea as the one by Barenco et al. (1995), though with additional modifications, described in Section 3.1.2, that makes the multicontrolled σ_x gates more economical. Until now, this means implementing this construction requires up to $32n - 74$ CNOTs. Still, there are further possible optimizations for this construction since the multicontrolled σ_x being inverse of itself allows both occurrences of this gate to use different schemes. Appropriately, the second occurrence of the gate is in reverse order. That way, any quantum gate that does not interact with the bottom two qubits can be canceled. The consequence for the quantum gate to the left of the dashed region

in Figure 13 is eliminating the last group of gates that reset the auxiliary qubits of the σ_x gate controlled by the bottom set of control qubits. The quantum gate to the right-hand side mirrors this effect.




Figure 13 – Region of the multicontrolled $SU(2)$ scheme by Iten et al. (2016) where the canceling of the resetting part of the $C_{k_2}(\sigma_x)$ gates happens. Source: Adapted from Iten et al. (2016).

Considering all these optimizations, the updated CNOT count can be calculated. Assuming there are $k_2 = \lceil n/2 \rceil$ controls in the bottom multicontrolled σ_x gates and $k_1 = n - k_2 - 1$ controls in the top multicontrolled σ_x gates, one middle $C_{k_2,n}(\sigma_x)$ gate produces $8k_2 - 6$ CNOTs minus the number of CNOTs in the auxiliary qubit resetting part of the circuit. That removes $4(k_2 - 3) + 3$ CNOTs, totaling $4k_2 + 3$ CNOTs in one of the two $C_{k_2,n}(\sigma_x)$ gates encircling the controlled B operator in the middle. Each top $C_{k_1,n}(\sigma_x)$ gate produces $8k_1 - 14$ CNOTs since they use only relative Toffoli gates. With the inclusion of the single-qubit gates, the expression to compute the improved CNOT count is then

$$4(8k_1 - 14) + 2(8k_2 - 6) + 2(k_2 + 3) + 6, \quad (3.9)$$

which amounts to $28n - 88$ when n is an even number of qubits or $28n - 92$ when n is odd.

4 PROPOSED DECOMPOSITION SCHEME FOR MULTICONTROLLED SPECIAL UNITARY SINGLE-QUBIT OPERATORS

This chapter is based substantially on the paper titled “Circuit Decomposition of Multicontrolled Special Unitary Single-Qubit Gates” by Vale et al. (2024), published in IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems (Volume: 43, Issue: 3, March 2024) © IEEE 2011, of which the thesis author is first author. Portions of the text, results, and figures are adapted from that article.

Before presenting this proposal, some definitions are necessary. We begin by returning to the decomposition scheme by He et al. (2017) explained in Section 3.1.3. With this scheme, a multicontrolled σ_x gate can be decomposed using a single auxiliary qubit. We first introduce an adaptation of this scheme for multicontrolled $SU(2)$ gates with at least one real-valued diagonal instead of σ_x gates. With our modification, the need for an auxiliary qubit is eliminated and all actions are performed on the target qubit.

4.1 DECOMPOSITION SCHEME FOR MULTICONTROLLED $SU(2)$ GATES WITH ONE REAL-VALUED DIAGONAL

We start by modifying the operator S from Equation (3.5) to a generic operator A . If A is not restricted to a specific gate, then the new scheme yields some operator U as follows.

$$A^\dagger \sigma_x A \sigma_x A^\dagger \sigma_x A \sigma_x = U \quad (4.1)$$

Note that, since each matrix is unitary, paired with their inverse, and the determinant of a product of matrices is the product of their determinants and each matrix is unitary, then the resulting matrix U possesses a determinant of 1. To simplify this approach, it is safe to assume that the operator A is an $SU(2)$ operator, which will not alter the property of Equation (4.1). Then, like Equation (2.9), A has the form

$$A = \begin{pmatrix} \alpha & -\beta^* \\ \beta & \alpha^* \end{pmatrix}, \quad (4.2)$$

Similarly as for the decomposition scheme by He et al. (2017), using the fact that the equations below

$$\begin{aligned} A^\dagger A A^\dagger A &= I \\ A^\dagger A \sigma_x A^\dagger A \sigma_x &= I \\ A^\dagger \sigma_x A A^\dagger \sigma_x A &= I \end{aligned} \quad (4.3)$$

mirror the behavior of Equations (3.2), (3.4), and (3.5) by not changing the qubit operated on, a similar circuit to the one shown in Figure 8 can be used to depict this adapted algorithm. This new quantum circuit is demonstrated in Figure 14. The only exception is that, unlike in the original circuit, no auxiliary qubit is shown because Hadamard operators are not being used to manipulate the outcome on the target qubit.




Figure 14 – Initial decomposition scheme, adapted from He et al. (2017), proposed for multicontrolled $SU(2)$ gates whose corresponding matrices contain a real-valued off diagonal. Source: Adapted from Vale et al. (2024) © 2024 IEEE.

Like before, the qubit register is depicted as being divided into two halves with $k_1 = \lceil k/2 \rceil$ qubits and $k_2 = \lfloor k/2 \rfloor$ qubits. No dedicated qubit is used as an auxiliary qubit, however this algorithm still makes use of dirty auxiliary qubits in its subroutines. Recalling from Section 3.1.1 that we can decompose k -controlled σ_x gates with no more than $k - 2$ dirty auxiliary qubits, each register uses qubits available from the other register as working qubits to perform the k_1 -controlled and k_2 -controlled σ_x gates. We use the improved version from Section 3.1.2 to reduce the gate count. Now we can turn our attention to the result of the sequence of operations in Equation (4.1) to understand what type of operator is yielded by it.

Lemma 1. *An operator given by Equation (4.1), where A is of the form of Equation (4.2), generates an $SU(2)$ operator with the restriction that its off-diagonal is real-valued.*

Proof. By calculating $A^\dagger \sigma_x A \sigma_x$, we obtain

$$A^\dagger \sigma_x A \sigma_x = \begin{pmatrix} (\alpha^*)^2 - (\beta^*)^2 & \alpha\beta^* + \alpha^*\beta \\ -(\alpha\beta^* + \alpha^*\beta) & \alpha^2 - \beta^2 \end{pmatrix}. \quad (4.4)$$

By inspecting the values $\alpha^*\beta$ and $\alpha\beta^*$, and denoting as $\Re()$ and $\Im()$ the real and imaginary parts of a complex number, respectively, we have

$$\begin{aligned}\alpha^*\beta &= [\Re(\alpha) - i\Im(\alpha)][\Re(\beta) + i\Im(\beta)] \\ &= \Re(\alpha)\Re(\beta) + \Im(\alpha)\Im(\beta) + i[\Re(\alpha)\Im(\beta) - \Im(\alpha)\Re(\beta)]\end{aligned}\quad (4.5)$$

and

$$\begin{aligned}\alpha\beta^* &= [\Re(\alpha) + i\Im(\alpha)][\Re(\beta) - i\Im(\beta)] \\ &= \Re(\alpha)\Re(\beta) + \Im(\alpha)\Im(\beta) - i[\Re(\alpha)\Im(\beta) - \Im(\alpha)\Re(\beta)].\end{aligned}\quad (4.6)$$

Consequently,

$$\alpha^*\beta + \alpha\beta^* = 2\Re(\alpha^*\beta). \quad (4.7)$$

Then we can rewrite $A^\dagger\sigma_x A\sigma_x$ as

$$A^\dagger\sigma_x A\sigma_x = \begin{pmatrix} \omega_1^* & \omega_2 \\ -\omega_2 & \omega_1 \end{pmatrix}, \quad (4.8)$$

where $\omega_1 = \alpha^2 - \beta^2 \in \mathbb{C}$ and $\omega_2 = 2\Re(\alpha^*\beta) \in \mathbb{R}$. Given that the determinant of $A^\dagger\sigma_x A\sigma_x$ is the product of the determinants of its matrices and $A \in SU(2)$, then

$$\begin{aligned}\det(A^\dagger\sigma_x A\sigma_x) &= \det(A^\dagger) \times \det(\sigma_x) \times \det(A) \times \det(\sigma_x) \\ &= (1) \times (-1) \times (1) \times (-1) \\ &= 1,\end{aligned}\quad (4.9)$$

therefore the resulting matrix is also special unitary. Its off-diagonal is formed by the real elements ω_2 and $-\omega_2$. With that, we can say

$$\begin{aligned}(A^\dagger\sigma_x A\sigma_x)^2 &= \begin{pmatrix} (\omega_1^*)^2 - \omega_2^2 & \omega_1^*\omega_2 + \omega_1\omega_2 \\ -\omega_1^*\omega_2 - \omega_1\omega_2 & \omega_1^2 - \omega_2^2 \end{pmatrix} \\ &= \begin{pmatrix} (\omega_1^*)^2 - \omega_2^2 & (\omega_1^* + \omega_1)\omega_2 \\ -(\omega_1^* + \omega_1)\omega_2 & \omega_1^2 - \omega_2^2 \end{pmatrix}.\end{aligned}\quad (4.10)$$

This is also an $SU(2)$ matrix with real off-diagonal, since $\det(A^\dagger\sigma_x A\sigma_x)^2 = (\det(A^\dagger\sigma_x A\sigma_x))^2 = 1$, $\omega_1^* + \omega_1 = 2\Re(\omega_1)$, and $\omega_2 \in \mathbb{R}$. Then the final form of Equation (4.1) can be represented as

$$(A^\dagger\sigma_x A\sigma_x)^2 = \begin{pmatrix} z^* & x \\ -x & z \end{pmatrix} = U, \quad (4.11)$$

with $z = \omega_1^2 - \omega_2^2 \in \mathbb{C}$ and $x = 2\Re(\omega_1)\omega_2 \in \mathbb{R}$, demonstrating that the adapted decomposition scheme yields $SU(2)$ matrices with real off-diagonal elements. $\square$

The lemma leads to the following theorem for the real off-diagonal case.

Theorem 1. *Every $SU(2)$ operator whose matrix representation contains real-valued elements in its off-diagonal can be generated by an operator A in the form of Equation (4.1) described in Lemma 1.*

Proof. We now want to show that any matrix $U \in SU(2)$ with real off-diagonal can always be generated by the operator sequence in question. From Lemma 1,

$$U = \begin{pmatrix} z^* & x \\ -x & z \end{pmatrix} = \begin{pmatrix} (\omega_1^*)^2 - \omega_2^2 & 2\Re(\omega_1)\omega_2 \\ -2\Re(\omega_1)\omega_2 & \omega_1^2 - \omega_2^2 \end{pmatrix} \quad (4.12)$$

The goal is to rewrite ω_1 and ω_2 in terms of $x \in \mathbb{R}$ and $z \in \mathbb{C}$. First we expand

$$z = \omega_1^2 - \omega_2^2 \quad (4.13)$$

as

$$\Re(z) + i\Im(z) = \Re(\omega_1)^2 - \Im(\omega_1)^2 - \omega_2^2 + 2i\Re(\omega_1)\Im(\omega_1), \quad (4.14)$$

since $\omega_1 \in \mathbb{C}$ and $\omega_2 \in \mathbb{R}$. We know that $x = 2\Re(\omega_1)\omega_2$, so for now we assume $x \neq 0$ and write

$$\omega_2 = \frac{x}{2\Re(\omega_1)}. \quad (4.15)$$

We can also isolate the real elements of Equation (4.14).

$$\Re(z) = \Re(\omega_1)^2 - \Im(\omega_1)^2 - \omega_2^2 \quad (4.16)$$

From the unitarity restriction, it is true that $|\omega_1|^2 + \omega_2^2 = 1$, which can also be expanded to

$$\Re(\omega_1)^2 + \Im(\omega_1)^2 + \omega_2^2 = 1. \quad (4.17)$$

By reorganizing the terms in Equation (4.17) as

$$\Re(\omega_1)^2 - 1 = -\Im(\omega_1)^2 - \omega_2^2, \quad (4.18)$$

Equation (4.16) may be rewritten as

$$\Re(z) = 2\Re(\omega_1)^2 - 1. \quad (4.19)$$

We isolate $\Re(\omega_1)$ and choose the positive solution and continue selecting the positive solutions where appropriate.

$$\Re(\omega_1) = \sqrt{\frac{\Re(z) + 1}{2}} \quad (4.20)$$

From Equation (4.14), after isolating the imaginary elements, $\Im(\omega_1)$ can be determined by substituting Equation (4.20) into

$$\Im(\omega_1) = \frac{\Im(z)}{2\Re(\omega_1)}, \quad (4.21)$$

yielding

$$\Im(\omega_1) = \frac{\Im(z)}{\sqrt{2(\Re(z) + 1)}}. \quad (4.22)$$

A solution for ω_1 can now be obtained from Equations (4.20) and (4.22), and a solution for ω_2 is possible by substituting Equation (4.20) into Equation (4.15). By doing this, we obtain

$$\omega_1 = \sqrt{\frac{\Re(z) + 1}{2}} + i \frac{\Im(z)}{\sqrt{2(\Re(z) + 1)}} \quad (4.23)$$

and

$$\omega_2 = \frac{x}{\sqrt{2(\Re(z) + 1)}}. \quad (4.24)$$

The next step is to construct A from Equation (4.2) by determining α and β from these solutions. We know from Lemma 1 that

$$\begin{aligned} \omega_1 &= \alpha^2 - \beta^2, \\ \omega_2 &= 2\Re(\alpha^*\beta). \end{aligned} \quad (4.25)$$

We choose β to be a real number to find a solution with real off-diagonal. From the unitarity constraint, $|\alpha|^2 + \beta^2 = 1$, so this expands to rewrite β as $\beta^2 = 1 - \Re(\alpha)^2 - \Im(\alpha)^2$. From this and the definition of ω_1 from Equation (4.25),

$$\begin{aligned} \omega_1 &= \Re(\alpha)^2 - \Im(\alpha)^2 + 2i\Re(\alpha)\Im(\alpha) - 1 + \Re(\alpha)^2 + \Im(\alpha)^2 \\ &= 2\Re(\alpha)^2 - 1 + 2i\Re(\alpha)\Im(\alpha) \end{aligned} \quad (4.26)$$

whose real part is $\Re(\omega_1) = 2\Re(\alpha)^2 - 1$. By comparing the isolated real part with the solution from Equation (4.23), we find

$$\Re(\alpha) = \sqrt{\frac{\sqrt{\frac{\Re(z)+1}{2}} + 1}{2}}. \quad (4.27)$$

With this result, a solution for β can be found. Since we assume $\beta \in \mathbb{R}$, then the definition of ω_2 in Equation (4.25) can be restated as $\omega_2 = 2\Re(\alpha^*\beta) = 2\Re(\alpha)\beta$. From the definition of x from Equation (4.12),

$$\omega_2 = \frac{x}{2\Re(\omega_1)}. \quad (4.28)$$

Using this, then

$$\begin{aligned} \beta &= \frac{\omega_2}{2\Re(\alpha)} \\ &= \frac{x}{4\Re(\omega_1)\Re(\alpha)} \\ &= \frac{x}{4\sqrt{\frac{\Re(z)+1}{2}}\sqrt{\frac{\Re(z)+1}{2}+1}}, \end{aligned} \quad (4.29)$$

which simplifies to

$$\beta = \frac{x}{2\sqrt{(\Re(z)+1)(\sqrt{\frac{\Re(z)+1}{2}+1)}}}. \quad (4.30)$$

Now only the imaginary part of α remains unsolved. Returning to Equation (4.26), we take the imaginary part to obtain $\Im(\omega_1) = 2\Re(\alpha)\Im(\alpha)$. Substituting the imaginary part of the definition of ω_1 from Equation (4.25) and the solution for $\Re(\alpha)$ from Equation (4.27) yields

$$\begin{aligned} \Im(\alpha) &= \frac{\Im(\omega_1)}{2\Re(\alpha)} \\ &= \frac{\frac{\Im(z)}{\sqrt{2(\Re(z)+1)}}}{2\sqrt{\frac{\Re(z)+1}{2}+1}}, \end{aligned} \quad (4.31)$$

which simplifies to

$$\Im(\alpha) = \frac{\Im(z)}{2\sqrt{(\Re(z)+1)(\sqrt{\frac{\Re(z)+1}{2}+1)}}}. \quad (4.32)$$

□

There are some additional considerations to be made about this result. In the case we are faced with $x = 0$, then it must be assumed that $\Re(\omega_1) = 0$ or $\omega_2 = 0$. We can choose ω_2 to be zero, which will then imply that $\Re(\alpha)$ or β is equal to zero. Since we assumed β is a real number, we can choose β to be zero. From the magnitude constraint, a possible solution is $\alpha = z^{\frac{1}{4}}$. Equation 4.32 is singular when $z = -1$, but this case necessarily implies $x = 0$, so it is covered by the separate branch $\beta = 0$, $\alpha = z^{1/4}$.

The quantum circuit as shown in Figure 14 is satisfied as a consequence of the conditions determined by Equation (4.3), which, as explained, depend on the configuration of control

qubits and are responsible for enabling or disabling the controlled operations. The final result, from Lemma 1 and Theorem 1, is a quantum circuit for a decomposition scheme for any n -controlled $SU(2)$ gate with real-valued off-diagonal. We can offer a small modification to the existing decomposition scheme to generate $SU(2)$ matrices with real-valued elements in its main diagonal. It is not mandatory in this case for the off-diagonal elements to be real-valued.

Having demonstrated the off-diagonal case, this allows us to state the following lemma for the condition where the desired $SU(2)$ has real main diagonal elements.

Lemma 2. *A modification of Lemma 1 can be made such that the resulting matrix is $SU(2)$ and contains real-valued main diagonal elements. By adding a pair of Hadamard operators enclosing the expression in Equation (4.1), as shown in*

$$HA^\dagger \sigma_x A \sigma_x A^\dagger \sigma_x A \sigma_x H = U, \quad (4.33)$$

the matrix generated by the sequence of operators meets the requirements.

Proof. Using Equation (4.11), we can rewrite Equation (4.33) and produce the $SU(2)$ matrix

$$\begin{aligned} \begin{pmatrix} \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} & -\frac{1}{\sqrt{2}} \end{pmatrix} \begin{pmatrix} z^* & x \\ -x & z \end{pmatrix} \begin{pmatrix} \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} & -\frac{1}{\sqrt{2}} \end{pmatrix} &= \begin{pmatrix} \Re(z) & -x - i\Im(z) \\ x - i\Im(z) & \Re(z) \end{pmatrix} \\ &= \begin{pmatrix} x' & -z' \\ z'^* & x' \end{pmatrix}. \end{aligned} \quad (4.34)$$

□

Importantly, the conditions required by Equation (4.3) are still satisfied due to H being its own inverse. This lemma leads to the next theorem to complete the decomposition scheme for any $SU(2)$ matrix that contains at least one real-valued diagonal.

Theorem 2. *Every $SU(2)$ operator whose matrix representation contains real-valued elements in its main diagonal can be generated by an operator A in the form of Equation (4.33) described in Lemma 2.*

Proof. From Lemma 2, Equation (4.34) has given us the expressions for $x' \in \mathbb{R}$ and $z' \in \mathbb{C}$ that generate an $SU(2)$ matrix with real main diagonal elements. We return to Theorem 1, where we show we can find possible values for α and β to construct a valid A , such as in shown in Equation (4.2), that gives us any desired $SU(2)$ matrix with real off-diagonal elements. In

other words, α and β give us a way to obtain x and z as in Equation (4.11). To do this for a target $SU(2)$ matrix with real-valued main diagonal, we first take the desired $x' \in \mathbb{R}$ and $z' \in \mathbb{C}$ from Equation (4.34) and encode them in $x \in \mathbb{R}$ and $z \in \mathbb{C}$ as follows.

$$\begin{aligned} x &= \Re(z') \\ z &= x' + i\Im(z') \end{aligned} \tag{4.35}$$

This encoding allows Equation (4.34) to yield the desired matrix. A suitable A that produces the intermediary matrix in Equation (4.34) can be found following the procedure shown in Theorem 1. By applying the two Hadamard operators at the beginning and at the end of the same expression, the target is yielded. This shows a general procedure to generate any $SU(2)$ matrix with real main diagonal elements. $\square$

As both theorems only have a strict requirement for one diagonal (the main or the off-diagonal) to have real elements, consequently both can be used to generate any $SO(2)$ operator, which are represented by the group of 2×2 orthogonal matrices with determinant 1. This leads to the following corollary.

Corollary 1. *Every $SO(2)$ matrix can be generated by Theorem 1 or Theorem 2.*

4.1.1 Cost

To help in calculating the cost of the decomposition scheme for $SU(2)$ operators with one real diagonal, we refer to the quantum circuit in Figure 14. All the controlled operators implement the decomposition for multicontrolled σ_x gates with multiple dirty auxiliary qubits from Section 3.1.1 with the optimizations from Section 3.1.2. The quantum circuit in Figure 7 shows the gate canceling of relative phase Toffoli gates that occurs in this scheme. Then, due to these optimizations, each multicontrolled σ_x with at least five qubits in Figure 14, meaning there are $k \geq 3$ control qubits and at least $k - 2$ auxiliary qubits, need at most $8k - 6$ CNOTs. Because we have subdivided the control qubits into two registers with k_1 qubits and k_2 qubits, then the maximum total number of CNOTs operators N_{CNOT} is given by

$$\begin{aligned} N_{\text{CNOT}} &= 2(8k_1 - 6) + 2(8k_2 - 6) \\ &= 16(k_1 + k_2) - 24 \\ &= 16(n - 1) - 24 \\ &= 16n - 40, \end{aligned} \tag{4.36}$$

where n is the number of qubits including all controls and the target qubit. With this and from Theorems 1 and 2, the following theorem can be stated.

Theorem 3. *The decomposition scheme for a multicontrolled $SU(2)$ matrix with one real-valued main or off-diagonal, represented in Figure 14, can be implemented as an n -qubit quantum circuit, where $n \geq 5$, with at most $16n - 40$ CNOT operators.*

4.1.2 Application to multicontrolled R_x , R_y , and R_z gates

So far, the decomposition scheme has been shown to apply to the case of a matrix with one real diagonal. The rotation operators, which were introduced in Equations (2.10), (2.11), and (2.12), satisfy this requirement, allowing us to state the following corollary.

Corollary 2. *The rotation operators R_y and R_z can be generated by the sequence of operators described in Lemma 1. The rotation operator R_x can be generated by the sequence of operators described in Lemma 2.*

Proof. We can apply the procedure explained in Theorem 1 to produce the rotation operators R_y and R_z . To produce R_x , the procedure described in Theorem 2 can be applied instead. Alternatively, by defining the operator A from Equation (4.33) as $A = R_z(\theta/4)$, then

$$R_x(\theta) = H(R_z(-\theta/4)\sigma_x R_z(\theta/4)\sigma_x)^2 H \quad (4.37)$$

is an expression that generates R_x . Similarly for R_y and R_z , by choosing $A = R_y(-\theta/4)$ and $A = R_z(-\theta/4)$ in Equation (4.1), respectively, then

$$\begin{aligned} R_y(\theta) &= (R_y(\theta/4)\sigma_x R_y(-\theta/4)\sigma_x)^2 \\ R_z(\theta) &= (R_z(\theta/4)\sigma_x R_z(-\theta/4)\sigma_x)^2. \end{aligned} \quad (4.38)$$

□

4.1.3 Application example

Acasiete et al. (2020) report the experimental implementation of discrete-time quantum walks on IBM quantum computers. Quantum walks have classical random walks as analogous and can be used to design quantum algorithms. The study demonstrates the realization of the staggered quantum walk (SQW) model, a coinless model of quantum walks defined on

arbitrary graphs and constrained by graph locality. It is noted by the authors that for certain problems coinless quantum walk models require fewer qubits to implement and are, therefore, more suitable for NISQ devices. For graphs with 2^n nodes, implementing quantum walks on quantum computers requires n qubits.

In this example, the proposed scheme for controlled rotation operators is applied to SQWs on the cycle. The SQW evolution operator on the cycle can be written as

$$U = U_1 U_0 = P^{-1} U_0 P U_0, \quad (4.39)$$

where

$$U_0 = I_2^{\otimes n-1} \otimes R_x(2\theta) \quad (4.40)$$

and P is the permutation operator

$$P = \sum_x |x+1\rangle \langle x|, \quad (4.41)$$

shifting the walker to the right and its inverse P^{-1} shifting it to the left.

The quantum circuit shown in Figure 15 is a representation of the U_1 operator. The improvement suggested to decompose P and P^{-1} is to replace the use of σ_x by $R_x(\pi)$ and use the identity

$$R_x(\theta) = H R_z(\theta) H. \quad (4.42)$$




Figure 15 – Quantum circuit of the U_1 operator restricted to 4 qubits (ACASIETE et al., 2020). Source: Adapted from Acasiete et al. (2020).

An alternative construction for multicontrolled R_z gate circuits is offered as seen in Listing 4.1. In this construction, the total number of CNOT gates grows exponentially, with 2^{n-1} CNOTs needed for $n - 1$ control qubits. Using our proposed approach, the cost $16n - 40$ in number of CNOTs starts to become advantageous from $n = 8$ qubits.

Listing 4.1 – Python method to generate the new multicontrolled R_z decomposition by Acasiete et al. (2020).

Source: Adapted from Acasiete et al. (2020).

```
1 def new_mcrz(qc, theta, q_controls, q_target):
```

```

n = len(q_controls)
3 newtheta = -theta/2**n
a = lambda n: log(n-(n*(n-1)),2)
5 qc.cx(q_controls[n-1], q_target)
qc.u1(newtheta, q_target)
7
for i in range(1, 2**(n)):
9     qc.cx(q_controls[int(a(i))], q_target)

```

Assuming no further optimizations and the number of required CNOTs for the Toffoli gate as six, we can compute the total number of CNOT gates of the quantum circuit for the operator U_1 using the authors' decomposition for multicontrolled R_z as

$$\begin{aligned}
N_{\text{theirs}} &= \times \left(\sum_{k=0}^{n-1} 2^k - (2^0 + 2^1 + 2^2) + 6 + 1 \right) \\
&= 2 \times \left(1 \left(\frac{1 - 2^n}{1 - 2} \right) \right) \\
&= 2 \times (2^n - 1) \\
&= 2^{n+1} - 2.
\end{aligned} \tag{4.43}$$

Similarly, with no other optimizations introduced, our proposed scheme for multicontrolled R_z gates yields the following CNOT count for the U_1 circuit, further inhibiting the gate count increase.

$$\begin{aligned}
N_{\text{ours}} &= 2 \times \left(\left(\sum_4^n (16k - 40) \right) + 6 + 1 \right) \\
&= 2 \times \left(\left(16 \sum_4^n k \right) - 40(n - 3) + 7 \right) \\
&= 2 \times \left(16 \left(\frac{(n - 3)(4 + n)}{2} \right) - 40n + 127 \right) \\
&= 2 \times (8n^2 + 8n - 96 - 40n + 127) \\
&= 16n^2 - 64n + 62.
\end{aligned} \tag{4.44}$$

4.2 GENERALIZED DECOMPOSITION SCHEME FOR ANY MULTICONTROLLED $SU(2)$ OPERATOR

We have proposed a decomposition scheme for any $SU(2)$ operator with real main diagonal or off-diagonal, and have also shown a quantum circuit that can perform this decomposition for multicontrolled versions of this type of $SU(2)$ operator. We now present the general case encompassing any $SU(2)$ operator.

Theorem 4. *Using Theorem 1, Theorem 2, and eigendecomposition, any $SU(2)$ operator can be generated.*

Proof. Given $U \in SU(2)$, there exists an eigendecomposition $QDQ^\dagger$ for U , where Q is unitary and its columns are formed by the eigenvectors of U , and D is the diagonal unitary matrix

$$D = \text{diag}(\lambda_1, \lambda_2) \quad (4.45)$$

with $|\lambda_1| = |\lambda_2| = 1$. Since $\det(U) = 1$, then

$$\det(D) = \lambda_1 \lambda_2 = 1. \quad (4.46)$$

Hence $D \in SU(2)$ and, in particular, D is a diagonal $SU(2)$ matrix. By Theorem 1, D can be generated by the proposed decomposition scheme.

It remains to show that the same holds for the diagonalizing matrix Q . The matrix Q obtained from the eigendecomposition is unitary, but not necessarily special unitary. However, this causes no loss of generality. Let

$$\det(Q) = e^{i\phi}, \quad (4.47)$$

and define

$$\tilde{Q} = e^{-i\phi/2} Q. \quad (4.48)$$

Then

$$\det(\tilde{Q}) = e^{-i\phi} \det(Q) = 1, \quad (4.49)$$

so $\tilde{Q} \in SU(2)$. Moreover,

$$\tilde{Q}D\tilde{Q}^\dagger = e^{-i\phi/2} QDQ^\dagger e^{i\phi/2} \quad (4.50)$$

$$= QDQ^\dagger \quad (4.51)$$

$$= U, \quad (4.52)$$

since the scalar phases cancel. Thus, without loss of generality, the diagonalizing matrix may be taken to lie in $SU(2)$.

Now note that each eigenvector may still be multiplied by an arbitrary phase factor of unit modulus without changing the eigenspace. Therefore, by choosing suitable phases for the eigenvectors, we may arrange for the main diagonal entries of $\tilde{Q}$ to be real. Since $\tilde{Q} \in SU(2)$ and has real-valued main diagonal elements, Theorem 2 applies to $\tilde{Q}$. The same argument also

applies to $\tilde{Q}^\dagger$, because $\tilde{Q}^\dagger \in SU(2)$ and its main diagonal entries are the complex conjugates of those of $\tilde{Q}$, hence still real. Therefore, both $\tilde{Q}$ and $\tilde{Q}^\dagger$ can be generated by the proposed decomposition scheme.

Combining these facts, we obtain a decomposition of U of the form

$$U = \tilde{Q} D \tilde{Q}^\dagger, \quad (4.53)$$

where D is generated by Theorem 1 and $\tilde{Q}, \tilde{Q}^\dagger$ are generated by Theorem 2. Hence any matrix $U \in SU(2)$ can be generated using Theorem 1, Theorem 2, and eigendecomposition. $\square$

A quantum circuit for a multicontrolled version of this decomposition can be constructed. It can be noted that there is a symmetry of gates in $Q D Q^\dagger$, i.e., Q at the start and its inverse $Q^\dagger$ at the end. Also note that the conditions for Equation (4.3) are still respected, and no computation is performed when one of the controls is zero. Beginning from the operator Q , when using Theorem 2 a sequence of operators $H(B^\dagger \sigma_x B \sigma_x)^2 H = Q$ in the form of Equation (4.33) is needed. However, since we want the multicontrolled version of the operators to decompose a multicontrolled $SU(2)$ gate, due of the existing symmetry only half of the operators $H B^\dagger \sigma_x B \sigma_x H$ are needed because of gate cancelations. By constructing a quantum circuit for multicontrolled Q as in Figure 16, we can use the present single-qubit gates to further reduce the circuit and maintain the overall symmetry of the decomposition scheme.




Figure 16 – The decomposition of multicontrolled Q and a possible circuit optimization. Source: Adapted from Vale et al. (2024) © 2024 IEEE.

This is shown as swapping the positions of the first H gate with the k_1 -controlled σ_x , which requires the gate

$$\tilde{H} = \frac{1}{\sqrt{2}} \begin{pmatrix} -1 & 1 \\ 1 & 1 \end{pmatrix}, \quad (4.54)$$

so that $\sigma_x H = \tilde{H} \sigma_x$. The consecutive single-qubit gates are then combined, represented by the gates C_1 and C_2 . The same can be done for the multicontrolled $Q^\dagger$ circuit, which mirrors

the circuit for multicontrolled Q and uses $C_1^\dagger$ and $C_2^\dagger$, as they are inverse gates. The quantum circuit for multicontrolled D follows the decomposition scheme for the case of multicontrolled $SU(2)$ gates with real off-diagonal, as in Figure 14.

The quantum circuit to decompose any multicontrolled $SU(2)$ gate can be realized, as Figure 17 illustrates, by concatenating the three circuits for multicontrolled Q , D , and $Q^\dagger$. Because of the way the circuit is structured, other optimizations are possible. The k_1 -controlled σ_x gate from the $Q^\dagger$ part of the circuit cancels with the same gate from the D part of the circuit. The k_2 -controlled σ_x gate from the $Q^\dagger$ part of the circuit matches with the first k_2 -controlled σ_x gate from the D part of the circuit to allow the relative phase Toffoli canceling optimizations of Section 3.1.1.




Figure 17 – Proposed generalized decomposition scheme for multicontrolled $SU(2)$ gates. Source: Adapted from Vale et al. (2024) © 2024 IEEE.

4.2.1 Cost

Referring to the quantum circuit in Figure 17, the multicontrolled σ_x gates that interact with the top k_1 control qubits each need $8k_1 - 6$ CNOTs. Since two of the four gates are cancelable, one from the $Q^\dagger$ circuit and one from the D circuit, the contribution of the k_1 -controlled σ_x gates is equal to $16k_1 - 12$ CNOTs. Due to gate canceling on the first two σ_x gates controlled by the bottom k_2 control qubits (one from the $Q^\dagger$ circuit and the other from the D circuit), these gates only perform the target flipping part of their decompositions. The target flipping part of one of these gates needs twelve CNOTs to apply two Toffoli gates and up to three CNOTs per application of each relative phase Toffoli gate. There are $k_2 - 3$ pairs of relative phase Toffoli gates where gate canceling can occur, so each pair contributes with four CNOTs. One relative phase Toffoli remains, resulting in a total number of CNOTs equal to $12 + (k_2 - 3)4 + 3 = 4k_2 + 3$ for the two reduced k_2 -controlled σ_x gates. Finally, there is no additional gate canceling in the remaining two k_2 -controlled σ_x gates. Therefore, the total

cost in CNOTs of the circuit N_{CNOT} is at most

$$\begin{aligned}
 N_{\text{CNOT}} &= (16k_1 - 12) + 2(4k_2 + 3) + (16k_2 - 12) \\
 &= 16(k_1 + k_2) + 8k_2 - 18 \\
 &= 16(n - 1) + 8\left\lfloor \frac{n-1}{2} \right\rfloor - 18 \\
 &= 16n + 8\left\lfloor \frac{n-1}{2} \right\rfloor - 34.
 \end{aligned} \tag{4.55}$$

Given that $\lfloor \frac{n-1}{2} \rfloor = \frac{n-1}{2}$ for odd n and $\lfloor \frac{n-1}{2} \rfloor = \frac{n}{2} - 1$ for even n , we have the following final CNOT count.

$$N_{\text{CNOT}} = \begin{cases} 20n - 38, n \text{ odd} \\ 20n - 42, n \text{ even} \end{cases} \tag{4.56}$$

The result from Equation (4.56) can be formalized in the following theorem.

Theorem 5. *The decomposition scheme for a multicontrolled $SU(2)$ matrix, represented in Figure 17, can be implemented as an n -qubit quantum circuit, where $n \geq 3$, with at most $20n - 38$ CNOT operators if n is odd or $20n - 42$ CNOT operators if n is even.*

5 DISCUSSION

From the approach to decompose multicontrolled σ_x gates by He et al. (2017) exposed in Section 3.1.3, we hoped to modify the quantum circuit to a more general construction. In our preliminary investigations, we noticed that the overall structure of the scheme shown in Figure 8 presented some symmetries in the disposition of quantum gates that allowed Equations (3.2), (3.3), (3.4), and (3.5) to hold. This property revealed the possibility of constructing other decompositions from this scheme by plugging in different gates in the positions of the single-qubit gates from the original circuit, as long as we maintain the pairings of inverse gates.

This change led to the first contribution of this work, which improved decompositions for each multicontrolled rotation operator, namely R_x , R_y , and R_z , compared to the known literature at the time. By using our approach, however, auxiliary qubits are not needed. Since the auxiliary qubit was included in the original scheme by He et al. (2017) to account for a phase manipulation of the target qubit to produce the controlled σ_z operator, in the modified decomposition schemes, we can ignore the auxiliary qubit as the results produced do not need a specific phase. By incorporating the optimizations for multicontrolled σ_x with auxiliary qubits (ITEN et al., 2016), we can generate the quantum circuits for multicontrolled rotation gates using only $O(16n)$ CNOTs, which is an improvement over the $O(28n)$ CNOTs produced using the decomposition scheme for multicontrolled $SU(2)$ gates proposed by Barenco et al. (1995) with optimizations presented by Iten et al. (2016) included.

Due to the pairs of gates and their conjugate transpose and the even number of σ_x gates applied to the target qubit, the determinant of the transformation on that qubit is equal to 1. In other words, no matter which gates we choose to replace the original S and $S^\dagger$ pair, the new circuit always produces multicontrolled $SU(2)$ gates. Since the modified scheme had a successful result, our later studies focused on generalizing the quantum circuit and learning what gate class—specifically, the subset of $SU(2)$ with multiple controls—we could construct using the general scheme. During this examination, we proposed the decomposition scheme for multicontrolled $SU(2)$ gates with real-valued off-diagonal, requiring the results presented in Lemma 1 and Theorem 1 of Section 4.1. In this scheme, illustrated as a quantum circuit in Figure 14, the gates A and $A^\dagger$ have a similar role to the S , $S^\dagger$ pair from the original scheme by He et al. (2017), but without the need for auxiliary qubits, no phase manipulation occurs.

From prior experience in exploring constructions for the multicontrolled rotation gates,

we already knew a way to transform a real-valued off-diagonal matrix into a real-valued main diagonal matrix by enclosing the matrix with two Hadamard operators. We used this knowledge to derive the decomposition scheme for multicontrolled $SU(2)$ gates with real main diagonal, which uses what is formally stated in Lemma 2 and Theorem 2 of Section 4.1. In Figure 14, adding one Hadamard gate to the beginning and one to the end on the target qubit constructs the new circuit. Consequently, $O(16n)$ CNOTs are needed to implement a multicontrolled $SU(2)$ operator with at least one real diagonal, as claimed in Theorem 3. The advantage obtained in the first result stayed true for any gate of this type and not only multicontrolled rotation operators.

The constructions so far mentioned served as a component for a larger circuit that extends the function of the proposed decomposition scheme to encompass any multicontrolled $SU(2)$ gate, regardless of the nature of the elements of its diagonals. Theorem 4 of Section 4.2 formally states how to decompose any $SU(2)$ operator, and is required for the generalized $SU(2)$ decomposition scheme. The quantum circuit in Figure 17 illustrates our proposal. The circuit has three parts, shown as the multicontrolled gates $Q^\dagger$, D , and Q , which come from the eigendecomposition of the $SU(2)$ gate. In our approach, we synthesize the middle part of the circuit, represented by the diagonal matrix D , using the approach from Theorem 1. Initially, the first and third parts are synthesized by following the approach from Theorem 2, but after optimizations the number of gates each part uses is roughly halved, leading to further gate cancelations with the middle part of the circuit. A total of $O(20n)$ CNOT gates is necessary to generate the circuit, still providing an advantage over the $O(28n)$ CNOTs in the literature at the time of designing our approach.

A possible further optimization that was missed in the published work and reduces the CNOT count was offered by Rosa, Duzzioni and Santiago (2025). Since the first and last parts of the generalized decomposition contain a multicontrolled gate and its conjugate transpose, only Q and $Q^\dagger$ and no controls are needed. The intent is for the action on the target qubit to either result in a transformation that is equivalent to the desired $SU(2)$ matrix or all operators should cancel each other with no action being performed. With this change, the gate cancelations described in Section 4.2 are no longer possible, but the CNOT cost is reduced from $20n$ to the $16n - 40$ of the middle circuit, which uses the decomposition scheme for one real diagonal.

The proposal as presented was used by Silva et al. (2024) to construct an approximate decomposition for multicontrolled $U(2)$ gates. The approximate decomposition is based on the

method introduced by Silva and Park (2022), seen in Section 3.2.2, an exact decomposition for the same type of operator with linear depth but a quadratic of gates. In taking consecutive roots of the operator U , whose multicontrolled version is being decomposed, the new approach chooses the number of control qubits n_c determining the 2^{n_c-1} -th root. This number is chosen such that $U^{1/2^{n_c-1}} = I$ up to an error ϵ . In a modification of the original circuit, the central multicontrolled roots of U in each half of the circuit is then ignored. With this, the overhead in cost of the algorithm is decomposing the multicontrolled R_x gates using the scheme presented in Section 4.1, thus making it linear. Moreover, Silva et al. (2024) also extend Lemmas 1 and 2 for multiple $SU(2)$ gates with one real diagonal, each with a distinct target, controlled by the same set of qubits. With $n \geq 8$ qubits and n_t distinct target qubits, an $(n - 2)$ -controlled gate controlling multiple $SU(2)$ gates has a total number of CNOTs equal to $16n + 16(n_t - 1) - 40$. This is used to improve the approximate decomposition so that each multicontrolled R_x , when sharing control qubits, uses this extended approach. If the individual decompositions for multicontrolled $SU(2)$ from Section 4.1 were used instead, $n_t(16n - 40)$ CNOTs would be required.

Sato et al. (2024) introduced an algorithm to solve linear hyperbolic partial differential equations using Hamiltonian simulation. The method delineates a quantum circuit representation of time evolution operators for Hamiltonian simulation by differential operators. This is done by diagonalizing each term of the Hamiltonian using the Bell basis. In the approximated expression of the time evolution operator for which a quantum circuit was built, the multicontrolled R_z gate occurs. The implementation of the decomposition for this gate follows the procedure described in Section 4.1. The approximated time evolution operator achieves a quadratic number of CNOT gates.

Finally, Zindorf and Bose (2024) reported a more efficient decomposition for multicontrolled $SU(2)$ operators with no auxiliary qubit, with a significantly reduced CNOT cost. The quantum circuit of the new approach is constructed with R_z and R_x single-qubit gates as well as diagonal and multicontrolled Z gates. The authors show an efficient way to decompose a multicontrolled Z gates combined with a diagonal gate in terms of relative phase Toffoli and a single iZ gate with two controls. This allows them to leverage gate commutations to make way for circuit depth reduction due to applying more gates in one timestep or canceling CNOT gates. The result is a decomposition for multicontrolled $SU(2)$ gates without auxiliary qubits that costs $12.5n$ CNOTs, which is more efficient than the decomposition proposed in this work, even including the gate reduction noted by Rosa, Duzzioni and Santiago (2025). Both full and

linear nearest-neighbor qubit connectivity were considered, unlike in the proposed method.

To reiterate these accomplishments, Table 3 summarizes the results achieved by the proposed method and the related decomposition schemes for multicontrolled single-qubit gates, including $U(2)$ and $SU(2)$ gates.

Table 3 – Comparison of CNOT gate count and required number of auxiliary qubits of this proposal and the different decomposition schemes for multicontrolled $U(2)$ and $SU(2)$ gates in the related literature.

Gate	Reference	# of aux. qubits	# of CNOTs
$C_{n-1,n}(U), U \in U(2)$	Barenco et al. (1995)	0	$16n^2 - 60n + 42$
$C_{n-1,n}(U), U \in U(2)$	Silva and Park (2022)	0	$4n^2 - 12n + 10$
$C_{n-1,n}(W), W \in SU(2)$	Barenco et al. (1995)	0	$48n - 162$
$C_{n-1,n}(W), W \in SU(2)$	(ITEN et al., 2016)	0	$28n - 88$ (n even) or $28n - 92$ (n odd)
$C_{n-1,n}(W), W \in SU(2)$	Ours	0	$20n - 38$ (n odd) or $20n - 42$ (n even); $16n - 40$ (real diag.)
$C_{n-1,n}(W), W \in SU(2)$	Rosa, Duzzioni and Santiago (2025)	0	$16n - 40$
$C_{n-1,n}(W), W \in SU(2)$	Zindorf and Bose (2024)	0	$12.5n$

Source: The author (2026).

6 CONCLUSION

This work was an effort to reduce the CNOT gate count of quantum circuits for a familiar type of multicontrolled gate, with focus solely on $SU(2)$ gates. The implementation of the proposed decomposition scheme for multicontrolled $SU(2)$ operators is publicly available in the qclib library (ARAUJO et al., 2022) for the Python programming language. Furthermore, the particular cases for multicontrolled rotation operators R_x , R_y , and R_z are implemented in the Qiskit software development kit (CONTRIBUTORS, 2023) according to this proposal as the method `_mcsu2_real_diagonal()` in the `multi_control_rotation_gates.py` file of the `qiskit.synthesis.multi_controlled` library. All of the contributions in our proposal have been published in IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems (Volume: 43, Issue: 3, March 2024) as “Circuit Decomposition of Multicontrolled Special Unitary Single-Qubit Gates”, Vale et al. (2024) © 2011 IEEE.

The proposed decomposition scheme for multicontrolled $SU(2)$ operators surpasses the approaches for general single-qubit gates (BARENCO et al., 1995; SILVA; PARK, 2022) and the methods designed for the same gate class (BARENCO et al., 1995; ITEN et al., 2016) that were known at the time of publication, including the previous state-of-the-art in the literature (ITEN et al., 2016). The proposal by Rosa, Duzzioni and Santiago (2025) contributed in offering an improvement to our decomposition scheme, and the novel approach by Zindorf and Bose (2024) outperformed ours with a reduced CNOT count. Like the methods used for comparison, and in contrast with the approach we based ours on, no auxiliary qubits are necessary.

Some limitations of the current work include the lack of adaptability of the circuit for particular hardware architectures. We took no steps, for example, to explicitly exploit device properties such as different configurations of the qubit connectivity graph. Therefore, in real-world environments, the expected improvements may be reduced. Most of the other methods compared with this proposal also rely on hardware-agnostic approaches, as is typical for many gate decomposition methods. One exception is the method by Zindorf and Bose (2024), which also takes into account linear nearest-neighbor connectivity.

Another attribute we can explore is the addition of auxiliary qubits. Since we noted as early as during initial investigations that a new decomposition method for this type of gate would not need auxiliary qubits, this was not within the scope of our work so far. However, it is worth noting that if more qubits become available for computation, more efficient methods

could take advantage of the additional resources.

Finally, we did not consider theoretical bounds on the number of gates and the circuit depth when decomposing multicontrolled $SU(2)$ gates. Because of this, at this point in the research, it is impossible to say whether the proposed method, using the provided method of decomposition, approaches or achieves optimality. Thus, it remains unknown to us if better techniques that make no use of auxiliary qubits could exist. We do know that reduced gate count and depth can be achieved, however, as evidenced by the alternative decomposition method by Zindorf and Bose (2024).

So far, our results include an important set of multicontrolled single-qubit gates encompassing the rotation operators R_x , R_y , and R_z . Although not universally applicable to all single-qubit gates with multiple controls, our proposal can be advantageous when the phase of a $U(2)$ operator can be ignored during computation. Furthermore, many quantum algorithms can already employ this new decomposition scheme for multicontrolled $SU(2)$ operators (HARROW; HASSIDIM; LLOYD, 2009; LLOYD; MOHSENI; REBENTROST, 2014; LOW; YODER; CHUANG, 2014; SCHULD; SINAYSKIY; PETRUCCIONE, 2016; BIAMONTE et al., 2017; PARK; PETRUCCIONE; RHEE, 2019; ORÚS; MUGEL; LIZASO, 2019; VERAS et al., 2020; GLEINIG; HOEFLER, 2021; MOZAFARI et al., 2021; SOUZA; CARVALHO; FERREIRA, 2021; GUO et al., 2022; LI et al., 2023). With all things considered, at the time of this study, this contribution benefits quantum circuit design by providing a decomposition scheme with predictable and lower cost for NISQ devices. Inspired by the new result from Zindorf and Bose (2024), it remains a worthwhile pursuit to find new methods to further reduce computational cost for common operations in NISQ devices.

Up to this point in this study, a promising result for a large group of quantum gates inspires further explorations into improving the synthesis of circuits for any multicontrolled single-qubit gates. It is desirable to have a linear decomposition scheme in a circuit design toolkit, both in the depth and number of gates, for any controlled $U(2)$ operator. Preliminary observations of problems that arise when focusing on $U(2)$ operators reveal the need to use additional qubits as auxiliaries. This need seems particularly true for decompositions that employ roots of gates. That said, the decomposition method for multicontrolled single-qubit gates proposed by Silva and Park (2022) uses no auxiliary qubits. Both that construction and the proposal by Barenco et al. (1995) have quadratic depth but involve computations that yield increasingly significant numerical errors as the number of qubits increases.

Due to the difficulties that emerged from how the mentioned methods were designed, relying on roots of gates is not viable in practice for large numbers of control qubits. Prevention

of numerical errors is an issue that must be addressed. In addition, linear gate count and circuit depth is worth achieving for this more general type of quantum operator, namely multicontrolled single-qubit operators. As an open problem, it remains to be seen if it is possible to decompose this larger group of operators in such a way that the quantum circuits have linear gate count and depth without resorting to approximate decompositions. A linear, albeit approximate, decomposition for multicontrolled $U(2)$ was proposed by Silva et al. (2024). Nevertheless, an exact and linear approach remains an open challenge in future explorations of quantum circuit design. Towards this goal, a new method to decompose multicontrolled gates should avoid the pitfalls mentioned.

This thesis has presented a decomposition strategy for multicontrolled $SU(2)$ gates without auxiliary qubits that significantly reduces the CNOT count while incidentally yielding shallower circuit depths. By prioritizing the reduction of the number of CNOT gates, which is an essential entangling gate frequently implemented in NISQ architectures, the significance of this work is to serve as an efficiency baseline and show that targeted algorithmic enhancements can offer performance gains during this transitional period, as resource availability increases and quantum architectures evolve towards fault tolerance. One lesson that can be taken from this work is that the exploration of particular gate structures can be a valuable methodology to discover optimization opportunities that cannot be provided by general synthesis schemes. As the field advances, these specialized decompositions can be iterated and improved upon and act as building blocks, alleviating compilation overhead. In this way, the proposed decomposition is a contribution to the journey of making complex quantum algorithms practically viable while operating within the constraints of current hardware. The main goal of this thesis is not to achieve optimality, but to establish a new structure of decompositions for multicontrolled $SU(2)$ operators, with asymptotic improvement in the CNOT count compared to the available literature approaches at the time, and with special relevance for subclasses of rotation operators.

REFERENCES

- AARONSON, S.; GOTTESMAN, D. Improved simulation of stabilizer circuits. *Physical Review A—Atomic, Molecular, and Optical Physics*, APS, v. 70, n. 5, p. 052328, 2004.
- ACASIETE, F.; AGOSTINI, F. P.; MOQADAM, J. K.; PORTUGAL, R. Implementation of quantum walks on ibm quantum computers. *Quantum Information Processing*, Springer, v. 19, n. 12, p. 426, 2020.
- ACHARYA, R. and; ABANIN, D. A.; AGHABABAIE-BENI, L.; ALEINER, I.; ANDERSEN, T. I.; ANSMANN, M.; ARUTE, F.; ARYA, K.; ASFAW, A.; ASTRAKHANTSEV, N.; ATALAYA, J.; BABBUSH, R.; BACON, D.; BALLARD, B.; BARDIN, J. C.; BAUSCH, J.; BENGTSSON, A.; BILMES, A.; BLACKWELL, S.; BOIXO, S.; BORTOLI, G.; BOURASSA, A.; BOVAIRD, J.; BRILL, L.; BROUGHTON, M.; BROWNE, D. A.; BUCHEA, B.; BUCKLEY, B. B.; BUELL, D. A.; BURGER, T.; BURKETT, B.; BUSHNELL, N.; CABRERA, A.; CAMPERO, J.; CHANG, H.-S.; CHEN, Y.; CHEN, Z.; CHIARO, B.; CHIK, D.; CHOU, C.; CLAES, J.; CLELAND, A. Y.; COGAN, J.; COLLINS, R.; CONNER, P.; COURTNEY, W.; CROOK, A. L.; CURTIN, B.; DAS, S.; DAVIES, A.; LORENZO, L. D.; DEBROY, D. M.; DEMURA, S.; DEVORET, M.; PAOLO, A. D.; DONOHOE, P.; DROZDOV, I.; DUNSWORTH, A.; EARLE, C.; EDLICH, T.; EICKBUSCH, A.; ELBAG, A. M.; ELZOUKA, M.; ERICKSON, C.; FAORO, L.; FARHI, E.; FERREIRA, V. S.; BURGOS, L. F.; FORATI, E.; FOWLER, A. G.; FOXEN, B.; GANJAM, S.; GARCIA, G.; GASCA, R.; GENOIS, E.; GIANG, W.; GIDNEY, C.; GILBOA, D.; GOSULA, R.; DAU, A. G.; GRAUMANN, D.; GREENE, A.; GROSS, J. A.; HABEGGER, S.; HALL, J.; HAMILTON, M. C.; HANSEN, M.; HARRIGAN, M. P.; HARRINGTON, S. D.; HERAS, F. J. H.; HESLIN, S.; HEU, P.; HIGGOTT, O.; HILL, G.; HILTON, J.; HOLLAND, G.; HONG, S.; HUANG, H.-Y.; HUFF, A.; HUGGINS, W. J.; IOFFE, L. B.; ISAKOV, S. V.; IVELAND, J.; JEFFREY, E.; JIANG, Z.; JONES, C.; JORDAN, S.; JOSHI, C.; JUHAS, P.; KAFRI, D.; KANG, H.; KARAMLOU, A. H.; KECHEDZHI, K.; KELLY, J.; KHAIRE, T.; KHATTAR, T.; KHEZRI, M.; KIM, S.; KLIMOV, P. V.; KLOTS, A. R.; KOBRIN, B.; KOHLI, P.; KOROTKOV, A. N.; KOSTRITSA, F.; KOTHARI, R.; KOZLOVSKII, B.; KREIKEBAUM, J. M.; KURILOVICH, V. D.; LACROIX, N.; LANDHUIS, D.; LANGE-DEI, T.; LANGLEY, B. W.; LAPTEV, P.; LAU, K.-M.; GUEVEL, L. L.; LEDFORD, J.; LEE, J.; LEE, K.; LENSKY, Y. D.; LEON, S.; LESTER, B. J.; LI, W. Y.; LI, Y.; LILL, A. T.; LIU, W.; LIVINGSTON, W. P.; LOCHARLA, A.; LUCERO, E.; LUNDAHL, D.; LUNT, A.; MADHUK, S.; MALONE, F. D.; MALONEY, A.; MANDRà, S.; MANYIKA, J.; MARTIN, L. S.; MARTIN, O.; MARTIN, S.; MAXFIELD, C.; MCCLEAN, J. R.; MCEWEN, M.; MEEKS, S.; MEGRANT, A.; MI, X.; MIAO, K. C.; MIESZALA, A.; MOLAVI, R.; MOLINA, S.; MONTAZERI, S.; MORVAN, A.; MOVASSAGH, R.; MRUCZKIEWICZ, W.; NAAMAN, O.; NEELEY, M.; NEILL, C.; NERSISYAN, A.; NEVEN, H.; NEWMAN, M.; NG, J. H.; NGUYEN, A.; NGUYEN, M.; NI, C.-H.; NIU, M. Y.; O'BRIEN, T. E.; OLIVER, W. D.; OPREMCAK, A.; OTTOSSON, K.; PETUKHOV, A.; PIZZUTO, A.; PLATT, J.; POTTER, R.; PRITCHARD, O.; PRYADKO, L. P.; QUINTANA, C.; RAMACHANDRAN, G.; REAGOR, M. J.; REDDING, J.; RHODES, D. M.; ROBERTS, G.; ROSENBERG, E.; ROSENFELD, E.; ROUSHAN, P.; RUBIN, N. C.; SAEI, N.; SANK, D.; SANKARAGOMATHI, K.; SATZINGER, K. J.; SCHURKUS, H. F.; SCHUSTER, C.; SENIOR, A. W.; SHEARN, M. J.; SHORTER, A.; SHUTTY, N.; SHVARTS, V.; SINGH, S.; SIVAK, V.; SKRUZNY, J.; SMALL, S.; SMELYANSKIY, V.; SMITH, W. C.; SOMMA, R. D.; SPRINGER, S.; STERLING, G.; STRAIN, D.; SUCHARD, J.; SZASZ, A.; SZTEIN, A.; THOR, D.; TORRES, A.; TORUNBALCI, M. M.; VAISHNAV, A.; VARGAS,

- J.; VDOVICHEV, S.; VIDAL, G.; VILLALONGA, B.; HEIDWEILLER, C. V.; WALTMAN, S.; WANG, S. X.; WARE, B.; WEBER, K.; WEIDEL, T.; WHITE, T.; WONG, K.; WOO, B. W. K.; XING, C.; YAO, Z. J.; YE, P.; YING, B.; YOO, J.; YOSRI, N.; YOUNG, G.; ZALCMAN, A.; ZHANG, Y.; ZHU, N.; ZOBRIST, N. Quantum error correction below the surface code threshold. *Nature*, Nature Publishing Group UK London, v. 638, n. 8052, p. 920–926, 2025.
- ALMUDEVER, C. G.; LAO, L.; FU, X.; KHAMMASSI, N.; ASHRAF, I.; IORGA, D.; VARSAMOPOULOS, S.; EICHLER, C.; WALLRAFF, A.; GECK, L. et al. The engineering challenges in quantum computing. In: IEEE. *Design, Automation & Test in Europe Conference & Exhibition (DATE)*, 2017. Lausanne, Switzerland: IEEE, 2017. p. 836–845.
- ARAUJO, I. F.; ARAÚJO, I. C. S.; SILVA, L. D. da; BLANK, C.; SILVA, A. J. da. *Quantum computing library*. 2022. Available at: <<https://github.com/qclib/qclib>>.
- ARUTE, F.; ARYA, K.; BABBUS, R.; BACON, D.; BARDIN, J. C.; BAREND, R.; BISWAS, R.; BOIXO, S.; BRANDAO, F. G.; BUELL, D. A. et al. Quantum supremacy using a programmable superconducting processor. *Nature*, Nature Publishing Group, v. 574, n. 7779, p. 505–510, 2019.
- BAKER, J. M.; DUCKERING, C.; HOOVER, A.; CHONG, F. T. Decomposing quantum generalized Toffoli with an arbitrary number of ancilla. *arXiv preprint arXiv:1904.01671*, 2019.
- BALAJI, P.; CONEFREY-SHINOZAKI, C.; DRAPER, P.; ELHADERI, J. K.; GUPTA, D.; HIDALGO, L.; LYTLE, A.; RINALDI, E. *Quantum Circuits for SU(3) Lattice Gauge Theory*. 2025. Available at: <<https://arxiv.org/abs/2503.08866>>.
- BARENCO, A.; BENNETT, C. H.; CLEVE, R.; DIVINCENZO, D. P.; MARGOLUS, N.; SHOR, P.; SLEATOR, T.; SMOLIN, J. A.; WEINFURTER, H. Elementary gates for quantum computation. *Physical Review A*, American Physical Society, v. 52, p. 3457–3467, 1995.
- BAREND, R.; KELLY, J.; MEGRANT, A.; VEITIA, A.; SANK, D.; JEFFREY, E.; WHITE, T. C.; MUTUS, J.; FOWLER, A. G.; CAMPBELL, B. et al. Superconducting quantum circuits at the surface code threshold for fault tolerance. *Nature*, Nature Publishing Group UK London, v. 508, n. 7497, p. 500–503, 2014.
- BARNES, J. P.; TROUT, C. J.; LUCARELLI, D.; CLADER, B. Quantum error-correction failure distributions: Comparison of coherent and stochastic error models. *Physical Review A*, APS, v. 95, n. 6, p. 062338, 2017.
- BERGHOLM, V.; VARTIAINEN, J. J.; MÖTTÖNEN, M.; SALOMAA, M. M. Quantum circuits with uniformly controlled one-qubit gates. *Physical Review A*, APS, v. 71, n. 5, p. 052330, 2005.
- BIAMONTE, J.; WITTEK, P.; PANCOTTI, N.; REBENTROST, P.; WIEBE, N.; LLOYD, S. Quantum machine learning. *Nature*, Nature Publishing Group UK London, v. 549, n. 7671, p. 195–202, 2017.
- BLATT, R.; WINELAND, D. Entangled states of trapped atomic ions. *Nature*, Nature Publishing Group UK London, v. 453, n. 7198, p. 1008–1015, 2008.
- BRAVYI, S.; CROSS, A. W.; GAMBETTA, J. M.; MASLOV, D.; RALL, P.; YODER, T. J. High-threshold and low-overhead fault-tolerant quantum memory. *Nature*, Nature Publishing Group UK London, v. 627, n. 8005, p. 778–782, 2024.

- BRUGIÈRE, T. G. de; BABOULIN, M.; VALIRON, B.; ALLOUCHE, C. Quantum circuits synthesis using Householder transformations. *Computer Physics Communications*, Elsevier, v. 248, p. 107001, 2020.
- CAMPBELL, E. T.; TERHAL, B. M.; VUILLOT, C. Roads towards fault-tolerant universal quantum computation. *Nature*, Nature Publishing Group UK London, v. 549, n. 7671, p. 172–179, 2017.
- CAO, Y.; ROMERO, J.; OLSON, J. P.; DEGROOTE, M.; JOHNSON, P. D.; KIEFEROVÁ, M.; KIVLICHAN, I. D.; MENKE, T.; PEROPADRE, B.; SAWAYA, N. P. et al. Quantum chemistry in the age of quantum computing. *Chemical Reviews*, ACS Publications, v. 119, n. 19, p. 10856–10915, 2019.
- CONTRIBUTORS, Q. *Qiskit: An Open-source Framework for Quantum Computing*. 2023.
- CÓRCOLES, A. D.; KANDALA, A.; JAVADI-ABHARI, A.; MCCLURE, D. T.; CROSS, A. W.; TEMME, K.; NATION, P. D.; STEFFEN, M.; GAMBETTA, J. M. Challenges and opportunities of near-term quantum computing systems. *Proceedings of the IEEE*, IEEE, v. 108, n. 8, p. 1338–1352, 2019.
- DEUTSCH, D. Quantum theory, the church–turing principle and the universal quantum computer. *Proceedings of the Royal Society of London. A. Mathematical and Physical Sciences*, The Royal Society London, v. 400, n. 1818, p. 97–117, 1985.
- DEUTSCH, D.; JOZSA, R. Rapid solution of problems by quantum computation. *Proceedings of the Royal Society of London. Series A: Mathematical and Physical Sciences*, The Royal Society London, v. 439, n. 1907, p. 553–558, 1992.
- FARHI, E.; NEVEN, H. Classification with quantum neural networks on near term processors. *arXiv preprint arXiv:1802.06002*, 2018.
- FÖSEL, T.; NIU, M. Y.; MARQUARDT, F.; LI, L. Quantum circuit optimization with deep reinforcement learning. *arXiv preprint arXiv:2103.07585*, 2021.
- FOWLER, A. G. Coping with qubit leakage in topological codes. *Physical Review A—Atomic, Molecular, and Optical Physics*, APS, v. 88, n. 4, p. 042308, 2013.
- GIDNEY, C. Halving the cost of quantum addition. *Quantum*, Verein zur Förderung des Open Access Publizierens in den Quantenwissenschaften, v. 2, p. 74, 2018.
- GISIN, N.; RIBORDY, G.; TITTEL, W.; ZBINDEN, H. Quantum cryptography. *Reviews of Modern Physics*, APS, v. 74, n. 1, p. 145, 2002.
- GLEINIG, N.; HOEFLER, T. An efficient algorithm for sparse quantum state preparation. In: IEEE. *2021 58th ACM/IEEE Design Automation Conference (DAC)*. San Francisco, CA, USA: IEEE, 2021. p. 433–438.
- Google Quantum AI and Collaborators; ABANIN, D. A.; ACHARYA, R.; AGHABABAIE-BENI, L.; AIGELDINGER, G.; AJAY, A.; ALCARAZ, R.; ALEINER, I.; ANDERSEN, T. I.; ANSMANN, M.; ARUTE, F.; ARYA, K.; ASFAW, A.; ASTRAKHANTSEV, N.; ATALAYA, J.; BABBUSH, R.; BACON, D.; BALLARD, B.; BARDIN, J. C.; BENGS, C.; BENGTSOON, A.; BILMES, A.; BOIXO, S.; BORTOLI, G.; BOURASSA, A.; BOVAIRD, J.; BOWERS, D.; BRILL, L.; BROUGHTON, M.; BROWNE, D. A.; BUCHEA, B.; BUCKLEY, B. B.;

BUELL, D. A.; BURGER, T.; BURKETT, B.; BUSHNELL, N.; CABRERA, A.; CAMPERO, J.; CHANG, H.-S.; CHEN, Y.; CHEN, Z.; CHIARO, B.; CHIH, L.-Y.; CHIK, D.; CHOU, C.; CLAES, J.; CLELAND, A. Y.; COGAN, J.; COHEN, S.; COLLINS, R.; CONNER, P.; COURTNEY, W.; CROOK, A. L.; CURTIN, B.; DAS, S.; LORENZO, L. D.; DEBROY, D. M.; DEMURA, S.; DEVORET, M.; PAOLO, A. D.; DONOHUE, P.; DROZDOV, I.; DUNSWORTH, A.; EARLE, C.; EICKBUSCH, A.; ELBAG, A. M.; ELZOUKA, M.; ERICKSON, C.; FAORO, L.; FARHI, E.; FERREIRA, V. S.; BURGOS, L. F.; FORATI, E.; FOWLER, A. G.; FOXEN, B.; GANJAM, S.; GARCIA, G.; GASCA, R.; GENOIS, E.; GIANG, W.; GIDNEY, C.; GILBOA, D.; GOSULA, R.; DAU, A. G.; GRAUMANN, D.; GREENE, A.; GROSS, J. A.; GU, H.; HABEGGER, S.; HALL, J.; HAMAMURA, I.; HAMILTON, M. C.; HANSEN, M.; HARRIGAN, M. P.; HARRINGTON, S. D.; HESLIN, S.; HEU, P.; HIGGOTT, O.; HILL, G.; HILTON, J.; HONG, S.; HUANG, H.-Y.; HUFF, A.; HUGGINS, W. J.; IOFFE, L. B.; ISAKOV, S. V.; IVELAND, J.; JEFFREY, E.; JIANG, Z.; JIN, X.; JONES, C.; JORDAN, S.; JOSHI, C.; JUHAS, P.; KABEL, A.; KAFRI, D.; KANG, H.; KARAMLOU, A. H.; KECHEDZHI, K.; KELLY, J.; KHAIRE, T.; KHATTAR, T.; KHEZRI, M.; KIM, S.; KING, R.; KLIMOV, P. V.; KLOTS, A. R.; KOBRIN, B.; KOROTKOV, A. N.; KOSTRITSKA, F.; KOTHARI, R.; KREIKEBAUM, J. M.; KURILOVICH, V. D.; KYOSEVA, E.; LANDHUIS, D.; LANGE-DEI, T.; LANGLEY, B. W.; LAPTEV, P.; LAU, K.-M.; GUEVEL, L. L.; LEDFORD, J.; LEE, J.; LEE, K.; LENSKEY, Y. D.; LEON, S.; LESTER, B. J.; LI, W. Y.; LILL, A. T.; LIU, W.; LIVINGSTON, W. P.; LOCHARLA, A.; LUCERO, E.; LUNDAHL, D.; LUNT, A.; MADHUK, S.; MALONE, F. D.; MALONEY, A.; MANDRÀ, S.; MANYIKA, J. M.; MARTIN, L. S.; MARTIN, O.; MARTIN, S.; MATIAS, Y.; MAXFIELD, C.; MCCLEAN, J. R.; MCEWEN, M.; MEEKS, S.; MEGRANT, A.; MI, X.; MIAO, K. C.; MIESZALA, A.; MINEV, Z.; MOLAVI, R.; MOLINA, S.; MONTAZERI, S.; MORVAN, A.; MOVASSAGH, R.; MRUCZKIEWICZ, W.; NAAMAN, O.; NEELEY, M.; NEILL, C.; NERSISYAN, A.; NEVEN, H.; NEWMAN, M.; NG, J. H.; NGUYEN, A.; NGUYEN, M.; NI, C.-H.; NIU, M. Y.; OAS, L.; O'BRIEN, T. E.; OLIVER, W. D.; OPREMCAK, A.; OTTOSSON, K.; PETUKHOV, A.; PIZZUTO, A.; PLATT, J.; POTTER, R.; PRITCHARD, O.; PRYADKO, L. P.; QUINTANA, C.; RAMACHANDRAN, G.; RAMANATHAN, C.; REAGOR, M. J.; REDDING, J.; RHODES, D. M.; ROBERTS, G.; ROSENBERG, E.; ROSENFELD, E.; ROUSHAN, P.; RUBIN, N. C.; SAEI, N.; SANK, D.; SANKARAGOMATHI, K.; SATZINGER, K. J.; SCHMIDHUBER, A.; SCHURKUS, H. F.; SCHUSTER, C.; SCHUSTER, T.; SHEARN, M. J.; SHORTER, A.; SHUTTY, N.; SHVARTS, V.; SIVAK, V.; SKRUZNY, J.; SMALL, S.; SMELYANSKIY, V.; SMITH, W. C.; SOMMA, R. D.; SPRINGER, S.; STERLING, G.; STRAIN, D.; SUCHARD, J.; SUCHSLAND, P.; SZASZ, A.; SZTEIN, A.; THOR, D.; TOMITA, E.; TORRES, A.; TORUNBALCI, M. M.; VAISHNAV, A.; VARGAS, J.; VDOVICHEV, S.; VIDAL, G.; VILLALONGA, B.; HEIDWEILLER, C. V.; WALTMAN, S.; WANG, S. X.; WARE, B.; WEBER, K.; WEIDEL, T.; WESTERHOUT, T.; WHITE, T.; WONG, K.; WOO, B. W. K.; XING, C.; YAO, Z. J.; YEH, P.; YING, B.; YOO, J.; YOSRI, N.; YOUNG, G.; ZALCMAN, A.; ZHANG, C.; ZHANG, Y.; ZHU, N.; ZOBRIST, N. Observation of constructive interference at the edge of quantum ergodicity. *Nature*, v. 646, n. 8086, p. 825–830, Oct. 2025. ISSN 0028-0836, 1476-4687. Available at: <<https://www.nature.com/articles/s41586-025-09526-6>>.

GOSSET, D.; KLIUCHNIKOV, V.; MOSCA, M.; RUSSO, V. An algorithm for the t-count. *arXiv preprint arXiv:1308.4134*, 2013.

GOTTESMAN, D. Theory of fault-tolerant quantum computation. *Physical Review A*, APS, v. 57, n. 1, p. 127, 1998.

GROVER, L.; RUDOLPH, T. Creating superpositions that correspond to efficiently integrable probability distributions. *arXiv preprint quant-ph/0208112*, 2002.

GROVER, L. K. A fast quantum mechanical algorithm for database search. In: *Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing*. New York, NY, USA: Association for Computing Machinery, 1996. (STOC '96), p. 212–219.

GUO, M.; LIU, H.; LI, Y.; LI, W.; GAO, F.; QIN, S.; WEN, Q. Quantum algorithms for anomaly detection using amplitude estimation. *Physica A: Statistical Mechanics and its Applications*, Elsevier, v. 604, p. 127936, 2022.

HÄNER, T.; STEIGER, D. S.; SVORE, K.; TROYER, M. A software methodology for compiling quantum programs. *Quantum Science and Technology*, IOP Publishing, v. 3, n. 2, p. 020501, 2018.

HARROW, A. W.; HASSIDIM, A.; LLOYD, S. Quantum algorithm for linear systems of equations. *Physical Review Letters*, APS, v. 103, n. 15, p. 150502, 2009.

HARROW, A. W.; NIELSEN, M. A. Robustness of quantum gates in the presence of noise. *Physical Review A*, APS, v. 68, n. 1, p. 012308, 2003.

HAVLÍČEK, V.; CÓRCOLES, A. D.; TEMME, K.; HARROW, A. W.; KANDALA, A.; CHOW, J. M.; GAMBETTA, J. M. Supervised learning with quantum-enhanced feature spaces. *Nature*, Nature Publishing Group UK London, v. 567, n. 7747, p. 209–212, 2019.

HE, Y.; LUO, M.-X.; ZHANG, E.; WANG, H.-K.; WANG, X.-F. Decompositions of n-qubit Toffoli gates with linear circuit complexity. *International Journal of Theoretical Physics*, Springer, v. 56, n. 7, p. 2350–2361, 2017.

HERMAN, D.; GOOGIN, C.; LIU, X.; GALDA, A.; SAFRO, I.; SUN, Y.; PISTOIA, M.; ALEXEEV, Y. A survey of quantum computing for finance. *arXiv preprint arXiv:2201.02773*, 2022.

HIETALA, K.; RAND, R.; HUNG, S.-H.; WU, X.; HICKS, M. A verified optimizer for quantum circuits. *Proceedings of the ACM on Programming Languages*, ACM New York, NY, USA, v. 5, n. POPL, p. 1–29, 2021.

ITEN, R.; COLBECK, R.; KUKULJAN, I.; HOME, J.; CHRISTANDL, M. Quantum circuits for isometries. *Physical Review A*, APS, v. 93, n. 3, p. 032318, 2016.

KANDALA, A.; MEZZACAPO, A.; TEMME, K.; TAKITA, M.; BRINK, M.; CHOW, J. M.; GAMBETTA, J. M. Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets. *Nature*, Nature Publishing Group UK London, v. 549, n. 7671, p. 242–246, 2017.

KILLORAN, N.; IZAAC, J.; QUESADA, N.; BERGHOLM, V.; AMY, M.; WEEDBROOK, C. Strawberry fields: A software platform for photonic quantum computing. *Quantum*, Verein zur Förderung des Open Access Publizierens in den Quantenwissenschaften, v. 3, p. 129, 2019.

KIM, T.; CHOI, B.-S. Efficient decomposition methods for controlled-R_n using a single ancillary qubit. *Scientific Reports*, Nature Publishing Group UK London, v. 8, n. 1, p. 5445, 2018.

- KIM, Y.; EDDINS, A.; ANAND, S.; WEI, K. X.; BERG, E. V. D.; ROSENBLATT, S.; NAYFEH, H.; WU, Y.; ZALETEL, M.; TEMME, K. et al. Evidence for the utility of quantum computing before fault tolerance. *Nature*, Nature Publishing Group UK London, v. 618, n. 7965, p. 500–505, 2023.
- KROL, A. M.; SARKAR, A.; ASHRAF, I.; AL-ARS, Z.; BERTELS, K. Efficient decomposition of unitary matrices in quantum circuit compilers. *Applied Sciences*, MDPI, v. 12, n. 2, p. 759, 2022.
- LI, C.-K.; ROBERTS, R.; YIN, X. Decomposition of unitary matrices and quantum gates. *International Journal of Quantum Information*, World Scientific, v. 11, n. 01, p. 1350015, 2013.
- LI, G.; DING, Y.; XIE, Y. Tackling the qubit mapping problem for nisq-era quantum devices. In: *Proceedings of the twenty-fourth international conference on architectural support for programming languages and operating systems*. [S.l.: s.n.], 2019. p. 1001–1014.
- LI, J.; GAO, F.; LIN, S.; GUO, M.; LI, Y.; LIU, H.; QIN, S.; WEN, Q. Quantum k-fold cross-validation for nearest neighbor classification algorithm. *Physica A: Statistical Mechanics and its Applications*, Elsevier, v. 611, p. 128435, 2023.
- LIU, H.-L.; SU, H.; GONG, S.-Q.; GU, Y.-C.; TANG, H.-Y.; JIA, M.-H.; WEI, Q.; SONG, Y.; WANG, D.; ZHENG, M. et al. Robust quantum computational advantage with programmable 3050-photon gaussian boson sampling. *arXiv preprint arXiv:2508.09092*, 2025.
- LIU, W.-Q.; WEI, H.-R. Optimal synthesis of the Fredkin gate in a multilevel system. *New Journal of Physics*, IOP Publishing, v. 22, n. 6, p. 063026, 2020.
- LIU, W.-Q.; WEI, H.-R.; KWEK, L.-C. Low-cost Fredkin gate with auxiliary space. *Physical Review Applied*, APS, v. 14, n. 5, p. 054057, 2020.
- LLOYD, S.; MOHSENI, M.; REBENTROST, P. Quantum principal component analysis. *Nature Physics*, Nature Publishing Group UK London, v. 10, n. 9, p. 631–633, 2014.
- LOW, G. H.; YODER, T. J.; CHUANG, I. L. Quantum inference on Bayesian networks. *Physical Review A*, APS, v. 89, n. 6, p. 062315, 2014.
- MALVETTI, E.; ITEN, R.; COLBECK, R. Quantum circuits for sparse isometries. *Quantum*, Verein zur Förderung des Open Access Publizierens in den Quantenwissenschaften, v. 5, p. 412, 2021.
- MARIN-SANCHEZ, G.; GONZALEZ-CONDE, J.; SANZ, M. Quantum algorithms for approximate function loading. *Physical Review Research*, APS, v. 5, n. 3, p. 033114, 2023.
- MASLOV, D. Advantages of using relative-phase Toffoli gates with an application to multiple control Toffoli optimization. *Physical Review A*, American Physical Society, v. 93, p. 022311, Feb 2016. ISSN 2469-9926, 2469-9934.
- MASLOV, D.; FALCONER, S. M.; MOSCA, M. Quantum circuit placement. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, IEEE, v. 27, n. 4, p. 752–763, 2008.

MILLER, D. M.; WILLE, R.; DRECHSLER, R. Reducing reversible circuit cost by adding lines. In: IEEE. *2010 40th IEEE International Symposium on Multiple-Valued Logic*. Barcelona, Spain: IEEE, 2010. p. 217–222.

MÖTTÖNEN, M.; VARTIAINEN, J. J.; BERGHOLM, V.; SALOMAA, M. M. Quantum circuits for general multiqubit gates. *Physical Review Letters*, APS, v. 93, n. 13, p. 130502, 2004.

MOTTONEN, M.; VARTIAINEN, J. J.; BERGHOLM, V.; SALOMAA, M. M. Transformation of quantum states using uniformly controlled rotations. *arXiv preprint quant-ph/0407010*, 2004.

MOZAFARI, F.; RIENER, H.; SOEKEN, M.; MICHELI, G. D. Efficient boolean methods for preparing uniform quantum states. *IEEE Transactions on Quantum Engineering*, IEEE, v. 2, p. 1–12, 2021.

MURALI, P.; BAKER, J. M.; JAVADI-ABHARI, A.; CHONG, F. T.; MARTONOSI, M. Noise-adaptive compiler mappings for noisy intermediate-scale quantum computers. In: *Proceedings of the 24th International Conference on Architectural Support for Programming Languages and Operating Systems*. New York, NY, USA: Association for Computing Machinery, 2019. (ASPLOS '19), p. 1015–1029.

MURALI, P.; LINKE, N. M.; MARTONOSI, M.; ABHARI, A. J.; NGUYEN, N. H.; ALDERETE, C. H. Full-stack, real-system quantum computer studies: Architectural comparisons and design insights. In: *Proceedings of the 46th International Symposium on Computer Architecture*. New York, NY, USA: Association for Computing Machinery, 2019. (ISCA '19), p. 527–540.

NIELSEN, M. A.; CHUANG, I. L. *Quantum computation and quantum information*. [S.l.]: Cambridge University Press, 2010.

ORÚS, R.; MUGEL, S.; LIZASO, E. Quantum computing for finance: Overview and prospects. *Reviews in Physics*, Elsevier, v. 4, p. 100028, 2019.

PAETZNICK, A.; SILVA, M. P. da; RYAN-ANDERSON, C.; BELLO-RIVAS, J. M.; III, J. P. C.; CHERNOGUZOV, A.; DREILING, J. M.; FOLTZ, C.; FRACHON, F.; GAEBLER, J. P.; GATTERMAN, T. M.; GRANS-SAMUELSSON, L.; GRESH, D.; HAYES, D.; HEWITT, N.; HOLLIMAN, C.; HORST, C. V.; JOHANSEN, J.; LUCCHETTI, D.; MATSUOKA, Y.; MILLS, M.; MOSES, S. A.; NEYENHUIS, B.; PAZ, A.; PINO, J.; SIEGFRIED, P.; SUNDARAM, A.; TOM, D.; WERNLI, S. J.; ZANNER, M.; STUTZ, R. P.; SVORE, K. M. *Demonstration of logical qubits and repeated error correction with better-than-physical error rates*. 2024. Available at: <<https://arxiv.org/abs/2404.02280>>.

PARK, D. K.; PETRUCCIONE, F.; RHEE, J.-K. K. Circuit-based quantum random access memory for classical data. *Scientific Reports*, Nature Publishing Group UK London, v. 9, n. 1, p. 3949, 2019.

PERUZZO, A.; MCCLEAN, J.; SHADBOLT, P.; YUNG, M.-H.; ZHOU, X.-Q.; LOVE, P. J.; ASPURU-GUZI, A.; O'BRIEN, J. L. A variational eigenvalue solver on a photonic quantum processor. *Nature Communications*, Nature Publishing Group UK London, v. 5, n. 1, p. 4213, 2014.

- PIRANDOLA, S.; ANDERSEN, U. L.; BANCHI, L.; BERTA, M.; BUNANDAR, D.; COLBECK, R.; ENGLUND, D.; GEHRING, T.; LUPO, C.; OTTAVIANI, C. et al. Advances in quantum cryptography. *Advances in Optics and Photonics*, Optica Publishing Group, v. 12, n. 4, p. 1012–1236, 2020.
- PRESKILL, J. Fault-tolerant quantum computation. In: *Introduction to quantum computation and information*. [S.l.]: World Scientific, 1998. p. 213–269.
- PRESKILL, J. Quantum computing in the NISQ era and beyond. *Quantum*, Verein zur Förderung des Open Access Publizierens in den Quantenwissenschaften, v. 2, p. 79, Aug. 2018.
- RESCH, S.; KARPUZCU, U. R. Benchmarking quantum computers and the impact of quantum noise. *ACM Computing Surveys (CSUR)*, ACM New York, NY, USA, v. 54, n. 7, p. 1–35, 2021.
- ROSA, E. C.; DUZZIONI, E. I.; SANTIAGO, R. de. Optimizing gate decomposition for high-level quantum programming. *Quantum*, Verein zur Förderung des Open Access Publizierens in den Quantenwissenschaften, v. 9, p. 1659, 2025.
- SAEEDI, M.; PEDRAM, M. Linear-depth quantum circuits for n-qubit Toffoli gates with no ancilla. *Physical Review A*, APS, v. 87, n. 6, p. 062318, 2013.
- SAKI, A. A.; ALAM, M.; GHOSH, S. Impact of noise on the resilience and the security of quantum computing. In: IEEE. *2021 22nd International Symposium on Quality Electronic Design (ISQED)*. [S.l.], 2021. p. 186–191.
- SATO, Y.; KONDO, R.; HAMAMURA, I.; ONODERA, T.; YAMAMOTO, N. Hamiltonian simulation for hyperbolic partial differential equations by scalable quantum circuits. *Physical Review Research*, APS, v. 6, n. 3, p. 033246, 2024.
- SCHULD, M.; SINAYSKIY, I.; PETRUCCIONE, F. Prediction by linear regression on a quantum computer. *Physical Review A*, APS, v. 94, n. 2, p. 022342, 2016.
- SELINGER, P. Quantum circuits of t-depth one. *Physical Review A—Atomic, Molecular, and Optical Physics*, APS, v. 87, n. 4, p. 042302, 2013.
- SHENDE, V. V.; BULLOCK, S. S.; MARKOV, I. L. Synthesis of quantum logic circuits. In: *Proceedings of the 2005 Asia and South Pacific Design Automation Conference*. New York, NY, USA: Association for Computing Machinery, 2005. (ASP-DAC '05), p. 272–275.
- SHOR, P. W. Fault-tolerant quantum computation. In: IEEE. *Proceedings of 37th Conference on Foundations of Computer Science*. Burlington, VT, USA: IEEE, 1996. p. 56–65.
- SHOR, P. W. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM Review*, SIAM, v. 41, n. 2, p. 303–332, 1999.
- SILVA, A. J. da; PARK, D. K. Linear-depth quantum circuits for multiqubit controlled gates. *Physical Review A*, v. 106, p. 042602, 2022.
- SILVA, J. D.; AZEVEDO, T. M. D.; ARAUJO, I. F.; SILVA, A. J. da. Linear decomposition of approximate multi-controlled single qubit gates. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, IEEE, 2024.

- SIVARAJAH, S.; DILKES, S.; COWTAN, A.; SIMMONS, W.; EDGINGTON, A.; DUNCAN, R. $t|ket\rangle$: A retargetable compiler for NISQ devices. *Quantum Science and Technology*, IOP Publishing, v. 6, n. 1, p. 014003, 2020.
- SOUZA, L. S. de; CARVALHO, J. H. de; FERREIRA, T. A. Classical artificial neural network training using quantum walks as a search procedure. *IEEE Transactions on Computers*, IEEE, v. 71, n. 2, p. 378–389, 2021.
- STEIGER, D. S.; HÄNER, T.; TROYER, M. ProjectQ: An open source software framework for quantum computing. *Quantum*, Verein zur Förderung des Open Access Publizierens in den Quantenwissenschaften, v. 2, p. 49, 2018.
- TANNU, S. S.; QURESHI, M. K. Not all qubits are created equal: A case for variability-aware policies for NISQ-era quantum computers. In: *Proceedings of the 24th International Conference on Architectural Support for Programming Languages and Operating Systems*. New York, NY, USA: Association for Computing Machinery, 2019. (ASPLOS '19), p. 987–999.
- TOMESH, T.; ALLEN, N.; SALEEM, Z. Quantum-classical tradeoffs and multi-controlled quantum gate decompositions in variational algorithms. *arXiv preprint arXiv:2210.04378*, 2022.
- VALE, R.; AZEVEDO, T. M. D.; ARAÚJO, I. C. S.; ARAUJO, I. F.; SILVA, A. J. da. Circuit decomposition of multicontrolled special unitary single-qubit gates. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, v. 43, n. 3, p. 802–811, 2024.
- VARTIAINEN, J. J.; MÖTTÖNEN, M.; SALOMAA, M. M. Efficient decomposition of quantum gates. *Physical Review Letters*, APS, v. 92, n. 17, p. 177902, 2004.
- VERAS, T. M. D.; ARAUJO, I. C. D.; PARK, D. K.; SILVA, A. J. D. Circuit-based quantum random access memory for classical data with continuous amplitudes. *IEEE Transactions on Computers*, IEEE, v. 70, n. 12, p. 2125–2135, 2020.
- WOOD, C. J.; GAMBETTA, J. M. Quantification and characterization of leakage errors. *Physical Review A*, APS, v. 97, n. 3, p. 032306, 2018.
- YANOFSKY, N. S.; MANNUCCI, M. A. *Quantum computing for computer scientists*. [S.l.]: Cambridge University Press, 2008.
- ZINDORF, B.; BOSE, S. Efficient implementation of multi-controlled quantum gates. *arXiv preprint arXiv:2404.02279*, 2024.