

Depth-Optimal Synthesis of Clifford Circuits with SAT Solvers

Tom Peham*
Chair for Design Automation
Technical University of Munich
Germany
tom.peham@tum.de

Nina Brandl*
Institute for Integrated Circuits
Johannes Kepler University Linz
Austria
nina.brandl@jku.at

Richard Kueng
Institute for Integrated Circuits
Johannes Kepler University Linz
Austria
richard.kueng@jku.at

Robert Wille
Chair for Design Automation
Technical University of Munich
Germany
Software Competence Center Hagenberg GmbH
Austria
robert.wille@tum.de

Lukas Burgholzer
Institute for Integrated Circuits
Johannes Kepler University Linz
Austria
lukas.burgholzer@jku.at

*Both authors contributed equally to this work.

Abstract—Circuit synthesis is the task of decomposing a given logical functionality into a sequence of elementary gates. It is (depth-)optimal if it is impossible to achieve the desired functionality with even shorter circuits. Optimal synthesis is a central problem in both quantum and classical hardware design, but also plagued by complexity-theoretic obstacles. Motivated by fault-tolerant quantum computation, we consider the special case of synthesizing blocks of Clifford unitaries. Leveraging entangling input stimuli and the stabilizer formalism allows us to reduce the Clifford synthesis problem to a family of poly-size satisfiability (SAT) problems – one for each target circuit depth. On a conceptual level, our result showcases that the Clifford synthesis problem is contained in the first level of the polynomial hierarchy (NP), while the classical synthesis problem for logical circuits is known to be complete for the second level of the polynomial hierarchy (Σ_2^P). Based on this theoretical reduction, we formulate a SAT encoding for depth-optimal Clifford synthesis. We then employ SAT solvers to determine a satisfying assignment or to prove that no such assignment exists. From that, the shortest depth for which synthesis is still possible (optimality) as well as the actual circuit (synthesis) can be obtained. Empirical evaluations show that the optimal synthesis approach yields a substantial depth improvement for random Clifford circuits and Clifford+T circuits for Grover search.

I. INTRODUCTION

Quantum computing is a computational paradigm that might offer computational advantages over classical algorithms for certain problems. State-of-the-art quantum computing hardware is still limited in scale, featuring a relatively small number of qubits prone to errors. While finding short-term applications for these *noisy intermediate scale quantum* (NISQ) computers is an ongoing and vibrant research field [1], they cannot be used for advanced applications such as integer factoring (Shor; [2]), unstructured search (Grover; [3]), solving linear systems (HHL; [4]), or convex optimization [5]–[7].

To scale up quantum computing to longer computations with many qubits, *fault-tolerant* computation schemes have to be used that protect the information of a qubit against errors and allow the application of quantum gates with a low error-rate, see e.g. [8]–[10]. One way to do this is using quantum error correcting codes and performing all computations in the Clifford+T gate-set [11]. The benefit of restricting the gate-set is that every gate can be performed in a fault-tolerant fashion on an appropriate code.

Fault-tolerant quantum circuits can become quite large due to the overhead from breaking down every quantum computation into this restricted gate-set, with error syndrome extraction and error correction steps. For near-term applications they are usually optimized to have a minimal two-qubit gate count as these gates tend to have the highest error rates. For fault-tolerant computations, a suitable performance metric is circuit depth because it directly correlates with the runtime of the computation.

The problem of synthesizing optimal circuits in the fault-tolerant regime has commonly been considered holistically, trying to minimize the entire circuit, or partially, by reducing the T-gate count only [13]–[16]. The general synthesis problem turns out to be very hard. Instead of minimizing the entire Clifford+T circuit, one can consider only the Clifford parts to try to make the general synthesis problem easier.

In this work, we show that the problem of synthesizing depth-optimal Clifford circuits is, in fact, at most as hard as Boolean satisfiability (SAT). We achieve this by considering the synthesis problem for a larger circuit that receives a pairwise maximally entangled state as input. Although this theoretically increases the problem size, the utilization of the maximally entangled state allows for breaking down the synthesis problem in terms of stabilizers (Gottesman-Knill).

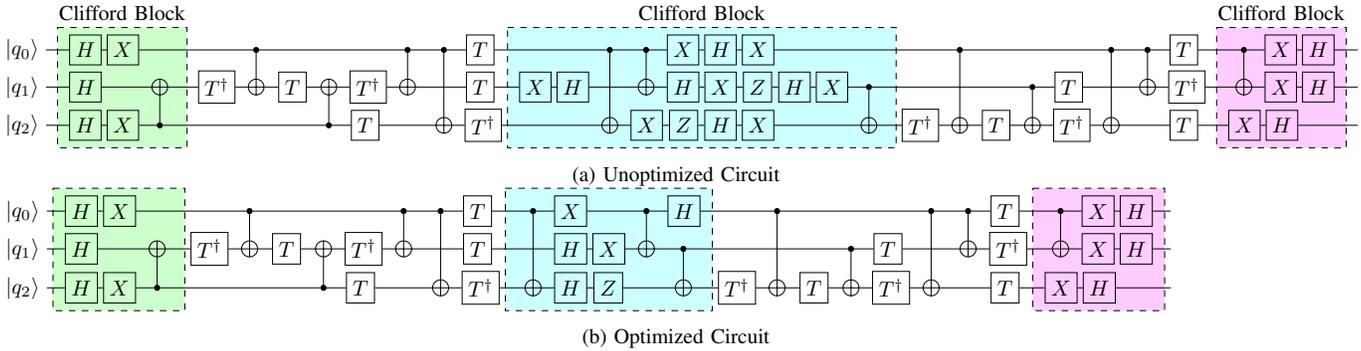


Fig. 1. *Two Clifford+T circuits for 3-qubit Grover search*: the oracle corresponds to a random 3-SAT formula with 3 variables and 5 clauses. **(top)** decomposition of one Grover block into Clifford+T gates, the actual synthesis is achieved using Qiskit’s `Grover` class using a `PhaseOracle` [12]. **(bottom)** depth-optimized synthesis of all (nontrivial) Clifford blocks using the methods from this work. Our SAT approach certifies that the green and purple blocks are already depth-optimal. This is not the case for the central teal block, where our method yields considerable improvements (depth 9 vs. depth 5).

Based on this perspective, we formulate a SAT encoding for synthesizing and optimizing n -qubit Clifford circuits of maximal depth d_{\max} with $O(n^2 d_{\max})$ variables and $O(n^4 d_{\max})$ constraints. For larger circuits, we furthermore develop heuristic optimization routines based on optimal SAT encodings that reduce Clifford circuit depth in a divide-and-conquer approach.

We implemented the proposed methods and compared them to a state-of-the-art and openly available Clifford synthesis technique. Results show that while the optimal synthesis scales only up to 5 qubits, the state of the art is, on average, two orders of magnitude away from the optimal depth. The heuristic approaches also yield better results than the state of the art and are much more scaleable than the optimal approach. The efficacy of the methods for fault-tolerant quantum computations is illustrated by showing how the depth of Clifford+T circuits can be reduced by optimizing Clifford sub-circuits. All implementations are publicly available via the quantum circuit compilation tool QMAP [17], which is part of the *Munich Quantum Toolkit* (MQT) and accessible at <https://github.com/cda-tum/qmap>.

The remainder of this work is structured as follows: **Sec. II** motivates the need for synthesizing depth-optimal Clifford circuits and provides context to the classical circuit synthesis problem. Then, **Sec. III** briefly goes over previous work on quantum circuit synthesis, especially in the Clifford case. In **Sec. IV** we show the main construction of this work and the reduction of Clifford synthesis to SAT before giving the full details on the SAT encoding in **Sec. V**. Based on that, we derive a heuristic optimization routine in **Sec. VI**. The evaluations of the methods introduced in this work are presented in **Sec. VII**. Finally, **Sec. VIII** concludes this work.

II. BACKGROUND AND MOTIVATION

To keep this work self-contained, the following sections review the central concepts of fault-tolerant quantum computing and circuit synthesis that are required throughout the rest of this work.

A. Fault Tolerant Quantum Computing

Quantum computations are prone to errors from multiple sources, be it decoherence, information leakage, or other kinds of noise. In order to protect quantum information, *quantum error correcting codes* (QECCs) [8], [11], [18] have been introduced. Essentially the information of a single *logical qubit* is encoded into multiple *physical qubits* for which various encoding schemes have been discovered.

Executing a quantum gate on a logical qubit is not a straightforward matter anymore as it is not at all obvious how to extend the functionality of a physical gate to the logical level. Furthermore, applying a gate to a logical qubit should be possible in a *fault-tolerant* fashion, such that applying a gate should introduce errors that can be corrected with an arbitrarily high probability.

The *Clifford gate-set* is a set of gates that can be applied *transversally* to many QECCs, i.e., applying a Clifford operation to a logical qubit can be done by performing that Clifford gate on every physical qubit individually. Therefore, if a gate application introduces an error onto a physical qubit, this error cannot spread throughout the encoded qubit and is simpler to correct. Any Clifford unitary can be obtained from a sequence of *Hadamard*, *Phase*, and *CNOT* gates. The Clifford group also includes the Pauli X -, Y - and Z -gates.

Clifford gates alone are not sufficient to perform universal quantum computation and can be classically simulated in polynomial time [11], [19]. However, adding a T-gate $T = |0\rangle\langle 0| + e^{i\frac{\pi}{8}}|1\rangle\langle 1|$ suffices to achieve universality. A T-gate can be realized fault-tolerantly using the magic state distillation protocol [20] and gate teleportation.

A fault-tolerant quantum computation can therefore be realized by alternating blocks of Clifford gates and blocks of T-gates where error syndromes are detected and corrected throughout the computation. Fig. 1 shows the quantum circuit for a 3-qubit Grover search that is implemented using the Clifford+T gate-set on the logical qubits $|q_0\rangle$, $|q_1\rangle$ and $|q_2\rangle$.

In this fault-tolerant regime, the most important performance metric is the *circuit depth*. Assuming gates on different qubits can be executed in parallel, the depth directly deter-

mines the runtime of the computation. As the *cycle time* can vary considerably depending on the technology the quantum computer is built upon, reducing the depth of a circuit can reduce the cost of a computation drastically.

Example 1. *The circuit in Fig. 1a has 3 Clifford blocks of interest. Fig. 1b shows the same circuit where each block was synthesized to have minimal depth. As one can see, the two outermost blocks were already depth-optimal in the original circuit. Although this looks intuitively true, it still needs some form of proof. The Clifford block in the middle, however, has a depth of 9 (the first Hadamard and CNOT-gate are parallel) and can be optimized to depth 5. High-depth Clifford blocks like this appear frequently when synthesizing circuits using state-of-the-art synthesis tools for quantum computing¹.*

The quest of finding such depth-optimal realizations of Clifford circuits will be our primary motivation in this work as the restricted nature of these circuits lends itself nicely to classical design automation methods such as *SAT* and *SMT solvers*.

B. From Classical to Quantum Synthesis

Before diving into the details of our quantum protocol, it is worthwhile to briefly review classical equivalence checking of logical circuits (or functions), as well as classical circuit synthesis. This will provide guidance and also motivate our use of SAT solvers for quantum circuit synthesis.

1) *Classical Equivalence Checking and SAT:* Let C, C' be two circuits with n input bits. We say that these circuits are equivalent ($C \simeq C'$) if and only if they produce the same output for each conceivable input. In formulas, $\forall x \in \{0, 1\}^n Cx = C'x$. Taking the contrapositive yields

$$C \not\simeq C' \Leftrightarrow \exists x \in \{0, 1\}^n \text{ s.t. } Cx \neq C'x.$$

The logical not-equality $Cx \neq C'x$ can readily be converted into a Boolean function $\phi_{C,C'}(x)$ with input x . This highlights an interesting one-to-one correspondence between (the negation of) circuit equivalence and the *satisfiability problem* (SAT). On a theoretical level, SAT is a hard problem that is complete for the problem class NP. Nevertheless, a plethora of heuristic SAT solvers [21] perform very well in practice. Correspondingly, how to utilize SAT for quantum computing also received attention recently (cf. [22]).

2) *Classical Optimal Synthesis and QBF:* Circuit synthesis aims to find a logical circuit that implements a desired n -bit target functionality with as few elementary gates as possible. Here, we will focus on *circuit depth* (i.e., layers of gates). A decision version of this optimization problem looks as follows:

$$\exists C_d, \text{depth}(C_d) \leq d_{\max} \forall x \in \{0, 1\}^n C_d x = Cx, \quad (1)$$

where C_d is a placeholder for another logical circuit. In words, this formula evaluates to true if and only if it is possible to exactly reproduce the functionality of circuit C with another logical circuit C_d that obeys $\text{depth}(C_d) \leq d_{\max}$. Multiple queries to this logical function with different values

of d_{\max} allow us to determine the optimal depth of *any* circuit synthesis, e.g. via binary search.

It is also possible to rephrase Eq. (1) as a *quantified Boolean formula* (QBF). For starters, note that we can represent any logical circuit C_d by a binary encoding y of length (at most) $\text{poly}(nd_{\max})$. Different bit strings y give rise to different circuits C_d and vice versa. With this binary encoding at hand, we can adapt the Boolean function reformulation of equivalence checking to the case at hand: $\phi_C(y, x) = 1$ if $Cx = C_d x$ and C_d is the circuit encoded by y . Otherwise, this formula evaluates to 0. Putting everything together, we obtain the following QBF reformulation of logical circuit synthesis:

$$\exists y \in \{0, 1\}^{\text{poly}(nd_{\max})} \forall x \in \{0, 1\}^n \phi_C(y, x) \stackrel{!}{=} 1. \quad (2)$$

The exist quantifier (\exists) ranges over all possible logical circuit encodings with depth at most d_{\max} ($C_d \leftrightarrow y$) while the forall quantifier (\forall) ranges over all 2^n possible input bitstrings.

Such QBFs are, in general, much harder to handle than a mere SAT problem with only one type of quantifier. In fact, the reformulated circuit synthesis problem Eq. (2) is complete for the problem class Σ_2^P – one branch of the second level of the polynomial hierarchy [23]. Unless the polynomial hierarchy collapses to the first level (which is widely believed to be false), these problems are much harder than SAT and, by extension, equivalence checking.

QBFs do arise naturally in a variety of contexts [24]. Solvers do exist, see e.g. [25]–[27] and typically rely on the *counter-example-guided inductive synthesis principle* (CEGIS) [28] which has its roots in abstraction refinement (CEGAR) [29], [30]. For program synthesis, for example, CEGIS-style solvers alternate between generating candidate programs and checking them for counter-examples.

3) Going Quantum: Circuit Equivalence and Synthesis:

The two classical challenges we just discussed have natural counterparts in the quantum realm. Quantum circuits that act on n qubits are also comprised of elementary (quantum) gates, but their functionality is radically different. We say that two such circuits U, V are *equivalent* if and only if

$$U \simeq V \Leftrightarrow \forall |\psi\rangle \in \mathbb{C}^{2^n} U|\psi\rangle = e^{i\phi} V|\psi\rangle,$$

where $\phi \in [0, 2\pi)$ is a complex phase. In contrast to (classical) logical circuits, there are infinitely many possible input states $|\psi\rangle$ to be checked.² Consistency and mutual interrelations permit us to compress this number to (at most) 4^n disjoint input states [32]–[36]. This number is, however, still exponential in the total number of qubits. It is known that the (negated) unitary equivalence problem is QMA-complete [37] (QMA is the appropriate quantum generalization of the classical problem class NP). So, at face value, the quantum circuit equivalence problem looks even harder than its classical counterpart.

Suppose that we are given a target unitary U , e.g., in the form of a high-level quantum circuit, and we want to

²This is already our point of departure from earlier work, most notably [31]. There, a subset of the authors asked a related, but simpler, question that arises from only considering $|\psi\rangle = |0, \dots, 0\rangle$ (i.e. fix a single input state).

¹The circuit in question was generated using Qiskit 0.42.1

decompose it into as few gate layers as possible (i.e., circuit depth). This practical problem occurs whenever we want to execute a high-level unitary (e.g., a quantum algorithm) on an actual n -qubit quantum computer. The details of this synthesis problem depend on the type of elementary gate-set but always produces a two-fold quantified problem

$$\exists U_d, \text{depth}(U_d) \leq d_{\max} \forall |\psi\rangle U_d |\psi\rangle = e^{i\phi} U |\psi\rangle,$$

that resembles Eq. (1) with non-binary quantum states and a complex phase $\phi \in [0, 2\pi)$. This problem is at least as hard as logical gate synthesis because it includes (reversible) encodings of the latter one as a special case. Reversibility of quantum circuits allows us to slightly streamline this display:

$$\exists U_d, \text{depth}(U_d) \leq d_{\max} \forall |\psi\rangle U U_d^\dagger |\psi\rangle = e^{i\phi} |\psi\rangle, \quad (3)$$

where U_d^\dagger is the adjoint or reverse circuit of U_d .

For some special cases like Clifford circuits, we can recast Eq. (3) as a logical Boolean formula whose intrinsic difficulty is on par with logical circuit equivalence or SAT as shown in Sec. IV. The difficulty level is then *much* lower than general classical circuit synthesis. Empowered by this rigorous theoretical insight, we then employ state-of-the-art SAT solvers to address the full Clifford synthesis problem.

III. RELATED WORK

Quantum circuits comprised of universal gate-sets are universal in the sense that they can approximate *every* unitary evolution. Quantum gate synthesis can be viewed as a quantitative take on this very issue: what is the best way to implement a given unitary evolution, e.g., a quantum computation?

The celebrated Solovay-Kitaev theorem [8], [38] can be viewed as a very general synthesis protocol for arbitrary single-qubit unitaries ($n = 1$) and arbitrary universal gate-sets. Extensions to n qubits can be achieved by either combining single-qubit gate synthesis with certain entangling multi-qubit gates or by directly generalizing the Solovay-Kitaev algorithm to $d = 2^n$ -dimensional unitaries [38]. Although implementations do exist, see e.g. [39], this synthesis algorithm is typically too slow for practical purposes. As a result, the community has moved away from this rigorous meta-algorithm and towards more scalable heuristics. Many of them address the universal gate-set comprised of Clifford+T gates, see e.g. [13]–[16].

Clifford circuits (without T gates) are an interesting special case in this context. Mathematically speaking, they form a representation of a finite symplectic group [40]–[42] with additional structure. For instance, it is known that every Clifford unitary can be decomposed into a Clifford circuit of depth at most $O(n)$ [43]. Such insights highlight that Clifford circuits cannot be overly complex – a feature that also extends to Clifford synthesis. For the special case of $n = 6$ qubits, competitive Clifford synthesis protocols have been put forth in [44]. For general n , the algorithms by Koenig and Smolin can be used to associate a given Clifford circuit with exactly one element of the Clifford group. And subsequently, their algorithm can be used to synthesize this very group element. More recently, the group of Robert Calderbank developed

Clifford synthesis algorithms that even work on the logical level (i.e., on top of an error-correcting stabilizer code) [45], while Bravyi *et al.* discovered constant depth representations of arbitrary Clifford circuits, under the assumption that one is allowed to use global entangling operations of Ising type [46].

While the aforementioned approaches do produce a provably correct decomposition of Clifford circuits into elementary (Clifford) gates, it is not clear whether size and depth are (close to) optimal. This is where reformulations in terms of SAT/QBF can make a significant difference. They reformulate the gate synthesis problem as a family of quantified Boolean formulas (QBFs), one for each maximum circuit depth d_{\max} we allow. Such a QBF evaluates to true if and only if an exact circuit representation with depth (at most) d_{\max} is possible and returns the explicit representation. Otherwise, it evaluates to false. Attempting to solve these QBFs for different depths bears the potential of identifying the best circuit representation of a given functionality. This is why SAT/QBF-based synthesis approaches have long been a mainstay in classical design automation [47]–[52]. In fact, the idea of combining SAT-based synthesis with Clifford circuits is not entirely new. In Ref. [31], a subset of authors proposes this very idea for optimal stabilizer state preparation: find the shortest Clifford circuit that takes $|0, \dots, 0\rangle$ as input and produces a known target stabilizer state. The ideas presented here may be viewed as an extension of these earlier ideas to full Clifford circuit synthesis. In addition, we also supply rigorous proofs of correctness and provide additional context, as well as background.

IV. MAIN RESULT AND THEORETICAL UNDERPINNING

Our main conceptual result is a one-to-one correspondence between Clifford synthesis, on the one hand, and Boolean satisfiability (SAT), on the other. The main result is displayed in Theorem 1 and originates from two genuinely quantum twists to the original synthesis questions: (i) maximally entangled input stimuli and (ii) the Gottesman-Knill theorem. It forms a rigorous foundation for optimal Clifford circuit synthesis with SAT, the topic of Sec. V below.

A. Quantum Twist 1: Maximally Entangled Input States

The first step in our theoretical argument goes by many names, including the Choi-Jamiokowski isomorphism [53], [54], entanglement-assisted process tomography [55] and the flattening operation in tensor analysis [56], [57]. Conceptually we take two quantum circuits U and V on n -qubits. Instead of testing all possible input states, we create a single universally valid input state to check their equivalence. In a first thought process, we consider a test circuit twice as large and apply U to the top n qubits and the inverse of each gate in V to the remaining n qubits. By entangling the k th input qubit of U with the k th input qubit of V we create a new state $|\omega_{2n}\rangle$ of size $2n$. As U and V operate on entangled qubits, all changes applied by the first circuit will be reverted by the second circuit if and only if they have the same functionality up to a global phase. If they instead differ at any point of their unitary, the resulting state will not be equal to $|\omega_{2n}\rangle$ again. The idea of

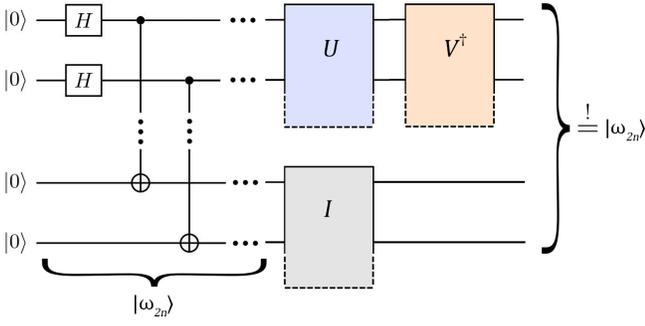


Fig. 2. Illustration of entanglement-assisted equivalence-checking: Two n -qubit circuits U, V have equivalent functionality (up to a global phase) if and only if the above circuit produces the pairwise maximally entangled state $|\omega_{2n}\rangle$. Here, V^\dagger is the reverse circuit (adjoint) of V , and I is the identity.

two circuits with identical unitaries reverting the changes of each other still applies when we let U and V^\dagger both work on the same n qubits, i.e., the control qubits of the pairwise entanglement. On the remaining target qubits, we apply the identity as shown in Fig. 2.

Lemma 1. Let U, V be two n -qubit quantum circuits and let $I^{\otimes n}$ be the n -qubit identity operation. Then, $U \simeq V$ if and only if

$$(UV^\dagger) \otimes I^{\otimes n} |\omega_{2n}\rangle \langle \omega_{2n}| (UV^\dagger)^\dagger \otimes I^{\otimes n} = |\omega_{2n}\rangle \langle \omega_{2n}|. \quad (4)$$

here $|\omega_{2n}\rangle$ is the tensor product of n 2-qubit Bell states that entangle the k th qubit with the $n+k$ th qubit for all $k \in [n]$.

Proof. Note that the Bell state is proportional to the vectorized identity matrix, i.e. $|\omega_{2n}\rangle \sqrt{2^n} = \text{vec}(I^{\otimes n})$. The operator-vector correspondence, see e.g. [56, Eq. (1.132)], then asserts $UV^\dagger \otimes I_n |\omega_{2n}\rangle = \text{vec}(UV^\dagger) / \sqrt{2^n}$. This ensures that $UV^\dagger = e^{i\phi} I_n$ (as a matrix-valued equality) if and only if

$$UV^\dagger \otimes I |\omega_{2n}\rangle = \frac{1}{\sqrt{2^n}} \text{vec}(UV^\dagger) \stackrel{!}{=} \frac{e^{i\phi}}{\sqrt{2^n}} \text{vec}(I^{\otimes n}) = e^{i\phi} |\omega_{2n}\rangle.$$

Eq. (4) takes the outer product of this vector-valued equation to absorb the global phase (mixed state formalism). \square

Applying Lemma 1 to the quantum circuit synthesis formula Eq. (3) gets rid of the complex phase and the forall quantifier over all possible input states:

$$\exists U_d, \text{depth}(U_d) \leq d_{\max} (UU_d^\dagger) \otimes I \Omega (UU_d^\dagger)^\dagger \otimes I = \Omega, \quad (5)$$

where we have used Ω as a shorthand notation for $|\omega_{2n}\rangle \langle \omega_{2n}|$. However, for the erasure of an entire forall quantifier, we have to go to the mixed state formalism and effectively double the number of qubits involved from n to $2n$, see Fig. 2.

B. Quantum Twist 2: Gottesman-Knill Theorem

On first sight, Eq. (5) looks much less daunting than Eq. (3) or even its classical counterpart Eq. (1). This is due to the fact that entanglement allows us to check for equivalence with a single input state (Lemma 1) instead of a forall (\forall) over exponentially many possibilities.

To exploit this reformulation, there are two broad avenues on how to proceed: (i) Use an actual quantum computer to

empirically check the single remaining equivalence in Eq. (5). This avenue was taken, for instance, in quantum assisted quantum compiling [58]. (ii) Use (strong) classical simulation of quantum circuits to check whether the two $2n$ -qubit states in Eq. (5) are really equivalent. Four broad simulation approaches come to mind: array-based [59], stabilizer-based [19], tensor networks [60] and decision diagrams [61], [62].

Here, we follow the second avenue and adopt stabilizer-based simulation. The main reason for this is that $\Omega = |\omega_{2n}\rangle \langle \omega_{2n}|$, the $2n$ -qubit state responsible for all these simplifications, is itself a stabilizer state with very desirable structure.

Fact 1. The $2n$ -qubit state $\Omega = |\omega_{2n}\rangle \langle \omega_{2n}|$ is a stabilizer state with generators $\langle (X I^{\otimes(n-1)})^{\otimes 2}, (Z I^{\otimes(n-1)})^{\otimes 2} \rangle$ where the parenthesis $()$ stand for cyclic permutation.

It is well-known that the 2-qubit Bell state is a stabilizer state with generators $\langle XX, ZZ \rangle$ [10]. Fact 1 follows from taking the n -fold tensor product of these generators at appropriate qubit locations. For $n=2$, for example, this means $\langle XIXI, ZIZI, IXIX, IZIZ \rangle$ or for $n=3$ $\langle XIIIXI, ZIIIZI, IXIIXI, IZIIZI, IIXIIX, IIZIIZ \rangle$, etc. Still, every k th qubit is entangled with the $(n+k)$ th qubit.

We can efficiently simulate stabilizer circuits (of which Clifford circuits are a part) in polynomial time on a classical computer according to the Gottesman-Knill theorem [63]. To simulate the circuit's action on the $2n$ -qubit stabilizer state Ω , we perform a logical mapping of stabilizers depending on the gates and again obtain a stabilizer state. We can conclude whether the applied circuits performed the identity based on the equality of the input and output generators.

Fact 2 (Gottesman-Knill). Suppose that both U and U_d are n -qubit Clifford circuits. Then, it is possible to efficiently check

$$(UU_d^\dagger) \otimes I \Omega (UU_d^\dagger)^\dagger \otimes I = \Omega. \quad (6)$$

This also follows from the conceptual idea of the entangled input: If U and U_d have the same functionality, U_d^\dagger would invert the stabilizer mapping done by U , which results in not altering the input stabilizers at all. Note that input and output generators must match exactly. As the bottom half of this $2n$ circuit applies the identity, it will not alter the generators on the last n qubits. Therefore we can also cut the generators we need to check in half, eliminating the \otimes^2 operation in Fact 1.

More precisely, it is possible to construct a Boolean function $\phi_{U,\Omega}(U_d)$ that evaluates to 1 if the input circuit U_d achieves Eq. (6) and 0 otherwise. We will present such an explicit construction in Sec. V below. All that matters at this point is that we can represent any (at most) depth- d_{\max} Clifford unitary U_d with a bitstring $y \in \{0,1\}^t$ that contains (at most) $l(n, d_{\max}) = O(n^2 d_{\max})$ Boolean variables. And, what is more, we can actually represent $\phi_{U,\Omega}(y)$ as a CNF with $O(n^4 d_{\max})$ clauses of constant length. Putting all this together ensures that we can rewrite the Clifford circuit synthesis problem as

$$\exists y \in \{0,1\}^{l(n,d_{\max})} \phi_{U,\Omega}(y) \stackrel{!}{=} 1. \quad (7)$$

C. Main Result and Synopsis

The insights culminating in Eq. (7) are worth a prominent display and a bit of additional context.

Theorem 1 (SAT reformulation of Clifford synthesis). *Let U be a n -qubit Clifford unitary (target) and fix a maximum depth $d_{\max} \in \mathbb{N}$. Then, the decision problem “is it possible to exactly reproduce U with (at most) d_{\max} Clifford layers?” can be rephrased as an instance of SAT with $O(n^2 d_{\max})$ variables and $O(n^4 d_{\max})$ clauses of constant size each.*

This insight has both conceptual and practical implications, especially if one keeps in mind that $d_{\max} \leq O(n)$ for any Clifford circuit [43]. In turn, *binary search* allows for exactly determining the minimum circuit depth for a given Clifford unitary U by solving (at most) $\lceil \log_2(n) \rceil + O(1)$ SAT reformulations for varying ansatz depths d_{\max} . What is more, satisfying assignments of the Boolean formula in Eq. (7) are bit encodings of actual Clifford circuits that exactly reproduce U and have depth at most d_{\max} (synthesis).

On a conceptual level, this approach provides a poly-time reduction of optimal Clifford synthesis to (logarithmically many instances of) SAT. This highlights that this special case is *much* easier than the general circuit synthesis problem (classical and quantum). In particular: optimal Clifford synthesis is at most as hard as classical circuit equivalence checking.

On a practical level, Theorem 1 provides a rigorous and context-specific motivation for employing state-of-the-art SAT solvers to rigorously address the Clifford synthesis problem. An actual step-by-step introduction to this encoding, as well as benchmarks, are the content of the remainder of this article.

V. OPTIMAL CLIFFORD CIRCUIT SYNTHESIS WITH SAT

The efficient simulability of Clifford circuits is based on the *stabilizer tableau* encoding of a stabilizer state, see, e.g., [10] and references therein. This polynomial-sized representation of a quantum state is the key to deriving a polynomially-sized SAT encoding for the considered synthesis problem. Hence, before going into details on the encoding itself, we will give a brief recap on how to work with stabilizer tableaux.

A. Stabilizer Tableau Representation of Stabilizer States

An n -qubit stabilizer state can be represented by a $(2n + 1) \times n$ binary matrix called the stabilizer tableau. The idea is that every stabilizer generator for a state can be written using $2n + 1$ bits of information. In the standard notation [19] for stabilizer tableaux, there are binary variables for Pauli X - and Z -type stabilizers $x_{i,j}, z_{i,j}$ with $i, j \in \{0, 1, \dots, n-1\}$ and $r_i = 1$ if the i -th stabilizer has a negative phase:

$$\left[\begin{array}{ccc|ccc|c} x_{0,0} & \dots & x_{0,n-1} & z_{0,0} & \dots & z_{0,n-1} & r_0 \\ \vdots & \ddots & \vdots & \vdots & \ddots & \vdots & \vdots \\ x_{n-1,0} & \dots & x_{n-1,n-1} & z_{n-1,0} & \dots & z_{n-1,n-1} & r_{n-1} \end{array} \right]$$

For a Pauli- Y type stabilizer at position i, j both $x_{i,j}$ and $z_{i,j}$ must be set to 1 at the corresponding position. The stabilizer states in this format can be altered by the usual Clifford gates

and the following update rules, where \oplus denotes a bitwise XOR operation:

- *Applying H on qubit j* : swaps the j -th X -type column with the j -th Z -type column and $r \oplus = x_j z_j$. This follows from the transformations $HXH^\dagger = Z$, $HZH^\dagger = X$ and $HYH^\dagger = -Y$, hence it switches $x_{i,j} \leftrightarrow z_{i,j} \forall i$ and flips the phase only in case of a Pauli- Y stabilizer;
- *Applying S on qubit j* : is a bitwise XOR of the j -th X -type column to the j -th Z -type column $z_j \oplus = x_j$ again with $r \oplus = x_j z_j$;
- *Applying $CNOT$ on control qubit c and target qubit t* : is a bitwise XOR of the c -th X -type column to the t -th X -type column $x_t \oplus = x_c$ and vice versa for Z -type $z_c \oplus = z_t$ as well as $r \oplus = x_c z_t (x_t \oplus z_c \oplus 1)$.

Further update rules for any other Clifford gate can be derived from these basic rules.

More information about a stabilizer state can be encoded into the tableau by including its *destabilizers* [19]. These are Pauli-strings that, together with the stabilizers, generate the entire Pauli group. They are treated identically to the stabilizer generators for the purpose of updating the stabilizer tableau.

B. Tableau and Gate Variables

In the following, let Q be the set of qubits acted on by a quantum circuit and d_{\max} the maximal depth of the circuit.

While all Clifford unitaries can be obtained from just H , S , and $CNOT$ gates, conveniently, the target gate-set used for compilation may also include other gates like the Pauli X , Y , and Z operations or two-qubit gates like the CZ gate. To reflect this flexibility in the encoding, we define two sets $SQGs$ and $TQGs$, the set of single-qubit gates and two-qubit gates, respectively, such that they can be used to implement any Clifford circuit.

At every layer of the quantum circuit, a certain gate can either be applied or not. This suggests introducing the variables

$$\begin{aligned} Svars &= \{g_q^d \mid g \in SQGs, q \in Q, 0 \leq d < d_{\max}\} \\ Tvars &= \{g_{q_0, q_1}^d \mid g \in TQGs, q_0 \in Q, q_1 \in Q \setminus \{q_0\}, \\ &\quad 0 \leq d < d_{\max}\} \end{aligned}$$

representing the application of a gate to a specific qubit (or pair of qubits) at depth d .

The possible stabilizer tableaux are encoded in a straightforward fashion according to their definition. The X -, Z - and R -part of the tableau use the variables

$$\begin{aligned} Xvars &= \{x_q^d \mid q \in Q, 0 \leq d < d_{\max}\}, \\ Zvars &= \{z_q^d \mid q \in Q, 0 \leq d < d_{\max}\}, \\ Rvars &= \{r^d \mid 0 \leq d < d_{\max}\}, \end{aligned}$$

where every element of the sets $Zvars$, $Xvars$ and $Rvars$ is a bitvector. These bitvectors encode how all stabilizers act on the Z -, X -, R - part for a particular qubit.

Based on the construction in Lemma 1, this encoding requires $2n$ qubits in order to guarantee that all circuits synthesized from these variables have the same unitary. But

having these n additional qubits has the undesirable side-effect that the synthesized circuit should act as the identity on the lower n qubits of the circuit. This unnecessarily blows up the search space as the identity can be implemented ambiguously. One could enforce constraints on these qubits, but this would unnecessarily increase the size of the encoding. We can avoid this complication by considering only the upper n qubits and switching from stabilizers of the entangled input state to stabilizers and destabilizers of the $|0\rangle^{\otimes n}$ state as stated by the following fact.

Fact 3. *For a Clifford unitary U the stabilizers of $(U \otimes I^{\otimes n})|\omega_{2n}\rangle$ on the first n qubits are identical to the stabilizers and destabilizers of $U|0\rangle^{\otimes n}$.*

Together with **Fact 1**, **Lemma 1** tells us that for a Clifford circuit U the $2n$ stabilizers of $U \otimes I_n|\omega_{2n}\rangle$ uniquely fix the unitary of the circuit. Given U , we can explicitly calculate these stabilizers by propagating the generators for Ω through $U \otimes I$ ignoring the lower n qubits. This boils down to only analyzing the first half of every stabilizer generator of the $2n$ -qubit state Ω . The result is a $2n(n+1)$ tableau for every given U . The initial tableau has diagonal entries with value 1 and coincides with the stabilizers of the $|0\rangle^{\otimes n}$ input state (Z -type) combined with the respective destabilizers (X -type).

Hence, we can encode our problems using only $|Q| = n$ qubits, and each bitvector for the tableau variables has size $2n$ since the information about the destabilizers has to be included as well.

C. Transition Relation

With the variables defined, we can now encode how gates act on the stabilizer tableaus as described in **Sub. V-A**. Naturally, the transition between tableaus would then be a constraint along the lines of

$$g_q^d \Rightarrow (\text{UpdateZ}(g, q, d) \wedge \text{UpdateX}(g, q, d) \wedge \text{UpdateR}(g, q, d)),$$

where the update formulas on the right encode the action of the gate on a qubit at a certain depth. For a Hadamard this would mean

$$\begin{aligned} \text{UpdateZ}(H, q, d) &= (z_q^{d+1} \Leftrightarrow x_q^d) \\ \text{UpdateX}(H, q, d) &= (x_q^{d+1} \Leftrightarrow z_q^d) \\ \text{UpdateR}(H, q, d) &= (r^{d+1} \Leftrightarrow (r^d \oplus x_q^d \wedge z_q^d)) \end{aligned}$$

While this encoding is correct, it is also wasteful in the sense that it would lead to $|SQGs| + |TQGs|$ number of implications of this type for every qubit and depth. This number can be decreased significantly by noting that many gates act identically on the different parts of the stabilizer tableau (quantum computation is local). The Pauli gates, for example, act as identity on the Z - and X -part of the tableau, only differing in how they change the R -part.

Since we know our gate-set, we can collect all possible transformations of the individual parts of the tableau a priori. Let $Z\text{-updates}(q, d)$ be the set of all possible updates to the Z -part of the stabilizer tableau on qubit q at depth d . The

elements of these sets are logical formulas over the tableau variables. We can then define a mapping $Z\text{-impliedby}(q, d) : Z\text{-updates}(q, d) \rightarrow \mathcal{P}(Svars)$ that maps every update formula to the set of single-qubit gate variables that act on the stabilizer tableau with that update rule.

The single-qubit changes to the Z -part are then encoded by introducing the following constraint for every qubit q , depth $0 \leq d < d_{max} - 1$ and $Z\text{-update} \in Z\text{-updates}(q, d)$:

$$\bigvee_{g_q^d \in Z\text{-impliedby}(Z\text{-update})} \implies (z_q^{d+1} \Leftrightarrow Z\text{-update}).$$

Obviously, this can be done in a similar fashion for the X - and R -parts of the tableau as well as for two-qubit gates.

While the constraints at this point encode all possible stabilizer tableaus, there are some variable assignments that lead to invalid circuits, e.g., when a qubit is acted on by two gates at the same depth. We, therefore, need to introduce another set of constraints for every depth d and qubit q to ensure consistency of the obtained solution.

$$\text{ExactlyOne} \left(\begin{array}{l} \{g_q^d \mid g \in SQGs\} \\ \{g_{q, q_1}^d \mid q_1 \in Q, g \in TQGs\} \\ \{g_{q_0, q}^d \mid q_0 \in Q, g \in TQGs\} \end{array} \cup \right)$$

D. Symmetry Breaking

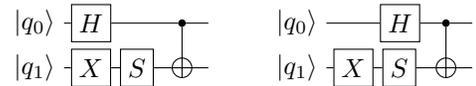
Symmetry breaking [64] is a widespread technique from the SAT-solving community. It introduces additional constraints to an existing CNF formula to avoid searching in symmetric parts of the search space. This can be done in an automated fashion by analyzing the formula for automorphisms to obtain, so-called ‘‘symmetry breakers’’. Doing this automatically has the downside that it is not clear which symmetries are found and whether the deduced constraints actually make the SAT instance any easier to solve or even harder. In the case of the SAT formulation above, we can obtain symmetry breakers manually by using knowledge specific to Clifford synthesis.

We can impose additional constraints on the SAT solver by eliminating valid solutions that could be expressed in a simpler manner. For example, the Hadamard gate is self-inverse. We can therefore add the constraint

$$H_q^d \implies \neg H_q^{d+1}$$

for $q \in Q$ and $0 \leq d < d_{max} - 1$. This eliminates all assignments to the gate variables that model a sequence of two consecutive Hadamard gates.

Another symmetry addresses possible degrees of freedom in the gate ordering and is best illustrated by means of the following equivalent circuits:



The Hadamard on the first qubit can either be parallel to the X gate or the S gate on the second qubit. We can break this symmetry by imposing that the identity single-qubit gate cannot be followed by a non-identity single-qubit gate. Otherwise, the non-identity gate could be moved to the left

without changing the Clifford unitary. More formally, for $q \in Q$, $0 \leq d < d_{\max} - 1$ we impose

$$I_q^d \implies \bigwedge_{g \in SQGs \setminus \{I\}} \neg g_q^{d+1},$$

where I is the identity gate.

A similar constraint can be imposed on two-qubit gates. If the identity is applied to a pair of qubits, no two-qubit gate can come after the identities. Again, we add the constraint

$$(I_{q_0}^d \wedge I_{q_1}^d) \implies \bigwedge_{g \in TQGs} \neg g_{q_0, q_1}^{d+1}.$$

for $q_0 \in Q$, $q_1 \in Q \setminus \{q_0\}$, $0 \leq d < d_{\max} - 1$.

There are many more symmetries that can be broken in the encoding. In principle, any Clifford gate identity can be used to derive a symmetry-breaking constraint. However, a trade-off has to be made between the number of constraints and the size of the solution space.

E. Optimizing Circuit Depth

The above encoding can be used to synthesize circuits that have *at most* a depth of d_{\max} but does not necessarily synthesize depth-optimal circuits. This is only guaranteed if d_{\max} is exactly the optimal depth, which has to be determined first.

One approach would be to start with an initial guess for d_{\max} and iteratively decrease it until the corresponding SAT instance has a solution but the SAT instance for $d_{\max} - 1$ does not. A (theoretically) efficient way to achieve just that is binary search. The original circuit's depth can be used as an upper bound. If this is far away from the optimum, however, it can lead to the generation of instances that are tough to solve.

Instead, the upper bound can be determined dynamically. For a related problem (state preparation circuits), binary search has been explored in Ref. [31] where it was also proposed to geometrically increase the depth horizon in which a solution is searched for. Unfortunately, after a few iterations, the SAT calls will be quite costly and only promises to speedup the entire optimization if the runtime to solve a SAT instance grows sub-exponentially. In the case of exponential growth, simply searching linearly or in an arithmetic progression is faster.

Yet another way of gauging the initial depth is from empirical knowledge. If it is known that, on average, the SAT method produces solutions that are 20% shallower than the output circuit by another optimization routine, we can simply try to run it and start with the expected depth as an initial guess. From this initial guess, linear or quadratic probing can be employed to find the optimal solution.

VI. HEURISTIC APPROACH VIA CIRCUIT DECOMPOSITION

Above, we have seen how to reformulate Clifford synthesis as a SAT problem. However, the search space for that problem grows exponentially with the maximal circuit depth d_{\max} and the number of qubits n . Depending on the specific SAT solver

being used, this exact synthesis approach can quickly become prohibitively expensive. One way to diminish these scaling issues is to split a big Clifford circuit into a collection of sub-circuits that can be synthesized in parallel. This splitting can be done both horizontally (to reduce qubit number) as well as vertically (to reduce circuit depth) and considerably reduces the size of the SAT search space. The result is a versatile heuristic (it cannot be guaranteed that the splitting into sub-circuits is optimal) that can be applied to larger Clifford circuits.

More precisely, let G be a target Clifford circuit on n qubits with maximum depth d_{\max} . Then, the associated SAT encoding features bitstrings of length $l = O(n^2 d_{\max})$, which corresponds to a search space of size $2^l = 2^{O(n^2 d_{\max})}$. We can now vertically split up $U = L_1 L_2$, where each L_i has depth $d' \approx d_{\max}/2$, and apply our Clifford synthesis to each Clifford block. The result is two parallel SAT instances with bitstrings of length $l' \approx l/2$ each. In turn, the size of the search space is only $2^{l'} \approx \sqrt{2^l}$. This quadratic improvement in search space size comes at the cost of a (potentially) non-optimal decomposition into two blocks.

This general idea extends to more than two vertical blocks. Let $G = L_1 \cdots L_m$ be a Clifford circuit with m layers, i.e., blocks of qubits $L_i = g_i^1 \cdots g_i^{k_i}$ such that all $g_i^{k_i}$ act on different qubits and can therefore be run in parallel. Given a *split size* s , we can partition the circuit into $\frac{m}{s}$ sub-circuits $G_i = L_{i \cdot s} \cdots L_{(i+1) \cdot s}$ for $0 \leq i < m$ (here it is assumed that m is divisible by s but extending the argument is straightforward). We then simply need to compute the target stabilizer tableau for each of these sub-circuits. For G_i , this can be done by simulating the circuit up to $L_{i \cdot s}$. We can then use the SAT reformulation proposed in Sec. V for each of the sub-circuits. The split size s can also provide a good initial guess as to the maximal depth needed for the encoding of each individual circuit.

Since no data needs to be shared between the individual instances, all the SAT instances can be run in parallel to obtain the optimized sub-circuits G'_i which are then concatenated for the final result, i.e., $G' = G'_1 \cdots G'_m$. Note that such a splitting approach is guaranteed to produce a correct Clifford gate decomposition. This follows from applying Theorem 1 to each of the subblocks involved. However, it may not achieve optimal circuit depth. After all, this divide-and-conquer heuristic treats different blocks of the target circuit completely independently and scales with circuit depth and qubit number of the initial circuit. Scaling issues can be countered by decreasing the split size, but the point of diminishing returns is eventually reached where the size of the split has a strong negative impact on the target metric.

Since any sub-circuit can be optimized using the SAT method without changing the circuit's functionality, we can take the divide-and-conquer approach even further. Given a maximal number of qubits n_{\max} , a circuit can be decomposed into sub-circuits $G = G_1 \cdots G_m$ such that the number of qubits in each circuit is bounded by n_{\max} and there are no two-qubit gates between any two sub-circuits for parallel

Table I
EXPERIMENTAL RESULTS FOR RANDOM CLIFFORD CIRCUITS.

n	Optimal			Heuristic Vertical			Heuristic Horizontal			Bravyi et al.		
	d	$ G $	t [s]	d	$ G $	t [s]	d	$ G $	t [s]	d	$ G $	t [s]
3	5.70	11.40	0.33	8.30	18.90	0.69	-	-	-	11.70	16.10	0.18
4	6.60	16.70	4.10	13.30	23.70	2.03	-	-	-	16.00	23.50	0.16
5	7.60	25.00	381.95	18.90	38.60	7.72	7.80	25.50	603.08	22.90	37.30	0.18
6	-	-	-	24.60	55.40	27.27	17.00	51.10	180.77	29.40	55.10	0.20
7	-	-	-	31.40	69.40	43.51	25.10	75.80	97.16	37.00	70.20	0.17
8	-	-	-	37.40	88.50	153.96	32.00	97.20	186.42	42.10	86.30	0.20
9	-	-	-	41.20	102.40	313.82	39.70	121.30	219.32	53.50	108.80	0.28
10	-	-	-	50.80	131.30	547.93	53.00	155.20	106.66	59.90	128.70	0.28
11	-	-	-	58.80	152.90	838.91	60.70	180.90	133.97	72.20	157.30	0.25
12	-	-	-	66.70	174.50	1464.46	75.80	217.40	117.73	78.90	170.10	0.26
13	-	-	-	75.50	206.50	2423.24	86.10	250.00	283.31	91.40	207.10	0.34
14	-	-	-	83.00	237.10	4412.97	99.80	305.70	111.89	100.30	235.20	0.33

n : Number of qubits d : Average depth $|G|$: Average gate count t : average runtime

optimization.

These two splitting techniques can be combined to make the approach as scalable as possible. Given a depth threshold d_{thr} , a split size $s < d_{thr}$, and a maximal number of qubits n_{max} , the circuit can be split into sub-circuits of at most n_{max} qubits. If any of these sub-circuits is deeper than d_{thr} , the circuits can be further split horizontally into blocks of depth s . These circuit blocks can then be optimized independently from each other (in parallel).

VII. EVALUATIONS

The methods proposed in Sec. V and Sec. VI have been implemented in C++ using the publicly available SMT solver Z3 [65]. The implementation is integrated into the quantum circuit compilation tool QMAP [17], which is part of the *Munich Quantum Toolkit* (MQT) and available at <https://github.com/cda-tum/qmap>.

To see how well the proposed methods perform in practice, we considered two types of benchmarks:

- (i) *Random Clifford circuits* (inspired by randomized benchmarking [66]–[68]): The circuits were obtained by sampling a random stabilizer tableau (including information about the destabilizers). Since our SAT encoding is based on the stabilizer tableau, this is already a valid input format for our method and no explicit circuit has to be generated. For every qubit number n , 10 random stabilizer tableaus have been generated and the results have been averaged over all these runs. The proposed methods are compared to the state of the art greedy Clifford synthesizer by Bravyi et al. [69] The timeout was set to 3 h.
- (ii) *Clifford+T implementations of Grover search* (inspired by fault-tolerant quantum computation [3], [8], [70]): we generated circuits for the Grover search algorithm using random Boolean functions as oracles. For each qubit number n , 10 circuits were generated in this fashion. Since these circuits contain T -gates, the circuit is partitioned into Clifford blocks and T blocks, and each (non-trivial) Clifford block is optimized separately.

All evaluations have been performed on a 3.6 GHz Intel Xeon W-1370P machine running Ubuntu 20.04 with 128 GiB of main memory and 16 hardware threads. For generating and synthesizing circuits as well as for the greedy optimizer, the quantum computing SDK Qiskit [12] (version 0.42.1) by IBM has been used.

The results of our experiments for synthesizing random Clifford circuits can be seen in Table I.

The data under column *Optimal* shows the results using the proposed optimal SAT approach. Unfortunately, the increase in variables and the scaling of the encoding in the number of qubits can be seen rather drastically here. It is only possible to synthesize random Clifford circuits up to 5 qubits within the given time limit. Nonetheless, we can see that other methods, especially the state of the art, synthesize n -qubit circuits that are far from depth-optimal, increasing the depth on average by 105.26% ($n = 3$), 142.42% ($n = 4$) and 201.32% ($n = 5$), respectively.

The column *Heuristic Vertical* shows the results using the proposed heuristic approach where the circuits were partitioned vertically, i.e., the resulting sub-circuits all had the same number of qubits as the original circuit. Compared to the vertical splitting heuristic, the state of the art still produces circuits that are 21.92% deeper on average.

Column *Heuristic Horizontal* shows the results using the proposed heuristic approach where the circuits were partitioned into sub-circuits of five qubits each. At this value, the optimal approach still yields results in an acceptable amount of time (thus, there are no entries for three and four qubits). This decomposition leads to much better results for lower qubit numbers but eventually produces worse results than the vertical decomposition after nine qubits – partially as an artifact of the decomposition scheme. The circuits cannot always be split perfectly into five qubit sub-circuits, and potentially parallel gates in the original circuit might not be parallel anymore after the optimization. Another reason is that it gets increasingly difficult to find deep sub-circuits of only five qubits for random Clifford circuits since the interactions

between qubits are bound to entangle more than five qubits rather quickly. A big upside of this synthesis method is its runtime. Since the runtime for synthesizing five qubit circuits is predictable, synthesizing all these sub-circuits can be done within a predictable time as well.

All in all, the results in [Table I](#) suggest that the heuristic approach could be improved with more sophisticated circuit decomposition techniques. The current state of the art produces circuits that are far from optimal, which leaves quite some room for improvement.

The Grover search benchmark was chosen to analyze the possible improvement of the depth of actual fault-tolerant quantum circuits. Guided by the results of the random Clifford benchmarks, we looked at Grover circuits for three to five qubits. The proposed optimization scheme resulted in circuits that were 13.38%, 21.71% and 16.72% shorter on average. The Grover benchmarks are made publicly available under <https://github.com/cda-tum/qmap>.

VIII. CONCLUSION AND OUTLOOK

Classical circuit synthesis, on the one hand, and Quantified Boolean Formulas (QBF), on the other, are two seemingly very different but related problems. This correspondence can be made one-to-one and forms the basis of several state-of-the-art approaches for optimal (logical) circuit synthesis: QBF solvers are employed to determine the shortest circuit representation of a desired logical functionality.

In this work, we have extended this general mindset to the quantum realm, considering the task of decomposing n -qubit Clifford circuits into as few elementary gates as possible. We then showed that deciding if it is possible to represent a given Clifford functionality with at most d_{\max} Clifford layers can be re-cast as a satisfiability (SAT) problem in $O(n^2 d_{\max})$ Boolean variables. The reduction uses maximally entangled input stimuli, as well as the Gottesman-Knill theorem. It highlights that Clifford synthesis contained in NP (the first level of the polynomial hierarchy) is easier than general logical circuit synthesis, which is complete for Σ_2^P (the second level of the polynomial hierarchy).

In the electronic design automation community, SAT solvers have been applied to tackle classical synthesis problems with great success. We showed that similar approaches apply to quantum computing and that there is a large potential to replicate the success of solving classical synthesis problems. While the optimal synthesis approach scales poorly in the number of qubits, it shows that there is a large gap between the optimal solution and the state of the art. To our knowledge, these are the first Clifford synthesis protocols that (i) are provably correct and (ii) come with a certificate of optimality. Furthermore, all of the proposed methods are publicly available within the *Munich Quantum Toolkit* (MQT) as part of the open-source quantum circuit compilation tool QMAP (<https://github.com/cda-tum/qmap>). QMAP already works natively with IBM’s Qiskit, and even tighter integration with quantum SDKs is left for future work.

These initial findings are encouraging and open the door for several interesting follow-up projects, i.e., a *refinement of the proposed numerical solver*. This will entail tweaks in the stabilizer encoding to squeeze out more performance, but also trying different SAT solvers and novel pre-processing techniques to determine sharper initial bounds on the maximum circuit depth. A *configurable gate-set* is also on our to-do list. For now, we only use $H, S, S^\dagger, CNOT$, as well as Pauli gates. In the future, this could be adapted to include additional single-qubit gates (e.g., the full single-qubit Clifford group) and two-qubit gates (e.g., $SWAP, CZ, CY, \dots$). Note that more ‘elementary’ gates directly translate into a more complex encoding and, therefore, a larger logical search space (more variables). On the other hand, this increased expressiveness per time step is bound to decrease the circuit depth and, therefore, result in shorter SAT formulas overall (fewer logical clauses). This trade-off might well be worthwhile.

The proposed encoding procedure is flexible enough to facilitate *architecture-aware* synthesis of Clifford circuits. Some quantum architectures only support certain interactions between their qubits, typically defined by a coupling map. We can respect this coupling map in the proposed SAT encoding by only permitting Clifford gates that are also native to the concrete architecture.

Virtually all these future research directions also extend to improve the proposed *heuristic solver* for near-optimal Clifford synthesis. We intend to explore different divide-and-conquer strategies (decomposing into sub-circuits) and explore (near-)optimal ways how to best synthesize each of these circuit blocks.

Last but not least, researchers are now beginning to suggest and explore the use of quantum computers to solve challenging subroutines in quantum synthesis. Quantum-assisted quantum compiling [58] falls into this category. For this work, the deliberate restriction to Clifford circuits has allowed us to not have to think along these lines (yet) – the Gottesman-Knill theorem ensures that classical simulation remains tractable throughout. But it is an interesting direction for future work to fruitfully combine quantum-assisted quantum-compiling ideas with conventional SAT-solving techniques. We leave such synergies for future work.

ACKNOWLEDGMENTS

The authors thank Armin Biere, David Gross, and Martina Seidl for inspiring discussions and valuable feedback.

T.P., R.W. and L.B. acknowledge funding from the European Research Council (ERC) under the European Unions Horizon 2020 research and innovation program (grant agreement No. 101001318), as well as financial support from the Munich Quantum Valley, which is supported by the Bavarian state government with funds from the Hightech Agenda Bayern Plus. All authors have been supported by the BMWK on the basis of a decision by the German Bundestag through the projects ProvideQ and QuaST, the Project QuantumReady (FFG 896217) and the State of Upper Austria in the frame of the COMET program (managed by the FFG).

REFERENCES

- [1] John Preskill, “Quantum computing in the NISQ era and beyond,” vol. 2, p. 79, 2018.
- [2] Peter W. Shor, “Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer,” *SIAM J. Comput.*, 1997.
- [3] Lov K. Grover, “A fast quantum mechanical algorithm for database search,” *Proc. of the ACM*, pp. 212–219, 1996.
- [4] Aram W. Harrow, Avinandan Hassidim, and Seth Lloyd, “Quantum algorithm for linear systems of equations,” *Physical Review Letters*, vol. 103, no. 15, 2009.
- [5] Fernando G.S.L. Brandao and Krysta M. Svore, “Quantum Speed-Ups for solving semidefinite programs,” in *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*, 2017, pp. 415–426. DOI: [10.1109/FOCS.2017.45](https://doi.org/10.1109/FOCS.2017.45).
- [6] Joran van Apeldoorn *et al.*, “Quantum SDP-Solvers: Better upper and lower bounds,” vol. 4, p. 230, 2020. DOI: [10.22331/q-2020-02-14-230](https://doi.org/10.22331/q-2020-02-14-230).
- [7] Fernando G. S. L. Brandão, Richard Kueng, and Daniel Stilck França, “Faster quantum and classical SDP approximations for quadratic binary optimization,” vol. 6, p. 625, 2022. DOI: [10.22331/q-2022-01-20-625](https://doi.org/10.22331/q-2022-01-20-625).
- [8] Alexei Kitaev, “Quantum computations: Algorithms and error correction,” *Russian Mathematical Surveys*, vol. 52, no. 6, pp. 1191–1249, 1997.
- [9] Peter W. Shor, “Fault-tolerant quantum computation,” in *Proceedings of 37th Conference on Foundations of Computer Science*, 1996, pp. 56–65. DOI: [10.1109/SFCS.1996.548464](https://doi.org/10.1109/SFCS.1996.548464).
- [10] Michael A. Nielsen and Isaac L. Chuang, *Quantum Computation and Quantum Information*. Cambridge University Press, 2010.
- [11] Daniel Gottesman, “Stabilizer codes and quantum error correction,” 1997.
- [12] Abraham Asfaw *et al.* “Learn quantum computation using Qiskit.” (2020).
- [13] Vadym Kliuchnikov. “Synthesis of unitaries with Clifford+T circuits.” arXiv: [1306.3200](https://arxiv.org/abs/1306.3200) [quant-ph]. (2013), preprint.
- [14] Matthew Amy *et al.*, “A meet-in-the-middle algorithm for fast synthesis of depth-optimal quantum circuits,” *IEEE Trans. on CAD of Integrated Circuits and Systems*, vol. 32, no. 6, pp. 818–830, 2013. DOI: [10.1109/TCAD.2013.2244643](https://doi.org/10.1109/TCAD.2013.2244643).
- [15] Olivia Di Matteo and Michele Mosca, “Parallelizing quantum circuit synthesis,” *Quantum Sci. Technol.*, vol. 1, no. 1, p. 015 003, 2016. DOI: [10.1088/2058-9565/1/1/015003](https://doi.org/10.1088/2058-9565/1/1/015003). arXiv: [1606.07413](https://arxiv.org/abs/1606.07413) [quant-ph].
- [16] Philipp Niemann, Robert Wille, and Rolf Drechsler, “Advanced exact synthesis of Clifford+T circuits,” *Quantum Inf Process*, vol. 19, no. 9, p. 317, 2020. DOI: [10.1007/s1128-020-02816-0](https://doi.org/10.1007/s1128-020-02816-0).
- [17] Robert Wille and Lukas Burgholzer, “MQT QMAP: Efficient quantum circuit mapping,” in *Int’l Symp. on Physical Design*, 2023.
- [18] A Robert Calderbank and Peter W Shor, “Good quantum error-correcting codes exist,” *Physical Review A*, vol. 54, no. 2, p. 1098, 1996.
- [19] Scott Aaronson and Daniel Gottesman, “Improved simulation of stabilizer circuits,” *Phys. Rev. A*, vol. 70, no. 5, p. 052 328, 2004. DOI: [10.1103/PhysRevA.70.052328](https://doi.org/10.1103/PhysRevA.70.052328).
- [20] Sergey Bravyi and Alexei Kitaev, “Universal quantum computation with ideal Clifford gates and noisy ancillas,” *Phys. Rev. A*, vol. 71, no. 2, p. 022 316, 2005. DOI: [10.1103/PhysRevA.71.022316](https://doi.org/10.1103/PhysRevA.71.022316).
- [21] Sahel Alounh *et al.*, “A comprehensive study and analysis on SAT-solvers: Advances, usages and achievements,” *Artif Intell Rev*, vol. 52, no. 4, pp. 2575–2601, 2019. DOI: [10.1007/s10462-018-9628-0](https://doi.org/10.1007/s10462-018-9628-0).
- [22] Lucas Berent, Lukas Burgholzer, and Robert Wille, “Towards a SAT encoding for quantum circuits: A journey from classical circuits to Clifford circuits and beyond,” in *International Conference on Theory and Applications of Satisfiability Testing*, 2022. arXiv: [2203.00698](https://arxiv.org/abs/2203.00698).
- [23] Sanjeev Arora and Boaz Barak, *Computational Complexity: A Modern Approach*. Cambridge University Press, 2009. DOI: [10.1017/CBO9780511804090](https://doi.org/10.1017/CBO9780511804090).
- [24] Ankit Shukla *et al.*, “A Survey on applications of quantified boolean formulas,” in *International Conference on Tools with Artificial Intelligence*, 2019, pp. 78–84. DOI: [10.1109/ICTAI.2019.00020](https://doi.org/10.1109/ICTAI.2019.00020).
- [25] Markus N. Rabe and Sanjit A. Seshia, “Incremental determination,” in *Conference on Theory and Applications of Satisfiability Testing*, vol. 9710, 2016, pp. 375–392.
- [26] Florian Lonsing and Uwe Egly, “DepQBF 6.0: A search-based QBF solver beyond traditional QCDCL,” in *International Conference on Automated Deduction*, vol. 10395, 2017, pp. 371–384.
- [27] Markus N. Rabe and Leander Tentrup, “CAQE: A Certifying QBF Solver,” in *Int’l Conf. on Formal Methods in CAD*, 2015, pp. 136–143. DOI: [10.1109/FMCAD.2015.7542263](https://doi.org/10.1109/FMCAD.2015.7542263).
- [28] Armando Solar-Lezama *et al.*, “Combinatorial sketching for finite programs,” in *Int’l Conf. On Architectural Support for Programming Languages and Operating Systems*, 2006, pp. 404–415. DOI: [10.1145/1168857.1168907](https://doi.org/10.1145/1168857.1168907).
- [29] Edmund Clarke *et al.*, “Counterexample-guided abstraction refinement,” in *Computer Aided Verification*, vol. 1855, 2000, pp. 154–169.
- [30] Susmit Jha and Sanjit A. Seshia, “A theory of formal synthesis via inductive learning,” *Acta Informatica*, vol. 54, no. 7, pp. 693–726, 2017. DOI: [10.1007/s00236-017-0294-5](https://doi.org/10.1007/s00236-017-0294-5).
- [31] Sarah Schneider, Lukas Burgholzer, and Robert Wille, “A SAT encoding for optimal Clifford circuit synthesis,” in *Asia and South Pacific Design Automation Conf.*, 2023.
- [32] David Gross *et al.*, “Quantum state tomography via compressed sensing,” *Phys. Rev. Lett.*, vol. 105, no. 15, p. 150 401, 2010. DOI: [10.1103/PhysRevLett.105.150401](https://doi.org/10.1103/PhysRevLett.105.150401).
- [33] Martin Kliesch *et al.*, “Guaranteed recovery of quantum processes from few measurements,” vol. 3, p. 171, 2019. DOI: [10.22331/q-2019-08-12-171](https://doi.org/10.22331/q-2019-08-12-171).
- [34] Ingo Roth *et al.*, “Recovering quantum gates from few average gate fidelities,” *Phys. Rev. Lett.*, vol. 121, no. 17, p. 170 502, 2018. DOI: [10.1103/PhysRevLett.121.170502](https://doi.org/10.1103/PhysRevLett.121.170502).
- [35] Lukas Burgholzer, Richard Kueng, and Robert Wille, “Random stimuli generation for the verification of quantum circuits,” in *Asia and South Pacific Design Automation Conf.*, 2021.
- [36] Madalin Gu *et al.*, “Fast state tomography with optimal error bounds,” *J. Phys. A: Math. Theor.*, vol. 53, no. 20, p. 204 001, 2020. DOI: [10.1088/1751-8121/ab8111](https://doi.org/10.1088/1751-8121/ab8111).
- [37] Dominik Janzing, Pawel Wocjan, and Thomas Beth, ““Non-identity check” is QMA-complete,” *Int. J. Quantum Inform.*, vol. 03, no. 03, pp. 463–473, 2005.
- [38] Christopher M. Dawson and Michael A. Nielsen, “The Solovay-Kitaev algorithm,” *Quantum Info. Comput.*, vol. 6, no. 1, pp. 81–95, 2006.
- [39] Tien Trung Pham, Rodney Van Meter, and Clare Horsman, “Optimization of the Solovay-Kitaev algorithm,” *Phys. Rev. A*, vol. 87, no. 5, p. 052 332, 2013. DOI: [10.1103/PhysRevA.87.052332](https://doi.org/10.1103/PhysRevA.87.052332).
- [40] Jeroen Dehaene and Bart De Moor, “Clifford group, stabilizer states, and linear and quadratic operations over GF(2),” *Phys. Rev. A*, vol. 68, no. 4, p. 042 318, 2003. DOI: [10.1103/PhysRevA.68.042318](https://doi.org/10.1103/PhysRevA.68.042318).

- [41] David Gross, “Hudsons theorem for finite-dimensional quantum systems,” *J. Math. Phys.*, vol. 47, no. 12, p. 122 107, 2006. DOI: [10.1063/1.2393152](https://doi.org/10.1063/1.2393152).
- [42] Huangjun Zhu *et al.* “The Clifford group fails gracefully to be a unitary 4-design.” arXiv: [1609.08172](https://arxiv.org/abs/1609.08172) [quant-ph]. (2016), preprint.
- [43] D. Maslov, “Linear Depth Stabilizer and Quantum Fourier Transformation Circuits with no Auxiliary Qubits in Finite Neighbor Quantum Architectures,” *Phys. Rev. A*, vol. 76, no. 5, p. 052 310, 2007. DOI: [10.1103/PhysRevA.76.052310](https://doi.org/10.1103/PhysRevA.76.052310). arXiv: [quant-ph/0703211](https://arxiv.org/abs/quant-ph/0703211).
- [44] Sergey Bravyi, Joseph A. Latone, and Dmitri Maslov. “6-qubit optimal Clifford circuits.” arXiv: [2012.06074](https://arxiv.org/abs/2012.06074). (2020), preprint.
- [45] Narayanan Rengaswamy *et al.*, “Logical Clifford synthesis for stabilizer codes,” *IEEE Transactions on Quantum Engineering*, vol. 1, pp. 1–17, 2020. DOI: [10.1109/TQE.2020.3023419](https://doi.org/10.1109/TQE.2020.3023419).
- [46] Sergey Bravyi, Dmitri Maslov, and Yunseong Nam, “Constant-cost implementations of Clifford operations and multiply-controlled gates using global interactions,” *Phys. Rev. Lett.*, vol. 129, no. 23, p. 230 501, 2022. DOI: [10.1103/PhysRevLett.129.230501](https://doi.org/10.1103/PhysRevLett.129.230501).
- [47] Roderick Bloem *et al.* “SAT-Based Methods for Circuit Synthesis.” arXiv: [1408.2333](https://arxiv.org/abs/1408.2333) [cs]. (2014), preprint.
- [48] Roderick Bloem, Robert Könighofer, and Martina Seidl, “SAT-Based synthesis methods for safety specs,” in *Verification, Model Checking, and Abstract Interpretation*, 2014, pp. 1–20.
- [49] Rüdiger Ehlers, Robert Könighofer, and Georg Hofferek, “Symbolically synthesizing small circuits,” in *Int’l Conf. on Formal Methods in CAD*, 2012, pp. 91–100.
- [50] Franz-Xaver Reichl, Friedrich Slivovsky, and Stefan Szeider, “Circuit minimization with QBF-Based exact synthesis,” *Proceedings of the AAI Conference on Artificial Intelligence*, vol. 37, no. 4, pp. 4087–4094, 2023. DOI: [10.1609/aaai.v37i4.25524](https://doi.org/10.1609/aaai.v37i4.25524).
- [51] Robert Wille *et al.*, “Quantified synthesis of reversible logic,” in *Design, Automation and Test in Europe*, 2008, pp. 1015–1020. DOI: [10.1109/DATE.2008.4484814](https://doi.org/10.1109/DATE.2008.4484814).
- [52] Heinz Riener *et al.*, “On-the-fly and DAG-aware: Rewriting boolean networks with exact synthesis,” in *Design, Automation and Test in Europe*, 2019, pp. 1649–1654. DOI: [10.23919/DATE.2019.8715185](https://doi.org/10.23919/DATE.2019.8715185).
- [53] Man-Duen Choi, “Completely positive linear maps on complex matrices,” *Linear Algebra and its Applications*, vol. 10, no. 3, pp. 285–290, 1975. DOI: [10.1016/0024-3795\(75\)90075-0](https://doi.org/10.1016/0024-3795(75)90075-0).
- [54] Andrzej Jamiolkowski, “Linear transformations which preserve trace and positive semidefiniteness of operators,” *Reports on Mathematical Physics*, vol. 3, no. 4, pp. 275–278, 1972. DOI: [10.1016/0034-4877\(72\)90011-0](https://doi.org/10.1016/0034-4877(72)90011-0).
- [55] Joseph B. Altepeter *et al.*, “Ancilla-assisted quantum process tomography,” *Phys. Rev. Lett.*, vol. 90, no. 19, p. 193 601, 2003. DOI: [10.1103/PhysRevLett.90.193601](https://doi.org/10.1103/PhysRevLett.90.193601).
- [56] John Watrous, *The Theory of Quantum Information*. Cambridge University Press, 2018, 590 pp.
- [57] Richard Kueng. “ACM 270: Quantum and classical information processing with tensors.” (2019).
- [58] Sumeet Khatri *et al.*, “Quantum-assisted quantum compiling,” vol. 3, p. 140, 2019.
- [59] Hans De Raedt *et al.*, “Massively parallel quantum computer simulator, eleven years later,” *Computer Physics Communications*, vol. 237, pp. 47–61, 2019. DOI: [10.1016/j.cpc.2018.11.005](https://doi.org/10.1016/j.cpc.2018.11.005).
- [60] Jacob C. Bridgeman and Christopher T. Chubb, “Hand-waving and interpretive dance: An introductory course on tensor networks,” *J. Phys. A: Math. Theor.*, 2017.
- [61] Stefan Hillmich, Igor L. Markov, and Robert Wille, “Just like the real thing: Fast weak simulation of quantum computation,” in *Design Automation Conf.*, 2020.
- [62] Philipp Niemann *et al.*, “QMDDs: Efficient quantum function representation and manipulation,” *IEEE Trans. on CAD of Integrated Circuits and Systems*, 2016.
- [63] Daniel Gottesman. “The Heisenberg representation of quantum computers.” arXiv: [quant-ph/9807006](https://arxiv.org/abs/quant-ph/9807006). (1998), preprint.
- [64] Armin Biere *et al.*, *Handbook of Satisfiability*. IOS Press, 2009.
- [65] Leonardo de Moura and Nikolaj Bjørner, “Z3: An efficient SMT solver,” in *Tools and Algorithms for the Construction and Analysis of Systems*, 2008, pp. 337–340.
- [66] Emanuel Knill *et al.*, “Randomized benchmarking of quantum gates,” *Phys. Rev. A*, vol. 77, no. 1, p. 012 307, 2008. DOI: [10.1103/PhysRevA.77.012307](https://doi.org/10.1103/PhysRevA.77.012307).
- [67] Easwar Magesan *et al.*, “Efficient measurement of quantum gate error by interleaved randomized benchmarking,” *Phys. Rev. Lett.*, vol. 109, no. 8, p. 080 505, 2012. DOI: [10.1103/PhysRevLett.109.080505](https://doi.org/10.1103/PhysRevLett.109.080505).
- [68] Richard Kueng *et al.*, “Comparing experiments to the fault-tolerance threshold,” *Phys. Rev. Lett.*, vol. 117, no. 17, p. 170 502, 2016.
- [69] Sergey Bravyi *et al.*, “Clifford circuit optimization with templates and symbolic Pauli gates,” vol. 5, p. 580, 2021. DOI: [10.22331/q-2021-11-16-580](https://doi.org/10.22331/q-2021-11-16-580). arXiv: [2105.02291](https://arxiv.org/abs/2105.02291).
- [70] Daniel Gottesman, “Theory of fault-tolerant quantum computation,” *Phys. Rev. A*, vol. 57, no. 1, pp. 127–137, 1998. DOI: [10.1103/PhysRevA.57.127](https://doi.org/10.1103/PhysRevA.57.127).