

Automated Synthesis of Fault-Tolerant State Preparation Circuits for Quantum Error-Correction Codes

Tom Peham^{1,*}, Ludwig Schmid^{1,†}, Lucas Berent^{1,‡}, Markus Müller^{2,3,§} and Robert Wille^{1,4,||}

¹Chair for Design Automation, *Technical University of Munich, Germany*

²Institute for Quantum Information, *RWTH Aachen University, Aachen D-52056, Germany*

³Peter Grünberg Institute, *Theoretical Nanoelectronics, Forschungszentrum Jülich, Jülich D-52425, Germany*

⁴Software Competence Center Hagenberg, *Austria*



(Received 30 August 2024; revised 4 December 2024; accepted 17 March 2025; published 14 May 2025; corrected 13 June 2025)

A central ingredient in fault-tolerant quantum algorithms is the initialization of a logical state for a given quantum error-correcting code from a set of noisy qubits. A scheme that has demonstrated promising results for small code instances that are realizable on currently available hardware composes a non-fault-tolerant state preparation circuit with a verification circuit that checks for spreading errors. Known circuit constructions of this scheme are mostly obtained manually, and no algorithmic techniques for constructing depth- or gate-optimal circuits exist. As a consequence, the current state-of-the-art exploits this scheme only for specific code instances and mostly for the special case of distance $d = 3$ codes only. In this work, we propose an automated approach for synthesizing fault-tolerant state preparation circuits for arbitrary CSS codes. We utilize methods based on satisfiability solving (SAT) to construct fault-tolerant state preparation circuits consisting of depth- and gate-optimal preparation and verification circuits. We also provide heuristics that can synthesize fault-tolerant state preparation circuits for code instances where no optimal solution can be obtained in an adequate time. Moreover, we give a general construction for nondeterministic state preparation circuits for codes beyond distance 3. Numerical evaluations using $d = 3$, $d = 5$, and $d = 7$ codes confirm that the generated circuits exhibit the desired scaling of the logical error rates. The resulting methods are publicly available as part of the *Munich Quantum Toolkit* (MQT) at <https://github.com/cda-tum/mqt-qecc>. Such methods are an important step in providing fault-tolerant circuit constructions that can aid in near-term demonstrations of fault-tolerant quantum computing.

DOI: [10.1103/PRXQuantum.6.020330](https://doi.org/10.1103/PRXQuantum.6.020330)

I. INTRODUCTION

Quantum bits and quantum operations suffer from unavoidable decoherence and noise. Therefore, quantum error-correction mechanisms need to be used to ensure that a quantum algorithm can be executed in a fault-tolerant manner to control the accumulation of errors. Error-correcting codes leverage redundancy by encoding quantum information in logical qubits formed by entangled states of noisy, physical qubits. Consequently, it is

important that universal operations can be carried out on the encoded information fault-tolerantly and that errors during the execution can be detected and corrected. A crucial step in any fault-tolerant circuit is the preparation of an encoded logical state.

To ensure that the logical state is initialized correctly—without containing too many errors—the preparation itself needs to be fault-tolerant. Intuitively, a requirement for fault-tolerant circuits is that errors on single qubits cannot uncontrollably “spread” to multiple qubits, for instance, through CNOT gates. For *Calderbank-Shor-Steane* (CSS) codes [1–3], one way of initializing an encoded state fault-tolerantly is to prepare all physical qubits in $|0\rangle$ or $|+\rangle$ and perform one round of error correction—projecting the product state onto the logical code space [4]. For large code instances, this overhead might be manageable, especially for QLDPC codes and single-shot codes [5–9].

This general procedure introduces significant overhead in circuit size, especially for small code instances that are relevant for near-to-midterm devices since all

*Contact author: tom.peham@tum.de

†Contact author: ludwig.s.schmid@tum.de

‡Contact author: lucas.berent@tum.de

§Contact author: markus.mueller@fz-juelich.de

||Contact author: robert.wille@tum.de

Published by the American Physical Society under the terms of the *Creative Commons Attribution 4.0 International* license. Further distribution of this work must maintain attribution to the author(s) and the published article’s title, journal citation, and DOI.

stabilizer generators must be measured multiple times in an error-correction round. Consequently, the near-term experimental realization of this scheme is difficult. An alternative that tries to reduce the number of measurements is a nondeterministic protocol, based on postselection [10–13]. Therein, the initialization is performed in two steps:

- (1) Prepare the logical state using a non-fault-tolerant state preparation circuit.
- (2) Conduct specific measurements on the prepared states using a so-called *verification circuit* that detects if an error spreads through the circuit. If at least one of these measurements measures a -1 eigenvalue, the state is discarded, and the protocol is restarted.

While this and related techniques were successfully demonstrated in recent experiments [11,12,14–19], there are clear drawbacks of current approaches: First, most of these techniques rely on manual construction, which is tedious and becomes infeasible when scaling to larger codes, where both the number of errors and the number of possible verification measurements increase quickly. Moreover, as for the general scheme outlined above, there is, in general, no guarantee that these manually constructed circuits are gate or depth optimal. Finally, these constructions have mostly been explored for specific code instances for small distance 3 (or 2) codes [10,12,14,20] that constitute a special case. Even for cases beyond distance 3, constructions of nondeterministic circuits were derived manually for specific states [13] or codes [21,22]. In general, constructing a verification circuit that ensures fault tolerance of the overall scheme is highly nontrivial and depends on the code (and the logical state to prepare). Thus, the manual construction in existing works is not generally applicable. Moreover, even though postselection schemes do not scale well to large distances due to the exponentially decreasing acceptance rate, optimal circuit constructions might still be viable for moderately sized codes on near-term hardware. Furthermore, the improvements to circuit size for small codes are amplified when concatenating codes, as the same state preparation circuit can be used for higher levels of concatenation where a single CNOT translates to a transversal CNOT between all data qubits of two code states. In particular, with the continuously improving error rates on current devices, experimental demonstrations of higher-distance codes might be possible with a verification-based scheme. To this end, the construction needs to be generalized beyond distance 3, and algorithmic methods that can be applied to a broad range of codes must be developed.

In this work, we address these issues and propose automated methods for the synthesis of fault-tolerant

state preparation circuits for *arbitrary* CSS codes. Our main contributions are as follows:

- (a) We separate the synthesis problems for a state preparation and verification circuit and tackle them individually. We frame both problems as binary optimization problems and discuss their computational complexity to provide a motivation for the techniques we use.
- (b) Based on that, we utilize satisfiability-solving (SAT) techniques, which, through the use of highly optimized solving algorithms, have been shown to be successful in dealing with small instances of optimization problems—both in the classical domain [23–29] and for quantum computing [30–34]. Thereby, we obtain exact solutions and, consequently, gate- or depth-optimal state preparation circuits and verification circuits that are gate optimal with respect to a given state preparation circuit.
- (c) To address the scalability issues of SAT techniques, we also propose *heuristic* solutions for synthesizing state preparation and verification circuits. While these do not guarantee optimal solutions, this will allow us to obtain solutions for larger problem instances.
- (d) To use the proposed methods for synthesizing logical basis states for CSS codes with distances greater than 3, we propose a generalized nondeterministic fault-tolerant state preparation scheme that guarantees fault tolerance by composing multiple verification circuits. The resulting scheme is summarized in Fig. 1.

In short, our contributions can be viewed as providing (optimal) circuits for every sub-circuit depicted in Fig. 1. To investigate the error suppression of the circuits obtained with the proposed methods, we synthesize nondeterministic fault-tolerant state preparation circuits for various distance-3 CSS codes and provide numerical simulations of the logical error rate under a circuit-level noise model using Stim [35]. Comparison with the state-of-the-art reinforcement learning-based synthesis method from Ref. [20] shows that our methods are competitive—producing equally good or better circuits than the reinforcement learning agent. Moreover, numerical evaluations show that the proposed scheme does indeed give circuits that exhibit the desired logical error-rate scaling for codes beyond distance 3. Finally, all proposed methods are made publicly available in the form of open-source software as part of the *Munich Quantum Toolkit* (MQT [36]) at <https://github.com/cda-tum/mqt-gecc>.

This paper is structured as follows. In Sec. II, we review notation, give fundamental notions from quantum coding, and define the notion of a fault-tolerant state preparation protocol. Section III describes the problem

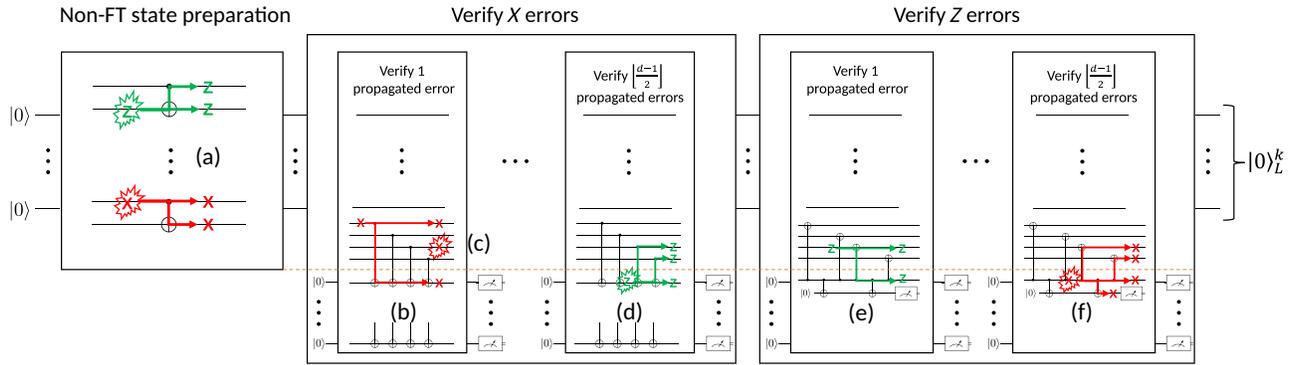


FIG. 1. Full nondeterministic fault-tolerant state preparation protocol for CSS codes. A sequence of stabilizer measurements follows a non-fault-tolerant state preparation circuit to check if any errors propagated through the circuit. The state is accepted if no measurement in the verification blocks indicates an error. (a) Errors in the non-fault-tolerant state preparation circuits propagate to higher-weight errors through CNOTs between data qubits. (b) If an error of at most weight i occurred in the state preparation circuit and propagated to a higher-weight error, measurements in the i th layer of the verification circuit detect this error. (c) Propagated errors are still detected by later layers of verification, even if further errors occur during a stabilizer measurement. (d) Z errors on the ancilla during the verification of X errors propagate to the data qubits. (e) If a Z error propagated through the non-FT state preparation circuit or the Z measurements of the previous verifications, it is later detected by flag fault-tolerant measurements. (f) Propagated X errors on the X measurement ancillas are detected by flag qubit measurements.

of synthesizing nondeterministic fault-tolerant state preparation circuits and positions our work in the context of existing approaches. In Sec. IV, we describe our optimal and heuristic synthesis methods for non-fault-tolerant state preparation for CSS codes. How to verify such circuits and make them fault-tolerant is then described in Sec. V, where we also discuss the computational complexity of this problem. Based on these building blocks, we then detail our generalized scheme for nondeterministic fault-tolerant state preparation in Sec. VI. Section VII shows the resulting circuit for our synthesis method for various CSS codes and numerically verifies that our proposed scheme prepares logical basis states of CSS codes in a fault-tolerant manner. Finally, Sec. VIII concludes this paper.

II. BACKGROUND

In this section, we review formal notation and discuss fundamental concepts that are needed throughout the work. Throughout we use $[i] = \{1, 2, \dots, i\} \subseteq \mathbb{Z}$. We make use of the *iverson bracket*, which maps Boolean expressions to

$$\text{integers } \llbracket P \rrbracket = \begin{cases} 1 & \text{if } P \text{ is true} \\ 0 & \text{otherwise,} \end{cases}$$

A. Quantum codes

An $\llbracket n, k, d \rrbracket$ -stabilizer code [37] is defined by a stabilizer group S , which is an Abelian subgroup of the n -qubit Pauli group \mathcal{P}_n . We require $-I \notin S$ such that the code is nontrivial. The set of codewords, i.e., the k -dimensional codespace, is defined as the simultaneous $+1$ eigenspace of all elements in S . The generators of the stabilizer group

$S = \langle g_1, g_2, \dots, g_m \rangle$ are called *checks*. They are used to infer whether a state is a codeword by measuring all m checks (using additional ancillas). The m measurement outcomes can be modeled as vector $s \in \mathbb{F}_2^m$ called the *syndrome*. If all syndrome bits are 0 (indicating a $+1$ measurement of the eigenvalue), the state is a codeword. Otherwise, the syndrome indicates which checks have failed and thus gives information about the approximate location of the error that occurred. The problem of finding a suitable recovery operation given the syndrome is called *decoding* and is a purely classical task.

A stabilizer code on n physical qubits with m checks encodes $k = n - m$ logical qubits. The *distance* d of the code is the minimum weight of an operator transforming one codeword into another.

A particularly important class of stabilizer codes are Calderbank-Shor-Steane (CSS) codes [1,2], whose stabilizer generators g_i are either all X , or all Z . The corresponding sets are denoted S_X and S_Z , respectively. Note that X errors only anticommute with Z stabilizers and Z errors only anticommute with X stabilizers. Therefore, X - and Z -type errors can be decoded independently.

Example 1. The distance-3 color code, also called *Steane code*, is the smallest instance of a family of two-dimensional (2D) planar color codes [38]. It can be visualized as depicted in Fig. 2, where qubits are associated with vertices and stabilizer generators with faces. Hence, it is a $\llbracket 7, 1, 3 \rrbracket$ CSS code defined by the stabilizer generators $S_X = \{X_1X_2X_3X_6, X_1X_3X_5X_7, X_4X_5X_6X_7\}$ and $S_Z = \{Z_1Z_2Z_5Z_6, Z_1Z_3Z_5Z_7, Z_4Z_5Z_6Z_7\}$. Logical operators are associated with the sides of the triangle. For

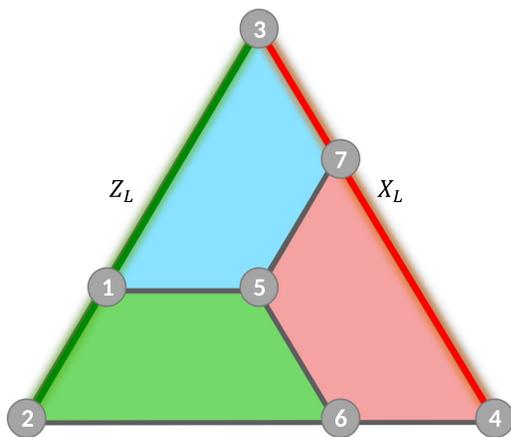


FIG. 2. Steane code. The qubits are placed on vertices and X - and Z -stabilizer checks are associated with faces, acting on adjacent vertices.

example $X_L = X_3X_4X_7$ and $Z_LZ_2Z_3$ would be minimal-weight logical X and Z operators (indicated by the red and green lines in Fig. 2). Due to the symmetry of the X and Z stabilizers, $X_L = X_1X_2X_3$ and $Z_L = Z_3Z_4Z_7$ are also valid logical operators.

The stabilizer generators in S_X, S_Z can be mapped to binary vectors that indicate the support of the respective Pauli operators, e.g., $XXII \in S_X, XXII \mapsto (1, 1, 0, 0) \in \mathbb{F}_2^n$. Hence, a CSS code can equivalently be defined as two binary matrices called *parity-check matrices*, $H_X \in \mathbb{F}_2^{n \times n}, H_Z \in \mathbb{F}_2^{n \times n}$ whose rows are the binary vectors corresponding to the stabilizer generators. Since the X - and Z -type generators have to commute, the matrices have to fulfill the orthogonality condition $H_X H_Z^T = 0$.

Example 2. The parity-check matrices of the Steane code are identical:

$$H_X = H_Z = \begin{pmatrix} 1 & 1 & & 1 & 1 & & \\ 1 & & 1 & & 1 & & 1 \\ & & & 1 & 1 & 1 & 1 \end{pmatrix}. \quad (1)$$

B. Noise model and fault tolerance

We use a standard, depolarizing circuit-level noise model parameterized by the noise strength p , as discussed in more detail in Sec. VII.

The weight of an error E is defined as $\text{wt}(E) = \min_{E' \in E \cdot S} |\text{supp}(E')|$, i.e., as the minimal support of a stabilizer reduced error with respect to the stabilizer group S of the code. We say that an error *propagates* through a two-qubit gate if an error on one qubit before the gate leads to a two-qubit error after the two-qubit gate. We call an error in a circuit \mathcal{C} a *propagated error* if it propagates through any gate in \mathcal{C} .

There are slightly different meanings to the term *fault tolerance* in the literature [39–41]. Intuitively, the definition we follow means that the uncontrolled spreading and accumulation of errors is avoided. More formally, we adhere to the following definition for a fault-tolerant state preparation protocol.

Definition 1 (Fault-tolerant state preparation [42,43]). A state preparation circuit is *strictly fault-tolerant* if for all $t \leq \lfloor d/2 \rfloor$, any error of probability order t propagates to an error of weight at most t .

III. PROBLEM DEFINITION AND MOTIVATION

In this section, we illustrate the problem of fault-tolerant state preparation with verification circuits and motivate the need for generalized automated approaches.

A. Considered problem

Before starting any fault-tolerant quantum computation, the logical qubits have to be initialized in some logical state—usually a Pauli eigenstate of the quantum error-correcting code used to encode the quantum information. These states are easily described for an $[[n, 1, d]]$ CSS code. The logical $|0\rangle_L$ state of the code is defined by the stabilizers of the code and the logical Z_L operator. Similarly, the logical $|+\rangle$ state is defined by the stabilizers of the code and the logical X_L operator. The Pauli eigenstates of a general $[[n, k, d]]$ CSS code are defined analogously for corresponding pairs of logical operators $X_L^i, Z_L^i, i \in [k]$. For the remainder of this work, we will focus on the preparation of the logical $|0\rangle_L^{\otimes k}$ state, but the methods transfer directly to preparing the $|+\rangle^{\otimes k}$ state by swapping the roles of X and Z stabilizers and operators and reversing the direction of every CNOT in the state preparation circuits.

A simple protocol for preparing Pauli eigenstates is to initialize all physical qubits in $|0\rangle$ and to project the product state onto the code space by measuring the stabilizers of the code. The logical Z_L^i operators are already stabilizers of the initial $|0\rangle^{\otimes n}$ and will still be a stabilizer of the projected state since the operators Z_L^i commute with the stabilizers of the code. To ensure fault tolerance in the presence of noisy syndrome measurements, the stabilizer measurements have to be repeated a number of times that is typically proportional to the code distance [38,44]. For Shor-type QEC, the number of required repetitions is in $O(d^2)$ [45] for instance. Additional CNOT overhead is incurred for some measurement schemes when using a flag-fault-tolerance scheme to measure the stabilizers [41]. For the Steane code, this procedure requires $6 \times (4 + 2) \times 3 = 108$ CNOTs in the worst case, i.e., the 6 weight-4 stabilizer generators of the Steane code have to be measured 3 times—the distance of the Steane code—and each measurement requires another two CNOTs for the flags. For experimental realizations in near-term quantum

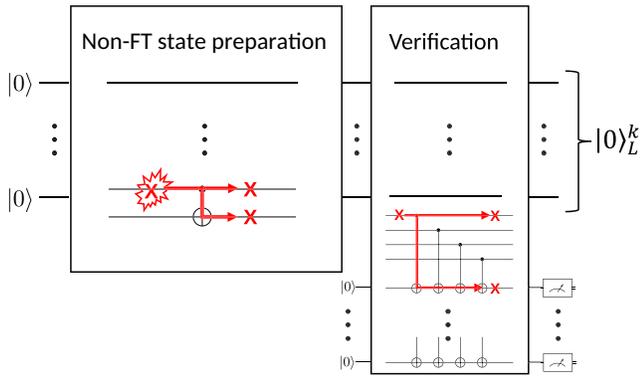


FIG. 3. Fault-tolerant Pauli-eigenstate preparation for stabilizer codes. A state is prepared with a non-fault-tolerant encoding circuit, and the state is postselected on a set of stabilizer measurements called the verification circuit. If one of these measurements fails, the state is discarded, and the process is restarted.

computers, the error rate of these CNOTs will completely overshadow any error suppression theoretically guaranteed by the Steane code.

For this reason, recent experimental demonstrations of fault tolerance have used a so-called *nondeterministic* scheme based on postselection to produce high-fidelity logical states [11, 12, 14–19]. This approach aims to prepare the logical state in a non-fault-tolerant manner and measure a set of stabilizers (including logical operators) that detect any uncorrectable error caused by a correctable error propagating through the state preparation circuit. If one of these measurements is triggered, the state is discarded, and the whole process is restarted. The measurement circuit is referred to as the *verification circuit*. The protocol is sketched in Fig. 3. The verification circuit depends on the specific state preparation circuit, as the errors that the verification measurements should detect are determined by the CNOT gate pattern in the state preparation circuit.

Example 3. For the $|0\rangle_L$ state of the Steane code, the optimal nondeterministic preparation scheme uses eight CNOTs for the preparation circuit and a single weight-three measurement of a logical Z_L operator [10] (cf. Fig. 4) in the verification circuit. The only problematic stabilizer-inequivalent errors that can occur in the circuit are X_5X_7 and X_2X_6 . These anticommute with the measured $Z_L = Z_1Z_6Z_7$ operator are thus detected in this verification measurement.

Ideally, the verification circuit only indicates an error if a propagated error occurs. Unfortunately, this is impossible in general. For example, any single-qubit error on the measurement ancilla will trigger the verification measurements. The *acceptance rate* r_A of a nondeterministic preparation scheme is the probability that a state will be accepted, i.e., that no measurement in the verification circuit is triggered. This probability decreases exponentially

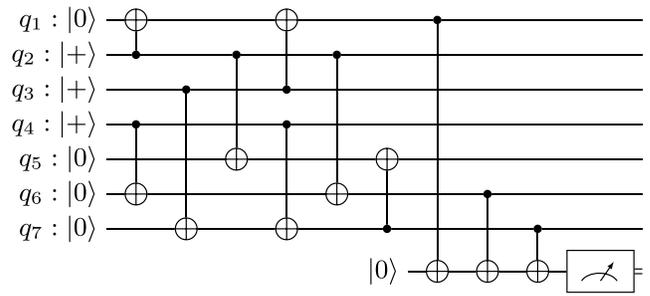


FIG. 4. Nondeterministic state preparation for the $|0\rangle_L$ of the Steane code from Ref. [10].

with the number of measurements even when ignoring errors in the state preparation circuit itself: $p_A \leq (1 - p)^m$ where m is the number of measurements.

This work aims to formulate a general algorithm that synthesizes a (near-)optimal circuit for each of these sub-circuits for any $[[n, k, d]]$ CSS code.

B. Related work

Recent work has tackled state preparation and verification circuit synthesis using reinforcement learning (RL) [20]. The biggest difference between our method and the RL approach is that we only focus on CSS codes, while the RL framework works for arbitrary stabilizer codes. Furthermore, we ignore qubit connectivity constraints in this work while this aspect is accounted for in Ref. [20]. Lastly, the RL framework also provides a method for tackling the state preparation and verification circuit synthesis problem *simultaneously*, whereas we treat them separately. The downside of the RL-based approach is that it does not give any optimality guarantees. Since the code is a parameter in the training of the RL agent, the runtimes for synthesizing a single circuit are high, even for small codes.

Since CSS code state preparation circuits consist only of CNOT gates [46], the state preparation circuit synthesis problem is closely related to the synthesis of linear reversible circuits. This is a well-researched problem with known algorithms reaching the theoretical lower bound on circuit size and depth [47–49]. These algorithms provide systematic ways that are *asymptotically optimal*. These asymptotic analyses hide a significant constant overhead that we cannot ignore for the smaller system sizes we are mainly interested in. That is why our state preparation circuit synthesis approach is more closely related to greedy CNOT circuit synthesis methods as proposed in Ref. [50]. The most important difference between the considered problems and linear circuit synthesis is that we have more degrees of freedom. Although both problems can be framed in terms of matrix elimination over \mathbb{F}_2 , the matrices in our setting are not necessarily invertible. Therefore, many different circuits prepare the same state despite implementing different linear maps.

Another closely related problem is that of general Clifford and stabilizer state preparation circuit synthesis. For these, there are also optimal [34,51,52] and asymptotically optimal [49,53–55] methods. Our SAT-based optimal state preparation circuit synthesis method is closely related to Refs. [34] and [52]. The optimal synthesis in Ref. [34] is focused only on the depth-optimal synthesis of Clifford circuits and does not provide an encoding to obtain (two-qubit) gate-optimal circuits. More importantly, as with the synthesis of linear circuits, there are more degrees of freedom that previous work does not take advantage of when focusing only on state preparation. The previous work on optimal stabilizer state preparation of Ref. [52] does not make use of this freedom either and does not guarantee the synthesis of a globally optimal circuit. Optimality is only guaranteed with respect to a specific stabilizer generator set. In the extreme case, Ref. [52] synthesizes a linear-sized circuit for the $|0\rangle^{\otimes n}$ state if a particularly unfortunate representation of its stabilizers is given as an input. For more complex states, it is not clear which generating set for the state's stabilizer group yields the best circuit. We address this specifically in Sec. IV.

IV. SYNTHESIS OF STATE PREPARATION CIRCUITS

In this section, we present a method for creating logical basis states for any CSS code given the stabilizer generators of the code. We focus on creating the $|0\rangle_L^{\otimes k}$ state, but the method works exactly the same for the $|+\rangle_L^{\otimes k}$ state by swapping the roles of the X and Z stabilizer and changing directions of CNOTs. For the rest of this section, we consider an $[[n, k, d]]$ CSS code with m X stabilizer generators.

A. Problem description and complexity

Since the orthogonal complement of the X -check matrix $H_X \in \mathbb{F}_2^{m \times n}$ is uniquely defined and coincides with the space spanned by the rows of H_Z and logical operators $Z_L^i, i \in [k]$, $|0\rangle_L^{\otimes k}$ is uniquely fixed by the X stabilizers of the code alone. An important property of logical X - or Z -basis state preparation circuits of CSS codes is that they consist only of physical qubits initialized in the X or Z basis followed by a sequence of CNOTs. In fact, there is a one-to-one correspondence between such circuits and CSS codes [46]. If we have m qubits initialized in $|+\rangle$ and $n - m$ qubits in $|0\rangle$ then the product state of these qubits has m X stabilizers and $n - m$ Z stabilizers. Applying a CNOT to this state maps X stabilizers to X stabilizers and Z stabilizers to Z stabilizers. Therefore, any CNOT circuit applied to this initial product state does not change the number of X or Z stabilizers. Combined with the uniqueness of the orthogonal complement, this tells us that if the state obtained by applying a CNOT circuit to the initial product state is stabilized by the X stabilizers of the code and no

other X stabilizers, it must necessarily also be stabilized by the Z stabilizers and the logical Z operators of the code and is, therefore, the $|0\rangle_L^{\otimes k}$ of the code. These properties simplify the considered problems significantly because, for state preparation, we only have to consider X stabilizers.

Any intermediate state during preparation is also a logical basis state of some CSS code with respective X checks H_X^a , where H_X^a denote the X checks of the partial state after the first a CNOTs. If the $a + 1$ st CNOT has control q_i and target q_j then H_X^{a+1} is obtained by adding column of H_X^a i onto column j of H_X^a . This comes from the fact that a Pauli X copies through the control of a CNOT. A CSS state preparation circuit \mathcal{C} on n qubits, therefore, implements an invertible matrix $M_{\mathcal{C}}$ over \mathbb{F}_2 converting H_X^0 to the final H_X .

Recall that the rows of H_X represent stabilizer generators for the X stabilizers of the given code. The choice of stabilizer generators is not unique in general. In particular, we can replace any two generators $\mathcal{X}_i, \mathcal{X}_j$ by equivalent generators $\mathcal{X}_i, \mathcal{X}_i \times \mathcal{X}_j$. We say that two check matrices H_X and H'_X are equivalent if $H_X = R \times H'_X$ for some invertible $\mathbb{F}_2^{m \times m}$ matrix R .

The task of CSS state preparation circuit synthesis, then, is to find a CNOT circuit \mathcal{C} such that $H_X M_{\mathcal{C}} = R \times [I \ 0]$ or, equivalently $H_X = R [I \ 0] M_{\mathcal{C}}^{-1}$. Here, we assumed without loss of generality that the qubits initialized in $|+\rangle$ are the first m qubits in the qubit ordering. The second formulation of the synthesis problem suggests that instead of constructing the circuit by successively adding CNOTs until we obtain H_X , we can start from H_X and apply CNOTs until we obtain the *reduced* check matrix $[A \ 0]$, where A is some rank $\text{rk}(H_X) \times \text{rk}(H_X)$ matrix. The state preparation circuit is then obtained by initializing the first m qubits in the $|+\rangle$ state and applying the CNOTs used in the reduction in *reverse* order.

Complexity wise, generating optimal CNOT circuits is closely related to the minimum circuit size problem, which is unlikely to be proven to be in P or to be NP-complete [56]. Therefore, it is unlikely to find a polynomial algorithm for state preparation circuit synthesis since it is simply a variant of that problem. In the following, we describe how to convert an instance of the state preparation circuit synthesis problem to a SAT formula and obtain gate- and depth-optimal preparation circuits. Since this reduction to SAT is not expected to scale due to the intractability of the SAT problem, we also provide a polynomial-time heuristic algorithm for this problem.

B. Optimal approach

We would like to find an optimal preparation circuit whenever possible. Fewer CNOTs naturally lead to lower error rates in any circuit, and shallower circuits lead to faster execution time, which in turn leads to a decrease in idling errors. Therefore, we want to synthesize *gate-* and *depth-*optimal circuits.

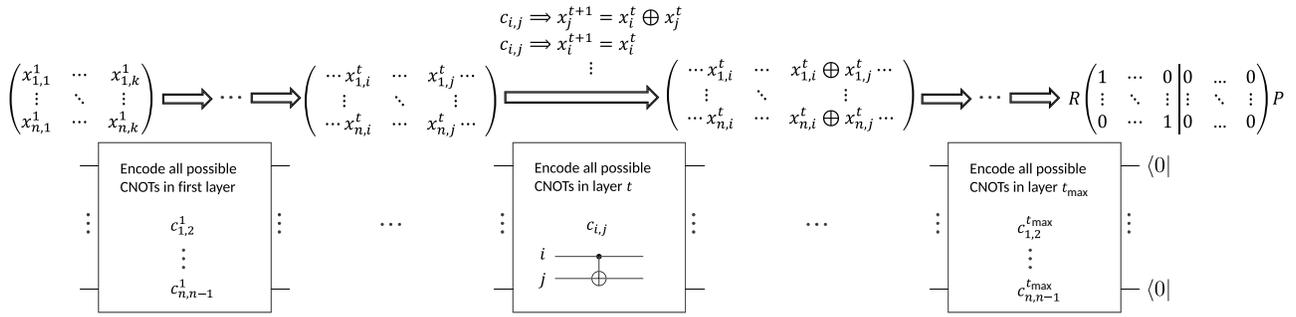


FIG. 5. Symbolic encoding for optimal state preparation circuit synthesis. The circuit is synthesized using at most t_{\max} layers of CNOTs. All possible column additions on the symbolic check matrices are related via variables $c_{i,j}^t$, which encode all possible CNOTs in layer t . Starting from the target check matrix H_X these constraints encode all possible valid circuits that prepare the logical $|0\rangle_L^{\otimes k}$ state.

Since the synthesis problem is formulated over \mathbb{F}_2 , SAT solvers [57] are a natural choice for this task. SAT solvers have previously been employed to synthesize general stabilizer circuits [34] and state preparation circuits in particular [52]. However, these approaches provide a solution only for a few qubits. We can improve upon these SAT encodings by taking advantage of the fact that we are dealing with CSS codes.

The idea behind the encoding is sketched in Fig. 5. Essentially, the task is to symbolically encode Gaussian elimination over \mathbb{F}_2 . To this end, we first introduce *matrix variables* $\{x_{i,j}^t \mid i \in [m], j \in [n], t \in [t_{\max} + 1]\}$, where t_{\max} is the given maximum number of elimination steps. The maximum number of elimination steps allows us to control how many CNOTs are applied at each step. Depending on the specific optimization task, we can use t_{\max} to limit the total number of CNOTs or the depth of the circuit.

Initially, we assert

$$\forall i \in [m]. \forall j \in [n]. x_{i,j}^1 = H_X[i, j],$$

where $H_X[i, j]$ denotes the matrix entry in row i and column j of H_X .

For the reduced matrix, we assert the pseudo-Boolean constraint

$$\left(\sum_{j=1}^n \prod_{i=1}^m x_{i,j}^{t_{\max}+1} \right) = \text{rk}(H_X).$$

In other words, the matrix must be reduced to a matrix $H_X^{t_{\max}+1}$ with exactly $\text{rk}(H_X)$ columns with nonzero entries. This is equivalent to asserting that there exists an invertible matrix R and a permutation matrix P such that $RH_X^{t_{\max}+1} = P[l_m \ 0]$. This constraint guarantees that the choice of stabilizer generators has no impact on the synthesized circuit's size.

Symbolic matrices in consecutive steps are obtained via column operations. Either a column is added onto another,

or a column does not change. Therefore, we introduce two Boolean formulas

$$\begin{aligned} \text{add}(i, j, t) &= \forall k \in [m]. x_{k,j}^{t+1} = x_{k,i}^t \oplus x_{k,j}^t \\ \text{id}(i, t) &= \forall k \in [m]. x_{k,i}^{t+1} = x_{k,i}^t. \end{aligned}$$

No matter which of the two metrics we optimize for, obtaining the optimal state preparation circuit is achieved by solving the encoded formula for different values of t_{\max} until two values t_1, t_2 are found such that

- (a) the formula is unsatisfiable for $t_{\max} = t_1$ and
- (b) the formula is satisfiable for $t_{\max} = t_2$.

The optimal solution is then the one obtained from the variable assignment to the satisfiable instance.

In the following, we will give the details of the encoding for synthesizing depth- and gate-optimal circuits.

1. Encoding for depth-optimal circuits

For synthesizing depth-optimal circuits, we interpret t_{\max} as the maximal circuit depth. We then need only to encode all possible layers of CNOTs at each reduction step. If the solver finds a satisfying assignment, we can extract the circuit directly from the variable assignment. If the solver shows no such assignment exists, we can incrementally vary t_{\max} until a satisfying assignment is found.

We introduce CNOT *variables* $\{c_{i,j}^t \mid i \in [n], j \in [n], t \in [t_{\max} + 1]\}$, which encode all possible CNOTs (or column additions) from control qubit (column) i to target qubit (column) j . These variables encode how the symbolic matrices at steps t and $t + 1$ are related:

$$\begin{aligned} \forall t \in [t_{\max}]. \forall i \in [n]. \forall j \in [n] \setminus \{i\}. \\ c_{i,j}^t \implies \text{add}(i, j, t). \end{aligned} \quad (2)$$

Additionally, any qubit can be involved in at most one CNOT at one timestep:

$$\forall t \in [t_{\max}]. \forall i \in [n]. \left(\sum_{j=1}^n \llbracket c_{i,j}^t \rrbracket + \sum_{j=1}^n \llbracket c_{j,i}^t \rrbracket \right) \leq 1.$$

Lastly, if a qubit is not the target of any CNOT, the corresponding column does not change:

$$\forall t \in [t_{\max}]. \forall i \in [n]. \left(\neg \bigwedge_{j=1}^n c_{j,i}^t \right) \implies \text{id}(i, t).$$

2. Encoding for gate-optimal circuits

For the synthesis of gate-optimal circuits, instead of encoding entire layers of CNOTs, symbolic matrices in consecutive elimination steps are related by only a single column addition, ensuring that at every time step, only one CNOT can be performed. The maximum number of time steps, then, corresponds to the maximum number of CNOTs allowed in the circuit construction. This can be achieved by augmenting the depth-optimal encoding with further constraints prohibiting more than one CNOT occurring at every step as done in Ref. [52]. Alternatively, these additional constraints can be avoided by encoding all possible control and target combinations for each CNOT instead of every possible CNOT. Using binary encoding for the control and target qubit indices, we require only $2\lceil \log_2 n \rceil$ instead of $n(n-1)$ variables per timestep.

Thus, we introduce *bit-vector* variables $\{\text{ctrl}^t \mid 1 \leq t \leq t_{\max}\}$ and $\{\text{tar}^t \mid 1 \leq t \leq t_{\max}\}$ of size $\lceil \log_2 n \rceil$. Since bit-vector logic is not supported by SAT solvers, the following formulation requires using a solver for *satisfiability modulo theories* (SMTs) that supports bit-vector logic. We can then encode CNOT *constraints* relating the control and target variables:

$$c_{i,j}^t \Leftrightarrow (\text{ctrl}^t = i \wedge \text{tar}^t = j).$$

We can symbolically encode the column addition constraints using Eq. (2).

Encoding the fact that a column does not change if the corresponding qubit is not the target of a CNOT is simpler in this setting:

$$\forall t \in [t_{\max}]. \forall i \in [n]. \text{tar}^t \neq i \implies \text{id}(i, t).$$

Of course, a qubit cannot be both the target and control of a CNOT:

$$\forall t \in [t_{\max}]. \text{ctrl}^t \neq \text{tar}^t.$$

Since the bit vectors might take on values outside the allowed qubit range, we impose the following constraint:

$$\forall t \in [t_{\max}]. (1 \leq \text{ctrl}^t \leq n) \wedge (1 \leq \text{tar}^t \leq n).$$

With these constraints encoded, we can try to obtain a solution using publicly available SMT solvers like Z3 [58].

C. Heuristic approach

CNOT and stabilizer circuit synthesis have been well researched in the past, and various asymptotically optimal algorithms exist [47,48,55]. For small circuits, the constant overhead in these constructions leads to unacceptably large circuits.

In Refs. [50,59], the authors propose a greedy path-finding-based method for synthesizing linear circuits. For an input \mathbb{F}_2 matrix H they define the cost functions

$$h_{\text{sum}}(H) = \sum_{i,j} H[i,j]$$

$$h_{\text{prod}}(H) = \sum_i \log\left(\sum_j H[i,j]\right),$$

where addition is taken over \mathbb{R} rather than \mathbb{F}_2 .

The first cost function simply counts the number of nonzero entries in the matrix. The idea behind the second cost function is that it weights columns that are almost completely zeroed higher.

We can use these cost functions to synthesize state preparation circuits by starting with the target check matrix H_X and greedily choosing the CNOT that minimizes the cost of the matrix resulting from applying the CNOT, i.e., the corresponding column addition. This is repeated until the cost can no longer be minimized, either because the reduced check matrix corresponds to a product state of $|0\rangle$ and $|+\rangle$ states or because the search is stuck in a local minimum. In Ref. [50], an iteration limit is set before stopping the calculation to avoid getting stuck in a local minimum indefinitely. In our case, we can use the fact that the stabilizers form a group by changing the basis of the row space to escape local minima when the synthesis algorithm is stuck. To this end, we can also greedily pick the row additions that remove the most ones [60]. If even this approach fails, we can always fall back to doing Gaussian elimination on the rows of the matrix. We make further use of this degree of freedom by preprocessing the matrix. The synthesized circuit will probably be quite deep if a matrix column has substantially more nonzero entries than other columns. This can be prevented by applying row operations that decrease the maximum column weight without decreasing the row weights.

References [50,59] also propose to use the combined cost of H and H^{-1} to guide the search. Since the matrices, in our case, are not invertible, this is not possible. Note that reducing the check matrix using such heuristics is a similar idea to the approach proposed in Ref. [42] where the ‘‘overlap’’ of stabilizers is used to reduce the CNOT count in state preparation circuits.

This greedy search strategy seems more promising for the smaller codes we focus on. Using the described cost functions for a heuristic search will optimize for gates, but we also want to optimize for depth as much as possible. A simple heuristic for this is to construct the circuit layer by layer, always greedily choosing CNOTs that minimize the cost function until all qubits are involved in a gate or until no possible CNOT would decrease cost anymore.

V. SYNTHESIS OF VERIFICATION CIRCUITS

State preparation circuits like the ones we synthesize above are generally not fault-tolerant. The key idea to ensure fault tolerance is to append a verification circuit to the preparation circuit that indicates whether an error has propagated in the circuit and postselecting on the verification measurement outcomes. More specifically, a verification circuit is the circuit implementation of a specific set of stabilizer or logical operator measurements of the state such that any propagated error anticommutes with one of the verification measurements. An error anticommutes with a stabilizer measurement if their support overlaps on an *odd* number of qubits. Thus, the propagated errors of the state preparation circuit impose a parity constraint on the measurements used in the verification circuit.

In the following, we will describe the problem of synthesizing verification circuits in detail. In particular, we describe the challenges that arise when generalizing verification circuit synthesis to codes with distances larger than 3. We show how the verification circuit synthesis problem can be broken down into synthesizing a sequence of verification circuits that verify a specific set of errors—significantly reducing the complexity of the problem.

To synthesize each sub-circuit, we show how to encode the synthesis problem into a SAT formula. Iterative queries to a SAT solver can then again be used to find the gate- or ancilla-optimal subverification circuit. As with the optimal synthesis of state preparation circuits, there will be limits to scaling such an optimal synthesis method. We, therefore, also propose a scalable, greedy synthesis scheme.

A. Problem description and complexity

A simple way to synthesize a verification circuit is to measure all stabilizer generators, which might include many superfluous measurements. Note again that since we prepare logical basis states, the respective logical operators are also stabilizers of the state.

Define the *i*-fault set $\mathcal{E}_i(\mathcal{C})$ of a circuit \mathcal{C} as the set of propagated errors E with $\text{wt}(E) > i$ that are caused by the propagation of an error of probability order i . Computing this set itself takes $O(\#\text{CNOTs}^i)$ time and is therefore fixed-parameter tractable in the number of CNOTs. Since we will

mainly look at small codes with a small number of CNOTs, computing this set is straightforward.

Let us continue to focus on the logical $|0\rangle_L^{\otimes k}$ state. Since we are looking at CSS codes, X and Z errors can be checked separately. Let us focus on verifying X errors for now, i.e., the verification circuit construction described in the following will only be strictly fault-tolerant against X -type errors. How this construction can be extended to be strictly fault-tolerant for general errors will be shown in Sec. VI. To check for X errors, any element of S_Z (including logical operators) can be measured in general. Searching for a minimally sized set of measurements seems infeasible considering the size of this group. This is stated by the following theorem.

Theorem 1. Consider the following decision problem: Given an X fault set \mathcal{E} , Z -stabilizer generators Z_1, \dots, Z_m and an integer k , is there a set of group elements $Z_{\text{ver},1}, \dots, Z_{\text{ver},t}$ such that $\sum_{i=1}^t \text{wt}(Z_{\text{ver},i}) \leq k$ and for all $E \in \mathcal{E}$ there is a $Z_{\text{ver},i}$ such that $E \times Z_{\text{ver},i} = -Z_{\text{ver},i} \times E$? This problem is NP-complete.

Proof. The problem is clearly in NP, since given a certificate for an instance it is straightforward to verify that it is a valid solution. Furthermore, this is just a more general version of the Coset leader problem, which asks for a minimum-weight representative of the coset of some linear code. Computing such minimal-weight representatives is NP-complete [61] and is essentially a decoding problem. ■

Consider an $[[n, k, d]]$ code, which protects against up to $t = \lfloor (d-1)/2 \rfloor$ errors and a state preparation circuit \mathcal{C} for the $|0\rangle_L^{\otimes k}$ of this code. Generating a verification circuit for all X errors $\bigcup_{i=1}^{\lfloor (d-1)/2 \rfloor} \mathcal{E}_i(\mathcal{C})$ does not ensure fault tolerance in and of itself because errors can also occur in the verification circuit itself. We distinguish the following cases.

- (1) t errors occur in the non-fault-tolerant state preparation circuit and propagate to an *uncorrectable* error $E_C \in \mathcal{E}_t$ at the output.
- (2) $i < t$ errors occur in the non-fault-tolerant state preparation circuit and propagate to a *correctable* error $E_C \in \mathcal{E}_i, E_C \notin \mathcal{E}_t$ at the output.
- (3) $i < t$ errors occur in the non-fault-tolerant state preparation circuit and propagate to an *uncorrectable* error $E_C \in \mathcal{E}_i \cap \mathcal{E}_t$ at the output.

The first case will always be detected by the verification circuit, assuming that at most t errors occur in the entire state preparation circuit. In the second case, if another error E_{ver} of weight $\text{wt}(E_{\text{ver}}) = t - i$ occurs in the verification circuit, then it could be that $E_C \times E_{\text{ver}} \notin \mathcal{E}_i$. If this error is not in the fault set \mathcal{E}_t , then this error is equivalent to some nonpropagated error. If it is in the fault set,

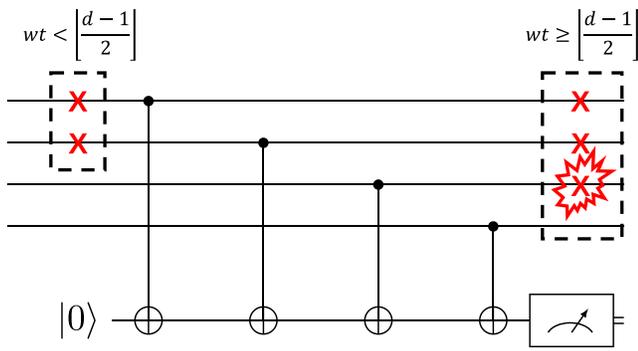


FIG. 6. A correctable error is (correctly) not detected by the verification circuit, but an error after the stabilizer measurement turns the error into an uncorrectable one. If this happens after the last measurement, which would detect this error, it goes by undetected.

however, it should be detected by some measurement in the verification circuit. But it could be that E_{ver} happens after every measurement that would detect this combined error. Then, the state will not be discarded, and the fault-tolerance condition of Definition 1 is violated (see also Fig. 6).

In the third case, the verification circuit would detect the error, but if no more than $t - i$ measurements detect this error, then all these measurements could be flipped with at most $t - i$ independent errors. We say that the error is *masked*. The verification circuit would then indicate no error, and Definition 1 is violated again (see Fig. 7).

For distance 3 codes for which fault-tolerant verification circuit-based state preparation schemes exist, this is not much of an issue since in cases 2 and 3, the state preparation circuit would be error free, i.e., $i = 0$. Any weight w X error in the verification circuit can lead to, at most, a weight w X error on the data qubits since the CNOTs of the measurement circuit only propagate X errors onto

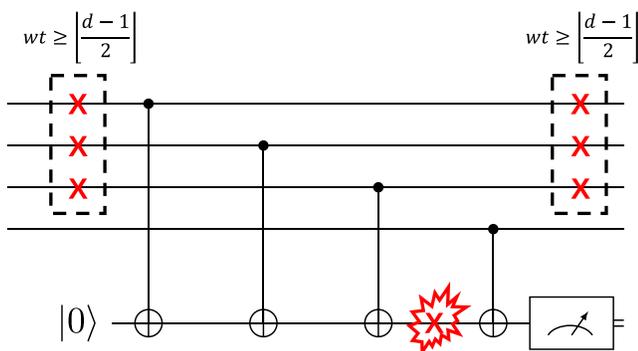


FIG. 7. An uncorrectable error would be detected by the verification circuit, but an error on the measurement ancilla flips the X parity, causing the measurement to falsely indicate no error (masked). If this happens after the last measurement that would detect this error, it goes by undetected.

ancilla qubits, so case 2 is unproblematic. Even an error on a measurement ancilla will only trigger one of the verification measurements, leading to the state being discarded. For distance-3 codes, the verification circuit can, therefore, be treated as error free without penalty.

Greater care must be put into synthesizing the verification circuit for larger distances. The case distinction above already suggests the solution. The verification circuit needs to detect not only uncorrectable propagated errors but correctable ones, too, as these might lead to a logical error together with additional errors in the rest of the circuit. Let ver_i be a verification circuit for \mathcal{E}_i . We achieve fault tolerance against one type of error (X or Z) with the verification circuit

$$\text{ver} = \text{ver}_t \circ \text{ver}_{t-1} \circ \dots \circ \text{ver}_1, \quad (3)$$

where \circ denotes circuit composition.

As before, if any measurement in this circuit indicates an error, the state is discarded, and the process is started anew. This ensures fault tolerance by the following reasoning.

Intuitively, if a layer in the verification circuit fails to detect a propagated error due to errors in the verification circuit itself, the resulting error will either be detected by a later verification or the error is equivalent to an unpropagated error.

More precisely, let E be a propagated error with $\text{wt}(E) > i$ caused by $i < t$ errors in the state preparation circuit. By construction, E will be detected by ver_i . Observe that a single-qubit X error on a qubit in the $|+\rangle$ state at the beginning of the circuit will propagate to a stabilizer operator of the code and act trivially on the prepared state. As a result, if $E \in \mathcal{E}_i(\mathcal{C})$ then

$$E \in \bigcap_{j=i}^{\text{wt}(E)} \mathcal{E}_j(\mathcal{C}),$$

meaning that if E is the result of i propagated errors in the state preparation circuit, it could also have been the result of $i + 1, \dots, \text{wt}(E)$ errors (the additional errors would have been trivial). This means that E will also be detected by each $\text{ver}_{i+1}, \dots, \text{ver}_{\text{wt}(E)}$.

The only way E can pass verification undetected, then, is that further errors in the verification circuit mask it. Any such error is either a single-qubit error on the data qubits or a single-qubit error on an ancilla qubit. A two-qubit X error on a CNOT is equivalent to a single-qubit X error on the control before the CNOT, so we do not have to give special consideration to this case.

From the perspective of ver_{i+1} , it is immaterial whether an error on the data qubits occurs in some previous verification layer or at the end of the state preparation circuit. This is because it does not matter if the combined error is detected by previous verification layers or not, as it

will definitely be detected by ver_{i+1} . Generally for an error E' resulting from u independent single-qubit errors on data qubits such that $\text{wt}(E \times E') > i + u$ it is the case that $E \cdot E' \in \mathcal{E}_{i+u}(\mathcal{C})$ and thus $E \times E'$ will be detected by ver_{i+u} .

Similarly, an error on an ancilla qubit is indistinguishable from a trivial error on the data qubits when considering the verification circuit ver_{i+1} . Again, if an error E' resulting from v independent errors on ancilla qubits occurs and $\text{wt}(E) > i + v$, then $E \in \mathcal{E}_{i+v}(\mathcal{C})$ and thus $E \times E'$ will be detected by ver_{i+v} .

Combining the two cases, for an error E' resulting from u independent single-qubit errors occur on data qubits and v single-qubit errors on ancilla qubits in the verification circuit such that $i + u + v \leq t$ and $\text{wt}(E \times E') > i + u + v$ then $E \times E' \in \mathcal{E}_{i+u+v}(\mathcal{C})$ and thus $E \cdot E'$ will be detected by ver_{i+u+v} . Therefore, the circuit ver is fault-tolerant.

The layered construction of the verification circuit simplifies the problem of verification circuit synthesis significantly. Instead of synthesizing the verification circuit wholesale, every layer can be constructed individually. Every \mathcal{E}_i gives rise to a new instance of the verification-circuit synthesis problem. For every \mathcal{E}_i we have to find a set of stabilizer measurements such that every error in \mathcal{E}_i anticommutes with at least one of these measurements.

This layered approach significantly simplifies reasoning about verification circuit constructions by breaking the synthesis down into the synthesis of smaller subproblems. This, however, potentially comes at the price of the ‘‘global’’ optimality of the full circuit. Since the different fault sets of the preparation and verification circuits are not necessarily independent, synthesizing the verification circuits separately might lead to certain errors redundantly being verified multiple times. Hence, in this work, we only consider optimal constructions for each sub-circuit.

Next, we will show optimal and heuristic algorithms to tackle these individual synthesis problems.

B. Optimal approach

We can frame the optimal verification-circuit synthesis problem again as a SAT problem with the following inputs:

- \mathcal{C} : a state preparation circuit over n_q qubits.
- $H_X \in \mathbb{F}_2^{m \times n}$: matrix representation of m X stabilizer generators of \mathcal{C} .
- $H_Z \in \mathbb{F}_2^{(n-m) \times n}$: matrix representation of $n - m$ Z stabilizer generators of \mathcal{C} .
- n_{anc} : maximum number of ancillas (measurements) used in the verification circuit.
- n_{CNOT} : maximum number of CNOTs used in the verification circuit.

Let $\mathcal{E}(\mathcal{C})$ denote the set of uncorrectable errors arising from propagating up to $\lfloor d/2 \rfloor$ independent X errors and let E_i

denote the \mathbb{F}_2^n representation of the i th error in $\mathcal{E}(\mathcal{C})$. The Z generators $\mathcal{Z} = \{Z_1, \dots, Z_{n-m}\}$ form a basis of the Z stabilizers, so we can express any verification circuit as a set of \mathbb{F}_2^{n-m} vectors indicating the contributing generators. We introduce the Boolean variables

$$\{v_s^i \mid i \in [n_{\text{anc}}], s \in [n - m]\}.$$

An assignment to these variables denotes the i th measurement as a linear combination of the $n - m$ Z -stabilizer generators.

The j th entry of the i th measurement can be symbolically computed via

$$\text{meas}_j^i = \left(\bigoplus_{s=1}^{n-m} v_s^i \wedge H_Z[s, j] \right).$$

Assuming that H_Z is sparse, it is inefficient to convert the Z stabilizers to Boolean constants and then perform these nested XOR operations. Since $a \oplus 0 = a$, $a \wedge 1 = a$ and $a \wedge 0 = 0$, we can simplify the computation by evaluating constants in Boolean functions before encoding the SMT formula. Such simplifications have been shown to lead to dramatic speedups in SMT solving [33].

Next, we need to ensure that every error is detected. Recall that an error is detected if it overlaps with a measurement on an odd number of qubits. Instead of computing the overlap on all qubits, we can again simplify this constraint by only considering qubits where an error has support. Every error E_k gives rise to a function

$$\text{detected}^{(k,i)} = \left(\bigoplus_{j \in \text{supp}(E_k)} \text{meas}_j^i \right).$$

Every error needs to be detected by at least one measurement, so we impose

$$\bigwedge_{k=1}^{|\mathcal{E}(\mathcal{C})|} \left(\bigvee_{i=1}^{n_{\text{anc}}} \text{detected}^{(k,i)} \right).$$

We can furthermore restrict the overall number of CNOTs used in the verification circuit with the pseudo-Boolean constraint:

$$\left(\sum_{i=1}^{n_{\text{anc}}} \left[\sum_{j=1}^{n_q} \llbracket \text{meas}_j^i \rrbracket \right] \right) \leq n_{\text{CNOT}}.$$

Since this encoding takes the maximal number of ancillae and the maximal number of CNOTs as parameters, we can iterate on both parameters to find either ancilla- or gate-optimal verification circuits.

C. Heuristic approach

A part of the difficulty of computing verification circuits is the size of the stabilizer group from which measurements need to be chosen. Instead of searching through the entire group, we can try to find a suitably large set of *candidate* measurements and try to find a good subset as a verification circuit. Indeed, if more than one measurement is required to verify a given fault set, the main difficulty is identifying a small set of verification measurements. The requirement of finding a small set of measurements such that every error is detected by at least one measurement is similar to the SET COVER [62] problem, which is defined as follows:

SET COVER

Input: A set U , a collection of subsets $P \subseteq \mathcal{P}(U)$ and an integer u

Problem: Is there a set $\text{Cov} \subseteq P$ with $|\text{Cov}| \leq u$ and $\bigcup_{A \in \text{Cov}} A = U$

The difference in our setting is that the number of potential stabilizers (covering sets) might be very large. We therefore consider this modified version of SET COVER:

SET COVER WITH SYMMETRIC DIFFERENCE

Input: A set U , a collection of subsets $P \subseteq \mathcal{P}(U)$ and an integers u and v .

Problem: Define $P_1 = P$ and $P_k = P_{k-1} \cup \{A \Delta B \mid A, B \in P_{k-1}\}$, where Δ denotes the *symmetric difference* of sets. Find $\text{Cov} \subseteq P_v$ such that $\bigcup_{A \in \text{Cov}} A = U$ and $|\text{Cov}| \leq u$.

By identifying $S \in \mathcal{Z}$ with the set $E(S, \mathcal{E}) = \{F \in \mathcal{E} \mid S \times F = -F \times S\}$, for two stabilizers $S_1, S_2 \in \mathcal{Z}$ we have $E(S_1 \times S_2, \mathcal{E}) = E(S_1, \mathcal{E}) \Delta E(S_2, \mathcal{E})$, i.e., any error in \mathcal{E} that anticommutes with both S_1 and S_2 will not anticommute with their product and therefore will not be detected by the respective measurement. With this translation, we can phrase an instance of verification circuit synthesis as an instance of the above-defined modified SET COVER problem.

The straight-forward greedy algorithm for SET COVER approximates the optimum quite well [63], so we can use a similar strategy for synthesizing a verification circuit. The idea is to start with the stabilizer generators and iteratively expand the set of candidate measurements by building products of stabilizers in the candidate set. In each iteration, we can try to find a small set of stabilizers such that all problematic errors are detected. If increasing the candidate set does not lead to a verification circuit with fewer stabilizer measurements, we stop. Alternatively, if the size of the candidate set reaches some predefined limit, we can take the best solution found until this point and terminate.

This greedy algorithm is formalized in Algorithm 1. This algorithm assumes a function $\text{set_cover}(U, P)$ that computes a solution to the SET COVER problem for the

ALGORITHM 1. Heuristic verification circuit

Data: Set of errors \mathcal{E} , Z -stabilizer generators \mathcal{Z} , maximal candidate set size c_{\max}

Result: Set of stabilizer measurements \mathcal{Z}_{ver} that detect \mathcal{E}

- 1 $P \leftarrow \{E(S, \mathcal{E}) \mid S \in \mathcal{Z}\};$
- 2 $\text{Cov} \leftarrow \text{set_cover}(P, \mathcal{E});$
- 3 **while** $|P| \leq c_{\max}$ **do**
- 4 $\text{Cov}_{\text{new}} \leftarrow \text{set_cover}(P, E(S));$
- 5 **if** $|\text{Cov}_{\text{new}}| = \text{Cov}$ **then**
- 6 **break**;
- 7 $P \leftarrow P \cup \{A \Delta B \mid A, B \in P\};$
- 8 $\text{Cov} \leftarrow \text{Cov}_{\text{new}};$
- 9 $\mathcal{Z}_{\text{ver}} \leftarrow$ Stabilizers corresponding to $\text{Cov};$
- 10 **return** $\mathcal{Z}_{\text{ver}};$

sets U and $P \subseteq \mathcal{P}$, which can be done in polynomial time using the greedy algorithm.

Let us make two remarks about the implementation of Algorithm 1:

- (a) When solving SET COVER in lines 2 and 4 there might be multiple choices for every greedy choice of set. In this case, we preferably pick the set that corresponds to lower-weight stabilizers.
- (b) Extracting the stabilizers from the set cover in line 9 can be done efficiently by maintaining a map from sets to stabilizer products.

In general, the candidate set can become large very quickly as, in the worst case, it doubles in size during every iteration of the while loop. For this reason, Algorithm 1 takes the parameter c_{\max} to limit the growth of this set.

By design, Algorithm 1 only aims at minimizing the number of measurements in the verification circuit and not the weight of those measurements. This approach is appropriate when minimizing the overall number of measurements but might result in a few high-weight stabilizer measurements. This problem can be somewhat mitigated by completely removing high-weight measurements from the set of measurement candidates. However, a few high-weight measurements could actually be less costly in terms of the number of CNOTs than many low-weight measurements. The preferred choice might, for example, depend on the noise characteristics of the respective hardware under consideration.

After an appropriate set of stabilizers has been found using Algorithm 1, we can try to minimize the weight of these stabilizers. If there are stabilizers $\mathcal{Z}_{\text{red}} = \{S \in \langle \mathcal{Z} \rangle \mid \mathcal{E}(S, \mathcal{E}) = \emptyset\}$, i.e., stabilizers that commute with all errors, then they form a group $\langle \mathcal{Z}_{\text{red}} \rangle$. Given a verification stabilizer S_{ver} and any $S_{\text{red}} \in \langle \mathcal{Z}_{\text{red}} \rangle$ we can replace S_{ver} by $S_{\text{ver}} \times S_{\text{red}}$ in the verification circuit without changing which errors are detected. Another way to phrase this

is to say that replacing S_{ver} by any element in the coset $S_{\text{ver}} \times \langle Z_{\text{red}} \rangle$ yields a valid verification circuit. Give the measurements output by Algorithm 1 we can therefore try to find the minimal representatives in their respective cosets, i.e., the so-called coset leader. As noted above, computing the minimal-weight representatives of these cosets is NP-hard [61]. Apart from constructing the standard array (i.e., listing all coset elements), there are methods using techniques from computer algebra [64,65], which still require exponential memory in the worst case [66]. As a postprocessing step, we can apply a heuristic search for the coset leaders of our verification stabilizers. One can also relegate this step to an SMT solver. For small codes, this can be expected to be efficient even if the global optimal verification circuit cannot be computed with the optimal method in an adequate amount of time.

VI. FULLY FAULT-TOLERANT NONDETERMINISTIC STATE PREPARATION

Until now, we have only considered verification circuits for one type of error, i.e., X errors in preparing the $|0\rangle_L^{\otimes k}$. But even though Z errors cannot cause a logical error on $|0\rangle_L^{\otimes k}$, propagating Z errors through the state preparation circuit still violates the fault-tolerant requirement of Definition 1. A straightforward fix to this is generating verification circuits ver_X and ver_Z according to Sec. V that check for X and Z errors, respectively, and combine them to the verification circuit $\text{ver} = \text{ver}_X \circ \text{ver}_Z$. This does not work however, because X (Z) errors on one of the ancillae in ver_X (ver_Z) can propagate to a higher weight error on the data qubits. This is commonly referred to as a *hook error* [44].

Unfortunately, there is no way to fix this predicament by introducing further stabilizer measurements since the new measurements can themselves be faulty. Instead, the hook errors have to be checked for specifically. Flag fault-tolerant stabilizer measurements have been introduced to solve this exact problem [41,67]. The idea is to add further “flag” ancillae that interact with the stabilizer measurement ancilla via a pair of CNOTs that—in the absence of errors—act as identity. However, if any error occurs on the measurement ancilla, it propagates to at least one of

the flags and is detected. See Fig. 8 for an example of a weight-4 flag fault-tolerant measurement.

Flag fault tolerance was conceived as a way to give decoders additional information to account for hook errors, which can lead to complicated decision trees for decoders depending on the exact pattern of flags that measure an error. We require no such finesse and opt to include the flag measurements as further measurements upon which we postselect. If any flag is measured in -1 , this indicates that an error was propagated from an ancilla to the data, and we discard the state.

Not all stabilizers necessarily have to be measured flag fault-tolerantly. If we measure ver_X first, we can account for any hook errors introduced by this circuit by adapting the measurements for ver_Z . Only ver_Z , then, has to be measured using flags. Depending on the code and state preparation circuit, reversing the order of verification circuits might be better in terms of CNOTs. Still, only the second verification circuit has to be flag fault-tolerant. With this flag-fault-tolerant measurement scheme for the second error type, we arrive at the proposed scheme sketched in Fig. 1.

Depending on the weight of the measurements in the first verification circuit, the number of additional hook errors might be pretty significant. These additional errors might necessitate complex measurements in the subsequent verification circuit. In that case, it might be better to also flag the first layer.

At this point, we have to revisit the issue of optimal circuits for the verification-circuit synthesis. The introduction of flag-fault-tolerant measurements complicates the gate optimality of the verification measurements even for the individual fault sets. First, no generally optimal flag scheme is known [41,67]. Secondly, there is also a trade-off between the size weight of the measured stabilizers and the size of the corresponding flag circuit. It might, therefore, be better to measure more low-weight stabilizers with low flag overhead instead of high-weight stabilizers with high flag overhead. Note that depending on the minimum-weight stabilizer generators of the given code, the use of flag qubits can not always be avoided.

We leave the problem of solving this global optimization problem for future work. In this work, we focus on providing locally optimal state preparation and verification circuits without considering flag overhead.

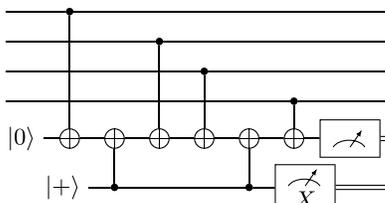


FIG. 8. Flag fault-tolerant measurement of a weight-4 stabilizer [41].

VII. NUMERICAL EVALUATIONS

The implementation of the methods presented in this work is publicly available as part of the Munich Quantum Toolkit [36,68]. To evaluate the efficacy of the proposed methods, we have synthesized state preparation circuits for various CSS codes.

A. Circuit construction

In this section, we focus on the investigation of circuit metrics.

1. Circuits for distance-3 and 5 codes

We generate nondeterministic fault-tolerant state preparation circuits for various $d = 3$ and $d = 5$ CSS codes using both the SMT and heuristic methods proposed in Secs. IV and V. This allows us to investigate the scalability of the exact approach as well as how close the heuristic approach is to the optimum. In addition to that, we also compare the results to solutions generated in Ref. [20] (as reviewed in Sec. III B, a synthesizer based on reinforcement learning). Note that on our machine the RL agent does not manage to find state preparation circuits within 24 h for codes with 12 or more physical qubits. In this case, we used the best circuit from the examples provided at <https://github.com/remmyzen/rlftqc>.

Tables I and II provide a summary of the circuits obtained for both gate and depth optimization [74], respectively. The tables list the CNOT count and depth of the non-fault-tolerant state preparation and verification circuits obtained from the proposed optimal and heuristic methods as well as the RL agent. More precisely, the first two columns denote the code instance and the prepared logical state, while the remaining columns denote the metrics of the obtained state preparation and verification circuits. “SAT” denotes the exact results, “heuristic” the heuristic results, and RL the solution obtained by the RL agent of Ref. [20].

For this comparison, we only focussed on verification circuits for the primary type of error (X errors for the

preparation of $|0\rangle_L^{\otimes k}$ and Z errors for the preparation of $|+\rangle_L^{\otimes k}$). Note that the number of measurements is equal to the number of additional ancilla qubits needed if qubits are not reused. Because hook errors are not considered for this comparison—and therefore, no additional flag qubits are needed—a single ancilla could be reused to perform all measurements.

These results show that for small codes, all methods produce similar results. For the cases where all methods produce the same circuit, however, the optimal method also yields an automatic proof that the obtained circuit is, in fact, gate or depth optimal. For larger codes ($n \geq 17$), both the optimal synthesis method and the RL method fail to yield a result within a 24-hour timeout. For the optimal synthesis methods specifically, the state preparation times out. The verification-circuit synthesis scales much more favorably. Moreover, the proposed approach clearly outperforms the RL synthesizer.

Interestingly, for the check matrices we investigated, optimizing for depth never comes at the cost of additional CNOT gates. This does not mean, however, that this cannot happen in general as one can construct check matrices where our implementation needs more CNOTs for the depth-optimal circuit.

For codes with more qubits ($n \geq 17$), the heuristic state preparation circuit starts to outperform the RL synthesis regarding CNOT count and circuit depth. For larger codes, the SMT and RL synthesis fail (indicated by “—” in the table) to produce any state preparation circuit within 24 hours. The heuristic synthesis, however, generates all circuits in under one second. Even in cases where the SMT solver cannot find an optimal state preparation circuit,

TABLE I. Circuit metrics for fault-tolerant state preparation circuits obtained using our optimal (SAT) and heuristic approach when optimizing for CNOT count, and comparison with the recently proposed RL method [20].

Code	State	State preparation circuit						Verification circuit					
		No. CNOTs			Depth			No. CNOTs			No. Measurements		
		SAT	Heuristic	RL	SAT	Heuristic	RL	SAT	Heuristic	RL	SAT	Heuristic	RL
[[15, 1, 3]] Reed-Muller (tetrahedral) code [69]	$ 0\rangle_L$	22	22	22	4	4	7	3	3	3	1	1	1
	$ +\rangle_L$	23	23	24	5	5	6	7	7	7	1	1	1
[[9, 1, 3]] Shor code [70]	$ 0\rangle_L$	8	8	8	5	6	5	3	3	3	1	1	1
	$ +\rangle_L$	6	6	6	2	2	2	0	0	0	0	0	0
[[9, 1, 3]] Rotated surface code [71]	$ 0\rangle_L$	8	8	8	5	6	3	3	4	3	1	1	1
[[15, 7, 3]] Hamming code [72]	$ 0\rangle_L$	—	22	—	—	12	—	6*	7	—	2	2	—
[[17, 1, 5]] 2D color code [73]	$ 0\rangle_L$	—	23	25	—	11	7	29*	32	—	6*	5	—
[[19, 1, 5]] 2D color code [73]	$ 0\rangle_L$	—	27	—	—	8	—	33*	38	—	6*	6	—
[[25, 1, 5]] Rotated surface code [71]	$ 0\rangle_L$	—	28	—	—	13	—	34*	34	—	5*	5	—

Entries marked with “*” denote circuits where the optimal verification circuit was generated for the heuristic state preparation circuit. Entries with a “—” denote instances where the respective method failed to provide a circuit within the timeframe of 24 h. Entries in bold denote instances where the obtained circuits using our methods yields improved metrics.

TABLE II. Circuit metrics for fault-tolerant state preparation circuits obtained using our optimal (SAT) and heuristic approach when optimizing for circuit depth, and comparison with the recently proposed RL method [20].

Code	State	State preparation circuit						Verification circuit					
		No. CNOTs			Depth			No. CNOTs			No. Measurements		
		SAT	Heuristic	RL	SAT	Heuristic	RL	SAT	Heuristic	RL	SAT	Heuristic	RL
[[15, 1, 3]] Reed-Muller (tetrahedral) code [69]	$ 0\rangle_L$	22	22	22	4	4	7	3	7	3	1	2	1
	$ +\rangle_L$	23	23	24	5	5	6	7	7	7	1	1	1
[[9, 1, 3]] Shor code [70]	$ 0\rangle_L$	8	9	8	3	3	5	3	5	3	1	2	1
	$ +\rangle_L$	6	6	6	2	2	2	0	0	0	0	0	0
[[9, 1, 3]] Rotated surface code [71]	$ 0\rangle_L$	8	8	8	3	3	3	3	4	3	1	1	1
[[15, 7, 3]] Hamming code [72]	$ 0\rangle_L$	23	22	—	4	4	—	8*	10	—	2*	2	—
[[17, 1, 5]] 2D color code [73]	$ 0\rangle_L$	—	23	25	—	5	7	28*	31	—	6*	5	—
[[19, 1, 5]] 2D color code [73]	$ 0\rangle_L$	—	27	—	—	5	—	33*	43	—	5*	6	—
[[25, 1, 5]] Rotated surface code [71]	$ 0\rangle_L$	—	28	—	—	5	—	35*	38	—	4*	5	—

Entries marked with “*” denote circuits where the optimal verification circuit was generated for the heuristic state preparation circuit. Entries with a “—” denote instances where the respective method failed to provide a circuit within the timeframe of 24 h. Entries in bold denote instances where the obtained circuits using our methods yields improved metrics.

it is still possible to find the optimal verification circuit (for the state preparation circuit generated by the heuristic synthesis). From these tables, we can conclude that synthesizing state preparation circuits using our greedy heuristic approach in combination with the optimal verification-circuit synthesis can produce competitive nondeterministic fault-tolerant state preparation circuits even for moderately sized codes.

Next, we use our methods to generate fully fault-tolerant circuits as described in Sec. VI. Again, we give the optimal synthesis method a 24-hour timeout. We use the heuristic method if the solver does not find a solution within this timeframe. All verification circuits were constructed gate-optimally for the $d = 3$ and $d = 5$ circuits. Some of the non-fault-tolerant state preparation circuits

were generated by the heuristic, either because the optimal synthesis method times out or because the heuristic circuit yielded a verification circuit that leads to an overall improved CNOT count. Table III shows circuit metrics for gate-optimized circuits found using this approach. This table shows explicitly how many additional ancillas are required to perform the verification circuit, the number of CNOTs needed for the entire circuit, and the circuit depth.

Table III furthermore lists the number of CNOTs required for the general scheme where the state is prepared through initialization of the data qubits in the all $|0\rangle$ ($|+\rangle$) state and subsequent projective measurements of the stabilizers. Note that this state preparation procedure only requires measuring one type of stabilizer since, for CSS codes,

TABLE III. Circuit metrics for the best circuits found using the proposed methods for the generalized nondeterministic fault-tolerant state preparation scheme.

Code	State	This work			Projective initialization		
		No. Ancillas w/o reuse	No. Ancillas with reuse	No. CNOTs	Depth	No. CNOTs worst case	No. CNOTs best case
[[15, 1, 3]] Reed-Muller (tetrahedral) code (see Fig. 16) [69]	$ 0\rangle_L$	1	1	25	8	96	64
	$ +\rangle_L$	5	2	42	14	120	80
[[9, 1, 3]] Rotated surface code [71]	$ 0\rangle_L$	1	1	11	6	36	24
[[12, 2, 4]] “Carbon” code (see Fig. 17) [18]	$ 0\rangle_L^2$	3	2	30	16	120	90
[[15, 7, 3]] Hamming code (see Fig. 18) [72]	$ 0\rangle_L^7$	2	1	28	11	96	64
[[17, 1, 5]] 2D color code (see Fig. 19) [73]	$ 0\rangle_L$	13	4	71	22	180	108
[[19, 1, 5]] 2D color code (see Fig. 20) [73]	$ 0\rangle_L$	14	4	120	25	210	126
[[25, 1, 5]] Rotated surface code [71]	$ 0\rangle_L$	25	4	127	31	200	120

the initial product state is already stabilized by all X or Z stabilizers of the code. Since measuring the stabilizers nondeterministically projects the state into the $+1$ or -1 eigenspace, multiple rounds of measurements are required to determine the correct sign of the stabilizers. For the codes considered here this requires d measurement rounds and in the best case $\lfloor d/2 \rfloor + 1$. Table III lists the CNOT count for both cases. The numbers for the projective initialization do not consider hook errors and flag qubits, which would add further CNOT gates to these circuits. Except for the distance-5 surface code, the proposed procedure requires fewer CNOTs than even the best-case scenario for the projective initialization (and even fewer for all cases if hook errors and flag qubits are taken into account) and significantly fewer CNOTs than the worst case. The benefit of projective initialization is, however, that the state is obtained deterministically, which may be favorable for larger code instances in general.

The circuits used for Table III are partially listed in Sec. B. All circuits listed in Table III are also publicly available [36,75]. An example of a fault-tolerant state preparation of $|0\rangle_L$ for a distance 5 code is depicted in Fig. 19. There, the verification circuit accounts for a large part of the entire circuit compared to the non-fault-tolerant state preparation sub-circuit.

2. Scaling to higher-distance codes

Although the optimal verification circuit synthesis approach was sufficient to synthesize circuits for the $d = 3$ and $d = 5$ codes considered in the previous section, naturally, there are limits to how far this method can be pushed. When scaling to higher-distance codes, the stabilizer groups, the number of qubits, and the size of the required state preparation circuits grow. In turn, the size of the fault sets that need verification grows rapidly. This directly impacts the SAT encoding since the size of the encoding grows as $O(|S| \cdot |\mathcal{E}|)$, where S is a generating set of the stabilizer group and \mathcal{E} is the set of errors to be verified. Given the layered construction of the verification circuit as illustrated in Fig. 1, and that the sizes

of the fault sets increase for later layers, the optimal synthesis algorithm can still generate the early layers of the verification circuit.

Consider the state preparation circuit for the zero state of the $[[31, 1, 7]]$ color code. A state preparation circuit with 46 CNOTs and depth 5 generated by the greedy algorithm is given in Fig. 21. To verify this state preparation circuit, we must construct three layers of X verification and three layers for the Z verification.

For one and two independent errors in the state preparation circuit, we obtain fault sets of sizes

$$\begin{aligned} |\mathcal{E}_1^X| &= 16 & |\mathcal{E}_2^X| &= 443 \\ |\mathcal{E}_1^Z| &= 14 & |\mathcal{E}_2^Z| &= 315, \end{aligned}$$

which are still small enough to attempt optimal synthesis for.

Table IV compares the verification circuits obtained for each fault set by the optimal synthesis, heuristic synthesis, and a naive verification circuit constructed by measuring all stabilizers of the code. Given the relatively high number of measurements in the X verification, we have opted to flag the measurements in both layers instead of adding potential hook errors to the fault sets for the Z verification. The table shows the stabilizer weights of the stabilizers measured in the verification circuit, the number of flag qubits needed to make these measurements flag-fault-tolerant, and the total number of CNOTs required to implement this circuit. In total, using the optimal verification circuit for the first two layers and the circuit obtained by the heuristic method for the third layer in the respective sub-circuits yields a verification circuit with 375 CNOTs, 119 ancilla qubits and depth 43. In comparison, the naive verification circuit has 720 CNOTs, 243 ancilla qubits, and a depth of 41. The lower depth comes from the fact that the measured stabilizers in the naive verification circuit generally have a lower weight than those in the synthesized measurements.

From Table IV we can see that quite a large overhead is incurred by flagging the stabilizer measurements. The weight-24 measurement the heuristic synthesis finds for \mathcal{E}_2^X is particularly bad in this regard. This suggests that

TABLE IV. Verification circuits for the $|0\rangle_L$ preparation circuit for the $[[31, 1, 7]]$ color code in Fig. 21.

Fault set	Optimal			Heuristic			Naive		
	Weights	No. Flags	No. CNOTs	Weights	No. Flags	No. CNOTs	Weights	No. Flags	No. CNOTs
\mathcal{E}_1^X	[7, 4, 7]	9	36	[12, 8]	10	40	$4 \times 11, 8 \times 3$	20	112
\mathcal{E}_2^X	[8, 7, 7, 7]	16	61	[11, 8, 7, 7]	18	69	$4 \times 11, 8 \times 3$	20	112
\mathcal{E}_3^X	—	—	—	[4, 8, 4, 4, 7, 8, 7]	19	80	$4 \times 11, 8 \times 3$	20	112
\mathcal{E}_1^Z	[12, 8]	10	40	[12, 12]	12	48	$4 \times 11, 8 \times 3$	20	112
\mathcal{E}_2^Z	[8, 8, 8, 8]	16	64	[12, 8, 4, 8, 24]	26	108	$4 \times 11, 8 \times 3$	20	112
\mathcal{E}_3^Z	—	—	—	[8, 4, 8, 8, 12, 8]	23	94	$4 \times 11, 8 \times 3$	20	112

for higher-distance codes with potential high-weight measurements, including a penalty for the stabilizer weight might yield better results. Even a relatively small increase in stabilizer weight could lead to a rise in the number of required flags and, thus, the overall number of CNOTs. For the $[[31, 1, 7]]$ color code, this effect is even more pronounced since the weight-4 measurements require only a single flag qubit and two CNOTs to flag. For example, combining 3 weight-4 measurements into one weight 12 measurements would require six flags instead of three and, therefore, 12 CNOTs instead of 6. As we focus on smaller codes that could be realized on near-term devices, we leave this avenue of investigation open for future work.

B. Logical error rate simulation

To show that the synthesized circuits indeed yield the expected logical error rates, we have simulated them using the open-source stabilizer simulator Stim [35] under a standard circuit-level depolarizing noise model as considered in previous works [15,41]. Noisy gates are modeled as ideal gates followed by a depolarizing noise channel parameterized by an error probability p . Since the considered circuits are composed only of two-qubit gates, we define the two-qubit depolarizing channel:

$$\epsilon_2(\rho) = (1 - p)\rho + \frac{p}{15} \sum_{E \in \mathcal{E}} E\rho E,$$

where $\mathcal{E} = \{P_1 \otimes P_2 \mid P_1, P_2 \in \{I, X, Y, Z\}\} \setminus \{I \otimes I\}$. In addition, there is a probability of $2p/3$ that a qubit will erroneously be initialized in the -1 eigenstate of the respective basis and measurements have a $2p/3$ probability of being flipped. If a qubit is idle during one layer of CNOTs, it is also subject to single-qubit depolarizing “idling” noise. What constitutes a layer of CNOTs depends on the computational model. We consider the two extreme cases: either all parallel CNOTs are executed in one layer, or only one CNOT can be performed in each layer. By default, we set $p_{\text{idle}} = p/100$, but since idle noise can have a significant impact on the performance of the state preparation circuits, we also consider other settings as described below.

Simulation runs for which any of the verification measurements indicate an error are discarded. For the remaining states, we calculate the logical error rate by measuring all data qubits in the Z basis, performing one round of ideal syndrome extraction and correction using a lookup-table decoder. A logical error is registered when the corrected classical bitstring anticommutes with any of the Z operators of the code. For each physical error rate p , we sample until we have collected 500 logical errors. For our simulations, this ensures that the uncertainties for the logical error rates are below 2% for all runs.

The results of the simulations with parallel gates and $p_{\text{idle}} = p/100$ are shown in Fig. 9. For all codes, we find that the observed error suppression is in agreement with the expected error rate relative to the distance d of the code [76]. All simulated distance $d = 3$ codes exhibit a $p_L \sim O(p^2)$ scaling of the logical error rate, the distance $d = 5$ codes exhibit a $p_L \sim O(p^3)$ scaling and the $[[31, 1, 7]]$ color code exhibits a $p_L \sim O(p^4)$ scaling. This shows that the circuits produced with our methods indeed generalize the postselection scheme previously known for distance $d = 3$ codes [10,12,14,20] to higher distances $d > 3$. Moreover, though the acceptance rates for larger codes are drastically worse for high physical error rates, this difference becomes less pronounced for the $d = 5$ codes, where all verification circuits accept the prepared state in $> 80\%$ of the cases at physical error rates of 10^{-3} .

To illustrate the performance of distance $d = 5$ state preparation more clearly, we compare the optimized state preparation circuits with a “naive” (canonical) verification circuits, where all stabilizers are measured twice. In addition, to visualize the impact of gate parallelism, we simulate the circuits with parallel gates and sequential gates. The respective results for the $[[17, 1, 5]]$ color code can be seen in Figs. 10 and 11. Similar plots for the $[[19, 1, 5]]$ color code and the $[[25, 1, 5]]$ rotated surface code are shown in Sec. A.

These results show that the difference between logical error rates of the optimized versus the naive scheme decreases with higher idling noise. This is reasonable since, depthwise, there is not much difference between the two circuits. In fact, the naive scheme has fewer idling qubits if CNOTs are executed in parallel and, therefore, if the idling noise is in the same order as the two-qubit gate noise, this behavior is expected. The story is a bit different when looking at acceptance rates. With $p > 10^{-3}$ the difference between the acceptance rates of the naive and optimized schemes is quite stark. For example, at $p = 4 \times 10^{-3}$ and $p_{\text{idle}} = p/100$, the optimized circuit has an acceptance rate of about 70%, whereas the naive verification circuit gives an acceptance rate of less than 20%.

When gates are executed sequentially, both circuits are subject to about the same amount of idling noise. In this case, the optimized circuit outperforms the naive circuit by orders of magnitude even for comparatively low idling noise.

These results show that optimized state preparation schemes are particularly suited for quantum computing platforms with a lower degree of parallelism. For architectures like neutral atoms [17,77] where two-qubit gates can be performed globally, the difference is less stark, assuming idling errors are manageable. For these systems, alternative schemes with a high degree of parallelism, like in the preparation of the Golay code [42], could be exploited.

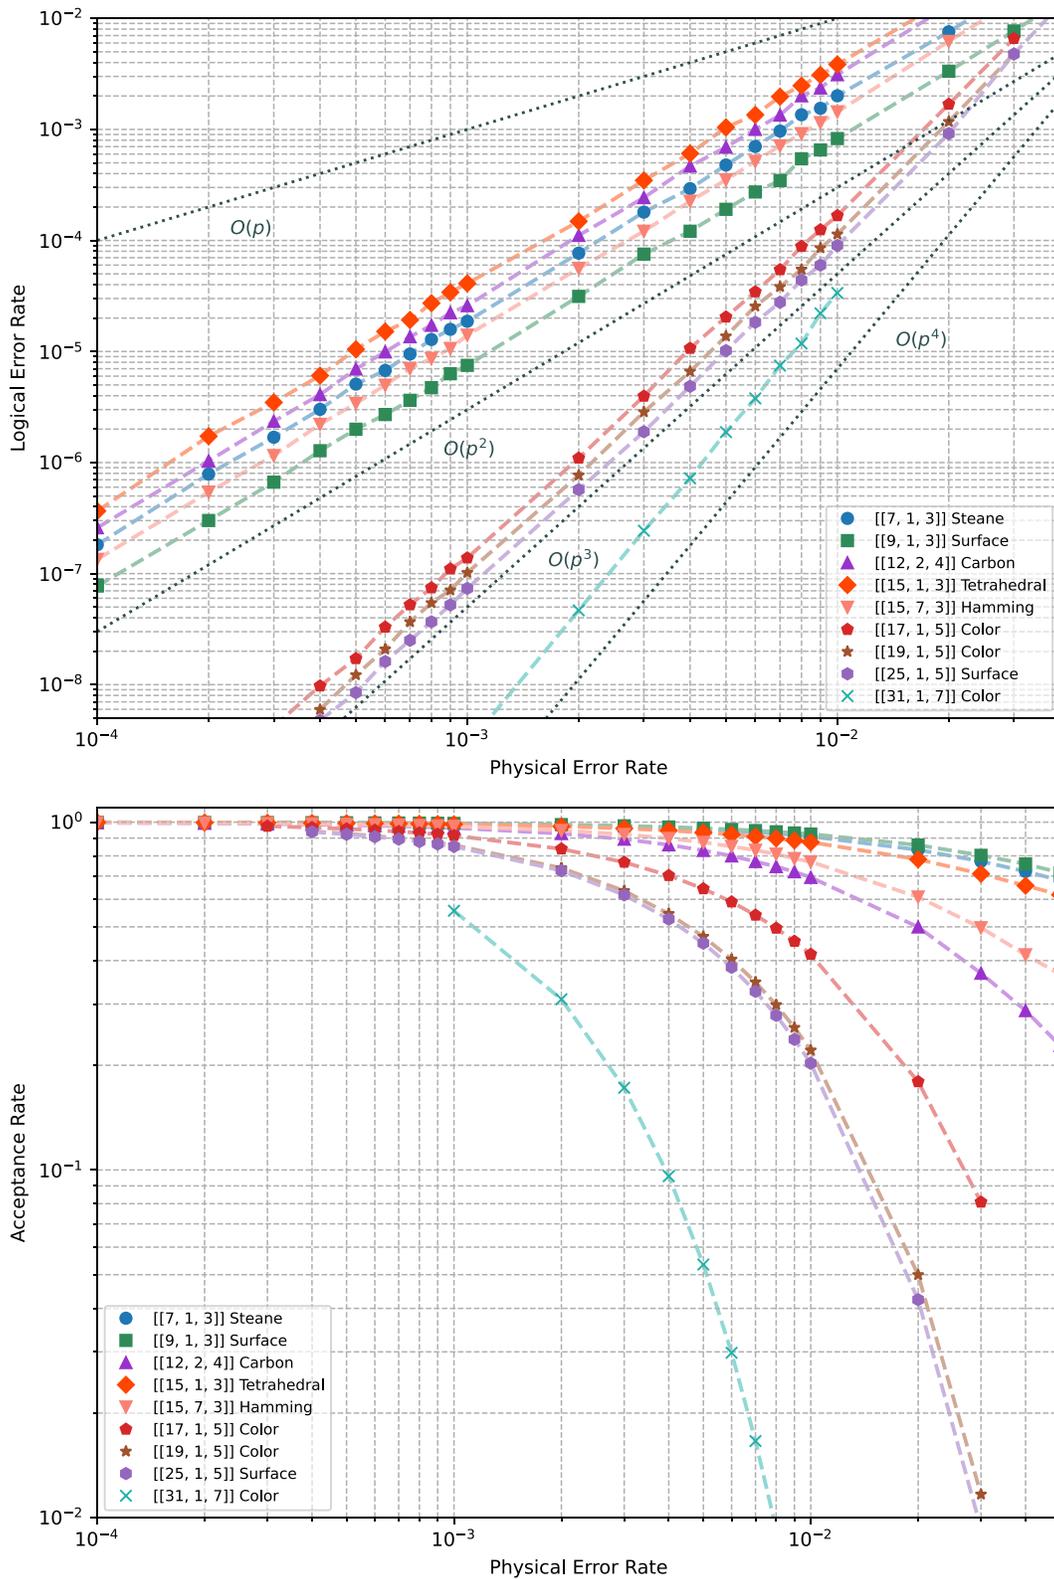


FIG. 9. Logical error rate and acceptance rate for nondeterministic fault-tolerant state preparation circuits for the logical $|0\rangle_L^{\otimes k}$ state constructed with our methods for various $d = 3, d = 5,$ and $d = 7$ CSS codes under circuit-level depolarizing noise ($p_{\text{idle}} = p/100$) using a LUT decoder.

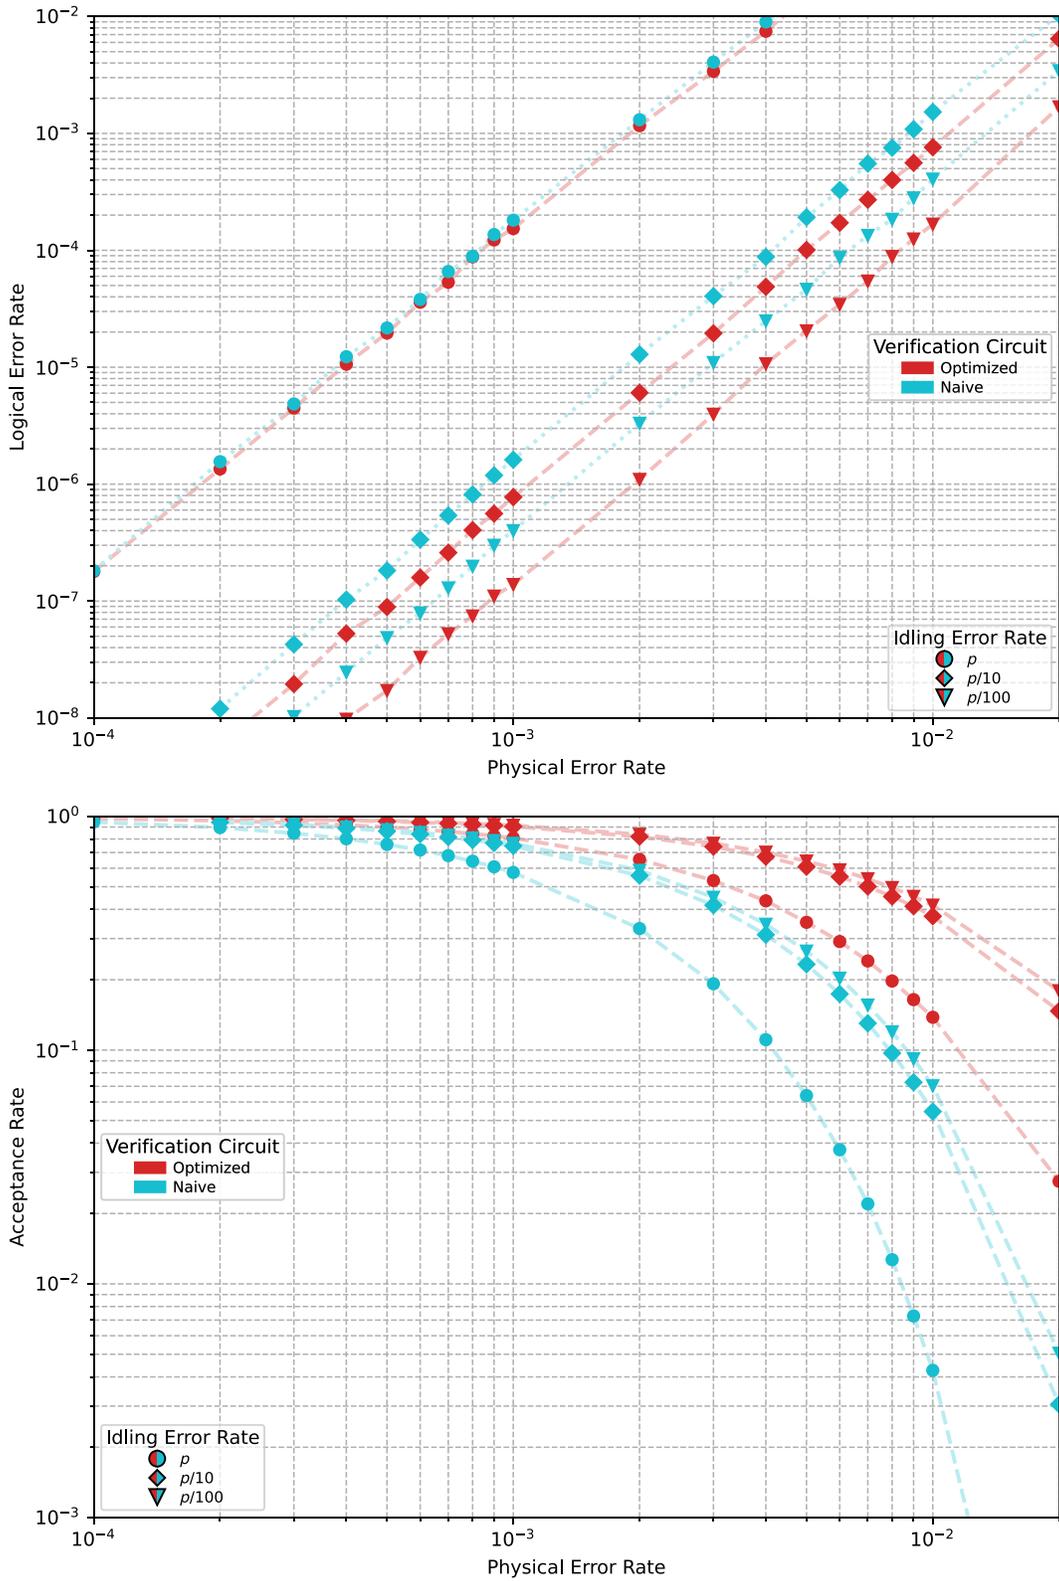


FIG. 10. Logical error rate and acceptance rate for nondeterministic fault-tolerant state preparation circuits for the logical $|0\rangle_L$ state of the $[[17, 1, 5]]$ color code constructed with our methods under circuit-level depolarizing noise with *parallel* gate execution using a LUT decoder.

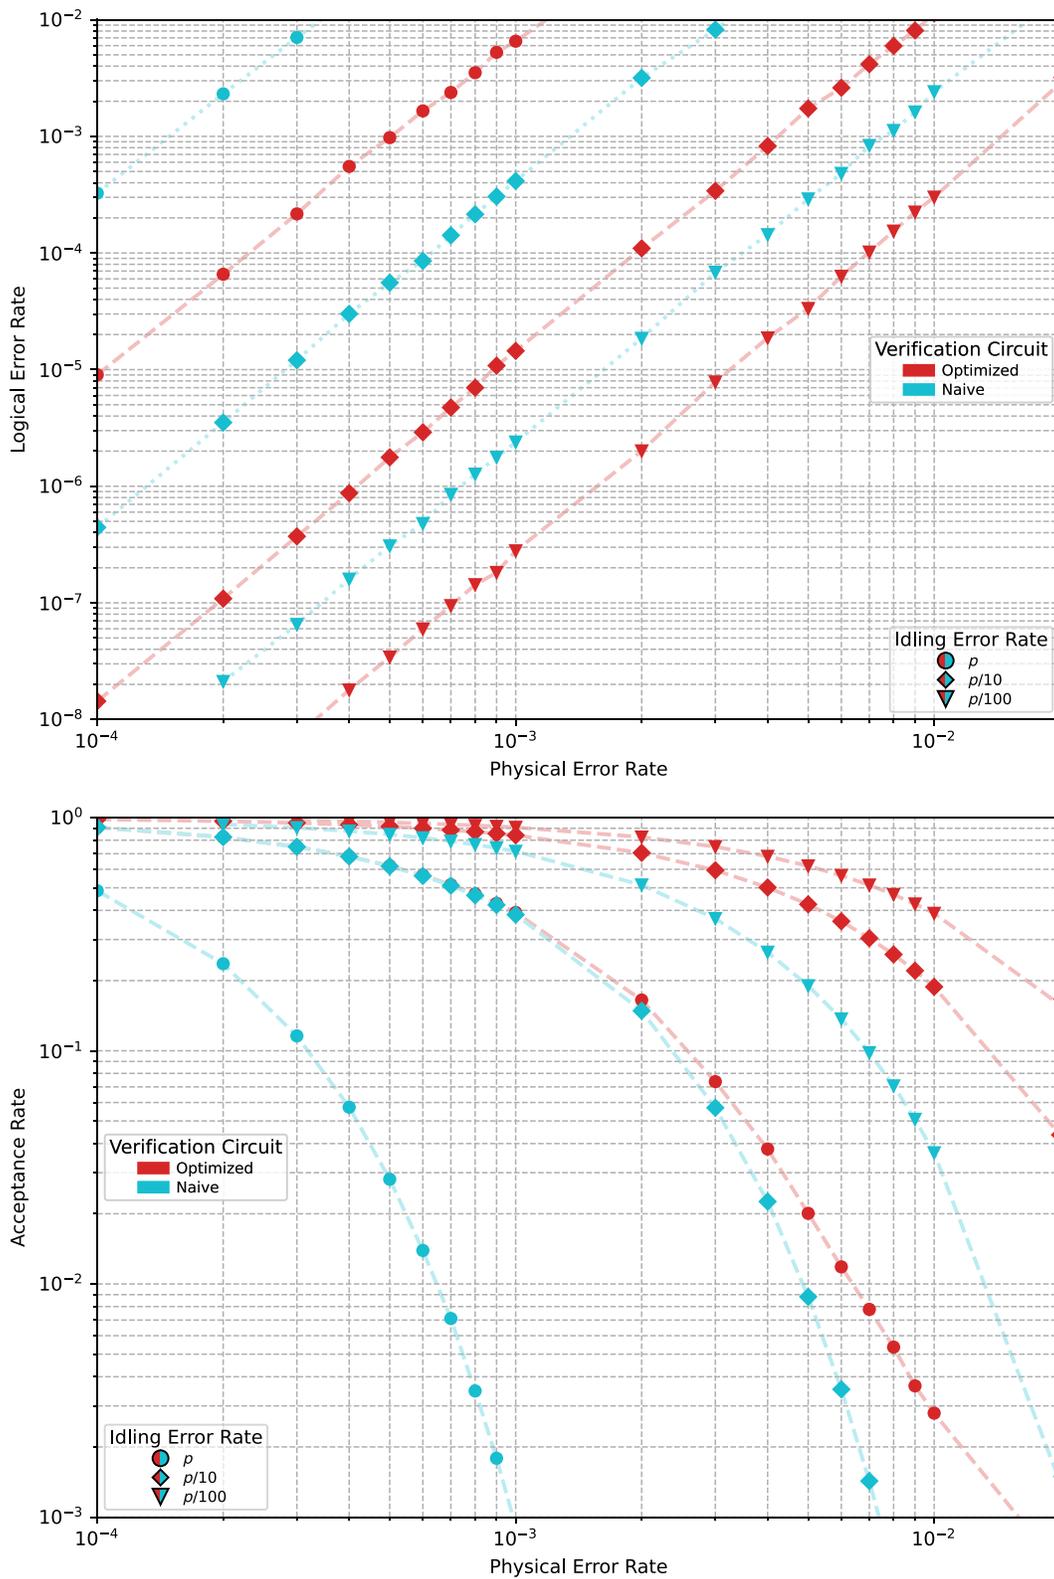


FIG. 11. Logical error rate and acceptance rate for nondeterministic fault-tolerant state preparation circuits for the logical $|0\rangle_L$ state of the $[[17, 1, 5]]$ color code constructed with our methods under circuit-level depolarizing noise with *sequential* gate execution using a LUT decoder.

VIII. CONCLUSION

In this work, we have introduced automated methods to synthesize (near-)optimal nondeterministic fault-tolerant state preparation circuits for logical computational basis states of any CSS code. We have shown the efficacy of the proposed methods by synthesizing fault-tolerant state preparation circuits for various small CSS codes—relevant for near-term devices—and demonstrated that these circuits have the desired fault-tolerance properties through numerical simulations. Even though we have framed the verification-circuit synthesis problem in the context of state preparation of certain logical basis states, our verification-circuit synthesis approach is more general. It can be used to verify other circuits as well.

In this work, we have focused on constraints on specific stabilizer measurements, but generally, in fault-tolerant circuit design, we are more concerned with constraints between different measurements. Looking forward, an interesting next step would be to apply methods used in this work to the design of fault-tolerant circuits from the perspective of detector-error models as in Ref. [43].

As discussed in Sec. VI, our methods do not guarantee a globally optimal circuit but rather give locally optimal sub-circuits for the state and verification-circuit synthesis problems. The natural next step would be to try to unify this into one synthesis problem. Given that the fault set of the non-fault-tolerant state preparation circuit has a huge impact on the verification circuit, we suspect there is a large room for improvement, especially when scaling to even higher distances. Alternatively, one could use these circuits in a concatenated scheme. Whether this can be extended to concatenation in a trivial matter or whether further measurements are needed to verify error propagation through the transversal CNOTs is left for future work.

In our work, we have assumed all-to-all connectivity between all data and ancilla qubits. Many quantum computing platforms do not provide such lenient locality constraints; however, further steps must be taken to ensure that nondeterministic fault-tolerant state preparation circuits can be embedded into such architectures without sacrificing fault tolerance.

All in all, this work is an important step towards providing automated design tools for fault-tolerant quantum computing research. The optimized circuits generated by our methods might play a role in further near-term experimental demonstrations of fault-tolerant quantum computing.

IX. USED SOFTWARE

The proposed methods have been implemented using the Python packages Stim [35], NumPy [78], Z3 [58], Qiskit

[79], LDPC [80], and Multiprocess [81,82]. For the simulations we have furthermore used GNU parallel [83]. The plots have been generated with Matplotlib [84].

ACKNOWLEDGMENTS

The authors would like to thank Timo Hillmann for valuable comments on an initial version of the manuscript.

T.P., L.S., L.B., and R.W. acknowledge funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation program (Grant Agreement No. 101001318) and Millenion, Grant Agreement No. 101114305). This work was part of the Munich Quantum Valley, which is supported by the Bavarian state government with funds from the Hightech Agenda Bayern Plus, and has been supported by the BMWK on the basis of a decision by the German Bundestag through project QuaST, as well as by the BMK, BMDW, and the State of Upper Austria in the frame of the COMET program (managed by the FFG).

Furthermore, M.M. gratefully acknowledges support by the European Union’s Horizon Europe research and innovation programme under Grant Agreement No. 101114305 (“MILLENION-SGA1” EU Project) and the ERC Starting Grant QNets through Grant No. 804247. This research is also part of the Munich Quantum Valley (K-8), which is supported by the Bavarian state government with funds from the Hightech Agenda Bayern Plus. He additionally acknowledges support by the BMBF project MUNIQC-ATOMS (Grant No. 13N16070), by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany’s Excellence Strategy “Cluster of Excellence Matter and Light for Quantum Computing (ML4Q) EXC 2004/1” 390534769.

DATA AVAILABILITY

The data that support the findings of this article are openly available [68,75].

APPENDIX A: FURTHER PLOTS FOR $d = 5$ CODES

This section provides further simulation results for the $[[19, 1, 5]]$ color code and the $[[25, 1, 5]]$ surface code. Simulations were conducted for idling noise of strengths p , $p/10$, and $p/100$ for two different scenarios: *parallel* and *sequential* gate execution.

Figure 12 shows simulation results for the $[[19, 1, 5]]$ color code with *parallel* gate execution, and Fig. 13 shows results for *sequential* gate execution for the $[[19, 1, 5]]$ code. Similarly, Fig. 14 shows simulation results for the $[[25, 1, 5]]$ color code with *parallel* gate execution, and Fig. 15 shows results for *sequential* gate execution for the $[[25, 1, 5]]$ code.

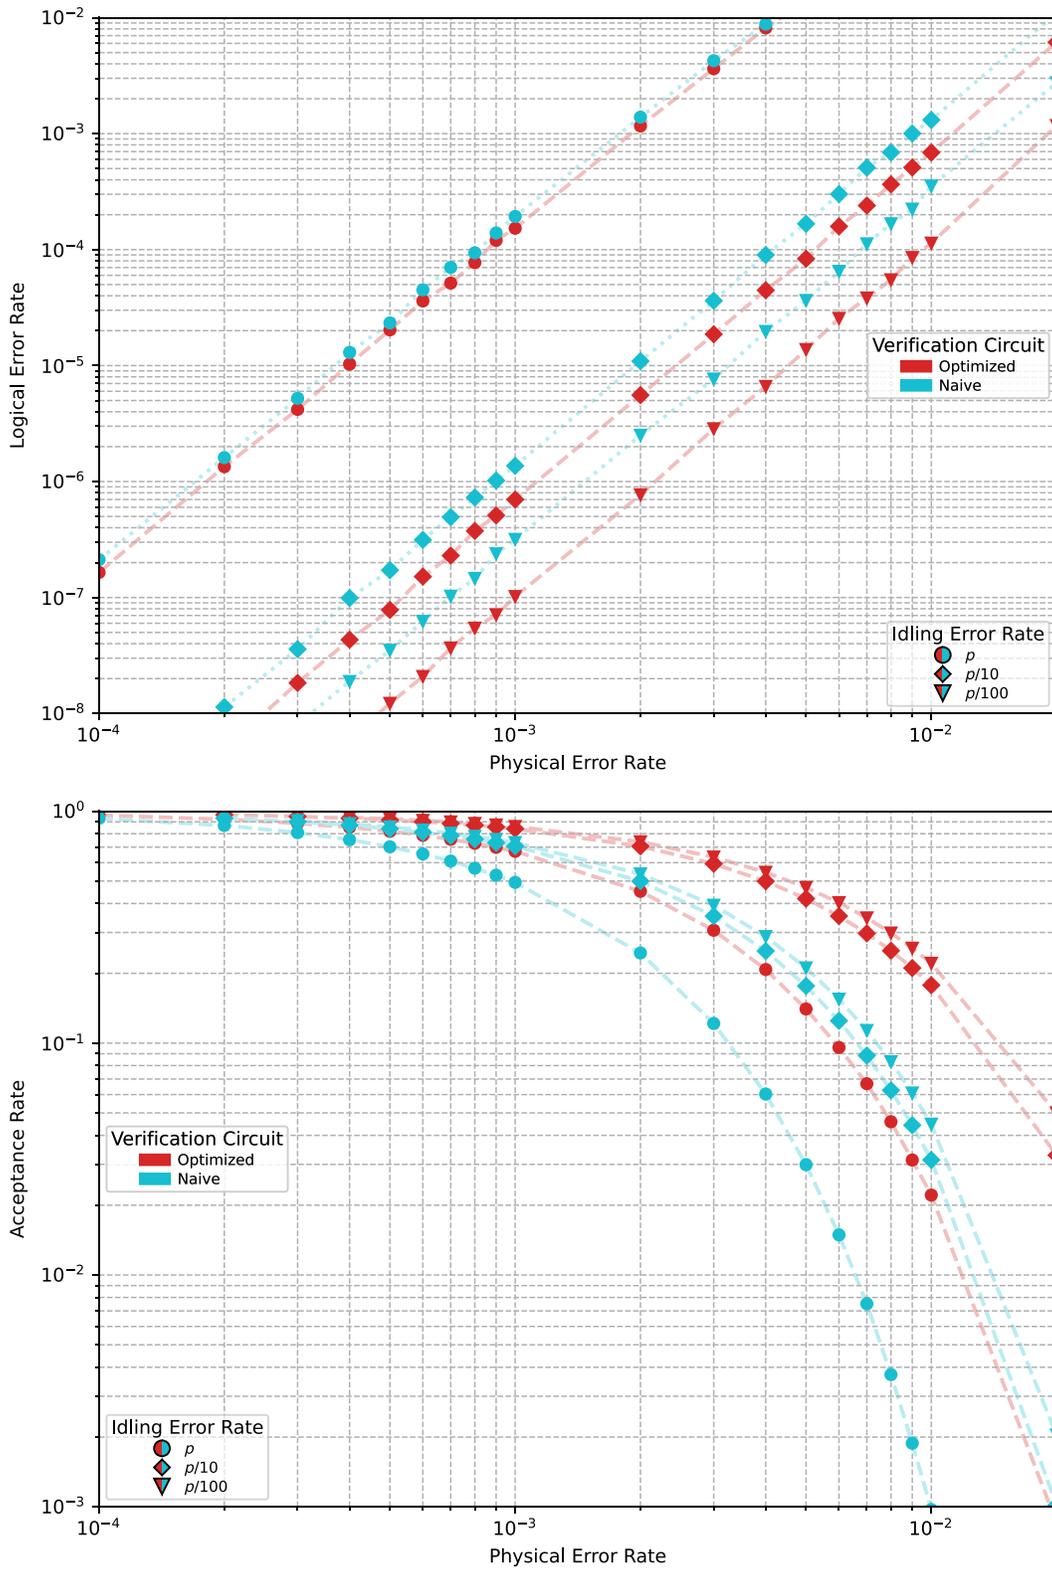


FIG. 12. Logical error rate and acceptance rate for nondeterministic fault-tolerant state preparation circuits for the logical $|0\rangle_L$ state of the $[[19, 1, 5]]$ color code constructed with our methods under circuit-level depolarizing noise with *parallel* gate execution using a LUT decoder.

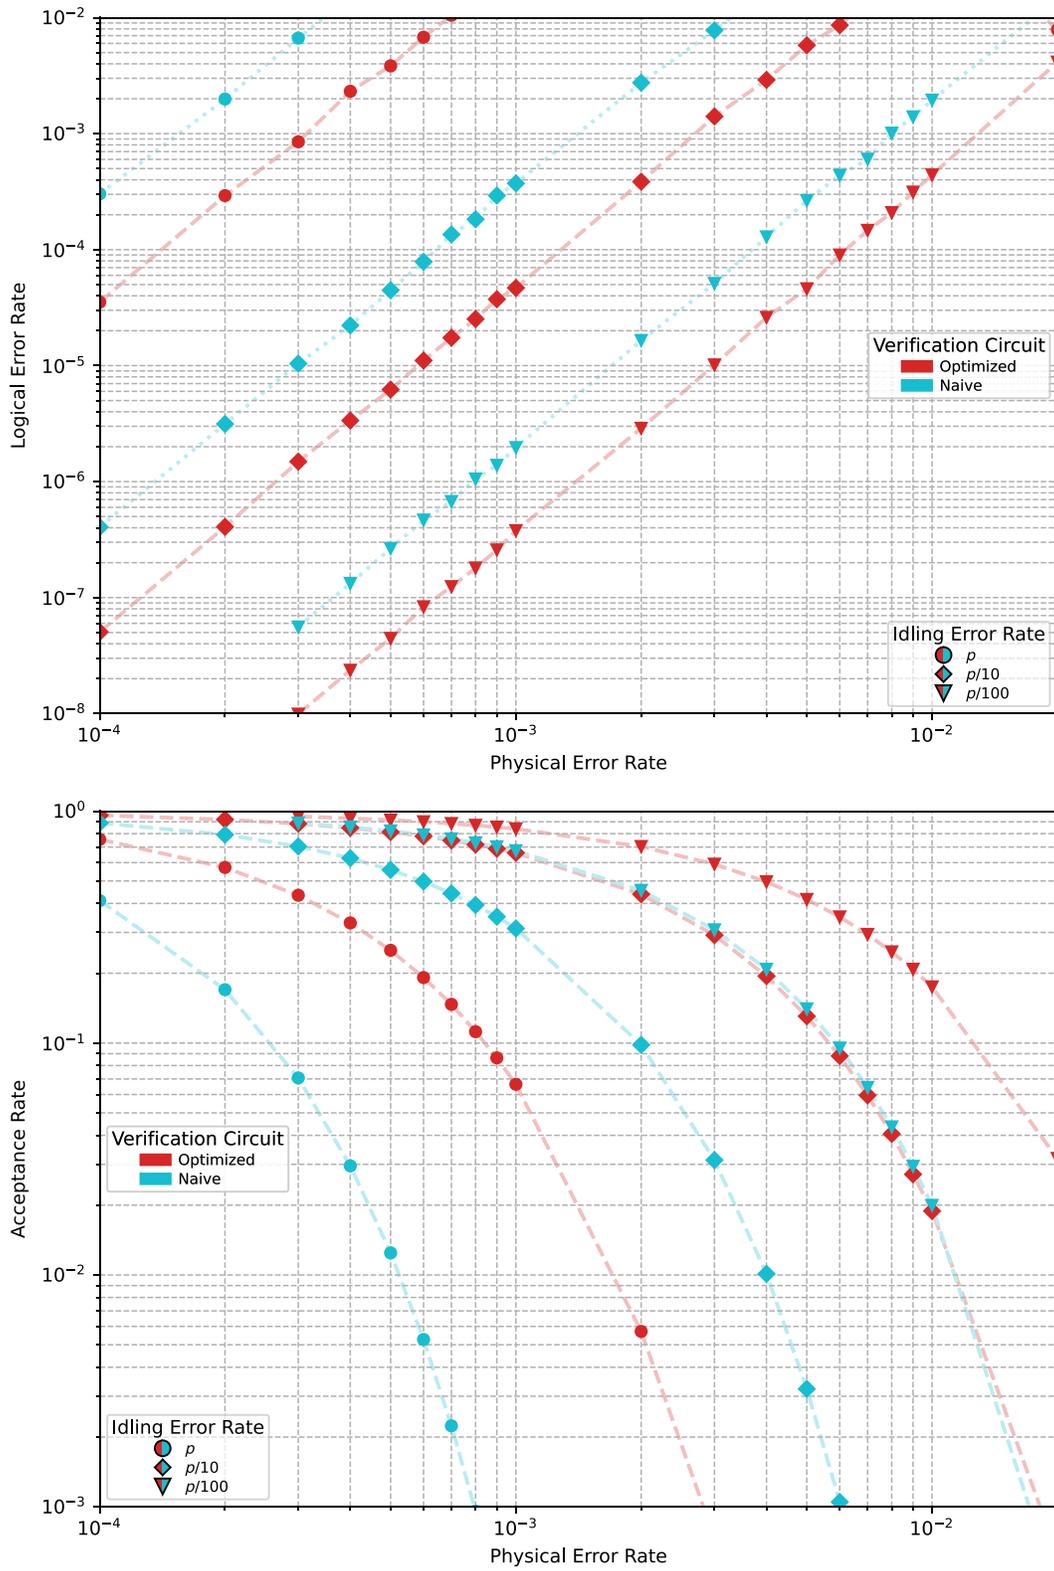


FIG. 13. Logical error rate and acceptance rate for nondeterministic fault-tolerant state preparation circuits for the logical $|0\rangle_L$ state of the $[[19, 1, 5]]$ color code constructed with our methods under circuit-level depolarizing noise with *sequential* gate execution using a LUT decoder.

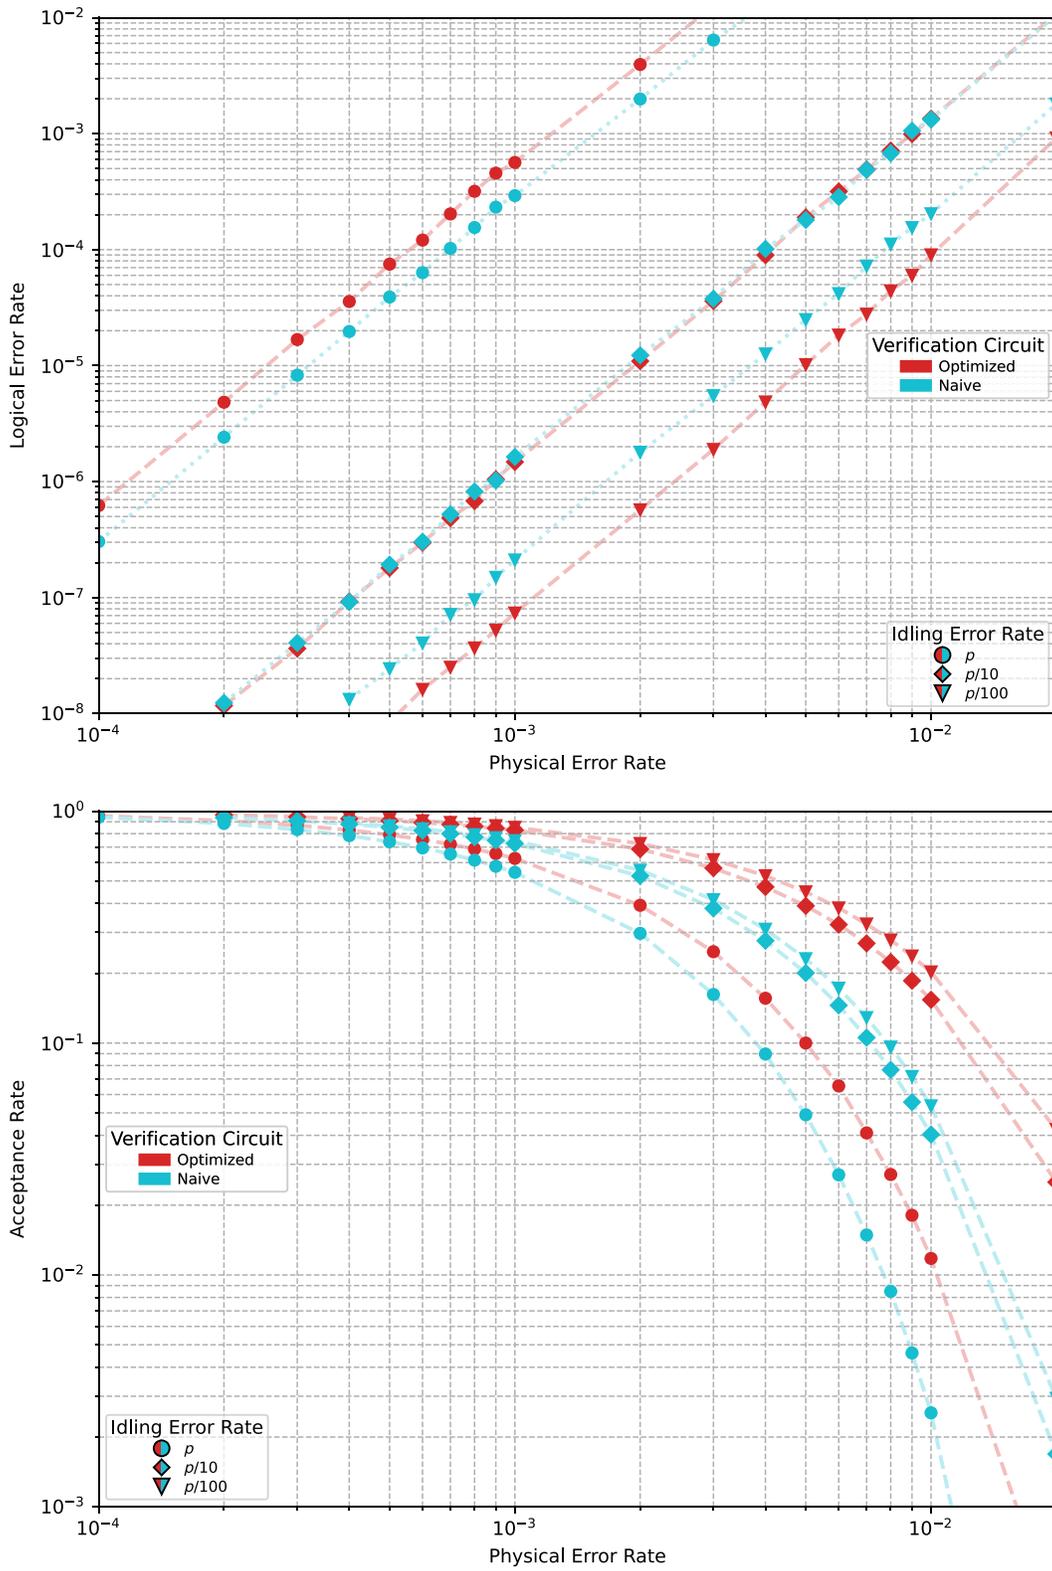


FIG. 14. Logical error rate and acceptance rate for nondeterministic fault-tolerant state preparation circuits for the logical $|0\rangle_L$ state of the $[[25, 1, 5]]$ rotated surface code constructed with our methods under circuit-level depolarizing noise with *parallel* gate execution using a LUT decoder.

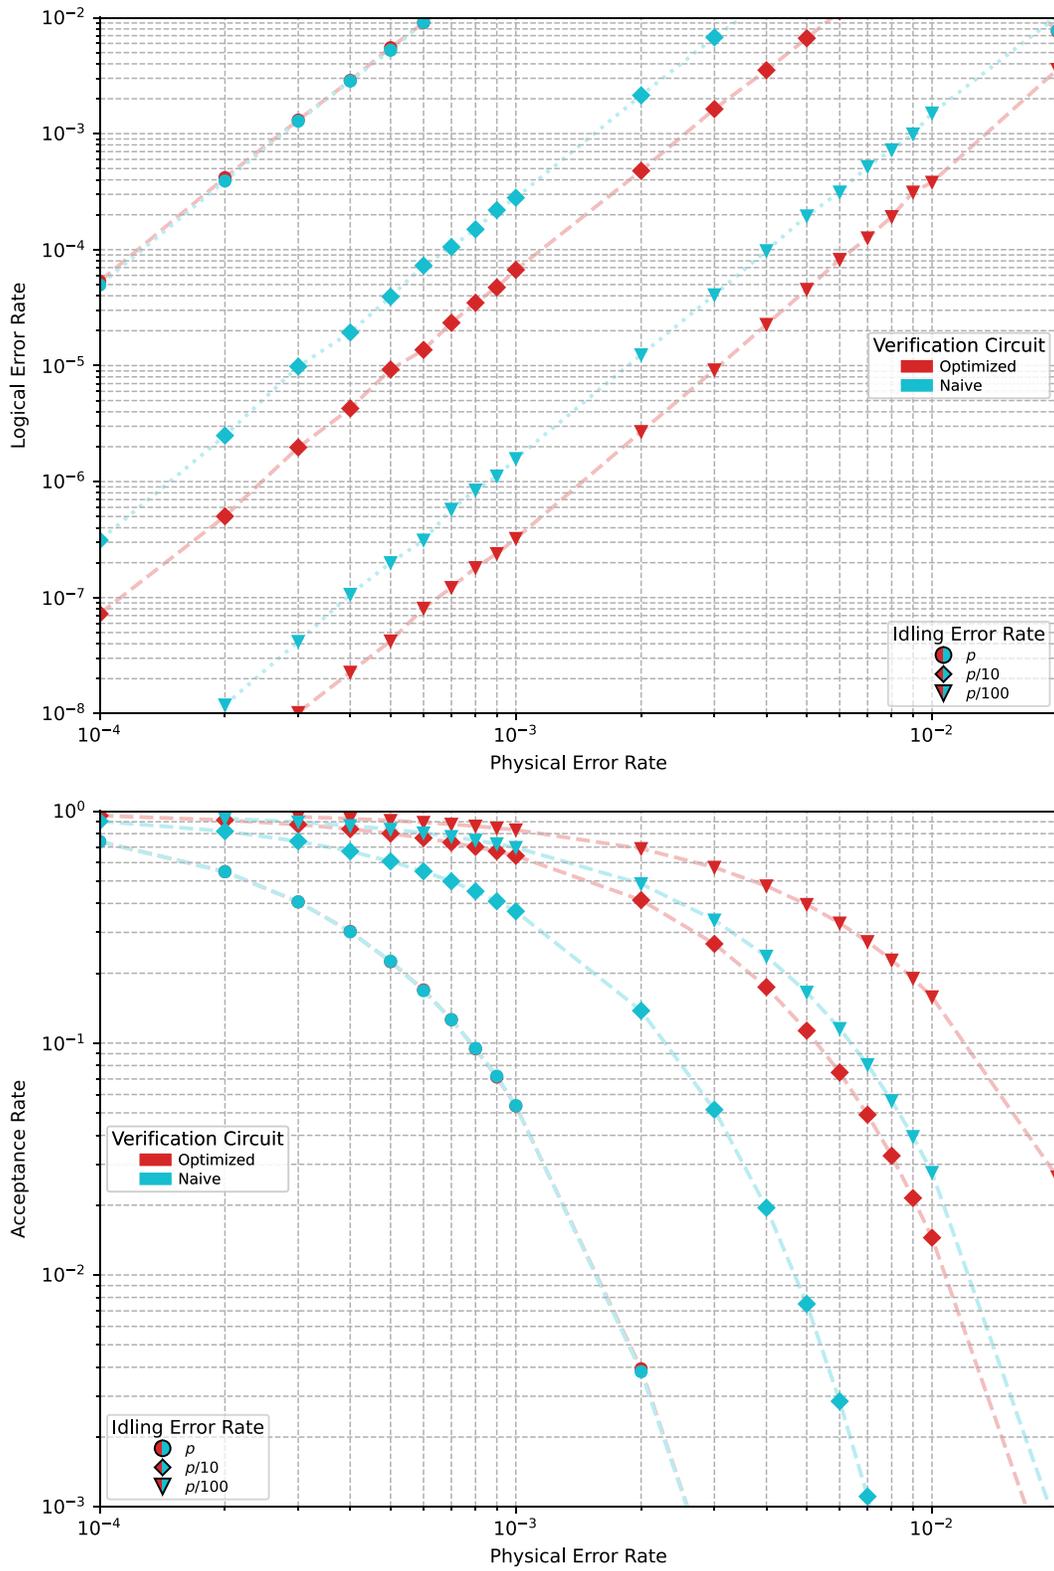


FIG. 15. Logical error rate and acceptance rate for nondeterministic fault-tolerant state preparation circuits for the logical $|0\rangle_L$ state of the $[[25, 1, 5]]$ rotated surface code constructed with our methods under circuit-level depolarizing noise with *sequential* gate execution using a LUT decoder.

APPENDIX B: CIRCUITS

This section lists state preparation circuits for various CSS codes synthesized with the proposed methods. The circuit metrics are shown in Table III.

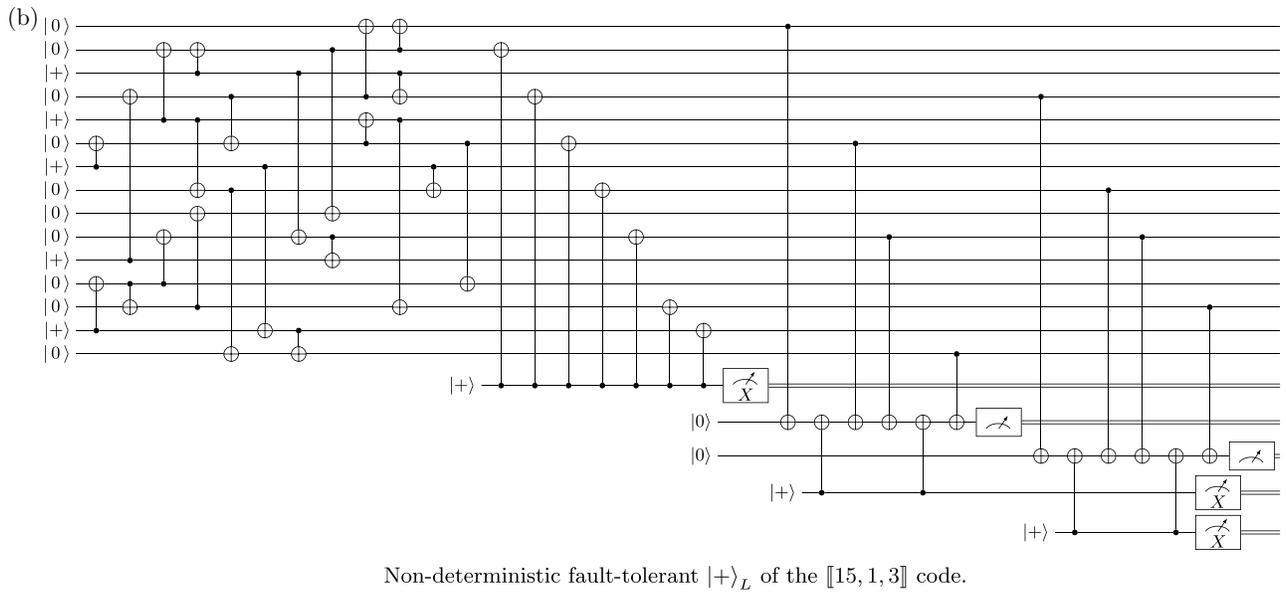
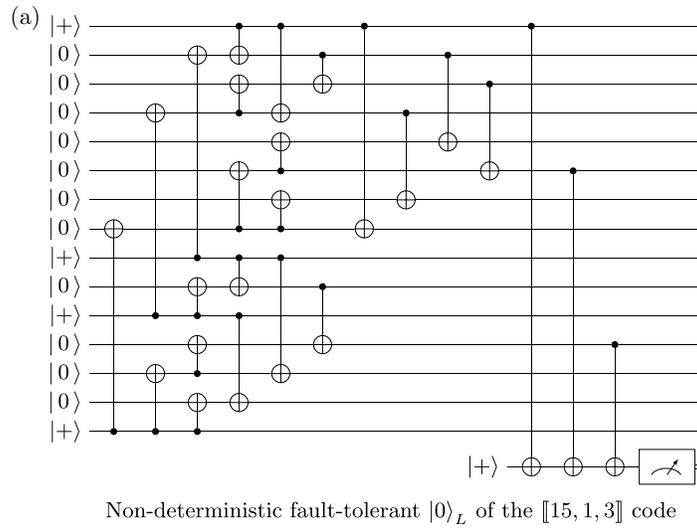


FIG. 16. Fault-tolerant state preparation circuits for the $[[15, 1, 3]]$ “tetrahedral” code. (a) Nondeterministic fault-tolerant $|0\rangle_L$ of the $[[15, 1, 3]]$ code. (b) Nondeterministic fault-tolerant $|+\rangle_L$ of the $[[15, 1, 3]]$ code.

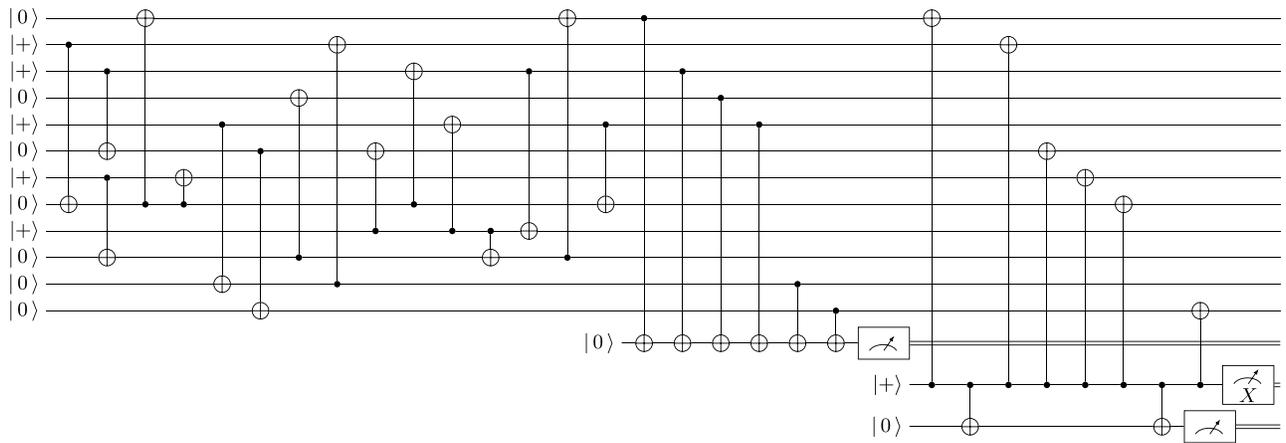


FIG. 17. Nondeterministic fault-tolerant $|0\rangle_L^2$ of the $[[12, 2, 4]]$ “Carbon” code.

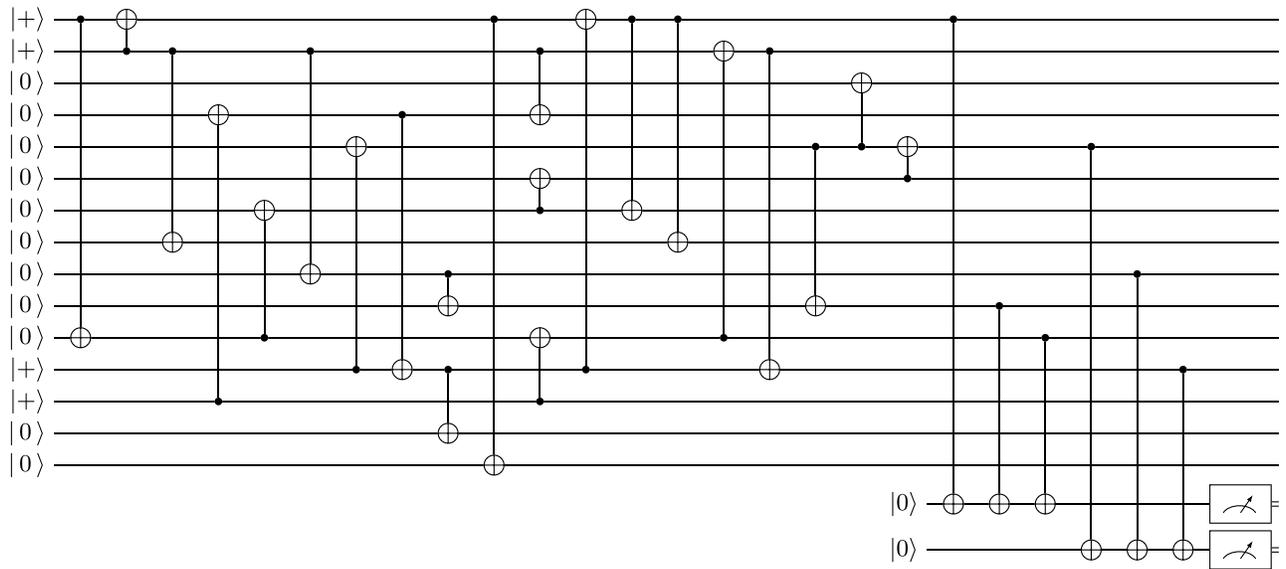


FIG. 18. Nondeterministic fault-tolerant $|0\rangle_L^7$ of the $[[15, 7, 3]]$ Hamming code.

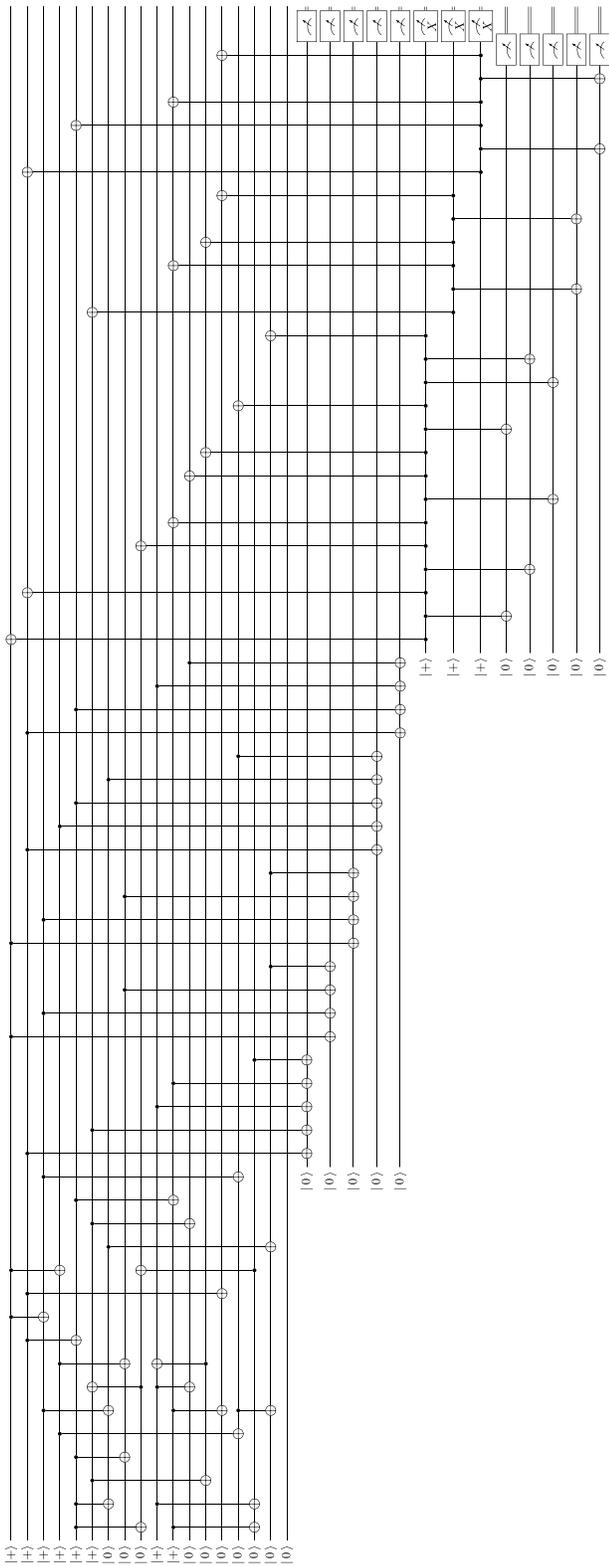


FIG. 19. Nondeterministic fault-tolerant $|0\rangle_L$ of the $[[17, 1, 5]]$ code.

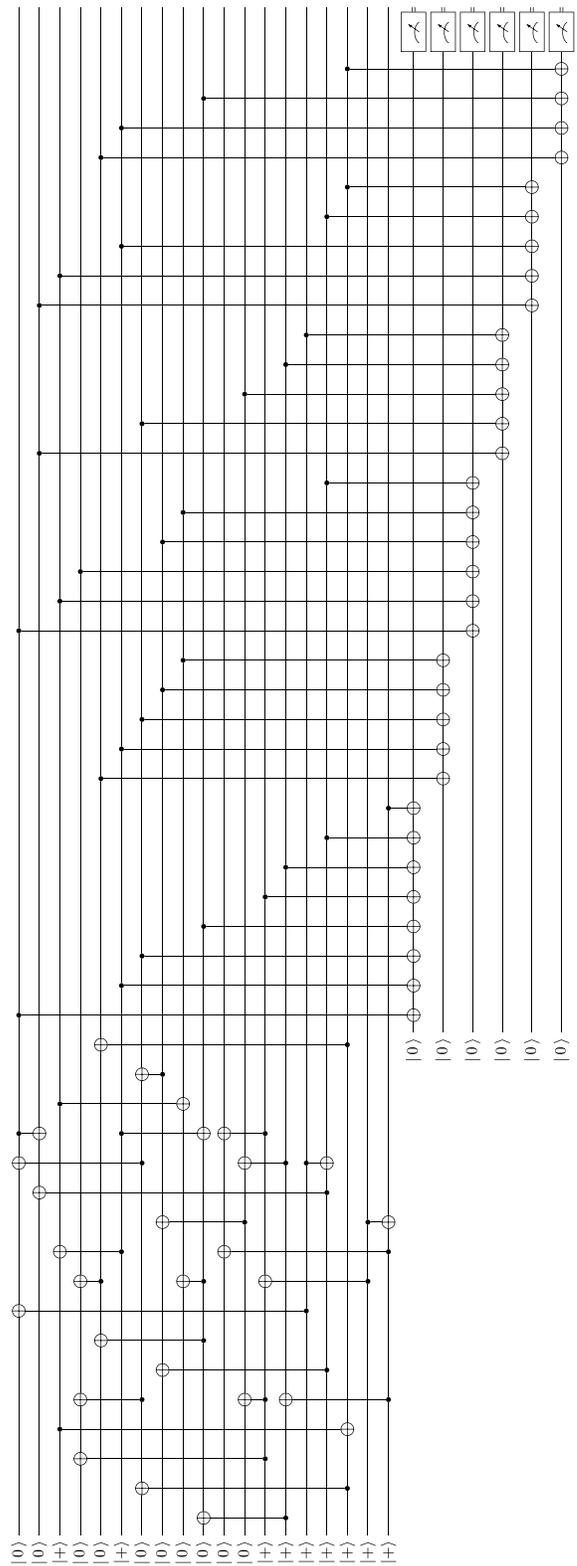


FIG. 20. Nondeterministic fault-tolerant $|0\rangle_L$ of the $[[19, 1, 5]]$ color code. Here only verification for X errors is shown.

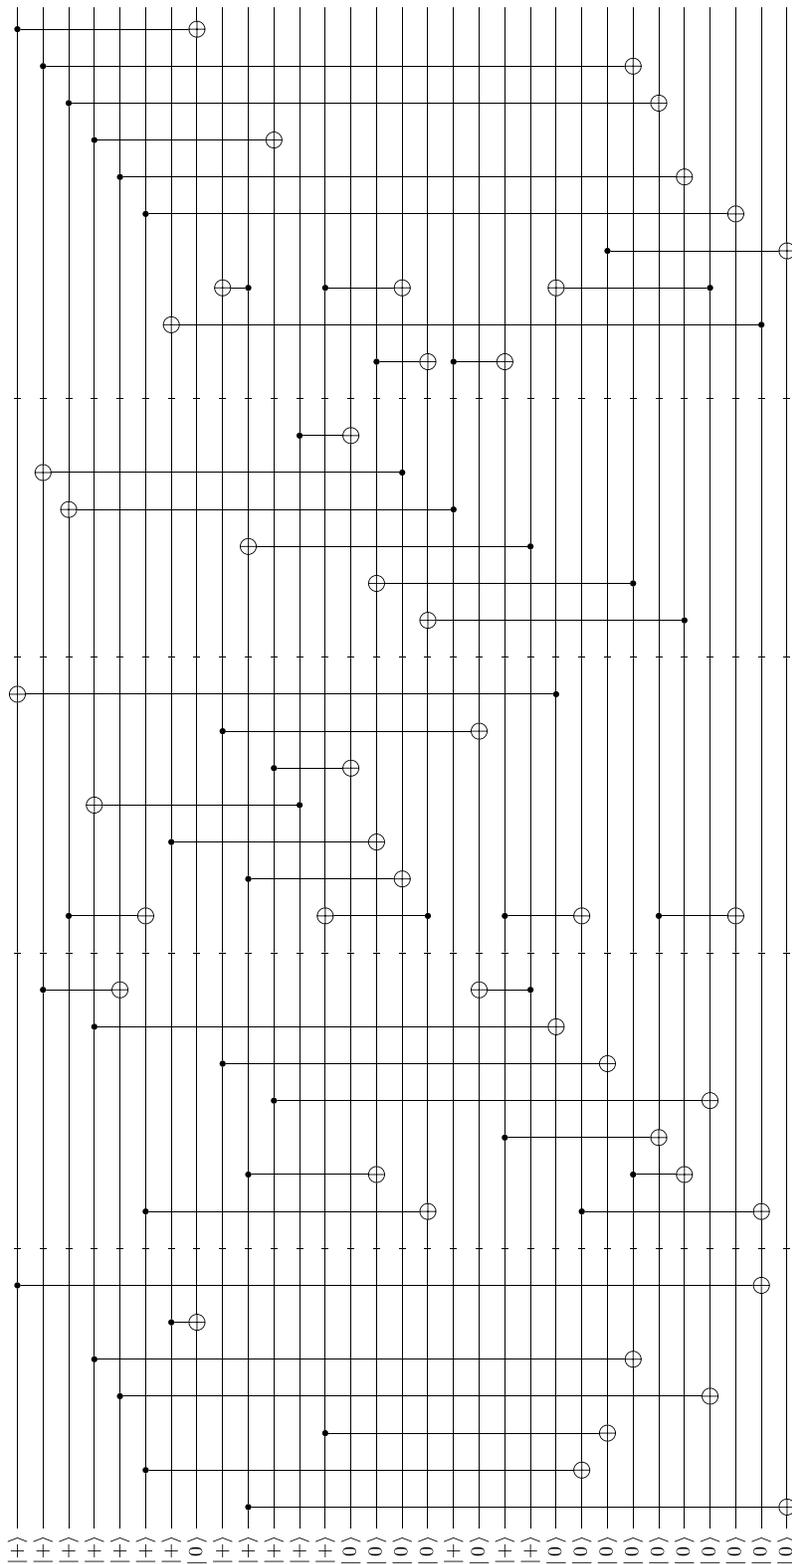


FIG. 21. Non-fault-tolerant $|0\rangle_L$ of the $[[31, 1, 7]]$ color code.

- [1] A. R. Calderbank and P. W. Shor, Good quantum error-correcting codes exist, *Phys. Rev. A* **54**, 1098 (1996).
- [2] A. Steane, Multiple particle interference and quantum error correction, *Proc. R. Soc. Lond. A* **452**, 2551 (1996).
- [3] A. M. Steane, Error correcting codes in quantum theory, *Phys. Rev. Lett.* **77**, 793 (1996).
- [4] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information* (Cambridge University Press, Cambridge, 2010).
- [5] N. P. Breuckmann and J. N. Eberhardt, Quantum low-density parity-check codes, *PRX Quantum* **2**, 040101 (2021).
- [6] H. Bombín, Single-shot fault-tolerant quantum error correction, *Phys. Rev. X* **5**, 031043 (2015).
- [7] N. Delfosse, B. W. Reichardt, and K. M. Svore, Beyond single-shot fault-tolerant quantum error correction, *IEEE Trans. Inf. Theory* **68**, 287 (2022).
- [8] A. O. Quintavalle, M. Vasmer, J. Roffe, and E. T. Campbell, Single-shot error correction of three-dimensional homological product codes, *PRX Quantum* **2**, 020340 (2021).
- [9] O. Fawzi, A. Grospellier, and A. Leverrier, Constant overhead quantum fault tolerance with quantum expander codes, *Commun. ACM* **64**, 106 (2020).
- [10] H. Goto, Minimizing resource overheads for fault-tolerant preparation of encoded states of the Steane code, *Sci. Rep.* **6**, 19578 (2016).
- [11] A. Bermudez, X. Xu, M. Gutiérrez, S. C. Benjamin, and M. Müller, Fault-tolerant protection of near-term trapped-ion topological qubits under realistic noise sources, *Phys. Rev. A* **100**, 062307 (2019).
- [12] L. Postler, S. Heußen, I. Pogorelov, M. Rispler, T. Feldker, M. Meth, C. D. Marciniak, R. Stricker, M. Ringbauer, R. Blatt, P. Schindler, M. Müller, and T. Monz, Demonstration of fault-tolerant universal quantum gate operations, *Nature* **605**, 675 (2022).
- [13] C. Chamberland and A. W. Cross, Fault-tolerant magic state preparation with flag qubits, *Quantum* **3**, 143 (2019).
- [14] F. Butt, S. Heußen, M. Rispler, and M. Müller, Fault-tolerant code-switching protocols for near-term quantum processors, *PRX Quantum* **5**, 020345 (2024).
- [15] S. Heußen, D. F. Locher, and M. Müller, Measurement-free fault-tolerant quantum error correction in near-term devices, *PRX Quantum* **5**, 010333 (2024).
- [16] C. Ryan-Anderson, J. G. Bohnet, K. Lee, D. Gresh, A. Hankin, J. P. Gaebler, D. Francois, A. Chernoguzov, D. Lucchetti, N. C. Brown, T. M. Gatterman, S. K. Halit, K. Gilmore, J. A. Gerber, B. Neyenhuis, D. Hayes, and R. P. Stutz, Realization of real-time fault-tolerant quantum error correction, *Phys. Rev. X* **11**, 041058 (2021).
- [17] D. Bluvstein, S. J. Evered, A. A. Geim, S. H. Li, H. Zhou, T. Manovitz, S. Ebadi, M. Cain, M. Kalinowski, D. Hangleiter, J. P. B. Ataiades, N. Maskara, I. Cong, X. Gao, P. S. Rodriguez, T. Karolyshyn, G. Semeghini, M. J. Gullans, M. Greiner, V. Vuletic, and M. D. Lukin, Logical quantum processor based on reconfigurable atom arrays, *Nature* **626**, 58 (2024).
- [18] M. P. da Silva *et al.*, Demonstration of logical qubits and repeated error correction with better-than-physical error rates, [arXiv:2404.02280](https://arxiv.org/abs/2404.02280).
- [19] I. Pogorelov, F. Butt, L. Postler, C. D. Marciniak, P. Schindler, M. Müller, and T. Monz, Experimental fault-tolerant code switching, [arXiv:2403.13732](https://arxiv.org/abs/2403.13732).
- [20] R. Zen, J. Olle, L. Colmenarez, M. Puviani, M. Müller, and F. Marquardt, Quantum circuit discovery for fault-tolerant logical state preparation with reinforcement learning, [arXiv:2402.17761](https://arxiv.org/abs/2402.17761).
- [21] A. Goswami, M. Mhalla, and V. Savin, Fault-tolerant preparation of quantum polar codes encoding one logical qubit, *Phys. Rev. A* **108**, 042605 (2023).
- [22] A. Goswami, M. Mhalla, and V. Savin, Factory-based fault-tolerant preparation of quantum polar codes encoding one logical qubit, *Phys. Rev. A* **110**, 012438 (2024).
- [23] A. Biere and D. Kröning, in *Handbook of Model Checking* (Springer, Cham, Switzerland, 2018), p. 277.
- [24] D. Brand, in *Int'l Conf. on CAD* (IEEE Computer Society Press, Santa Clara, CA, USA, 1993), p. 534.
- [25] S. Eggensgluß, R. Wille, and R. Drechsler, in *Int'l Conf. on CAD* (IEEE Computer Society Press, San Jose, California, 2013).
- [26] A. Gebregiorgis and M. B. Tahoori, in *2019 Des. Autom. Test Eur. Conf. Exhib. DATE* (IEEE Computer Society Press, Florence, Italy, 2019), p. 1028.
- [27] D. Kaufmann, A. Biere, and M. Kauers, in *2019 Form. Methods Comput. Aided Des. FMCAD* (IEEE Computer Society Press, San Jose, California, USA, 2019), p. 28.
- [28] R. Wille, D. Große, F. Haedicke, and R. Drechsler, in *Forum on Specification and Design Languages* (IEEE Computer Society Press, Sophia Antipolis, France, 2009).
- [29] V. Khomenko, M. Koutny, and A. Yakovlev, in *Proc. Fourth Int. Conf. Appl. Concurr. Syst. Des. 2004 ACSD 2004* (IEEE Computer Society Press, Hamilton, Ontario, Canada, 2004), p. 16.
- [30] R. Wille, L. Burgholzer, and A. Zulehner, in *Design Automation Conf.* (Association for Computing Machinery, Las Vegas, NV, USA, 2019).
- [31] B. Tan and J. Cong, in *Int'l Conf. on CAD* (Association for Computing Machinery, Virtual Event, USA, 2020).
- [32] D. B. Tan, M. Y. Niu, and C. Gidney, A SAT scalpel for lattice surgery: Representation and synthesis of subroutines for surface-code fault-tolerant quantum computing, [arXiv:2404.18369](https://arxiv.org/abs/2404.18369) [quant-ph].
- [33] N. Shutty and C. Chamberland, Decoding merged color-surface codes and finding fault-tolerant Clifford circuits using solvers for satisfiability modulo theories, *Phys. Rev. Appl.* **18**, 014072 (2022).
- [34] T. Peham, N. Brandl, R. Kueng, R. Wille, and L. Burgholzer, in *2023 IEEE Int. Conf. Quantum Comput. Eng. QCE* (IEEE Computer Society Press, Seattle, USA, 2023), Vol. 01, p. 802.
- [35] C. Gidney, Stim: A fast stabilizer circuit simulator, *Quantum* **5**, 497 (2021).
- [36] R. Wille, L. Berent, T. Forster, J. Kunasaikaran, K. Mato, T. Peham, N. Quetschlich, D. Rovara, A. Sander, L. Schmid, D. Schönberger, Y. Stade, and L. Burgholzer, *The MQT Handbook: A summary of design automation tools and software for quantum computing*, [arXiv:2405.17543](https://arxiv.org/abs/2405.17543).
- [37] D. Gottesman, *Stabilizer codes and quantum error correction.*, Ph.D. thesis, Caltech, 1997.
- [38] H. Bombin, An introduction to topological quantum codes, [arXiv:1311.0277](https://arxiv.org/abs/1311.0277).

- [39] P. Aliferis, D. Gottesman, and J. Preskill, Quantum accuracy threshold for concatenated distance-3 codes, [arXiv:quant-ph/0504218](#).
- [40] D. Gottesman, An introduction to quantum error correction and fault-tolerant quantum computation, [arXiv:0904.2557](#) [quant-ph].
- [41] C. Chamberland and M. E. Beverland, Flag fault-tolerant error correction with arbitrary distance codes, *Quantum* **2**, 53 (2018).
- [42] A. Paetznick and B. W. Reichardt, Fault-tolerant ancilla preparation and noise threshold lower bounds for the 23-qubit Golay code, [arXiv:1106.2190](#) [quant-ph].
- [43] P.-J. H. S. Derks, A. Townsend-Teague, A. G. Burchards, and J. Eisert, Designing fault-tolerant circuits using detector error models, [arXiv:2407.13826](#) [quant-ph].
- [44] E. Dennis, A. Kitaev, A. Landahl, and J. Preskill, Topological quantum memory, *J. Math. Phys.* **43**, 4452 (2002).
- [45] T. Tansuwannont, B. Pato, and K. R. Brown, Adaptive syndrome measurements for Shor-style error correction, *Quantum* **7**, 1075 (2023).
- [46] A. Kissinger, Phase-free ZX diagrams are CSS codes (\dots or how to graphically grok the surface code), [arXiv:2204.14038](#) [quant-ph].
- [47] J. Jiang, X. Sun, S.-H. Teng, B. Wu, K. Wu, and J. Zhang, Optimal space-depth trade-off of CNOT circuits in quantum logic synthesis, [arXiv:1907.05087](#) [quant-ph].
- [48] K. N. Patel, I. L. Markov, and J. P. Hayes, Optimal synthesis of linear reversible circuits, *Quantum Inf. Comput.* **8**, 282 (2008).
- [49] D. Maslov and B. Zindorf, Depth optimization of CZ, CNOT, and Clifford circuits, *IEEE Trans. Quantum Eng.* **3**, 1 (2022).
- [50] Timothée Goubault de Brugière, M. Baboulin, B. Valiron, S. Martiel, and C. Allouche, Gaussian elimination versus greedy methods for the synthesis of linear reversible circuits, *ACM Trans. Quantum Comput.* **2**, 1 (2021).
- [51] S. Bravyi, J. A. Latone, and D. Maslov, 6-qubit optimal Clifford circuits, *npj Quantum Inf* **8**, 1 (2022).
- [52] S. Schneider, L. Burgholzer, and R. Wille, in *Asia and South Pacific Design Automation Conf.* (IEEE Computer Society Press, Tokyo, Japan, 2023).
- [53] S. Aaronson and D. Gottesman, Improved simulation of stabilizer circuits, *Phys. Rev. A* **70**, 052328 (2004).
- [54] S. Bravyi, R. Shaydulin, S. Hu, and D. Maslov, Clifford circuit optimization with templates and symbolic Pauli gates, *Quantum* **5**, 580 (2021).
- [55] D. Maslov, Linear depth stabilizer and quantum Fourier transformation circuits with no auxiliary qubits in finite neighbor quantum architectures, *Phys. Rev. A* **76**, 052310 (2007).
- [56] V. Kabanets and J.-Y. Cai, in *Proc. Thirty-Second Annu. ACM Symp. Theory Comput.* (ACM, Prague, Czech Republic, 2000), p. 73.
- [57] A. Biere, M. Heule, Hans van Maaren, and T. Walsh, *Handbook of Satisfiability* (IOS Press, Clifton, VA, USA, 2009).
- [58] Leonardo de Moura and N. Bjørner, in *Tools Algorithms Constr. Anal. Syst.* (Springer, Dublin, Ireland, 2008), p. 337.
- [59] B. Schaeffer and M. Perkowski, A cost minimization approach to synthesis of linear reversible circuits, [arXiv:1407.0070](#).
- [60] In principle, these reductions can be tried between every column operation to reduce the number of 1s in the matrix even faster. We have not observed a definite benefit of doing this as it sometimes actually leads to worse solutions.
- [61] E. Berlekamp, R. McEliece, and H. van Tilborg, On the inherent intractability of certain coding problems, *IEEE Trans. Inf. Theory* **24**, 384 (1978).
- [62] R. M. Karp, in *Complexity of Computer Computations: Proceedings of a symposium on the Complexity of Computer Computations, held March 20–22, 1972, at the IBM Thomas J. Watson Research Center, Yorktown Heights, New York, and sponsored by the Office of Naval Research, Mathematics Program, IBM World Trade Corporation, and the IBM Research Mathematical Sciences Department*, edited by R. E. Miller, J. W. Thatcher, and J. D. Bohlinger (Springer US, Boston, MA, 1972), p. 85.
- [63] V. Chvatal, A greedy heuristic for the set-covering problem, *Math. Oper. Res.* **4**, 233 (1979).
- [64] M. Borges-Quintana, M. A. Borges-Trenard, and E. Martínez-Moro, Computing coset leaders of binary codes, [arXiv:1001.5074](#).
- [65] M. Borges-Quintana, M. A. Borges-Trenard, I. Márquez-Corbella, and E. Martínez-Moro, Computing coset leaders and leader codewords of binary codes, [arXiv:1211.5568](#).
- [66] E. W. Mayr, Some complexity results for polynomial ideals, *J. Complex.* **13**, 303 (1997).
- [67] R. Chao and B. W. Reichardt, Flag fault-tolerant error correction for any stabilizer code, *PRX Quantum* **1**, 010302 (2020).
- [68] L. Berent, L. Burgholzer, P.-J. H. S. Derks, T. Grurl, T. Peham, C. Pichler, and L. Schmid, MQT QECC, GitHub (2025), available at <https://github.com/cda-tum/mqt-qecc>, version 1.9.0.
- [69] A. Steane, Quantum Reed-Muller codes, *IEEE Trans. Inf. Theory* **45**, 1701 (1999).
- [70] P. W. Shor, Scheme for reducing decoherence in quantum computer memory, *Phys. Rev. A* **52**, R2493 (1995).
- [71] H. Bombin and M. A. Martin-Delgado, Optimal resources for topological two-dimensional stabilizer codes: Comparative study, *Phys. Rev. A* **76**, 012305 (2007).
- [72] A. M. Steane, Simple quantum error-correcting codes, *Phys. Rev. A* **54**, 4741 (1996).
- [73] H. Bombin and M. A. Martin-Delgado, Topological quantum distillation, *Phys. Rev. Lett.* **97**, 180501 (2006).
- [74] Note that the RL method does not minimize circuit depth and hence may produce different circuits of varying depth for multiple attempts. Here, we use the instances that have the lowest depth for a minimal CNOT count.
- [75] T. Peham, L. Schmid, L. Berent, M. Müller, and R. Wille, Data for “Automated Synthesis of Fault-Tolerant State Preparation Circuits for Quantum Error Correction Codes” (2025) available at <https://zenodo.org/records/15058397>.
- [76] Note that, as pointed out, e.g., in Ref. [18], if one operates the quantum circuit for the $[[12, 2, 4]]$ carbon code in

an error detecting regime, i.e., each run with any error (on data qubits) triggering stabilizer is discarded, all two-error events will be sorted out, and one obtains a approximately p^3 scaling for the logical error rate. In our simulations, we focus only on single errors in the preparation of the carbon code states.

- [77] S. J. Evered, D. Bluvstein, M. Kalinowski, S. Ebadi, T. Manovitz, H. Zhou, S. H. Li, A. A. Geim, T. T. Wang, N. Maskara, H. Levine, G. Semeghini, M. Greiner, V. Vuletić, and M. D. Lukin, High-fidelity parallel entangling gates on a neutral-atom quantum computer, *Nature* **622**, 268 (2023).
- [78] C. R. Harris *et al.*, Array programming with NumPy, *Nature* **585**, 357 (2020).
- [79] Qiskit contributors, Qiskit: An Open-source Framework for Quantum Computing, Zenodo (2023).
- [80] J. Roffe, LDPC: Python tools for low density parity check codes (2022).
- [81] M. M. McKerns, L. Strand, T. Sullivan, A. Fang, and M. A. G. Aivazis, Building a framework for predictive science, [arXiv:1202.1056](https://arxiv.org/abs/1202.1056).
- [82] Michael McKerns and Michael Aivazis, Pathos: A framework for heterogeneous computing (2010).
- [83] O. Tange, *GNU Parallel 2018* (Ole Tange, 2018).
- [84] J. D. Hunter, Matplotlib: A 2D graphics environment, *Comput. Sci. Eng.* **9**, 90 (2007).

Correction: The reference citations in the Data Availability statement were set incorrectly during the production cycle and have been fixed. Reference [75] has been corrected to include the link to the data repository.