

Automatic Implementation and Evaluation of Error-Correcting Codes for Quantum Computing

An Open-Source Framework for Quantum Error Correction

Thomas Grurl^{*†}

Christoph Pichler[†]

Jürgen Fuß^{*}

Robert Wille^{‡§}

^{*}University of Applied Sciences Upper Austria, Austria

[†]Johannes Kepler University Linz, Austria

[‡]Chair for Design Automation, Technical University of Munich, Germany

[§]Software Competence Center Hagenberg GmbH (SCCH), Austria

thomas.grurl@fh-hagenberg.at

christoph.pichler@jku.at

juergen.fuss@fh-hagenberg.at

robert.wille@tum.de

www.cda.cit.tum.de/research/quantum/

Abstract—Due to the fragility of quantum mechanical effects, real quantum computers are plagued by frequent noise effects that cause errors during computations. Quantum error-correcting codes address this problem by providing means to identify and correct corresponding errors. However, most of the research on quantum error correction is theoretical or has been evaluated for specific hardware models only. Moreover, the development of corresponding codes and the evaluation of whether they indeed solve the problem for a particular hardware model, still often rests on tedious trial-and-error thus far. In this work, we propose an open-source framework that supports engineers and researchers in these tasks by automatically applying error-correcting codes for a given application followed by an automatic noise-aware quantum circuit simulation. Case studies showcase that this allows for a substantially more efficient implementation and evaluation of error-correcting codes.

I. INTRODUCTION

Quantum computers can solve specific problems substantially faster than classical computers. Examples for such problems include algorithms that have been found in the areas of chemistry [1], machine learning [2], biology [3], and finance [4]. They achieve this by exploiting quantum mechanical effects during their computations, such as superposition (i.e., a qubit can be in a combination of 0 and 1) and entanglement (i.e., measurement outcomes of individual qubits are correlated). However, the same properties that make quantum computing so powerful also make it extremely prone to errors. For example, while it is a key feature of quantum bits that they can be in a superposition of 0 and 1, this also makes them fragile against operational or environmental noise. Accordingly, it is believed that (quantum) error correction is going to be an essential part of future quantum hardware [5].

But error correction is not easy. That it is even possible in the quantum realm is not straightforward. The first algorithm demonstrating this was published in 1995 by Peter Shor [6]. Since then there have been considerable developments in the research on quantum error correction. Today, there are multiple types of quantum error-correcting schemes available such as CSS codes [6]–[8], stabilizer codes [9], or surface codes [10], [11]. Additionally, there have also been developments in adaptive error-correcting codes optimized for specific quantum hardware [12], [13].

However, despite this progress, most of the corresponding work in this domain still heavily relies on manual labor and/or is based on theoretical results only. To the best of our knowledge, methods (or implementations thereof) that take a quantum circuit and *automatically* extend it with a correspondingly chosen error-correcting scheme (similar to a compiler that translates a given quantum algorithm to a corresponding hardware-applicable realization) do not exist, yet. This may be caused by the fact that, thus far, the focus clearly (and understandably) is on the development of concepts for error-correcting codes and proof-of-concept implementations on selected hardware models only. Additionally, the development of such a framework is not straightforward. For

example, the decoding and correction steps are exponentially hard [14] and, while this is addressed in recent work (such as [15]), still remain complex. Furthermore, applying quantum operations to qubits protected by an error-correcting code is often not directly possible but requires specific routines (as shown by the no-go theorem [16]).

As a consequence, evaluations on whether an error-correcting code is useful in different scenarios are cumbersome and, hence, often done with rather selected use cases only. Because of this, comprehensive case studies in which the usefulness of error-correcting codes are evaluated with respect to different circuits, how they are configured, and on what hardware model they are applied, do not exist.

In this work we address this issue: We propose an open-source framework¹ that supports engineers and researchers in the task of evaluating error-correcting codes. The framework allows to *automatically* apply error correction schemes to a given quantum circuit, followed by an automatic noise-aware quantum circuit simulation. To this end, we implemented different error-correcting codes and utilized existing methods for quantum circuit simulation. In this way, a framework is created that allows for error-correcting codes to be easily analyzed—with minimal manual effort. The proposed framework is implemented in such a modular way that it can be readily extended for new error-correcting codes or different simulation styles. Case studies showcase that the proposed framework allows for efficient evaluations of error-correcting codes depending on varying properties.

The remainder of the paper is organized as follows: In Section II, we review the basics of quantum computing and noise in quantum devices. Afterward, in Section III, we motivate the proposed framework by revisiting the basics of quantum error correction and illustrating the current problems. Section IV presents the proposed framework, followed by Section V demonstrating the application and usefulness of the framework. Finally, Section VI concludes the paper.

II. PRELIMINARIES

In this section, we review the basics of quantum computing as well as noise in quantum computing. We refer the interested reader to [17] for a thorough introduction to this topic.

A. Quantum Computing

The basic unit of information in the quantum world is called a *quantum bit* or *qubit*. Like a classical bit, a qubit can assume the states 0 and 1, but, additionally, it can also assume an arbitrary combination of those *basis states*. More generally—using Dirac notation—the state of a qubit $|\psi\rangle$ can be written as $|\psi\rangle = \alpha_0 \cdot |0\rangle + \alpha_1 \cdot |1\rangle$ with $\alpha_0, \alpha_1 \in \mathbb{C}$ and the normalization constraint $|\alpha_0|^2 + |\alpha_1|^2 = 1$. The values α_0, α_1 are called *amplitudes* and specify how strongly the qubit relates to each

¹The framework is available at www.github.com/cda-tum/qecc

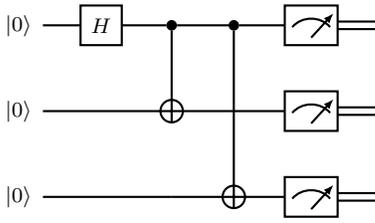


Fig. 1. Quantum circuit constructing the maximally entangled state

of the basis state $|0\rangle$ and $|1\rangle$. Measuring a qubit collapses it to a basis state. More precisely, with probability $|\alpha_0|^2$ the state $|\psi\rangle$ collapses to $|0\rangle$, and with probability $|\alpha_1|^2$ it collapses to $|1\rangle$. Only the basis state is returned by the measurement, the amplitudes cannot be directly measured. The state description can be extended for multi-qubit systems, e.g., a two-qubit state $|\phi\rangle$ is fully characterized by $2^2 = 4$ amplitudes and, hence, $|\phi\rangle = \alpha_{00} \cdot |00\rangle + \alpha_{01} \cdot |01\rangle + \alpha_{10} \cdot |10\rangle + \alpha_{11} \cdot |11\rangle$.

Having this formalism to describe quantum states, the next question becomes how to modify them. This is done by using so-called quantum operations. Important single-qubit operations are the Hadamard (H) operation, which transforms a qubit from a basis state into a superposition ($H|\psi\rangle = \alpha_0/\sqrt{2} \cdot (|0\rangle + |1\rangle) + \alpha_1/\sqrt{2} \cdot (|0\rangle - |1\rangle)$), the X operation, which is the quantum equivalent of flipping a bit ($X|\psi\rangle = \alpha_0 \cdot |1\rangle + \alpha_1 \cdot |0\rangle$), and the Z operation, which flips the phase of a qubit ($Z|\psi\rangle = \alpha_0 \cdot |0\rangle - \alpha_1 \cdot |1\rangle$). Two-qubit operations are also possible. An essential two-qubit operation is the *Controlled NOT* (CX), which flips the target qubit *iff* the control qubit is set to one ($CX|\phi\rangle = \alpha_{00} \cdot |00\rangle + \alpha_{01} \cdot |01\rangle + \alpha_{10} \cdot |11\rangle + \alpha_{11} \cdot |10\rangle$).

Quantum circuits are often illustrated as circuit diagrams. Diagrams consist of one or more vertical lines which represent qubits. Per convention, the top line represents qubit 0 and the remaining qubits are labeled sequentially. The lines are interrupted by operations that are applied to the respective qubit, which is usually indicated by labeled rectangles. In a quantum circuit, time flows from left to right, i.e., the leftmost operation on a qubit wire is applied first, and then the second left-most operation is applied, and so on.

Example 1. In Fig. 1, a quantum circuit is depicted that constructs a maximally entangled state over three qubits. The circuit consists of three qubits which are all initialized in basis state $|0\rangle$. Thus, the initial state is $|000\rangle$. Applying the H operation to qubit 0 (i.e., the topmost line) changes the state to

$$H|000\rangle = 1/\sqrt{2} \cdot (|000\rangle + |100\rangle)$$

The next operations are two CX gates. The \bullet symbol marks the control qubit and the \oplus specifies the target qubit. These yield

$$\begin{aligned} CX \ 1/\sqrt{2} \cdot (|000\rangle + |100\rangle) &= 1/\sqrt{2} \cdot (|000\rangle + |110\rangle) \text{ and} \\ CX \ 1/\sqrt{2} \cdot (|000\rangle + |110\rangle) &= 1/\sqrt{2} \cdot (|000\rangle + |111\rangle). \end{aligned}$$

Finally, all qubits are measured, indicated by the last operation on each lane. After the measurement, the qubits only contain classical information, which is indicated by the double wire. The final state is either $|000\rangle$ or $|111\rangle$, each with a probability of $1/2$. Note that, in this case, measuring one qubit affects the other qubits as well. This is an essential concept of quantum computing called entanglement.

B. Noise and Errors in Quantum Computing

Example 1 illustrates how *ideal* quantum computers behave. Unfortunately, *real* quantum computers are plagued by noise effects, which degrade computations and introduce errors. Although the error probability of quantum systems is constantly reduced, noise will always be a part of quantum computing [5].

In the literature, errors resulting from these noise effects are often distinguished between *operational errors* and *coherence errors* [18].

Operational errors are introduced whenever an operation is applied to a qubit [18]. They occur since quantum computers are mechanical apparatuses, susceptible to (tiny) errors. Hence, whenever an operation is applied to a qubit, it may be not executed at all, or in a (slightly) modified fashion. Operational errors are highly specific to each quantum computer, therefore they are often simulated using a depolarization error [19]. This error simulates that a qubit is set to a completely random state [17].

Coherence errors occur because of the fragile nature of qubits. This fragility leads to the problem that qubits can hold information only for a limited amount of time. Two types of coherence errors are usually distinguished [18]:

- Qubits can lose or gain energy from the environment, i.e., a qubit in a high-energy state ($|1\rangle$) tends to relax into a low-energy state ($|0\rangle$) and vice-versa. This is called *amplitude damping* or *T1 error*.
- When qubits interact with the environment it can happen that a phase flip effect occurs. This leads to a *phase-flip* or *T2 error*.

III. GENERAL IDEA AND MOTIVATION

As reviewed in Section II-B, quantum computers are plagued by noise, which drastically limits their usefulness in the real world. Quantum error correction tackles this problem, which makes it an essential part of building scalable and resilient quantum hardware. This section reviews the main ideas of the corresponding concepts and provides a motivation why an easy-to-use framework for their evaluation is needed.

A. Quantum Error Correction

To illustrate quantum error correction, we first revisit classical error correction, as it serves the same purpose. In the classical world, error correction is achieved by encoding information beyond its theoretical minimum. The redundant information is then used to identify and (if possible) correct errors. To illustrate this, consider the following example:

Example 2. Suppose a sender wants to transfer a single (classical) bit 0 to a receiver at a different location. The transmission channel is prone to bit-flip errors, which is why both sender and receiver agree to protect the bit, using the three-bit repetition code. For this code, each bit of information is encoded by tripling it, i.e., $0 \rightarrow 000$ and $1 \rightarrow 111$. Hence, the one-bit message 0 is encoded into the codeword ‘000’ before being transmitted to the receiver. During the transmission a bit-flip error occurs (distorting the message, e.g., to 001), the receiver could, through majority voting, still infer that the sent message most likely was 000 and could therefore correctly restore the original message 0.

In the quantum world, error correction is done similarly, but some properties of quantum mechanics make it more complex: A key difference between the quantum and the classical world is that measurements in the quantum world affect the observed system. So, while the classical system can be measured without risk of compromising the encoded information, special care must be taken in quantum error correction as to not destroy information by measuring it. This is worsened by the no-cloning theorem [20], which asserts that it is not possible to clone (i.e., copy) arbitrary quantum states. It is therefore not possible to simply clone qubits before measuring them to keep their information intact.

Instead, in order to add redundancy in the quantum world, we expand the Hilbert space, in which the information is encoded—effectively distributing the information of a single qubit among more qubits [6]. We illustrate this using the three-qubit bit-flip code, which allows detection and correction of single-qubit bit-flip errors. Using this code, a single qubit is

encoded by entangling it with two ancillary qubits (this can be achieved by two CX operations, similar to Example 1), i.e.,

$$|\psi\rangle = \alpha_0 \cdot |0\rangle + \alpha_1 \cdot |1\rangle \rightarrow |\psi_L\rangle = \alpha_0 \cdot |000\rangle + \alpha_1 \cdot |111\rangle. \quad (1)$$

After encoding the information of state $|\psi\rangle$, it is distributed among the 3-qubit state $|\psi_L\rangle$.

More precisely, the information of $|\psi\rangle$ (encoded into the 2-dimensional Hilbert space $\text{span}\{|0\rangle, |1\rangle\}$) is now encoded into the 8-dimensional Hilbert space $\text{span}\{|000\rangle, |001\rangle, \dots, |111\rangle\}$. This 8-dimensional Hilbert space can be split into 4 subspaces $\mathcal{C} = \text{span}\{|000\rangle, |111\rangle\}$, $\mathcal{F}_1 = \text{span}\{|001\rangle, |110\rangle\}$, $\mathcal{F}_2 = \text{span}\{|010\rangle, |101\rangle\}$, and $\mathcal{F}_3 = \text{span}\{|100\rangle, |011\rangle\}$, where the subspace \mathcal{C} represents the *logical code space* (indicating that the system is in a valid state) and the subspaces $\mathcal{F}_1, \mathcal{F}_2$, as well as \mathcal{F}_3 represent *logical error spaces* (indicating a bit-flip error). That is, \mathcal{F}_1 indicates that a bit-flip error in the first qubit occurred, \mathcal{F}_2 indicates a bit-flip at the second qubit, and \mathcal{F}_3 indicates a bit-flip in the third qubit. By using a special kind of measurement it is possible to infer which qubits are equal (we provide an example realization of such an indirect measurement later in Section IV-A). With this knowledge, it is possible to infer in which of the subspace $\mathcal{C}, \mathcal{F}_1, \mathcal{F}_2, \mathcal{F}_3$ $|\psi_L\rangle$ resides. Note that this measurement does not produce any information whether the encoded qubit is 0 or 1 and, therefore, does not change the encoded qubit.

Example 3. Suppose a sender wants to transfer the single qubit $|\psi'\rangle = \sqrt{1/3} \cdot |0\rangle + \sqrt{2/3} \cdot |1\rangle$ to a receiver at a different location. The transmission channel is prone to bit-flip errors, so sender and receiver agree to protect the qubit using the three-qubit bit-flip code (presented above). The sender encodes the qubit $|\psi'\rangle = \sqrt{1/3} \cdot |0\rangle + \sqrt{2/3} \cdot |1\rangle \rightarrow |\psi'_L\rangle = \sqrt{1/3} \cdot |000\rangle + \sqrt{2/3} \cdot |111\rangle$ and sends $|\psi'_L\rangle$ to the receiver.

During the transmission the first qubit flips—leaving it in state $\sqrt{1/3} \cdot |100\rangle + \sqrt{2/3} \cdot |011\rangle$. Before using the qubit, the receiver measures if all physical qubits making up $|\psi'_L\rangle$ are equal. The measurement shows that the first qubit is different from the other two qubits and, hence, the receiver infers that $|\psi'_L\rangle$ resides in the subspace \mathcal{F}_3 , indicating that a bit-flip error has occurred in the first qubit. The receiver therefore applies an X operation to the first qubit of $|\psi'_L\rangle$ —restoring it back to $|\psi'_L\rangle = \sqrt{1/3} \cdot |000\rangle + \sqrt{2/3} \cdot |111\rangle$. Afterwards, the receiver can decode $|\psi'_L\rangle$ to $|\psi'\rangle$ and use it for future computations.

While the concepts from above allow handling bit-flip errors (something that might have been sufficient in the classical world), the quantum world allows for a substantially larger continuum of errors. For example, recall the depolarization error introduced in Section II-B, which sets a qubit to a completely random state—this cannot be corrected using bit-flip correction only. Even worse, qubits also contain phase information, which has no classical counterpart at all and which must also be protected. Fortunately, it turns out that it is sufficient to correct only a subset of errors to correct *all* possible (unitary) errors that can occur in quantum computers [21]. In practice, it suffices to consider bit-flip and phase-flip errors only to have a universal error-correcting code [22]. In 1995 Peter Shor developed the first quantum error-correcting code that accomplished just that [6]. Since then, multiple types of quantum error-correcting schemes have been developed such as CSS codes ([6]–[8]), stabilizer codes ([9]), surface codes [10], [11], and adaptive error-correcting codes optimized for specific quantum hardware [12], [13].

B. Motivation

As shown above, quantum error correction addresses the problem of errors within quantum computations and there is a multitude of error correction schemes available. However, using these schemes is not as straightforward as Example 3

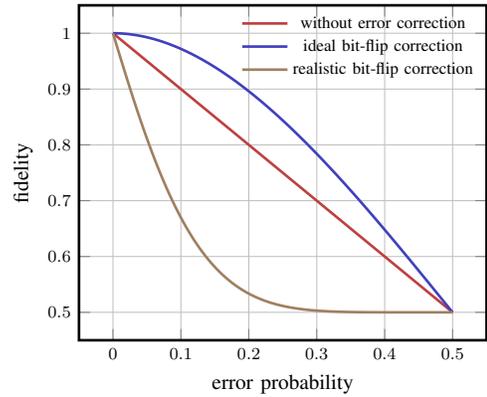


Fig. 2. State fidelity with increasing error probability

suggests. Noisiness is not restricted to specific erroneous quantum channels but affects *all* operations (including the operations necessary for error correction itself).

Example 4. The effects of noise are different depending on the simulated hardware model. To illustrate this, in Fig. 2, we plotted the fidelity² of state $|0\rangle$ depending on the probability of a bit-flip error in three different scenarios—without any kind of error protection (red curve), with the ideal bit-flip error protection presented in Example 3 (blue curve), and with a more realistic bit-flip error protection (brown curve). In the latter case, we assumed that the operations necessary to realize the bit-flip error correction are also affected by noise (this is reasonable to assume, as discussed in Section II-B). Without any kind of error correction, the fidelity of the encoded state more or less linearly declines with increasing error probability. Using the ideal bit-flip error protection increases the fidelity, as long as the error probability stays below 50%. Finally, in the realistic case the fidelity severely drops compared to the case where no error correction is applied at all.

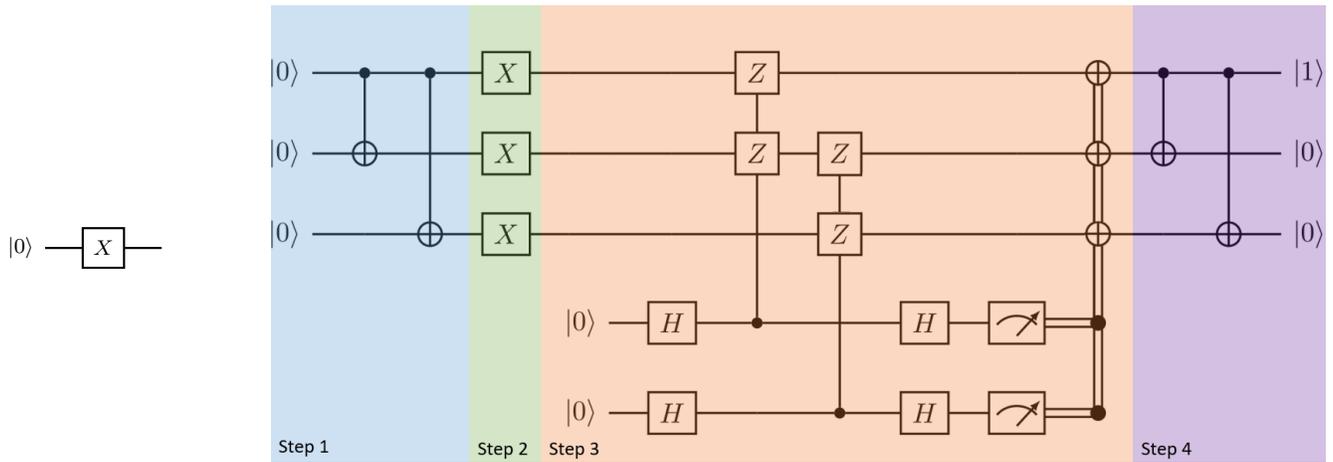
As illustrated by the example, naively applying error correction schemes to quantum circuits may even lower the fidelity (as proven by the threshold theorem [24]–[26]). The optimal error correction scheme depends on the properties of the simulated circuit as well as the considered quantum computer. Additionally, it is necessary to optimally tailor the error correction scheme to the specific use case. Thus a thorough evaluation is needed.

However, while there exists a lot of research on quantum error correction (e.g., [6]–[13]), the focus clearly (and understandably) is on the development of corresponding concepts for error-correcting codes and proof-of-concept implementations on selected quantum circuits. As illustrated above, however, evaluating error-correcting codes requires a broader consideration—thus far, involving extensive research and tedious manual implementation of said schemes. To address this problem, we propose a framework that automates the process of applying error correction to circuits and also allows for noise-aware quantum circuit simulation.

IV. PROPOSED FRAMEWORK

The proposed framework is supposed to support the process of evaluating error correction schemes. To this end, it does not only allow to apply error-correcting codes to a given quantum circuit but also supports the entire process of circuit preparation and simulation. More precisely, the framework (1) automates the circuit compilation flow (including the application of error-correcting codes), (2) provides quantum circuit simulation with noise-aware quantum circuit simulators, and (3) is flexible and easily extensible for new use cases.

²The fidelity is a metric that expresses the overlap between two states [23]. Here we use it to express the overlap between the state with and without errors.



(a) Original circuit

(b) Circuit with bit-flip error correction

Fig. 3. Applying the bit-flip error-correcting code to a circuit.

Note that, while circuit compilation and simulation are necessary parts of such a framework, there already exist available solutions (e.g., [19], [27]–[30]). Therefore, the proposed framework is realized in such a manner that it can be coupled with these solutions. In the remainder of this section, the proposed framework is presented.

A. Compiler

During the compilation process, the quantum circuit is prepared for the circuit simulation or execution on a real quantum computer. Here, we focus on the main functionality that is provided by the framework, namely the automatic application of error-correcting codes to the quantum circuit. For this, we exploit the fact that, while there exist many different error correction schemes, all of them must in practice realize the same steps when applied to a quantum circuit. These steps are:

1. *Qubit Encoding*: The basic idea of all error-correcting codes is that information is encoded in such a way that it allows detection and correction of occurring errors. While encoding depends on the respective approach, all schemes have in common that information is distributed among multiple qubits. At the end of this step, all qubits within the logical circuit are encoded into logical qubits consisting of multiple physical qubits.

2. *Operation Encoding*: Having a quantum circuit with encoded qubits, the next step accordingly adjusts the operations. More precisely, the originally intended operations are mapped to logical operations, i.e., operations that are functionally equivalent but consider the applied qubit encoding. How each operation is mapped is not always trivial [16] and depends on the used error correction scheme, which is why only specific quantum operations are supported for each error-correcting code.

3. *Error Correction*: The error detection and correction routines are added to the circuit. This consists of syndrome extraction followed by some corrective operations on the logical qubits—depending on the extracted syndrome. The optimal frequency of executing the detection and correction routine, i.e., after how many quantum operations this routine is applied to a logical qubit, depends on several factors, e.g., the hardware specifications of the designated quantum computer or the error-proneness of the application. To accommodate this, the frequency of the detection and correction step can be freely adjusted.

4. *Qubit Decoding*: Finally, in the last step, the logical qubit is decoded back to a single physical qubit, so that it can be measured. For convenience, this is realized in such a way, that the ordering of the output qubits is equal to the order of the qubits in the original circuit.

The following example illustrates the process of applying an error-correcting code using the bit-flip code presented in Example 3.

Example 5. In Fig. 3, a circuit without (Fig. 3a) and with (Fig. 3b) bit-flip error correction is presented. The color-coding in the circuit with error correction represents during which step the respective parts are generated. During the encoding, each qubit of the original circuit is encoded into a logical qubit by entangling it with two ancillary qubits. Next, each quantum operation is copied for each qubit making up the logical qubit, e.g., an X operation in the original circuit is mapped to an X operation onto each qubit making up the respective logical qubit. Afterwards, error detection and correction is added. For the error detection, a projective measurement is added, which is realized by two ancillary qubits then contain the syndrome, which indicates whether a bit-flip error has occurred. If that is the case, the flipped qubit is corrected by applying an X operation to it. Finally, the qubit is decoded so that it can be measured.

Since these steps are necessary for all error-correcting codes, the framework provides templates for realizing them. Due to this modular design, the framework can easily be extended, either by modifying already available error-correcting codes or by adding new ones. Currently the framework supports four error-correcting codes ranging in size to different error correction types, namely the Shor Code [6] (as this is the first code that corrects arbitrary single-qubit errors), the Laflamme Code [31] (as this is the smallest possible code protecting against arbitrary single-qubit errors), the Steane code [7] (as it is a well-studied code used in other experiments such as [13]), and a surface code [32], [33] (due to their popularity in the current research on error-correcting codes).

B. Simulator

An essential part of the framework is the ability to directly simulate circuits with (and without) error correction and different hardware models. To this end, several simulation approaches have recently been introduced in the literature (e.g., [19], [27]–[30]). Some of these are available as open-source, and hence can be readily integrated into the proposed framework. By this, we gain access to multiple simulation styles, which come with their advantages/drawbacks and the most appropriate simulator can be selected for each use case.

To illustrate these differences, we simulated the entanglement circuits (which construct the GHZ state over all qubits) with an increasing number of qubits. During all simulations,

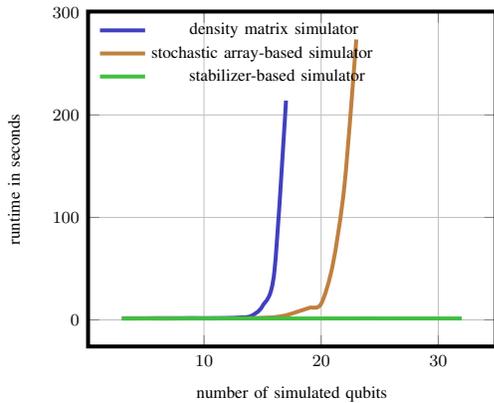


Fig. 4. Runtime of noise-aware quantum circuit simulation depending on the simulation style

we applied depolarization noise with 0.001 % probability to the qubit whenever it has been used. We ran this experiment with Qiskit [19] using three different simulation styles, namely, a density matrix simulator, a stochastic array-based simulator (with 2000 shots), and a stabilizer-based simulator (with 2000 shots). In Fig. 4 the simulation times are plotted as a function of the number of simulated qubits. The plot shows that the runtime of the density matrix simulator and the stochastic array-based simulator grows exponentially with increasing qubits, while the stabilizer-based simulator exhibits a polynomial runtime. However, while the density matrix simulator performs worst in runtime, it allows to simulate arbitrary quantum operations and gives a *full* description of the final state. The stochastic array-based simulator also allows to simulate arbitrary quantum operations but only approximates the final quantum state. In contrast, the stabilizer-based simulator is restricted to a non-universal gate set of quantum operations (i.e., the Clifford gate set).

This small experiment shows the advantage of having a variety of different quantum circuit simulators available for experiments. Currently Qiskit [19] is used for quantum circuit simulation. We chose Qiskit as it offers not only different noise-aware quantum circuit simulators but also access to real quantum hardware. Naturally, the quantum circuit simulator can easily be switched to different software.

V. APPLICATION AND DEMONSTRATION

We demonstrate the usefulness of the proposed framework by considering the Steane code [7] as a well-known representative of an error-correcting code. For this code, we evaluated in which scenarios the error correction actually improves the reliability and in which it does not³. To this end, we evaluated the effects of three different properties:

- 1) Size of the considered circuit: error correction is a costly procedure (see Fig. 3). Whether error correction improves the reliability of a circuit substantially depends on its size.
- 2) Frequency of error correction: Detecting and correcting errors is a costly operation. The proper frequency with which this routine is performed is critical when applying an error-correcting code.
- 3) Assumed hardware model: The considered hardware model (i.e., error types and error probability) also effects the usefulness of the quantum circuit. For example, if the error probability is too high, the additional operations introduced by error correction may actually degrade the quality of the results.

For the evaluation, we considered a benchmark that creates the GHZ state between all qubits (taken from [34]). We chose this benchmark because its simplicity makes it the ideal

³Using the proposed framework, similar considerations can be done with other codes as well.

candidate to test the Steane code in different scenarios, without having to account for fragments within the data. Besides that, the benchmark can be easily varied in size, and the created states are strongly correlated thereby making (uncorrected) errors easily noticeable. If not stated otherwise, the following default parameters have been used in the evaluation:

- Depolarization noise (applied whenever a qubit is used) has been considered with an error probability of 0.001 %.
- Error correction and detection has been applied to a logical qubit prior to the measurement and whenever it has been used 500 times.
- The entanglement benchmark is used with five qubits and 10.000 dummy operations added to the end of the circuit (the dummy operations do not change the result but are affected by noise).

For each property discussed above, this default use-case is adapted to evaluate the respective property. That is, for (1) we simulated entanglement circuits with an increasing number of dummy operations, for (2) we simulated entanglement circuits with a decreasing rate of error correction steps, and, finally, for (3) we simulated the benchmark with increasing error probability not only with depolarization noise (mimicking gate errors) but also with amplitude damping errors (mimicking coherence errors). In the latter case we used the entanglement benchmark with two qubits (instead of five), as we had to use a stochastic array-based simulator for simulating the amplitude damping noise.

We used the proposed framework to implement the Steane code to the circuits for all runs and simulated the circuits using Qiskit [19]—which, as described in Section IV-B, has been integrated into the framework. More precisely, we used the stabilizer-based simulator when considering depolarization errors and the state vector simulator when considering amplitude damping errors (as these cannot be simulated with the stabilizer simulator). We simulated all experiments stochastically with 2000 shots. The simulations yield a probability distribution for measuring specific basis states, which we use to calculate the “classical” fidelity (also known as the Hellinger coefficient [35]) between the expected probability distribution (i.e., the received probability distribution when no errors occur during the computation) and the actual distribution received by the noise-aware quantum circuit simulation. We chose this distance measure over others, such as the fidelity between the quantum states [23] or the average fidelity (as used for example in [36]), as we want to evaluate how a noisy quantum computer affects the “correctness” of the measured final states for *specific* quantum algorithms. For these evaluations, we conducted 285 simulations and applied error correction to the quantum circuits 137 times.

Fig. 5 summarizes the results of the experiments. Fig 5a depicts the fidelity with increasing depth of the entanglement circuits. It shows that the usefulness of error correction increases with the gate depth. Conducting error correction only pays off when the circuit is of sufficient depth. Next, in Fig. 5b the fidelity is plotted with decreasing error correction frequency. Contrary to intuition, it is not advantageous to correct errors as often as possible during the computation. This can be explained by the fact that, as discussed above, error correction is expensive and adds a substantial amount of extra operations. Hence, it has a positive effect if error correction application and resulting overhead is traded-off (a task for which the proposed framework offers a very helpful tool). Finally, Fig. 5c depicts the fidelity with increasing (physical) error probability and simulated error type. The plot shows that error correction is only useful, as long as the error probability stays low. Additionally, the simulated noise type noticeably affects the effectiveness of the error correction scheme. Thus, depolarization errors could be handled more effectively compared to amplitude damping errors.

Obviously, these demonstrations only provide a snippet of what kind of evaluations can be conducted with the proposed framework. But it clearly demonstrates the usefulness of the

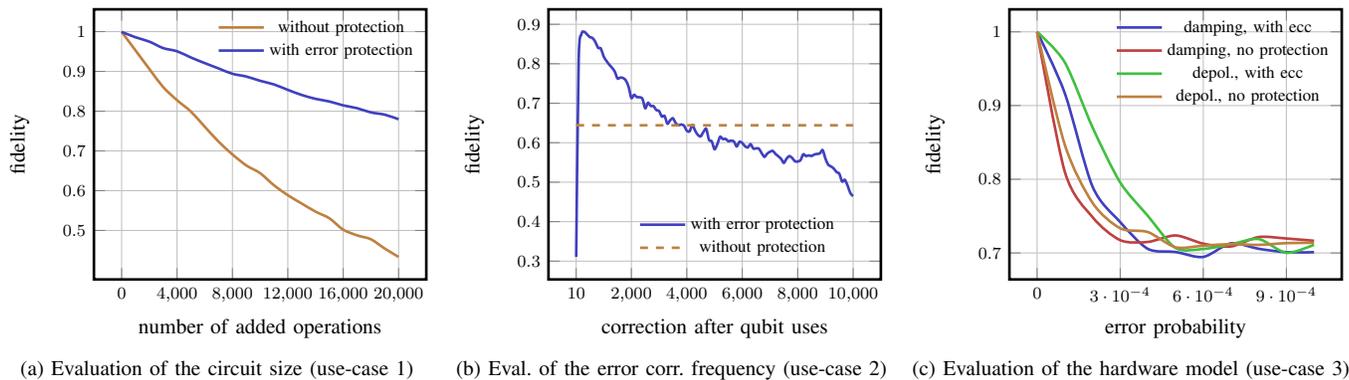


Fig. 5. Experiments

approach. While, until now, it required a lot of tedious manual work to apply error correction to circuits and evaluate the corresponding effects, now the proposed framework automates the task—making quantum error correction more broadly applicable and easier to analyze.

VI. CONCLUSIONS

Quantum error-correcting codes are an essential part of building scalable and resilient quantum hardware. However, currently most of the corresponding work in this domain heavily relies on manual labor and/or is based on theoretical results only. We address this problem by proposing an open-source framework that automates the process of applying error correction and also allows noise-aware simulation of the protected circuits. We demonstrate the advantages and usefulness of the proposed framework by evaluating the reliability of the Steane error-correcting code with respect to different parameters. The proposed framework is published as open-source (available at www.github.com/cda-tum/qecc) and implemented in a modular fashion, so that it can be easily configured and extended for new error-correcting codes and to make quantum error correction broadly applicable.

ACKNOWLEDGMENTS

This work received funding from the University of Applied Sciences PhD program of the State of Upper Austria (managed by the FFG), from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation program (grant agreement No. 101001318), 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.

REFERENCES

- [1] S. McArdle, S. Endo, A. Aspuru-Guzik, S. C. Benjamin, and X. Yuan, “Quantum computational chemistry,” *Reviews of Modern Physics*, vol. 92, no. 1, pp. 1–59, 2020.
- [2] J. Biamonte, P. Wittek, N. Pancotti, P. Rebentrost, N. Wiebe, and S. Lloyd, “Quantum machine learning,” *Nature*, vol. 549, no. 7671, pp. 195–202, 2017. arXiv: 1611.09347.
- [3] B. A. Cordier, N. P. D. Sawaya, G. G. Guerreschi, and S. K. McWeeney, “Biology and medicine in the landscape of quantum advantages,” 2021. arXiv: 2112.00760.
- [4] D. J. Egger *et al.*, “Quantum Computing for Finance: State-of-the-Art and Future Prospects,” *IEEE Transactions on Quantum Engineering*, vol. 1, pp. 1–24, 2021. eprint: 2006.14510.
- [5] J. Preskill, “Quantum Computing in the NISQ era and beyond,” *Quantum*, vol. 2, p. 79, 2018.
- [6] P. W. Shor, “Scheme for reducing decoherence in quantum computer memory,” *Physical Review Applied*, vol. 52, R2493–R2496, 4 1995.
- [7] A. M. Steane, “Error correcting codes in quantum theory,” *Phys. Rev. Lett.*, vol. 77, pp. 793–797, 5 1996.
- [8] A. R. Calderbank and P. W. Shor, “Good quantum error-correcting codes exist,” *Physical Review Applied*, vol. 54, pp. 1098–1105, 2 1996.
- [9] D. Gottesman, “Class of quantum error-correcting codes saturating the quantum hamming bound,” *Physical Review Applied*, vol. 54, pp. 1862–1868, 3 1996.

- [10] S. B. Bravyi and A. Y. Kitaev, “Quantum codes on a lattice with boundary,” no. 96, pp. 1–6, 1998. arXiv: 9811052.
- [11] E. Dennis, A. Kitaev, A. Landahl, and J. Preskill, “Topological quantum memory,” *Journal of Mathematical Physics*, vol. 43, no. 9, pp. 4452–4505, 2002. arXiv: 0110143.
- [12] A. S. Fletcher, P. W. Shor, and M. Z. Win, “Structured near-optimal channel-adapted quantum error correction,” *Physical Review Applied*, vol. 77, p. 012320, 1 2008.
- [13] A. Robertson, C. Granade, S. D. Bartlett, and S. T. Flammia, “Tailored Codes for Small Quantum Memories,” *Physical Review Applied*, vol. 8, no. 6, pp. 1–6, 2017. arXiv: 1703.08179.
- [14] M. Hsieh and F. Le Gall, “Np-hardness of decoding quantum error-correcting codes,” *Physical Review Applied*, vol. 83, p. 052331, 5 2011.
- [15] L. Berent, L. Burgholzer, and R. Wille, “Software Tools for Decoding Quantum Low-Density Parity Check Codes,” in *Asia and South Pacific Design Automation Conf.*, 2023. eprint: 2209.01180.
- [16] B. Eastin and E. Knill, “Restrictions on transversal encoded quantum gate sets,” *Physical review letters*, vol. 102, p. 110502, 2009.
- [17] M. Nielsen and I. Chuang, *Quantum Computation and Quantum Information*. Cambridge Univ. Press, 2000.
- [18] S. S. Tannu and M. K. Qureshi, “Not All Qubits Are Created Equal,” *Proceedings of the Twenty-Fourth International Conference on Architectural Support for Programming Languages and Operating Systems*, pp. 987–999, 2019.
- [19] G. Aleksandrowicz *et al.*, “Qiskit: An open-source framework for quantum computing,” *Zenodo*, 2019.
- [20] W. K. Wootters and W. H. Zurek, “A single quantum cannot be cloned,” *Nature*, vol. 299, no. 5886, pp. 802–803, 1982.
- [21] E. Knill and R. Laflamme, “Theory of quantum error-correcting codes,” *Physical Review Applied*, vol. 55, pp. 900–911, 2 1997.
- [22] J. Roffe, “Quantum Error Correction: An Introductory Guide,” *Contemporary Physics*, vol. 60, no. 3, pp. 226–245, 2019.
- [23] R. Jozsa, “Fidelity for mixed quantum states,” *Journal of Modern Optics*, vol. 41, no. 12, pp. 2315–2323, 1994.
- [24] E. Knill, R. Laflamme, and W. H. Zurek, “Resilient quantum computation,” *Science*, vol. 279, no. 5349, pp. 342–345, 1998.
- [25] D. Aharonov and M. Ben-Or, “Fault-tolerant quantum computation with constant error rate,” 4, vol. 38, 2008, pp. 1207–1282.
- [26] A. Y. Kitaev, “Quantum computations: Algorithms and error correction,” *Russian Mathematical Surveys*, vol. 52, no. 6, pp. 1191–1249, 1997.
- [27] C. Gidney, “Stim: A fast stabilizer circuit simulator,” *Quantum*, vol. 5, p. 497, 2021.
- [28] T. Grurl, J. Fuß, and R. Wille, “Noise-aware Quantum Circuit Simulation With Decision Diagrams,” *IEEE Trans. on CAD of Integrated Circuits and Systems*, 2022.
- [29] Cirq Developers, “Cirq,” version v0.14.1, *Zenodo*, 2022.
- [30] G. G. Guerreschi, J. Hogaboam, F. Baruffa, and N. P. D. Sawaya, “Intel Quantum Simulator: A cloud-ready high-performance simulator of quantum circuits,” *Quantum Sci. Technol.*, vol. 5, p. 034007, 2020.
- [31] R. Laflamme, C. Miquel, J. P. Paz, and W. H. Zurek, “Perfect quantum error correcting code,” *Physical Review Letters*, vol. 77, no. 1, pp. 198–201, 1996. arXiv: 9602019.
- [32] J. R. Wootton, A. Peter, J. R. Winkler, and D. Loss, “Proposal for a minimal surface code experiment,” *Physical Review A*, vol. 96, no. 3, 2017. arXiv: 1608.05053.
- [33] A. G. Fowler, A. M. Stephens, and P. Groszkowski, “High-threshold universal quantum computation on the surface code,” *Physical Review Applied*, vol. 80, p. 052312, 5 2009.
- [34] N. Quetschlich, L. Burgholzer, and R. Wille, “MQT Bench: Benchmarking software and design automation tools for quantum computing,” 2022.
- [35] A. Bhattacharyya, “On a measure of divergence between two multinomial populations,” *Sankhyā: The Indian Journal of Statistics*, vol. 7, no. 4, pp. 401–406, 1946.
- [36] A. Gilchrist, N. K. Langford, and M. A. Nielsen, “Distance measures to compare real and ideal quantum processes,” *Physical Review Applied*, vol. 71, p. 062310, 6 2005.