

# Digital Zero-Noise Extrapolation with Quantum Circuit Unoptimization

Elijah Pelofske<sup>\*1</sup> and Vincent Russo<sup>†2</sup>

<sup>1</sup>Los Alamos National Laboratory, Information Systems & Modeling

<sup>2</sup>Unitary Foundation

## Abstract

Quantum circuit unoptimization is an algorithm that transforms a quantum circuit into a different circuit that uses more gate operations while maintaining the same unitary transformation. We demonstrate that this method can implement digital zero-noise extrapolation (ZNE), a quantum error mitigation technique. By employing quantum circuit unoptimization as a form of circuit folding, noise can be systematically amplified. The key advantages of this approach are twofold. First, its ability to generate an exponentially increasing number of distinct circuit variants as the noise level is amplified, which allows noise averaging over many circuit instances with slightly different circuit structure which mitigates the effect of biased error propagation because of the significantly altered circuit structure from quantum circuit unoptimization, or highly biased local noise on a quantum processor. Second, quantum circuit unoptimization by design resists circuit simplification back to the original unmodified circuit, making it plausible to use ZNE in contexts where circuit compiler optimization is applied server-side. We evaluate the effectiveness of quantum circuit unoptimization as a noise-scaling method for ZNE in two test cases using depolarizing noise numerical simulations: random quantum volume circuits, where the observable is the heavy output probability, and QAOA circuits for the (unweighted) maximum cut problem on random 3-regular graphs, where the observable is the cut value. We show that using quantum circuit unoptimization to perform ZNE can approximately recover signal from noisy quantum simulations.

## 1 Introduction

Quantum error mitigation is a class of algorithms and techniques whose goal is to suppress errors in quantum computations on near-term noisy quantum computers [1–15]. Generally, error suppression can refer to techniques that do not incur any computational overhead, such as dynamical decoupling [16–22]. Of the techniques that require classical post-processing, there are several families of techniques that have been proposed, all of which typically refer to quantum error mitigation. Almost all post-processing-based quantum error mitigation strategies require a real number observable quantity to be measured from the samples obtained by the quantum computer—with some exceptions that use sampling-based methods and knowledge of the error model of the quantum hardware to also extract error-mitigated samples [23]. In general, however, the resources required to implement quantum error mitigation have an exponential overhead [24, 25]—and therefore these techniques are always a type of heuristic approximation and will likely not scale to large quantum circuit sizes due to their resource requirement scaling, but can improve the signal measured from current noisy quantum computers.

Numerous quantum error mitigation algorithms have been developed, including probabilistic error cancellation (PEC) [26–28], Clifford data regression (CDR) [29], and machine learning-based approaches [6, 30–33]. This study focuses on zero-noise extrapolation (ZNE) [3, 5, 7, 12, 14, 26], a quantum error mitigation technique that amplifies the noise in a quantum computation by a specified factor. The resulting measurements contain a controlled amount of noise or error. An observable quantity is extracted from these measurements, and an extrapolation is performed to estimate the zero-noise limit. ZNE in comparison to many other quantum error mitigation algorithms has a lower computational and sampling overhead, making it more accessible to use on current quantum computers. To summarize, ZNE is comprised of the following components:

1. A noise amplification method. This can be done in various ways, but the required property is that the quantum circuit is executed with a quantifiable amount of noise introduced into the computation.
2. The computation of interest must have a numerical quantity, an observable, that can be extracted from the measured qubit states (i.e., samples).

---

\*Email: epelofske@lanl.gov

†Email: vincent@unitary.foundation

3. A regression curve fits the noise-amplified observable data and then can be used to extrapolate that noise observable fit to the zero-noise point.

Generally, there are two methods for scaling the noise. The first is digital, meaning in the circuit representation of the unitary that is being simulated—typically, this comes in the form of an increased number of circuit instructions that are still implementing an identical unitary (under noiseless conditions). The second is by stretching the pulse-level waveforms on the hardware [5, 34] that are implementing the circuit-level instructions—in particular by still implementing the same operations, but over longer timescales, which thus creates more noise due to the limited qubit coherence times and an accumulation of errors over time.

In this study, we demonstrate that digital zero-noise extrapolation, specifically, the noise amplification portion of the algorithm, can be implemented using quantum circuit unoptimization, which was introduced in [35] to test the effectiveness of quantum circuit compiler optimization. Quantum circuit unoptimization has several valuable properties compared to the existing digital noise amplification methods.

First, quantum circuit unoptimization enables the generation of an exponentially growing number of potential circuit variants as the circuit noise is amplified. Circuit variant here means a circuit with a different set of elementary quantum gates, but still implements exactly the same unitary operation. In contrast, other methods, such as global unitary circuit folding, generate exactly one circuit per noise scale factor. This exponential growth in possible circuit variants can be particularly advantageous when the noise profile of the quantum hardware is highly variable or non-uniform. By averaging over multiple circuit variants—each representing a slightly different implementation of the same logical unitary operation—one can mitigate the impact of localized noise effects. This circuit variant averaging also mitigates noise biasing that can occur from the inherent differences in quantum circuit structure from the unoptimization routine. This approach reduces bias in the measured observables, as the diverse noise profiles, or specifically physical implementations, of the circuit effectively “sample” different noise environments, leading to more robust results.

Second, the unoptimization routines of [35] are designed to resist removal by compiler optimization passes in standard quantum compilers, such as Qiskit, at least with current versions. Compiler optimization often aims to simplify quantum circuits by eliminating redundant operations or reducing gate counts, which can inadvertently reverse noise amplification steps in techniques like unitary folding. Quantum circuit unoptimization circumvents this issue by introducing structural complexity that compilers do not recognize or simplify. As a result, the noise amplification achieved through quantum circuit unoptimization is more resilient to server-side circuit optimization, ensuring that the intended scaling remains intact during execution.

Third, quantum circuit unoptimization allows for fractional noise scale factors. While this is not a property unique to quantum circuit unoptimization, it is a significant advantage for ZNE. Fractional noise scaling enables finer control over the noise amplification process, particularly for scale factors close to 1.0. This capability is especially valuable in highly noisy computation scenarios where slight deviations from the original circuit may already introduce significant noise. Unoptimization facilitates more accurate extrapolation to the zero-noise limit by enabling precise control over noise scaling.

## 2 Methods

### 2.1 Quantum Circuit Unoptimization

The primary algorithm employed in this study is *quantum circuit unoptimization*, a novel quantum circuit compiler algorithm introduced in [35]. This algorithm systematically transforms a given quantum circuit into a more complex version while preserving its unitary equivalence under ideal, noiseless conditions. By doing so, quantum circuit unoptimization provides a mechanism for controlled noise amplification. The unoptimization process follows a four-step iterative “recipe” that is applied repeatedly to the input quantum circuit. Each iteration increases the circuit’s depth and gate count, introducing additional opportunities for noise while preserving the original computation.

At a high level, for a given quantum circuit, the recipe steps are:

1. Insert: Insert a two-qubit gate  $A$  and its Hermitian conjugate  $A^\dagger$  between two 2-qubit gates in the circuit denoted as  $B_1$  and  $B_2$ . We choose  $A$  such that it is a random two-qubit unitary.
2. Swap: Swap the  $B_1$  gate with the  $A^\dagger$  gate in the circuit, replacing  $A^\dagger$  with  $\widetilde{A^\dagger}$ .
3. Decompose: Decompose multi-qubit unitary gates into elementary gates.
4. Synthesize: Apply quantum circuit synthesis.

For the insert step, a random two-qubit unitary  $A$  and its Hermitian conjugate  $A^\dagger$  are introduced into the circuit. These gates are inserted between two existing two-qubit gates in the circuit, denoted as  $B_1$  and  $B_2$ . The insertion

of  $A$  and  $A^\dagger$  does not alter the logical functionality of the circuit, as the combination  $A^\dagger A = I$ . The selection of  $B_1$  and  $B_2$  can follow different strategies, such as choosing gates that share a common qubit (the *concatenated strategy*) or selecting them randomly (the *random strategy*). The choice of strategy affects the unoptimization procedure’s characteristics and its impact on circuit structure. The concatenated strategy, as we implement it, is specifically a deterministic two 2-qubit gate selection procedure, where we choose the first pair of two-qubit gates that share a common qubit. This means specifically that for the concatenated strategy, there is no randomness in the selection of this pair of 2-qubit gates. However, one could certainly utilize random selection of pairs of two qubit gates that share a common qubit.

After insertion, the  $B_1$  gate is swapped with the  $A^\dagger$  gate. This swapping operation introduces a new gate, denoted as  $\widetilde{A^\dagger}$ , which encapsulates the interaction between  $B_1$  and  $A^\dagger$ . Specifically,

$$\widetilde{A^\dagger} = (I \otimes B_1)^\dagger (I \otimes A^\dagger) (B_1 \otimes I), \quad (1)$$

representing the transformation of  $A^\dagger$  under the influence of  $B_1$ . This step introduces additional complexity in the circuit by altering the placement and nature of the gates while maintaining unitary equivalence.

Now containing additional gates from the first two steps, the circuit undergoes a decomposition and gate synthesis process. These last two steps can be thought of as effectively grouped together because this is really just a circuit compilation routine to compile down to some quantum instruction gateset. Critically, these steps are expressing the unoptimized circuit, which is now comprised of Multi-qubit unitary gates, including  $\widetilde{A^\dagger}$ , are broken down into elementary gates, such as single-qubit rotations and CNOT gates. This decomposition is performed using Qiskit [36], specifically a single decomposition pass. The final step optimizes the decomposed circuit using a synthesis procedure, in this case we implement this step in the form of the Qiskit transpiler [36], with optimization level 3 (which is the maximum), targeting a particular gateset. If this circuit is being implemented on a specific quantum computing architecture, then the native gateset of that quantum computer would be used as the target gateset when performing these decomposition and synthesis steps.

The unoptimization recipe can be applied iteratively to a quantum circuit with each application subsequently increasing the circuit’s depth, increasing the total number of instructions, and altering the overall circuit structure. Let  $R(C)$  represent the application of the unoptimization recipe to a circuit  $C$ . After  $i$  iterations, the resulting circuit can be expressed recursively as

$$R_i(C) := R(R(\dots R(C)\dots)), \quad (2)$$

where the recipe is applied  $i$  times. Each application of the recipe corresponds to a noise scale factor,  $\lambda_i$ , which quantifies the ratio of the depth of the original circuit to the depth of the circuit after  $i$  iterations. Formally, the noise scale factor is defined as

$$\lambda_i = \frac{n(C)}{n(R_i(C))}, \quad (3)$$

where  $n(C)$  and  $n(R_i(C))$  denote the total number of single- and two-qubit gates in the original and the  $i$ -th unoptimized circuits, respectively. This metric, rather than circuit depth alone, provides a more comprehensive measure of circuit complexity and the opportunities for noise to accumulate. Importantly, eq. (3) is only an approximation of the true noise amplification because it does not capture the differences in the circuit structure and error propagation. We use this metric because the circuit depth ratio corresponds to noise amplification and is easy to compute, whereas comprehensive error propagation measurement can be computationally intensive.

Quantum circuit unoptimization is used in this study as the noise-scaling mechanism within the zero-noise extrapolation (ZNE) error mitigation framework. The process is depicted in Figure 1, using an example input circuit of a 4-qubit graph state. After applying the unoptimization recipe iteratively to generate circuits with increasing noise levels, the resulting circuits are executed on a quantum device or simulator. For each circuit  $R_i(C)$ , the expectation value of a target observable is measured. The corresponding noise scale factor  $\lambda_i$  is associated with each measurement. Using these data points, an extrapolation method (e.g., linear, Richardson, or exponential extrapolation) is applied to estimate the zero-noise limit of the observable.

## 2.2 Test Cases

We evaluate quantum circuit unoptimization for ZNE using two different classes of quantum circuits. The first class consists of quantum volume (QV) circuits, where the observable is the heavy output probability (HOP). Quantum volume is a benchmark for noisy quantum computers that requires coherent sampling of dense, random quantum circuits of a specific form [37–40], where the goal of the benchmark is that the noisy quantum computer must




Figure 1: Depiction of quantum circuit unoptimization applied to ZNE. Panel **a)** shows a small, but representative, example of quantum circuit unoptimization being used. The example is a 4-qubit fully connected graph state circuit which is the input for a single round of the circuit unoptimization algorithm, using the concatenated pair of two qubit gates selection strategy (the common qubit is  $q_0$ ). Unoptimization consists of gate insert, swap, decompose, and synthesis operations. The right-most circuit shows the final state of the original circuit after a single round of unoptimization has been applied to it. The input and resulting circuits are unitarily equivalent to one another, which remains true even after repeated rounds of unoptimization. Recursive application of quantum circuit unoptimization results in a growing combinatorial explosion of possible circuit (equivalent) variants with different circuit structure as more iterations are applied. This large number of potential circuit variants is due to random choice of two 2-qubit gates during the Insert step (this is only true for the random strategy), and because the two qubit gate  $A$  in the Insert step is chosen as a random unitary. Measurement gates are not shown for brevity in the circuit renderings, but in this case the state of all 4 qubits at the end of the circuit are measured. Additional information about the gates used in the Qiskit-generated circuit diagrams are provided in the Appendix A. Panel **b)** shows a conceptual diagram of using quantum circuit unoptimization to amplify noise and then carry out zero noise extrapolation, given an arbitrary input circuit  $C$  consisting of single and two-qubit gates. For some choice of integers  $i_1 < i_2 < \dots < i_n$ , we apply  $i_k$  iterations of the recipe to the circuit  $C$ , denoted as  $R_{i_k}(C)$  for some  $k \in \{1, \dots, n\}$ . For each resulting circuit from  $R_{i_k}(C)$ , we obtain a corresponding noise-scale factor as a function of the resulting circuit's depth and the original circuit's depth. The greater the number of recipe iterations corresponds to a circuit with increasing depth (and hence increased noise). We compute the expectation value of each noise-scale factor and extrapolate to the zero-noise limit via some extrapolation method (e.g., linear, Richardson, exponential, etc.).

reliably sample a quantum circuit that has approximately the same depth (circuit depth) as its width (number of qubits being used in the circuit). ZNE was previously applied to quantum volume in a technique known as *effective quantum volume* [41, 42], which extends the quantum volume benchmark to include reliable estimation of noise mitigated quantities when the circuit is executed on a noisy quantum computer.

QV circuits are constructed such that the number of layers is equal to the number of qubits, creating random dense circuits with an inherent symmetry in depth and width. These circuits exhibit a key property: under ideal conditions, they concentrate amplitude into a subset of computational basis states, leading to a heavy output probability of approximately  $\frac{1+\ln(2)}{2} \approx 0.85$  in the asymptotic limit [43]. However, as decoherence and noise accumulate, this amplitude concentration diminishes, and the heavy output probability approaches 0.5, corresponding to uniform random sampling. Therefore, the heavy output probability measure is defined over the range  $[0, 1]$ .

The intuition behind the QV benchmark is that it requires the noisy quantum hardware to be able to sample dense circuits that have the same number of qubits as the gate depth is—meaning that the benchmark does not disproportionately favor very shallow-depth circuits or very high-depth circuits acting on a few qubits. Quantum volume, however, requires full classical state-vector simulation of each circuit, meaning that this is not scalable into a regime where full classical simulation of the quantum system can not be performed.

The second class of quantum circuits we evaluate consists of those used in the Quantum Approximate Optimization Algorithm (QAOA) [44, 45], specifically applied to the Maximum Cut (Max-Cut) problem on random 3-regular graphs [46–53]. Max-Cut is a well-known NP-hard combinatorial optimization problem [54–56].

QAOA circuits consist of alternating simulations of layers of a cost Hamiltonian  $H_C$ , called the phase separator, (which is derived from the problem’s objective function) and a mixer Hamiltonian  $H_M$ . The mixer gives parameterized interference, or state transitions, between solutions of different cost value. The number of layers, denoted as  $p$ , determines the depth of the circuit. The two Hamiltonians  $H_C$  and  $H_M$  do not commute, and each is parameterized by a set of real numbers  $\vec{\gamma}$  and  $\vec{\beta}$ , respectively, referred to as *QAOA angles*.

QAOA aims to approximate ground states of diagonal (classical) cost Hamiltonians corresponding to combinatorial optimization problems. However, for QAOA to perform well, a sufficiently large depth  $p$  is required, and the QAOA angles must be carefully chosen. These parameters are typically optimized variationally, or designed to approximate adiabatic quantum evolution [57]. In this study, we use the fixed-angle QAOA parameters proposed in [58], which are known to perform well on random 3-regular Max-Cut instances due to the parameter concentration property of the QAOA control parameters which have been seen in several problem classes [59–61]. The observable of interest in our ZNE simulations is the expected graph cut value given by the Hamiltonian,

$$H(z) = \sum_{(u,v) \in E} z_u z_v \quad (4)$$

where  $E$  is the edge set defining the interactions in the graph  $G = (V, E)$ , and  $z$  is a spin configuration represented as a vector of values  $z \in \{-1, +1\}^n$ . The minimum energy of this Hamiltonian corresponds to the optimal Max-Cut solution.

Applying quantum circuit unoptimization to QAOA circuits is a compelling test case for several reasons. First, the noise characteristics of QAOA circuits are well-studied, and the cost Hamiltonian’s expectation value is sensitive to noise, making it a suitable benchmark for evaluating ZNE. Second, the structure of QAOA circuits, which involves specific patterns of entangling gates and parameterized rotations, allows us to investigate how unoptimization interacts with structured quantum algorithms. Third, as QAOA circuits are actively explored for practical quantum applications, demonstrating the effectiveness of unoptimization as a noise-scaling technique in this context could be relevant to broader algorithmic developments.




Figure 2: Quantum circuit unoptimization scaling as a function of recursive iteration number (x-axis) applied to a 10 qubit quantum volume circuit. The top row shows single and two-qubit gate counts. The bottom row shows the overall circuit noise scale factor  $\lambda$ . The left-hand column uses a two-qubit gate selection based on a shared qubit index (concatenated strategy), and the right-hand column selects pairs of two-qubit gates at random (random strategy). The two-qubit gate selection is based on a shared qubit shows a more stable noise scale factor scaling.

## 2.3 Numerical Simulations

We present several numerical simulations to analyze the noise characteristics of using quantum circuit unoptimization for ZNE. All numerical simulations are performed using depolarizing noise (error) models using the Qiskit

quantum programming library [36] (version number 1.3.2). Specifically, a depolarizing error rate of 0.001 is applied to all single and two-qubit gates and a finite shot count of  $1e-6$  measurements are used for each circuit and noise scale factor execution. The quantum circuit unoptimization noise scale factor (eq. (3)) is computed after applying the Qiskit transpiler optimization pass to each unoptimization iteration. The unoptimization iteration is applied to the original circuit after applying the Qiskit transpiler optimization pass. `optimization_level 3` was used for all Qiskit optimization passes. All circuits are transpiled to the basis gates of `cx` (CNOT) and `U3` – reported single and two-qubit gate counts are based on only these two gates.

The extrapolation functions we test are linear and quadratic functions. The coefficients for these extrapolations are fit using least squares curve fitting. The quadratic function is  $ax^2 + bx + c$  where  $a, b$  and  $c$  are tunable coefficients. The linear function is  $ax + b$  where  $a$  and  $b$  are tunable coefficients.

### 3 Results




Figure 3: Zero-noise extrapolation using linear and quadratic regression for a Quantum Volume circuit with 10 qubits using quantum circuit unoptimization gate selection method random (right) and concatenated (left). There were 35 circuit unoptimization iterations used to amplify the noise and each step’s estimated observable is plotted as a blue point in each sub-figure. Each sub-plot shows the heavy output probability (y-axis) as a function of the noise scale factor  $\lambda$ . The red marker denotes the noiseless observable quantity, and the two yellow markers denote the quantum circuit unoptimization ZNE estimates of that observable. Note that each sub-plot used a single run of recursive quantum circuit unoptimization – meaning that the unoptimized circuit built on itself as the number of steps increased, as opposed to independently choosing a new unoptimization path (with a new random seed) for each noise level.

This section presents results using quantum circuit unoptimization for noise scaling. One of the crucial components of zero-noise extrapolation is the *noise scale factor*. Here, we define the noise scale factor  $\lambda$  as the ratio between the total number of circuit operations before quantum circuit unoptimization and after, see eq. (3). The total number of circuit operations is the sum of the single and two-qubit gates, not including qubit measurement or initialization operations. The intuition behind this choice of scaling factor is that it follows the circuit complexity of unitary circuit folding. However, this noise scale factor is only an *approximation*—the altered circuit structure can cause differences in error propagation.

The extrapolation functions we use in these simulations are only linear and quadratic regression. This is because the observable scaling as a function of  $\lambda$  does not drop off according to an exponential function. However, quantum-circuit unoptimization ZNE applied to other noise levels, quantum circuits or observables could require other function fits, such as Richardson or exponential extrapolation.

Figure 2 details gate count scaling and correspondingly  $\lambda$  scaling as a function of the number of quantum circuit unoptimization iterations. This shows that quantum circuit unoptimization gives a reasonably consistent, i.e., predictable, increase in circuit complexity. However, the concatenated strategy gives a *more* predictable scaling of circuit gate operation count, whereas the random strategy gives more variable gate operation counts.

Figure 3 shows quantum circuit unoptimization ZNE applied to a random quantum volume circuit. These results show that the noise can be amplified predictably and thus the noiseless signal can be recovered reasonably well.

Figure 4 demonstrates using quantum circuit unoptimization ZNE for a  $p = 2$  QAOA circuit applied to (un-weighted) Max-Cut on a random 3-regular graph, where the observable is the graph cut value. These simulations used fixed angles from ref. [58]. Notably, these results show a clear difference between the two different two-qubit gate selection methods; the random strategy, on average, results in more substantial noise amplification, but is also




Figure 4: Zero-noise extrapolation using linear and quadratic regression for a fixed-angle  $p = 2$  QAOA circuit applied to Max-Cut on a random 3-regular graph with 12 qubits using quantum circuit unoptimization gate selection method random (right) and concatenated (left). 35 quantum circuit unoptimization iterations were used to amplify the noise (the x-axis shows the computed  $\lambda$  quantity for each one of these iterations), and each step’s estimated observable is plotted in each sub-figure. Note that each sub-plot used a single run of recursive quantum circuit unoptimization – meaning that the unoptimized circuit built on itself as the number of steps increased, as opposed to independently choosing a new unoptimization path (with a new random seed) for each noise level.

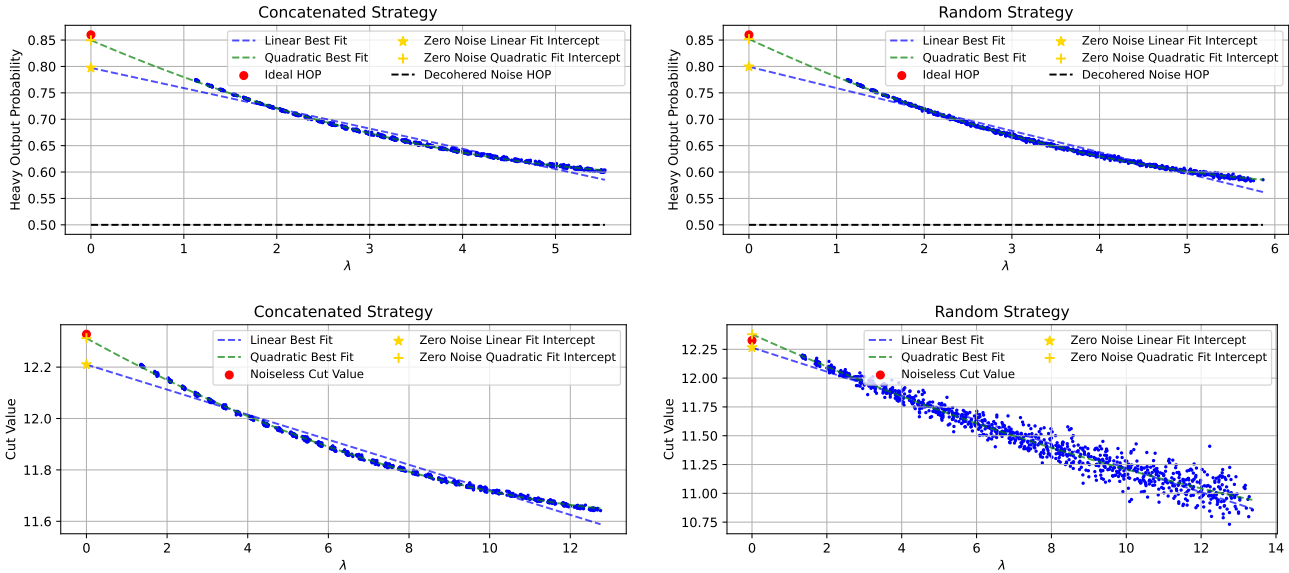


Figure 5: Averaged circuit variant zero-noise extrapolation numerical simulation tests. For each sub-plot, quantum circuit unoptimization is applied using a total of 30 independent recursive sequences (each from 1 to 35 iterations). This is unlike Figures 3 and 4, where only a single unoptimization sequence is used. The underlying quantum circuits (a fixed quantum volume circuit in the left column and a fixed Max-Cut QAOA circuit in the right-hand column), and observables, are the same as in Figures 3 and 4. Each separate quantum circuit unoptimization run can result in slightly different noise scale factors (eq. (3)). The random selection strategy applied to the  $p = 2$  QAOA circuit has the highest observable variability. Still, in all four instances, the zero-noise extrapolation can recover a good signal from the noisy quantum circuit simulation. Notice that the variance seen in the individual datapoints for the concatenated strategy plots is purely due to the random two qubit unitary selection during the insert stage of unoptimization.

more unpredictable, whereas the concatenated strategy results in a more controlled and predictable noise amplification. The optimal Max-Cut value for this specific problem instance is 16 – higher rounds of  $p$  would need to be used for QAOA to reach the optimal solution cut value.

Figure 6 reports distributions of heavy output probability (HOP) root mean squared error (RMSE) from an ensemble of 300, 10 qubit, random quantum volume circuits, when using quantum circuit unoptimization to perform ZNE to mitigate the errors from a depolarizing error model. Lower RMSE corresponds to better approximation of the noiseless heavy output probability observable. Figure 6 additionally shows the difference in error rate from






Figure 6: RMSE distributions between the ideal heavy output probability (HOP) and the ZNE extrapolated HOP using linear and quadratic extrapolation to mitigate the noise from a depolarizing error model when sampling a set of random 10-qubit quantum volume circuits. Each distribution contains HOP results from a set of fixed 300 randomly generated, and unique, quantum volume circuits. Each type of ZNE circuit unoptimization is applied to this set of random quantum volume circuits, and then, the two types of extrapolations are performed. Combined, these plots illustrate the differences in error rate between all combinations of the following 3 circuit unoptimization ZNE parameters: linear vs. quadratic extrapolation, concatenated vs random strategy, and averaged over many circuit unoptimization passes vs. a single circuit unoptimization pass. The mean of each distribution is marked by the vertical dashed line corresponding to the color of that distribution. The overall lowest error rate set of parameters is quadratic extrapolation with the random two-qubit gate selection strategy and averaged over 30 distinct recursions of circuit unoptimization (right plot, cyan distribution). Overall findings are as follows. Random strategy gives a lower error rate than concatenated strategy. Quadratic extrapolation performs better than linear extrapolation. Averaging the extrapolation over 30 independent recursions of circuit unoptimization results in lower error rates than using only a single recursion run to extrapolate from.

using a single circuit unoptimization recursion run compared to extrapolating based on many (in this case, 30) independent recursions. Here, independent recursion means explicitly that for each run we have chosen a different random seed for the pseudorandom number generator in the code to choose a slightly different path of two-qubit gates choices (only in the case of the random strategy), as well as different random two-qubit unitaries  $A$  (described in Section 2.1). The result is that performing noise averaging using the multiple circuit unoptimization recursions results in lower error rates. This intuitively makes sense because the circuit unoptimization procedure can cause many different circuit structures with varying error propagation, and therefore averaging over many different circuit structure cases will result in less biased observables. These results in Figure 6 also show that the quadratic extrapolation can approximate the correct heavy output probability observable better than linear extrapolation. The concatenated gate selection strategy has a higher error rate than the random gate selection strategy, which makes sense because the concatenated gate selection strategy does not use any random choice when choosing the pair of 2-qubit gates, which means that the resulting observable extrapolation has less noise variance to average over.

Quantum circuit unoptimization with the random strategy allows some freedom in random selection of the pair of two-qubit gates, one can generate many circuit variants (meaning they implement equivalent unitaries but contain a distinct set of quantum instructions). Additionally, for both the concatenated and random strategy we can choose random two-qubit unitaries during the insert stage of unoptimization, which additionally results in many potential circuit variants. Figure 5 explores this by repeating quantum circuit unoptimization on the same fixed test cases as Figure 3 and 4, but now running 30 complete recursions of the unoptimization (where each subsequent circuit is constructed based on the gates created from the previous unoptimization run). Each run used a new random seed so that the path of random gate selection and two-qubit unitary generation could be different. We observe that some strategies and circuit types can yield highly variable noise amplification, in this case, the random gate selection strategy applied to the QAOA circuit instance, whereas the other three cases resulted in more consistent noise amplification. Figures 5 and 6 shows overall that the circuit averaging does result in good observable recovery – although notably the non-averaged cases seen in Figure 3 and 4 performed similarly. This shows that the average of extrapolated observables from many circuit variants can give robust ZNE computation,



but also that even single iterations of circuit unoptimization still perform well.

Combined, these results show that the concatenated strategy is better for performing controlled (predictable) noise amplification using quantum circuit unoptimization, but also results in higher error rates than the random gate selection option as seen in Figure 6, due to the lack of random 2-qubit gate selection that can result in better noise averaging. Therefore, the best strategy that results in the overall lowest error rate is to use the random gate selection strategy, and to average over many independent recursions of the quantum circuit unoptimization, and then to use the extrapolation that gives the best curve fit to the observable that is being measured.

## 4 Discussion

We explored using the quantum circuit unoptimization routines introduced in [35] as a noise-scaling method for the zero-noise extrapolation (ZNE) quantum error mitigation technique. To validate our approach, we applied this noise-scaling method to random quantum volume and QAOA circuits, supported by numerical simulations, to demonstrate its effectiveness. These results highlight the potential of quantum circuit unoptimization to introduce controlled noise amplification while preserving the underlying logical functionality of the circuits.

It is worthwhile to consider the trade-offs of quantum circuit unoptimization as a noise-scaling technique compared to other known techniques, such as unitary folding. One of the primary benefits of quantum circuit unoptimization is that, as more iterations of the unoptimization are performed, it becomes progressively more difficult for the original circuit to be recovered via compilers running on real quantum device backends. This property is desirable because it ensures that the noise amplification introduced by the unoptimization process remains intact during execution, preventing compiler optimizations from inadvertently simplifying the circuit. In contrast, unitary folded circuits can sometimes be undone by the compiler, particularly when folding operations are explicitly designed to be reversible and are thus more susceptible to simplification during standard quantum circuit compilation. By comparison, the structural complexity introduced by unoptimization makes it more resistant to such optimizations. Note however that the circuit unoptimization routine could likely be identified and reversed by a compiler that was specifically designed to identify these circuit structures – it appears that current quantum circuit compilers are not able to identify and reverse these structures [35].

While this increased circuit complexity provides advantages in preserving noise scaling, it also introduces potential challenges. Most current quantum computers provide options to disable compiler optimizations, mitigating the risk of unwanted circuit simplifications. However, even when optimizations are turned off, certain low-level adjustments made by the compiler may still partially reverse aspects of the unoptimization process, potentially leading to less predictable noise amplification factors (for example, if the backend removes identity gates that introduce errors due to idling time). This unpredictability arises because unoptimization alters the circuit structure in more challenging ways to model or control than the more straightforward unitary folding approach. The original circuit structure is preserved in unitary folding, and the noise amplification factor is more precisely determined, making it a more predictable method in many scenarios.

Using quantum circuit unoptimization means a slight change in how ZNE is applied. Typically, when using ZNE, the user chooses a noise amplification factor, and then the algorithm attempts to produce a circuit representation with that level of noise amplification. When using quantum circuit unoptimization, instead, the noise scale factor (eq. (3)) is computed based on the circuits generated by the unoptimization, and that value is then used for the extrapolation. Figure 2 shows that the circuit complexity scaling is predictable, which means that a noise amplification factor in principle could be a priori requested. The closest integer required number of circuit unoptimization iterations could then be selected to produce that noise amplification amount. However, this study used a more precise measurement of the noise scale factor post-circuit unoptimization (including post Qiskit transpiler optimization).

Another consideration is the relative complexity of implementation. Unitary folding is conceptually simple and requires minimal additional overhead to apply. By contrast, quantum circuit unoptimization involves multiple steps, including gate insertion, swapping, decomposition, and synthesis, all of which must be implemented carefully to ensure correctness and compatibility with the input circuit. To partially offset this complexity, we provide a fully self-contained implementation of quantum circuit unoptimization in Python, compatible with IBM Qiskit circuits. This implementation can be adapted to other quantum circuit libraries, ensuring accessibility for researchers and practitioners interested in exploring unoptimization as a noise-scaling technique.

An advantage of quantum circuit unoptimization compared to existing noise amplification methods is that an exponentially growing number of potential, distinct, circuit variants can be generated for a particular noise amount. Most other existing methods allow only one or a limited number of circuit variants to be generated.

Generating many circuit variants may facilitate better noise averaging across many of these circuit variants on quantum hardware with significantly biased noise.

The primary limitation of quantum circuit unoptimization for noise amplification in ZNE is the highly variable changes to circuit structure. These circuit changes can result in very different error propagation compared to the unmodified circuit. However, the previously mentioned property that *many* random (equivalent) circuit variants can be generated with quantum circuit unoptimization to mitigate this issue via noise averaging across many instances with different circuit structures. Therefore, in practice on noisy quantum computers, we recommend averaging over many random circuits for each noise scale factor. A possible future subject of study is developing other circuit unoptimization algorithms, in addition to [35], which creates a quantum circuit structure that minimizes highly biased error propagation (e.g., a circuit structure that approximately mimics the original circuit structure).

Finally, it is essential to note that while quantum circuit unoptimization offers certain advantages for specific use cases, there is no theoretical reason to expect it to outperform more conventional approaches to perform consistently noise scaling, such as unitary folding or pulse stretching. Each method has unique strengths and limitations, and the choice of technique should depend on the specific application, the noise characteristics of the hardware, and the desired level of control over circuit structure and complexity.

## 5 Software

The software implementing the quantum circuit unoptimization recipe, its application as a noise-scaling technique for ZNE, as well as all data and plots used in this work are available in a public GitHub repository

<https://github.com/unitaryfund/circuit-unoptimization>.

## 6 Acknowledgments

The authors thank Yusei Mori for assisting in implementing the algorithm of quantum circuit unoptimization and Andrea Mari, Farrokh Labib, Nate T. Stemen, and Nathan Shammah for the helpful discussions on this approach. E.P. was supported by the U.S. Department of Energy through the Los Alamos National Laboratory the NNSA’s Advanced Simulation and Computing Beyond Moore’s Law Program at Los Alamos National Laboratory. Los Alamos National Laboratory is operated by Triad National Security, LLC, for the National Nuclear Security Administration of U.S. Department of Energy (Contract No. 89233218CNA000001). This research used resources provided by the Darwin testbed at Los Alamos National Laboratory (LANL), funded by the Computational Systems and Software Environments subprogram of LANL’s Advanced Simulation and Computing program (NNSA/DOE). LANL report number LA-UR-25-21046.

## A Appendix

In Figure 1, the circuit diagrams in the top panel may be unfamiliar to readers who are not users of the Qiskit package. This section provides an overview of the gates appearing in the diagrams and their corresponding matrix representations. The standard set of single-qubit gates used includes the Hadamard ( $H$ ), Pauli- $X$ , and Pauli- $Z$  gates, defined as

$$H = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}, \quad X = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \quad Z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}.$$

The two-qubit gates used are the controlled-NOT ( $CX$ ) and controlled-Z ( $CZ$ ) gates, defined as

$$CX = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix}, \quad CZ = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & -1 \end{pmatrix}.$$

The  $U$  and  $U_3$  gates are equivalent in these circuit diagrams due to the legacy naming convention from earlier versions of Qiskit and represents the most general single-qubit rotation. They are parameterized by three angles  $\theta$ ,  $\phi$ , and  $\lambda$ , and their matrix representation is given by

$$U(\theta, \phi, \lambda) = U_3(\theta, \phi, \lambda) = \begin{pmatrix} \cos \frac{\theta}{2} & -e^{i\lambda} \sin \frac{\theta}{2} \\ e^{i\phi} \sin \frac{\theta}{2} & e^{i(\phi+\lambda)} \cos \frac{\theta}{2} \end{pmatrix}.$$

The parameters  $\theta$ ,  $\phi$ , and  $\lambda$  determine rotations about the y-axis and z-axis of the Bloch sphere, making  $U_3$  universal for single-qubit operations.

The  $U_2$  gate is a special case of  $U_3$  where the rotation angle is fixed at  $\theta = \frac{\pi}{2}$ . This simplifies the matrix to

$$U_2(\phi, \lambda) = U_3\left(\frac{\pi}{2}, \phi, \lambda\right) = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & -e^{i\lambda} \\ e^{i\phi} & e^{i(\phi+\lambda)} \end{pmatrix}.$$

The  $U_2$  gate can be considered as a simplified version of  $U_3$  that performs a Hadamard-like operation with additional phase shifts along the z-axis. Its reduced parametrization makes it computationally efficient while retaining sufficient flexibility for many quantum algorithms.

In addition to standard gates, some diagrams feature gates labeled ‘circuit-X’ (e.g., ‘circuit-826’, ‘circuit-829’). These labels are automatically generated by Qiskit when custom or composite gates are included in the circuit. A custom gate represents a more complex operation that cannot be directly reduced to standard gates like  $U_3$  or  $CX$ . These custom gates typically correspond to either composite sub-circuits representing a sequence of gates grouped and treated as a single operation or as higher-dimensional unitary matrices applied to multiple qubits.

## References

- [1] Zhenyu Cai et al. “Quantum error mitigation”. In: *Reviews of Modern Physics* 95.4 (Dec. 2023). ISSN: 1539-0756. DOI: 10.1103/revmodphys.95.045005. URL: <http://dx.doi.org/10.1103/RevModPhys.95.045005>.
- [2] Ying Li and Simon C Benjamin. “Efficient variational quantum simulator incorporating active error minimization”. In: *Physical Review X* 7.2 (2017), p. 021050. DOI: 10.1103/PhysRevX.7.021050. URL: <https://journals.aps.org/prx/abstract/10.1103/PhysRevX.7.021050>.
- [3] Kristan Temme, Sergey Bravyi, and Jay M Gambetta. “Error mitigation for short-depth quantum circuits”. In: *Physical Review Letters* 119.18 (2017), p. 180509. DOI: 10.1103/PhysRevLett.119.180509. URL: <https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.119.180509>.
- [4] Suguru Endo, Simon C Benjamin, and Ying Li. “Practical quantum error mitigation for near-future applications”. In: *Physical Review X* 8.3 (2018), p. 031027. DOI: 10.1103/PhysRevX.8.031027. URL: <https://journals.aps.org/prx/abstract/10.1103/PhysRevX.8.031027>.
- [5] Abhinav Kandala et al. “Error mitigation extends the computational reach of a noisy quantum processor”. In: *Nature* 567.7749 (2019), pp. 491–495. DOI: 10.1038/s41586-019-1040-7. URL: <https://www.nature.com/articles/s41586-019-1040-7>.
- [6] Armands Strikis et al. “Learning-based quantum error mitigation”. In: *PRX Quantum* 2.4 (2021), p. 040330. DOI: 10.1103/PRXQuantum.2.040330. URL: <https://journals.aps.org/prxquantum/abstract/10.1103/PRXQuantum.2.040330>.
- [7] Tudor Giurgica-Tiron et al. “Digital zero noise extrapolation for quantum error mitigation”. In: *2020 IEEE International Conference on Quantum Computing and Engineering (QCE)*. 2020, pp. 306–316. DOI: 10.1109/QCE49297.2020.00045. URL: <https://arxiv.org/abs/2005.10921>.
- [8] Ryan LaRose et al. “Mitiq: A software package for error mitigation on noisy quantum computers”. In: *Quantum* 6 (2022), p. 774. DOI: 10.22331/q-2022-08-11-774. URL: <https://quantum-journal.org/papers/q-2022-08-11-774/>.
- [9] Suguru Endo et al. “Hybrid quantum-classical algorithms and quantum error mitigation”. In: *Journal of the Physical Society of Japan* 90.3 (2021), p. 032001. DOI: 10.7566/JPSJ.90.032001. URL: <https://journals.jps.jp/doi/10.7566/JPSJ.90.032001>.
- [10] Youngseok Kim et al. “Scalable error mitigation for noisy quantum circuits produces competitive expectation values”. In: *Nature Physics* (2023), pp. 1–8. DOI: 10.1038/s41567-022-01914-3. URL: <https://www.nature.com/articles/s41567-022-01914-3>.
- [11] Bálint Koczor. “Exponential error suppression for near-term quantum devices”. In: *Physical Review X* 11.3 (2021), p. 031057. DOI: 10.1103/PhysRevX.11.031057. URL: <https://journals.aps.org/prx/abstract/10.1103/PhysRevX.11.031057>.
- [12] Andre He et al. “Zero-noise extrapolation for quantum-gate error mitigation with identity insertions”. In: *Physical Review A* 102 (1 2020), p. 012426. DOI: 10.1103/PhysRevA.102.012426. URL: <https://link.aps.org/doi/10.1103/PhysRevA.102.012426>.
- [13] William J Huggins et al. “Virtual distillation for quantum error mitigation”. In: *Physical Review X* 11.4 (2021), p. 041036. DOI: 10.1103/PhysRevX.11.041036. URL: <https://journals.aps.org/prx/abstract/10.1103/PhysRevX.11.041036>.
- [14] Vincent R. Pascuzzi et al. “Computationally efficient zero-noise extrapolation for quantum-gate-error mitigation”. In: *Physical Review A* 105.4 (2022), p. 042406. DOI: 10.1103/PhysRevA.105.042406. URL: <https://journals.aps.org/pra/abstract/10.1103/PhysRevA.105.042406>.
- [15] Chao Song et al. “Quantum computation with universal error mitigation on a superconducting quantum processor”. In: *Science Advances* 5.9 (2019), eaaw5686. DOI: 10.1126/sciadv.aaw5686. URL: <https://www.science.org/doi/10.1126/sciadv.aaw5686>.
- [16] Lorenza Viola, Emanuel Knill, and Seth Lloyd. “Dynamical decoupling of open quantum systems”. In: *Physical Review Letters* 82 (12 1999), pp. 2417–2421. DOI: 10.1103/PhysRevLett.82.2417. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.82.2417>.
- [17] Alexandre M Souza, Gonzalo A Álvarez, and Dieter Suter. “Robust dynamical decoupling”. In: *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 370.1976 (2012), pp. 4748–4769. DOI: 10.1098/rsta.2011.0355. URL: <https://royalsocietypublishing.org/doi/10.1098/rsta.2011.0355>.
- [18] P. Facchi, D. A. Lidar, and S. Pascazio. “Unification of dynamical decoupling and the quantum Zeno effect”. In: *Physical Review A* 69 (3 2004), p. 032314. DOI: 10.1103/PhysRevA.69.032314. URL: <https://link.aps.org/doi/10.1103/PhysRevA.69.032314>.

- [19] Lea F Santos and Lorenza Viola. “Dynamical control of qubit coherence: Random versus deterministic schemes”. In: *Physical Review A* 72.6 (2005), p. 062303. DOI: 10.1103/PhysRevA.72.062303. URL: <https://journals.aps.org/prabstract/10.1103/PhysRevA.72.062303>.
- [20] Lorenza Viola and Emanuel Knill. “Random decoupling schemes for quantum dynamical control and error suppression”. In: *Physical Review Letters* 94.6 (2005), p. 060502. DOI: 10.1103/PhysRevLett.94.060502. URL: <https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.94.060502>.
- [21] Bibek Pokharel et al. “Demonstration of fidelity improvement using dynamical decoupling with superconducting qubits”. In: *Physical Review Letters* 121.22 (2018), p. 220502. DOI: 10.1103/PhysRevLett.121.220502. URL: <https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.121.220502>.
- [22] Pavel Sekatski, Michalis Skotiniotis, and Wolfgang Dür. “Dynamical decoupling leads to improved scaling in noisy quantum metrology”. In: *New Journal of Physics* 18.7 (2016), p. 073034. DOI: 10.1088/1367-2630/18/7/073034. URL: <https://iopscience.iop.org/article/10.1088/1367-2630/18/7/073034>.
- [23] Samantha V Barron et al. “Provable bounds for noise-free expectation values computed from noisy samples”. In: *Nature Computational Science* (2024), pp. 1–11. DOI: 10.1038/s43588-024-00709-1. URL: <https://www.nature.com/articles/s43588-024-00709-1>.
- [24] Ryuji Takagi et al. “Fundamental limits of quantum error mitigation”. In: *npj Quantum Information* 8.1 (2022), p. 114. DOI: 10.1038/s41534-022-00618-z. URL: <https://www.nature.com/articles/s41534-022-00618-z>.
- [25] Yihui Quek et al. “Exponentially tighter bounds on limitations of quantum error mitigation”. In: *Nature Physics* 20.10 (2024), pp. 1648–1658. DOI: 10.1038/s41567-024-02536-7. URL: <https://www.nature.com/articles/s41567-024-02536-7>.
- [26] Andrea Mari, Nathan Shammah, and William J Zeng. “Extending quantum probabilistic error cancellation by noise scaling”. In: *Physical Review A* 104.5 (2021), p. 052607. DOI: 10.1103/PhysRevA.104.052607. URL: <https://link.aps.org/accepted/10.1103/PhysRevA.104.052607>.
- [27] Ewout Van Den Berg et al. “Probabilistic error cancellation with sparse Pauli–Lindblad models on noisy quantum processors”. In: *Nature Physics* 19.8 (2023), pp. 1116–1121. DOI: 10.1038/s41567-023-02042-2. URL: <https://www.nature.com/articles/s41567-023-02042-2>.
- [28] Yue Ma and M. S. Kim. “Limitations of probabilistic error cancellation for open dynamics beyond sampling overhead”. In: *Physical Review A* 109 (1 2024), p. 012431. DOI: 10.1103/PhysRevA.109.012431. URL: <https://link.aps.org/doi/10.1103/PhysRevA.109.012431>.
- [29] Piotr Czarnik et al. “Error mitigation with Clifford quantum-circuit data”. In: *Quantum* 5 (Nov. 2021), p. 592. ISSN: 2521-327X. DOI: 10.22331/q-2021-11-26-592. URL: <http://dx.doi.org/10.22331/q-2021-11-26-592>.
- [30] Stefan H. Sack and Daniel J. Egger. “Large-scale quantum approximate optimization on nonplanar graphs with machine learning noise mitigation”. In: *Physical Review Research* 6.1 (Mar. 2024). ISSN: 2643-1564. DOI: 10.1103/physrevresearch.6.013223. URL: <http://dx.doi.org/10.1103/PhysRevResearch.6.013223>.
- [31] Haoran Liao et al. “Machine learning for practical quantum error mitigation”. In: *Nature Machine Intelligence* (2024), pp. 1–9. DOI: 10.1038/s42256-024-00927-2. URL: <https://www.nature.com/articles/s42256-024-00927-2>.
- [32] Changjun Kim, Kyungdeock Daniel Park, and June-Koo Rhee. “Quantum error mitigation with artificial neural network”. In: *IEEE Access* 8 (2020), pp. 188853–188860. DOI: 10.1109/ACCESS.2020.3031607. URL: <https://ieeexplore.ieee.org/document/9226505>.
- [33] Elizabeth R Bennewitz et al. “Neural error mitigation of near-term quantum simulations”. In: *Nature Machine Intelligence* 4.7 (2022), pp. 618–624. DOI: 10.1038/s42256-022-00509-0. URL: <https://www.nature.com/articles/s42256-022-00509-0>.
- [34] Youngseok Kim et al. “Evidence for the utility of quantum computing before fault tolerance”. In: *Nature* 618.7965 (2023), pp. 500–505. DOI: 10.1038/s41586-023-06096-3. URL: <https://www.nature.com/articles/s41586-023-06096-3>.
- [35] Yusei Mori et al. “Quantum circuit unoptimization”. In: (2024). DOI: 10.48550/arXiv.2311.03805. arXiv: 2311.03805 [quant-ph]. URL: <https://arxiv.org/abs/2311.03805>.
- [36] Ali Javadi-Abhari et al. “Quantum computing with Qiskit”. In: (2024). DOI: 10.48550/arXiv.2405.08810. arXiv: 2405.08810 [quant-ph]. URL: <https://arxiv.org/abs/2405.08810>.
- [37] Andrew W. Cross et al. “Validating quantum computers using randomized model circuits”. In: *Physical Review A* 100 (3 2019), p. 032328. DOI: 10.1103/PhysRevA.100.032328. URL: <https://link.aps.org/doi/10.1103/PhysRevA.100.032328>.

- [38] Charles H. Baldwin et al. “Re-examining the quantum volume test: Ideal distributions, compiler optimizations, confidence intervals, and scalable resource estimations”. In: *Quantum* 6 (May 2022), p. 707. ISSN: 2521-327X. DOI: 10.22331/q-2022-05-09-707. URL: <https://doi.org/10.22331/q-2022-05-09-707>.
- [39] Elijah Pelofske, Andreas Bärttschi, and Stephan Eidenbenz. “Quantum Volume in Practice: What Users Can Expect From NISQ Devices”. In: *IEEE Transactions on Quantum Engineering* 3 (2022), 1–19. ISSN: 2689-1808. DOI: 10.1109/tqe.2022.3184764. URL: <http://dx.doi.org/10.1109/TQE.2022.3184764>.
- [40] Petar Jurcevic et al. “Demonstration of quantum volume 64 on a superconducting quantum computing system”. In: *Quantum Science and Technology* 6.2 (2021), p. 025020. DOI: 10.1088/2058-9565/abe519. URL: <https://iopscience.iop.org/article/10.1088/2058-9565/abe519>.
- [41] Elijah Pelofske et al. “Increasing the Measured Effective Quantum Volume with Zero Noise Extrapolation”. In: *ACM Transactions on Quantum Computing* 5.3 (Sept. 2024). DOI: 10.1145/3680290. URL: <https://doi.org/10.1145/3680290>.
- [42] Ryan LaRose et al. “Error mitigation increases the effective quantum volume of quantum computers”. In: (2022). DOI: 10.48550/arXiv.2203.05489. arXiv: 2203.05489 [quant-ph]. URL: <https://arxiv.org/abs/2203.05489>.
- [43] Scott Aaronson and Lijie Chen. “Complexity-Theoretic Foundations of Quantum Supremacy Experiments”. In: *Proceedings of the 32nd Computational Complexity Conference. CCC '17*. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2017. ISBN: 9783959770408. DOI: 10.4230/LIPIcs.CCC.2017.22. URL: <https://dl.acm.org/doi/10.5555/3135595.3135617>.
- [44] Edward Farhi, Jeffrey Goldstone, and Sam Gutmann. “A quantum approximate optimization algorithm”. In: *arXiv preprint arXiv:1411.4028* (2014). DOI: 10.48550/arxiv.1411.4028. URL: <https://arxiv.org/abs/1411.4028>.
- [45] Edward Farhi, Jeffrey Goldstone, and Sam Gutmann. “A quantum approximate optimization algorithm applied to a bounded occurrence constraint problem”. In: *arXiv preprint arXiv:1412.6062* (2015). DOI: 10.48550/arXiv.1412.6062. URL: <https://arxiv.org/abs/1412.6062>.
- [46] Matthew P Harrigan et al. “Quantum approximate optimization of non-planar graph problems on a planar superconducting processor”. In: *Nature Physics* 17.3 (2021), pp. 332–336. DOI: 10.1038/s41567-020-01105-y. URL: <https://www.nature.com/articles/s41567-020-01105-y>.
- [47] Leo Zhou et al. “Quantum approximate optimization algorithm: Performance, mechanism, and implementation on near-term devices”. In: *Physical Review X* 10.2 (2020), p. 021067. DOI: 10.1103/PhysRevX.10.021067. URL: <https://journals.aps.org/prx/abstract/10.1103/PhysRevX.10.021067>.
- [48] Joao Basso et al. “The quantum approximate optimization algorithm at high depth for MaxCut on large-girth regular graphs and the Sherrington-Kirkpatrick model”. In: *arXiv preprint arXiv:2110.14206* (2021). DOI: 10.4230/LIPICS.TQC.2022.7. URL: <https://quantum-journal.org/papers/q-2022-07-07-759/>.
- [49] Zhihui Wang et al. “Quantum approximate optimization algorithm for MaxCut: A Fermionic view”. In: *Physical Review A* 97.2 (2018), p. 022304. DOI: 10.1103/PhysRevA.97.022304. URL: <https://journals.aps.org/prx/abstract/10.1103/PhysRevA.97.022304>.
- [50] Gavin E Crooks. “Performance of the quantum approximate optimization algorithm on the maximum cut problem”. In: *arXiv preprint arXiv:1811.08419* (2018). DOI: 10.48550/arXiv.1811.08419. URL: <https://ui.adsabs.harvard.edu/abs/2019APS..MARK27002C/abstract>.
- [51] Gian Giacomo Guerreschi and Anne Y Matsuura. “QAOA for Max-Cut requires hundreds of qubits for quantum speed-up”. In: *Scientific Reports* 9.1 (2019), p. 6903. DOI: 10.1038/s41598-019-43176-9. URL: <https://www.nature.com/articles/s41598-019-43176-9>.
- [52] Kunal Marwaha. “Local classical Max-Cut algorithm outperforms  $p = 2$  QAOA on high-girth regular graphs”. In: *Quantum* 5 (2021), p. 437. DOI: 10.22331/q-2021-04-20-437. URL: <https://quantum-journal.org/papers/q-2021-04-20-437/>.
- [53] Matthew B Hastings. “Classical and quantum bounded depth approximation algorithms”. In: *arXiv preprint arXiv:1905.07047* (2019). DOI: 10.26421/QIC19.13-14-3. URL: <https://arxiv.org/abs/1905.07047>.
- [54] Michel X. Goemans and David P. Williamson. “Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming”. In: *J. ACM* 42.6 (Nov. 1995), 1115–1145. ISSN: 0004-5411. DOI: 10.1145/227683.227684. URL: <https://doi.org/10.1145/227683.227684>.
- [55] Luca Trevisan. “Max cut and the smallest eigenvalue”. In: *Proceedings of the forty-first annual ACM symposium on Theory of computing*. 2009, pp. 263–272. DOI: 10.1137/090773714. URL: <https://epubs.siam.org/doi/10.1137/090773714>.

- [56] Paola Festa et al. “Randomized heuristics for the Max-Cut problem”. In: *Optimization methods and software* 17.6 (2002), pp. 1033–1058. DOI: 10.1080/1055678021000090033. URL: <https://www.tandfonline.com/doi/abs/10.1080/1055678021000090033>.
- [57] Stefan H. Sack and Maksym Serbyn. “Quantum annealing initialization of the quantum approximate optimization algorithm”. In: *Quantum* 5 (July 2021), p. 491. ISSN: 2521-327X. DOI: 10.22331/q-2021-07-01-491. URL: <http://dx.doi.org/10.22331/q-2021-07-01-491>.
- [58] Jonathan Wurtz and Danylo Lykov. “Fixed-angle conjectures for the quantum approximate optimization algorithm on regular MaxCut graphs”. In: *Phys. Rev. A* 104 (5 Nov. 2021), p. 052419. DOI: 10.1103/PhysRevA.104.052419. URL: <https://link.aps.org/doi/10.1103/PhysRevA.104.052419>.
- [59] V. Akshay et al. “Parameter concentrations in quantum approximate optimization”. In: *Phys. Rev. A* 104 (1 July 2021), p. L010401. DOI: 10.1103/PhysRevA.104.L010401. URL: <https://link.aps.org/doi/10.1103/PhysRevA.104.L010401>.
- [60] Fernando G. S. L. Brandao et al. “For Fixed Control Parameters the Quantum Approximate Optimization Algorithm’s Objective Function Value Concentrates for Typical Instances”. In: (2018). DOI: 10.48550/arXiv.1812.04170. arXiv: 1812.04170 [quant-ph]. URL: <https://arxiv.org/abs/1812.04170>.
- [61] Joao Basso et al. “The Quantum Approximate Optimization Algorithm at High Depth for MaxCut on Large-Girth Regular Graphs and the Sherrington-Kirkpatrick Model”. en. In: Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2022. DOI: 10.4230/LIPICS.TQC.2022.7. URL: <https://drops.dagstuhl.de/entities/document/10.4230/LIPICS.TQC.2022.7>.