

Improving Quantum Approximate Optimization by Noise-Directed Adaptive Remapping

Filip B. Maciejewski^{1,2}, Jacob Biamonte^{1,2}, Stuart Hadfield^{1,2}, and Davide Venturelli^{1,2}

¹Quantum Artificial Intelligence Laboratory (QuAIL), NASA Ames Research Center, CA, USA

²USRA Research Institute for Advanced Computer Science (RIACS), CA, USA

2025.11.01

We present Noise-Directed Adaptive Remapping (NDAR), a heuristic algorithm for approximately solving binary optimization problems by leveraging certain types of noise. We consider access to a noisy quantum processor with dynamics that features a global attractor state. In a standard setting, such noise can be detrimental to the quantum optimization performance. Our algorithm bootstraps the noise attractor state by iteratively gauge-transforming the cost-function Hamiltonian in a way that transforms the noise attractor into higher-quality solutions. The transformation effectively changes the attractor into a higher-quality solution of the Hamiltonian based on the results of the previous step. The end result is that noise aids variational optimization, as opposed to hindering it. We present an improved Quantum Approximate Optimization Algorithm (QAOA) runs in experiments on Rigetti’s quantum device. We report approximation ratios 0.9–0.96 for random, fully connected graphs on $n = 82$ qubits, using only depth $p = 1$ QAOA with NDAR. This compares to 0.34–0.51 for standard $p = 1$ QAOA with the same number of function calls.

1 Introduction

Noise in state-of-the-art quantum processors limits quantum circuit depth. Significant effort has been spent on improving the implementation of quantum approximate optimization [1, 2, 3, 4,

5, 6, 7, 8, 9, 10, 11, 12, 13, 14], and developing various error mitigation [15] and error correction [16, 17] techniques. In this work, we propose an algorithmic framework that does not mitigate the noise directly, but instead attempts to align the quantum optimization with the dissipative part of the evolution by using the knowledge about the noise.

To this aim, we exploit the concept of bitflip transforms: problem remappings that exchange the logical definitions of a set of binary variables, or equivalently that exchange the interpretation of $|0\rangle$ and $|1\rangle$ basis states for a set of physical qubits. Assuming that the device noise dynamics is biasing outputs towards a specific classical “attractor” state (denoted as $|0\dots 0\rangle$), we propose an algorithm that iteratively gauge-transforms the problem in a way that consistently assigns better cost-function values to the noise attractor. The remappings are adaptively selected based on the best solution obtained through variational optimization in the previous step (see Fig. 1 for illustration).

We refer to the new algorithm as Noise-Directed Adaptive Remapping (NDAR). The actual variational optimization method is a subroutine of NDAR that can be chosen in a way best suited for a given experimental setup (in fact, it does not necessarily have to be quantum). In this work, we focus mainly on the canonical example of Quantum Approximate Optimization Algorithm (QAOA) [1, 2] though the adaptive remapping approach we introduce applies to other approaches. We implement the NDAR adaptive loop for $p = 1$ QAOA targeting the minimization of fully-connected, randomly-weighted Sherrington-Kirkpatrick Hamiltonians on $n = 82$ -qubit subsystem of Rigetti Computing’s quantum processing unit (QPU) based on

Filip B. Maciejewski: fmaciejewski@usra.edu

Davide Venturelli: davide.venturelli@nasa.gov




Fig. 1: Illustration of a Noise-Directed Adaptive Remapping run for a single Hamiltonian instance. The X-axis represents iterations of the algorithm. The purple dots correspond to the approximation ratio (AR) of the $|0 \dots 0\rangle$ (attractor) state, while cyan dots are the best-observed solution (as a result of QAOA runs). In consecutive steps, the Hamiltonian is re-transformed in a way that the best solution from the previous step effectively becomes the (approximate) attractor of the new step, resulting in a stairs-like climb towards higher approximation ratios until convergence. The distributions of ARs are stacked horizontally for each NDAR step, showing how the tail of the distribution from a previous step, becomes a point near the center of the distribution in the next step (in early iterations). The presented data corresponds to a single instance from the experimentally obtained dataset (for $n = 82$ qubits) in Fig. 2b. The inset illustrates how binary (spin) variables are associated with the qubit states of a physical QPU system. Qubits are represented by anharmonic oscillators (e.g. transmons) with ground state $|0\rangle$ and first excited state $|1\rangle$. Gates drive coherent transitions (blue arrows) while noise (red arrow) induces asymmetric incoherent transitions.

superconducting transmons. We observe solutions with approximation ratios above ≈ 0.9 , compared to ≈ 0.5 of standard QAOA, when both methods are given the same number of circuit runs.

Related Work The use of bitflip gauges (a.k.a. spin-reversal or spin-inversion transforms) has been considered for quantum annealing in the context of embedding problem instances to “unbias” individual qubits and generally to reduce performance errors [18, 19, 20, 21, 22, 23, 24, 25, 26]. In practice, bitflip transforms are often chosen at random, see, e.g., [18, 23, 26]. In Ref. [23], the authors developed a heuristic algorithm that searches for the best gauge, and presented experimental results on a quantum annealer, treating the bitflip gauge as an additional parameter to be tuned. Our approach is con-

ceptually different – the NDAR choice of bitflip transforms is informed by the noise model, and it is iteratively updated in the course of implementation. For quantum gate model applications, some studies investigate and potentially exploit various problem symmetries (such as qubit ordering); see, for example, Refs. [27, 28, 29, 30, 31, 32, 8, 33, 34]. Effects of various types of noise on QAOA performance are studied, for example, in Refs. [35, 36, 37, 38]. The idea of adaptively changing the ansatz is found in some recent works, such as the ADAPT-QAOA [6] protocol which adaptively constructs the mixer Hamiltonian. In the case of NDAR applied to QAOA, applying a bitflip transform is equivalent to adaptively changing the physically implemented phase separation (cost) Hamiltonians, and adding a simple post-processing procedure (re-labeling). Generally speaking, we note that Noise-Directed

Adaptive Remapping might fit into a more general framework of noise-assisted quantum algorithms that use open systems dynamics as a resource [39, 40, 41, 42, 43, 44, 45, 46, 47, 48, 49]. Finally, we note that the NDAR algorithm is conceptually related but distinct from so-called “warm start” strategies for quantum approximate optimization [50, 51, 52, 53], where the ansatz or initial state are modified based on approximate solutions obtained classically. In the case of NDAR, the ansatz is adaptively modified based on the approximate solutions obtained in *quantum* optimization (and based on our knowledge of noise dynamics). Finally, we emphasize that NDAR is applicable to a wide variety of quantum and classical algorithms, especially those categorized as Ising machines [54].

2 Problem Remappings

A large variety of industrially relevant optimization tasks can be formulated as optimization problems over binary variables [55, 2]. For example, numerous important problems [56, 57] can be directly translated or reduced to quadratic unconstrained optimization problems over Boolean variables $x_i = 0, 1$, or, equivalently, instances of the Ising model: $\min \sum_i h_i s_i + \sum_{i < j} J_{ij} s_i s_j$ for classical spin variables $s_i = \pm 1$, and real coefficients h_i, J_{ij} . While generalizations to non-Ising cost functions for NDAR are possible, for the rest of the paper we will suppose we are given such a problem cost function we wish to minimize, along with a suitable quantum stochastic computational method for sampling from possible problem instance solutions.

For quantum optimization methods, such as those based on QAOA and generalizations, variable values are typically mapped to qubit basis states $|0\rangle, |1\rangle$. The cost function is then mapped [56, 57] to a diagonal cost Hamiltonian

$$H = \sum_i h_i Z_i + \sum_{i < j} J_{ij} Z_i Z_j + \dots \quad (1)$$

where Z_i is Pauli Z operator acting on qubit i , and $J_{i,j} \in \mathbb{R}$ is interaction strength between qubits i, j , while h_i is a local field magnitude.

Bitflip transforms Consider the unitary “bitflip” operator $P_{\mathbf{y}} = \bigotimes_{i=0}^{n-1} X_i^{y_i}$ that acts to flip

the $|0\rangle, |1\rangle$ basis states as specified by the bit-string $\mathbf{y} \in \{0, 1\}^n$. This change-of-basis can be applied to H by incorporating the change of signs to the weights of h_i and $J_{i,j}$ as $H \rightarrow H^{\mathbf{y}}$ with

$$\begin{aligned} H^{\mathbf{y}} &= P_{\mathbf{y}} H P_{\mathbf{y}} \\ &= \sum_i (-1)^{y_i} h_i Z_i + \sum_{i < j} (-1)^{y_i + y_j} J_{i,j} Z_i Z_j + \dots \end{aligned} \quad (2)$$

This transformation preserves cost Hamiltonian eigenvalues, with eigenvectors (candidate problem solutions) permuted under $P_{\mathbf{y}}$. In particular, under this transformation the $|0 \dots 0\rangle$ state is mapped to $|y_0 \dots y_{n-1}\rangle$.

Gauge-transformed QAOA circuits Consider a standard QAOA setup, where the quantum circuit consists of alternating layers of phase and mixer operators applied to the initial state $|+\rangle^{\otimes n}$, followed by measurement in the computational basis. The phase and mixer operators at layer l are given by $U_P(\gamma_l) = \exp(-i\gamma_l H)$ and $U_M(\beta_l) = \exp(-i\beta_l \sum_j^n X_j)$, respectively. The parameters γ_l and β_l control the algorithm performance and they are assumed to be set by means of a *parameter setting strategy* [58]. The physical implementation of QAOA implicitly assumes a particular choice of mapping of the problem variables and H . In the noiseless setting, QAOA circuits are robust to the choice of basis in the sense that bitflip transformations of the cost Hamiltonian $H \rightarrow H^{\mathbf{y}}$ are equivalent to transforming the entire circuit $U \rightarrow U^{\mathbf{y}}$. This follows easily from the properties $U_P(H^{\mathbf{y}}) = P_{\mathbf{y}} U_P(H) P_{\mathbf{y}}$, $P_{\mathbf{y}} U_M P_{\mathbf{y}} = U_M$, and $P_{\mathbf{y}} |+\rangle^{\otimes n} = |+\rangle^{\otimes n}$ as shown in Appendix A. Hence all resulting probabilities of measuring particular outcomes (via Born’s rule) remain the same, up to a known bitflip transformation, and so cost Hamiltonian expectation values are invariant. Nevertheless, in general, implementing $U_P(H^{\mathbf{y}})$ gives different unitaries for different $\mathbf{y}$.

Noise breaks bitflip gauge symmetry Clearly, we do not expect such bitflip gauge invariance of QAOA to hold in the presence of noise, since generally the non-controlled part of the evolution has no a priori reason to respect a given symmetry. Consider a simple example of the independent amplitude-damping channel. In the Kraus operator representation, this corresponds to a probabilistic application of the

transition operator $|0\rangle\langle 1|$ to each qubit. Since $|0\rangle\langle 1|$ does not commute with bitflips, the QAOA circuit probabilities become dependent on the particular problem mapping and are no longer invariant under bitflip transformations. Hence QAOA’s performance will depend on the particular choice of $\mathbf{y}$ for the cost Hamiltonian $H^{\mathbf{y}}$, even though different $H^{\mathbf{y}}$ all encode the same problem. One way of interpreting this example is that applying a bitflip gauge transformation effectively changes the stationary state of the noisy dynamics (with respect to the problem to be solved), and thus can potentially improve the algorithm by aligning the attractor state with high-quality solutions of the cost function. For example, it is intuitively clear that in the case of amplitude damping, the states featuring low Hamming weights (more 0s than 1s) are favored by noisy evolution. This suggests that problem mappings for which low-Hamming-weight states have low cost might be beneficial. Indeed, for example, for the well-studied Maximum Cut problem, the state $|00\dots 0\rangle$ corresponds to the worst possible cut (i.e., the empty set) in the standard formulation. If we instead, say, bitflipped the mapping of half the variables, then $|00\dots 0\rangle$ would instead correspond to a cut near the random guessing expectation value (i.e., half of the edges), a dramatic improvement over zero. We exploit these observations in the following sections to improve QAOA runs through NDAR, where we utilize an adaptive approach to account for the noise present in the Ankaa-2 QPU.

3 Noise-Directed Adaptive Remapping Algorithm

Inspired by the above observations, we propose the Noise-Directed Adaptive Remapping (NDAR) summarized as Algorithm 1 (and illustrated in Fig. 1). The NDAR consists of an external loop, in which each iteration involves implementing a stochastic optimization (such as QAOA) for a gauge-transformed Hamiltonian. The gauge transformation in each step is chosen based on the results of the optimization in the previous iteration – it is defined by the best-cost bitstring obtained so far. The NDAR ensures that the (postulated) steady state of the noise in each iteration step is the lowest energy state observed so far (see Fig. 1). This adap-

tively changes the optimized cost function in a way that greedily improves the energy of the attractor state, potentially leading to improved optimization.

In Algorithm 1, we leave a generic **stochastic optimizer** that encompasses some black-box optimization subroutine. In each NDAR iteration, the **stochastic optimizer** takes $H^{\mathbf{y}}$ and a number of allowed samples M as an input, and returns a list of M proposed solutions to the problem $H^{\mathbf{y}}$ (represented conventionally as bitstrings). This can correspond, for example, to a standard QAOA circuit runs coupled with an arbitrary parameter setting strategy. We also allow a general **TERMINATION RULE** to denote a stopping condition for the NDAR external loop. Various heuristics can be used – for example, terminate if the cost value did not decrease sufficiently relative to the previous iteration steps [59]. Note that in the Algorithm 1, we use $|0\dots 0\rangle$ as the attractor state. The algorithm can be easily modified for a different attractor which doesn’t have to be known a priori but could also be discovered experimentally (for example, by running identity circuits via techniques similar to unitary folding [60]).

4 Experimental demonstration

In this section, we present the results of experiments performed on 82-qubit subsystems of Rigetti’s superconducting transmon device Ankaa-2. Technical details of implementation are provided in Appendix C. Furthermore, in Appendix B we report numerical demonstrations of NDAR (with simulated amplitude damping) for small system sizes.

We start by empirically investigating how the choice of the bitflip gauge affects the results of variational optimization. We define the approximation ratio (AR) as the value of the cost function applied on a specific output (i.e., a bitstring) of the QAOA circuit after measurement in the computational basis, divided by the best-known cost of the problem (E_{GS}). Multiple runs result in multiple outputs which define an AR distribution.

We conjecture that QAOA optimization will perform better for bitflip gauges that transform a Hamiltonian in such a way that the




Fig. 2: Investigation of the effects of applying bitflip gauges to improve an $n = 82$ -qubit $p = 1$ QAOA in experiments on Rigetti's Ankaa-2 superconducting quantum device. a) Correlation between the approximation ratio (AR) of the attractor state $|0 \dots 0\rangle$ (X-axis) and the AR attained via QAOA (Y-axis) using transformed Hamiltonian. Differently colored regions of the histogram correspond to the AR estimators constructed from different percentiles of the AR distribution. We computed Pearson's correlation r_q and Spearman's rank correlation ρ_q coefficients at each percentile q . We obtained $r_{0.1\%} = 0.87^{+0.03}_{-0.04}$, $r_{10\%} = 0.964^{+0.009}_{-0.012}$, and $r_{100\%} = 0.984^{+0.004}_{-0.005}$; and $\rho_{0.1\%} = 0.86^{+0.04}_{-0.05}$, $\rho_{10\%} = 0.96^{+0.01}_{-0.02}$, and $\rho_{100\%} = 0.981^{+0.006}_{-0.007}$. Confidence intervals are 95% bias-corrected and accelerated (BC_a) paired-bootstrap intervals [61] with 10^5 resamples, calculated using the standard `scipy.stats` module [62]. Two-sided tests yielded p-values $< 10^{-58}$ for all cases. Taken together, these results indicate a strong, nearly linear association between the AR of the $|0 \dots 0\rangle$ state and the quality of optimization. The shaded region corresponds to the $y \leq x$ where the optimized AR is lower than the AR of the attractor state (indicating no performance gain). Data points are aggregated for 10 Hamiltonian instances. b) Overview of Noise-Directed Adaptive Remapping algorithm (see Algorithm 1) performance as a function of iterations (X-axis). The Y-axis shows the AR of the best-found solution (out of $1000 \times 150 \times (i + 1)$, with i being the iteration index). Each dotted line corresponds to a different random Hamiltonian instance, solid lines indicate the maximal and minimal performance across 10 instances. Besides best-found solutions obtained via NDAR applied to QAOA (orange), we plot standard QAOA optimization (light blue), and a baseline that corresponds to uniform random sampling (grey). Since in the case of QAOA and random baseline the protocols are not iterative, the iteration index corresponds to increasing the number of cost function evaluations/gathered samples to match those of NDAR.

approximation ratio (AR) of the attractor state $\langle 0 \dots 0 | H^y | 0 \dots 0 \rangle / E_{GS}$ (or equivalently $\langle \mathbf{y} | H | \mathbf{y} \rangle / E_{GS}$) is higher. Hence we first investigate a correlation between said AR and the performance of the QAOA on the cost Hamiltonian after running the $\mathbf{y}$ -transformed circuit with the chosen parameter setting strategy. The Hamiltonian cost function we optimize corresponds to fully connected, weighted graphs

$$H = \sum_{i < j} J_{i,j} Z_i Z_j, \quad (3)$$

with 2-body interaction coefficients $J_{i,j}$ drawn uniformly at random from the bimodal distribution of values ± 1 (so-called Sherrington-Kirkpatrick (SK) model [63]).

We implement $p = 1$ QAOA optimization for $n = 82$ qubits with Tree-Structured Parzen Estimators (TPE) [64, 65] parameter setting strategy using a budget of $t = 100$ cost function evaluations (trials). Each trial corresponds to estimat-

ing the cost function value from $s = 1000$ samples at fixed parameters. The QAOA is implemented for a set of 10 distinct Hamiltonian cost functions. For each Hamiltonian instance, the optimization is run for 20 random bitflip gauges $P_{\mathbf{y}}$.

For each QAOA run, we collect the distribution of ARs and compare the result with $\langle \mathbf{y} | H | \mathbf{y} \rangle / E_{GS}$. We construct the estimators from three quantiles of the total distribution over samples (for the best trial), corresponding to averaging only a subset of all AR outputs [66] after sorting them in descending order (1, 100, and 1000 samples, which constitute the top 0.1%, 10%, and 100% of all of the output samples for the best trial, respectively). We aggregate the data for all test Hamiltonians and plot the results as 2-dimensional histograms on the left side of Fig 2.

We observe a clear positive correlation between the AR of the optimized distribution and the AR of the attractor state evaluated on the cost

Algorithm1 Noise-Directed Adaptive Remapping (Greedy version, attractor $\equiv |0 \dots 0\rangle$)

Input:

$H_{(0)}$: Hamiltonian encoding the cost function
stochastic optimizer: a sub-algorithm that performs a stochastic optimization of $H_{(l)}$ (gauge-transformed Hamiltonian in iteration step $l + 1$). Example: standard (noisy) QAOA ansatz.

M : total number of samples in each optimization step

TERMINATION RULE: a rule specifying termination of the loop. Example: the energy did not decrease relative to the previous iteration.

Pseudocode for iteration $j \geq 1$:

1. Run optimization for Hamiltonian $H_{(j-1)}$ using **stochastic optimizer**. The result is M bitstrings (measurement outcomes) $\{\mathbf{x}_{(j)}^i\}_{i=1}^M$.
 2. Calculate cost values $\{E_{(j)}^i = \langle \mathbf{x}_{(j)}^i | H_{(j-1)} | \mathbf{x}_{(j)}^i \rangle\}_{i=1}^M$ corresponding to the Hamiltonian $H_{(j-1)}$.
 3. Select the bitstring $\mathbf{x}^*$ corresponding to the minimum of the cost function $E_{(j)}^* = \min_i E_{(j)}^i$.
if **TERMINATION RULE**:
 return lowest energy solution found overall
else:
 proceed to step 4.
 4. Apply bitflip gauge transform $P_{\mathbf{x}^*}$ to $H_{(j-1)}$, obtaining new Hamiltonian $H_{(j)} := H_{(j-1)}^{\mathbf{x}^*}$ (cf. Eq. (2)). This ensures that the transformed Hamiltonian has the property $\langle 00 \dots 0 | H_{(j)} | 00 \dots 0 \rangle = E_{(j)}^*$. Note that $H_{(j)}$ is equivalent to the original cost Hamiltonian $H_{(0)}$ via the product of gauge transformations applied so far, and so the solution to the original problem is easily recovered. Set $j = j + 1$. Proceed to step 1.
-

Hamiltonian for all quantiles, as confirmed by the estimated values of Pearson’s and Spearman’s correlation coefficients [67] close to 1.0, reported in the Figure’s caption. Importantly, there is a clear separation between the ARs of best-found solutions (top 0.1%) and the ARs of the attractor state, as seen by the fact that all data points lie above the shadowed $x \leq y$ region. This variance suggests that choosing a proper bitflip gauge and collecting sufficient samples can indeed aid the algorithm performance in experiments.

We exploit this observation further by implementing the Noise-Directed Adaptive Remapping algorithm applied to find solutions to the same class of $n = 82$ -qubit Hamiltonians with the same ansatz ($p = 1$ QAOA). Each step of the NDAR loop entails implementing TPE parameter search using $t = 150$ trials, with $s = 1000$ samples each. Following Ref. [8], we additionally allow the TPE loop to optimize over categorical variable controlling the gates’ ordering of the ansatz (see Appendix C for details). We specify **TERMINATION RULE** for NDAR (recall Algorithm 1) that terminates the loop if neither the minimal found energy nor the cost function (i.e., mean energy estimator) decreases compared to the previous step.

The results are presented on the right side of Fig. 2. The data points correspond to the best-found solution, meaning the best bitstring out of $150 * 1000 * (i + 1)$ samples (i being iteration step counting from 0). The NDAR QAOA results from QPU are colored orange. We also implement the standard QAOA (light blue) with the same number of function calls as NDAR. For reference, we provide the best approximation ratios obtained from the same number of shots using uniformly random sampling (grey).

We observe that Noise-Directed Adaptive Remapping improves the best-found solutions from range $\approx [0.34, 0.51]$ (variation between Hamiltonians for standard QAOA, blue lines) to $\approx [0.9, 0.96]$. In particular, without NDAR, standard QAOA (light blue) performs worse than random sampling (grey) when we consider the best sample, while NDAR (orange) allows us to outperform both the standard QAOA and random sampling already after 3 iterations. Importantly, the NDAR algorithm with QAOA converged after at most 20 optimization runs (achieving high-quality results much earlier),

suggesting that the overhead from using NDAR is not prohibitively high for $n = 82$ qubits. Note that we did not fine-tune the meta-parameters of the TPE loop, and it is likely that both performance and convergence can be improved.

In Fig. 3a, we present a 3-dimensional visualization of how the AR distributions evolve as a function of NDAR iteration (aggregating all instances, differently from Fig. 1). We observe how the distributions converge towards higher approximation ratio regions, following the increasing quality of the solution corresponding to the attractor state. One of the important effects of noise on QAOA optimization is that it highly contracts the attainable search space (for example, in the case of strong amplitude damping, this corresponds to restricting the search to only low hamming weights basis states). At the same time, we expect that NDAR, due to the re-labeling of basis elements, will effectively expand that space (by effectively remapping the Hamming sectors), allowing to concentrate the search in regions where better solutions lie.

To investigate that effect, in Figs. 3b and 3c, we plot an evolution of two types of Hamming weight (HW) distributions. The first type (middle plot) corresponds to physically registered measurement outcomes ('raw' data) that are highly concentrated close to an attractor state (which experimentally appears to be closer to $|1 \dots 1\rangle$ rather than $|0 \dots 0\rangle$), a fact we attribute to possible overheating of the chip, see also Appendix C.2). This demonstrates how physical noise restricts the search space and explains low approximation ratios attained in Fig. 2b, even if QAOA (without NDAR) is given many cost function evaluations. The second type (right plot) corresponds to the *effective* (or remapped) Hamming weight that results from gauge transformations applied to the measurement basis elements. The samples distributions are initially concentrated close to very high HW (initial optimization, iteration 0, red), and then exhibit oscillatory behavior with increasing iterations, to finally converge around $\frac{n}{2} = 41$ Hamming weights (last iteration, blue). Note that for generic random SK Hamiltonians, we expect the ground states to have Hamming weight close to $\frac{n}{2}$, thus the NDAR convergence to that region is a desirable behavior.

5 Conclusions and future work

We introduced Noise-Directed Adaptive Remapping (NDAR), a new algorithmic framework for quantum approximate optimization in the presence of certain types of noise. NDAR iteratively changes the encoding of the target Hamiltonian in a way that adaptively attempts to align the quantum evolution with the noise dynamics. We believe that the Noise-Directed Adaptive Remapping algorithm has the potential to become a standard method in the general toolbox of quantum optimization in the non-fault-tolerant regime, and, more generally, a technique to exploit the interplay of symmetries and open-system dynamics to improve algorithm performance. In particular, as many superconducting, photonic, and electronic quantum information processing platforms are affected by amplitude-damping-like noise, we expect NDAR to be particularly useful for optimization implemented in those devices (including quantum annealers and analog simulators). At the same time, it is important to emphasize the limitations of our method. It relies heavily on assumptions about the noise model—specifically, that it features a classical attractor state. Such a model does not typically apply to certain platforms, such as ion traps or neutral atoms, which can make the vanilla NDAR impractical. An interesting open question is how to generalize NDAR (for example, by using more general gauge transformations) so that it can be applied to such hardware. Alternatively, one could artificially introduce noise into the protocol to enforce additional bias, thereby allowing NDAR to exploit knowledge of previously found solutions (see Ref. [68] for a promising study in this direction).

In the previous sections, we empirically demonstrated the effectiveness of NDAR in a proof-of-concept $p = 1$ standard QAOA optimization on Rigetti's 82-qubit QPU, observing large performance gains over vanilla QAOA without NDAR. It is worth stressing that those gains are *not* due to noise-mitigation effects (as we are not getting closer to the ideal QAOA), but apparently come from the feedback-loop structure of the protocol – in fact, we suspect that noiseless QAOA with $p = 1$ could perform worse than the presented NDAR implementation. While the presented experimental results seem promising, more analytical and numerical work is required to better un-




Fig. 3: Evolution of the samples' distributions properties as a function of NDAR iteration. The Y-axis is NDAR iteration, and the Z-axis represents a (normalized) number of occurrences. The X-axis corresponds to the a) Approximation Ratio, b) the physically recorded, raw Hamming weight, and c) the effective Hamming weight attained due to NDAR bitstrings relabeling. First and last iterations are colored red and blue, respectively; while intermediate iterations are graduated green. The plots show aggregated data for 10 Hamiltonians. For each Hamiltonian, there is a total of 1000 samples shown for each iteration (taken from the optimized trials from the same data as in Fig. 2b).

derstand the general applicability of our method, as well as its dependence on the noise type and strength. Moreover, our study unveils numerous avenues for further research, a few of which we now outline.

First, we note that the proposed formulation of NDAR is intentionally very basic, and many heuristics can be applied to extend it. For example, the outer NDAR loop is currently greedy, meaning we always choose the new bitflip gauge based on the best-found solution. This has clear downsides, as follows from studies of classical greedy algorithms. It is possible that more flexible strategies would be beneficial – such as allowing for optimization over multiple low-energy gauges, performing the choice of the bitflip gauge based on its Hamming weight or typicality, or performing an additional local (classical) search in the neighborhood of the best-found solutions.

Second, we tailored the presentation on the Ising model, which is very general in its utility and has the convenience of supporting an easy remapping of the Hamiltonian via a simple change of the signs of its coefficients. However, similar simple transformations exist in other more complex models (such as the q-state Potts model that can be related to constrained optimization of the coloring type [69]). In general, NDAR could be formulated for polynomial cost functions of integer variables if an appropriate basis is defined, which could also include qudits encodings [70, 71]. The study of this rich space

of mapping transformations and Hamiltonian encodings is likely a very exciting route of investigation.

Third, we believe it would be interesting to implement NDAR with different algorithms than QAOA. For example, NDAR might be conceptually related to recently studied dissipative quantum algorithms that aim to mimic the cooling processes in Nature to find low-energy states, see Refs. [45, 46, 47, 48, 49]. Bath-engineering techniques, which have shown promise to avoid the barren plateau problem, might also be coupled with our remappings [72]. Moreover, experimental tests with different types of unitary ansätze (including those that are not necessarily bitflip-gauge invariant, such as Quantum Mixer-Phaser Ansatz [73]) and optimization algorithms are warranted [2, 6, 74, 75, 76, 77, 8, 78, 79, 80, 81]. In particular, it would be intriguing to combine NDAR with methods such as ADAPT-QAOA [6]; with iterative/recursive versions of QAOA [82, 5, 83]; with multilevel/multigrid quantum optimization strategies [84, 85, 86, 87, 88, 89]; with alternative qubit-efficient mappings/encodings [90, 91, 92]; with warm-start techniques [50, 51, 52, 53] or with recently-proposed quantum optimization schemes that generate candidate solutions via expectation-values statistics [93, 94].

Finally, we should note that NDAR idea could be applied to a variety of non-quantum Ising machines [54], including digital algorithms, where

the noise can be present because of hardware design, or could be added as a functional driver to exploit the attractor effect.

Code availability

GitHub repository containing the Python code for implementation of NDAR is available under github.com/usra-riacs/quantum-approximate-optimization [95].

Acknowledgments

This work was supported by the Defense Advanced Research Projects Agency (DARPA) under Agreement No. HR00112090058 and IAA8839, Annex 114. Authors acknowledge support under NASA Academic Mission Services under contract No. NNA16BD14C. We thank Maxime Dupont, Bram Evert, Mark Hodson, Stephen Jeffrey, Aaron Lott, Salvatore Mandrà, and Bhuvanesh Sundar for fruitful discussions at various stages of this project.

References

- [1] Edward Farhi, Jeffrey Goldstone, and Sam Gutmann. “A quantum Approximate Optimization Algorithm” (2014). [arXiv:1411.4028](https://arxiv.org/abs/1411.4028).
- [2] Stuart Hadfield, Zhihui Wang, Bryan O’gorman, Eleanor G Rieffel, Davide Venturelli, and Rupak Biswas. “From the quantum approximate optimization algorithm to a quantum alternating operator ansatz”. *Algorithms* **12**, 34 (2019).
- [3] Amira Abbas, Andris Ambainis, Brandon Augustino, Andreas Bäertschi, Harry Buhrman, Carleton Coffrin, Giorgio Corbiana, Vedran Dunjko, Daniel J Egger, Bruce G Elmgreen, et al. “Challenges and opportunities in quantum optimization”. *Nature Reviews Physics* **Page 718–735** (2024).
- [4] Natasha Sachdeva, Gavin S. Harnett, Smarak Maity, Samuel Marsh, Yulun Wang, Adam Winick, Ryan Dougherty, Daniel Canuto, You Quan Chong, Michael Hush, Pranav S. Mundada, Christopher D. B. Bentley, Michael J. Biercuk, and Yuval Baum. “Quantum optimization using a 127-qubit gate-model Ibm quantum computer can outperform quantum annealers for nontrivial binary optimization problems” (2024). [arXiv:2406.01743](https://arxiv.org/abs/2406.01743).
- [5] Maxime Dupont, Bram Evert, Mark J Hodson, Bhuvanesh Sundar, Stephen Jeffrey, Yuki Yamaguchi, Dennis Feng, Filip B Maciejewski, Stuart Hadfield, M Sohaib Alam, et al. “Quantum-enhanced greedy combinatorial optimization solver”. *Science Advances* **9**, eadi0487 (2023).
- [6] Linghua Zhu, Ho Lun Tang, George S. Barron, F. A. Calderon-Vargas, Nicholas J. Mayhall, Edwin Barnes, and Sophia E. Economou. “Adaptive quantum approximate optimization algorithm for solving combinatorial problems on a quantum computer”. *Phys. Rev. Res.* **4**, 033029 (2022).
- [7] Maxime Dupont and Bhuvanesh Sundar. “Extending relax-and-round combinatorial optimization solvers with quantum correlations”. *Physical Review A* **109** (2024).
- [8] Filip B Maciejewski, Stuart Hadfield, Benjamin Hall, Mark Hodson, Maxime Dupont, Bram Evert, James Sud, M Sohaib Alam, Zhihui Wang, Stephen Jeffrey, et al. “Design and execution of quantum circuits using tens of superconducting qubits and thousands of gates for dense Ising optimization problems”. *Physical Review Applied* **22**, 044074 (2024).
- [9] Maxime Dupont, Bhuvanesh Sundar, Bram Evert, David E. Bernal Neira, Zedong Peng, Stephen Jeffrey, and Mark J. Hodson. “Benchmarking quantum optimization for the maximum-cut problem on a superconducting quantum computer”. *Phys. Rev. Appl.* **23**, 014045 (2025).
- [10] S. Ebadi, A. Keesling, M. Cain, T. T. Wang, H. Levine, D. Bluvstein, G. Semeghini, A. Omran, J.-G. Liu, R. Samajdar, X.-Z. Luo, B. Nash, X. Gao, B. Barak, E. Farhi, S. Sachdev, N. Gemelke, L. Zhou, S. Choi, H. Pichler, S.-T. Wang, M. Greiner, V. Vuletić, and M. D. Lukin. “Quantum optimization of maximum independent set using Rydberg atom arrays”. *Science* **376**, 1209–1215 (2022).

- [11] Andrew Byun, Minhyuk Kim, and Jaewook Ahn. “Finding the Maximum Independent Sets of Platonic Graphs Using Rydberg Atoms”. *PRX Quantum* **3**, 030305 (2022).
- [12] Minhyuk Kim, Kangheun Kim, Jaeyong Hwang, Eun-Gook Moon, and Jaewook Ahn. “Rydberg quantum wires for maximum independent set problems”. *Nature Physics* **18**, 755–759 (2022). [arXiv:2109.03517](#).
- [13] Andrew D. King, Jack Raymond, Trevor Lanting, Richard Harris, Alex Zucca, Fabio Altomare, Andrew J. Berkley, Kelly Boothby, Sara Ejtemaee, Colin Enderud, Emile Hoskinson, Shuiyuan Huang, Eric Ladizinsky, Allison J. R. MacDonald, Gaelen Marsden, Reza Molavi, Travis Oh, Gabriel Poulin-Lamarre, Mauricio Reis, Chris Rich, Yuki Sato, Nicholas Tsai, Mark Volkmann, Jed D. Whittaker, Jason Yao, Anders W. Sandvik, and Mohammad H. Amin. “Quantum critical dynamics in a 5,000-qubit programmable spin glass”. *Nature* **617**, 61–66 (2023).
- [14] Minh-Thi Nguyen, Jin-Guo Liu, Jonathan Wurtz, Mikhail D. Lukin, Sheng-Tao Wang, and Hannes Pichler. “Quantum Optimization with Arbitrary Connectivity Using Rydberg Atom Arrays”. *PRX Quantum* **4**, 010316 (2023).
- [15] Zhenyu Cai, Ryan Babbush, Simon C. Benjamin, Suguru Endo, William J. Huggins, Ying Li, Jarrod R. McClean, and Thomas E. O’Brien. “Quantum error mitigation”. *Rev. Mod. Phys.* **95**, 045005 (2023).
- [16] Emanuel Knill and Raymond Laflamme. “Theory of quantum error-correcting codes”. *Physical Review A* **55**, 900 (1997).
- [17] Joschka Roffe. “Quantum error correction: an introductory guide”. *Contemporary Physics* **60**, 226–245 (2019).
- [18] Sergio Boixo, Tameem Albash, Federico M Spedalieri, Nicholas Chancellor, and Daniel A Lidar. “Experimental signature of programmable quantum annealing”. *Nature communications* **4**, 2067 (2013).
- [19] Andrew D King and Catherine C McGeoch. “Algorithm engineering for a quantum annealing platform” (2014). [arXiv:1410.2628](#).
- [20] Alejandro Perdomo-Ortiz, Joseph Fluegeman, Rupak Biswas, and Vadim N. Smelyanskiy. “A performance estimator for quantum annealers: Gauge selection and parameter setting” (2015). [arXiv:1503.01083](#).
- [21] Alejandro Perdomo-Ortiz, Bryan O’Gorman, Joseph Fluegeman, Rupak Biswas, and Vadim N. Smelyanskiy. “Determination and correction of persistent biases in quantum annealers”. *Scientific Reports* **6**, 18628 (2016).
- [22] Kristen L Pudenz. “Parameter setting for quantum annealers”. In 2016 IEEE high performance extreme computing conference (HPEC). *Pages 1–6*. IEEE (2016).
- [23] Elijah Pelofske, Georg Hahn, and Hristo Djidjev. “Optimizing the spin reversal transform on the d-wave 2000q”. In 2019 IEEE International Conference on Rebooting Computing (ICRC). *Pages 1–8*. IEEE (2019).
- [24] Aaron Barbosa, Elijah Pelofske, Georg Hahn, and Hristo N Djidjev. “Optimizing embedding-related quantum annealing parameters for reducing hardware bias”. In Parallel Architectures, Algorithms and Programming: 11th International Symposium, PAAP 2020, Shenzhen, China, December 28–30, 2020, Proceedings 11. *Pages 162–173*. Springer (2021).
- [25] Zoe Gonzalez Izquierdo, Shon Grabbe, Stuart Hadfield, Jeffrey Marshall, Zhihui Wang, and Eleanor Rieffel. “Ferromagnetically shifting the power of pausing”. *Physical Review Applied* **15**, 044013 (2021).
- [26] D-Wave Systems. “QPU Solvers: Spin-Reversal (Gauge) Transforms, https://docs.dwavesys.com/docs/latest/handbook_qpu.html#spin-reversal-gauge-transforms” (2025).
- [27] Google Quantum AI. “Qubit Picking, <https://quantumai.google/cirq/hardware/qubit-picking>” (2025).
- [28] IBM Quantum. “NoiseAdaptive-Layout, <https://quantum.cloud>.

- ibm.com/docs/en/api/qiskit/0.46/qiskit.transpiler.passes.NoiseAdaptiveLayout" (2025).
- [29] Prakash Murali, Jonathan M. Baker, Ali Javadi-Abhari, Frederic T. Chong, and Margaret Martonosi. "Noise-adaptive compiler mappings for noisy intermediate-scale quantum computers". In Proceedings of the Twenty-Fourth International Conference on Architectural Support for Programming Languages and Operating Systems. [Page 1015–1029](#). ASPLOS '19New York, NY, USA (2019). Association for Computing Machinery.
 - [30] Yanjun Ji, Kathrin F. Koenig, and Ilia Polian. "Improving the performance of digitized counterdiabatic quantum optimization via algorithm-oriented qubit mapping". [Phys. Rev. A](#) **110**, 032421 (2024).
 - [31] Atsushi Matsuo, Shigeru Yamashita, and Daniel J Egger. "A SAT approach to the initial mapping problem in SWAP gate insertion for commuting gates". [IEICE Transactions on Fundamentals of Electronics, Communications and Computer Sciences](#)[Page 2022EAP1159](#) (2023).
 - [32] Ruslan Shaydulin, Stuart Hadfield, Tad Hogg, and Ilya Safro. "Classical symmetries and the Quantum Approximate Optimization Algorithm". [Quantum Information Processing](#)**20** (2021).
 - [33] Alexey Galda, Xiaoyuan Liu, Danylo Lykov, Yuri Alexeev, and Ilya Safro. "Transferability of optimal QAOA parameters between random graphs". In 2021 IEEE International Conference on Quantum Computing and Engineering (QCE). [Pages 171–180](#). IEEE (2021).
 - [34] Ruslan Shaydulin and Alexey Galda. "Error mitigation for deep quantum optimization circuits by leveraging problem symmetries". In 2021 IEEE International Conference on Quantum Computing and Engineering (QCE). [Pages 291–300](#). IEEE (2021).
 - [35] Mahabubul Alam, Abdullah Ash-Saki, and Swaroop Ghosh. "Analysis of quantum approximate optimization algorithm under realistic noise in superconducting qubits" (2019). [arXiv:1907.09631](#).
 - [36] Cheng Xue, Zhao-Yun Chen, Yu-Chun Wu, and Guo-Ping Guo. "Effects of quantum noise on quantum approximate optimization algorithm". [Chinese Physics Letters](#) **38**, 030302 (2021).
 - [37] Jeffrey Marshall, Filip Wudarski, Stuart Hadfield, and Tad Hogg. "Characterizing local noise in QAOA circuits". [IOP SciNotes](#) **1**, 025208 (2020).
 - [38] Filip B. Maciejewski, Flavio Baccari, Zoltán Zimborás, and Michał Oszmaniec. "Modeling and mitigation of cross-talk effects in readout noise with applications to the Quantum Approximate Optimization Algorithm". [Quantum](#) **5**, 464 (2021).
 - [39] Frank Verstraete, Michael M. Wolf, and J. Ignacio Cirac. "Quantum computation and quantum-state engineering driven by dissipation". [Nature Physics](#) **5**, 633–636 (2009).
 - [40] M. J. Kastoryano, F. Reiter, and A. S. Sørensen. "Dissipative Preparation of Entanglement in Optical Cavities". [Physical Review Letters](#)**106** (2011).
 - [41] F. Reiter, A. S. Sørensen, P. Zoller, and C. A. Muschik. "Dissipative quantum error correction and application to quantum sensing with trapped ions". [Nature Communications](#) **8**, 1822 (2017).
 - [42] Jonathan Foldager, Arthur Pesah, and Lars Kai Hansen. "Noise-Assisted Variational Quantum Thermalization" (2021). [arXiv:2111.03935](#).
 - [43] Chenfeng Cao and Xin Wang. "Noise-Assisted Quantum Autoencoder". [Physical Review Applied](#)**15** (2021).
 - [44] Juha Leppäkangas, Nicolas Vogt, Keith R. Fratus, Kirsten Bark, Jesse A. Vaitkus, Pascal Stadler, Jan-Michael Reiner, Sebastian Zanker, and Michael Marthaler. "Quantum algorithm for solving open-system dynamics on quantum computers using noise". [Physical Review A](#)**108** (2023).
 - [45] Toby S. Cubitt. "Dissipative ground state preparation and the Dissipative Quantum Eigensolver" (2023). [arXiv:2303.11962](#).

- [46] Zhiyan Ding, Chi-Fang Chen, and Lin Lin. “Single-ancilla ground state preparation via lindbladians”. *Physical Review Research* **6** (2024).
- [47] Chi-Fang Chen, Michael J. Kastoryano, Fernando G. S. L. Brandão, and András Gilyén. “Quantum Thermal State Preparation” (2023). [arXiv:2303.18224](#).
- [48] X. Mi, A. A. Michailidis, S. Shabani, et al. “Stable quantum-correlated many-body states through engineered dissipation”. *Science* **383**, 1332–1337 (2024).
- [49] Chi-Fang Chen, Hsin-Yuan Huang, John Preskill, and Leo Zhou. “Local minima in quantum systems” (2023). [arXiv:2309.16596](#).
- [50] Daniel J Egger, Jakub Mareček, and Stefan Woerner. “Warm-starting quantum optimization”. *Quantum* **5**, 479 (2021).
- [51] Reuben Tate, Majid Farhadi, Creston Herold, Greg Mohler, and Swati Gupta. “Bridging classical and quantum with Sdp initialized warm-starts for QAOA”. *ACM Transactions on Quantum Computing* **4**, 1–39 (2023).
- [52] Reuben Tate, Jai Moondra, Bryan Gard, Greg Mohler, and Swati Gupta. “Warm-Started QAOA with Custom Mixers Provably Converges and Computationally Beats Goemans-Williamson’s Max-Cut at Low Circuit Depths”. *Quantum* **7**, 1121 (2023).
- [53] Jonathan Wurtz and Peter J Love. “Classically optimal variational quantum algorithms”. *IEEE Transactions on Quantum Engineering* **2**, 1–7 (2021).
- [54] Naeimeh Mohseni, Peter L McMahon, and Tim Byrnes. “Ising machines as hardware solvers of combinatorial optimization problems”. *Nature Reviews Physics* **4**, 363–379 (2022).
- [55] Giorgio Ausiello, Pierluigi Crescenzi, Giorgio Gambosi, Viggo Kann, Alberto Marchetti-Spaccamela, and Marco Protasi. “Complexity and approximation: Combinatorial optimization problems and their approximability properties”. *Springer Science & Business Media*. (2012).
- [56] Andrew Lucas. “Ising formulations of many NP problems”. *Frontiers in physics* **2**, 74887 (2014).
- [57] Stuart Hadfield. “On the representation of Boolean and real functions as Hamiltonians for quantum computing”. *ACM Transactions on Quantum Computing* **2**, 1–21 (2021).
- [58] David E. Bernal Neira, Robin Brown, Pratik Sathe, Filip Wudarski, Marco Pavone, Eleanor G. Rieffel, and Davide Venturelli. “Benchmarking the Operation of Quantum Heuristics and Ising Machines: Scoring Parameter Setting Strategies on Optimization Applications” (2024). [arXiv:2402.10255](#).
- [59] Walter Vinci and Daniel A Lidar. “Optimally stopped optimization”. *Physical Review Applied* **6**, 054016 (2016).
- [60] Tudor Giurgica-Tiron, Yousef Hindy, Ryan LaRose, Andrea Mari, and William J Zeng. “Digital zero noise extrapolation for quantum error mitigation”. In 2020 IEEE International Conference on Quantum Computing and Engineering (QCE). Pages 306–316. IEEE (2020).
- [61] Bradley Efron and Robert J Tibshirani. “An introduction to the bootstrap”. *Chapman and Hall/CRC*. (1994).
- [62] Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, and SciPy 1.0 Contributors. “SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python”. *Nature Methods* **17**, 261–272 (2020).
- [63] David Sherrington and Scott Kirkpatrick. “Solvable model of a spin-glass”. *Physical review letters* **35**, 1792 (1975).
- [64] James Bergstra, Daniel Yamins, and David Cox. “Making a science of model search: Hyperparameter optimization in hundreds of dimensions for vision architectures”. In Proceedings of the 30th International Conference on Machine Learning. Volume 28 of Proceedings of Machine Learning Research, pages 115–123. Atlanta, Georgia, USA (2013). PMLR. url: <https://proceedings.mlr.press/v28/bergstra13.html>.

- [65] Takuya Akiba, Shotaro Sano, Toshihiko Yanase, Takeru Ohta, and Masanori Koyama. “Optuna: A next-generation hyperparameter optimization framework” (2019). [arXiv:1907.10902](#).
- [66] Troels F Rønnow, Zhihui Wang, Joshua Job, Sergio Boixo, Sergei V Isakov, David Wecker, John M Martinis, Daniel A Lidar, and Matthias Troyer. “Defining and detecting quantum speedup”. [science](#) **345**, 420–424 (2014).
- [67] Daniel Zwillinger and Stephen Kokoska. “CRC standard probability and statistics tables and formulae”. [Crc Press](#). (1999).
- [68] Wai-Hong Tam, Hiromichi Matsuyama, Ryo Sakai, and Yu Yamashiro. “Enhancing NDAR with delay-gate-induced amplitude damping” (2025). [arXiv:2504.12628](#).
- [69] Kensuke Inaba, Takahiro Inagaki, Koji Igarashi, Shoko Utsunomiya, Toshimori Honjo, Takuya Ikuta, Koji Enbutsu, Takeshi Umeki, Ryoichi Kasahara, Kyo Inoue, et al. “Potts model solver based on hybrid physical and digital architecture”. [Communications Physics](#) **5**, 137 (2022).
- [70] Sergey Bravyi, Alexander Kliesch, Robert Koenig, and Eugene Tang. “Hybrid quantum-classical algorithms for approximate graph coloring”. [Quantum](#) **6**, 678 (2022).
- [71] A Barış Özgüler and Davide Venturelli. “Numerical gate synthesis for quantum heuristics on bosonic quantum processors”. [Frontiers in Physics](#) **10**, 900612 (2022).
- [72] Antonio Sannia, Francesco Tacchino, Ivano Tavernelli, Gian Luca Giorgi, and Roberta Zambrini. “Engineered dissipation to mitigate barren plateaus”. [npj Quantum Information](#) **10** (2024).
- [73] Ryan LaRose, Eleanor Rieffel, and Davide Venturelli. “Mixer-phaser ansätze for quantum optimization with hard constraints”. [Quantum Machine Intelligence](#) **4**, 17 (2022).
- [74] Rebekah Herrman, Phillip C Lotshaw, James Ostrowski, Travis S Humble, and George Siopsis. “Multi-angle quantum approximate optimization algorithm”. [Scientific Reports](#) **12**, 6781 (2022).
- [75] Jonathan Wurtz and Peter J. Love. “Counterdiabaticity and the quantum approximate optimization algorithm”. [Quantum](#) **6**, 635 (2022).
- [76] Alicia B. Magann, Kenneth M. Rudinger, Matthew D. Grace, and Mohan Sarovar. “Feedback-based quantum optimization”. [Physical Review Letters](#) **129** (2022).
- [77] Xiaoyuan Liu, Anthony Angone, Ruslan Shaydulin, Ilya Safro, Yuri Alexeev, and Lukasz Cincio. “Layer Vqe: A variational Approach for Combinatorial Optimization on Noisy Quantum Computers”. [IEEE Transactions on Quantum Engineering](#) **3**, 1–20 (2022).
- [78] Kostas Blekos, Dean Brand, Andrea Ceschini, Chiao-Hui Chou, Rui-Hao Li, Komal Pandya, and Alessandro Summer. “A review on quantum approximate optimization algorithm and its variants”. [Physics Reports](#) **1068**, 1–66 (2024).
- [79] Anthony Wilkie, Igor Gaidai, James Ostrowski, and Rebekah Herrman. “Quantum approximate optimization algorithm with random and subgraph phase operators”. [Physical Review A](#) **110** (2024).
- [80] Ruslan Shaydulin, Changhao Li, Shouvanik Chakrabarti, Matthew DeCross, et al. “Evidence of scaling advantage for the quantum approximate optimization algorithm on a classically intractable problem”. [Science Advances](#) **10** (2024).
- [81] V Vijendran, Aritra Das, Dax Enshan Koh, Syed M Assad, and Ping Koy Lam. “An expressive ansatz for low-depth quantum approximate optimisation”. [Quantum Science and Technology](#) **9**, 025010 (2024).
- [82] Sergey Bravyi, Alexander Kliesch, Robert Koenig, and Eugene Tang. “Obstacles to Variational Quantum Optimization from Symmetry Protection”. [Physical Review Letters](#) **125**, 260505 (2020).
- [83] Lucas T Brady and Stuart Hadfield. “Iterative quantum algorithms for maximum independent set”. [Physical Review A](#) **110**, 052435 (2024).
- [84] Hayato Ushijima-Mwesigwa, Ruslan Shaydulin, Christian FA Negre, Susan M Mniszewski, Yuri Alexeev, and Ilya

- Safro. “Multilevel combinatorial optimization across quantum architectures”. *ACM Transactions on Quantum Computing* **2**, 1–29 (2021).
- [85] Anthony Angone, Xiaoyuan Liu, Ruslan Shaydulin, and Ilya Safro. “Hybrid quantum-classical multilevel approach for maximum cuts on graphs”. In 2023 IEEE High Performance Extreme Computing Conference (HPEC). Pages 1–7. IEEE (2023).
- [86] Filip B Maciejewski, Bao G Bach, Maxime Dupont, P Aaron Lott, Bhuvanesh Sundar, David E Bernal Neira, Ilya Safro, and Davide Venturelli. “A multilevel approach for solving large-scale qubo problems with noisy hybrid quantum approximate optimization”. In 2024 IEEE High Performance Extreme Computing Conference (HPEC). Pages 1–10. IEEE (2024).
- [87] Bao Bach, Jose Falla, and Ilya Safro. “MLQAOA: Graph learning accelerated hybrid quantum-classical multilevel qaoa”. In 2024 IEEE International Conference on Quantum Computing and Engineering (QCE). Volume 1, pages 1–12. IEEE (2024).
- [88] Atithi Acharya, Romina Yalovetzky, Pierre Minssen, Shouvanik Chakrabarti, et al. “Decomposition Pipeline for Large-Scale Portfolio Optimization with Applications to Near-Term Quantum Computing” (2024). [arXiv:2409.10301](https://arxiv.org/abs/2409.10301).
- [89] Bao G Bach, Filip B. Maciejewski, and Ilya Safro. “Solving large-scale QUBO with transferred parameters from multilevel QAOA of low depth” (2025). [arXiv:2505.11464](https://arxiv.org/abs/2505.11464).
- [90] Bhuvanesh Sundar and Maxime Dupont. “Qubit-efficient quantum combinatorial optimization solver” (2024). [arXiv:2407.15539](https://arxiv.org/abs/2407.15539).
- [91] Marco Sciorilli, Lucas Borges, Taylor L Patti, Diego Garcia-Martin, Giancarlo Camilo, Anima Anandkumar, and Leandro Aolita. “Towards large-scale quantum optimization solvers with few qubits”. *Nature Communications* **16**, 476 (2025).
- [92] Marco Sciorilli, Giancarlo Camilo, Thiago O. Maciel, Askery Canabarro, Lucas Borges, and Leandro Aolita. “A competitive NISQ and qubit-efficient solver for the labs problem” (2025). [arXiv:2506.17391](https://arxiv.org/abs/2506.17391).
- [93] Maxime Dupont and Bhuvanesh Sundar. “Extending relax-and-round combinatorial optimization solvers with quantum correlations”. *Physical Review A* **109** (2024).
- [94] Maxime Dupont, Tina Oberoi, and Bhuvanesh Sundar. “Optimization via Quantum Preconditioning”. *Physical Review Applied* **24** (2025).
- [95] F. B. Maciejewski, B. G. Bach, J. Biamente, S.A. Hadfield, and D. Venturelli. “quapopt – open source GitHub repository for quantum approximate optimization”. <https://github.com/usra-riacs/quantum-approximate-optimization> (2025).
- [96] Yuichi Hirata, Masaki Nakanishi, Shigeru Yamashita, and Yasuhiko Nakashima. “An Efficient Method to Convert Arbitrary Quantum Circuits to Ones on a Linear Nearest Neighbor Architecture”. In 2009 Third International Conference on Quantum, Nano and Micro Technologies. Pages 26–33. (2009).
- [97] Quantum AI team and collaborators. “qsim – Optimized quantum circuit simulators”. <https://quantumai.google/qsim> (2025).
- [98] Sergei V. Isakov, Dvir Kafri, Orion Martin, et al. “Simulations of Quantum Circuits with Approximate Noise using qsim and Cirq” (2021). [arXiv:2111.02396](https://arxiv.org/abs/2111.02396).
- [99] Flavio Baccari, Christian Gogolin, Peter Wittek, and Antonio Acín. “Verifying the output of quantum optimizers with ground-state energy lower bounds”. *Physical Review Research* **2** (2020).
- [100] Salvatore Mandra, Ata Akbari Asanjan, Lucas Brady, Aaron Lott, David E. Bernal Neira, and Humberto Munoz Bauza. “PySA: Fast Simulated Annealing in Native Python”. <https://github.com/nasa/pysa> (2023).

Supplementary Materials for
Improving Quantum Approximate Optimization
by Noise-Directed Adaptive Remapping

A Gauge transformations of QAOA circuits

Recall an n -qubit QAOA circuit of depth p and parameters $\gamma, \beta \in \mathbb{R}^p$ is given by

$$U(H; \gamma, \beta) := \prod_{l=p}^1 U_M(\beta_l) U_P(\gamma_l), \quad (4)$$

with phase separator

$$U_P(\gamma_l) := U_P(H; \gamma_l) = \exp(-i\gamma_l H) \quad (5)$$

and mixer operator

$$U_M(\beta_k) := \bigotimes_{j=0}^{n-1} \exp(-i\beta_k X_j). \quad (6)$$

The QAOA operator $U(H; \gamma, \beta)$ is applied to the standard initial state $|s\rangle = |+\rangle^{\otimes n}$, where $|+\rangle = \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$. We now formalize the observation that bitflip gauge transformations of the input cost Hamiltonian do not affect QAOA performance in the ideal setting. The same result does not generally apply to other ansätze. Let $Q_{\mathbf{x}} := |\mathbf{x}\rangle\langle\mathbf{x}|$ such that $\langle\psi|Q_{\mathbf{x}}|\psi\rangle = \langle Q_{\mathbf{x}}\rangle_{\psi}$ gives the probability of measuring bitstring $\mathbf{x}$ for any quantum state $|\psi\rangle$. In what follows, we denote $Q_{\mathbf{x}^{\mathbf{y}}} = |\mathbf{x}^{\mathbf{y}}\rangle\langle\mathbf{x}^{\mathbf{y}}| := P_{\mathbf{y}}|\mathbf{x}\rangle\langle\mathbf{x}|P_{\mathbf{y}}$.

Lemma 1. *Bitflip (spin-reversal) gauge transformations $H \rightarrow H^{\mathbf{y}}$ do not affect the performance of QAOA in the noise-free setting, i.e. for any fixed p, γ, β and $\mathbf{x}, \mathbf{y} \in \{0, 1\}^n$ we have*

$$\langle Q_{\mathbf{x}^{\mathbf{y}}} \rangle_{U(H^{\mathbf{y}}; \gamma, \beta)|s\rangle} = \langle Q_{\mathbf{x}} \rangle_{U(H; \gamma, \beta)|s\rangle},$$

i.e. the probability of measuring the transformed bitstring $\mathbf{x}^{\mathbf{y}}$ is the same as that of measuring $\mathbf{x}$ originally, and hence

$$\langle H^{\mathbf{y}} \rangle_{U(H^{\mathbf{y}}; \gamma, \beta)|s\rangle} = \langle H \rangle_{U(H; \gamma, \beta)|s\rangle}.$$

Proof. Here we take advantage of the useful properties $P_{\mathbf{y}}^{\dagger} = P_{\mathbf{y}}$, $P_{\mathbf{y}}^2 = I$, $P_{\mathbf{y}}X_jP_{\mathbf{y}} = X_j$, and $P_{\mathbf{y}}|s\rangle = |s\rangle$, for any $\mathbf{y} \in \{0, 1\}^n$. Let $p = 1$; the arbitrary p case follows similarly. Fix γ, β and observe that

$$U(H^{\mathbf{y}}) = e^{-i\beta_l B} e^{-i\gamma_l P_{\mathbf{y}} H P_{\mathbf{y}}} = e^{-i\beta_l B} P_{\mathbf{y}} e^{-i\gamma_l H} P_{\mathbf{y}} = P_{\mathbf{y}} e^{-i\beta_l B} e^{-i\gamma_l H} P_{\mathbf{y}} = P_{\mathbf{y}} U(H) P_{\mathbf{y}},$$

and so applying the properties above gives

$$\langle \mathbf{x} | U(H^{\mathbf{y}}) | s \rangle = \langle \mathbf{x} | P_{\mathbf{y}} U(H) P_{\mathbf{y}} | s \rangle = \langle \mathbf{x}^{\mathbf{y}} | U(H) | s \rangle,$$

i.e., quantum amplitudes are preserved which implies measurement probabilities are too. $\square$

We remark that related techniques are applied to obtain results concerning the interplay of QAOA and classical cost function symmetries in [32]. In the setting of that work we are given a transformation A which commutes with a particular cost Hamiltonian H , i.e. $AH - HA = 0$, and the affect on QAOA dynamics and outcomes is studied. In our distinct setting $P_{\mathbf{y}}$ and H do not commute as this would imply $H^{\mathbf{y}} = H$.

The lemma shows that on an ideal quantum device, gauge transformations would have no effect on QAOA performance. In contrast, the proofs breaks down in the presence of noise.

Remark 1. Note that the above is mathematically equivalent to applying basis transformation defined by P_y to every unitary in the circuit, but is operationally different from it – change of basis would correspond to applying P_y to every unitary, as opposed to just phase separators. The equivalence between running QAOA for H^y and performing a global change of basis does not hold in the same way for arbitrary quantum algorithms. For other ansätze that similarly do not depend on the choice of bitflip gauge, the algorithm parameters may have to transform also if, for example, the mixer of a more general QAOA circuit [2] does not commute with the particular P_y . Other cases of parameterized quantum circuits may depend on the particular ordering (permutation) of variables in the chosen encoding, such as Quantum Alternating Mixer-Phaser Ansatz [73]. Clearly, NDAR can be applied in such scenarios as well. In particular, one can imagine a black-box algorithm where the circuit does not depend on the target problem. Then the NDAR feedback loop would not modify the circuit at all, but only the measured gauge-transformed Hamiltonian.

B Numerical studies




Fig. 4: Investigation of the effects of applying bitflip gauges to improve an $n = 16$ -qubit variational optimization with $p = 1$ QAOA in simulations. The data presented is of the same type as Fig. 2 in the main text, see description of that figure for details.

Recall that in Fig 2 of the main text, we presented the results of two experiments on the Ankaa-2 device. The left side of Fig 2a presents a correlation investigation between the approximation ratio of the postulated attractor state (denoted as $|0 \dots 0\rangle$) and QAOA’s performance, while the right side of Fig. 2 shows the overview of Noise-Directed Adaptive Remapping algorithm implementation.

We now consider the results of an analogous investigation of what we have presented in Fig. 2 for simulated systems. In our simulations, we use non-identical, uncorrelated strong amplitude-damping noise. A tensor-product, local two-qubit channel is applied to the circuit after every application of every gate. The 2-qubit gates are set to be more noisy than the 1-qubit gates. We assume linear connectivity of a device and implement an efficient SWAP network from [96] (same as in experiments in the main text). This significantly increases circuit depth, making the simulations closer to what happens in experiments. We assume that the native gate set consists of arbitrary ZZ and XY rotations. All simulations are performed using Qsim framework [97] that implements the quantum trajectories method [98] to approximate sampling from noisy circuits.

We implement $p = 1$ QAOA optimization for $n = 16$ qubits for the same class of Hamiltonians as in the main text (SK model). The QAOA is implemented for a set of 10 random Hamiltonians. The aggregated results are presented in Fig. 4.

For the correlations plot (left side of Fig. 4), for each Hamiltonian, the optimization is run independently for 20 random bitflip gauges, i.e., at a given optimization, the original Hamiltonian is gauge-transformed via a random bitflip P_y and such process is repeated independently for 20 random y s. Importantly, the optimizer seed is fixed in all simulations, meaning that the only sources of performance difference should be statistical noise, the gauge transformation, and the the varying hardness

of optimization of particular Hamiltonian instances. For each optimization run, we estimate the approximation ratio of the final output and compare the result with the corresponding approximation ratio. The optimization is run with $t = 100$ cost function evaluations, each involving gathering $s = 256$ samples. For reference, we plot also the results of noiseless optimization.

The plot on the left of Fig. 4 shows, for the noiseless setting, no correlation between the AR of the attractor state and QAOA’s performance. This is expected since without noise all gauges are equivalent (recall discussion in Appendix A). In the case of amplitude damping, however, we observe a positive correlation for mid-quantiles (256 best samples, and 50 best samples, corresponding to 100% and $\simeq 20\%$ of all samples). Interestingly, for poor gauge choices (AR below 0.0), the performance is mostly flat for those quantiles, indicating that the optimizer gets stuck. For the best-found solution (top 0.4% of results), we do not observe a correlation between the two variables. We expect this to be a small-size effect ($n = 16$). Indeed, note that the actual ground state (approximation ratio 1) is found in multiple cases regardless of the gauge.

Due to the above observations, we decided to highly reduce the number of trials and samples for the NDAR implementation – otherwise the NDAR would often converge at the first iteration (as indicated by results in Fig 4a), rendering the results uninteresting. The plot on the right of Fig. 4 shows the performance overview of the Noise-Directed Adaptive Remapping algorithm applied to noisy simulated QAOA. The QAOA is guided by the TPE optimizer with $t = 20$ cost function evaluations, each involving gathering $s = 100$ samples. We observe that the NDAR applied to noisy QAOA (orange) finds the ground state in all tested cases after 3 iterations. Results of standard QAOA (light blue) show that in that case, the performance is spread out across Hamiltonians, indicating that for some instances, the strong amplitude damping noise prohibits finding the ground state. The reference random uniform baseline (grey), given the same small number of samples, achieves comparable performance. The above results indicate that NDAR can aid the noisy QAOA even when the computational resources (function calls and samples) are highly limited.

C Additional experimental data

C.1 QAOA circuit implementation




Fig. 5: Schematics of the Rigetti’s Ankaa-2 superconducting transmon chip topology with respect to 2-qubit iSWAP operation. Missing edges not shown. Source: rigetti.com.

The Rigetti’s Ankaa-2 (Fig. 5) device’s native gate set consists of iSWAP, $\pm \frac{\pi}{2}$ RX rotations and arbitrary RZ rotations. The device’s connectivity (w.r.t. iSWAP gate) is that of a rectangular 7 by 12 grid (with some edges missing, not shown in Fig. 5). To implement QAOA circuits for fully connected Hamiltonians, we implement an efficient SWAP network from Ref. [96] on linear chains of qubits. Since the phase separator (the $\exp(-i\gamma H)$ operator) of the QAOA ansatz implemented with SWAP network has a structure of multiple RZZ rotations followed by SWAP gates, we compile the

product of the two gates jointly (a similar trick was applied for efficient compilation in Ref. [8]). Each implementation of RZZ+SWAP requires applying 3 ISWAP gates and at most 4 $\text{RX}(\pm\frac{\pi}{2})$ and 5 RZ rotations (the exact number depends on the qubit). Implementing a single layer of QAOA ansatz for n qubits requires consecutive application of n linear chains, each with $\approx \frac{n}{2}$ RZZ+SWAP gates. Thus, for the biggest experiments shown in the text with $n = 82$, the $p = 1$ QAOA ansatz consists of around $\approx 10,000$ ISWAP gates and $\approx 30,000$ single-qubit rotations.

C.2 Parameter setting strategy details

To implement variational optimization (for Figs. 1, 2, and 3) we use the Tree-Structured Parzen Estimators (TPE) implementation from the `optuna` package [65]. We use default hyperparameters (but with turned-off pruning trials). Each optimization was run with a random seed, and each optimization started from the same initial points (all angles set to 0.1).

As indicated in the main text, besides 2 variational angles (for $p = 1$ QAOA), we allow the optimizer to also vary between 10 random gates' orderings (GOs) of the ansatz (specified as categorical variables). In the noiseless scenario, allowed GOs do not affect the circuit. However, in real-life settings, it was demonstrated, for example in Ref. [8], that due to the noise, the optimization over GOs can highly improve performance. When implementing the Noise-Directed Adaptive Remapping algorithm, to allow more versatility, we change the allowed set of GOs at random in each loop step. To make a fair comparison when implementing standard QAOA (shown in Fig. 2b), we allow the optimizer to choose over the number of GOs equal to $10 * i^{\max}$, where $i^{\max}$ is iteration index at which NDAR converged (for the corresponding Hamiltonian instance). This gives the TPE optimizer access to the same total number of GOs for QAOA optimization as for the NDAR implementation. Note that standard QAOA is given also proportionally more function calls compared to a single NDAR step (i.e., $t = i^{\max} * 150$).

As indicated in the main text, the experimental samples analysis for the 82-qubit system revealed that, interestingly, the bitstrings were concentrated close to $|1\dots 1\rangle$ physical state, as opposed to $|0\dots 0\rangle$ (the steady state of the amplitude damping channel). We suspect that this is due to overheating of the chip, but we intend to investigate the physical causes of observed effects in future work. We note that since the Hamiltonians we implemented consisted of only ZZ interactions, they exhibit Z_2 symmetry, meaning that the energy remains invariant under a global $X^{\otimes n}$ transformation. This means that, for Hamiltonians considered in this work, $|0\dots 0\rangle$ has always the same energy as $|1\dots 1\rangle$. The consequence for the NDAR implementation is such that in practice the algorithm was not affected by whether we assumed that the physical attractor is $|0\dots 0\rangle$ or $|1\dots 1\rangle$.

Whenever the approximation ratio is reported, we calculate $\frac{E}{E_{\text{GS}}}$, where E_{GS} is the best-found solution obtained either via branch-and-bound SDP technique from Ref. [99] or via simulated annealing package PySA [100].