Investigating quantum approximate optimization algorithms under bang-bang protocols

Daniel Liang^{1,2}, Li Li (李力)³, and Stefan Leichenauer¹¹*X, The Moonshot Factory, Mountain View, California 94043, USA*²*Quantum Information Center, University of Texas at Austin, Austin, Texas 78712, USA*³*Google Research, Mountain View, California 94043, USA*

(Received 28 May 2020; accepted 13 August 2020; published 14 September 2020)

The quantum approximate optimization algorithm (QAOA) is widely seen as a possible usage of noisy intermediate-scale quantum (NISQ) devices. We analyze the algorithm as a bang-bang protocol with fixed total time and a randomized greedy optimization scheme. We investigate the performance of bang-bang QAOA on MAX-2-SAT, finding the appearance of phase transitions with respect to the total time. As the total time increases, the optimal bang-bang protocol experiences a number of jumps and plateaus in performance, which match up with an increasing number of switches in the standard QAOA formulation. At large times, it becomes more difficult to find a globally optimal bang-bang protocol and performances suffer. We also investigate the effects of changing the initial conditions of the randomized optimization algorithm and see that better local optima can be found by using an adiabatic initialization.

DOI: [10.1103/PhysRevResearch.2.033402](https://doi.org/10.1103/PhysRevResearch.2.033402)

I. INTRODUCTION

The use of quantum computation to solve problems deemed hard for classical computation is an area of massive interest in both the physics and computer science communities. One candidate algorithm for practical speedups on noisy intermediate-scale quantum (NISQ) devices is the quantum approximate optimization algorithm (QAOA) proposed by Farhi *et al.* [1]. The QAOA involves switching between two Hamiltonians, with the number of switches being defined by a parameter called p , as well as an optimization process to control how long each Hamiltonian should be applied.

As stated in the original QAOA paper, the common belief is that p controls the approximation ratio of QAOA, and so p should be as large as possible before the circuit becomes too deep and is overwhelmed by hardware noise [2]. Indeed, Farhi *et al.* [1] were able to show that as $p \rightarrow \infty$, QAOA is able to achieve a perfect approximation ratio, since in that limit QAOA is as powerful as adiabatic quantum computation [3,4].

However, a recent paper from Shaydulin and Alexeev [5] gave evidence to the contrary, stating that the optimization of variational parameters is difficult at large p , and performance improvements at large p are marginal when dealing with bounded computation in the optimization process. We provide further evidence of this. Inspired by Day *et al.* [6], we give data from a large-scale classical simulation of a modification to QAOA, which we call bang-bang QAOA, applied to the

problem of MAX-2-SAT. While not necessarily practical for NISQ devices, this modification acts as a thought experiment to show that even in the case where p is allowed to be fairly large, while the total time is instead bounded, one does not see large improvements with greater values of p . Similar to Shaydulin and Alexeev [5], we assert that this is because of a proliferation of local optima, making it difficult to find optima that are close to the global optima.

While we know that as $p \rightarrow \infty$ that one can choose the QAOA parameters to correspond to a Trotterized adiabatic quantum computation and achieve a perfect approximation ratio [1], in the finite p regime it is not fully understood whether or not the optimal parameters for QAOA should appear adiabatic [7–10]. We emphasize that while some of these works studied QAOA through the lens of bang-bang control theory, they ultimately focused on applying bang-bang control theory to the adiabatic algorithm and seeing if the result resembles QAOA. In contrast, we assume a bang-bang structure and use a randomized greedy optimization algorithm to study the solution space of bang-bang QAOA algorithms with a fixed total time. In the bang-bang QAOA model, we see that when total time is small, the best protocols do not appear adiabatic, but rather correspond to finite- p implementations of standard QAOA. When the total time is large, the proliferation of local optima means that our optimization procedure depends strongly on the initialization. Though we cannot say much about any global optima, we do see that adiabatic initialization provides a good heuristic for finding better protocols.

II. QAOA AND BANG-BANG PROTOCOLS

As its name suggests, the QAOA is a quantum-based algorithm for combinatorial optimization designed to find the set of inputs that approximately optimizes an efficiently

FIG. 1. Example of a bang-bang QAOA protocol with $N_b = 8$. It is equivalent to a $p = 2$ QAOA protocol with $\gamma_1 = T/N_b$, $\beta_1 = 3T/N_b$, $\gamma_2 = 2T/N_b$, and $\beta_2 = 2T/N_b$. The protocols are read from left-to-right in order of applying the Hamiltonians.

computable objective function. At a high level, it does so by encoding this objective function along the diagonal of a Hamiltonian. The algorithm then tries to find a circuit that efficiently brings the state $|0^n\rangle$ as close as possible to the ideal state by applying two different Hamiltonians. We will first describe the standard QAOA as given by Farhi *et al.* [1], followed by our bang-bang QAOA modification. Note that there exist a wide variety of other interesting modifications to the standard QAOA [11–13].

A. Standard QAOA

Let $E = \sum_{x \in \{0,1\}^n} f(x)|x\rangle\langle x|$ be the Hamiltonian that encodes the objective function f along its diagonal. E will be referred to as the *constraint Hamiltonian* while $X^{\otimes n} = (|0\rangle\langle 1| + |1\rangle\langle 0|)^{\otimes n}$ will be referred to as the *mixing Hamiltonian*. Now let $\beta_1, \dots, \beta_p, \gamma_1, \dots, \gamma_p$ be positive real parameters for QAOA with depth p .¹ The state produced by QAOA is then

$$|\psi\rangle = e^{i\beta_p X^{\otimes n}} e^{i\gamma_p E} \dots e^{i\beta_1 X^{\otimes n}} e^{i\gamma_1 E} H^{\otimes n} |0^n\rangle,$$

where $H = (|0\rangle\langle 0| + |1\rangle\langle 0| + |0\rangle\langle 1| - |1\rangle\langle 1|)/\sqrt{2}$ is the Hadamard operator. The $2p$ parameters are optimized based on the expectation value $\langle \psi | E | \psi \rangle$ in order to increase the chance of measuring a good input when $|\psi\rangle$ is measured in the computational basis.

B. Bang-bang QAOA

A bang-bang control scheme is a system that switches abruptly between two different modes and is an important part of optimal control theory [14]. Here, the two modes will be the application of the Hamiltonians E and $X^{\otimes n}$ respectively. In order to explore the space of protocols computationally, we break up the total time T into N_b blocks of time of T/N_b time each, with each block assigned to one or the other Hamiltonian. A bang-bang QAOA protocol then involves iterating through the blocks applying the corresponding Hamiltonian for T/N_b amount of time. This simply involves applying either e^{iET/N_b} or $e^{iX^{\otimes n}T/N_b}$, respectively. See Fig. 1 for an example.

If one were to translate a bang-bang QAOA protocol into the language of the standard QAOA, the p value of said protocol could be as large as $N_b/2$. However, the total amount of time is at most T . In addition, in the large N_b limit this bang-bang QAOA model can approximate any standard QAOA

protocol such that $\sum_{i=1}^p (\beta_i + \gamma_i) \approx T$. Later on we will also argue that it is not worthwhile to consider QAOA protocols for large T or $\sum_{i=1}^p (\beta_i + \gamma_i)$ in the bang-bang and standard QAOA, respectively, due to the difficulty of optimization.

III. MAX-2-SAT

For boolean expressions a conjunction is a logical AND and is typically represented as $\wedge$. A disjunction is a logical OR that is represented as $\vee$. Finally, $\neg$ as a unary operator represents logical negation. Given boolean values $x_0, \dots, x_{n-1}$, a k -CNF (Conjunctive Normal Form) is the conjunction over disjunctive clauses of size k . More intuitively, a k -CNF is an AND-of-ORs where each OR involves k boolean values. 2-SAT is then the problem of determining if there exists a assignment of $x_0, \dots, x_{n-1}$ such that a given 2-CNF is satisfied. The natural optimization version of the problem, MAX-2-SAT, is then the problem of determining the maximum number of clauses satisfiable by an assignment of $x_0, \dots, x_{n-1}$.

It is important to make the distinction between 2-SAT and MAX-2-SAT; 2-SAT is in P [15–17], while MAX-2-SAT is NP-Hard [18]. Hardness of approximation results have shown that no Polynomial-Time Approximation Schemes (PTAS) exist for MAX-2-SAT with approximation ratios better than $\frac{21}{22} \approx 0.955$ [19] assuming $P \neq NP$ and ~ 0.943 [20] when also assuming the unique games conjecture. Here the approximation ratio of an algorithm refers to a guarantee that an algorithm with approximation ratio r for a problem instance with optimal solution $C_{\max}$ achieves a result of at least $rC_{\max}$ (potentially only with high probability if randomized or quantum). There does, however, exist an efficient algorithm based on semidefinite programming that achieves an approximation ratio of 0.94 [21]. It is worth noting that a uniformly random assignment of literals will satisfy $\frac{3}{4}$ of the clauses in expectation. By the probabilistic method this also ensures that at least $\frac{3}{4}$ of the clauses are always satisfiable.

Finally, we will show how to encode MAX-2-SAT as a Hamiltonian E . Given a disjunctive clause of 2 literals, there is exactly one assignment that does not satisfy the clause. We then design a diagonal Hamiltonian for the clause that is 1 for every literal assignment where the clause is not satisfied. For instance, given the clause $\mathcal{C} = (x_i \vee \neg x_j)$ where $i < j$ the only assignment that does not satisfy $\mathcal{C}$ is $x_i = 0, x_j = 1$, where 1 is True and 0 is False. We can then define the Hamiltonian for $\mathcal{C} = (x_i \vee \neg x_j)$

$$E_{\mathcal{C}} = I^{\otimes n} - I^{\otimes i-1} \otimes |0\rangle\langle 0| \otimes I^{\otimes j-i-1} \otimes |1\rangle\langle 1| \otimes I^{n-j}. \quad (1)$$

The diagonal of this Hamiltonian is then 1 for every computational basis state x such that $x_i \neq 0$ or $x_j \neq 1$. The Hamiltonian of MAX-2-SAT is then the sum over all Hamiltonians induced by the clauses $\{\mathcal{C}\}$ in the 2-CNF

$$E = \sum_{\mathcal{C}} E_{\mathcal{C}}. \quad (2)$$

One can see that the diagonal encodes the number of clauses satisfied by the assignment of literals. The objective function of QAOA is the expected number of satisfied clauses, $\langle \psi | E | \psi \rangle$. If $C_{\max}$ is the maximum number of satisfiable clauses, by linearity of expectation this leads to the expected

¹The QAOA parameters are not always restricted to be positive, but the β parameters are naturally periodic, and so too are the γ parameters when (as will be the case here) the objective function takes on integer values.

approximation ratio

$$f_{\text{obj}} = \frac{\langle \psi | E | \psi \rangle}{C_{\text{max}}}. \quad (3)$$

While we cannot normally directly compute the approximation ratio without knowing C_{max} , maximizing $\langle \psi | E | \psi \rangle$ will also maximize f_{obj} due to being related by a constant factor.

IV. METHODS

In this section, we will simply outline preliminary information to understand our results. We will henceforth set the number of variables in our MAX-2-SAT instances to be 10. This is due to the exponential nature of the dimension of the Hilbert space with respect to the variables (i.e., qubits). This means that adding a single variable will double the amount of computation needed.

A. Stochastic descent

QAOA optimizes the quantum circuit in order to increase the probabilities of getting a good measurement. Given a bang-bang QAOA protocol P that produces state $|\psi_P\rangle$, our objective function $f_{\text{obj}}(P)$ will be the expected approximation ratio of the resulting state in Eq. (3). In the bang-bang QAOA our protocols fall into a discrete space. As such, we use the following greedy randomized optimization approach introduced by Day *et al.* [6] with $k = 1$.

Algorithm 1. Stochastic descent (SD_k).

```

1: Input:  $N_b, T, k$ 
2: Routines: RandomProtocol, FindAllUpdate,
   RandomShuffle, UpdateProtocol,  $f_{\text{obj}}$ 
3: initialize:
   //Initialize protocol at random
4:  $P_{\text{old}} \leftarrow \text{RandomProtocol}(N_b)$ 
   //Finds the list of updates with at most  $k$ -flips
5:  $\text{ListOfAllUpdates} \leftarrow \text{FindAllUpdate}(N_b, k)$ 
6: shuffle:
   //Shuffle updates in a random order
7:  $\text{ListOfAllUpdates} \leftarrow \text{RandomShuffle}(\text{ListOfAllUpdates})$ 
   //Iterate over all possible update
8: for  $\text{update}$  in  $\text{ListOfAllUpdates}$  do
   //Update protocol given the specified update
9:  $P_{\text{new}} \leftarrow \text{UpdateProtocol}(P_{\text{old}}, \text{update})$ 
   //Evaluates the objective function of each protocol
10:   and compares them
11:   if  $f(P_{\text{new}}) > f(P_{\text{old}})$  then
12:      $P_{\text{old}} \leftarrow P_{\text{new}}$ 
   //If update accepted, then restart for loop
13:   goto shuffle
return  $P_{\text{old}}$ 
    
```

If bang-bang protocols are viewed as bit strings, the algorithm randomly iterates through all protocols of Hamming distance at most k away from the current one and updates itself to the first protocol it finds that performs better. If no protocol is better, then we say that the current protocol is a k -local optimum. Note that the number of protocols that need to be considered at each update grows with k as $\sum_{i=1}^k \binom{N_b}{i}$, which is

near exponentially for $k \leq N_b/2$. Thus, increasing k quickly becomes very computationally expensive.

B. Random protocol initialization

An interesting aspect of stochastic descent (SD) is the distribution with which the initial random protocol is drawn from. The original algorithm proposed by Day *et al.* [6] uniformly samples at random. In this paper, we propose two new initialization to study the relation between bang-bang QAOA with adiabatic quantum computation.

We define $\xi(q)$ to be the Bernoulli random variable with probability q of being 1 and $1 - q$ of being -1 . Varying q as a function over blocks generates three different random initialization methods:

Adiabatic: $\prod_{i=1}^{N_b} \xi(i/N_b)$. It favors $X^{\otimes n}$ in early blocks and E in late blocks of the protocols.

Uniform: $\prod_{i=1}^{N_b} \xi(0.5)$. The probabilities of $X^{\otimes n}$ and E being sampled are equal. This is the default distribution used in this paper.

Antiadiabatic: $\prod_{i=1}^{N_b} \xi(1 - i/N_b)$. It favors E in early blocks and $X^{\otimes n}$ in late blocks of the protocols.

C. Correlator

Given a set S of protocols with N_b blocks, we define the correlator of S as the following: View a protocol as a collection of values $P \in \{-1, 1\}^{N_b}$, where P_i refers to the value at block i . Let $\bar{P}_i = \frac{1}{|S|} \sum_{P \in S} P_i$ represent the empirical average of block i over all protocols in the set S . The correlator is defined as a certain variance of the protocol values,

$$\sigma = \frac{1}{N_b |S|} \sum_{i=1}^{N_b} \sum_{P \in S} (P_i - \bar{P}_i)^2 = 1 - \frac{1}{N_b} \sum_{i=1}^{N_b} \bar{P}_i^2. \quad (4)$$

A small correlator means the protocols in S are similar to each other. Intuitively, it is really an “anticorrelator” as the value is small when the protocols are similar. We choose to keep the same name as Day *et al.* [6] for consistency.

D. Protocol smoothing

It is also important to analyze the actual structure of bang-bang QAOA protocols after stochastic descent. While one could plot the protocols themselves as $\{-1, 1\}$ values along a time-scale, this does little to see the effects of how a protocol may favor one Hamiltonian over the other at different points in time. As such, we also opt to smooth the protocols by taking a rolling average. In addition, this smoothing allows us to properly see how close many of these bang-bang QAOA protocols are to being standard QAOA protocols for small total time by smoothing over minor deviations.

More formally, let w be a positive integer known as the *window size*. Similar to the correlator, we will view bang-bang QAOA protocols as $P \in \{-1, 1\}^{N_b}$ where P_i refers to the value at block i . We then defined the *smoothed protocol*

$$P' \in [-1, 1]^{N_b - w + 1},$$

FIG. 2. Top panel represents aggregate statistics on the expected approximation ratio with respect to total time T . Percentiles are given respectively. 10 000 protocols are sampled per time-step with $N_b = 200$ blocks. Fill grey region represents protocols without SD_1 applied. Shaded regions for each color denote the percentile regions as indicated in the legend. The initial protocols are sampled uniformly at random. The bottom panel contains the corresponding correlator. The vertical lines and thumbnails indicate 10 randomly sampled protocols found by SD_1 at corresponding total times. More details of the profiles of protocols are in Figs. 3 and 4. In the text we discuss the features of this figure, such as the rapid increase in median expected approximation ratio at specific times and the increase in the correlator over time.

where

$$P'_i = \frac{1}{w} \sum_{j=0}^{w-1} P_{i+j}.$$

E. Problem instances

The 2-DNFs used were constructed by randomly generating a clause with two unique indices drawn uniformly, as well as whether or not to negate each variable. Several of these clauses are then independently created, with the number of random clauses n_c being a parameter specified at runtime. Note that it is possible that two identical clauses are created, and by the Birthday paradox we expect this to happen when $n_c \gtrsim n$. While this is the regime that we end up creating our problem instances with, one can simply repeat the process an expected constant number of times until success. We ensure that there are no identical clauses in our problem instances. See Appendix for the actual problem instances used of 10, 20, and 30 clauses, respectively. For clarity, we focus on the 10 clause problem instance in the proceeding results section.

V. RESULTS

In Fig. 2, we can see how the protocols drawn uniformly at random perform without being optimized with SD (in grey) as compared to SD_1 (in green), with a substantial increase in expected approximation ratio. Even with SD_1 it is easy to see the benefit of a greedy optimization strategy for bang-bang QAOA. Interestingly, without SD_1 , protocols perform worse

than the naive classical algorithm of a uniformly random assignment of variables, which achieves at least a $3/4$ approximation ratio.

A. Small time regime

We will refer to the small time regime as $T \lesssim 6$, though this value is likely problem instance specific. The important aspect of these figures in this regime is the rapid increase in median expected approximation ratio around $T \approx 1.5$ and ≈ 3.5 , which we refer to as a *phase transition* in performance. We attribute this to be the minimal total time needed for protocols to start enacting nontrivial behavior, corresponding to SD_1 converging on a $p \approx 2$ and ≈ 3 protocol, respectively, when viewed as a standard QAOA protocol.

We can see from Figs. 3(a) and 3(b) that at very small time the protocols only apply each Hamiltonian once. Then as the total time increases, the protocols then transition into two switchbacks as seen in Figs. 3(d) and 3(e) leading to the median expected approximation ratio to increase substantially. This again repeats with $p \approx 3$ like in Fig. 3(f), however the increase in median expected approximation ratio is not as great as before.

Looking at the correlator in Fig. 2, at $T = 0$ it starts off around 1 since every starting protocol is a local optimum with a uniform probability of being selected. As SD_1 begins to optimize towards specific protocols, the value then quickly drops as there are only a few and very similar local optima. It is then when transitioning to a new local optima that the correlator begins to temporarily spike, as there is a mixture of

FIG. 3. Illustrates how bang-bang QAOA protocols converge towards standard QAOA protocols with small p values for $N_b = 200$ and 10 clauses with uniform random initialization. Drawing 10 protocols uniformly at random with a window size $w = 1$. (a) $T = 0.5$ and (b) 1.0 are before the first transition point and are similar to standard QAOA with $p = 1$. (c) $T = 1.5$ is during the transition, where the rapid increase in the correlator is due to the mixture of two kinds of protocols. (d) $T = 2.2$ is after the transition where protocols resemble $p = 2$. Likewise, (e) $T = 3.0$ is before the second transition, (f) $T = 3.5$ is during, and (g) $T = 4.2$ is after. (h) $T = 6.0$ is at the point where the median expected approximation ratio begins to plateau. Color map is based on the expected approximation ratios of protocols and the numbers on the right indicate their corresponding values. See Fig. 2 to see the transition points.

protocols as with Fig. 3(f). Once the transition has finished, the correlator then quickly decreases again. However, there is a general trend towards protocols becoming uncorrelated as the number of local optima start increasing with total time.

B. Large time regime

It is then at large time that increasing total time no longer becomes as beneficial for the global optima and the number of local optima starts to increase rapidly. Here, despite the fact

FIG. 4. Demonstrates the profiles of local optima at large time ($T = 9.5$) with [(a) and (b)]adiabatic, [(c) and (d)]uniform, and [(e) and (f)]antiadiabatic initialization for $N_b = 200$ and 10 clauses. Protocols are sampled randomly and smoothed with a window size $w = 1$ and 50 , respectively. Within each subfigure, there is no qualitative difference in the profiles between the best and worst protocols and no strong underlying standard QAOA-like structure behind the protocols like with Fig. 3. Additionally, the shape of the smoothed protocols being similar to the expected initial protocol even after SD_1 indicates that local optima can be found with qualitatively different protocols. Colormap is based on the expected approximation ratios of protocols and the numbers on the right indicate their corresponding values. Note that the same 10 protocols are shown with different smoothing for each initialization. Initialization explained in Sec. IV B.

that the best protocols continue to do marginally better at large T , looking carefully at Fig. 2 the median trends downward. We believe this is due to the local optima no longer being close to the global as it becomes more and more difficult to find better optima using SD_1 . This tells us that within the realms of greedily optimized bang-bang QAOA, there is more than enough time necessary for a near-optimal protocol and any extra total time contributes to extra degrees of freedom that make optimization more difficult. This becomes even more apparent as the number of clauses increases and the median protocol begins to quickly fall off as total time increases.

Looking at the protocols in Fig. 4(c) without smoothing, there is no apparent discernible profile with the protocols in how they relate to their expected approximation ratio. However, we do see in Fig. 4(d) that the protocols tend to favor the constraint Hamiltonian and appear neither adiabatic nor anti-adiabatic. Together with Figs. 4(b) and 4(f), we find the trend of protocols remain qualitatively similar to their initialization.

FIG. 5. (a) 100 vs (b) 200 blocks are compared side-by-side for the same problem instance containing 10 clauses. With 200 blocks one sees slightly better median expected approximation ratios.

C. Summary

Below a certain total time T , no bang-bang QAOA protocol does well since the resulting unitary of the circuit will still be close to identity. After a certain point, a select few protocols start exhibiting nontrivial behavior, which is then found by SD_1 and the protocols transition to a much better expectation. The protocols do not initially benefit from further increases in T , which causes performance to plateau. Then at some point, T becomes large enough to allow for another set of nontrivial behavior, and this process continues until the transition to large time. At this point T becomes too large and it becomes too difficult to find a solution near the global optima. The protocols then start exhibiting less structure and become very different from each other quantitatively based on the correlator, but qualitatively do not deviate far from their initialization.

D. 100 versus 200 blocks

In Fig. 5, we illustrate how the number of blocks N_b affects the expected approximation ratio of bang-bang QAOA. While the shape of both graphs are very similar, one can see that between Figs. 5(a) and 5(b) $N_b = 200$ tends to give better expected approximation ratios, especially when the total time T becomes large. Additionally, the correlator dips lower around the phase transitions, indicating that the protocols actually

FIG. 6. Expected approximation ratio and the correlator when run on the (a) 10, (b) 20, and (c) 30 clauses problem instances. 10 000 protocols are sampled per time-step with $N_b = 200$ blocks. Grey plots represent protocols without SD_1 applied. The initial protocols are sampled uniformly at random. Because the ratio of satisfiable clauses tends to decrease as the total number of clauses increases, the baseline of random guessing is able to perform better as it is always able to satisfy $3/4$ of all clauses.

FIG. 7. Expected approximation ratio and the correlator when using (a) adiabatic, (b) uniform, and (c) antiadiabatic initialization. The probability distribution that the initial protocol is drawn from does have minor small T and more pronounced large T effects on how the protocols perform, even after SD_1 . Here 10 clauses are used with $N_b = 200$.

concentrate better with larger N_b around the phase transitions, despite the fact that there are exponentially more protocols available as N_b increases.

E. Larger number of clauses

It is of course important to analyze more than a single problem instance. In Fig. 6, we find that the overall behavior remains relatively consistent between the problem instances with 10, 20, and 30 clauses, respectively. More specifically, we see that the median expected approximation ratio increases in jumps in the small time regime, before trailing off at large time. This decay in median expected approximation ratio is especially pronounced in Fig. 6(c). Similarly, the correlator moves up and down in the small time regime, though increasing to a value of nearly 1 as time increases.

F. Effects of random initialization

Looking at Fig. 7, they all perform similarly at small T . This is to be expected as there are few local optima such that the initialization only effects the starting distance to the global. However, at large T , Fig. 8 shows us that the uniform random initialization tends to do poorly with respect to the median protocol. The adiabatic initialization however tends to consistently do well at large T , while the anti-adiabatic initialization exhibits large variance in its median expected approximation ratio over time (see Sec. IV B for the definition

FIG. 8. Comparison of median expected approximation ratio of protocols based on initialization. (a) 10, (b) 20, and (c) 30 clause problem instances are presented. Adiabatic initialization tends to consistently do well even at large total time, as opposed to uniform random which tends to drop off. Antiadiabatic initialization leads to high variance in the median expected approximation ratio.

FIG. 9. Top ten bang-bang QAOA protocols using (a) adiabatic, (b) uniform, and (c) anti-adiabatic initialization at large time ($T = 9.5$), smoothed with window size $w = 50$ for $N_b = 200$ with 10 clauses. Color map is based on the expected approximation ratios of protocols and the numbers on the right indicate their corresponding values. Since 1 corresponds to the objective Hamiltonian, an adiabatic protocol will gradually increase in value.

of adiabatic and antiadiabatic initialization). As T becomes large, the adiabatic theorem, the driving force behind adiabatic quantum computation [3], starts becoming relevant. If intuition from the adiabatic theorem and adiabatic quantum computation extends to the local optima found by SD_1 , the local optima found around the adiabatic initialization are then likely to perform better on average than those initialized from uniform or anti-adiabatic distributions.

Further evidence of this can be seen by once again examining the protocols themselves. Looking at Fig. 4, we see that the randomly drawn protocols from each initialization are very similar to their expected starting protocol. Finally, if we instead look at the best protocols of each initialization, Fig. 9 shows the best protocols appear to be qualitatively more adiabatic: adiabatic initialization leads to strongly adiabatic algorithms, anti-adiabatic is largely uniform, and random initialization slightly favors adiabaticity. Thus even though the local optima themselves seem to be unbiased, the best protocols seem to be found in the space around qualitatively more adiabatic protocols than the initialization.

G. Iterations plots

As a final demonstration of the difficulty of finding the globally optimal protocol as the number of local optima increases with T , we examine the average number of iterations SD_1 needs to find a local optima. Looking at Fig. 10, we can see that the number of iterations needed increases at the first phase transitions, before decaying as T increases. This is true for all three initialization. What this effectively means is that the distance from a random protocol to its nearby local optima decreases with T regardless of the three starting positions.

VI. CONCLUSION

Ultimately, it is not clear the bang-bang QAOA should be used in practice with NISQ devices. As stated in Sec. II B, the depth of the circuit can potentially be as large as $N_b/2$, which is exactly the reason $p = O(1)$ is used in standard QAOA to avoid this. However, as a thought experiment as to the value of the p parameter itself, this serves as further

FIG. 10. Shows the average number of iterations of SD_1 are needed before a local optima are found on the (a) 10, (b) 20, and (c) 30 clause problem instances. As the number of local optima increase with total time, it becomes easier to find one such that the number of iterations quickly decreases.

evidence that larger p values are not necessary to achieve the best approximation ratios when the optimization process is limited to bounded computation. We see that while a minimal amount of time is needed for bang-bang QAOA protocols to achieve nontrivial approximation ratios, they fail to substantially improve in the median expected approximation ratio for larger T . It is also not clear from the data alone how good of an approximation ratio one can get using bang-bang QAOA efficiently.

Due to the nature of classical simulation of quantum mechanics, collecting data is incredibly time intensive even with parallelization of sample collection. For example, because SD_k takes time exponential in k for small k , we were restricted to SD_1 . Additionally, the number of variables was only set to 10, creating very small 2-SAT instances. It will be interesting to see if these behaviors remain the same even with larger problem instances and/or using SD_k for $k > 1$. Additionally, though we examine $N_b = 100$ and 200 in Fig. 5 there is not currently enough data to draw strong conclusions between the relationship between N_b and performance.

Another consideration is that various other modifications to QAOA such as Li *et al.* [11], which modifies the objective function, can be combined with bang-bang QAOA. One compelling modification could involve the ability to apply a Hamiltonian for negative time, corresponding to negative $\{\beta_i\}$ and $\{\gamma_i\}$ parameters in standard QAOA such that total time becomes $T = \sum_i |\beta_i| + |\gamma_i|$ [22]. Changes to SD are necessary, such as redefining the distance metric between protocols beyond Hamming distance, as well as preventing the cancellation of Hamiltonians. How these modifications work in tandem with bang-bang QAOA may lead to interesting phenomena that could potentially lead to a more practical algorithm.

Open source code of this paper is available in github repository [23].

ACKNOWLEDGMENTS

The authors thank Zan Armstrong for suggestions in data visualization, and Murphy Yuezhen Niu, John Platt for their review and comments. Daniel Liang would like to thank the Simons It from Qubit Collaboration and Scott Aaronson for supporting him. X, formerly known as Google[x], is part of the Alphabet family of companies, which includes Google, Verily, Waymo, and others [24]. Quantum simulation and SD in this paper were implemented using Cirq [25] and Apache Beam [26].

APPENDIX: PROBLEM INSTANCES USED IN THIS PAPER

Notation: $\vee$ represents a disjunction, $\wedge$ represents a conjunction, and $\neg$ represents logical negation.

1. 10 clauses

$$f(x) = (\neg x_8 \vee x_9) \wedge (\neg x_5 \vee x_7) \wedge (x_0 \vee \neg x_6) \wedge (\neg x_4 \vee x_5) \wedge (x_4 \vee \neg x_5) \wedge \\ (\neg x_0 \vee x_2) \wedge (x_0 \vee \neg x_4) \wedge (\neg x_0 \vee x_7) \wedge (\neg x_4 \vee \neg x_7) \wedge (x_7 \vee x_8)$$

2. 20 clauses

$$f(x) = (\neg x_6 \vee \neg x_9) \wedge (\neg x_1 \vee \neg x_4) \wedge (x_0 \vee x_6) \wedge (\neg x_6 \vee \neg x_7) \wedge (x_2 \vee \neg x_6) \wedge \\ (x_3 \vee x_8) \wedge (\neg x_1 \vee x_4) \wedge (\neg x_0 \vee x_6) \wedge (x_5 \vee \neg x_9) \wedge (x_0 \vee \neg x_8) \wedge \\ (x_1 \vee \neg x_8) \wedge (x_1 \vee x_8) \wedge (x_5 \vee \neg x_7) \wedge (x_2 \vee \neg x_7) \wedge (\neg x_0 \vee \neg x_5) \wedge \\ (x_6 \vee \neg x_9) \wedge (\neg x_0 \vee \neg x_7) \wedge (x_0 \vee x_3) \wedge (\neg x_1 \vee x_6) \wedge (\neg x_0 \vee x_3)$$

3. 30 clauses

$$f(x) = (\neg x_3 \vee x_7) \wedge (\neg x_4 \vee x_8) \wedge (\neg x_3 \vee \neg x_9) \wedge (\neg x_6 \vee \neg x_7) \wedge (x_1 \vee x_5) \wedge \\ (\neg x_0 \vee x_6) \wedge (\neg x_3 \vee \neg x_4) \wedge (x_0 \vee x_8) \wedge (x_5 \vee x_7) \wedge (x_3 \vee \neg x_8) \wedge \\ (x_1 \vee x_8) \wedge (\neg x_0 \vee \neg x_6) \wedge (x_1 \vee x_2) \wedge (x_0 \vee \neg x_1) \wedge (\neg x_5 \vee x_9) \wedge \\ (x_4 \vee \neg x_6) \wedge (\neg x_2 \vee \neg x_8) \wedge (x_8 \vee \neg x_9) \wedge (x_7 \vee \neg x_9) \wedge (x_1 \vee \neg x_4) \wedge \\ (x_6 \vee \neg x_9) \wedge (x_3 \vee x_4) \wedge (\neg x_5 \vee x_6) \wedge (x_1 \vee \neg x_9) \wedge (x_1 \vee \neg x_3) \wedge \\ (x_2 \vee \neg x_5) \wedge (\neg x_0 \vee x_7) \wedge (x_0 \vee x_2) \wedge (\neg x_0 \vee \neg x_1) \wedge (\neg x_7 \vee x_9)$$

-
- [1] E. Farhi, J. Goldstone, and S. Gutmann, A quantum approximate optimization algorithm, [arXiv:1411.4028](#).
- [2] F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, R. Barends, S. Boixo, M. Broughton, B. B. Buckley, D. A. Buell *et al.*, Quantum approximate optimization of non-planar graph problems on a planar superconducting processor, [arXiv:2004.04197](#).
- [3] E. Farhi, J. Goldstone, S. Gutmann, and M. Sipser, Quantum computation by adiabatic evolution, [arXiv:quant-ph/0001106](#).
- [4] T. Kadowaki and H. Nishimori, Quantum annealing in the transverse ising model, *Phys. Rev. E* **58**, 5355 (1998).
- [5] R. Shaydulin and Y. Alexeev, Evaluating quantum approximate optimization algorithm: A case study, in *Proceedings of Tenth International Green and Sustainable Computing Conference (IGSC)* (IEEE, Piscataway, NJ, 2019).
- [6] A. G. R. Day, M. Bukov, P. Weinberg, P. Mehta, and D. Sels, Glassy Phase of Optimal Quantum Control, *Phys. Rev. Lett.* **122**, 020601 (2019).
- [7] L. T. Brady, C. L. Baldwin, A. Bapat, Y. Kharkov, and A. V. Gorshkov, Optimal protocols in quantum annealing and qaoa problems, [arXiv:2003.08952](#).
- [8] Z.-C. Yang, A. Rahmani, A. Shabani, H. Neven, and C. Chamon, Optimizing Variational Quantum Algorithms Using Pontryagin's Minimum Principle, *Phys. Rev. X* **7**, 021027 (2017).
- [9] A. Bapat and S. Jordan, Bang-bang control as a design principle for classical and quantum optimization algorithms, *Quantum Info. Comput.* **19**, 424 (2019).
- [10] G. B. Mbeng, R. Fazio, and G. Santoro, Quantum annealing: A journey through digitalization, control, and hybrid quantum variational schemes, [arXiv:1906.08948](#).
- [11] L. Li, M. Fan, M. Coram, P. Riley, and S. Leichenauer, Quantum optimization with a novel gibbs objective function and ansatz architecture search, *Phys. Rev. Research* **2**, 023074 (2020).
- [12] S. Hadfield, Z. Wang, B. O'Gorman, E. Rieffel, D. Venturelli, and R. Biswas, From the quantum approximate optimization algorithm to a quantum alternating operator ansatz, *Algorithms* **12**, 34 (2019).
- [13] J. Cook, On the relationships between Z-, C-, and H-local unitaries, [arXiv:1907.11368](#).

- [14] L. S. Pontryagin, *Mathematical Theory of Optimal Processes* (CRC Press, Boca Raton, FL, 1987).
- [15] M. R. Krom, The decision problem for a class of first-order formulas in which all disjunctions are binary, *Math. Logic Quarterly* **13**, 15 (1967).
- [16] B. Aspvall, M. F. Plass, and R. E. Tarjan, A linear-time algorithm for testing the truth of certain quantified boolean formulas, *Information Processing Lett.* **8**, 121 (1979).
- [17] S. Even, A. Itai, and A. Shamir, On the complexity of timetable and multicommodity flow problems, *SIAM J. Comput.* **5**, 691 (1976).
- [18] M. Garey, D. Johnson, and L. Stockmeyer, Some simplified np-complete graph problems, *Theor. Comput. Sci.* **1**, 237 (1976).
- [19] J. Håstad, Some optimal inapproximability results, *J. ACM* **48**, 798 (2001).
- [20] S. Khot, G. Kindler, E. Mossel, and R. O'Donnell, Optimal inapproximability results for max-cut and other 2-variable CSPs? *SIAM J. Comput.* **37**, 319 (2007).
- [21] M. Lewin, D. Livnat, and U. Zwick, Improved rounding techniques for the MAX 2-SAT and MAX DI-CUT problems, in *Integer Programming and Combinatorial Optimization, IPCO 2002*, Lecture Notes in Computer Science, Vol. 2337, edited by W. J. Cook and A. S. Schulz (Springer, Berlin, Heidelberg, 2002).
- [22] E. Farhi, J. Goldstone, and S. Gutmann, A quantum approximate optimization algorithm applied to a bounded occurrence constraint problem, [arXiv:1412.6062](https://arxiv.org/abs/1412.6062).
- [23] D. Liang, L. Li, and S. Leichenauer, "Investigating quantum approximate optimization algorithms under bang-bang protocols", github.com/google-research/google-research/tree/master/bangbang_qaoa (2020).
- [24] www.x.company.
- [25] Cirq: A python framework for creating, editing, and invoking noisy intermediate scale quantum (nisq) circuits, <https://github.com/quantumlib/Cirq>.
- [26] Apache beam, <https://beam.apache.org/>.