



**Universiteit
Leiden**
The Netherlands

Algorithm selection and configuration for Noisy Intermediate Scale Quantum methods for industrial applications

Moussa, C.

Citation

Moussa, C. (2023, October 11). *Algorithm selection and configuration for Noisy Intermediate Scale Quantum methods for industrial applications*. Retrieved from <https://hdl.handle.net/1887/3643423>

Version: Publisher's Version

License: [Licence agreement concerning inclusion of doctoral thesis in the Institutional Repository of the University of Leiden](#)

Downloaded from: <https://hdl.handle.net/1887/3643423>

Note: To cite this publication please use the final published version (if applicable).

Chapter 5

Tabu-driven quantum neighborhood samplers

Combinatorial optimization is an important application targeted by quantum computing. However, near-term hardware constraints make quantum algorithms unlikely to be competitive when compared to high-performing classical heuristics on large practical problems. The small size of instances used in previous chapters 3 and 4 illustrates such an issue. One option to achieve advantages with near-term devices is to use them in combination with classical heuristics. In particular, we propose using quantum methods to sample from classically intractable distributions – which is the most probable approach to attain a true provable quantum separation in the near term – which are used to solve optimization problems faster. In this chapter, we numerically study this enhancement by an adaptation of Tabu Search using QAOA as a neighborhood sampler¹. We show that QAOA provides a flexible tool for exploration-exploitation in such hybrid settings and can provide evidence that it can help in solving problems faster by saving many tabu iterations and achieving better solutions.

¹Contents of this chapter are based on [139]; Charles Moussa, Hao Wang, Henri Calandra, Thomas Bäck, and Vedran Dunjko. Tabu-driven quantum neighborhood samplers. In Christine Zarges and Sébastien Verel, editors, *Evolutionary Computation in Combinatorial Optimization*, pages 100–119, Cham, 2021. Springer International Publishing.

5.1 Introduction

While numerous works have been studying various theoretical and empirical properties of QAOA [26, 46, 207, 186, 135], many practical challenges remain. Indeed, only small-sized problems and very limited p can be run on real hardware, which severely limits the quality of the solution obtained empirically [79, 14, 114]. Many open questions still remain, e.g., regarding the comparison of QAOA with other heuristic methods on various cases of instances that stem from particular problem domains, with different optimizers, and with varying levels of experimental (or simulated) noise. One reason of why so many uncertainties remain is that classical simulation is computationally very expensive, and quantum devices are still scarce to prevent large real-world tests [127, 79].

In contrast to optimization problems, quantum advantage has been demonstrated in sampling [11]. Indeed, theoretical results establish a quantum advantage in producing samples according to certain distributions of constant-depth quantum circuits [27]. In this direction, it has been demonstrated that the sampling of the QAOA circuit, even at $p = 1$, cannot be efficiently simulated classically [58]. The above considerations point to a possibility of utilizing sampling features of QAOA for neighborhood explorations with the added benefit that, since the neighborhood may be limited to fewer variables, a smaller quantum device may already lead to improved performance of a large instance.

It is interesting to delve into sampling aspects in the domain of classical local search algorithms, where we seek the optimum in the vicinity of the current solution with respect to either the original optimization problem or a subproblem thereof, using a deterministic or stochastic sampling strategy [12]. Such a sampling-based local procedure is typically realized by the combination of some parametric distribution family for drawing local trial points (e.g., the binomial or a power-law distribution [48]) and a selection method for choosing good trial points, and hence the overall outcome of this procedure results in the family of sampling distributions [47, 199, 112].

In this chapter, we use QAOA circuits as local neighborhood samplers, having malleable support in (many) good local optima but still allowing a level of exploration (which is desirable since local optima may not lead to global optima). This introduces the topics of sampling and multiobjective aspects of QAOA that allow balancing between exploration and exploitation. To this end, we study its combination with tabu search (TS), a metaheuristic that has been successfully applied in practice for combinatorial optimization by local search. Moreover, to control the trade-off between

exploration and exploitation, we add this critical component in TS to the specification of the standard QAOA circuit.

Contributions - We construct an algorithm incorporating QAOA in TS with the usual attribute-based short-term memory structure (a.k.a the tabu list). With our approach, we kill two birds with one stone: we gain quantum enhancements, while the local properties of tabu search can make the required quantum computations naturally economic in terms of needed qubit numbers, which is vital in the near-term quantum era. We analyze and benchmark this incorporation with *small QAOA depths* against a classical TS procedure on QUBO problems of up to 500 variables. We also propose a penalized version of QAOA incorporating knowledge from a current solution. We find that QAOA is often beneficial in terms of saved iterations, and can find shorter paths toward better solutions. The structure of the chapter is as follows. Section 5.2 provides the necessary background on TS and interplay with quantum techniques. In section 5.3, we detail the TS procedure incorporating a short-term memory structure with QAOA. The results of our simulations are presented in section 5.4. We conclude with a discussion in section 5.5.

5.2 Background

Tabu Search (TS) [66] is a meta-heuristic that guides a local heuristic search procedure to explore the search space beyond local optimality. One of the main components of TS is its use of adaptive memory, which creates a more flexible search behavior. Such a framework allows using a quantum algorithm as a local search tool, for solving large instances with limited-sized quantum devices. Various works leveraged TS for solving QUBOs [105, 65, 147, 67, 148] using short-term and long-term strategies used during the search. We note also different hybrid settings that combine a basic TS procedure with another framework such as genetic search [120] and Path Relinking [198]. TS was also incorporated with quantum computers to tackle larger problems beyond their limitations. Indeed, finding methods to leverage smaller devices is of main importance. Many divide-and-conquer approaches have been designed for quantum circuits and algorithms [50, 159, 30, 150]. In this paper, the size of the QC comes into play more naturally as a hyperparameter defining the «radius» of the search space.

With respect to the interplay between TS and quantum techniques, to our knowledge TS has only been considered from the perspective of D-Wave quantum annealers, a different non-universal model of computation than the gate-based model. The first approach of this kind is an algorithm called qbsolv [25]. It starts with an initial TS

5.3. Tabu-driven QAOA sampling

run on the whole QUBO. Then the problem is partitioned into several subproblems solved independently with the annealer. Subproblems are created randomly, by selecting variables. Non-selected ones have their values fixed (clamping values) from the TS solution. The subsolutions are then merged and a new TS is run as an improvement method. The second approach is an iterative solver designed in [162]. At each iteration, a subproblem is submitted to the annealer. The subproblem is obtained by clamping values from a current solution. A tabu list is used in which each element is a list of variables of length k . Each element is kept tabu for a user-defined number of iterations. In contrast in this work, we consider using QAOA in combination with TS.

5.3 Tabu-driven QAOA sampling

Inspired by the above-mentioned works, we use a simple TS procedure where QAOA is added in the neighborhood generation phase to solve QUBO problems. Note that we could also apply QAOA in more sophisticated frameworks, but a simpler approach is easier for understanding the benefits of QAOA with TS.

Local search algorithms explore a search space by generating sequences of possible solutions which are refined. At each step, we generate a so-called neighborhood from a current solution. In particular, if we denote the current solution x , a generated neighborhood corresponds to candidates x' that differ by at most k bits. We denote this set as $N_k(x) = \{x' \in \{0,1\}^n | \delta_H(x', x) \leq k\}$, where δ_H denotes the Hamming distance. For a simple one-bit-flip generation strategy, this corresponds to $k = 1$ and TS uses a modified neighborhood due to tabu conditions. Although increasing k could help exploration, the neighborhood generation comes at an exponential cost. But this could mean finding better solutions in fewer TS iterations, and thus also in principle overall faster if a fast good method for neighborhood exploration is devised.

This motivates the use of a quantum algorithm as a proxy for exploring $N_k(x)$. Specifically, we will use QAOA which $2p$ real parameters are tweaked in a continuous optimization scheme resulting in a probability distribution on N_k . Increasing p (assuming the optima are found over the parameters) will improve the quality of the output (likelihood of returning an actual global optimum).

However, in the case of local search, a greedy strategy that tends to select the best point in the neighborhood would not only lead to potential stagnation but also result in longer optimization time (unless the neighborhoods are already the size of the overall problem) [20]. Indeed, one may also consider modifications that impose (various) notions of locality, which are usually not considered in standard QAOA

uses where it is used for the entire instance, with the sole goal of finding optima. To this end, we first outline the basic TS procedure generally used to solve QUBO problems [67, 120, 198]. Then, we show how QAOA can be combined with the latter. Finally, we propose a modification of QAOA that balances between going for the global optimum and prioritizing local improvements relative to the current TS solution.

5.3.1 The basic TS algorithm

The basic TS procedure for solving a QUBO with objective function $f(x)$ is described in Alg. 2, for $k = 1$ excluding the green-highlighted part. It uses a simple *tabu list* recording the number of iterations a variable remains tabu during the search. A variable can be set tabu for a fixed number of iterations (denoted Tabu tenure TT) but also with a random tenure. Each iteration can be considered as updating a current solution denoted x , exploring a modified neighborhood N'_1 due to the tabu considerations. Generally, x is chosen greedily when evaluating the objective function over candidates $x \in N'_1$.

For large problem instances, there exists an efficient evaluation technique for QUBO solvers leveraging one-bit flip move [64]. Let $\Delta_x = f(x') - f(x)$ be a move value, that is the effect in the objective of going to x' from x . For one-bit flip moves, we denote as $\Delta_x(i)$ the move value upon flipping the i -th variable, which can be computed using only the QUBO coefficients. The procedure records a data structure storing those move values, which is updated after each TS iteration.

Initially, all variables can be flipped (line 5). At each iteration, the tabu solution x is updated by flipping the variable that minimizes the objective over the neighborhood obtained by one-bit flip moves over non-tabu variables (lines 6-8). If the new tabu solution improves over the best-recorded solution, the aspiration is activated. In this case, the tabu attribute of the flipped variable is removed. The tabu list is finally updated (lines 18-23) and iterations continue until the stopping criterion is reached. This can be either a maximum number of TS iterations and/or a maximum number of TS iterations allowed without improvement of the best solution (*improvement cutoff*).

5.3.2 QAOA neighborhood sampling

In the usual TS algorithms, the neighborhood consists of candidates with Hamming distance one relative to the current tabu solution x . We note that sometimes considering also neighbors that are at most k -Hamming distance away from x helps in finding better solutions. The number k can be set in our case as large as the (limited) number

5.3. Tabu-driven QAOA sampling

of available qubits in quantum hardware.

To study the exploration of such neighborhoods, a brute-force generation approach is initially tested, and thereafter replaced by QAOA. As stated before, getting the optimum for subproblems in TS may lead to getting stuck during the search. QAOA, by definition, is a flexible framework as an exploration-exploitation tool. On the one hand, QAOA generates better solutions the deeper the circuit (p), and the better the classical optimization procedure within QAOA is. It is known it can have advantages over various standard algorithms, e.g. Simulated and Quantum Annealing [186]. To extract all advantages from the capacities of QAOA, we can further modulate the distribution of outputs it produces by limited depth or, as we present next, modifications to the QAOA objective to prioritize a more local behavior. Such flexibility is important, not only for the exploration-exploitation trade-off as it provides interesting ways to fine-tune the algorithm depending on the instance to solve.

First, the choice of variables to run QAOA on needs to be addressed. Considering the $\binom{N}{k}$ possibilities would be intractable. Variables can be chosen randomly but an approach incorporating one-bit flip move values can help in guiding toward an optimum. The k variables can be chosen amongst the non-tabu ones at each step. Plus, this means QAOA is an attempt at improving the solution one would get with the one-bit flip strategy outlined previously. One can either select the k variables greedily or add randomness by using the one-flip gains as weights for defining a probability to be chosen. For simplification, we consider the greedy selection based on one-bit flip move values. If we consider the chosen variables that were flipped, the update step of the incremental evaluation strategy can be applied. Let $l \leq k$ be the number of different bits. The newly generated candidate can be considered as a result of l sequential one-bit flips. Thus, l calls to the above-mentioned efficient procedure are required.

A second consideration concerns the tabu strategy for updating the tabu list. We choose to set as tabu the variables amongst the k chosen ones that were flipped. Choosing to flip all chosen ones could be problematic as it could lead to all variables being tabu very early during TS. An aspiration criterion can be used if the new candidate gives the best evaluation found during the search.

Finally, the question of how to run QAOA is of main importance. In our first scenario, QAOA will be run as a proxy for brute-force (with exploration properties) to optimize the subproblem defined over the k chosen variables. This is done by fixing in the QUBO the non-chosen ones from the current tabu solution x . The depth p of QAOA can be user-defined. In this work, we limit p to 2 to showcase sampling aspects

of QAOA at small depths.

Our QAOA-featured TS is outlined in Alg. 2 where QAOA addition is indicated in green shades. It starts with the same steps as with the standard Tabu search algorithm until line 8. The QAOA part kicks in from line 9 by first choosing a subset of k variables, and executing QAOA on the sub-QUBO problem where we optimize over the chosen k variables while keeping the remaining bits the same with the current point x . After obtaining the best point from QAOA (lines 11-13), we select the better one from the QAOA outcome $x^{k\text{-bit}}$ and the best one-bit flip point $x^{1\text{-bit}}$ and use it to update the current search point. The move values are then updated by l calls of the fast incremental method (line 14). Finally, if the best-so-far point is improved by the updated search point, we drop the accepted bit flips from the tabu list (line 21). Otherwise, we reset their tabu value to the sum of tabu tenure and a random tenure (line 23).

5.3.3 Enforcing locality with penalized QAOA

As a tool in local search algorithms, QAOA may be useful with modifications that impose notions of locality. We incorporate these notions in the cost hamiltonian so that they are captured during the QAOA evolution. This can be done through the cost function by adding a penalty term. A possibility is to consider the Hamming distance with a current tabu solution x . Hence, the objective for QAOA becomes:

$$\min_{x'} [f(x') + A\delta_H(x', x)], \quad (5.1)$$

where δ_H corresponds to the Hamming distance and A is a constant. The right-hand side additive term of Eq. (5.1) aims to encourage the output of candidates that differ by a few bits from the current solution if $A > 0$, and vice-versa. As a neighborhood sampler, the penalty may help in enforcing locality. However, setting the parameter A is non-trivial for activating the effect of the extra term.

There is also a possibility to add information about the fitness gain in differing from x . If switching a bit is an improved move, it would be prioritized. Our algorithm uses the one-bit flip move values Δ_x in order to select the variables to run QAOA on, from which we can construct a weighted penalty term:

$$-\frac{1}{2} \sum_{j=1}^N \Delta_x(j) (-1)^{x_j} x'_j. \quad (5.2)$$

For a minimization problem, $\Delta_x(j) < 0$ characterizes encouraging flipping the j -th

5.4. Tabu-driven QAOA sampling

Algorithm 2: QAOA-featured Tabu Search for solving QUBO.

Input: An initial solution x_0 , Cost function $f(x)$

Parameter : Tabu tenure TT, Random tabu tenure rTT, subproblem size k

Output: The best solution achieved x^*

```

1  $x^* \leftarrow x \leftarrow x_0$ ;
2  $Tabu(i) \leftarrow 0, \Delta_x(i) \leftarrow 0$  for  $i \in [1..N]$ ;
3 while stopping criterion not reached do
4    $x^{pre} \leftarrow x$ ;
5   for  $i \in [1..N] \cap Tabu(i) = 0$  do
6      $x_i^{(i)} \leftarrow 1 - x_i, x^{(i)} \leftarrow x$ ;
7     one bit-flip gain:  $\Delta_x(i) = f(x^{(i)}) - f(x)$ ;
8    $x^{1-bit} \leftarrow x^{(j)}, j \leftarrow \arg \min \Delta_x(i)$ ;
9   Select greedily or randomly a subset of variables  $K \subseteq I$  s.t.  $|K| = k$ ;
10  Get a new QUBO by fixing the  $N - k$  other variables in  $x$ ;
11  Run QAOA and get the best sample  $\hat{x}$  minimizing the new QUBO;
12   $x_i^{k-bit} \leftarrow x_i$  for  $i \in I \setminus K$  and  $x_i^{k-bit} \leftarrow \hat{x}_i$  for  $i \in K$ ;
13   $x \leftarrow$  the better out of  $x^{1-bit}$  and  $x^{k-bit}$ ;
14  Update move values  $\Delta_x(i)$  for  $i \in [1..N] \cap x_i^{pre} \neq x_i$ ;
15   $aspiration \leftarrow \text{False}$ ;
16  if  $f(x) < f(x^*)$  then
17     $x^* \leftarrow x, aspiration \leftarrow \text{True}$ ;
18   $Tabu(i) \leftarrow Tabu(i) - 1$  for  $i \in [1..N] \cap Tabu(i) > 0$ ;
19  for  $j \in [1..N] \cap x_j^{pre} \neq x_j$  do
20    if  $aspiration$  then
21       $Tabu(j) \leftarrow 0$ ;
22    else
23       $Tabu(j) \leftarrow TT + \text{Random}(rTT)$ ;

```

variable in new candidates. For a candidate, in this case, $\Delta_x(j)$ is added to the cost. Conversely, $\Delta_x(j) > 0$ would result in penalizing candidates with the j -th bit value flipped. One can also multiply by a positive constant A for enforcing more the locality effects. The penalty translates in an additional operator term H_{penalty} , following an application of a usual cost operator in QAOA, where:

$$H_{\text{penalty}} = -\frac{1}{2} \sum_{j=1}^N \Delta_x(j) (-1)^{x_j} \sigma_z^j \quad (5.3)$$

The corresponding quantum circuit of depth one is very simple and given by $\bigotimes_i R_Z((-1)^{x_i} \gamma \Delta_x(i)), \gamma \in \mathbb{R}$.

5.4 Simulations

We performed extensive simulations over instances of QUBO problems publicly available in the well-known OR-Library [65, 15, 16]. In designing them, our objectives are 1) investigating whether exploring larger neighborhoods can facilitate faster convergence (in terms of TS iterations), 2) elucidating the effect of locality on QAOA output given by the introduced penalty in Eq. (5.3), and 3) studying the utility of QAOA as a proxy for brute-force.

5.4.1 Larger neighborhood exploration benefits

To study the first objective we replaced QAOA with brute-force search in Alg. 2. Starting from the all-zero initial solution, we first run the basic TS with different constant tabu tenures ($rTT = 0$). Then, we do the same with brute-force TS for different values of k up to 20. We study two regimes that differ in how TS search results, namely when k is not comparatively small to N on instances where $N = 20$, and when it is on instances of 100, 200, and 500 variables.

k/N relatively large

In the first regime, we assume that we can explore a large percentage (more than 25%) of the instance size greedily. As instances, we take the first eight instances of *bqpgka* (named 1a-8a consisting of 30-100 variables), for which we solve by TS. Then we select randomly 20 variables out of N and clamp values of the non-selected ones from the solutions. We do so five times per instance, resulting in 40 instances of size 20. Then TS with different k/N values for subproblems (0.9, 0.75, 0.5, 0.25) were tried on this suite. When using brute-force, we set $TT = 2$ and tried many values for the basic TS.

Fig. 5.1 shows the proportion of (run, target value) pairs aggregated over all functions for 10 targets generated by linear spacing using the benchmarking and profiling tool IOHprofiler [49] for iterative optimization heuristics. The target values were normalized by the optimum of the problems. We observed that for $k/N \geq 0.5$, these instances are straightforward to solve (in 3 iterations). The case $k = 5$ required 8 iterations on one instance but managed to achieve optimality on all of them. However, the basic TS, run for 20000 iterations, failed to solve the same instance. Letting this instance aside, 5 iterations would be required for $k = 5$, and 22 for the classical TS procedure. Hence, we clearly observe, as expected, degrading performances as k/N decreases. This also enabled us to confirm numerically that a flip-gain-based approach

5.4. Simulations

when considering subproblems is in general beneficial for solving QUBOs.




Figure 5.1: Empirical cumulative distribution function (ECDFs) of tabu iterations for each algorithm aggregated over all 20-variable problems with 10 target values evenly spanning the range of all observed function values.

k/N relatively small

In this regime, we study the case when k/N is relatively small (less than 20%). The simulations are carried out on 15 QUBO instances from bqpqka: 100-variable (named 1d-5d), 200-variable (1e-5e) and 500-variable problems (1f-5f). Algorithms are run until the best objective value found in [65] is reached or a maximum number of TS iterations is reached.

We run the algorithms with different TT values ranging from 2 to 10 and adding 15, and for $k = 10, 15, 20$. We limit the number of iterations to 20000 for the basic version and 1000 for the brute-force approach (200 for $k = 20$ though). Table 5.1 shows for the different values of TT which ones achieved the best performances in terms of target only. In these cases, we observe that we can find a k such that better solutions are found using fewer iterations, especially on the densest instances 4f and 5f. These instances, when considering the underlying graph given by the coefficients connecting different variables, have a density of respectively 0.75 and 1 (a non-zero coefficient for each pair of variables). For $k = 20$ the proposed approach achieved optimality where the basic TS failed. Again, we observe performances, in terms of target achieved, depending on setting well the tabu tenure in accordance with k . Intuitively, one could

think that the larger k , the smaller TT has to be to save iterations and achieve a better objective value. But we clearly observe counter-examples.

Table 5.1: Best values of tabu tenure (ranging from 2 to 10, adding 15) achieving the optimum obtained with the first TS iteration(s) to reach the corresponding maximum given in [65]. The best performances per instance are highlighted in bold. The mention All means all TT values reached the same solution. The NA mention means no run returns the optimum, with the best value obtained in parenthesis.

ALGO.	1D (6333)	2D (6579)	3D (9261)	4D (10727)	5D (11626)
BASIC	15 / 300	ALL / 71	<9 / 90	2 / 67	6 / 171
$k = 10$	ALL / 13	>2 / 61	6 / 39	2 / 31	ALL / 19
$k = 15$	ALL / 10	2+5 / 26	ALL / 8	3 / 18	5 / 42
$k = 20$	ALL / 7	5 / 37	ALL / 5	ALL / 11	2 / 18
ALGO.	1E (16464)	2E (23395)	3E (25243)	4E (35594)	5E (35154)
BASIC	9 / 419	6 / 493	10 / 190	15 / 170	6 / 238
$k = 10$	8 / 485	7 / 111	ALL / 21	6 / 55	ALL / 25
$k = 15$	ALL / 13	6+7 / 31	ALL / 13	3 / 41	2 / 20
$k = 20$	2 / 35	NA (23370) / 82	3 / 44	3 / 25	4 / 31
ALGO.	1F (61194)	2F (100161)	3F (138035)	4F (172771)	5F (190507)
BASIC	8 / 970	7 / 795	4 / 449	NA (172734) / 1148	NA (190502) / 647
$k = 10$	15 / 213	7 / 337	15 / 77	NA (172449) / 110	NA (190502) / 126
$k = 15$	8 / 225	8 / 173	15 / 139	NA (172734) / 46	NA (190502) / 383
$k = 20$	NA (61087) / 184	NA (100158) / 101	4 / 57	4 / 57	9 / 150

Larger k exploration, in this regime, turns out to not always be beneficial. This seems counter-intuitive when considering a target objective only, on an instance-to-instance basis comparison. Fig. 5.2 shows the proportion of (run, target value) pairs aggregated over all functions for 1000 targets generated by linear spacing. The target values were normalized by the optimum of the problems. Again, we observe that, with larger k , the proportion of successes is higher, when measured at the same number of iterations. Note that this can only be observed for $k = 20$ up to 200 iterations. The proportion for the basic TS was close to 0.9 while the brute-force approach was superior to 0.99. Hence, we can reach very good solutions with fewer iterations as k increases.

In summary, as opposed to the previous regime, the structure of the problems becomes very important and we have to look at performances in an aggregated way to

5.4. Simulations

witness the benefits of exploring larger neighborhoods. Having outlined some performances given by the brute-force approach on subproblems, we switch to QAOA and study its sampling effect as a proxy.




Figure 5.2: Empirical cumulative distribution function (ECDFs) of tabu iterations for each algorithm aggregated over all dimensions, all problems, and 1000 target values evenly spanning the range of all observed function values.

5.4.2 QAOA as a proxy for brute-force

The second part of our simulations studies the output of QAOA as a proxy for brute-force. To this end, we first study an example TS run from our previous simulations. We take the subproblem QUBOs obtained at each step (except the first one), and run QAOA at $p = 1$ and 2, and we study the distribution of the energy given by $|\gamma, \beta\rangle$, after optimization, with and without the penalty term.

Having outlined the properties of the QAOA output, we run Alg. 2 and study its performances in comparison to the basic TS. From the optimized angles, we try different sampling strategies to generate a candidate per iteration: just sampling once, sampling 10 times and choosing a candidate greedily, and finally considering all samples (even during optimization) greedily. The latter corresponds to a quasi brute-force (BF) approach.




Figure 5.3: QUBO evaluation distribution given by sampling 10^5 times from the QAOA quantum state, and by running simulated annealing 1000 times (bottom), for the last TS iteration done on instance 1e.

Energy distribution of QAOA

As a first step, we study how the QAOA output distribution looks like at small depth, with the purpose of elucidating how it can help in avoiding detrimental greedy search behavior. We consider, as an example, instance 1e for which $k = 15, TT = 5$ used 13 iterations greedily. The subproblem QUBOs are kept and we run QAOA on them as previously stated. The third and last iterations are interesting when considering the penalty term. The former is a case where the optimum is obtained by flipping all bits except one and where all flip moves are favorable. The last iteration has flip moves from the current tabu solution discouraging flipping all bits. Plus, very few candidates (0.2%) improve over the tabu solution. For these iterations, we look at the quantum state given by QAOA and analyze the distribution of the QUBO evaluation (or energy in an Ising context).

Fig. 5.3 shows the distribution given by 10^5 samples from the last iteration's quantum state at $p = 1$ and 2. At $p = 1$, we observe a homogeneous spread with two major humps on the left and right sides of the tabu solution evaluation. There is a probability of 23.9% of improving from it by the quantum state. At $p = 2$, we see the distribution is shifted to lower energies, yielding an improved probability of 75%. The average energy is 16350 for QAOA $p = 1$ and 16428 for $p = 2$. Moreover, the standard

5.4. Simulations

deviation of the output decreases from 79.2 to 42. This is expected as to the limit of infinite depth, QAOA converges to the optimum with less variance.

When running simulated annealing 1000 times with a temperature of 17.5 using 100 steps, we observed that, unlike QAOA, the energy spread is restrained to a few points, the optimum being most present. Decreasing slightly the temperature or the number of steps would always yield the optimum. However, in terms of exploration opportunities, QAOA could allow visiting different paths that may lead to fewer iterations required towards improved solutions.

Theory shows that increasing the depth would permit QAOA to find the optimum assuming optimal parameters are found. But by limiting the depth, we can control how other good candidates are spread from the optimum. This could engender new paths to solve a problem differently, where a suboptimal solution on a subproblem leads to an easier one for QAOA towards better candidates. Note this can be done in different ways. One way would be using brute-force and perturbing or mixing the solution, which is not efficient. We could also have the same effect with simulated annealing. But in many cases, we can fail to find the optimum and even less find a bunch of candidates around it. Finally, due to its flexibility, QAOA permits in leveraging modifications to introduce locality notions in a multiobjective scheme for local search, such as the previously mentioned penalty. In the following, we study its effect on the QAOA distribution.

Penalty effect

In section 5.3.3, we introduced a penalty term to impose notions of locality in QAOA as a local search tool. An extra operator based on the hamiltonian given in Eq. (5.3), translates to a circuit of depth one concatenated with a QAOA layer. We study its effect on the QAOA distribution obtained on the resulting sub-qubos at the third and last iteration.

On iteration 3, for both QAOA and its penalized version, the distribution tends to output candidates with the largest Hamming distances to the current point as expected. Also, the most likely candidate is the one that completely differs from the current point, which is more favored by the penalty effect. No significant changes were observed when penalizing at $p = 2$.

For the last iteration, Fig. 5.4 shows that the original QAOA at $p = 1$ results in many probability peaks compared to the penalized version, which evolves to a major peak with an increased depth. The penalized version demonstrates two major peaks, from which the optimum and a close candidate to x are preferred.

This characterizes the interplay between optimizing and penalizing. At $p = 1$, the penalized version has a better probability of improving x (0.34 vs 0.239) and a higher probability of finding the optimum (0.1 vs 0.02). However, the unpenalized $p = 2$ version was more likely to output the optimum (respectively 0.75 and 0.26, where the penalty at $p = 2$ yields 0.62 and 0.26).

In summary, using the penalty creates a balance between the greedy approach and the one-bit-flip gains knowledge from the current solution. This could result in smoothening the distribution while favouring interesting candidates for both objectives. This will modify the search path taken during TS depending on the outcome. Having studied numerically the output of the quantum state one can get with QAOA, and the penalty effect, we switch to less idealized simulations where the subproblems depend on the QAOA output during the TS search.




Figure 5.4: Distribution of the evaluations (or energies) for the last iteration obtained on instance 1e given by the QAOA output state, with and without the penalty term. A colormap is given for the Hamming distance with the current tabu solution.

QAOA exploration possibilities

After looking at examples of QAOA output and outlining a few possible exploration opportunities, we carried out a few extra simulations but not considering the subproblems obtained with brute-force. Hence, QAOA (and its penalized version) was called

5.4. Simulations




Figure 5.5: Median of best normalized evaluations achieved over TS iterations for instance 1d. The mentions D1, D2, mean respectively $p = 1, 2$ and the penalized version is indicated by «pen». A higher curve corresponds to better solutions reached. At iteration 20, the basic TS value would be 0.38, while the lowest QAOA curve value is 0.967. The $m = 1000$ runs and BF are over 0.99 starting at the 8th iteration. At iteration 5, the penalized $p = 2, m = 10$ version is slightly better than the others, even BF (0.9049 vs 0.9007). At the 6th, the $m = 10$ versions are above BF (respectively 0.9645 and 0.9578 for unpenalized and penalized $p = 2$, 0.9555 and 0.9470 for $p = 1$, and 0.9409 for BF).

once per iteration with BIPOP-CMAES [145, 75] optimizing from one set of angles ². In the following, we give a few examples to illustrate the exploration possibilities.

We take instances where BF simulations required few iterations and where the basic TS was beaten in target. Namely, instance 1d for $k = 15$ and $TT = 5$ and 1e for $k = 15$ and $TT = 10$. We run Alg. 2 for 10 times, limiting them to 20 TS iterations. Different numbers of samples are used for generating a new candidate per iteration: just once, 10 times, and 1000 times. The latter could be considered a quasi-brute-force approach in these runs. We denote as m the number of samples used in the following.

²When using BIPOP-CMAES, we run circuits with 1000 measurements to estimate expectation values. The optimizer stops when it has reached 2000 evaluations. We obtained great performances in terms of averaged ratios (as the evaluations divided by the optimum of the subproblem), superior to 0.97 at the considered depths.




Figure 5.6: Median of best normalized evaluations achieved over TS iteration for instance 1e. At iteration 20, the lowest QAOA curve value is 0.991 (while 0.19 for the basic TS). The $m = 1000$ runs and BF are over 0.99 starting at the 11th iteration. At iteration 7, the original $p = 2$ QAOA using 10 samples point is higher than the others, even BF (0.7827 vs 0.7808). This also happens at iteration 12, with the penalized $p = 2, m = 1000$ (0.997085 vs 0.996902).

Fig. 5.6 and Fig. 5.5 show the median of the best normalized evaluation obtained per run by iteration. We observe from iteration 7 for 1d and 10 for 1e that $m = 1000$ is equivalent (in median) to the BF generation. In general, the more samples used, the better the solution found. However, we had at a few iterations median runs that achieved higher values than BF. For instance, this happened for the penalized $p = 2$ QAOA at the 5th iteration with $m = 10$ for 1d and $m = 1000$ at the 12th for 1e. We consider also the frequency of runs for which the basic TS was beaten, and the optimum was found. Table 5.2 summarizes our results. Increasing m improves the frequency of successful runs. We observe that new paths were found, mainly with an extra iteration or two but exceptionally one run or two over 10 could save one iteration. This was the case on instance 1e with 12 iterations instead of 13, exclusively with the penalized version. These examples are numerically in favor of a greedy (or quasi) approach to solving subproblems. However, QAOA allows, through a trade-off

5.5. Conclusion and outlook

between exploration and exploitation, to discover new paths toward optimality that are still interesting in terms of number of iterations.

Table 5.2: Frequency of successful runs in beating the basic TS and finding the optima for instance 1d for $k = 15$ and $TT = 5$, and instance 1e for $k = 15$ and $TT = 10$ (separated by /), in terms of QAOA settings (with the number of measurements noted m). We report also the number of iterations that led to the optimum.

QAOA	m	FREQUENCY BEATING TS (/10)	FREQUENCY OPTIMUM (/10)	ITERATIONS TO OPTIMUM
D1	1	0 / 0	0 / 0	
D1PEN		0 / 0	0 / 0	
D2		0 / 1	0 / 1	/ 14
D2PEN		0 / 0	0 / 0	
D1	10	7 / 2	1 / 2	18 / 15,16
D1PEN		4 / 2	0 / 2	/ 15,15
D2		5 / 1	2 / 1	11 / 14
D2PEN		4 / 0	1 / 0	10 /
D1	1000	8 / 7	8 / 7	9,10,12 / 13
D1PEN		8 / 5	8 / 5	10 / 12,13
D2		9 / 9	9 / 9	10,11 / 13,14
D2PEN		10 / 7	10 / 7	9,10 / 12,13

5.5 Conclusion and outlook

In this chapter, we studied sampling aspects when quantum approaches, specifically the QAOA algorithm, are considered in combinatorial optimization. We considered a practically relevant setting where a gate-based quantum algorithm, limited in the number of qubits, is utilized in a hybrid quantum-classical framework to solve large optimization instances faster. Our framework constitutes a powerful yet simple heuristic, Tabu Search, in tandem with QAOA as a local neighborhood sampler.

As a starting point, numerical experiments over open-source QUBO problems up to 500 variables validate using QAOA as a proxy to explore larger neighborhoods, under the assumption that subproblems are solved optimally. Continuing, we investigated the exploration possibilities given by QAOA output at small depth. User-defined parameters such as depth and number of measurements used to generate a candidate can be increased to favor exploitation. In our examples, solving subproblems emphasizing more on the latter gave better general performances. Yet, we found that exploration

can be beneficial. Iterations can be saved with our QAOA procedure, illustrating that missing to generate the solution of a subproblem in previous iterations could yield faster paths towards better solutions. Hence, the QAOA-based algorithm we introduce in this work becomes a very flexible tool in such hybrid quantum-classical settings.

We see numerous possibilities for future work. First, our model allows for many hyperparameters whose function needs to be explored, and, as is usually done in many local search methods, the exploration/exploitation trade-offs can be made online-adaptive. Second, the effect of real-world limitations, most importantly noise, and hardware connectivity, calls for further investigation. Although QAOA can be run on real hardware [79, 201], its output quality will improve as the quantum devices decrease in error. Angle-setting strategies such as the ones studied in chapter 4 could be incorporated to reduce the runtime. Finally, it would be interesting to propose different frameworks (e.g. [147, 148, 67, 120]) with special emphasis on the exploration possibilities given by small-depth quantum algorithms, and cross-compare with standard techniques in future works. We believe our approach combined with these types of analyses will provide new promising ways to maximize the use of limited near-term quantum computing architectures for real-world and industrial optimization problems.

5.5. Conclusion and outlook
