
This is an electronic reprint of the original article.
This reprint may differ from the original in pagination and typographic detail.

Rindell, Tom; Yenilen, Berat; Halonen, Niklas; Pönni, Arttu; Tittonen, Ilkka; Raasakka, Matti
Exploring the optimality of approximate state preparation quantum circuits with a genetic algorithm

Published in:
Physics Letters A

DOI:
[10.1016/j.physleta.2023.128860](https://doi.org/10.1016/j.physleta.2023.128860)

Published: 01/05/2023

Document Version
Publisher's PDF, also known as Version of record

Published under the following license:
CC BY

Please cite the original version:
Rindell, T., Yenilen, B., Halonen, N., Pönni, A., Tittonen, I., & Raasakka, M. (2023). Exploring the optimality of approximate state preparation quantum circuits with a genetic algorithm. *Physics Letters A*, 475, Article 128860. <https://doi.org/10.1016/j.physleta.2023.128860>

This material is protected by copyright and other intellectual property rights, and duplication or sale of all or part of any of the repository collections is not permitted, except that material may be duplicated by you for your research use or educational purposes in electronic or print form. You must obtain permission for any other use. Electronic or print copies may not be offered, whether for sale or otherwise to anyone who is not an authorised user.



Exploring the optimality of approximate state preparation quantum circuits with a genetic algorithm

Tom Rindell^{a,b}, Berat Yenilen^c, Niklas Halonen^a, Arttu Pönni^a, Ilkka Tittonen^a,
Matti Raasakka^{a,*}

^a Micro and Quantum Systems group, Department of Electronics and Nanoengineering, Aalto University, Finland

^b Department of Physics, University of Helsinki, Finland

^c Department of Physics, RWTH Aachen University, Germany

ARTICLE INFO

Article history:

Received 30 January 2023

Received in revised form 4 April 2023

Accepted 18 April 2023

Available online 24 April 2023

Communicated by M.G.A Paris

Keywords:

Quantum state preparation

Genetic algorithm

Quantum circuit complexity

Noisy intermediate-scale quantum

NISQ

ABSTRACT

We study the approximate state preparation problem on noisy intermediate-scale quantum (NISQ) computers by applying a genetic algorithm to generate quantum circuits for state preparation. The algorithm can account for the specific characteristics of the physical machine in the evaluation of circuits, such as the native gate set and qubit connectivity. We use our genetic algorithm to optimize the circuits provided by the low-rank state preparation algorithm introduced by Araujo et al., and find substantial improvements to the fidelity in preparing Haar random states with a limited number of CNOT gates. Moreover, we observe that already for a 5-qubit quantum processor with limited qubit connectivity and significant noise levels (IBM Falcon 5T), the maximal fidelity for Haar random states is achieved by a short approximate state preparation circuit instead of the exact preparation circuit. We also present a theoretical analysis of approximate state preparation circuit complexity to motivate our findings. Our genetic algorithm for quantum circuit discovery is freely available at <https://github.com/beratyenilen/qc-ga>.

© 2023 The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Quantum computing has gained a lot of attention in the recent years mainly due to the continued development of quantum hardware by both academic research groups and companies. Significant speed-ups in certain computational tasks by quantum computers were theoretically demonstrated already in the 1990's by several authors [1–4]. However, a quantum advantage for specific computational problems was only recently experimentally demonstrated [5,6]. This has brought the field past the proof-of-concept stage to a point, where real-world applications for quantum computing are starting to emerge, although still severely limited by the hardware. First applications are expected to be simulations of quantum systems and quantum machine learning, as these typically have lower hardware requirements compared to many other potential applications, which require higher computational fidelity and thus quantum error correction with additional overhead [7].

The application of genetic algorithms to quantum circuit discovery has already been discussed in the literature for quite a while (see e.g. [8–13]). Automating the design of quantum algorithms is an appealing idea, as the behavior of quantum systems, such as a quantum processor, is often unintuitive and challenging to understand. However, only recently the growth of classical and quantum computational resources has started to allow for the practical feasibility of this approach. Genetic algorithms have recently been successfully applied to quantum circuit optimization, e.g., in quantum machine learning [14–16] and chemistry [17], and specialized software packages have been developed for the task [18]. The original inspiration for our work was however the paper [19] by Potoček et al., who obtained promising results in applying a multi-objective genetic algorithm to discover optimized quantum circuits for the implementation of quantum Fourier transform and Grover search. With the multi-objective approach not only one finds circuits that reproduce the exact desired unitary, but the algorithm can also produce a set of alternative circuits with variable trade-offs between accuracy and circuit depth. This can be highly desirable, when running the circuit on physical hardware, as the error rate grows exponentially as a function of the circuit depth in the absence of error correction.

* Corresponding author.

E-mail address: matti.raasakka@aalto.fi (M. Raasakka).

State preparation is among the most fundamental tasks in quantum computing. Indeed, essentially all quantum algorithms proceed by preparing some quantum state for the qubits, whose statistics contains information for solving the computational problem at hand. The state preparation problem can be stated as follows: Given some set of basis gates, find a quantum circuit composed of these gates, which transforms the all-zeros state $|00\dots 0\rangle$ to some target state $|\psi\rangle$ [20,21]. Besides being a suitable first problem to gauge the performance of our genetic algorithm for circuit discovery, this problem also has practical significance. If the same state needs to be prepared a large number of times during the running of the algorithm, as is often the case due to the probabilistic nature of quantum computing, small gains in the preparation process of the state may lead to significant savings in the running time of the algorithm.

The minimum number of gates required to reach a certain state (possibly with a margin of error) from a fixed reference state is called the quantum circuit complexity of a quantum state [22,23]. Our genetic algorithm for state preparation provides estimates of the quantum circuit complexity of multi-qubit states. On the theory side, quantum circuit complexity is of fundamental importance for several reasons. Most fundamentally, the difficulty of a computational task can be quantified by how the circuit complexity scales as a function of the input size. Understanding the quantum circuit complexity of different computational problems has far-reaching practical consequences, e.g., for the safety of cryptographic protocols. In condensed matter theory, phase transitions have recently been studied with quantum circuit complexity [24–26]. On the other hand, in quantum gravity circuit complexity has been the topic of intense research, in particular in relation to the anti de Sitter/conformal field theory (AdS/CFT) duality through the ‘complexity=volume’ [27,28] and ‘complexity=action’ [29] conjectures, with attempts to relate quantum circuit complexity of quantum field theory states to spacetime geometry.

In this work, we developed a genetic algorithm enabling us to find quantum circuits, which provide trade-offs between weighted circuit depth (‘circuit cost’) and the fidelity of the output state. We then applied the genetic algorithm to find approximate state preparation circuits for 1000 Haar random states of 5 qubits, using the qubit connectivity and the noise model of a 5-qubit IBM Falcon 5T processor provided by IBM Qiskit. More specifically, we first applied the low-rank state preparation (LRSP) algorithm introduced in [30] to the states, and then further optimized the circuits with the genetic algorithm to study their optimality. For all the states, the fidelity with noise is maximized by an approximate circuit. Theoretical analysis shows that, in the absence of noise, the fidelity of approximate circuits approaches unity exponentially as the circuit depth increases, while noise causes an exponential decay in fidelity as a function of the circuit depth. It is the delicate balance of these two effects, which allows to find approximate circuits that maximize the noisy fidelity. We find that the maximal noisy fidelity of the output states is increased by 9.4% on average with the circuits found by the genetic algorithm as compared to the circuits given by the LRSP algorithm, with maximum observed relative improvement of 38%.

The structure of the rest of the paper is as follows. In Section 2 we discuss the state preparation problem in more detail. We also discuss the effect of noise on the state preparation problem on physical hardware, and provide a theoretical analysis of the relation between circuit depth and state preparation fidelity. Section 3 describes the detailed implementation of our genetic algorithm. In Section 4 we explore the performance of our algorithm, and show that we improve on the LRSP algorithm results. Section 5 provides first a summary and discussion of the results. Finally, we provide some pointers for the direction of future work.

2. State preparation problem in the presence of noise

2.1. Problem description

Let us restate what we call in this paper the *exact state preparation problem* more precisely. Let n be a positive integer. We are given

- an n -qubit state $|\text{target}\rangle$, the *target state*, and
- a universal set of quantum gates G .

Quantum gates are typically assumed to be simple unitaries, which act only on a few qubits at a time.¹ Some gates, such as a 1-qubit rotation gate $R_z(\theta)$, may be parametrized by one or more real parameters. Universality of the gate set implies that we can approximate any n -qubit unitary to an arbitrary precision $\epsilon > 0$ by a finite sequence of gates from the gate set. Several examples of universal gate sets exist. However, in this work we mainly concentrate on the universal gate set

$$G_{\text{IBM}} = \{R_z, \sqrt{X}, X, \text{CNOT}\},$$

which has been experimentally implemented as the native gate set of IBM quantum processors. The exact state preparation problem then asks us to find a sequence of gates $g_1 g_{l-1} \dots g_2 g_1$, $g_k \in G \forall k$, such that

$$g_1 g_{l-1} \dots g_2 g_1 |00\dots 0\rangle = |\text{target}\rangle.$$

Upper bounds on the number of CNOT gates for exact state preparation without ancilla qubits have been derived and improved upon by several different authors [20,31,32,21,33].² Typically only the CNOT count is considered for near-term applications, as the error rates of 1-qubit gates are an order of magnitude smaller. Often the connectivity of the qubits is not restricted in deriving these results, which means that the CNOT gates may act on any pair of qubits. The strictest bounds to our knowledge on the number of CNOT gates for all-to-all connectivity are given by [21,33]:

$$C_{\text{ub}} = \begin{cases} \frac{23}{24} 2^n - \frac{3}{2} 2^{\frac{n+1}{2}} + \frac{4}{3}, & n \text{ odd } (n \geq 5) \\ \frac{23}{24} 2^n - 2^{\frac{n}{2}+1} + \frac{5}{3}, & n \text{ even} \end{cases}.$$

However, the qubit connectivity may also be restricted without changing the problem significantly. This brings the problem closer to practice, as the connectivity of qubits is usually heavily restricted on physical hardware, but makes theoretical analysis more challenging. The only work we found considering restricted connectivity is [32], where the authors consider the state preparation problem on a linear chain of qubits. The difference in CNOT counts between different topologies is always polynomial, as one may permute qubits using a polynomial number of SWAP gates, but it may still have a significant contribution in practice.

Also, the approximate state preparation problem may be considered, which asks us to find a sequence of gates $g_1 g_{l-1} \dots g_2 g_1$, $g_k \in G \forall k$, such that

$$|\langle \text{target} | g_1 g_{l-1} \dots g_2 g_1 | 00\dots 0 \rangle|^2 < \delta$$

¹ Some implementations of quantum processors, such as ion traps, may allow for gates, which act on all the qubits simultaneously. However, here we mostly concentrate on superconducting devices, such as transmon qubits.

² In this paper we only consider the case without ancilla qubits. With the use of ancilla qubits the circuit depth can be significantly reduced, see e.g. [34,35]. The closely related task of unitary synthesis has also been the subject of active research efforts recently, see e.g. [33,36–38].




Fig. 1. Histogram of noisy fidelities for 1000 simulated state preparation circuits for 2, 3, 4 and 5 qubits provided by the standard Qiskit library method. The simulations were performed using the IBM Qiskit FakeVigo backend.

for some fixed $\delta > 0$. Upper bounds on the number of CNOT gates in these cases are much less known. A recent approach to numerically estimate the bound for a small number of qubits is found in [39]. However, we have been unable to find a theoretical analysis of the approximate state preparation problem. In the next subsection (and in Appendix A) we provide a back-of-the-envelope theoretical analysis of the optimal trade-off between fidelity and CNOT count, which we later confirm with our numerical results obtained with the genetic algorithm.

The state preparation problem for a physical quantum computer is somewhat different from the abstract version of the problem due to the presence of noise. Even if we could find the absolute minimum depth circuit, which prepares the target state exactly, this may still not be the best circuit in practice. Producing a generic state of n qubits exactly requires an exponential $\sim 2^n$ number of gates, which makes the circuit basically useless with the non-negligible levels of noise in the current hardware, for just a moderate number of qubits. This is demonstrated in Fig. 1, where we have plotted the performance of the standard IBM Qiskit exact state preparation method (based on the algorithm presented in [20]) for an increasing number of qubits in the presence of simulated noise. Therefore, what we ideally want to find is a circuit, which offers an optimal trade-off between the circuit depth (or the CNOT count) and the noise-free error in the target state such that the total error in the presence of noise is minimized. This is where a genetic algorithm becomes particularly useful, as it can perform simultaneous multi-objective optimization on these parameters. In Section 3 we will explore in detail the application of our genetic algorithm to the state preparation problem.

2.2. Theoretical analysis

On a theoretical level, the trade-off between circuit length and fidelity can be understood in the following way. First of all, let us consider the dependence of the final state fidelity on the length of the circuit in the presence of noise. For simplicity imagine that all our gates have the same error rate p . Then the probability of no errors for a circuit of length l is $P(\text{no error}) = (1 - p)^l$. Assuming that any error takes us to an approximately orthogonal state, we get for the final state fidelity

$$F(\rho_{\text{final}}, U_{\text{circ}}|00\dots 0\rangle) \approx (1 - p)^l,$$

which shows an exponential decrease in fidelity as the circuit length increases. In practice, at least for a superconducting NISQ

device, we may consider l to be the number of CNOT gates in the circuit to get a reasonable estimate, since the single-qubit gates have error rates that are typically an order of magnitude smaller.

On the other hand, we may reduce the effects of noise by employing a shorter circuit, which approximates the target state. In Appendix A we estimate the dependence of the noise-free error ϵ on the CNOT count l in preparing an arbitrary state. We find that (given certain assumptions discussed in more detail in the appendix) we may approximate

$$\epsilon \approx \exp \left[-\frac{A(n)}{2} \frac{l}{l_{\text{ph}}} - \frac{B(n)}{2} \right],$$

where

$$A(n) = \frac{(2 \ln 2)n + \ln \left(\frac{n(n-1)}{2} \right)}{2^n - n - 1}, \quad B(n) = \frac{(\ln 2)n^2}{2^n - n - 1},$$

and n is the number of qubits. Moreover, l_{ph} is the average number of physical CNOT gates required to implement one logical CNOT gate, which reflects the connectivity of the qubits. We observe that the error ϵ decreases exponentially as a function of the number of CNOT gates l . This is essentially due to the fact that the number of states or, more precisely, the fraction of the Hilbert space volume reached with a fixed error $\epsilon \ll 1$ increases exponentially in l : Each CNOT gate allows for two more single-qubit rotations (each with two independent parameters) to be added after it, thus increasing the dimensionality of the subspace of states, which can be reached without error, by 4. Since the dimensionality grows linearly in l , the volume grows exponentially. Thus, the error ϵ needed to cover the whole Hilbert space (i.e., to reach the volume of the full Hilbert space) decreases exponentially in l . Accordingly, we should be able to get a decent approximation to any state with a relatively small number of CNOT gates as compared to the circuit producing the exact state.

Now, imagine we want to produce the target state $|\text{target}\rangle$ with our noisy quantum computer. We may reach noise-free fidelity

$$F(|\text{approx}\rangle, |\text{target}\rangle) = 1 - \epsilon^2 \approx 1 - \exp \left[-A(n) \frac{l}{l_{\text{ph}}} - B(n) \right]$$

with a circuit containing at most l CNOT gates and outputting the approximation state $|\text{approx}\rangle$. When we run this circuit on the noisy quantum computer, the fidelity reduces approximately by the factor $(1 - p)^l$, where p is the error rate of the CNOT gates.³ Thus, we reach a total fidelity

$$F(\rho_{\text{noisy}}, |\text{target}\rangle) \approx (1 - p)^l \left(1 - \exp \left[-A(n) \frac{l}{l_{\text{ph}}} - B(n) \right] \right).$$

The total fidelity in the presence of noise has a global maximum at

$$\frac{l_*}{l_{\text{ph}}} = \frac{1}{A(n)} \ln \left(1 - \frac{A(n)}{\ln(1 - p)} \right) - \frac{B(n)}{A(n)},$$

meaning that at length l_* the trade-off between circuit accuracy and noise is optimized. (See Appendix B for an asymptotic analysis of these formulae.)

Unfortunately, in practice the number l_{ph} of physical CNOT gates per one logical CNOT gate is not easy to compute exactly, since the distribution is unknown. However, a lower bound is provided by the average distance $\langle d \rangle$ between qubit pairs, $l_{\text{ph}} \geq \langle d \rangle$,

³ As we later note in fitting these formulae to our numerical results, the error rate p also accounts implicitly for 1-qubit gates errors, decoherence and other possible errors that crop up during the computation. Thus, we may expect p to be larger in practice than the bare CNOT error rate.

since in order to implement a CNOT between an arbitrary pair of qubits we need to at least connect the qubits via CNOT gates. An upper limit is provided by $l_{\text{ph}} \leq 1 + 6((d) - 1)$, since we may always move distant qubits next to each other by permuting qubits with SWAP gates before applying the CNOT gate, and then move the qubits back to their original positions afterwards. (However, this is usually not the optimal way to perform the logical CNOT gate.) Each SWAP gate requires 3 CNOT gates, so altogether we need 6 CNOT gates to move a qubit one step forward and then back. Using these estimates, we find that for a small number of qubits a significant reduction of noise can be achieved by using short approximate circuits for state preparation tasks. For example, for a linear chain of 10 qubits ($n = 10$, $l_{\text{ph}} \sim (d) \approx 3.7$) with a CNOT error rate $p = 0.01$ we find the maximum value of $F(\rho_{\text{noise}}, |\text{target}\rangle)_{\text{max}} \approx 0.163$ at circuit length $l_* = 67$. For the exact state preparation circuit of length $l \approx 2^{10} l_{\text{ph}}$ the total fidelity drops essentially down to zero. In Section 4 we will see that these formulae are reproduced reasonably well by our numerical experiments.

3. Multi-objective optimization using genetic algorithms

Genetic algorithms are heuristic optimization algorithms motivated by the evolutionary processes in nature [11,40]. The basic idea is that we have a population of candidate solutions to the problem described by their genomes, which are iteratively refined through selective breeding. Individual solutions are selected from the population for breeding (i.e., to mutate or combine with other selected solutions) based on their fitness, which measures how well they solve the desired problem. Diversity in the population is maintained by mutations. The selection pressure on the population guides it towards more optimal solutions. Genetic algorithms seem to be particularly well suited to global optimization problems, which present a complex fitness landscape.

Since genetic algorithms deal with populations of candidate solutions, they are also naturally adapted to multi-objective optimization problems [40]. In this case there exist several different properties of the solutions, which we try to optimize simultaneously, but which cannot be simultaneously optimized to their global minima/maxima. In this case, the best we can do is to find solutions with a reasonable trade-off between the different desired properties. Several multi-objective genetic algorithms have been developed in the past, such as Fast Non-dominated Sorting Genetic Algorithm (NSGA2) and Improved Strength Pareto Evolutionary Algorithm (SPEA2), which perform very well against many different kind of problems. In this work we apply the NSGA2 algorithm [41] for ranking and selection of the candidate solutions, as it is implemented in the DEAP Python package [42], while genetic representation and the evolutionary operations of mutation and cross-over for quantum circuits are adapted from the work of Potoček et al. [19]. Let us now describe some of the details of the implementation. Our full implementation can be found at [43].

The problem instance is defined by the target state $|\text{target}\rangle$ and the set of allowed gates. In this work, we focus on the IBM quantum processors with the native gate set

$$G_{\text{IBM}} = \{R_z(\theta), X, \sqrt{X}, \text{CNOT}\}$$

where $R_z(\theta) = e^{iZ\theta/2}$, $\theta \in [0, 2\pi)$ is the parametrized rotation by angle θ around the z -axis on the Bloch sphere. This set of gates constitutes a universal gate set, where the only native multi-qubit gate is the CNOT gate. Moreover, the allowed CNOT gates are limited by the qubit connectivity of the quantum processor. A quantum circuit C is represented in the genetic algorithm by a genome, which is just a list of tuples of form (g, q) , where $g \in G_{\text{IBM}}$ and q is a list of indices specifying the qubits on which the gate g acts.

The operations, which are used to modify circuits, are defined as follows:

- Operations acting on single gates within one circuit:
 1. **Discrete uniform mutation:** Iterate over all gates in the circuit and randomly change the target and/or control qubits with probability inversely proportional to the length of the circuit.
 2. **Continuous uniform mutation:** Iterate over all gates in the circuit and adjust the parameter for the $R_z(\theta)$ gates.
 3. **Move gate:** A randomly chosen gate is moved to a new randomly chosen position in the circuit.
 4. **Insert mutate invert:** Apply a discrete mutation on a single gate, after which a randomly selected gate and its inverse are placed immediately before and immediately after, respectively.
- Operations acting on sequences of gates within one circuit:
 5. **Sequence insertion:** Generate a random gate sequence, which is then inserted at a random point in the gate list.
 6. **Sequence and inverse insertion:** Apply a sequence insertion at one point and insert the inverse sequence at a later point in the gate list.
 7. **Sequence deletion:** Delete a randomly selected interval in the gate list.
 8. **Sequence replacement:** Sequence deletion followed by sequence insertion.
 9. **Sequence swap:** Swap two randomly chosen sequences in the gate list.
 10. **Sequence scramble:** Pick a sequence in the gate list randomly, and perform a random permutation on that sequence.
- Operations acting on two circuits:
 11. **Crossover:** Select a random number of gates from the beginning of the first parent circuit, disregard the same number of gates from the second parent, and combine them into a child circuit. Additionally, a second child with the opposite combination is returned as well.
- Operations modifying other properties of a single circuit:
 12. **Permutation mutation:** Permute the mapping of logical qubits to physical qubits.
 13. **Clean:** Remove all pairs of subsequent gates that are inverses of each other, and combine subsequent pairs of z -rotations. Iterate until no simplifications are found.

Single-gate operations 1–4 depend on the expected mutation count (EMC) parameter, which is a fixed constant during the evolution. For these operations the probability of a mutation on a single gate is given by EMC/l , where l is the length of the gate list. Operation 2 also takes an additional continuous mutation width (CMW) parameter for the adjustment of parametrized gates. The gate parameters are adjusted by adding a value from Gaussian distribution with standard deviation CMW/ϵ , where ϵ the error of the circuit. With this approach the parameter will be adjusted less for more optimal circuits. Sequence operations 5–10 pick the length of the sequence randomly from a geometric distribution with the mean value given by the expected sequence length (ESL) parameter, which is another fixed constant. Constants EMC, CMW and ESL are all set to 2.0 by default, following Potoček et al. [19]. While the evolutionary operations above were mostly adapted from Potoček et al. [19], some have been altered for the purposes of this implementation.

Before running the genetic algorithm, parameters such as the population size, the number of generations, the number of qubits, and the simulation backend must be initialized. Some parameters, such as the noise model, basis gates, and the connectivity can be taken, e.g., from the desired simulated Qiskit backend, but can also be adjusted if needed. The population is initialized with circuits

obtained from the LRSP algorithm [30]. The remaining population is filled with randomly generated circuits. The size of the population stays constant during the evolution.

Candidate circuits in every generation are evaluated with respect to two criteria:

1. **Circuit cost:** Defined as

$$\text{cost} = (\text{number of single-qubit gates}) \\ + 10 \times (\text{number of CNOT gates}).$$

2. **Fidelity:** Apply the unitary U_C corresponding to the circuit C to the initial state $|00\dots 0\rangle$, and compute the fidelity

$$F(|\text{target}\rangle, U_C|00\dots 0\rangle) = |\langle \text{target} | U_C | 00\dots 0 \rangle|^2$$

with respect to the target state.⁴

In comparing circuits, a circuit is said to dominate those circuits, which have both a lower fidelity and a higher cost. The optimal 'non-dominated' circuits are those, which are not dominated by any other circuit. After evaluation the circuits are sorted in a list from best to worst using the multi-objective NSGA-2 non-dominated sorting algorithm. Essentially, the circuits are ranked according to their dominance: Non-dominated circuits get rank 1; those circuits, which are dominated only by rank 1 circuits, get rank 2, etc.

The production of the next generation proceeds as follows. First, the best 10% of the population according to their rank are moved directly into the next generation. If the circuits from a certain rank do not all fit into the 10%, the chosen ones are picked randomly. Then, the remaining 90% of the next generation is filled with modified versions, i.e., offspring of the circuits in the current generation. First, the rank of the parent circuit is picked according to a Boltzmann-like distribution, where the probability of a rank r to be chosen is proportional to e^{-r} . After the rank has been chosen, any circuit is picked from that rank with an equal probability. New candidates are generated from the chosen circuits with evolutionary operations listed above until the population is filled for the next generation. Evolutionary operations are picked randomly with equal probability. This procedure results in a new generation of circuits with an altered population. The process is iterated over for a given number of generations.

Once the algorithm has been run for a given number of generations, a final 'clean' operation is performed on every circuit of the final population to remove all the gates that trivially cancel each other. The final population is then saved alongside a DEAP logbook file that contains statistical data for every generation.

4. Results

For our target states we generated 1000 random statevectors for 5 qubits, sampled from the uniform Haar measure. For each state, the genetic algorithm was seeded with circuits obtained from the LRSP algorithm (transpiled using the Qiskit transpiler), and then run on the Triton supercomputer cluster [44] with population of 400 for 30 000 generations. Without multithreading, the duration of a single run for one state varied from 32 to 50 hours. The backend simulation and noise model were implemented from IBM

Qiskit's 'FakeVigo' backend, which models a 5-qubit IBM Falcon 5T processor.

The best non-dominated circuits found for a given state reproduce the approximately exponential convergence to the exact state as a function of the number of CNOT gates (in the absence of noise), as predicted by our theoretical analysis. See Fig. 2 for examples of results for four arbitrarily chosen individual random states. For each state the genetic algorithm was run 100 times, and the non-dominated solutions of the final populations are plotted in Fig. 2. When we compare the circuits with the best noise-free fidelity for a given number of CNOT gates generated by the genetic algorithm and LRSP algorithm for a given state, we can see a significant rise in fidelity over LRSP results in most cases. This shows that the approximate state preparation circuits obtained from the LRSP algorithm are not near-optimal, at least when transpiled with the Qiskit transpiler for the specific quantum processor.

For Fig. 3, we ran the genetic algorithm once for each one of 1000 random states, and plotted the non-dominated solutions of the final populations. Fitting of the theoretical fidelity curve derived in Subsection 2.2 to the data is also shown with $n = 5$ and l_{ph} ranging between the limiting values $\langle d \rangle = 1.8$ and $1 + 6(\langle d \rangle - 1) = 5.8$ with unit steps. The best fit with the optimal circuits is achieved with $l_{ph} \approx 3.8$ indicating that about this many physical CNOT gates are needed to implement one logical CNOT gate in the state preparation circuits, on average. The CNOT error rate of the backend is $p = 0.0088$. We see that the actual error rate seems to be somewhat higher, as the theoretical fit overestimates the noisy fidelities by a small margin. This is explained by the fact that also other errors not accounted for by our theoretical analysis, such as single-qubit gate errors and identity errors, contribute to the exponential decay of fidelity as a function of the circuit length.

We may then consider including the noise in the simulation of the non-dominated solutions. Running the simulation with noise for the final populations shows that the highest total fidelity is achieved for shorter approximate circuits, as predicted by our theoretical analysis. For most of the runs improvements over LRSP circuits were achieved. Averaging over single genetic algorithm runs for 1000 random states gives an average absolute improvement of about 0.057 ($\sigma \approx 0.031$) in the maximum fidelity with noise, or an average relative improvement of about 9.4% ($\sigma \approx 5.2\%$). The largest observed absolute improvement was about 0.22, and the largest relative improvement 38%. In Fig. 4 the distributions for the absolute and relative differences between the LRSP and genetic algorithm circuits are plotted. We see that in most cases the genetic algorithm is able to significantly improve on the circuits found by the LRSP algorithm. In a few cases the genetic algorithm was unable to improve on the performance of the LRSP circuit. In these cases it would be possible either to use the LRSP circuit directly, or to run the genetic algorithm again. It is likely that running the genetic algorithm more than once would improve the results, since (as we saw already in Fig. 2) there is quite a lot of variation between the runs for an individual state.

5. Summary & outlook

We implemented a multi-objective genetic algorithm for quantum circuit discovery, and applied it to perform a numerical analysis of the approximate state preparation problem. We focused particularly on the native gate set, connectivity and noise profile of the IBM 5-qubit processors. The genetic algorithm was able to improve on the circuits found by the low-rank state preparation (LRSP) algorithm introduced in [30] and compiled with the Qiskit compiler in that it found preparation circuits with the same number of CNOT gates but a higher fidelity for the output state. As a result, also the maximum fidelity with realistic noise (implemented from the IBM Qiskit 'FakeVigo' backend) was significantly

⁴ In the current implementation the application of the unitary U_C and the computation of the fidelity is done via classical simulation. This is the most resource-intensive part of the algorithm. However, one could imagine moving this part of the algorithm to be performed by a quantum computer if one had an access to a good enough machine.

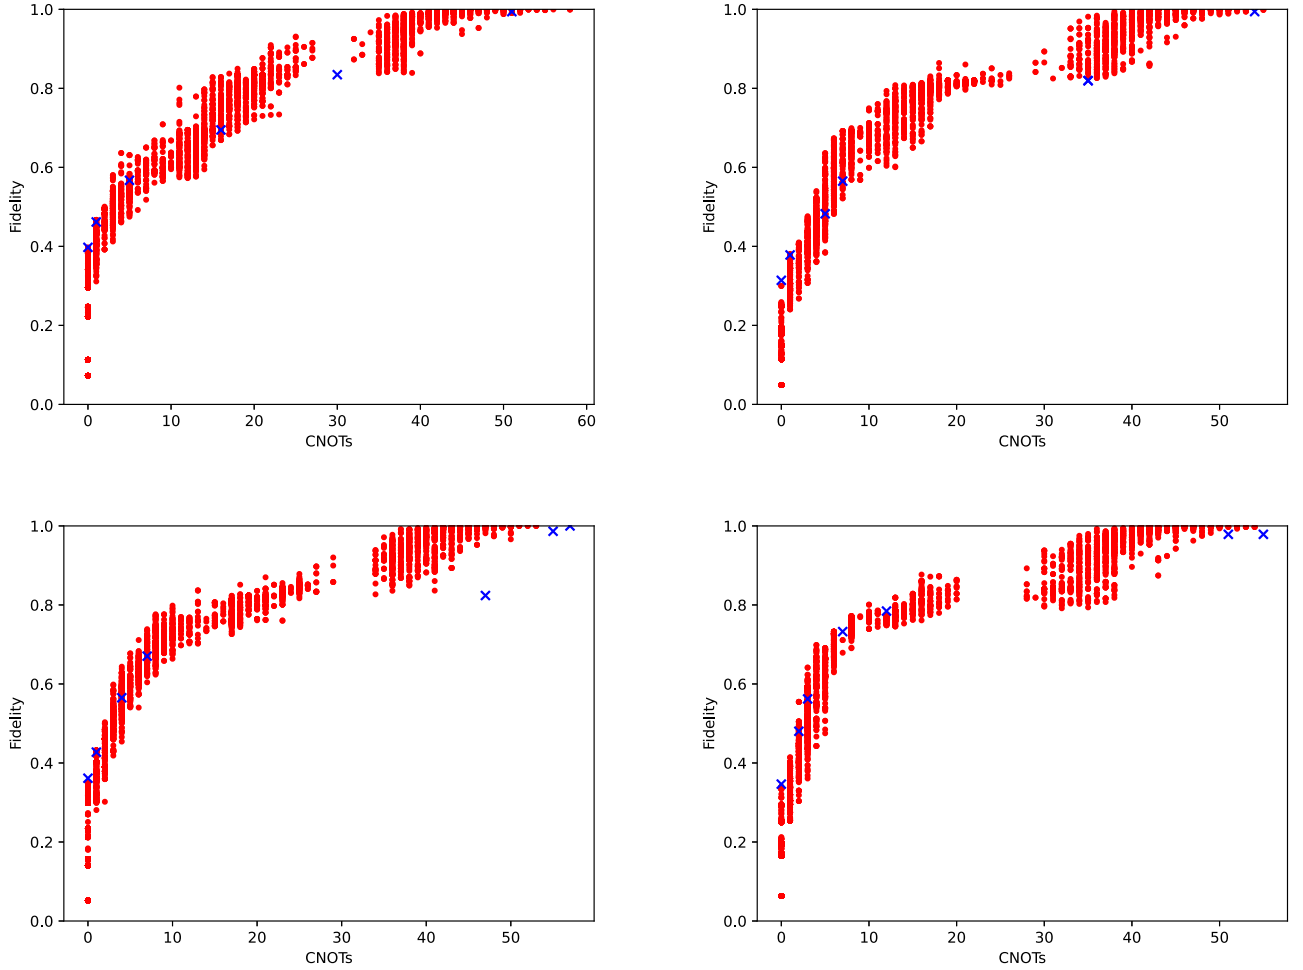


Fig. 2. Example results for four arbitrarily picked random states (indices 160, 559, 748 and 932 in our database for the states in [43]). Blue crosses: Circuits obtained by the LRSP algorithm. Red dots: Non-dominated circuits from 100 runs of the genetic algorithm for each state.




Fig. 3. Numerical results and theoretical fidelity curves for 1000 random states without (left) and with (right) noise. Red dots: Non-dominated solutions of genetic algorithm runs for 1000 different random states. Parameters for the runs were 30 000 generations, population of 400, 5 qubits, and FakeVigo backend connectivity and noise model. Curves: Theoretical fidelity curves with $n = 5$ and $p = 0.0088$. I_{ph} varies from $\langle d \rangle = 1.8$ (blue) to $1 + 6(\langle d \rangle - 1) = 5.8$ (purple) with unit steps.

improved in most cases with an average relative improvement of about 9.4%, and maximum observed improvement of 38%. We also provided a back-of-the-envelope theoretical analysis of the approximate state preparation problem to support our numerical findings, and derived the optimal fidelity curve, which was reproduced by our numerical results. Our results support the overall conclusions

that (1) there is still room to improve in current state-of-the-art algorithms for approximate state preparation, (2) the maximum fidelity for generic state preparation on a NISQ computer may be achieved by using a shorter approximate circuit rather than the exact state preparation circuit, and (3) genetic algorithms continue to provide a useful tool for quantum circuit discovery and optimiza-




Fig. 4. Left: The histogram of absolute differences in noisy fidelities between the genetic algorithm circuits and the LRSP circuits, together with a Gaussian fit. Right: The same for the relative differences in fidelity.

tion, especially as the classical computing power at the disposal of researchers continues to increase.

In the future, our genetic algorithm can be applied to find near-optimal circuits for the state preparation task on NISQ computers. It would be interesting to search for optimized circuits for the preparation of some practically more relevant states than Haar random states, and see if the improvements can be as significant as for random states. Even though running the genetic algorithm is computationally expensive, it may still be beneficial to optimize the state preparation circuit to save the scarce quantum computational resources on a NISQ machine, especially if the particular state needs to be prepared many times during the desired algorithm. However, for a larger quantum processor the output state fidelity evaluation step will need to be implemented on the quantum machine itself, since the classical simulation becomes prohibitively expensive. Indeed, this seems to be a key challenge to overcome in quantum circuit compilation in general.

Finally, the idea of automated quantum circuit discovery is very appealing, and genetic algorithms seem to be well suited for the task. However, it would be great if we could somehow generalize the search from individual circuits, as done in this work, to families of circuits with different numbers of qubits, perhaps defined in some iterative manner. In this way it would be possible to search for general solutions to problems with different input sizes, as is often required. This might make it possible to numerically explore more directly the scaling of circuit depth as a function of the input size – a key relation in determining the computational complexity of a problem – and thus shed more light on the computational capabilities of quantum computers.

CRediT authorship contribution statement

Tom Rindell: Writing – original draft, Visualization, Validation, Software, Methodology, Investigation, Data curation. **Berat Yenilen:** Visualization, Software, Methodology, Investigation, Data curation. **Niklas Halonen:** Visualization, Software, Methodology, Investigation, Data curation. **Arttu Pönni:** Software, Methodology, Investigation, Data curation. **Ilkka Tittonen:** Supervision, Resources, Methodology, Investigation, Funding acquisition, Conceptualization. **Matti Raasakka:** Writing – review & editing, Writing – original draft, Supervision, Software, Resources, Project administration, Methodology, Investigation, Funding acquisition, Formal analysis, Conceptualization.

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: Tom Rindell reports financial support was provided by Business Finland. Arttu Pönni reports financial support was provided by Business Finland.

Data availability

Link to the code can be found in the manuscript.

Acknowledgements

We are grateful to Business Finland for supporting the ‘Post-Quantum Cryptography (PQC) Finland’ project (project number 9154/31/2019) [45], which partially funded this research. We acknowledge the computational resources provided by the Aalto Science-IT project, in particular, the Triton cluster. The authors also wish to acknowledge CSC – IT Center for Science, Finland, for computational resources.

Appendix A. Error estimate for approximate state preparation circuits

We may estimate the dependence of the noise-free error on the CNOT count by considering the fraction of the total volume of the Hilbert space, which can be reached with error ϵ by using at most l CNOT-gates. The Hilbert space \mathcal{H}_n of n qubits is geometrically the $2(2^n - 1)$ -dimensional complex projective space \mathbb{CP}^{2^n-1} . (See, e.g., [46] for a review of complex projective geometry.) A complex projective space comes with a canonical metric, the Fubini-Study metric.⁵ The volume of a $2d$ -dimensional complex projective space, as defined by this metric, is given by

$$\text{Vol}(\mathbb{CP}^d) = \frac{\pi^d}{d!}.$$

Accordingly, the total volume of \mathcal{H}_n is given by $\text{Vol}(\mathcal{H}_n) = \pi^{2^n-1}/(2^n-1)!$.

⁵ We normalize the metric so that the geodesic length $D_{FS}(|\phi\rangle, |\psi\rangle) = \arccos(|\langle\phi|\psi\rangle|)$ between any two (normalized) state vectors $|\phi\rangle, |\psi\rangle \in \mathbb{CP}^d$. With this normalization, any two orthogonal states have distance $\frac{\pi}{2}$, and the unique closed geodesic passing through those states has circumference π .

The subspace we reach from the initial state $|00\dots 0\rangle$ without any CNOT gates is a tensor product $\mathcal{H}_1^{\otimes n}$ of single-qubit Hilbert spaces \mathcal{H}_1 , each of which is isomorphic to a 2-dimensional complex projective space of volume $\text{Vol}(\mathbb{CP}_1) = \pi$. Thus, we get for the volume of the $2n$ -dimensional subspace of states reachable by single-qubit unitaries $\text{Vol}(\mathcal{H}_1^{\otimes n}) = \pi^n$. Now, imagine ‘thickening’ this subspace by distance ϵ into the total Hilbert space. For $\epsilon \ll 1$ we can approximate the thickened subspace as the Cartesian product $\mathbb{CP}_1^{\times n} \times B_{d_\perp}(\epsilon)$, where $B_d(R)$ denotes the d -dimensional ball of radius R , and $d_\perp = 2(2^n - n - 1)$ is the codimension of the subspace $\mathcal{H}_1^{\otimes n}$ in \mathcal{H}_n . The volume of a d -dimensional ball is given by

$$\text{Vol}(B_d(R)) = \frac{\pi^{\frac{d}{2}}}{\Gamma(\frac{d}{2} + 1)} R^d.$$

Accordingly, the volume of the thickened subspace is given by

$$\begin{aligned} \text{Vol}(\mathbb{CP}_1^{\times n} \times B_{d_\perp}(\epsilon)) &= \text{Vol}(\mathbb{CP}_1)^n \text{Vol}(B_{d_\perp}(\epsilon)) \\ &= \frac{\pi^{2^n - 1}}{(2^n - n - 1)!} \epsilon^{2(2^n - n - 1)}. \end{aligned}$$

Without restrictions on the connectivity of the qubits, a single CNOT gate can be applied to any of the $\frac{n(n-1)}{2}$ pairs of qubits. Each allowed CNOT gate in a circuit increases the dimension of the subspace of states that can be reached without an error: After each CNOT gate two more single-qubit rotations can be added, which lead to distinct states. However, z -rotations of the control qubit and x -rotations of the target qubit commute with the CNOT gate, reducing the number of free parameters for each one of the rotations from 3 to 2. Each space of single-qubit rotations after a CNOT is isomorphic to a 2-sphere $S_2(1) \cong \mathbb{CP}_1$, so we get an extra direct product factor of \mathcal{H}_1^2 for our subspace from each CNOT gate. Accordingly, with l CNOT gates the volume of the subspace of states reached without an error is given by $\text{Vol}(\mathcal{H}_1^{\otimes(n+2l)}) = \pi^{n+2l}$. Allowing for error ϵ now gives volume

$$\text{Vol}(\mathbb{CP}_1^{\times(n+2l)} \times B_{d_\perp}(\epsilon)) = \frac{\pi^{2^n - 1} \epsilon^{2(2^n - (n+2l) - 1)}}{(2^n - (n+2l) - 1)!}.$$

In the leading order, the number of different circuits containing l CNOT gates is given by $(n(n-1)/2)^l$. Some of these circuits are actually equivalent due to commuting CNOT gates and the saturation of the Hilbert space, but this only affects the sub-leading order behavior. Circuits with larger l also reproduce the subspaces of states given by circuits with smaller l as we may cancel out CNOT gates from the beginning of the circuit by setting the preceding single-qubit unitaries to identity, given that the initial state of the circuit is $|00\dots 0\rangle$. Assuming for simplicity that each different combination of CNOT gates leads to a distinct subspace of states,⁶ we reach a volume

$$\begin{aligned} \text{Vol}(\mathbb{CP}_1^{\times(n+2l)} \times B_{d_\perp}(\epsilon)) &\left(\frac{n(n-1)}{2}\right)^l \\ &= \frac{\pi^{2^n - 1} \epsilon^{2(2^n - (n+2l) - 1)}}{(2^n - (n+2l) - 1)!} \left(\frac{n(n-1)}{2}\right)^l \end{aligned}$$

with circuits containing at most l CNOT gates. Dividing by the total volume of the Hilbert space, we get for the fraction of the volume covered by the subspace of states reached within an error ϵ

$$\begin{aligned} \eta_{n,l}(\epsilon) &= \frac{\text{Vol}(\mathbb{CP}_1^{\times(n+2l)} \times B_{d_\perp}(\epsilon))}{\text{Vol}(\mathcal{H}_n)} \left(\frac{n(n-1)}{2}\right)^l \\ &= \frac{(2^n - 1)! \epsilon^{2(2^n - (n+2l) - 1)}}{(2^n - (n+2l) - 1)!} \left(\frac{n(n-1)}{2}\right)^l. \end{aligned}$$

Using the approximation $\Gamma(x + \alpha) \approx \Gamma(x)x^\alpha$ for $x \gg 1$, for large n we may approximate

$$\frac{(2^n - 1)!}{(2^n - (n+2l) - 1)!} \approx (2^n)^{n+2l} = 2^{n(n+2l)}.$$

We then get

$$\begin{aligned} \eta_{n,l}(\epsilon) &\approx 2^{n(n+2l)} \epsilon^{2(2^n - (n+2l) - 1)} \left(\frac{n(n-1)}{2}\right)^l \\ &= 2^{n^2} \epsilon^{2(2^n - n - 1)} \left(\frac{4^n n(n-1)}{\epsilon^4}\right)^l. \end{aligned}$$

We may then obtain at least a crude estimate for the error ϵ , which is needed to cover the full Hilbert space by setting $\eta_{n,l}(\epsilon) = 1$ and solving for ϵ . This gives

$$\epsilon \approx \exp \left[-\frac{(2 \ln 2)n + \ln \left(\frac{n(n-1)}{2}\right)}{2(2^n - n - 1)} l - \frac{(\ln 2)n^2}{2(2^n - n - 1)} \right].$$

In the above analysis we did not yet consider the possible restrictions on the connectivity of qubits. We may estimate the effect of restricted qubit connectivity by dividing l , the number of logical CNOT gates in the circuit, by the average number l_{ph} of physical CNOT gates required to implement one logical CNOT gate. Thus, we arrive at an estimate for the maximum error with l CNOT gates:

$$\epsilon \approx \exp \left[-\frac{(2 \ln 2)n + \ln \left(\frac{n(n-1)}{2}\right)}{2(2^n - n - 1)} \frac{l}{l_{\text{ph}}} - \frac{(\ln 2)n^2}{2(2^n - n - 1)} \right].$$

A linear scaling of the number of CNOT gates is obviously a very simplistic way to take the connectivity into account. However, we know that different connectivities differ only polynomially, so the overall exponential behavior of ϵ must be preserved. We also get a decent fit to the numerical results using this approach, as we see in Section 4.

Appendix B. Asymptotic analysis of the fidelity formula

Let us consider the asymptotic scaling of l_* according to the formula

$$\frac{l_*}{l_{\text{ph}}} = \frac{1}{A(n)} \ln \left(1 - \frac{A(n)}{\ln(1-p)} \right) - \frac{B(n)}{A(n)},$$

where

$$A(n) = \frac{(2 \ln 2)n + \ln \left(\frac{n(n-1)}{2}\right)}{2^n - n - 1}, \quad B(n) = \frac{(\ln 2)n^2}{2^n - n - 1}$$

as we increase the number of qubits n . We have

$$x \ln \left(1 - \frac{1}{cx} \right) \rightarrow -\frac{1}{c}$$

as $x \rightarrow \infty$, so the first term converges to the constant $\left(\ln \frac{1}{1-p}\right)^{-1}$ as $n \rightarrow \infty$, since $A(n) \rightarrow 0$ in the limit. The second term

$$\frac{B(n)}{A(n)} = \frac{(\ln 2)n^2}{(2 \ln 2)n + \ln \left(\frac{n(n-1)}{2}\right)} \approx \frac{n}{2}$$

⁶ This assumption will presumably hold approximately true for $\epsilon \ll 1$, until we reach $\sim 2^n$ CNOT gates, which is required for the exact preparation of an arbitrary state.

when $n \gg 1$. Thus, for large n , we have the approximation

$$\frac{l_*}{l_{\text{ph}}} \approx \left(\ln \frac{1}{1-p} \right)^{-1} - \frac{n}{2}.$$

First of all, to get any benefit from the computation and reach a better average fidelity than with the initial state $|00 \dots 0\rangle$, we must have $l_* > 0$, which gives the condition $p < 1 - e^{-2/n} \approx \frac{2}{n}$ for $n \gg 1$. Thus, the error rate of CNOT gates must decrease inversely proportionally to the number of qubits to have any chance of preparing an arbitrary target state to any degree of approximation. On the other hand, counter-intuitively the optimum length l_* of a state preparation circuit actually *decreases* linearly at fixed error rate as the number of qubits is increased. This occurs because it gets exponentially hard to get close to an arbitrary state. It doesn't pay off to try that hard with noisy gates, because the errors build up as we try to increase the length of the circuit to get closer to the target state. Thus, we are better off with a shorter circuit, which at least gives somewhat a decent approximation, rather than trying to climb up the exponentially long ladder with faulty rungs towards the target state.

When $p \ll 1$, we may further approximate $\ln \frac{1}{1-p} \approx p$. Assuming again that $n \gg 1$, the maximum value of total fidelity is found to be approximated by

$$\begin{aligned} F(\rho_{\text{noise}}, |\text{target}\rangle)_{\text{max}} &\approx (1-p)^{l_*} \left(1 - \exp \left[-A(n) \frac{l_*}{l_{\text{ph}}} - B(n) \right] \right) \\ &\approx (1-p)^{\left(\frac{1}{p} - \frac{n}{2}\right)l_{\text{ph}}} \left(1 - \frac{2^n \ln(1-p)}{(2 \ln 2)n} \right)^{-1}. \end{aligned}$$

In order to keep the maximum total fidelity sufficiently high, both of the factors in this expression must remain large enough as n grows. In order to keep the first term constant, $(1-p)^{\left(\frac{1}{p} - \frac{n}{2}\right)l_{\text{ph}}} = f_1$, we find that the error rate must satisfy $p \approx 2(1 + \frac{\ln f_1}{l_{\text{ph}}})/n$, which again implies inverse proportionality of the error rate to the number of qubits. However, the second factor imposes a stricter requirement. In order for it to stay constant,

$$\left(1 - \frac{2^n \ln(1-p)}{(2 \ln 2)n} \right)^{-1} = f_2,$$

the error rate must satisfy

$$p \approx (\ln 2)(f_2^{-1} - 1)2^{-n + \log_2 n}.$$

We find a nearly inverse exponential scaling requirement for the error rate in the number of qubits, which reflects the general difficulty of the arbitrary state preparation task in the presence of noise.

References

- [1] D. Deutsch, R. Jozsa, Rapid solution of problems by quantum computation, *Proc. R. Soc. Lond. A* 439 (1992) 553–558, <https://doi.org/10.1098/rspa.1992.0167>.
- [2] D. Simon, On the power of quantum computation, *SIAM J. Comput.* 26 (1997) 1474–1483, <https://doi.org/10.1109/SFCS.1994.365701>.
- [3] P. Shor, Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer, *SIAM J. Comput.* 26 (1997) 1484–1509, <https://doi.org/10.1137/s0097539795293172>, arXiv:quant-ph/9508027.
- [4] L. Grover, A fast quantum mechanical algorithm for database search, in: *Annual ACM Symposium on Theory of Computing*, ACM, 1996, pp. 212–219, arXiv:quant-ph/9605043.
- [5] F. Arute, K. Arya, R. Babbush, et al., Quantum supremacy using a programmable superconducting processor, *Nature* 574 (2019) 505–510, <https://doi.org/10.1038/s41586-019-1666-5>, arXiv:1910.11333.
- [6] H.-S. Zhong, H. Wang, Y.-H. Deng, et al., Quantum computational advantage using photons, *Science* 370 (2020) 1460–1463, <https://doi.org/10.1126/science.abe8770>, arXiv:2012.01625.
- [7] J. Preskill, Quantum computing in the NISQ era and beyond, *Quantum* 2 (2020) 79, <https://doi.org/10.22331/q-2018-08-06-79>, arXiv:1801.00862.
- [8] L. Spector, H. Barnum, H. Bernstein, Genetic programming for quantum computers, in: J. Koza, W. Banzhaf, K. Chellapilla, et al. (Eds.), *Genetic Programming 1998: Proceedings of the Third Annual Conference*, 1998, pp. 365–374.
- [9] C. Williams, A. Gray, Automated design of quantum circuits, in: C. Williams (Ed.), *Lecture Notes in Computer Science*, vol. 1509, Springer, 1999, pp. 113–125.
- [10] M. Lukac, M. Perkowski, Evolving quantum circuits using genetic algorithm, in: *Proceedings 2002 NASA/DoD Conference on Evolvable Hardware*, 2002, pp. 177–185.
- [11] L. Spector, *Automatic Quantum Computer Programming: A Genetic Programming Approach*, Kluwer, 2004.
- [12] P. Massey, J. Clark, S. Stepney, Human-competitive evolution of quantum computing artefacts by genetic programming, *Evol. Comput.* 14 (2006) 21–40, <https://doi.org/10.1162/evco.2006.14.1.21>.
- [13] C. Ruican, M. Udrescu, L. Prodan, M. Vladutiu, Automatic synthesis for quantum circuits using genetic algorithms, in: B. Beliczynski, A. Dzieliński, M. Iwanowski, B. Ribeiro (Eds.), *Lecture Notes in Computer Science*, vol. 4431, Springer, 2007, pp. 174–183.
- [14] S. Altares-López, A. Ribeiro, J. García-Ripoll, Automatic design of quantum feature maps, *Quantum Sci. Technol.* 6 (2021) 045015, <https://doi.org/10.1088/2058-9565/ac1ab1>, arXiv:2105.12626.
- [15] Z. Lu, P.-X. Shen, D.-L. Deng, Markovian quantum neuroevolution for machine learning, *Phys. Rev. Appl.* 16 (2021) 044039, <https://doi.org/10.1103/PhysRevApplied.16.044039>, arXiv:2012.15131.
- [16] L. Arufe, M. Gonzalez, A. Oddi, R. Rasconi, R. Varella, Quantum circuit compilation by genetic algorithm for quantum approximate optimization algorithm applied to MaxCut problem, *Swarm Evol. Comput.* 69 (2022) 101030, <https://doi.org/10.1016/j.swevo.2022.101030>.
- [17] D. Chivilikhin, A. Samarin, V. Ulyantsev, I. Iorsh, A.R. Oganov, O. Kyriienko, Mogvqe: multiobjective genetic variational quantum eigensolver, arXiv:2007.04424, 2020.
- [18] L. Sünkel, D. Martyniuk, D. Mattern, J. Jung, A. Paschke, GA4QCO: genetic algorithm for quantum circuit optimization, arXiv:2302.01303, 2023.
- [19] V. Potoček, A. Reynolds, A. Fedrizzi, D. Corne, Multi-objective evolutionary algorithms for quantum circuit discovery, arXiv:1812.04458, 2018.
- [20] V. Shende, S. Bullock, I. Markov, Synthesis of quantum-logic circuits, *IEEE Trans. Comput.-Aided Des. Integr. Circuits Syst.* 25 (2006) 1000–1010, <https://doi.org/10.1109/tcad.2005.855930>, arXiv:quant-ph/0406176.
- [21] M. Plesch, C. Brukner, Quantum-state preparation with universal gate decompositions, *Phys. Rev. A* 83 (2011) 032302, <https://doi.org/10.1103/physreva.83.032302>, arXiv:1003.5760.
- [22] A. Chi-Chih Yao, Quantum circuit complexity, in: *Proceedings of 1993 IEEE 34th Annual Foundations of Computer Science*, 1993, pp. 352–361.
- [23] J. Watrous, Quantum computational complexity, in: R. Meyers (Ed.), *Encyclopedia of Complexity and Systems Science*, Springer, 2009, pp. 7174–7201, arXiv:0804.3401.
- [24] F. Liu, S. Whitsitt, J. Curtis, R. Lundgren, P. Titum, Z.-C. Yang, J. Garrison, A. Gorshkov, Circuit complexity across a topological phase transition, *Phys. Rev. Res.* 2 (2020) 013323, <https://doi.org/10.1103/PhysRevResearch.2.013323>, arXiv:1902.10720.
- [25] Z. Xiong, D.-X. Yao, Z. Yan, Nonanalyticity of circuit complexity across topological phase transitions, *Phys. Rev. B* 101 (2020) 174305, <https://doi.org/10.1103/PhysRevB.101.174305>, arXiv:1906.11279.
- [26] S. Roca-Jerat, T. Sancho-Lorente, J. Román-Roche, D. Zueco, Circuit complexity through phase transitions: consequences in quantum state preparation, arXiv:2301.04671, 2023.
- [27] L. Susskind, Computational complexity and black hole horizons, arXiv:1402.5674, 2014.
- [28] J. Couch, S. Eccles, T. Jacobson, P. Nguyen, Holographic complexity and volume, *J. High Energy Phys.* 2018 (2018) 1–39, [https://doi.org/10.1007/JHEP11\(2018\)044](https://doi.org/10.1007/JHEP11(2018)044), arXiv:1807.02186.
- [29] A. Brown, D. Roberts, L. Susskind, B. Swingle, Y. Zhao, Holographic complexity equals bulk action?, *Phys. Rev. Lett.* 116 (2016) 191301, <https://doi.org/10.1103/PhysRevLett.116.191301>, arXiv:1509.07876.
- [30] I. Araujo, C. Blank, A. da Silva, Approximated quantum-state preparation with entanglement dependent complexity, arXiv:2111.03132, 2021.
- [31] M. Möttönen, J. Vartiainen, V. Bergholm, M. Salomaa, Transformation of quantum states using uniformly controlled rotations, *Quantum Inf. Comput.* 5 (2005) 467–473, arXiv:quant-ph/0407010.
- [32] V. Bergholm, J. Vartiainen, M. Möttönen, M. Salomaa, Quantum circuits with uniformly controlled one-qubit gates, *Phys. Rev. A* 71 (2005) 052330, <https://doi.org/10.1103/physreva.71.052330>, arXiv:quant-ph/0410066.
- [33] R. Iten, R. Colbeck, I. Kukuljan, J. Home, M. Christandl, Quantum circuits for isometries, *Phys. Rev. A* 93 (2016) 032318, <https://doi.org/10.1103/physreva.93.032318>, arXiv:1501.06911.
- [34] X.-M. Zhang, M.-H. Yung, X. Yuan, Low-depth quantum state preparation, *Phys. Rev. Res.* 3 (2021) 043200, <https://doi.org/10.1103/PhysRevResearch.3.043200>, arXiv:2102.07533.
- [35] X.-M. Zhang, T. Li, X. Yuan, Quantum state preparation with optimal circuit depth: implementations and applications, *Phys. Rev. Lett.* 129 (2022) 230504, <https://doi.org/10.1103/PhysRevLett.129.230504>, arXiv:2201.11495.

- [36] M. Davis, E. Smith, A. Tudor, K. Sen, I. Siddiqi, C. Iancu, Heuristics for quantum compiling with a continuous gate set, arXiv:1912.02727, 2019.
- [37] D. Camps, R. Van Beeumen, Approximate quantum circuit synthesis using block encodings, Phys. Rev. A 102 (2020) 052411, <https://doi.org/10.1103/physreva.102.052411>, arXiv:2007.01417.
- [38] E. Younis, K. Sen, K. Yelick, C. Iancu, Qfast: quantum synthesis using a hierarchical continuous circuit space, arXiv:2003.04462, 2020.
- [39] S. Ashhab, N. Yamamoto, F. Yoshihara, K. Semba, Numerical analysis of quantum circuits for state preparation and unitary operator synthesis, Phys. Rev. A 106 (2022) 022426, <https://doi.org/10.1103/physreva.106.022426>, arXiv:2204.13524.
- [40] A. Konak, D. Coit, A. Smith, Multi-objective optimization using genetic algorithms: a tutorial, J. RESS 91 (2006) 992–1007, <https://doi.org/10.1016/j.ress.2005.11.018>.
- [41] K. Deb, A. Pratap, S. Agarwal, T. Meyarivan, A fast and elitist multiobjective genetic algorithm: NSGA-II, IEEE Trans. Evol. Comput. 6 (2002) 182–197, <https://doi.org/10.1109/4235.996017>.
- [42] F.-A. Fortin, F.-M. De Rainville, M.-. Gardner, M. Parizeau, C. Gagné, DEAP: evolutionary algorithms made easy, J. Mach. Learn. Res. 13 (2012) 2171, <https://pypi.org/project/deap/>.
- [43] T. Rindell, B. Yenilen, N. Halonen, A. Pönni, I. Tittonen, M. Raasakka, <https://github.com/beratyenilen/qc-ga>.
- [44] Triton cluster, <https://scicomp.aalto.fi/triton/>.
- [45] 'Post-Quantum Cryptography (PQC) Finland' project, <https://www.pqc.fi/>.
- [46] I. Bengtsson, K. Życzkowski, Geometry of Quantum States, 2nd ed., Cambridge University Press, 2017.