

Utilizing Resource Estimation for the Development of Quantum Computing Applications

Nils Quetschlich*

Mathias Soeken†

Prakash Murali‡

Robert Wille*§

*Chair for Design Automation, Technical University of Munich, Germany

†Microsoft Quantum, Switzerland

‡University of Cambridge, United Kingdom

§Software Competence Center Hagenberg GmbH (SCCH), Austria

nils.quetschlich@tum.de

mathias.soeken@microsoft.com

pm830@cam.ac.uk

robert.wille@tum.de

<https://www.cda.cit.tum.de/research/quantum>

<https://quantum.microsoft.com>

Abstract—Quantum computing has made considerable progress in recent years in both software and hardware. But to unlock the power of quantum computers in solving problems that cannot be efficiently solved classically, quantum computing at scale is necessary. Unfortunately, quantum simulators suffer from their exponential complexity and, at the same time, the currently available quantum computing hardware is still rather limited (even if roadmaps make intriguing promises). Hence, in order to evaluate quantum computing applications, end-users are still frequently restricted to toy-size problem instances (which additionally often do not take error correction into account). This substantially hinders the development and assessment of real-world quantum computing applications. In this work, we demonstrate how to utilize *Resource Estimation* to improve this situation. We show how the current workflow (relying on simulation and/or execution) can be complemented with an estimation step, allowing that end-users (1) actually can consider real-world problem instances already today (also considering error correction schemes and correspondingly required hardware resources), (2) can start exploring possible optimizations of those instances across the entire design space, and (3) can incorporate hypotheses of hardware development trends to derive more informed and, thus, better design space parameters. Overall, this enables end-users already today to check out the promises of possible future quantum computing applications, even if the corresponding hardware to execute them is not available yet.

I. INTRODUCTION

In recent years, quantum computing has witnessed remarkable advancements in both software and hardware. As a result, an expanding range of quantum devices with progressively enhanced qubit quality has become available, leading to increased interest in academia and industry for tackling diverse problems across multiple application domains. This trend is indicative of the growing potential of quantum computing to offer transformative solutions to complex problems that are beyond the capabilities of classical computing [1].

The utilization of quantum computers to solve complex problems involves a multi-stage process: The first step involves selecting or developing a quantum algorithm that offers a quantum advantage, i.e., can solve the problem *better* compared to the best-known classical algorithm (e.g., in terms of algorithmic complexity, execution time, or solution quality). The second step entails the encoding of the problem

in terms of a quantum program that can be compiled into machine instructions for a quantum computer. This process factors in various constraints on quantum gate sets and qubit connectivities that arise from different quantum computing technologies being explored, such as superconducting [2], ion traps [3], neutral atoms [4], or Majorana qubits [5]. The third step involves executing the compiled program on the quantum computer (or a corresponding simulator), and the final step covers the interpretation of the measurement results from the quantum computer as the solution to the original problem. Developing such solutions is an active area of research in domains such as finance [6], chemistry [7], machine learning [8], and optimization [9].

Unfortunately, this process fails as soon as practically relevant applications are considered. This is because the application of the described workflow currently relies on the use of quantum simulators and near-term quantum computers. The former deploy classical computers to simulate quantum algorithms—a computationally complex task which is limited to a few dozen qubits. The latter supports a larger number of qubits, but their scalability is limited as well. Importantly, the noise rates in these near-term quantum computers currently are still too high to support *high-fidelity* executions of practically relevant applications.

Consequently, a comprehensive evaluation of the potential of quantum computing for a considered problem necessitates the consideration of quantum error correction to scale to larger problem instances closer to real-world problem sizes. This induces an overhead in the required resources (such as, e.g., the number of qubits) which effectively exceeds the limits of current devices by far. All of that could lead to a situation in which certain quantum algorithms might not provide any practical quantum advantage once the presumed characteristics of the fault-tolerant quantum computer and the efficacy of error correction protocols are factored in [10]. Because of this, it is often still not clear what applications are suitable to be solved using quantum computing—constituting a major bottleneck in the progress of quantum computing application development.

Recently, a complementary approach to the execution on simulators or quantum computing hardware emerged: *Resource Estimation* (RE) such as, e.g., proposed in [10], [11].

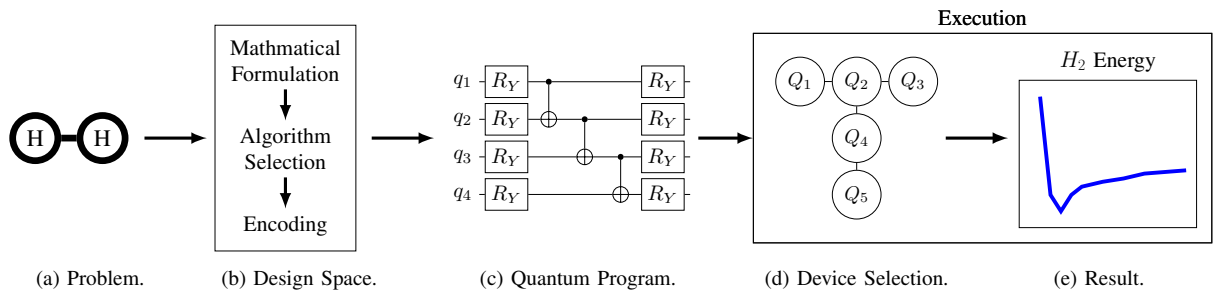


Fig. 1: Current workflow for the development of quantum computing applications.

Instead of actually executing a given quantum program, RE gives an estimate of the resources necessary to execute it in a fault-tolerant fashion. Although this approach does not return the actual solution, it allows one to determine an early resource estimate—a procedure that has been exploited in conventional computing for decades, e.g., in classical HW/SW co-design where cost estimates are used to guide which functionalities are realized in hardware and which in software (see, e.g., [12], [13]).

This paper illustrates how the current workflow to solve problems using quantum simulators and noisy near-term quantum computers can be complemented by substituting the actual execution step with an estimation step—leading to the following three improvements: Firstly, end-users are enabled to consider real-world problem instances already today (also considering error correction schemes to explore the hardware requirements of a chosen set of design space parameters and determine the required hardware resources). These estimates can then be compared with the quantum hardware vendor’s roadmaps to obtain a sense of the chosen approach’s feasibility. Secondly, end-users can start exploring optimizations across the entire design space of a quantum computing application and, by that, develop and facilitate trade-offs between different design space parameters. Thirdly, end-users can even incorporate various hypotheses of hardware development trends and their potential improvements into the methodology.

On the basis of that, more informed and, thus, better design space parameters are derived. Overall, the workflow complementation eventually provides valuable insights for the development of future quantum technologies and promising applications for them due to its versatility and applicability across different stages of quantum computing application development. In other words, this enables end-users already today to check out the promises of possible future quantum computing applications, even if the corresponding hardware to execute them is not available yet.

The remainder of this work is structured as follows: Section II describes the current workflow to solve problems using quantum computing. Section III outlines how resource estimation can complement the flow in the absence of large-scale quantum computers. Then, the complemented workflow is applied in a representative case study in Section IV based on a real-world problem instance from the domain of chemistry to demonstrate its advantages and is discussed in Section V. Section VI concludes this work.

II. MOTIVATION

In this section, the quantum solution workflow is reviewed and illustrated with a running example.

A. Quantum Solution Workflows

The workflow to solve a problem using quantum computing is illustrated in Fig. 1 (based on the workflow proposed in [14]). Starting with the problem itself as sketched in Fig. 1a, it must be translated into a form suitable for quantum computing as indicated in Fig. 1b. This comprises (1) the mathematical problem formulation, (2) the selection of a quantum algorithm that is generally capable of solving the problem considered, and (3) its encoding into a quantum program based on that—requiring multiple design decisions that eventually form the *design space*. Any combination of suitable design choices leads to a quantum program as shown in Fig. 1c.

Example 1. A prominent example application in the domain of chemistry is to calculate the ground state energy of a molecule such as the H_2 molecule illustrated in Fig. 1a. To mathematically describe the problem, its Hamiltonian must be derived. This is usually done with the help of already existing software tools, such as PySCF [15]. Subsequently, the resulting Hamiltonian must be mapped to a format suitable for quantum computing, e.g., using the Jordan-Wigner mapping. Next, the Variational Quantum Eigensolver (VQE, [7]) is selected as the solving algorithm since it is one of the prime algorithms for this kind of problem. For VQE, a respective ansatz must be chosen for the encoding as a quantum program—in this case, the TwoLocal ansatz¹ with a linear entanglement pattern is chosen, resulting in the quantum program being sketched in Fig. 1c.

Then, the resulting quantum program must be executed. To this end, both quantum simulators (based on different data structures [16] such as, e.g., decision diagrams [17]–[21], tensor networks [22], [23], matrix-product states [24], [25], or sparsity [26]) and near-term quantum computers (based on various technologies such as, e.g., superconducting [2], ion traps [3], neutral atoms [4], or Majorana qubits [5]) are available and a representative must be selected as shown in Fig. 1d. This representative must be capable of executing the quantum program—both in terms of the number of required

¹See <https://qiskit.org/documentation/stubs/qiskit.circuit.library.TwoLocal.html> for details.




Fig. 2: Current execution limits for quantum programs.

qubits and also on its qubits’ characteristics such as the gate execution error rates or decoherence time.

Example 2. *The quantum program derived in Example 1 requires four qubits. Therefore, the ibmq_perth device with 5 qubits is selected which has a limited connectivity as depicted in Fig. 1d.*

To eventually determine the desired solution to the initial problem, the encoded quantum program can now be executed on the chosen device. Afterwards, the result is extracted from the measurement results.

Example 3. *When executing the program resulting from Example 1 on the device chosen from Example 2, the ground state energy is approximated as insinuated in Fig. 1e. Determining the minimum value of this graph returns the desired value of the H_2 molecule.*

B. Current Limitations and General Idea

There are two options to execute quantum programs: Either the currently available quantum simulators or near-term quantum computers. However, both options are limited in their capabilities as illustrated in Fig. 2 and do not provide means for exploring practically relevant problems.

Quantum simulators provide ideal qubits but have a rather small capacity—usually limited to a few dozen qubits—while, on the other hand, the currently available so-called *Noisy Intermediate Scale Quantum* (NISQ) computers [27] with up to several hundreds of qubits do not provide a sufficient qubit quality for a reliable execution (and, thus, effectively restricting the quantum program depth). As a consequence, the development of quantum computing applications at the moment considers mostly toy-size problem instances whose quantum computing solutions still fit the current simulators and NISQ computers—not considering the scalability to larger problem sizes closer to real-world scenarios.

Most quantum device vendors have published roadmaps to significantly scale their device capabilities over the next years. However, it would be disadvantageous to wait for the availability of sufficiently large quantum devices before considering real-world problem sizes when developing quantum computing applications—leading to a situation where, in the worst case, powerful devices are available, but no suitable applications are available to make use of them.

C. Resource Estimation

Resource Estimation (RE) is a promising approach to overcome this bottleneck and to provide guidance on implementing important quantum algorithm instances today for the hardware of tomorrow. Instead of executing a given quantum program, RE estimates the required resources (such as the number of qubits and the runtime) based on assumed hardware characteristics such as gate execution times, fidelity rates, and, the underlying error correction scheme. This methodology allows one to consider quantum programs orders of magnitude larger than the current limits of both quantum simulators and computers. Although this obviously does not result in an actual execution and, hence, solution, it already gives insights into the applicability and scalability of the chosen design space parameters and assumed hardware characteristics.

RE is an emerging topic within the quantum computing community and, recently, various methodologies have been proposed. While some works (such as, e.g., [28]) focus on rather manual generation of resource estimates, more and more software tools are proposed to automate this procedure—e.g., Microsoft’s Azure Quantum Resource Estimator [11], [29], Google’s Qualtran [30], Zapata’s BenchQ [31], and MIT Lincoln Lab’s pyLIQTR [32]. These software tools provide simple means of estimating what resources are necessary to reliably execute a given quantum program based on assumed qubit hardware characteristics.


III. RESOURCE ESTIMATION-DRIVEN DEVELOPMENT

In this work, we describe how resource estimation complements the quantum solution workflow to allow the development of quantum computing applications for practically relevant problems already today, even when fault-tolerant quantum hardware is not available yet.

A. Resource Estimation for Scalability Exploration

Fig. 3 illustrates how we integrate resource estimation into the workflow by substituting the *Execution* step (indicated in Fig. 1d/Fig. 1e) by an *Estimation* step (indicated in Fig. 3d/Fig. 3e). More precisely, instead of returning the solution of the considered application instance, the complemented workflow returns an estimate of the required resources such as how many physical qubits would be needed for execution together with the estimated runtime.

Example 4. *Consider again the problem from Example 1 and its respective quantum program derived in Example 2. Although this program represents a small problem instance that is easily executable on current NISQ computers, VQE generally does not scale well with respect to the size of the problem considered compared to other approaches, e.g., based on Quantum Phase Estimation (QPE) [34]. Hence, larger problem instances such as the one depicted in Fig. 3a would lead to different design space parameter choices (sketched in Fig. 3b)—e.g., QPE-based algorithms—and, thus, result in a different quantum program as shown in Fig. 3c. However, those approaches assume a (close-to) error-free execution—requiring error correction schemes to overcome the erroneous*



(a) Problem (taken from [33]). (b) Design Space. (c) Quantum Program. (d) Resource Estimation. (e) Result.

Fig. 3: Resource estimation-driven development of quantum computing applications.

hardware and, thus, control the errors during computation in a way to stay within a provided error budget. Using RE, the number of required physical qubits and the corresponding runtime can be determined as depicted in Fig. 3d based on assumptions on the hardware characteristics of the qubits. This ultimately leads to an assessment of whether a practical quantum advantage is likely for the provided problem as indicated in Fig. 3e.

Following this approach gives an estimate of the required resources for not-yet-executable problem sizes. The resulting estimates of, e.g., required physical qubits and the runtime, might give an indication of the time horizon needed to solve those problem instances using actual quantum computers by comparing the required estimated resources with the roadmaps of device vendors. This does not give an actual problem result, but a reference whether the considered problem is worth considering in the foreseeable future—an information highly relevant for various stakeholders from application developers in industry up to decision makers in funding agencies and politics to, e.g., define their national quantum strategy.

B. Resource Estimation for Design Space Exploration

Often, problems can be solved using different quantum computing approaches. For example, there exist various algorithms for simulating quantum chemistry, based on Trotterization [35], qubitization [36], or tensor hypercontraction [37], to name a few. On top of that, there are usually various possibilities to implement the quantum algorithm as a quantum program—and these choices have a significant influence on the required resources. Therefore, a workflow that supports end-users in exploiting these design choices would certainly be helpful.

Using RE, the workflow described above can even be further complemented to provide this support for the design space exploration for a given problem by introducing a feedback loop as depicted in Fig. 3 from Fig. 3d to Fig. 3b. This loop allows end-users to determine the most promising set of design space parameters in a guided fashion. Finding such a set helps to assess the practicality of the execution of an algorithm instance on a future computer as sketched in Fig. 3e.

Example 5. Although the asymptotic computational complexity of quantum algorithms for simulating quantum chemistry can be compared, implementing them as quantum programs for a specific instance may introduce overhead that could alter the ranking of the algorithms and favor one over the other. Moreover, while a particular algorithm may be favorable for one set of instances, other algorithms may be preferable for other sets. When considering various implementations for building blocks used in the quantum program, such as quantum arithmetic or table lookup, the search space expands even further.

This approach requires an effort to generate the initial quantum program that solves the problem considered based on the selected design space parameters before evaluating its scalability. However, this is not always necessary and often it is sufficient to know how the selected quantum algorithm scales in terms of logical counts such as the required number of qubits and the number of gates with the problem size. Using these numbers is already sufficient to run a resource estimate—simplifying the design exploration process even further.

Following this approach provides powerful means to explore and evaluate the design space how best to solve the considered problem already today without having to wait for the availability of sufficiently large quantum devices. This early start of developing sophisticated quantum computing applications for real-world problem instances accelerates the race to catch up with classical solutions that have been developed and optimized for decades—a process that will take years, and, therefore, should be started as early as possible.

C. Resource Estimation for Hardware Characteristics Exploration

Finally, not only does the design space for how to encode an application into a quantum program offer a large degree of freedom, the assumed hardware characteristics used as input for the RE does as well. So far, this freedom has been utilized by current RE tools by providing pre-defined configurations (such as the Azure Quantum Resource Estimator [11] provides multiple configurations for both gate-based

TABLE I: Resource estimates.

	Distance	Factories	Phys. qubits	Runtime
G μ s, 10^{-3}	31	12	5.62M	156 years
G μ s, 10^{-4}	15	12	1.31M	75 years
G ns, 10^{-3}	31	14	6.04M	38 days
G ns, 10^{-4}	15	15	1.47M	18 days
M ns, 10^{-4}	17	13	4.20M	16 days
M ns, 10^{-6}	9	14	1.30M	8 days

and Majorana qubits) of assumed hardware characteristics. Additionally, these characteristics can be modified by end-users—and, hence, opens up the possibility of evaluating how further progress in hardware development might affect the resource estimates.

Example 6. *Similarly to how classical computers have evolved dramatically from early small-scale to current systems, quantum computing hardware is also expected to improve over time in terms of the number of qubits available, their error rates, decoherence times, and others. By employing RE, various hypotheses, such as an increase in gate execution speed or a decrease in gate execution error, can be examined and assessed.*

Evaluating different hypotheses on how the quantum hardware may develop significantly aids to create a deeper understanding of the scalability of the application considered and its design space parameters. Therefore, this evaluation should also be incorporated in the workflow of how quantum computing solutions are derived—leading to more informed and, thus, better design space parameters.

IV. CASE STUDY

In this section, we demonstrate various RE scenarios in order to showcase the benefits of the workflow described above for a practical quantum chemistry application.² We evaluate the resources to calculate the ground state energy of a Hamiltonian to chemical accuracy of 1 mHartree using the qubitization quantum simulation algorithm [36] on top of a double-factorized representation of the Hamiltonian [33]. The Hamiltonian describes the 64 electron and 56 orbital active space of one of the stable intermediates in the ruthenium-catalyzed carbon fixation cycle [38] shown in Fig. 3a.

A. Scalability Exploration

As a first experiment, we estimate the resources for six examples of qubit parameters which represent various regimes of interest [11, Table II] as shown in Table I. The label contains the operation times regime (either μ s or ns), and the limiting error rate of its Clifford operations, and is prefixed by whether the physical instruction sets are either gate-based (G) or have Majorana (M) instruction sets. To ensure a chemical accuracy of 1 mHartree, we provide an error budget of 1% to the resource estimator.

Next, we focus on the data point ‘M ns, 10^{-6} ’ as a base estimate and evaluate space/time trade-offs by allowing a

²The source code for this experiment can be found at <https://github.com/da-tum/mqt-problemsolver>.

TABLE II: Influence of T factories.

Factories	Distance	Fraction	Phys. qubits	Runtime
14	9	17.97%	1.30M	8 days
13	9	16.91%	1.28M	9 days
12	9	15.81%	1.26M	10 days
11	9	14.69%	1.25M	11 days
10	9	13.53%	1.23M	12 days
9	9	12.35%	1.21M	13 days
8	9	11.13%	1.20M	15 days
7	9	9.87%	1.18M	17 days
6	9	8.58%	1.16M	19 days
5	9	7.26%	1.15M	23 days
4	9	5.89%	1.13M	29 days
3	9	4.49%	1.11M	39 days
2	9	3.04%	1.10M	58 days
1	9	1.54%	1.08M	116 days

longer runtime of the algorithm. In the base estimate, 17.97% of the total number of physical qubits is used to run 14 T factories in parallel that will produce the required 270 billion T states over the runtime of the algorithm. By slowing down the execution of the algorithm using logical idle operations, we require fewer T factories to run in parallel to produce the same number of required T states. We analyze how the runtime and the number of overall physical qubits is affected by limiting the number of T factories to $F = 1, \dots, 14$. The results are shown in Table II.

Assuming that one month is still a reasonable runtime, we find that with 4 factories we are below that threshold but are able to save about 170,000 physical qubits compared to the base estimate.

Although the current quantum computing application development usually relies on available quantum simulators and NISQ computers, the proposed methodologies allow end-users to get at least a resource estimate—and, by that, an assessment of its practicality—for their problem instances closer to real-world scenarios.

B. Design Space Exploration

Next, we make hypothetical assumptions on what impact optimizations to space (number of qubits) and time (number of operations) on the logical abstraction layer may have—potentially caused by different design space parameters. For this purpose, we first extract the logical pre-layout counts from the previous RE results. These are independent of the qubit parameters and space and time constraints, and are: 1,318 qubits, 67,474,931,068 Toffoli gates, 96 T gates, 11,987,084 rotation gates in depth of 11,986,482, as well as 63,472,407,520 single-qubit measurements. In addition to the baseline (using the data point ‘M ns, 10^{-6} ’) we consider three scenarios: 1) $\frac{1}{2}/2$: use half the number of qubits, but all gate and depth counts are double, 2) $2/\frac{1}{2}$: use twice the number of qubits, but all gate counts and depths are halved, 3) $\frac{3}{4}/\frac{3}{4}$: save 25% both in qubit counts, gate counts, and depth. The respective resulting estimates are shown in Table III.

As we can see, hypothetical savings do not fully propagate to the physical estimates, which emphasizes the importance of full-stack physical resource estimates. For the $\frac{1}{2}/2$, the physical qubits did not decrease by 50%, even though the

TABLE III: Influence of different design space parameters.

Scenario	Distance	Factories	Phys. qubits	Runtime
baseline	9	14	1.30M	8 days
$\frac{1}{2}/2$	9	14	773.06k	17 days
$2/\frac{1}{2}$	9	14	2.34M	4 days
$\frac{3}{4}/\frac{3}{4}$	7	18	804.11k	5 days

runtime increased by more than twice. On the other hand, for the $2/\frac{1}{2}$, the number of physical qubits did not increase by 2, even though the logical number of qubits is twice as high as the baseline. Finally, in the optimistic scenario, both space and time savings are possible. Since the savings allow for decreasing the code distance from 9 to 7, final savings of more than 25% to physical qubits and runtime are possible.

Even without the availability of sufficiently large quantum devices, different design space parameters can already be explored and their influences can be quantified and assessed. Consequently, this allows end-users to keep improving their quantum computing applications such that optimized solutions are at hand once the corresponding quantum computing devices become available.

C. Hardware Characteristics Exploration

In the final experiment, we want to model how a potential change in error rates would affect the resulting resource estimates. To this end, we start with the data point ‘G ns, 10^{-3} ’ for a base estimate. Here, the assumptions for the qubit parameters are $t_{\text{meas}}^{(0)} = 100$ ns measurement time, and $t_{\text{gate}}^{(0)} = 50$ ns gate operation times, as well $p_T^{(0)} = 10^{-3}$ measurement and gate error rates [11]. We assume these parameters at some time point t in the future, and model an improvement of $t_{\text{meas}}^{(t)} = 0.9^t \cdot t_{\text{meas}}^{(0)}$, $t_{\text{gate}}^{(t)} = 0.9^t \cdot t_{\text{gate}}^{(0)}$, and $p_T^{(t)} = 0.1^t \cdot p_T^{(0)}$ after t equally spaced time steps. These constants were arbitrarily chosen for the sake of the example and the resulting estimates are shown in Table IV. The experiments can be repeated using more accurate numbers from observing the rate of past improvements.

Using RE, the influence of such (rather simple) assumptions on the progress in quantum computing hardware can easily be determined for a given quantum program that realizes any application. This provides end-users with powerful means to explore the design space for their problems—leading to more informed and, therefore, more scalable solutions.

V. DISCUSSION

Although RE does not provide a solution to the problem considered, it can be used in a complementary fashion to the workflow of developing quantum computing applications. RE is a powerful tool for exploring the scalability to real-world problem sizes under assumed hardware characteristics and error correction schemes—and, optionally, even considering hypotheses of how the hardware might improve. By that, the most promising design space parameters can be determined although sufficiently large hardware is not yet available. However, the approach is based on various assumptions, such as the assumed hardware characteristics and the methodology used

TABLE IV: Influence of different hardware characteristics.

t	Distance	Factories	Phys. qubits	Runtime
0	31	14	6.04M	38 days
1	15	15	1.47M	17 days
2	11	13	733.28k	11 days
3	7	14	295.96k	6 days
4	7	10	278.52k	6 days

to actually derive the resource estimates. Due to the currently missing hardware on the required scale, it is impossible to validate those assumptions in an experimental evaluation to prove the methodology.

VI. CONCLUSIONS

Many of the current design flows for developing quantum computing applications focus on simulation and execution on near-term devices, which prohibits the exploration of large-scale quantum programs for quantum computing at scale. This leads to the situation where quantum computing applications focus on toy-size problem instances without taking error correction into account. In this paper, we describe how *Resource Estimation* can readily complement such design flows. Instead of execution a given quantum program, its required hardware resources are estimated. Since this procedure does not rely on existing quantum simulators and computers, it is *not* restricted to toy-size problem instances—resulting in three improvements: (1) End-users are enabled to consider real-world problem instances already today (also taking error correction schemes into account and determining the required hardware resources). These estimates can then be compared with the quantum hardware vendor’s roadmaps to obtain a sense of the chosen approach’s feasibility. (2) End-users can start exploring optimizations across the entire *design space* of a quantum computing application to determine the most efficient parameters. (3) End-users can even incorporate various hypotheses of hardware development trends to derive more informed and, thus, better design space parameters. Overall, the described workflow enables the development of quantum computing applications for real-world problem instances already today without having to wait for the availability of sufficiently large quantum devices.

ACKNOWLEDGMENTS

N.Q. and R.W. acknowledge funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation program (DA QC, grant agreement No. 101001318 and MILLENION, grant agreement No. 101114305), the Munich Quantum Valley, which is supported by the Bavarian state government with funds from the High-tech Agenda Bayern Plus, and the BMWK on the basis of a decision by the German Bundestag through project QuaST, as well as the BMK, BMDW, the State of Upper Austria in the frame of the COMET program, and the QuantumReady project within Quantum Austria (managed by the FFG).

REFERENCES

- [1] T. L. Scholten *et al.*, “Assessing the benefits and risks of quantum computers,” 2024. arXiv: 2401.16317.
- [2] M. Kjaergaard *et al.*, “Superconducting qubits: Current state of play,” *Annual Review of Condensed Matter Physics*, 2020.
- [3] C. D. Bruzewicz, J. Chiaverini, R. McConnell, and J. M. Sage, “Trapped-ion quantum computing: Progress and challenges,” *Applied Physics Reviews*, 2019.
- [4] L. Henriët *et al.*, “Quantum computing with neutral atoms,” *Quantum*, 2020.
- [5] T. Karzig *et al.*, “Scalable designs for quasiparticle-poisoning-protected topological quantum computation with Majorana zero modes,” *Phys. Rev. B*, 2017.
- [6] N. Stamatopoulos *et al.*, “Option pricing using quantum computers,” *Quantum*, 2020.
- [7] A. Peruzzo *et al.*, “A variational eigenvalue solver on a photonic quantum processor,” *Nature Communications*, 2014.
- [8] C. Zoufal, A. Lucchi, and S. Woerner, “Quantum Generative Adversarial Networks for learning and loading random distributions,” *npj Quantum Information*, 2019.
- [9] S. Harwood, C. Gambella, D. Tenev, A. Simonetto, D. Bernal Neira, and D. Greenberg, “Formulating and Solving Routing Problems on Quantum Computers,” *IEEE Transactions on Quantum Engineering*, 2021.
- [10] T. Hoefler, T. Haener, and M. Troyer, “Disentangling hype from practicality: On realistically achieving quantum advantage,” 2023. arXiv: 2307.00523.
- [11] M. E. Beverland *et al.*, “Assessing requirements to scale to practical quantum advantage,” 2022. arXiv: 2211.07629.
- [12] Y.-K. Choi, Y. Chi, J. Wang, and J. Cong, “FLASH: Fast, parallel, and accurate simulator for HLS,” *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 2020.
- [13] L. Servadei *et al.*, “Accurate Cost Estimation of Memory Systems Utilizing Machine Learning and Solutions from Computer Vision for Design Automation,” *IEEE Transactions on Computers (TC)*, 2020.
- [14] N. Quetschlich, L. Burgholzer, and R. Wille, “Towards an Automated Framework for Realizing Quantum Computing Solutions,” in *Int’l Symp. on Multi-Valued Logic*, 2023.
- [15] Q. Sun *et al.*, “The python-based simulations of chemistry framework (PySCF),” 2017. arXiv: 1701.08223.
- [16] R. Wille, L. Burgholzer, S. Hillmich, T. Grurl, A. Ploier, and T. Peham, “The Basis of Design Tools for Quantum Computing: Arrays, Decision Diagrams, Tensor Networks, and ZX-Calculus,” in *Design Automation Conf.*, 2022.
- [17] A. Zulehner and R. Wille, “Advanced simulation of quantum computations,” *IEEE Trans. on CAD of Integrated Circuits and Systems*, 2019.
- [18] X. Hong, X. Zhou, S. Li, Y. Feng, and M. Ying, “A tensor network based decision diagram for representation of quantum circuits,” *ACM Trans. Des. Autom. Electron. Syst.*, 2022.
- [19] G. F. Viamontes, I. L. Markov, and J. P. Hayes, “Improving gate-level simulation of quantum circuits,” *Quantum Information Processing*, 2003.
- [20] T. Grurl, J. Fuß, and R. Wille, “Noise-aware Quantum Circuit Simulation With Decision Diagrams,” *IEEE Trans. on CAD of Integrated Circuits and Systems*, 2022.
- [21] A. Zulehner, S. Hillmich, and R. Wille, “How to Efficiently Handle Complex Values? Implementing Decision Diagrams for Quantum Computation,” in *Int’l Conf. on CAD*, 2019.
- [22] J. Biamonte and V. Bergholm, “Tensor networks in a nutshell,” 2017. arXiv: 1708.00006.
- [23] I. L. Markov and Yaoyun. Shi, “Simulating quantum computation by contracting tensor networks,” *SIAM Journal on Computing*, 2008.
- [24] R. Orús, “A practical introduction to tensor networks: Matrix product states and projected entangled pair states,” *Annals of Physics*, 2014.
- [25] D. Perez-Garcia, F. Verstraete, M. Wolf, and J. Cirac, “Matrix product state representations,” *Quantum Information and Computation*, 2007.
- [26] S. Jaques and T. Häner, “Leveraging state sparsity for more efficient quantum simulations,” 2021. arXiv: 2105.01533.
- [27] J. Preskill, “Quantum computing in the NISQ era and beyond,” *Quantum*, 2018.
- [28] D. Litinski, “How to compute a 256-bit elliptic curve private key with only 50 million Toffoli gates,” 2023. arXiv: 2306.08585.
- [29] W. van Dam, M. Mykhailova, and M. Soeken, “Using azure quantum resource estimator for assessing performance of fault tolerant quantum computation,” 2023. arXiv: 2311.05801.
- [30] *Qualtran — Qualtran documentation*. [Online]. Available: <https://qualtran.readthedocs.io/en/latest/> (visited on 11/14/2023).
- [31] *Benchq, Zapata AI*. [Online]. Available: <https://github.com/zapatacomputing/benchq> (visited on 11/14/2023).
- [32] *pyLIQTR*, USC Information Sciences Institute. [Online]. Available: <https://github.com/isi-usc-edu/pyLIQTR>.
- [33] V. von Burg *et al.*, “Quantum computing enhanced computational catalysis,” *Phys. Rev. Res.*, 2021.
- [34] H. Liu, G. H. Low, D. S. Steiger, T. Häner, M. Reiher, and M. Troyer, “Prospects of quantum computing for molecular sciences,” *Materials Theory*, 2022.
- [35] E. Campbell, “Random compiler for fast hamiltonian simulation,” *Phys. Rev. Lett.*, 2019.
- [36] G. H. Low and I. L. Chuang, “Hamiltonian Simulation by Qubitization,” *Quantum*, 2019.
- [37] J. Lee *et al.*, “Even more efficient quantum computations of chemistry through tensor hypercontraction,” *PRX Quantum*, 2021.
- [38] S. Wesselbaum *et al.*, “Hydrogenation of carbon dioxide to methanol using a homogeneous ruthenium–triphos catalyst: From mechanistic investigations to multiphase catalysis,” *Chem. Sci.*, 2015.