

# GPU-accelerated simulations of quantum annealing and the quantum approximate optimization algorithm <sup>☆</sup>

Dennis Willsch <sup>a,\*</sup>, Madita Willsch <sup>a,b</sup>, Fengping Jin <sup>a</sup>, Kristel Michielsens <sup>a,b,c</sup>, Hans De Raedt <sup>a,d</sup>

<sup>a</sup> Institute for Advanced Simulation, Jülich Supercomputing Centre, Forschungszentrum Jülich, 52425 Jülich, Germany

<sup>b</sup> AIDAS, 52425 Jülich, Germany

<sup>c</sup> RWTH Aachen University, 52056 Aachen, Germany

<sup>d</sup> Zernike Institute for Advanced Materials, University of Groningen, Nijenborgh 4, 9747 AG Groningen, the Netherlands

## ARTICLE INFO

### Article history:

Received 24 June 2021

Received in revised form 4 May 2022

Accepted 5 May 2022

Available online 10 May 2022

### Keywords:

Quantum computing

Quantum annealing

Approximate quantum annealing

QAOA

High performance computing

Computer simulation

Parallelization

## ABSTRACT

We study large-scale applications using a GPU-accelerated version of the massively parallel Jülich universal quantum computer simulator (JUQCS-G). First, we benchmark JUWELS Booster, a GPU cluster with 3744 NVIDIA A100 Tensor Core GPUs. Then, we use JUQCS-G to study the relation between quantum annealing (QA) and the quantum approximate optimization algorithm (QAOA). We find that a very coarsely discretized version of QA, termed approximate quantum annealing (AQA), performs surprisingly well in comparison to the QAOA. It can either be used to initialize the QAOA, or to avoid the costly optimization procedure altogether. Furthermore, we study the scaling of the success probability when using AQA for problems with 30 to 40 qubits. We find that the case with the largest discretization error scales most favorably, surpassing the best result obtained from the QAOA.

© 2022 The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

## 1. Introduction

The simulation of universal quantum computers requires a large number of matrix-vector updates, most of which are 2-component and 4-component tensor operations. As such, the task of simulating quantum computers is an ideal candidate to profit from recent developments in the GPU industry. We use a GPU-accelerated version of our in-house software JUQCS [1,2], termed JUQCS-G, to benchmark JUWELS Booster, a cluster of 3744 NVIDIA A100 Tensor Core GPUs, integrated in the modular supercomputer JUWELS [3]. A dockerized version of JUQCS is available online [4].

JUWELS Booster is part of the JUWELS cluster-booster architecture at the Jülich Supercomputer Centre (JSC) in which a cluster of multi-core nodes is connected via a high-speed network to a cluster of GPUs, the booster, which forms the basis of the modular supercomputer at JSC. The modular supercomputer architecture generalizes the cluster-booster concept by potentially interconnecting a variety of modules with, among others, different acceleration technologies, AI-adapted nodes and storage devices. The modu-

lar supercomputer concept allows for a seamless integration of quantum computing architectures and future neuromorphic systems to realize the vision of a holistic future hybrid supercomputer [5]. Such a system enables hybrid simulations involving quantum and/or neuromorphic devices that open up new possibilities for demanding computing tasks in science and industry. This will eventually allow for hybrid computing paradigms in a production environment.

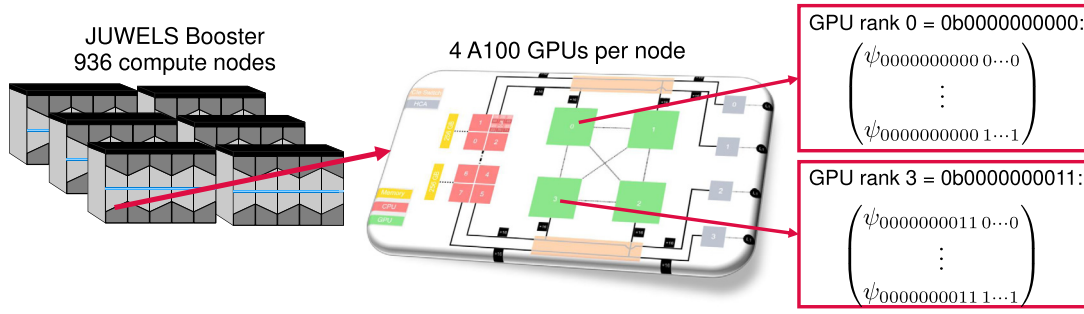
JUQCS is a massively parallel simulator [1,2,4,6] that has also been used for Google's quantum supremacy demonstration [7]. Using JUQCS-G, we study the quantum approximate optimization algorithm (QAOA) [8,9], a popular variational algorithm for near-term gate-based quantum computers, also known as noisy intermediate-scale quantum (NISQ) devices [10]. The prospect of producing useful results for NISQ devices has stimulated considerable interest in the scientific community [11–23].

The QAOA simulations, which were performed on the JUWELS Booster, used the CPUs to carry out the classical (optimization) part of the QAOA and the GPUs to carry out the quantum part formulated in terms of a quantum circuit. On the modular supercomputer architecture with a quantum module, the optimization could be performed on the CPUs of the JUWELS cluster or booster and the operations in the quantum circuit on the QPUs (quantum

<sup>☆</sup> The review of this paper was arranged by Prof. David W. Walker.

\* Corresponding author.

E-mail address: [d.willsch@fz-juelich.de](mailto:d.willsch@fz-juelich.de) (D. Willsch).



**Fig. 1. Distribution of the complex amplitudes of the state vector  $|\psi\rangle$  on the GPUs across the compute nodes.** Each GPU is handled by one MPI process. For each GPU, the leftmost qubit indices of the coefficients (the *global* qubits, separated from the *local* qubits by a space) represent the MPI rank that uniquely identifies the GPU in the supercomputer. This is indicated for the GPUs belonging to MPI rank 0 and 3 for a case with 10 global qubits. On each GPU, the complex amplitudes for each index of the remaining local qubits are stored. During non-local quantum gate operations, typically half of all complex amplitudes need to be transferred once between  $N_{\text{GPU}}/2$  pairs of GPUs (often across different compute nodes). For these transfers, the MPI communication scheme of JUQCS-G follows the original one described in [6], with the only qualitative change being that a CUDA-aware MPI implementation is used to transfer memory between the GPUs.

processing unit), enabling efficient quantum-classical hybrid computations.

The QAOA can be related to a discretized version of quantum annealing (QA) [14,19,24,25]. QA is another popular paradigm of quantum computation [26–35] that is studied alongside the gate-based model of quantum computation [36]. Special devices built to perform QA are the D-Wave quantum annealers. The largest existing quantum annealer is the D-Wave Advantage, which has 5000+ physical qubits [37] and has been used for quantum support vector machines [38,39] (see also [40]), in studies of stock markets [41], for computer vision [42], and for lattice gauge theory [43]. It has recently been benchmarked with 3D spin glass problems [44], garden optimization problems [45] and exact cover problems [46]. In the present work, the same exact cover problems as in Ref. [46], derived from simplified optimization problems encountered in airplane scheduling, are used to analyze the large-scale simulation results produced by different physical models designed to solve such problems.

In this paper, we scrutinize the overlapping region between QA and the QAOA. We start from a coarse, second-order time-discretization of QA that we call *approximate quantum annealing* (AQA). We increase the time step that controls the discretization error (sometimes referred to as the *Trotter error* [47,48], although the formalism goes well beyond Trotter's investigation [49], see [50–52]). Furthermore, we use JUQCS-G to study the scaling of the success probability when using AQA for exact cover problems with 30 to 40 qubits. Surprisingly, we find that, while the cases with smaller discretization error provide useful initializations for the QAOA, the cases with largest discretization error scale much better when increasing the number of qubits.

Ideas that are similar to AQA have been investigated before [19, 24,25,53]. In particular, in [25] a first-order discretized version of QA, referred to as *Trotterized quantum annealing*, was used as initialization for the QAOA. The authors studied the performance for  $p \leq 10$  QAOA steps and relatively small systems with  $N \leq 12$  qubits. Here, we study a *second-order* discretization of QA. We study not only the QAOA initialization but also the *dynamics* of AQA. Furthermore, we consider much larger systems with up to  $N = 40$  qubits and up to  $n = 100$  steps (corresponding to  $p = 101$ ).

While it is almost trivial to simulate short QAOA gate circuits for less than 26 qubits on a modern PC, simulating the fairly lengthy circuits (5000+ gates) for the 40 qubits exact cover problems requires substantial supercomputer resources (and more than 16TB of random access memory). The GPU-enabled software that we have developed in house enables us to perform such simulations in a reasonable time span. Having data for 30–40 qubits allows us to assess the potential, e.g. the scaling behavior, of the

QAOA and AQA in a regime that was previously inaccessible (in practice).

This paper is structured as follows. In Section 2, we describe the GPU-accelerated universal quantum computer simulator JUQCS-G and show benchmarks of JUWELS Booster. In Section 3, we present applications to QA, AQA, and the QAOA. We summarize our findings in Section 4.

## 2. JUQCS-G

In this section, we outline the central task performed by universal quantum computer simulators such as JUQCS in general, and its GPU-accelerated version JUQCS-G in particular. After this, we present benchmark results for JUWELS Booster.

### 2.1. Simulating quantum computers on GPUs

The basic unit of computation for a gate-based quantum computer is a single qubit, described by two complex numbers  $|\psi\rangle = (\psi_0, \psi_1)$  that are normalized so that  $\langle\psi|\psi\rangle = |\psi_0|^2 + |\psi_1|^2 = 1$ . By definition, an  $N$ -qubit system is described by  $2^N$  complex numbers

$$|\psi\rangle = \psi_{0\dots 00} |0\dots 00\rangle + \psi_{0\dots 01} |0\dots 01\rangle + \dots + \psi_{1\dots 11} |1\dots 11\rangle, \quad (1)$$

where  $|0\dots 00\rangle, \dots, |1\dots 11\rangle$  are the computational basis states [36] and the coefficients  $\psi_{0\dots 00}, \dots, \psi_{1\dots 11}$  are normalized such that  $\langle\psi|\psi\rangle = 1$ . For clarity, we explicitly write the  $2^N$  complex coefficients in the state  $|\psi\rangle$  as a rank- $N$  tensor  $\psi_{q_{N-1}\dots q_1 q_0}$  with indices  $q_j \in \{0, 1\}$ . In other words, an  $N$ -qubit system is described by a complex-valued, rank- $N$  tensor, a tensor product of  $N$  two-dimensional vectors.

For large-scale universal quantum computer simulations, the main difficulty lies in the management of all  $2^N$  complex numbers. For instance, for  $N = 42$  using double precision floating-point numbers, the tensor  $\psi_{q_{N-1}\dots q_1 q_0}$  occupies  $16 \times 2^{42} \text{ B} = 64 \text{ TiB}$  of distributed memory.

JUQCS-G distributes the complex numbers over the memory of the GPUs as indicated in Fig. 1. Each GPU stores  $2^M$  coefficients of  $|\psi\rangle$  in its local memory, i.e., each GPU stores the coefficients  $(\psi_{q_{N-1}\dots q_M 0\dots 0}, \dots, \psi_{q_{N-1}\dots q_M 1\dots 1})$ . For this reason, we call the rightmost  $M$  qubits  $q_{M-1}\dots q_0$  *local* qubits. As a consequence, the total number of required GPUs is given by  $N_{\text{GPU}} = 2^{N-M}$ .

Since the complex numbers are distributed over multiple GPUs on different compute nodes, data has to be transferred over the network. This is necessary, for instance, if a particular part of the data on one node is required for the computation on another node. To exchange data between the GPUs on different compute nodes,

the *Message Passing Interface* (MPI) is used. MPI provides a standard for distributed memory computation and takes care of the communication, i.e., the sending and receiving of data between different compute nodes. For details, we refer the reader to the literature [54].

JUQCS-G uses CUDA-aware MPI to manage the distributed memory. Each GPU is controlled by one MPI process, whose rank  $r \in \{0, \dots, N_{\text{GPU}} - 1\}$  is initially given by the leftmost  $N - M$  indices of  $|\psi\rangle$  in binary notation. Thus, the GPU with rank  $\text{bin}(r) = q_{N-1} \dots q_M$  holds the coefficients  $(\psi_{\text{bin}(r)0\dots 0}, \dots, \psi_{\text{bin}(r)1\dots 1})$ . For this reason, we call the leftmost  $N - M$  qubits *global* qubits.

A quantum gate is a unitary operation that transforms some of the coefficients of  $|\psi\rangle$ . The most elementary quantum gate is a single-qubit gate, i.e., a  $2 \times 2$  unitary matrix  $U = (u_{qq'})$ . It transforms the coefficients of  $|\psi\rangle$  in terms of 2-component updates. For instance, a single-qubit gate on qubit  $j$  transforms the tensor  $|\psi\rangle$  according to

$$\psi_{q_{N-1} \dots q_{j+1} q q_{j-1} \dots q_0} \leftarrow \sum_{q'=0}^1 u_{qq'} \psi_{q_{N-1} \dots q_{j+1} q' q_{j-1} \dots q_0}, \quad (2)$$

for  $q = 0, 1$ . Similarly, a two-qubit gate is a  $4 \times 4$  unitary matrix that operates on two indices of  $|\psi\rangle$ , and a three-qubit gate operates on three indices, etc. A suitable set of one- and two-qubit gates suffices to construct a universal quantum computer (simulator) [55,56]. The set of quantum gates implemented by JUQCS-G is documented in [1].

We do not use sparse matrix techniques but exploit the special structure of single-, two- and three-qubit operations. We never store or operate on large dense matrices. To perform 2-component updates as expressed in Eq. (2), we loop over all pairs of elements of  $|\psi\rangle$  and multiply each pair of elements with the same  $2 \times 2$  matrix (which depends on the particular gate). The grouping in pairs depends on the qubit that is being operated on. For the two-qubit operations, we loop over quadruples of elements of  $|\psi\rangle$  and multiply each quadruple of elements with the same  $4 \times 4$  matrix (which depends on the particular gate). The time it takes to perform all these arithmetic operations is counted as “compute time”. As the size of the quantum computer increases, we need more and more compute nodes to store  $|\psi\rangle$ , and although the MPI communication is very efficient by itself, it takes an increasingly larger part of the elapsed time (but still scales approximately linearly, not exponentially).

If a quantum gate acts on a global qubit, coefficients of  $|\psi\rangle$  that are stored on different GPUs need to be combined with each other. This requires MPI communication between the GPUs. For circuits with many quantum gates involving global qubits, the MPI communication may take a large part of the simulation time (cf. Fig. 3 below). For instance, a single-qubit gate on a global qubit requires the transfer of  $2^N/2$  complex numbers (i.e., half of all memory) between pairs of GPUs. JUQCS-G minimizes the communication overhead by *relabeling* global and local qubits after such a global quantum gate. Thereby, the complex numbers need to be transferred over the network only *once*, and not back again after the transformation. Each GPU keeps track of the labeling of global and local qubits in a local permutation array. Further details of this optimal MPI communication scheme are explained in [6].

The keyword in the large-scale simulations performed by JUQCS is *universal*. It means that any quantum circuit for an  $N$ -qubit system can in principle be simulated, as long as the circuit depth is not unreasonably long (*unreasonably* because then it would also not be executable on a gate-based quantum computer device). In the literature, this kind of simulation method is sometimes referred to as the *Schrödinger simulation method*, because the whole tensor  $\psi_{q_{N-1} \dots q_1 q_0}$  (i.e., the whole *wave function*  $|\psi\rangle$ ) is propagated

through the quantum circuit. The simulation time grows linearly in the total number of gates.

In contrast to the Schrödinger simulation method, there is also the so-called *Feynman simulation method* [57–62]. Here, tensor networks are used to obtain only one (or a few) amplitude(s) of the final quantum state. One then sums over each path through the quantum circuit that would contribute to this amplitude. In principle, much larger qubit systems can then be simulated (e.g., a 128-qubit circuit was simulated in [2]). Of course, the kinds of circuits that can be simulated by such an approach are very restricted and not universal. The simulation time grows exponentially in the circuit depth and depends strongly on the number of Schmidt coefficients of multi-qubit gates (see the supplementary material of [7]). However, truncating Schmidt coefficients opens the possibility to simulate circuits with smaller fidelity. An overview of the limits of such simulations is given in [63].

A combination of both Schrödinger and Feynman approaches can be used to simulate larger circuits of the quantum supremacy experiment [7], and has recently been used on a cluster of GPUs to spoof the quantum supremacy test [64].

## 2.2. Benchmarks and scalings

The large amount of MPI communication required for simulating universal quantum computations makes simulating quantum computers an ideal candidate to benchmark large supercomputers. Combined with the many tensor operations required (cf. Section 2.1), JUQCS-G is a very versatile application to benchmark Tensor Core GPUs. In this section, we report benchmark results for JUQCS-G running on JUWELS Booster, a cluster with 3744 NVIDIA A100 Tensor Core GPUs distributed over 936 compute nodes (see Fig. 1).

Each A100 GPU has a local memory of 40 GiB, so the maximum number of local qubits is 31. For quantum circuits with  $N \geq 32$  qubits, MPI communication between the GPUs is necessary. For the present benchmark study, we simulate quantum circuits for 32–42 qubits on 2–2048 GPUs.

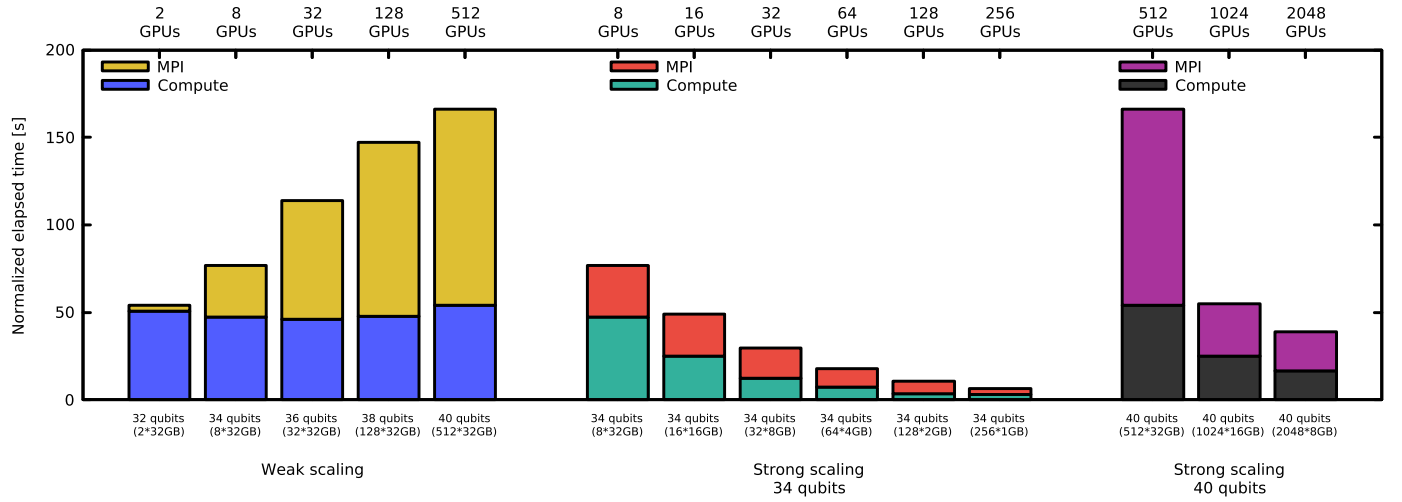
In Fig. 2, we show simulation results for QAOA circuits for 32–40 qubit exact cover problems (the details of which are described in the following section). We see that the computation time (i.e., the run time excluding the time required for the MPI communication) stays approximately constant with increasing system size, indicating ideal weak scaling. The MPI communication time increases roughly linearly. Most importantly, none of these simulation times grow *exponentially* in the number of qubits. In this sense, JUQCS-G beats the exponential growth associated with quantum circuit simulations.

To compare timing data of different runs of problems belonging to the same class, it is expedient to express this data in a way that takes into account that the number of gates depends on the problem size  $N$ . In the present case, we take the number of gates  $n_{\text{gates}}(32)$  for the smallest corresponding problem instance as reference and define

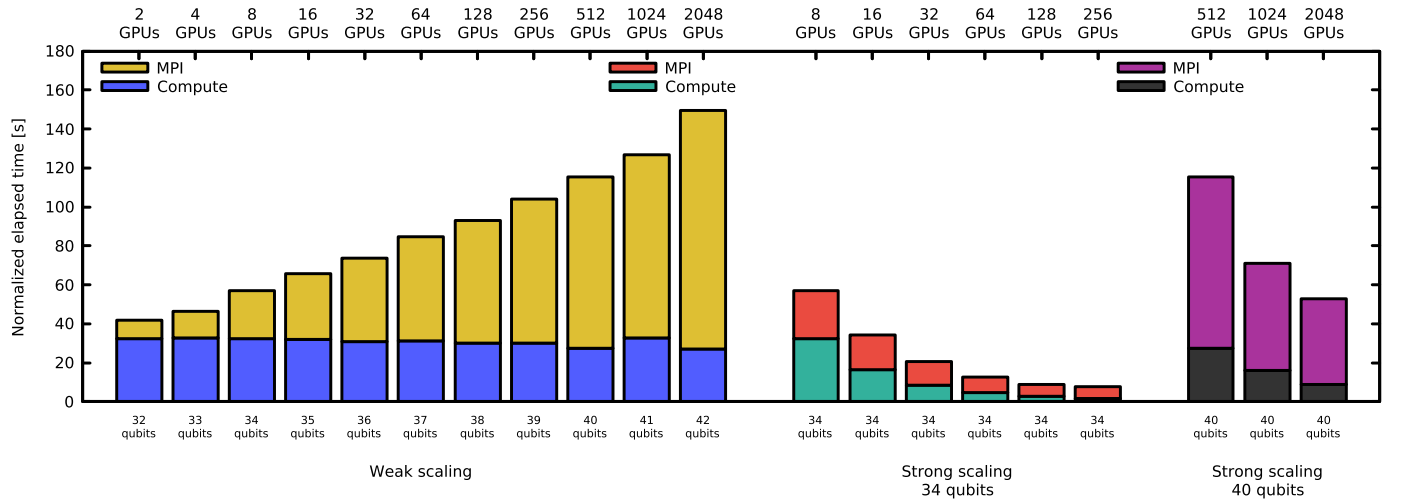
$$\text{Normalized elapsed time} = \frac{n_{\text{gates}}(32)}{n_{\text{gates}}(N)} T_{\text{elapsed}}(N). \quad (3)$$

Studying the strong scaling results for 34 qubits, we find ideal strong scaling. As the number of GPUs increases, the normalized elapsed time decreases exponentially. When doubling the number of GPUs used, the normalized elapsed time is (almost perfectly) halved.

Looking closely at the 40-qubit strong scaling results in Fig. 2 (rightmost bars), we see that the drop in simulation time from 512 to 1024 GPUs is in fact better than expected. For perfect strong scaling, we would expect the simulation time to decrease by a factor of 2 when doubling the number of GPUs (in practice, this



**Fig. 2.** Weak and strong scaling results for QAOA on JUWELS Booster using 4 NVIDIA A100 GPUs per node. Shown is the normalized elapsed time given by Eq. (3) as a function of the number of GPUs. The problem size given by the number of qubits and the memory per GPU are indicated on the bottom axis. “Compute” refers to the elapsed time for executing the quantum circuit only. “MPI” refers to the elapsed time for communication plus the elapsed time to prepare and postprocess MPI buffers. There is no overlap between computation and communication. (For interpretation of the colors in the figure(s) and table(s), the reader is referred to the web version of this article.)



**Fig. 3.** The same as Fig. 2 but for the Hadamard benchmark circuits  $(H^{\otimes N})^{11}$ . In this case, the largest runs for 40–42 qubits were repeated several times to estimate the fluctuations due to different node allocations; they were on the order of 1 second and thus negligible (data not shown). “Compute” refers to the elapsed time for executing the quantum circuit only. “MPI” refers to the elapsed time for communication plus the elapsed time to prepare and postprocess MPI buffers. There is no overlap between computation and communication.

decrease would be expected to be even a little less). Going from 512 to 2048 GPUs, i.e., using 4 times as many GPUs, brings the normalized elapsed time down by almost a factor of 4 as expected. This holds for the computing time as well as for the MPI communication time. However, we observe that the time needed with 1024 GPUs is only a third of the time needed with 512 GPUs, so much better than the theoretical optimum. Note that the unexpected behavior can be attributed to the MPI communication part only. Considering only the computing time, we still observe the expected scaling. As this run was performed in October 2020 during the early testing period of JUWELS Booster, we assumed that an explanation for the behavior might be found in an irregularity in the DragonFly+ topology of the communication network.

Therefore, we repeated the large-scale benchmark in February 2021 after JUWELS Booster went into production. This time, we used quantum circuits consisting only of Hadamard gates on each qubit, repeated 11 times in a row,  $(H^{\otimes N})^{11}$ . Such circuits have been found to be well-suited for both benchmarking gate-based quantum computers [65] and universal quantum computer simulators [1]. They create uniform superpositions over all  $N$  qubits

and require exchanging  $2^N/2$  complex numbers over the whole GPU network for each global single-qubit  $H$  gate. Since the total number of gates as a function of  $N$  is not constant, we need to make the benchmark results for different  $N$  relatable by normalizing the run times w.r.t. the 32-qubit version. For instance, as the 32-qubit circuit has 352  $H$  gates and the 42-qubit circuit has 462  $H$  gates, the corresponding normalization factor is given by  $462/352 \approx 1.31$ . The 11-fold repetition of the Hadamard gates makes potential GPU/CUDA/MPI initialization times negligible.

The results of this second benchmark are shown in Fig. 3. We see that in this case, the computation times show nearly ideal scaling, i.e., the elapsed time for increasing qubit number and number of used GPUs stays approximately constant (ideal weak scaling) and for constant qubit numbers, doubling the number of used GPUs halves the computation time (ideal strong scaling) in the 34-qubit case as well as in the 40-qubit case. Also for the MPI communication time, the results follow the theoretical expectation.

To compare the speedup over the CPU-based version of JUQCS, JUQCS-E [1], we also report results for the normalized run times

**Table 1**

**Comparison of the GPU-based simulator JUQCS-G (first row) and the CPU-based simulator JUQCS-E [1] (second to last row) for the largest systems using the Hadamard benchmark circuits ( $H^{\otimes N}$ )<sup>11</sup>.** The time  $t_{\text{total}}$  is the run time spent for the total simulation, normalized by the number of gates with respect to the 32-qubit case (see Eq. (3)). The time  $t_{\text{MPI}}$  is the elapsed time for communication plus the elapsed time to prepare and postprocess MPI buffers. JUQCS-E uses all cores of the CPUs on each node.

qubits	nodes	processes	hardware	normal.	$t_{\text{total}}$ [s]	$t_{\text{MPI}}$ [s]
42	256	1024	GPU	1.31	149.4	122.1
42	256	2048	CPU	1.31	2632.4	1297.7
42	512	4096	CPU	1.31	1500.4	763.4
43	512	4096	CPU	1.34	2714.4	1343.3

for the largest circuits in Table 1 using only CPUs. In this mode of operation, JUWELS Booster can also run 43-qubit circuits.

For the 42-qubit case, we see that the normalized run time on 2048 CPUs,  $t_{\text{total}} = 2632.4$  s, is a factor of 18 larger than the GPU-accelerated version with  $t_{\text{total}} = 149.4$  s (also shown in Fig. 3). Furthermore, after subtracting the MPI communication time  $t_{\text{MPI}}$ , the speedup due to the GPU acceleration for the computation-only part is 49. This is a very significant improvement in terms of the computational resources required for the simulations. Clearly, large-scale quantum circuit simulations can tremendously benefit from recent GPU developments.

### 3. Applications

In this section, we use JUQCS-G to study the quantum computer applications QA, AQA, and the QAOA. The QAOA work presented is, in spirit, similar to the work reported in Ref. [14]. However, the largest optimization problems (16 variable MaxCut and 18 variable 2-SAT) studied in Ref. [14] are much smaller (recall the exponential dependence on the number of variables) than the 40-variable exact cover problems studied in the present manuscript.

We first outline the mathematical background and its implementations, and then present the simulation results.

#### 3.1. Background

In this section, we discuss the methods that we used in our studies. First, we briefly review the most important aspects of QA and the QAOA in Sections 3.1.1 and 3.1.2, respectively. The definition of the exact cover problem, which is the class of problems that we study in this paper, is given in Section 3.1.3.

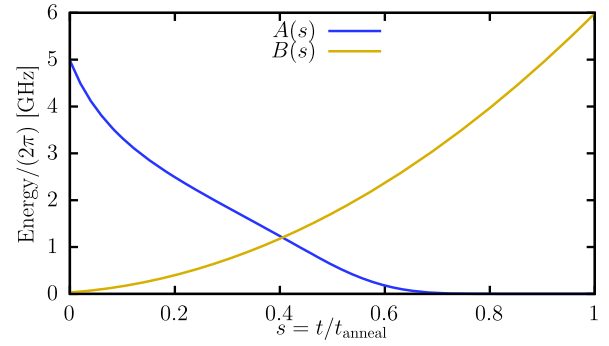
##### 3.1.1. Quantum annealing

QA was initially intended as an algorithm for conventional computers [26–28]. Over time, it has evolved into the idea of a quantum computing device that works fundamentally different from the gate-based quantum computer.

The concept of QA is based on the adiabatic theorem [67,68]. For this reason, a QA device is also called *adiabatic quantum computer*. Although slightly different concepts are sometimes associated with QA and adiabatic quantum computation, the basic working principle is the same: The quantum system (consisting of qubits) is prepared in the ground state of an initial Hamiltonian such as

$$H_I = - \sum_{i=0}^{N-1} \sigma_i^x, \quad (4)$$

whose ground state is given by  $|\psi_{\text{init}}\rangle = |+\rangle^{\otimes N}$  where  $|+\rangle = (|0\rangle + |1\rangle)/\sqrt{2}$  is the uniform superposition of  $|0\rangle$  and  $|1\rangle$ . During the time evolution, the Hamiltonian changes according to



**Fig. 4.** Annealing schedule of the DW\_2000Q\_6 quantum annealer, taken from [66]. The annealing functions  $A(s)$  (blue line) and  $B(s)$  (yellow line) describe the evolution of the QA Hamiltonian given in Eq. (5).

$$H(s) = A(s)H_I + B(s)H_C, \quad s = t/t_{\text{anneal}}, \quad (5)$$

where  $t_{\text{anneal}}$  is the time used for the annealing process, and the two annealing functions  $A(s)$  and  $B(s)$  fulfill  $A(0) \gg B(0)$  and  $A(1) \ll B(1)$ . An example annealing schedule is shown in Fig. 4, which is also used to initialize the variational QAOA parameters (see below).

The final Hamiltonian in Eq. (5),  $H_C$ , represents an optimization problem that is to be solved. This means that the ground state of  $H_C$  encodes the solution of a certain optimization problem. Here, we choose  $H_C$  to be the Ising Hamiltonian

$$H_C = \sum_{i=0}^{N-1} h_i \sigma_i^z + \sum_{i<j} J_{ij} \sigma_i^z \sigma_j^z. \quad (6)$$

The idea is that if the annealing process described by Eq. (5) is carried out at zero temperature and sufficiently slowly so that the adiabatic theorem holds, then the quantum system stays in its instantaneous ground state. Thus, at the end of the annealing process, the quantum system ends up in the ground state of the Hamiltonian  $H_C$ . Measuring the qubits would then yield the answer to the initial optimization problem.

In practice, on a quantum annealer not only the annealing time  $t_{\text{anneal}}$  determines the probability of success (i.e., the probability that the system ends in its ground state and not in an excited state), but also an environment at finite temperature, control errors and precision limits have an influence on it [30,69–75].

##### 3.1.2. The quantum approximate optimization algorithm

The QAOA was introduced by Farhi et al. [8]. It is a variational method that is suitable for execution on a gate-based quantum computer. The objective is to find the ground state (or a low energy state) of a problem Hamiltonian such as  $H_C$  given by Eq. (6) that represents an optimization problem. The state that is prepared by the QAOA quantum circuit is given by

$$|\beta, \gamma\rangle = \prod_{k=1}^p e^{-i\beta_k H_D} e^{-i\gamma_k H_C} |+\rangle^{\otimes N}, \quad (7)$$

where  $\gamma = (\gamma_1, \dots, \gamma_p)$  and  $\beta = (\beta_1, \dots, \beta_p)$  are the  $2p$  variational parameters that have to be optimized, and  $H_D$  is a *mixing Hamiltonian* that is commonly chosen as  $H_D = -H_I$  (cf. Eq. (4)), i.e.,

$$H_D = \sum_{i=0}^{N-1} \sigma_i^x. \quad (8)$$

Note that other choices have also been proposed [76,77].

It is worth mentioning that for this choice, the QAOA parameters  $\beta_k$  can be reduced to the range  $[0, \pi)$ . For  $\gamma_k$ , however, such

a periodicity condition depends on the minimum spacing between the eigenvalues of  $H_C$ . In other words, the range of values for  $\gamma_k$  depends on the particular problem instance defined by  $h_i$  and  $J_{ij}$  (see below).

For a given number of steps  $p$ , the energy of the optimized variational state ( $E_p^* = \min_{\beta, \gamma} \langle \beta, \gamma | H_C | \beta, \gamma \rangle$ ) is lower than the energy of the optimized variational state with  $p - 1$  steps [8].

However, it has been found that the optimization of the variational parameters can be rather inefficient. Therefore, one often tries to use the observation that the optimal parameters  $\beta_k$  and  $\gamma_k$  seem to follow certain patterns [14,15,19,78–80]. Here, we investigate these patterns in relation to QA and their interpretation as an optimized annealing scheme (see also [14,19,25]).

### 3.1.3. Exact cover

The exact cover problem is an NP-complete problem [81] that has become a popular choice to study optimization using quantum computing systems [15–17,46,82–86]. Exact cover problems belong to the class of set covering and partitioning problems that is covered in a vast amount of literature in Operations Research (see e.g. [87,88]).

In this paper, we study the instances of exact cover problems used in [46]. In matrix form, they are written as

$$\min_{x_i=0,1} \sum_{f=0}^{F-1} \left( \sum_{i=0}^{N-1} a_{if} x_i - 1 \right)^2, \quad (9)$$

where  $a \in \{0, 1\}^{N \times F}$  is the Boolean problem matrix that defines the exact cover instance, and  $x_i$  are the problem variables. Intuitively, the solution  $x \in \{0, 1\}^N$  of Eq. (9) selects rows of  $a$  in such a way that in each column of the selected rows, the entry 1 is *covered exactly once*, and all other entries are 0.

We study exact cover problems with 30 to 40 variables (qubits) and  $F = 472$  terms. For each problem size, we have four different instances. Problem instances are labeled by their qubit number and an additional label from 0 to 3 in brackets, such as problem 30(0).

To find a problem Hamiltonian  $H_C$  of the form of Eq. (6), whose ground state represents the solution to Eq. (9), we replace the problem variables according to

$$x_i \mapsto (1 + \sigma_i^z)/2. \quad (10)$$

Denoting the  $-1$  ( $+1$ ) eigenstate of  $\sigma_i^z$  as  $|0\rangle$  ( $|1\rangle$ ), we can represent the problem variable  $x_i$  by the qubit state  $|x_i\rangle$ . Thus, the replacement Eq. (10) yields a diagonal Hamiltonian whose eigenvalues take all possible values of the objective function in Eq. (9). Consequently, the ground state of  $H_C$  (i.e., the state with minimum eigenvalue) is the solution to Eq. (9). For this reason, we also define the *success probability* for these problems as the probability to find the system in the ground state (note that for all problem instances that we study in this paper, the ground state is unique [46]).

After multiplying out the square, the Hamiltonian can be expressed in the form of Eq. (6) plus an additive constant  $C$  (see [46] for the calculation), yielding

$$h_i = \sum_j \frac{1}{2} (aa^T)_{ij} - (ab)_i, \quad (11)$$

$$J_{ij} = \frac{1}{2} (aa^T)_{ij}, \quad (12)$$

$$C = b^T b + \frac{1}{2} \sum_{i < j} (aa^T)_{ij} + \frac{1}{2} \sum_i ((aa^T)_{ii} - (2ab)_i), \quad (13)$$

where  $b = (1, \dots, 1)^T$  is an  $F$ -dimensional vector of ones.

As  $a$  and  $b$  in Eqs. (11)–(13) contain only zeros and ones, we know that  $h_i$  and  $J_{ij}$  vary at most by half integers. Therefore, the range of values for  $\gamma_k$  can be reduced to  $[0, 2\pi)$  (because  $\gamma_k \mapsto \gamma_k + 2\pi$  only causes a global phase in the QAOA state in Eq. (7)).

For all AQA and QAOA applications (except the grid scan in Fig. 5 below), however, we rescale the parameters  $\{h_i\}$ ,  $\{J_{ij}\}$  and  $C$  to a uniform parameter range by dividing them by

$$r = \max \left\{ \max \left[ \frac{\max\{h_i\}}{h_{\max}}, 0 \right], \max \left[ \frac{\min\{h_i\}}{h_{\min}}, 0 \right], \max \left[ \frac{\max\{J_{ij}\}}{J_{\max}}, 0 \right], \max \left[ \frac{\min\{J_{ij}\}}{J_{\min}}, 0 \right] \right\}, \quad (14)$$

where  $h_{\max} = -h_{\min} = 2$  and  $J_{\max} = -J_{\min} = 1$ . Note that the same normalization is also performed when solving such problems on the D-Wave quantum annealer [46,89]. This does not change the solutions of the problems. However, it brings the energies of different problem instances on a uniform scale. This in turn improves the optimization of the QAOA parameters, and it also allows the use of the same AQA time step  $\tau$  (see below) for different problem instances.

## 3.2. Implementations

In this section, we discuss how quantum physics simulations are used to carry out the QAOA and AQA.

### 3.2.1. QAOA

We initialize the  $2p$  variational QAOA parameters  $\beta_k$  and  $\gamma_k$  in Eq. (7) according to the second-order Suzuki-Trotter decomposition. This amounts to (see Appendix A)

$$\beta_k = -\tau(A(s_{k+1}) + A(s_k))/2, \quad k = 1, \dots, p-1, \quad (15)$$

$$\beta_p = -\tau A(s_p)/2, \quad (16)$$

$$\gamma_k = \tau B(s_k), \quad k = 1, \dots, p, \quad (17)$$

where  $s_k = (k-1)/(p-1)$  (slightly different from Ref. [14]) and we take  $A(s)$  and  $B(s)$  from the DW\_2000Q\_6 annealing schedule [66] (see Fig. 4). This procedure is motivated by the relation between the QAOA and QA as discussed in more detail in [14] (see also [25] where the first-order case is discussed).

Given values for the variational parameters  $\beta$  and  $\gamma$ , JUQCS-G computes the variational state  $|\beta, \gamma\rangle$  given in Eq. (7), where the exponentials  $\exp(-i\beta_k H_D) = \prod_i \exp(-i\beta_k \sigma_i^x)$  are computed as a sequence of rotations around the  $x$  axis with angle  $2\beta_k$ , and the exponentials  $\exp(-i\gamma_k H_C)$  are computed as a sequence of rotations around the  $z$  axis with angle  $2\gamma_k$  and controlled- $Z$  gates.

JUQCS-G also computes the probability of the solution state in the variational state  $|\beta, \gamma\rangle$  and the energy expectation value  $E_p(\beta, \gamma) = \langle \beta, \gamma | H_C | \beta, \gamma \rangle$ . In the optimization phase of the QAOA, this energy is passed to the optimizer (we use several optimizers from the *scipy* library [90]; see below). The optimizer then proposes new values for the variational parameters which are in turn passed to JUQCS-G. If this optimization loop does not reach convergence, we use an additional stopping criterion of a maximum of 200 calls to JUQCS-G.

We note that in practice, it is only possible to optimize for the energy  $E_p$  and not for the success probability. Since we use a state-vector simulator and know the ground state, we could in principle also optimize for the probability to observe the ground state. However, in our benchmark, we consider the realistic situation that we do not know the ground state and thus optimize for the energy. We compute the probability of the ground state in a given variational state only as a measure of success.

For optimization problems, often the approximation ratio is also considered as a measure for the performance of the QAOA. In our

case, however, only finding the unique ground state is considered as success since none of the excited states encodes a valid solution to the exact cover problem.

### 3.2.2. AQA

To introduce the basic idea of AQA, we first review how the time evolution of a QA process is simulated. This allows us to describe in what sense the description becomes “approximate”, and when the simulation of QA enters the regime of what we call AQA.

In essence, a simulation of QA requires the solution of the time-dependent Schrödinger equation (TDSE),  $i\partial_t |\psi(t)\rangle = H(t) |\psi(t)\rangle$ , with a time-dependent Hamiltonian  $H(t)$ , such as the QA Hamiltonian given in Eq. (5). In principle, but also in practice, the time-discretized TDSE can be expressed as a quantum gate circuit which can then be processed by JUQCS-G. However, for convenience (and also as a check on the JUQCS data), we often solve the TDSE with the quantum spin dynamics simulator (QSDS) (in house software with the MPI communication scheme taken from JUQCS but without GPU implementation). QSDS solves the TDSE for the generic spin-1/2 Hamiltonian

$$H_{\text{QSDS}}(t) = - \sum_{\alpha=x,y,z} \left( \sum_{i<j} \tilde{J}_{ij}^{\alpha}(t) \sigma_i^{\alpha} \sigma_j^{\alpha} + \sum_{i=0}^{N-1} \tilde{h}_i^{\alpha}(t) \sigma_i^{\alpha} \right), \quad (18)$$

where  $N$  is the number of spins (qubits). For the optimization problems at hand, we have

$$\tilde{h}_i^x(t) = A(t/t_{\text{anneal}}), \quad (19)$$

$$\tilde{h}_i^z(t) = -B(t/t_{\text{anneal}})h_i, \quad (20)$$

$$\tilde{J}_{ij}^z(t) = -B(t/t_{\text{anneal}})J_{ij}, \quad (21)$$

where the annealing functions  $A$  and  $B$  are shown in Fig. 4, and  $h_i$  and  $J_{ij}$  encode the problem instance as before.

QSDS solves the TDSE by time stepping using the second-order Suzuki-Trotter formula [91,92]

$$\begin{aligned} |\Psi((l+1)\tau)\rangle = & \left\{ \exp \left[ \frac{i\tau}{2} \sum_{\alpha=x,z} \sum_{i=0}^{N-1} \tilde{h}_i^{\alpha}(l\tau) \sigma_i^{\alpha} \right] \right. \\ & \times \exp \left[ i\tau \sum_{i<j} \tilde{J}_{ij}(l\tau) \sigma_i^z \sigma_j^z \right] \\ & \left. \times \exp \left[ \frac{i\tau}{2} \sum_{\alpha=x,z} \sum_{i=0}^{N-1} \tilde{h}_i^{\alpha}(l\tau) \sigma_i^{\alpha} \right] \right\} |\Psi(l\tau)\rangle, \quad (22) \end{aligned}$$

for  $l = 0, \dots, n$  (such that  $t_{\text{anneal}} = (n+1)\tau$ ). Note that the action of each of the matrix exponentials in Eq. (22) on any state vector can be computed exactly. For the initial state, we take  $|\Psi(0)\rangle = |+\rangle^{\otimes N}$ . Apart from collecting all single-spin terms of the Hamiltonian Eq. (18) into the same matrix exponential, the structure of the QAOA (cf. Eq. (7)) is the same as that of Eq. (22).

The basic idea of AQA is to solve the TDSE with a time step  $\tau$  which is too large to yield an accurate time evolution of a genuine QA process. Moreover, the number of time steps  $n$  is taken to be rather small. Therefore the corresponding “annealing time” is rather short in which case the time evolution is unlikely to be adiabatic.

In other words, we do not rely on the adiabatic theorem but hope that with a relatively small number of factors in the product formula with a relatively large time step, we can nevertheless generate a final state which is close to the ground state of the problem Hamiltonian. Clearly, AQA is a heuristic method, partially motivated by findings [14,15,19,78,79] that optimal values for the

variational parameters  $\beta_k$  and  $\gamma_k$  were often found to follow curves which resemble such an approximate annealing schedule. For AQA, we use again Eqs. (15)–(17) but with the convention:  $s_k = k/n$  for  $k = 0, \dots, n$ .

For each step, AQA and QAOA perform exactly the same number of single-qubit and two-qubit gates. Only the single-qubit gates may require exchange of data among MPI processes. In our AQA simulations, we also compute the spin expectation values during the time evolution. In terms of computational effort, AQA for a fixed  $n$  (i.e.  $n+1$  steps since we start counting at 0) is equivalent to a single evaluation of a QAOA circuit with  $p = n+1$ .

### 3.3. Results

In this section, we present the simulation results on the QAOA, AQA, and a comparison between them.

#### 3.3.1. QAOA

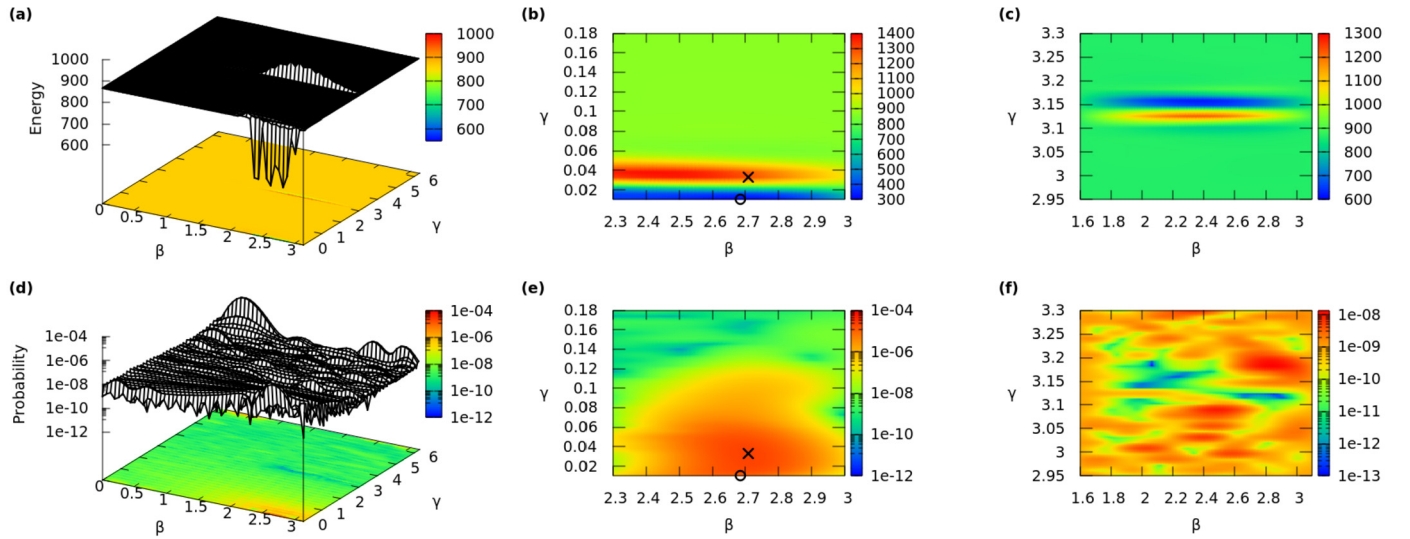
We start with QAOA for  $p = 1$  by studying the energy landscape and the success probability for a 30-qubit problem instance (named 30(0)). For this purpose, we perform a scan of the parameters  $\beta \in [0, \pi)$  and  $\gamma \in [0, 2\pi)$  and compute the energy as well as the success probability obtained for the QAOA circuit (note that, as argued in Section 3.1.3, these parameter intervals cover the range of different QAOA trial states). The results are shown in the left column of Fig. 5 for the energy (top) and success probability (bottom). The middle column shows zooms with finer grids into regions around the energy minimum and the success probability maximum. The right column shows zooms into another region of interest noticeable in Fig. 5(a).

The point with the highest success probability is marked with a cross and the point with the lowest energy is marked with a circle. Although these points are relatively close, the energy is very different. The point with the highest success probability is even close to an energy maximum. However, the point at the energy minimum still has a relatively high success probability. In the right column, which shows the zoom in the vicinity of another local energy minimum, we find that the success probability is quite low (see the scale of the color bars). If, during the optimization process, the minimizer gets stuck in such a local minimum, the probability to observe the ground state will often be very small.

We find that the optimal parameters  $\beta^*$  and  $\gamma^*$  in this case are large (almost  $\pi$ ) and small (almost 0), respectively, as would be the values for  $A(0)$  and  $B(0)$  in an annealing scheme. This is encouraging for our annealing scheme initialization for QAOA with  $p > 1$ .

Fig. 6 shows the paths that different optimization algorithms take when starting from the point with minimum energy found in the scan (the black circle in Fig. 6(b)). The optimization algorithms are standard optimizers provided by `scipy` [90]: sequential least squares programming (SLSQP), the gradient-free Nelder-Mead algorithm [93] (NM), the L-BFGS-B algorithm [94,95], and the conjugate gradient algorithm [96] (CG) (see [22] for a thorough comparison of different optimizers for the QAOA). Note that for the optimization, we use the rescaled version of the problem according to Eq. (14) (here  $r = 36.75$ ); otherwise small variations in  $\gamma$  led to large fluctuations in the energy and no optimizer except NM was able to converge to the energy minimum (data not shown).

Fig. 6 shows that with rescaling, all optimizers converge to the energy minimum, although L-BFGS-B and CG require 3–6 times more quantum circuit simulations than SLSQP and NM. Note, however, that the convergence to the energy minimum depends crucially on the good initial point; other random initial points produced much worse results (data not shown). Furthermore, it is worth noting that none of the optimizers comes across the point with the largest success probability (the cross in Fig. 5(b)); only



**Fig. 5.** (a)–(c) Energy landscape and (d)–(f) success probability landscape for QAOA with  $p = 1$  for the 30-qubit exact cover problem 30(0). The left column shows the landscapes obtained by scanning a  $64 \times 64$  grid  $\beta \in [0, \pi)$  and  $\gamma \in [0, 2\pi)$ . The middle column shows a zoom into the area around the minimum energy found in the scan. The largest success probability (cross) and the energy minimum (circle) in this area are indicated. Interestingly, these two points are not at the same location. The right column shows a zoom into another area of interest around  $\gamma \approx \pi$  and  $\beta \lesssim \pi$ .

NM and CG venture once into a region with better success probability (the right panels in Fig. 6). Obviously, this is not a flaw of the optimizers (which can only optimize for the energy in practice), but rather a deficiency of variational algorithms in general.

We take a number of QAOA steps  $p \leq 13$  and minimize the cost function  $\langle \beta, \gamma | H_C | \beta, \gamma \rangle$  w.r.t.  $\beta_k$  and  $\gamma_k$ , as one would do for QAOA running on genuine quantum hardware. The hope is then that by minimizing the cost function, we will also obtain relatively large values for the success probabilities. In Table 2, we present the results for a set of exact cover instances.

The QAOA results for  $p = 7, 13$  are encouraging in the sense that the success probabilities are relatively large, i.e., much larger than  $2^{-N}$  which would be the probability to pick the correct solution from a uniform distribution at random. However, the number of JUQCS-G calls required to obtain such values is also fairly large. The numbers in parentheses indicate the number of JUQCS-G calls corresponding to the highest observed success probability. Almost all runs have been terminated after 200 JUQCS-G calls (black entries in Table 2). Runs which were terminated by the minimizer (red entries) have a substantially lower success probability (smaller than 4%), suggesting that the minimizer became stuck in a local minimum.

We also performed some QAOA simulations with  $p = 3$ . We observed that for problem instance 30(0), the achieved success probability was smaller by a factor of 10–20 than in the  $p = 7$  and  $p = 13$  cases after using a similar number of JUQCS-G calls. For larger problem instances, the minimizer seemed to get stuck in local optima as the obtained success probabilities were smaller than 1%. We thus concluded that  $p = 3$  would be too small for larger problem instances and we did not proceed with  $p = 3$ .

### 3.3.2. AQA

A representative AQA result for a 40-variable exact cover problem is shown in Fig. 7. In this simulation, we chose  $n = 50$  and the time step  $\tau = 0.4$  ns, corresponding to a total annealing time of  $t_{\text{anneal}} = 20.4$  ns. This annealing time is very short compared to the annealing times commonly used by D-Wave quantum annealers (orders of  $\mu$ s).

In Table 3 we present the data of the AQA simulations with  $n = 50$  and  $\tau = 0.4$  ns for exact cover problems with 30, 32, ..., 40 variables. Column six of Table 3 shows that the success probability systematically decreases as the number of qubits increases. This

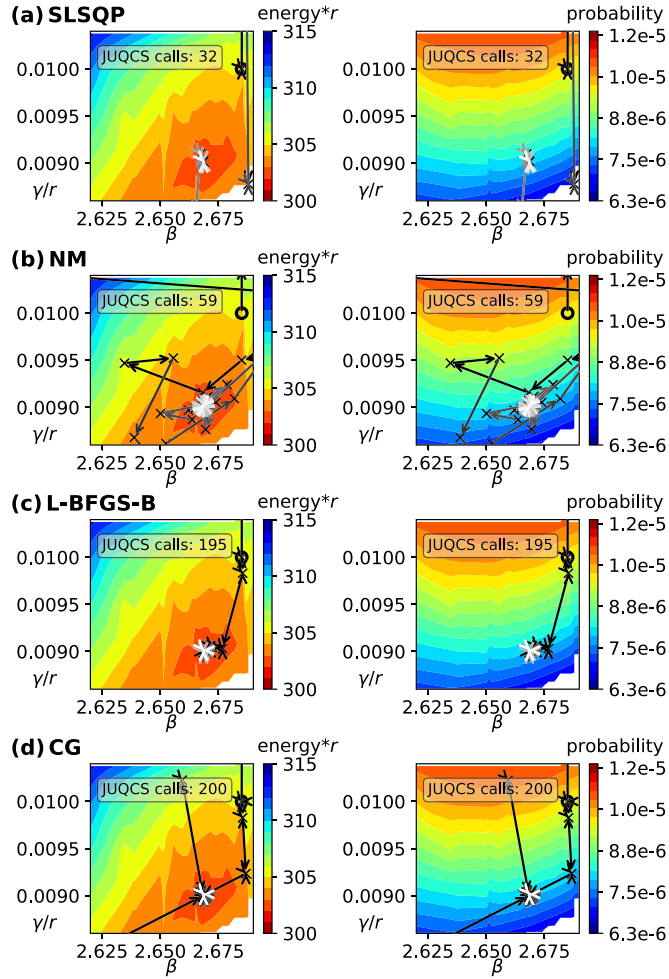
decrease is what one would expect on the basis of the Landau-Zener model and the assumption that the minimal spectral gap decreases with the system size. However, AQA uses a time step of  $\tau = 0.4$  ns that may actually be too large to justify an interpretation in terms of the Landau-Zener model. Table 3 also shows that the computational resources required for QSDS to perform these AQA simulations can be considerable.

As already observed earlier [14] and also observed in the AQA simulations, solving the TDSE for model parameters that pertain to D-Wave quantum annealers requires annealing times of the order of nanoseconds to obtain success probabilities of 1% or better. This observation leads to the conclusion that for the exact cover problems studied here, the annealing time required by TDSE solvers is much shorter than the typical annealing times used by D-Wave quantum annealers, which are of the order of microseconds (see also [97]). Of course, the TDSE simulations deal with a closed quantum system, free of the interactions with other degrees of freedom which are affecting the operation of real QA devices. Nevertheless, if technically possible, it would be of interest to perform this kind of very fast annealing on genuine quantum annealer hardware. Finally, it should be mentioned that the wall-clock time required by QSDS (or JUQCS) to cover the nanosecond time span is much larger than a few microseconds, see Table 3. Therefore, D-Wave quantum annealers are very fast simulators in comparison to the software simulators running on conventional semiconductor hardware.

### 3.3.3. Comparison of QAOA and AQA

Results to compare QAOA and AQA are presented in Fig. 8(a). It shows the success probability as a function of the number of AQA steps  $5 \leq n \leq 100$  for different values of  $\tau$ . Additionally, the arrows for  $p = 6$  and  $p = 13$  show the results obtained after optimizing the corresponding  $\beta$  and  $\gamma$  with the QAOA (using SLSQP) after 200 JUQCS-G calls. The initial and final values for  $\beta$  and  $\gamma$  are shown in Fig. 9.

We compare QAOA and AQA in terms of computational work. Performing QAOA with  $p$  steps and  $m$  optimization cycles (i.e.,  $m$  calls to JUQCS-G) needs computational work proportional to  $m \times p$ . Performing AQA with  $n$  is equivalent to performing QAOA with  $p = n + 1$  and  $m = 1$  (as AQA only needs a single call to JUQCS-G), so the computational work equivalent for AQA is  $n + 1$ . Thus, for QAOA



**Fig. 6.** Comparison of different classical optimizers used to optimize the variational parameters  $\beta$  and  $\gamma$  for QAOA with  $p = 1$  for the 30-qubit exact cover problem 30(0). The optimizers are (a) SLSQP, (b) NM, (c) L-BFGS-B, and (d) CG (see main text). The starting point for the optimization is the point with minimal energy found in the initial scan (black circles, corresponding to the black circles in Figs. 5(b) and (e)). For each optimizer, the left (right) panel shows the energy (success probability) landscape. The number of JUQCS-G calls used by each optimizer is indicated in the top left corner of each panel. The parameters for each call are shown with black crosses. The order in which the parameters are evaluated by the optimizers is indicated with arrows with colors evolving from black (for the first JUQCS-G call) to white (for the last JUQCS-G call). Note that, although the rescaled version of the problem was used for the optimization (see Eq. (14); here  $r = 36.75$ ), we plot  $\gamma/r$  and energy  $* r$  to make the scale comparable with the grid scan in Fig. 5.

to compete with AQA, it should use  $m < (n + 1)/p$  optimization cycles to reach a similar success probability.

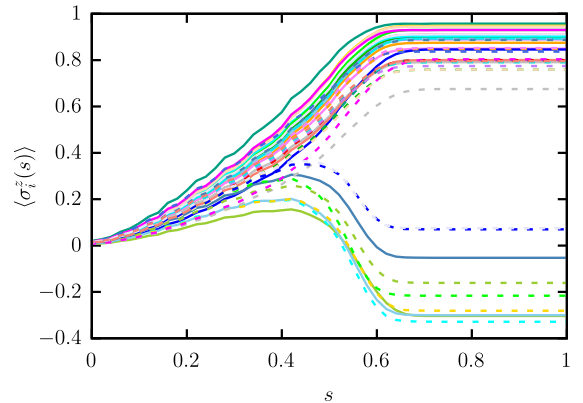
However, as already mentioned and seen in Table 2, the number of JUQCS-G calls  $m$  required to obtain high success probabilities with the QAOA is fairly large. And as Fig. 8(a) shows, sometimes even  $m = 200$  optimization cycles (with computational work  $200p$ ) are not enough to reach the success probabilities that AQA reaches already after  $n = 100$  steps. Hence, for the exact cover instances considered, QAOA cannot compete with AQA in terms of computational efficiency.

For AQA, we find that the success probability increases for increasing number of steps  $n$ . As Fig. 8(b) shows, the main increase in the success probability is due to the increased annealing time. The success probability also increases with  $\tau$ , up to a certain point where no further improvement is made. For  $\tau = 0.8$  ns, we find that for a fixed annealing time the probability is substantially lower than for the other values of  $\tau$ . Here, the time step  $\tau = 0.8$  ns

**Table 2**

**QAOA results for exact cover instances, obtained by minimizing the energy expectation value using SLSQP.** QAOA quantum gate circuits were executed using JUQCS-G. The success probability is determined by computing the probability of the ground state after each iteration and searching for the iteration number (given in parentheses) for which this probability is largest. The number of JUQCS-G calls was limited to 200. Red colored entries: The run was terminated by the minimizer that was probably stuck in a local minimum; black colored entries: The run was terminated when the 200 JUQCS-G calls had been reached. For  $p = 13$ , the calculations are too costly to warrant filling all missing entries.

qubits (instance)	success probability (JUQCS-G calls)	
	$p = 7$	$p = 13$
30(0)	0.3398 (165)	0.6214 (187)
30(3)	0.3708 (196)	-
32(0)	0.2841 (195)	-
32(3)	0.2745 (192)	0.4741 (193)
34(0)	0.1924 (190)	-
34(3)	0.2251 (196)	0.5075 (187)
36(0)	0.1081 (191)	-
36(3)	0.1545 (175)	0.0387 (94)
38(0)	0.0901 (187)	-
38(3)	0.1200 (174)	0.0159 (124)
40(0)	0.0068 (71)	0.0088 (123)
40(3)	0.0061 (38)	-



**Fig. 7.** AQA results for the 40-variable exact cover instance 40(0). Shown are the spin expectation values  $\langle \sigma_i^z(s) \rangle$  during the time evolution generated by Eq. (5), as a function of the normalized annealing time  $s = t/t_{\text{anneal}}$ . Different lines correspond to different qubits  $i = 0, \dots, N - 1$  for  $N = 40$ . The success probability to generate the state corresponding to the solution of this instance is 0.038. The number of steps is  $n = 50$  and the time step is  $\tau = 0.4$  ns.

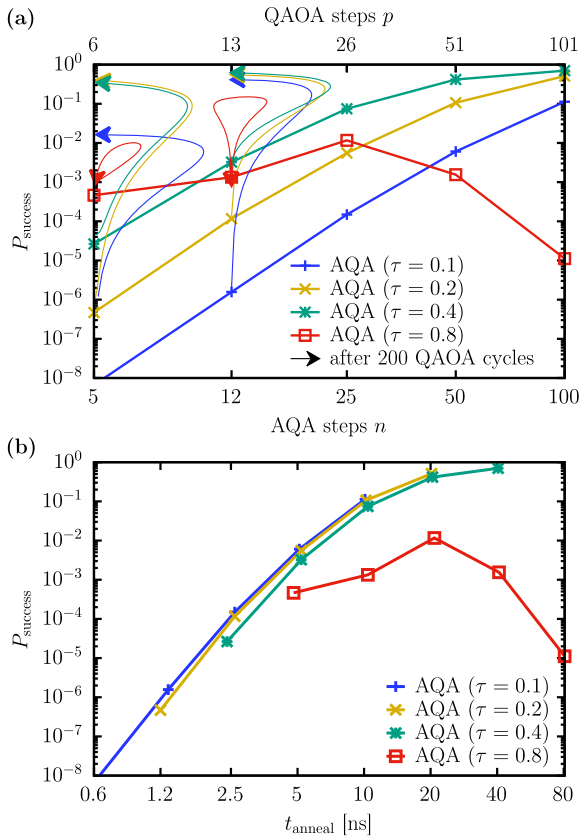
**Table 3**

**AQA results (success probabilities  $P_{\text{success}}$ ) obtained by solving the TDSE for Hamiltonians derived from exact cover instances.** Required hardware resources as well as the total run time  $t$  are also listed. The annealing scheme is obtained by discretizing the DW\_2000Q\_6 annealing scheme. QSDS was used with  $n = 50$  and  $\tau = 0.4$  ns, corresponding to an annealing time  $t_{\text{anneal}} = 20.4$  ns. All data was generated on JUWELS-CLUSTER [3], except column seven which lists the elapsed times  $t_{\text{FE}}$  that it took four A100 GPUs to solve the exact cover instances by full enumeration.

instance	nodes	processes	cores	$t$ [hh:mm]	$P_{\text{success}}$	$t_{\text{FE}}$
30(0)	64	1024	3072	00:08	0.417	1.7 s
32(3)	256	4096	12288	00:14	0.237	1.7 s
34(3)	256	4096	12288	00:52	0.193	2.4 s
36(3)	256	4096	12288	03:50	0.110	6.0 s
38(3)	256	4096	12288	16:40	0.085	22.3 s
40(0)	1024	16384	49152	24:40	0.038	91.8 s

is too large to justify even a crude approximation of an annealing schedule.

We believe that AQA is best seen as a viable heuristic, requiring a few numerical experiments to optimize the parameters (in contrast to the QAOA which usually needs many iterations to obtain a reasonable result). To some extent, Fig. 8(b) gives a hint for an explanation why AQA works well. For small time steps ( $\tau = 0.1$ ),

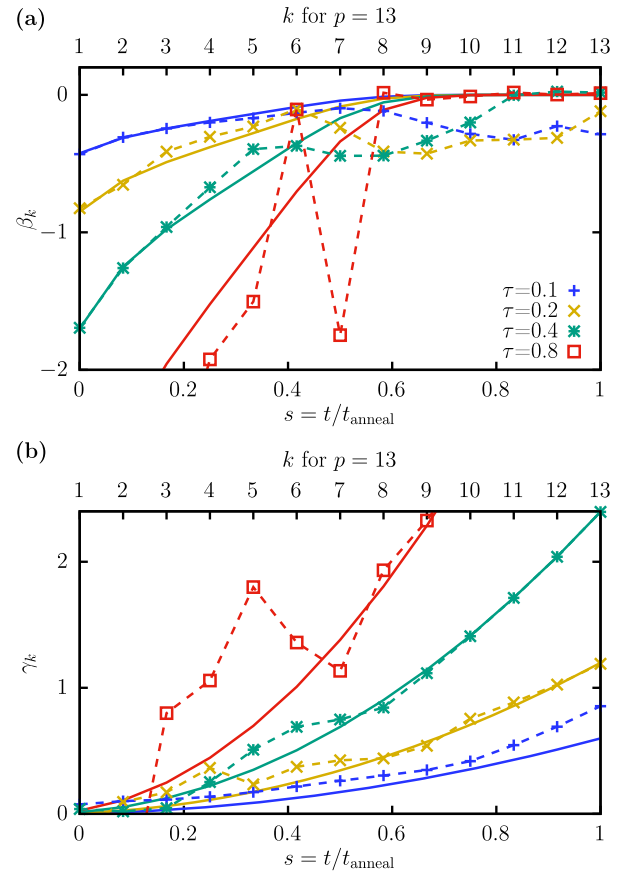


**Fig. 8. Comparison of AQA and QAOA, using exact cover instance 30(0).** (a) Success probability as a function of  $n$  (AQA, bottom axis) and  $p$  (QAOA, top axis). In terms of computational effort, AQA with  $n$  steps is equivalent to a single JUQCS-call for the QAOA circuit with  $p = n + 1$ , so they are shown together. AQA results are indicated by markers (lines are guides to the eye). QAOA results are indicated by the arrows showing the improvement due to the optimization of the variational parameters from the AQA initialization. (b) Success probability obtained by AQA as a function of “annealing time”  $t_{\text{anneal}} = (n + 1)\tau$  for different values of  $\tau$ . Lines are guides to the eye.

AQA is essentially the same as slow quantum annealing which, according to the adiabatic theorem, should give the ground state of the problem Hamiltonian. If the time step is too large ( $\tau = 0.8$ ), we lose contact with the idea of quantum annealing. We still get reasonable success probabilities for  $t < 20$ , but if we then anneal longer, the success probability drops. We have started a new project that specifically studies AQA to address this aspect.

The fact that QAOA is able to optimize the cases  $\tau \in \{0.1, 0.2, 0.4\}$  ns can be interpreted as follows: For  $\tau$  up to 0.4 ns, AQA still resembles QA with a very short annealing time (e.g. by rendering the system in a low energy state as indicated in [98]), so optimization can increase the success probability (as indicated by the arrows in Fig. 8(a)). The case  $\tau = 0.8$  ns, however, does not seem to yield suitable initial values for the parameters  $\beta_k$  and  $\gamma_k$  as is clear from the fact that the optimization during QAOA does not yield a significant improvement. However, for AQA with a small number of steps  $n$  (e.g.  $n = 5$  where  $P_{\text{success}} \approx 10^{-3}$  in Fig. 8(a)), it may still be a reasonable choice (see also the surprisingly good scaling in Fig. 11 below).

On the one hand, we find that with AQA for a large number of steps ( $n \approx 50$ – $100$ ), we obtain similar success probabilities as with QAOA for smaller  $p \approx 6$ – $13$ , but the QAOA optimization requires many calls to JUQCS-G. Moreover, we also observed that the minimizer can get stuck in a local optimum which then does not lead to an improved performance over AQA even for the same number of steps and many more circuit evaluations. However, also for AQA, we have to search for a good value of  $\tau$  which optimizes



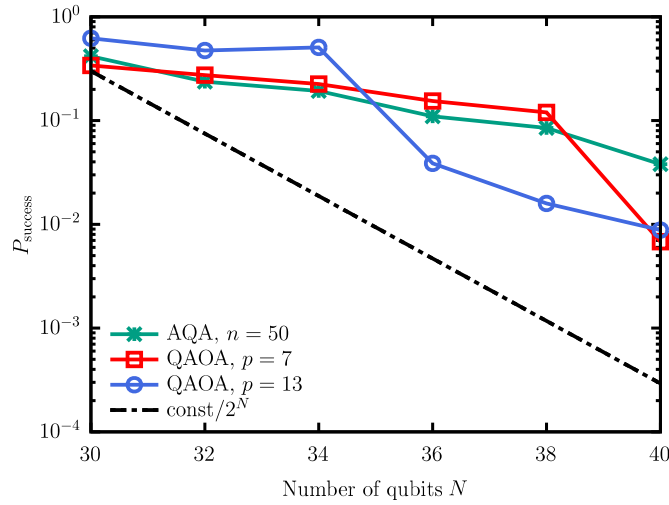
**Fig. 9. Visualization of the variational QAOA parameters (a)  $\beta_k$  and (b)  $\gamma_k$  for  $p = 13$  using exact cover instance 30(0).** Solid lines show the initial values taken from the second-order QA initialization (see Fig. 4 and Eqs. (15)–(17)), corresponding to the start of the lines with arrows in Fig. 8(a). Dashed lines show the final parameters after 200 QAOA optimization cycles (i.e., 200 JUQCS-G calls), corresponding to the end of the lines with arrows in Fig. 8(a). For  $\tau = 0.8$  ns (red squares), not the full set of final QAOA parameters is shown to keep the scale reasonable for the other cases and because Fig. 8(a) shows that the optimization brings no improvement in the performance. Note that the fact that the optimized  $\beta_k$  and  $\gamma_k$  still roughly follow the initialization from QA suggests that it was an effective modification of the annealing schedule that could so dramatically improve the success probability in Fig. 8(a).

the success probability for a given number of steps  $n$ . The same  $\tau$  that leads to an optimal success probability for a certain value of  $n$  may not be optimal for other values of  $n$ . Still, for AQA, we basically have to optimize a single parameter only (if  $n$  is fixed) and not  $2p$  parameters as is the case for QAOA.

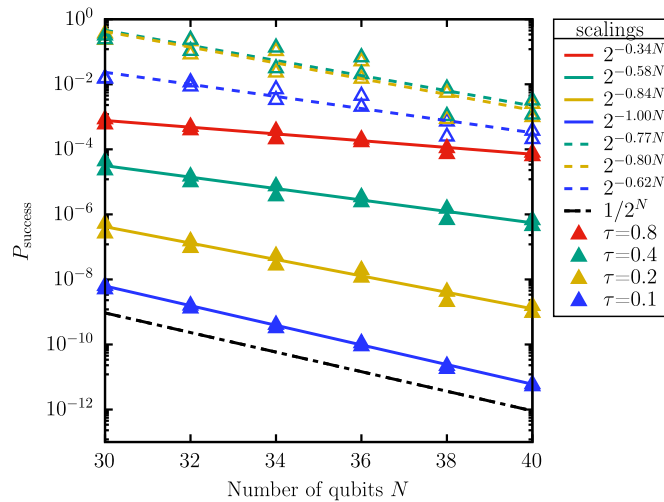
On the other hand, for NISQ devices, AQA with a large number of steps  $n$  (and equivalently QAOA with a large number of steps  $p$ ) will probably suffer from accumulated errors during the relatively long quantum circuit. Thus, NISQ devices may cope better with QAOA with small  $p$  than AQA with large  $n$ . Perhaps, building on the result that the optimized  $\beta_k$  and  $\gamma_k$  in Fig. 9 were not far from the annealing initialization, the best solution might be to indeed use AQA with small  $n$  but with better effective (maybe problem-dependent) annealing schedules. Comparing the performance of AQA and QAOA on NISQ devices in practice would be an interesting study for the future.

### 3.3.4. Scaling as a function of the problem size $N$

In Figs. 10 and 11, we show the scaling of the success probabilities obtained for different problem instances with increasing number of qubits using AQA and QAOA. In Fig. 10, the scalings of AQA and QAOA with the system size look quite similar up to  $N = 34$ . For larger  $N$ , the drops in the success probability for the



**Fig. 10.** Scaling of the success probability as a function of the system size  $N$ . The different markers correspond to AQA with  $n = 50$  and  $\tau = 0.4$  ns (green asterisks), QAOA for  $p = 7$  (red squares), and QAOA for  $p = 13$  (blue circles), taken from Tables 2 and 3. The dash-dotted line indicates the scaling of a uniform probability distribution. The green asterisk at  $N = 30$  is the same point shown in Fig. 8 at  $n = 50$ ,  $\tau = 0.4$  ns and  $t_{\text{anneal}} = 20.4$  ns. Lines are guides to the eye.



**Fig. 11.** Scaling of the success probability as a function of the system size  $N$ , using AQA with  $n = 5$  (filled triangles) and pre-optimized QAOA with  $p = 6$  (open triangles). Here, pre-optimized means that for all instances, the same values for  $\beta_k$  and  $\gamma_k$  (obtained from the optimization of the 30-qubit problem instance 30(0), see Fig. 8) are used. We ran two problem instances for each system size, so all triangles appear in pairs. Solid (dashed) lines show fits to the AQA (QAOA) results. Different colors correspond to different values for  $\tau$  as indicated in Fig. 8. The dash-dotted line indicates the probability to find the solution when picking from a uniform distribution at random. The data for all runs was obtained with JUQCS-G, using a quantum circuit that performs the time evolution simulated by QSDS (see Eq. (22)), thereby leveraging the computational power of the GPUs.

QAOA data are where the minimizer probably got stuck in a local optimum (red entries in Table 2). We note that for QAOA, we ran the optimization procedure for each system size. For AQA, we did not perform any optimization but we used a relatively large step size  $\tau$ .

In Fig. 11, we always use the same  $\beta_k$  and  $\gamma_k$  obtained from the QAOA optimization for problem instance 30(0). In other words, we take the variational parameters obtained by optimizing the 30-qubit instance 30(0), and we use the *same* parameters for *different* problem instances of *different* size. In this way, we test how well the effective “modified annealing schedule” (cf. Fig. 9) generalizes to other problems of larger size.

Fig. 11 shows that the QAOA parameters generalize systematically, but as expected, the success probability still drops exponentially with increasing qubit number. Furthermore, the fits to the data (lines) show that the exponential scaling is of the form  $2^{-\alpha N}$  for  $\alpha \gtrsim 0.6$ . In contrast, we observe that for AQA, although the drop in success probability is also exponential, the exponent  $\alpha$  behaves more favorably. Remarkably, this favorable scaling is especially pronounced for the large value of  $\tau = 0.8$  ns (where  $\alpha = 0.34$ ), which is very far in the AQA regime.

#### 4. Summary

The first part of this paper was devoted to the study of the weak and strong scaling behavior of a GPU-accelerated version (JUQCS-G) of the Jülich Universal Quantum Computer Simulator (JUQCS) [1] by performing benchmarks on JUWELS Booster, a supercomputer with 3744 NVIDIA A100 Tensor Core GPUs. Our data shows that JUQCS-G exhibits nearly perfect weak and strong scaling for systems up to 42 qubits. Comparing elapsed times for JUQCS-G and for JUQCS-E, a non-GPU version of JUQCS, shows that the former is a factor of 10–18 faster than the latter. As the number of qubits reaches the maximum that the available memory allows, the larger fraction of the elapsed time goes into MPI communication, for both the GPU and non-GPU version. In any case, using the GPU version significantly reduces the computing time required to simulate quantum computers and quantum systems.

In the second part of the paper, we have used JUQCS-G to solve exact cover problems with up to 40 variables (qubits). Hereby the focus was on the assessment of the potential of the quantum approximate optimization algorithm (QAOA) as a vehicle to solve optimization problems involving 30–40 qubits. Due to the minimization of parameters reflecting the variational nature of the QAOA, it is necessary to execute the quantum circuit many times. In most cases, at least for the 30–40 qubit instances that we have studied, the number of repetitions (with different sets of parameters) has a negative impact on the efficiency of the QAOA.

As an alternative, we also studied the performance of what we called approximate quantum annealing (AQA). AQA is a discretized version of quantum annealing which is approximate in the sense that we use only a few, relatively large time steps, possibly beyond the regime where quantum annealing is theoretically justified through the adiabatic theorem. Nevertheless, we found that, without any optimization, we already obtain success probabilities  $\gg 1\%$  for problem instances up to  $N = 40$  qubits. These promising results suggest that for future gate-based quantum computers which can cope with a larger circuit depth, direct AQA may provide a better alternative to the QAOA as it avoids the costly optimization procedure. As a matter of fact, from a computational viewpoint, AQA is much more efficient than the QAOA.

It is self-evident that all the simulation results that we have presented in this paper have been obtained by simulating the ideal mathematical model of a gate-based quantum computer. In this sense, the 30–40 qubit results presented in this paper are the “best case”, very unlikely to be achieved by using a real quantum processor. Of course, it is possible to incorporate noise and errors into our simulations (left for future work), but accounting for the intrinsic quantum gate errors of 30–40 qubit systems requires simulation times that are currently prohibitive [99]. Clearly, to get a view on the errors involved, it would be very interesting to run say a 30-qubit exact cover quantum circuit on a NISQ device and compare the experimental data with the simulation results. Furthermore, as our conclusions are drawn from results obtained for 30–40 variable exact cover problems, it might be of interest to investigate how generic these conclusions are by studying different types of optimization problems.

## Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

## Acknowledgements

The authors gratefully acknowledge the Gauss Centre for Supercomputing e.V. ([www.gauss-centre.eu](http://www.gauss-centre.eu)) for funding this project by providing computing time on the GCS Supercomputer JUWELS at Jülich Supercomputing Centre (JSC). We would like to thank M. Svensson for providing the exact cover problem instances. We gratefully acknowledge support during the JUWELS Booster Early Access period by A. Herten, M. Hrywniak, J. Kraus, A. Koehler, P. Messmer, M. Knobloch as well as the JUWELS Booster Project Team (JSC, Atos, ParTec, NVIDIA). D.W.'s work was partially supported by the Q(AI)<sup>2</sup> project. D.W. and M.W. acknowledge support from the project Jülich Unified Infrastructure for Quantum computing (JU-NIQ) that has received funding from the German Federal Ministry of Education and Research (BMBF) and the Ministry of Culture and Science of the State of North Rhine-Westphalia.

## Appendix A. Second-order initialization of the QAOA parameters

The QAOA state with  $2p$  variational parameters  $\beta_k$  and  $\gamma_k$  reads (see Eq. (7))

$$|\beta, \gamma\rangle = e^{-i\beta_p H_D} e^{-i\gamma_p H_C} \dots e^{-i\beta_1 H_D} e^{-i\gamma_1 H_C} |+\rangle^{\otimes N}. \quad (\text{A.1})$$

Inserting the values for the QAOA parameters given in Eqs. (15)–(17), and replacing  $|+\rangle^{\otimes N}$  by  $e^{i\tau A(s_1)H_D/2} |+\rangle^{\otimes N}$  (which only differs from  $|+\rangle^{\otimes N}$  by a global phase and is thus physically equivalent) yields

$$|\beta, \gamma\rangle = e^{i\tau A(s_p)H_D/2} e^{-i\tau B(s_p)H_C} \dots e^{i\tau A(s_2)+A(s_1)H_D/2} e^{-i\tau B(s_1)H_C} e^{i\tau A(s_1)H_D/2} |+\rangle^{\otimes N}. \quad (\text{A.2})$$

Here we see that  $|\beta, \gamma\rangle$  can be expressed as

$$|\beta, \gamma\rangle = U(s_p) \dots U(s_1) |+\rangle^{\otimes N}, \quad (\text{A.3})$$

where  $U(s_k)$  for  $k = 1, \dots, p$  is the second-order Suzuki-Trotter decomposition [51,52],

$$U(s_k) = e^{i\tau A(s_k)H_D/2} e^{-i\tau B(s_k)H_C} e^{i\tau A(s_k)H_D/2}, \quad (\text{A.4})$$

of the discretized time-evolution operator generated by the QA Hamiltonian  $H(s) = A(s)(-H_D) + B(s)H_C$  (see Eq. (5)). We note that besides the choice  $s_k = (k-1)/(p-1)$  taken in this paper, also the mid-point decomposition used in [14] is a good choice for the discretization (cf. [91]).

## References

- [1] H. De Raedt, F. Jin, D. Willsch, M. Willsch, N. Yoshioka, N. Ito, S. Yuan, K. Michielsen, *Comput. Phys. Commun.* 237 (2019) 47.
- [2] D. Willsch, H. Lagemann, M. Willsch, F. Jin, H. De Raedt, K. Michielsen, in: M. Müller, K. Binder, A. Trautmann (Eds.), *NIC Symposium 2020*, in: Publication Series of the John von Neumann Institute for Computing (NIC) NIC Series, vol. 50, Forschungszentrum Jülich GmbH Zentralbibliothek, Verlag, Jülich, 2020, pp. 255–264.
- [3] Jülich Supercomputing Centre, *J. Large-Scale Res. Facil.* 5 (2019) A135.
- [4] H. De Raedt, D. Willsch, Jülich Universal Quantum Computer Simulator (docker container), <https://jugit.fz-juelich.de/qip/juqcs-docker.git>, 2021.
- [5] E. Suarez, W. Frings, N. Attig, S. Achilles, J. De Amicis, T. Eickermann, E. Gregory, B. Hagemeier, A. Herten, J. Jitsev, D. Krause, E. Di Napoli, J. Meinke, K. Michielsen, B. Mohr, D. Pleiter, A. Strube, Th. Lippert, in: *NIC Symposium 2020*, in: Publication Series of the John von Neumann Institute for Computing (NIC) NIC Series, vol. 50, 2020, pp. 1–19.
- [6] K. De Raedt, K. Michielsen, H. De Raedt, B. Trieu, G. Arnold, M. Richter, T. Lippert, H. Watanabe, N. Ito, *Comput. Phys. Commun.* 176 (2007) 121.
- [7] F. Arute, K. Arya, R. Babbush, D. Bacon, J.C. Bardin, R. Barends, R. Biswas, S. Boixo, F.G.S.L. Brandao, D.A. Buell, B. Burkett, Y. Chen, Z. Chen, B. Chiaro, R. Collins, W. Courtney, A. Dunsworth, E. Farhi, B. Foxen, A. Fowler, C. Gidney, M. Giustina, R. Graff, K. Guerin, S. Habegger, M.P. Harrigan, M.J. Hartmann, A. Ho, M. Hoffmann, T. Huang, T.S. Humble, S.V. Isakov, E. Jeffrey, Z. Jiang, D. Kafri, K. Kechedzhi, J. Kelly, P.V. Klimov, S. Knysh, A. Korotkov, F. Kostritsa, D. Landhuis, M. Lindmark, E. Lucero, D. Lyakh, S. Mandrà, J.R. McClean, M. McEwen, A. Megrant, X. Mi, K. Michielsen, M. Mohseni, J. Mutus, O. Naaman, M. Neeley, C. Neill, M.Y. Niu, E. Ostby, A. Petukhov, J.C. Platt, C. Quintana, E.G. Rieffel, P. Roushan, N.C. Rubin, D. Sank, K.J. Satzinger, V. Smelyanskiy, K.J. Sung, M.D. Trevithick, A. Vainsencher, B. Villalonga, T. White, Z.J. Yao, P. Yeh, A. Zalcman, H. Neven, J.M. Martinis, *Nature* 574 (2019) 505.
- [8] E. Farhi, J. Goldstone, S. Gutmann, A quantum approximate optimization algorithm, *arXiv:1411.4028*, 2014.
- [9] E. Farhi, A.W. Harrow, Quantum supremacy through the quantum approximate optimization algorithm, *arXiv:1602.07674*, 2016.
- [10] J. Preskill, *Quantum* 2 (2018) 79.
- [11] Z. Wang, S. Hadfield, Z. Jiang, E.G. Rieffel, *Phys. Rev. A* 97 (2018) 022304.
- [12] J.S. Otterbach, R. Manenti, N. Alidoust, A. Bestwick, M. Block, B. Bloom, S. Caldwell, N. Didier, E.S. Fried, S. Hong, P. Karalekas, C.B. Osborn, A. Papageorge, E.C. Peterson, G. Prawiroatmodjo, N. Rubin, C.A. Ryan, D. Scarabelli, M. Scheer, E.A. Sete, P. Sivarajah, R.S. Smith, A. Staley, N. Tezak, W.J. Zeng, A. Hudson, B.R. Johnson, M. Reagor, M.P. da Silva, C. Rigetti, *Unsupervised machine learning on a hybrid quantum computer*, *arXiv:1712.05771*, 2017.
- [13] X. Qiang, X. Zhou, J. Wang, C.M. Wilkes, T. Loke, S. O'Gara, L. Kling, G.D. Marshall, R. Santagati, T.C. Ralph, J.B. Wang, J.L. O'Brien, M.G. Thompson, J.C.F. Matthews, *Nat. Photonics* 12 (2018) 534.
- [14] M. Willsch, D. Willsch, F. Jin, H. De Raedt, K. Michielsen, *Quantum Inf. Process.* 19 (2020) 197.
- [15] P. Vikstål, M. Grönkvist, M. Svensson, M. Andersson, G. Johansson, G. Ferrini, *Phys. Rev. Appl.* 14 (2020) 034009.
- [16] A. Bengtsson, P. Vikstål, C. Warren, M. Svensson, X. Gu, A.F. Kockum, P. Krantz, C. Krizan, D. Shiri, I.-M. Svensson, G. Tancredi, G. Johansson, P. Delsing, G. Ferrini, J. Bylander, *Phys. Rev. Appl.* 14 (2020) 034010.
- [17] N.acroix, C. Hellings, C.K. Andersen, A. Di Paolo, A. Remm, S. Lazar, S. Krinner, G.J. Norris, M. Gabureac, J. Heinsoo, A. Blais, C. Eichler, A. Wallraff, *PRX Quantum* 1 (2020) 110304.
- [18] G. Pagano, A. Bapat, P. Becker, K.S. Collins, A. De, P.W. Hess, H.B. Kaplan, A. Kyprianidis, W.L. Tan, C. Baldwin, L.T. Brady, A. Deshpande, F. Liu, S. Jordan, A.V. Gorshkov, C. Monroe, *Proc. Natl. Acad. Sci.* 117 (2020) 25396, <https://www.pnas.org/content/117/41/25396.full.pdf>.
- [19] L. Zhou, S.-T. Wang, S. Choi, H. Pichler, M.D. Lukin, *Phys. Rev. X* 10 (2020) 021067.
- [20] V. Akshay, H. Philathong, M.E.S. Morales, J.D. Biamonte, *Phys. Rev. Lett.* 124 (2020) 090504.
- [21] M.P. Harrigan, K.J. Sung, M. Neeley, K.J. Satzinger, F. Arute, K. Arya, J. Atalaya, J.C. Bardin, R. Barends, S. Boixo, M. Broughton, B.B. Buckley, D.A. Buell, B. Burkett, N. Bushnell, Y. Chen, Z. Chen, B. Chiaro, R. Collins, W. Courtney, S. Demura, A. Dunsworth, D. Eppens, A. Fowler, B. Foxen, C. Gidney, M. Giustina, R. Graff, S. Habegger, A. Ho, S. Hong, T. Huang, L.B. Ioffe, S.V. Isakov, E. Jeffrey, Z. Jiang, C. Jones, D. Kafri, K. Kechedzhi, J. Kelly, S. Kim, P.V. Klimov, A.N. Korotkov, F. Kostritsa, D. Landhuis, P. Laptev, M. Lindmark, M. Leib, O. Martin, J.M. Martinis, J.R. McClean, M. McEwen, A. Megrant, X. Mi, M. Mohseni, W. Mruczkiewicz, J. Mutus, O. Naaman, C. Neill, F. Neukart, M.Y. Niu, T.E. O'Brien, B. O'Gorman, E. Ostby, A. Petukhov, H. Putterman, C. Quintana, P. Roushan, N.C. Rubin, D. Sank, A. Skolik, V. Smelyanskiy, D. Strain, M. Streif, M. Szalay, A. Vainsencher, T. White, Z.J. Yao, P. Yeh, A. Zalcman, L. Zhou, H. Neven, D. Bacon, E. Lucero, E. Farhi, R. Babbush, *Nat. Phys.* (2021), <https://doi.org/10.1038/s41567-020-01105-y>.
- [22] M. Fernández-Pendás, E.F. Combarro, S. Vallecorsa, J. Ranilla, I.F. Rúa, *J. Comput. Appl. Math.* 113388 (2021).
- [23] M. Medvidović, G. Carleo, *npj Quantum Inf.* 7 (2021) 101.
- [24] M. Streif, M. Leib, *Quantum Sci. Technol.* 5 (2020) 034008.
- [25] S.H. Sack, M. Serbyn, Quantum annealing initialization of the quantum approximate optimization algorithm, *arXiv:2101.05742 [quant-ph]*, 2021.
- [26] B. Apolloni, C. Carvalho, D. de Falco, *Stoch. Process. Appl.* 33 (1989) 233.
- [27] A. Fennila, M. Gomez, C. Sebenik, C. Stenson, J. Doll, *Chem. Phys. Lett.* 219 (1994) 343.
- [28] T. Kadowaki, H. Nishimori, *Phys. Rev. E* 58 (1998) 5355.
- [29] J. Brooke, D. Bitko, T.F. Rosenbaum, G. Aeppli, *Science* 284 (1999) 779.
- [30] R. Harris, M.W. Johnson, T. Lanting, A.J. Berkley, J. Johansson, P. Bunyk, E. Tolkacheva, E. Ladizinsky, N. Ladizinsky, T. Oh, F. Cioata, I. Perminov, P. Spear, C. Enderud, C. Rich, S. Uchaikin, M.C. Thom, E.M. Chapple, J. Wang, B. Wilson, M.H.S. Amin, N. Dickson, K. Karimi, B. Macready, C.J.S. Truncik, G. Rose, *Phys. Rev. B* 82 (2010) 024511.
- [31] M.W. Johnson, M.H.S. Amin, S. Gildert, T. Lanting, F. Hamze, N. Dickson, R. Harris, A.J. Berkley, J. Johansson, P. Bunyk, E.M. Chapple, C. Enderud, J.P. Hilton, K. Karimi, E. Ladizinsky, N. Ladizinsky, T. Oh, I. Perminov, C. Rich, M.C. Thom, E.

- Tolkacheva, C.J.S. Truncik, S. Uchaikin, J. Wang, B. Wilson, G. Rose, *Nature* 473 (2011) 194.
- [32] P.I. Bunyk, E.M. Hoskinson, M.W. Johnson, E. Tolkacheva, F. Altomare, A.J. Berkley, R. Harris, J.P. Hilton, T. Lanting, A.J. Przybysz, J. Whittaker, *IEEE Trans. Appl. Supercond.* 24 (2014) 1.
- [33] J. Job, D. Lidar, *Quantum Sci. Technol.* 3 (2018) 030501.
- [34] P. Hauke, H.G. Katzgraber, W. Lechner, H. Nishimori, W.D. Oliver, *Rep. Prog. Phys.* 83 (2020) 054401.
- [35] R.K. Nath, H. Thapliyal, T.S. Humble, *SN Comput. Sci.* 2 (2021) 365.
- [36] M. Nielsen, I. Chuang, *Quantum Computation and Quantum Information*, 10th ed., Cambridge University Press, Cambridge, 2010.
- [37] C. McGeoch, P. Farré, The D-Wave Advantage System: an Overview, Tech. Rep., D-Wave Systems Inc, Burnaby, BC, Canada, 2020, D-Wave Technical Report Series 14-1049A-A.
- [38] H.S. Bhatia, F. Phillipson, in: M. Paszynski, D. Kranzlmüller, V.V. Krzhizhanovskaya, J.J. Dongarra, P.M.A. Sloot (Eds.), *Computational Science – ICCS 2021*, Springer International Publishing, Cham, 2021, pp. 84–97.
- [39] F. Phillipson, R.S. Wezeman, I. Chiscop, *Computers* 10 (2021) 71.
- [40] D. Willsch, M. Willsch, H. De Raedt, K. Michielsen, *Comput. Phys. Commun.* 248 (2020) 107006.
- [41] J. Cohen, C. Alexander, Picking efficient portfolios from 3,171 US common stocks with new quantum and classical solvers, arXiv:2011.01308 [quant-ph], 2020.
- [42] T. Birdal, V. Golyanik, C. Theobalt, L. Guibas, Quantum permutation synchronization, arXiv:2101.07755 [quant-ph], 2021.
- [43] S.A. Rahman, R. Lewis, E. Mendicelli, S. Powell, SU(2) lattice gauge theory on a quantum annealer, arXiv:2103.08661 [hep-lat], 2021.
- [44] A.D. King, W. Bernoudy, Performance benefits of increased qubit connectivity in quantum annealing 3-dimensional spin glasses, arXiv:2009.12479 [quant-ph], 2020.
- [45] C.D. Gonzalez Calaza, D. Willsch, K. Michielsen, *Quantum Inf. Process.* 20 (2021) 305.
- [46] D. Willsch, M. Willsch, C.D.G. Calaza, F. Jin, H. De Raedt, M. Svensson, K. Michielsen, Benchmarking advantage and D-wave 2000Q quantum annealers with exact cover problems, arXiv:2105.02208 [quant-ph], 2021.
- [47] M. Heyl, P. Hauke, P. Zoller, *Sci. Adv.* 5 (2019) eaau8342.
- [48] L.M. Sieberer, T. Olsacher, A. Elben, M. Heyl, P. Hauke, F. Haake, P. Zoller, *npj Quantum Inf.* 5 (2019) 78.
- [49] H.F. Trotter, *Proc. Am. Math. Soc.* 10 (1959) 545.
- [50] M. Suzuki, *Commun. Math. Phys.* 51 (1976) 83.
- [51] H. De Raedt, B. De Raedt, *Phys. Rev. A* 28 (1983) 3575.
- [52] M. Suzuki, *J. Math. Phys.* 26 (1985) 601, <https://doi.org/10.1063/1.526596>.
- [53] M. Streif, M. Leib, Comparison of QAOA with quantum and simulated annealing, arXiv:1901.01903 [quant-ph], 2019.
- [54] Message Passing Interface Forum, MPI: A Message-Passing Interface Standard Version 4.0, 2021.
- [55] D. Deutsch, A. Barenco, A. Ekert, *Proc. R. Soc. Lond. A* 449 (1995) 669.
- [56] D.P. DiVincenzo, *Phys. Rev. A* 51 (1995) 1015.
- [57] E. Pednault, J.A. Gunnels, G. Nannicini, L. Hoesch, T. Magerlein, E. Solomonik, R. Wisnieff, Breaking the 49-qubit barrier in the simulation of quantum circuits, arXiv:1710.05867, 2017.
- [58] S. Boixo, S.V. Isakov, V.N. Smelyanskiy, H. Neven, Simulation of low-depth quantum circuits as complex undirected graphical models, arXiv:1712.05384, 2017.
- [59] Z. Chen, Q. Zhou, C. Xue, X. Yang, G. Guo, G. Guo, *Sci. Bull.* 964 (2018).
- [60] I.L. Markov, A. Fatima, S.V. Isakov, S. Boixo, Quantum supremacy is both closer and farther than it appears, arXiv:1807.10749, 2018.
- [61] B. Villalonga, D. Lyakh, S. Boixo, H. Neven, T.S. Humble, R. Biswas, E.G. Rieffel, A. Ho, S. Mandrà, Establishing the quantum supremacy frontier with a 281 Pflop/s simulation, arXiv:1905.00444, 2019.
- [62] B. Villalonga, S. Boixo, B. Nelson, C. Henze, E. Rieffel, R. Biswas, S. Mandrà, *npj Quantum Inf.* 5 (2019), <https://doi.org/10.1038/s41534-019-0196-1>.
- [63] Y. Zhou, E.M. Stoudenmire, X. Waintal, *Phys. Rev. X* 10 (2020) 041038.
- [64] F. Pan, P. Zhang, Simulating the Sycamore quantum supremacy circuits, arXiv:2103.03074 [quant-ph], 2021.
- [65] K. Michielsen, M. Nocon, D. Willsch, F. Jin, Th. Lippert, H. De Raedt, *Comput. Phys. Commun.* 220 (2017) 44.
- [66] D-Wave Systems, QPU-specific anneal schedules, [https://docs.dwavesys.com/docs/latest/doc\\_physical\\_properties.html](https://docs.dwavesys.com/docs/latest/doc_physical_properties.html).
- [67] E. Farhi, J. Goldstone, S. Gutmann, M. Sipser, Quantum computation by adiabatic evolution, arXiv:quant-ph/0001106, 2000.
- [68] A.M. Childs, E. Farhi, J. Preskill, *Phys. Rev. A* 65 (2001) 012322.
- [69] N.G. Dickson, M.W. Johnson, M.H. Amin, R. Harris, F. Altomare, A.J. Berkley, P. Bunyk, J. Cai, E.M. Chapple, P. Chavez, F. Cioata, T. Cirip, P. deBuen, M. Drew-Brook, C. Enderud, S. Gildert, F. Hamze, J.P. Hilton, E. Hoskinson, K. Karimi, E. Ladizinsky, N. Ladizinsky, T. Lanting, T. Mahon, R. Neufeld, T. Oh, I. Perminov, C. Petróff, A. Przybysz, C. Rich, P. Spear, A. Tcaciuc, M.C. Thom, E. Tolkacheva, S. Uchaikin, J. Wang, A.B. Wilson, Z. Merali, G. Rose, *Nat. Commun.* 4 (2013) 1903.
- [70] Z. Bian, F. Chudak, R. Israel, B. Lackey, W.G. Macready, A. Roy, *Front. Phys.* 2 (2014) 56.
- [71] M.H. Amin, *Phys. Rev. A* 92 (2015) 052323.
- [72] A. Mishra, T. Albash, D.A. Lidar, *Nat. Commun.* 9 (2018) 2917.
- [73] J. Marshall, D. Venturelli, I. Hen, E.G. Rieffel, *Phys. Rev. Appl.* 11 (2019) 044083.
- [74] A. Pearson, A. Mishra, I. Hen, D.A. Lidar, *npj Quantum Inf.* 5 (2019) 107.
- [75] P. Weinberg, M. Tylutki, J.M. Rönkkö, J. Westerholm, J.A. Åström, P. Manninen, P. Törmä, A.W. Sandvik, *Phys. Rev. Lett.* 124 (2020) 090502.
- [76] S. Hadfield, Z. Wang, B. O’Gorman, E.G. Rieffel, D. Venturelli, R. Biswas, *Algorithms* 12 (2019) 34.
- [77] L. Zhu, H.L. Tang, G.S. Barron, N.J. Mayhall, E. Barnes, S.E. Economou, An adaptive quantum approximate optimization algorithm for solving combinatorial problems on a quantum computer, arXiv:2005.10258, 2020.
- [78] G.E. Crooks, Performance of the quantum approximate optimization algorithm on the maximum cut problem, arXiv:1811.08419v1, 2018.
- [79] F.G.S.L. Brandao, M. Broughton, E. Farhi, S. Gutmann, H. Neven, For fixed control parameters the quantum approximate optimization algorithm’s objective function value concentrates for typical instances, arXiv:1812.04170, 2018.
- [80] E. Farhi, J. Goldstone, S. Gutmann, L. Zhou, The quantum approximate optimization algorithm and the Sherrington-Kirkpatrick model at infinite size, arXiv:1910.08187, 2019.
- [81] R.M. Karp, in: R.E. Miller, J.W. Thatcher, J.D. Bohlinger (Eds.), *Complexity of Computer Computations: Proceedings of a Symposium on the Complexity of Computer Computations*, held March 20–22, 1972, at the IBM Thomas J. Watson Research Center, Yorktown Heights, New York, and sponsored by the Office of Naval Research, Mathematics Program, IBM World Trade Corporation, and the IBM Research Mathematical Sciences Department, Springer US, Boston, MA, 1972, pp. 85–103.
- [82] E. Farhi, J. Goldstone, S. Gutmann, J. Lapan, A. Lundgren, D. Preda, *Science* 292 (2001) 472.
- [83] V. Choi, Adiabatic quantum algorithms for the NP-complete maximum-weight independent set, exact cover and 3SAT problems, arXiv:1004.2226 [quant-ph], 2010.
- [84] A. Lucas, *Front. Phys.* 2 (2014) 5.
- [85] Y. Cao, S. Jiang, D. Perouli, S. Kais, *Sci. Rep.* 6 (2016) 33957.
- [86] I. Sax, S. Feld, S. Zielinski, T. Gabor, C. Linnhoff-Popien, W. Maurer, in: *Proceedings of the 17th ACM International Conference on Computing Frontiers*, CF ’20, Association for Computing Machinery, New York, NY, USA, 2020, pp. 108–117.
- [87] A. Ernst, H. Jiang, M. Krishnamoorthy, D. Sier, *Eur. J. Oper. Res.* 153 (2004) 3, timetabling and Rostering.
- [88] A. Tahir, G. Desaulniers, I. El Hallaoui, *EURO J. Transp. Logist.* 8 (2019) 713.
- [89] D-Wave Systems, D-Wave Solver Properties and Parameters, Tech. Rep., D-Wave Systems Inc., Burnaby, BC, Canada, 2021, D-Wave User Manual 09-1169A-S.
- [90] E. Jones, T. Oliphant, P. Peterson, et al., *SciPy: Open source scientific tools for Python*, 2001.
- [91] M. Suzuki, *Proc. Jpn. Acad. Ser. B* 69 (1993) 161.
- [92] H. De Raedt, K. Michielsen, in: M. Rieth, W. Schommers (Eds.), *Handbook of Theoretical and Computational Nanotechnology*, American Scientific Publishers, Los Angeles, 2006, pp. 2–48.
- [93] J.A. Nelder, R. Mead, *Comput. J.* 7 (1965) 308.
- [94] C. Zhu, R.H. Byrd, P. Lu, J. Nocedal, *ACM Trans. Math. Softw.* 23 (1997) 550.
- [95] J. Morales, J. Nocedal, *ACM Trans. Math. Softw.* 38 (2011) 7.
- [96] W.H. Press, S.A. Teukolsky, W.T. Vetterling, B.P. Flannery, *Numerical Recipes: The Art of Scientific Computing*, 3rd edition, Cambridge University Press, New York, USA, 2007.
- [97] M. Willsch, D. Willsch, F. Jin, H. De Raedt, K. Michielsen, *Phys. Rev. A* 101 (2020) 012327.
- [98] T.-J. Hsu, F. Jin, C. Seidel, F. Neukart, H. De Raedt, K. Michielsen, *Commun. Comput. Phys.* 26 (2019) 928.
- [99] D. Willsch, Supercomputer simulations of transmon quantum computers, Ph.D. thesis, RWTH Aachen University, Aachen, 2020.