

This is the accepted manuscript made available via CHORUS. The article has been published as:

## Hybrid annealing: Coupling a quantum simulator to a classical computer

Tobias Graß and Maciej Lewenstein

Phys. Rev. A **95**, 052309 — Published 3 May 2017

DOI: [10.1103/PhysRevA.95.052309](https://doi.org/10.1103/PhysRevA.95.052309)

# Hybrid annealing: coupling a quantum simulator to a classical computer

Tobias Graß<sup>1,2</sup> and Maciej Lewenstein<sup>2,3</sup>

<sup>1</sup>*Joint Quantum Institute, University of Maryland, College Park, MD 20742, U.S.A.*

<sup>2</sup>*ICFO-Institut de Ciències Fotòniques, The Barcelona Institute of Science and Technology, 08860 Castelldefels, Spain and*

<sup>3</sup>*ICREA-Institució Catalana de Recerca i Estudis Avançats, Lluís Companys 23, 08010 Barcelona, Spain*

Finding the global minimum in a rugged potential landscape is a computationally hard task, often equivalent to relevant optimization problems. Annealing strategies, either classical or quantum, explore the configuration space by evolving the system under the influence of thermal or quantum fluctuations. The thermal annealing dynamics can rapidly freeze the system into a low-energy configuration, and it can well be simulated on a classical computer, but it easily gets stuck in local minima. Quantum annealing, on the other hand, can be guaranteed to find the true ground state, and can be implemented in modern quantum simulators, however quantum-adiabatic schemes become prohibitively slow in the presence of quasidegeneracies. Here we propose a strategy which combines ideas from simulated annealing and quantum annealing. In such hybrid algorithm, the outcome of a quantum simulator is processed on a classical device. While the quantum simulator explores the configuration space by repeatedly applying quantum fluctuations and performing projective measurements, the classical computer evaluates each configuration and enforces a lowering of the energy. We have simulated this algorithm for small instances of the random energy model, showing that it potentially outperforms both simulated thermal annealing and adiabatic quantum annealing. It becomes most efficient for problems involving many quasi-degenerate ground states.

## I. INTRODUCTION

Constructing a quantum computer has been an outstanding goal and immense driving force for physical research during the past decades. In recent years, a series of notable advances has been made towards this goal: Using superconducting qubits several companies, such as D-Wave [1], IBM [2, 3], and Google [4], have developed information-processing devices. Academic research has produced prototypes of ion-based quantum computers [5, 6]. Other systems which may allow for performing quantum computations are neutral atoms, such as Rydberg atoms [7], atoms in cavities [8] or photonic crystals [9], as well as quantum dots or artificial atoms implemented as defects in a solid-state system [10, 11]. Apart from finding the best amongst these different hardware architectures, research also has to determine the algorithmic design of future quantum computers. Universal quantum computers, as realized in Refs. [4–6] or even accessible via a cloud service [3], break any computational task into simple gate operations. While this is a powerful approach, in particular as it also admits quantum error correcting codes [12], the best scalability so far has been achieved with so-called analog quantum computers [1]. These are specific-purpose devices, designed to address one problem of interest. As there are mappings between different problems, such special-purpose computer would still be useful for addressing a whole class of problems. For instance, if an analog quantum computer can solve just a single NP-complete problem in polynomial time, by the definition of NP-completeness any NP problem would become solvable in polynomial time. Such quantum device would then outperform any known classical algorithm for a diverse variety of problems, containing the traveling salesman problem, various satisfiability problems, graph-coloring problems, the number partitioning

problem, to name just a few.

Typically, analog quantum computers implement an algorithm known as quantum annealing [13] or adiabatic quantum computing [14]. The idea behind this is the following: Starting from an initial Hamiltonian with an easy-to-prepare ground state one can reach the ground state of a complex target Hamiltonian by slowly varying a Hamiltonian parameter. The target Hamiltonian should be chosen such that its ground state provides the solution to the computational problem of interest. Here, one can exploit the NP-completeness of Ising spin glass models [15, 16], allowing for any NP problem to be mapped onto an Ising model. In this Ising picture, the complicated target Hamiltonian is usually approached from an initial Hamiltonian, in which a strong magnetic field simply polarizes all spins in a transverse direction.

However, the computational time needed for the adiabatic computation increases quadratically with the inverse energy gap. For several relevant models, the gap has been shown to become exponentially small on the annealing path [17, 18]. In these cases, approaches which do not rely on the adiabatic principle are desirable. In this paper, we describe an annealing scheme which avoids the requirement of adiabaticity. Instead, we envisage a hybrid solution where classical information processing is paired with quantum technologies, see Fig. 1(b). More concretely, the quantum part of our protocol consists of (i) the preparation of an input state chosen from the computational basis, (ii) a short unitary time evolution of this state, performed by a quantum simulator which implements the target Hamiltonian exposed to quantum fluctuations, (iii) a measurement which projects the evolved quantum state onto a state from the computational basis. The outcome of this measurement is then processed to a classical computer which evaluates the configuration, and determines the next input state. In the simplest version,




FIG. 1: (a) In the depicted energy landscape, the global minimum at  $C$  can be reached from  $A$  via a combination of thermal annealing and quantum tunneling. The high potential barrier from  $A$  to  $B$  can be tunneled quantum-mechanically, and the wide barrier between  $B$  and  $C$  can be surpassed by thermal fluctuations. (b) One annealing cycle in the hybrid algorithm: given an input state from the computational basis, a quantum simulator searches for nearby minima by time-evolving the state in the presence of quantum fluctuations. After a projective measurement, a classical device evaluates the energy of the output state. Based on this evaluation, which might also take into account the result of a simultaneously performed simulated annealing (SA) scheme, a new input state for the next cycle is defined.

this evaluation may be restricted to comparing the energy in the initial and the final configuration. Repeating these steps allows for systematically lowering the energy of the system.

Such protocol resembles a simulated annealing algorithm which was pioneered in Ref. [19]. Simulated annealing mimics the evolution of a classical system in the presence of thermal noise, for instance implemented by random spin flips. Updates are accepted or rejected according to a probability distribution characterized by some temperature. Slowly decreasing the temperature maximizes the probability of reaching a low-energy state.

In the protocol described above, thermal noise is replaced by quantum fluctuations. As illustrated in Fig. 1(a), such an algorithm based on quantum fluctuations might surpass high energy barriers for which the likelihood of escaping is strongly suppressed in the case of thermal annealing. In contrast, surpassing wide energy barriers might be no problem for thermal noise, and combining both quantum and thermal fluctuations would enhance the chances of finding the global minimum. In the scheme described above, such combination can straightforwardly be achieved by running a classical simulated annealing algorithm in parallel to the quantum simulation.

As a hardware requirement to run such algorithm one needs a quantum simulator which implements the Hamiltonian of interest (e.g. a classical Ising-like model), and exposes it to quantum fluctuations (e.g. to a transverse magnetic field). Moreover, one should be able to initialize the quantum simulator in an arbitrary classical state, and to project the evolved quantum state back onto a state from the classical basis via a projective measurement. All information processing is done by a classical computer, that is, evaluating the energies (corresponding to the cost function of the optimization problem), deciding about whether to reject or accept the measured state, etc. The mentioned requirements to the quantum simulator are already met, for instance, in existing trapped ion experiments [20, 21]. There also exist proposals how to implement NP-complete problems in an ionic quantum simulator [22, 23].

In this manuscript, we numerically test the envisaged protocol. To do so, we focus on a random energy model, for which simulated annealing fails because similar spin configurations may have completely different energies. In such random energy landscape, also adiabatic quantum annealing has been shown to suffer from an exponentially small gap [17]. After introducing the model, and describing the details of our algorithm, we simulate the algorithm in small Hilbert spaces of dimension up to  $2^{12}$ , and compare it to other annealing strategies. For the random energy model, our annealing scheme, even if it is restricted to quantum fluctuations only, outperforms the simulated thermal annealing. It can further be improved by taking into account also thermal fluctuations. However, for the small-sized problems considered here, the quantum adiabatic annealing turns out to be the best choice. The situation changes if we artificially introduce quasi-degeneracies in the energy landscape. While this even enhances the performance of our hybrid algorithm, it prevents the adiabatic scheme from finding the correct state. In a final discussion, we point out the main advantages of the algorithm.

## II. MODEL AND METHOD

In the random energy model a  $D$ -dimensional Hilbert space, spanned by states  $|\alpha\rangle$ , is associated with a random energy landscape,  $H_0 = \sum_{\alpha} E_{\alpha} |\alpha\rangle \langle\alpha|$ , where the

$E_\alpha$  are chosen from some probability distribution. Here, we consider the Hilbert space of  $N$  spin-1/2 particles,  $D = 2^N$ , and the energies are normally distributed with mean zero,  $\langle E_\alpha \rangle = 0$ , and variance  $\langle E_\alpha^2 \rangle = \epsilon^2 = 1$ , introducing a unit  $\epsilon$  for energy. Under these premises, the random energy model provides a simple toy model for a spin glass transition. Above a critical temperature, an exponentially large number of states contributes almost equally to the free energy, while in a low-temperature phase the free energy is dominated by the ground state [24].

Choosing the  $|\alpha\rangle$  to be eigenstates of  $\sigma_z^{(i)}$  operators, we introduce quantum fluctuations through a transverse field term  $H_1 = B \sum_i \sigma_x^{(i)}$ . While in an adiabatic annealing algorithm, the field strength  $B$  decreases with time, starting from a large value which polarizes all spins along  $\sigma_x$ , our hybrid algorithm will evolve an initial state  $|\alpha\rangle$  under a constant Hamiltonian  $H = H_0 + H_1$ . The field strength  $B$  is supposed to be of the order of  $\epsilon$ . On the one hand, such field is sufficiently weak, such that it does not disturb very much the classical energy landscape for which we are interested to find its minimum. On the other hand, it is strong enough to introduce enough quantum fluctuations for a fast tunneling to other classical states. Accordingly, the proper choice enhances our chance of finding a new local minimum after projecting the time-evolved state onto the classical basis. A second parameter which is crucial for the performance of our algorithm is the duration of the time evolution. On short time scales, the probability of reaching a configuration of less energy increases with the evolution time  $t$ . After some time, however, the improvement rate oscillates around an average value, and, on average, one does not benefit from a longer evolution anymore. In our simulations, a good timing for making the projective measurement was found to be around  $t = 10$  (in units  $\tau \equiv \hbar/\epsilon$ ).

In a practical application, the classical energy landscape shall reproduce the cost function of a given optimization problem. In this case, evaluation of the energies can easily be done using a classical computer. If the quantum simulator, after the projective measurement, produces an output with decreased energy, the algorithm will accept this output as the new input state. Otherwise, it shall restore the previous configuration, and repeat the time evolution under the quantum field. In the latter case, even if the parameters  $B$  and  $t$  remain unchanged, the subsequent output may be different from the previous one, since the projective measurement introduces intrinsic randomness to the algorithm. Nevertheless, by varying one of these parameters, the performance of the algorithm is strongly enhanced. In our simulation, we therefore have chosen a random field  $B$  in the interval  $[0, 1]$ . As in simulated annealing, the algorithm also benefits from a non-zero acceptance rate for steps towards higher energies, to avoid getting stuck in metastable configurations. In our simulation, we have chosen an acceptance probability  $p(\Delta E) = \min\{\exp(-\beta\Delta E), 1\}$  with  $\Delta E$  the

$N$	average number of steps	median number of steps
8	100	60
9	160	100
10	240	150
11	390	250
12	650	390

TABLE I: For different Hilbert space sizes (parametrized by the number of qubits  $N$ ), we list the average and median number of annealing steps needed by the hybrid algorithm to reach the ground state. For the averaging, we have taken into account 1000 instances of the random energy model (667 in the case of  $N = 12$ ). In each run, the algorithm was initialized in the middle of the energy spectrum, i.e. in the  $2^{N-1}$ th level.

energy difference between new and old configuration. We worked with a constant low temperature,  $\beta = 10/\epsilon$ .

As described so far, our algorithm only involves quantum fluctuations, but it resembles a simulated annealing scheme in how information is processed. We therefore refer to it as “hybrid annealing” (HA), to distinguish it from the traditional “simulating annealing” (SA) where updates are produced by mimicking thermal noise. An algorithm where hybrid annealing and simulating annealing are run in parallel will be called “hybrid-simulating annealing” (HA+SA). We first consider the performance of the HA algorithm, and then compare it with other strategies.

### III. RESULTS

In our simulation of the HA algorithm, we need to mimic the time evolution of a quantum system on a classical computer. This necessarily restricts our simulation to small samples, and we have considered Hilbert spaces spanned by  $N = 8$  up to  $N = 12$  qubits. Despite the self-averaging nature of the random energy model, we therefore need to consider averages over many instances. The average over many instances will also compensate for the random nature of the algorithm itself, due to the fact that it consists of repeated projective measurements.

First, we have determined how many evolution steps it takes until the ground state is reached. The results are presented in Table I. Although the annealing time scales exponentially with  $N$ , it grows significantly slower than the size of the Hilbert space: Both average and median number of steps fit well to  $f(N) \propto b^{N-8}$ , with  $b = 1.60 \pm 0.02$ . In contrast, a random search would require an average number of steps that grows linearly with the Hilbert space, that is,  $\propto 2^N$ .

### A. Comparison with quantum-adiabatic scheme

Next, we compare the performance of the HA algorithm to other annealing approaches. To this aim, our simulation now fixes the runtime, and we compare the different algorithms by evaluating their success rate. In algorithms consisting of repeated updating steps (HA, SA, HA+AS), we fix the number of steps to 200. For the HA, as detailed above, this amounts to a minimum total runtime of  $t_1 = 2000\tau$ , where instantaneous state preparation and immediate classical data processing are assumed. We can then directly compare the HA runtime to the runtime of a quantum adiabatic annealing approach. In the latter, a single, uninterrupted quantum time evolution is performed for some given amount of time. Its final outcome is the evolved quantum state at time  $t_1$ , and the success probability is given by the squared overlap of this state with the ground state of  $H_0$ .

Apart from fixing the runtime, other details determine the protocol for quantum adiabatic annealing: The standard strategy is to prepare the system in an initial state which is fully polarized by a transverse field term  $H_1$ , and evolve under a time-dependent Hamiltonian  $H(t) = \alpha_0(t)H_0 + \alpha_1(t)H_1$ . In our simulation, we have considered a constant target Hamiltonian, i.e.  $\alpha_0(t) = \text{const.} = 1$ , and an exponentially decaying quantum field  $\alpha_1(t) = \exp(-t/t_0)$ , with  $t_0 = t_1/10$ . In this way, a single time-dependent parameter controls the evolution. We can still initialize the system in the ground state of  $H_1$  without exciting higher levels if the initial field strength  $B$  is chosen large enough. With our choice of  $B = 10\epsilon$ , the ground state of  $H(0)$  and the ground state of  $H_1$  differed from 1 by less than  $10^{-4}$  for all instances considered here.

The success probability  $p_{AA}$  of the annealing is given by the squared overlap of the final state with the desired target state. Averaged over 100 realizations with  $N = 10$  qubits, we find  $p_{AA} = 0.83$ . This is better than the corresponding success probability in the hybrid scheme,  $p_{HA} = 0.71$  (averaged over 1000 runs), obtained by counting the number of times the ground state is reached within the given runtime.

However, it is known that the adiabatic annealing becomes prohibitively expensive for large  $N$  due to exponentially small gaps [17]. In an attempt to mimic this problem in a small sample, we have artificially introduced a quasi-degenerate manifold above the true ground state. This makes it impossible for an adiabatic algorithm to determine the true ground state. In our simulation, we have chosen the four lowest excitations to have energies  $E = E_{GS} + 0.001\epsilon$ . This reduces the success rate of the adiabatic scheme to  $p_{AA} = 0.26$ . On the other hand, the performance of the hybrid annealing scheme is even improved,  $p_{HA} = 0.88$ .

This improvement can be understood from an enhanced tunneling between levels at equal energies, and provides a major benefit of the hybrid algorithm com-




FIG. 2: (a) An energy landscape for two spins: The global minimum  $|\downarrow\downarrow\rangle$  is detuned from the local minimum  $|\uparrow\uparrow\rangle$  by an energy offset  $\Delta$ , and separated by a barrier of height  $u$ . (b) For the depicted energy landscape (with  $u = 100$ ), we plot the tunneling probability  $T$  in a transverse field of strength  $B = 1$  as a function of the detuning  $\Delta$ . Asymptotically,  $T \sim 1/\Delta$ , as indicated by the fit (dashed line). The model explains the enhanced performance of our algorithm in quasi-degenerate energy landscapes.

pared to adiabatic annealing schemes. The behavior is captured by a simple two-spin model illustrated in Fig. 2(a): In this model, the local minimum  $|\uparrow\uparrow\rangle$  is separated from the global minimum  $|\downarrow\downarrow\rangle$  by a high energy cost  $u$  in configurations  $|\uparrow\downarrow\rangle$  and  $|\downarrow\uparrow\rangle$ . Within a transverse field  $B$ , there is a finite probability of tunneling from the local to the global minimum,  $p(t) = |\langle\downarrow\downarrow|e^{iHt}|\uparrow\uparrow\rangle|$ . In Fig. 2(b), we plot the time-averaged tunneling rate  $T = \lim_{\tau \rightarrow \infty} \frac{1}{\tau} \int_0^\tau p(t) dt$  as a function of the detuning  $\Delta$ . Asymptotically, the tunneling rate  $T$  is inversely proportional to the energy offset  $\Delta$ . This observation explains the enhanced performance of our algorithm in a quasi-degenerate scenario.

### B. Comparison with simulated thermal annealing

We have compared the hybrid scheme also with the outcome of simulated annealing (SA), where updates are produced by randomly flipping one spin. As HA, the SA scheme accepts an update following a probability distribution  $p(\Delta E, \beta) = \min\{\exp(-\beta\Delta E), 1\}$ . However, to bring the system from one potential well to another one, in the SA scheme one typically needs a series of accepted updates (of the order  $N$  flips). To make such a process likely, a high temperature is required in the beginning. In the course of the evolution, the temperature is then slowly decreased in order to relax into a (meta)stable configuration. Thus, in contrast to the HA algorithm, for SA it is crucial that the parameter  $\beta$  is varied during the algorithm. In the data shown in Fig. 3, we use  $\beta(k) = \beta_0 c^{-k}$ , with  $\beta_0 = 1/\epsilon$  and  $c = 0.98$ . Starting from  $k = 0$ , the variable  $k$  increases in steps of one after ten successful updates, or after twenty steps. We find that in only less than ten percent of the instances, the SA scheme is able to find the ground state of the random energy model within 200 steps. In most cases, it quickly gets stuck in local minima for long times. Clearly, one



reason for the bad performance of the simulated annealing is the lack of any structure in the random energy landscape.

### C. Improved hybrid annealing

An advantage of the HA scheme is the fact that it can easily be combined with purely classical algorithms. For instance, we are able to improve the HA algorithm by combining it with the SA algorithm. As depicted in Fig. 1(a), high barriers between two minima suppress thermal tunneling and make simulated annealing fail. On the other hand, large energy offsets between minima can reduce quantum tunneling. Following this picture, an algorithm which allows for both processes should have the best performance.

Accordingly, we have designed an algorithm which performs a hybrid-simulated annealing protocol (HA+SA). In each step, this algorithm chooses between two updates: One update is based on a random spin flip mimicking thermal fluctuations. A second update is produced by the unitary time evolution in a quantum field and the projective measurement in the hybrid algorithm. Among these two updates, the HA+SA algorithm chooses the one of lower energy, which is then accepted or not depending on the energy balance.

In Fig. 3 we analyze the performance of each algorithm. As a function of the number of steps, the figure plots the index of the lowest energy level reached so far. From this figure, it is seen that simulated annealing (SA) gets stuck in a relatively high level (on average 83rd level after 200 steps for  $N = 11$ ), while the hybrid annealing (HA) on average reaches the first-excited level. The combination of hybrid annealing and simulated annealing (HA+SA) yields only a slightly better performance on long time scales, but its advantages are apparent on short time scales. Here, HA+SA is the scheme for which the energy drops fastest. Since the combined algorithm produces two updates per step, we may compare it with a simulated annealing algorithm which chooses the better update among two random spin flips (SA2). It is seen that SA2 improves the simulated annealing performance on short time scales, but it is inefficient on the long run.

## IV. DISCUSSION

Summarizing the results presented above, our simulations show that the proposed hybrid annealing scheme outperforms the traditional simulated annealing approach on the random energy model, though without achieving polynomial scaling. The algorithm can be combined with the simulated annealing approach, giving rise to a setup where classical and quantum algorithm supplement each other (HA+SA). Accordingly, the HA+SA




FIG. 3: Performance of different algorithms applied to the random energy model with  $2^{11} = 2048$  levels: Averaged over 1000 instances, we plot the lowest level which each algorithm has found as a function of the runtime, that is, of the number of update steps the algorithm has gone through. Each algorithm starts from the 1024th state. At any runtime, the combination of hybrid annealing and simulated annealing (HA+SA) gives the best result. On the long run, the hybrid algorithm (HA) performs equally well, whereas the two purely thermal algorithms, SA and SA2, remain stuck in relatively high levels.

scheme will always outperform a strategy based on SA alone.

Comparing the hybrid approach with a quantum-adiabatic scheme, the latter shows better success rates in a small sample of the random energy model. However, small samples favor the adiabatic scheme due to pronounced finite-size gaps. Therefore, we have artificially increased the complexity by introducing quasi-degeneracies to the otherwise randomly generated energy landscape. While the success rate of the adiabatic algorithm dramatically drops, the performance of the hybrid algorithm even improves in the presence of quasi-degeneracies. This is a major benefit of the hybrid scheme compared to adiabatic annealing schemes. It suggests that the algorithm's performance could be enhanced for large samples.

Apart from the final output of the algorithm, other criteria are worth being considered: In contrast to the adiabatic scheme, the hybrid approach does not require long coherence times. In view of physical realizations, this favors its scalability. Moreover, since all the relevant data produced in the hybrid algorithm is classical, no quantum-error correction is needed. Since measurements are taken repeatedly, the evolution of the system is permanently tracked. This allows for some feedback control, for instance by automatically stopping the algorithm when the cost function drops below a desired value. Clearly, such strategy produces much more data than a single adiabatic run, providing information also about excited states. This information could be useful

if thermodynamic potentials are to be evaluated, and it also may help to evaluate the quality of the final result.

Moreover, such information could be useful if a large-scale problem is tackled piecewise or in a modular way: Assume a problem which is too big for a single quantum simulator. While there is no obvious way how to perform the adiabatic approach for a problem which exceeds the size permitted by the quantum simulator, the hybrid approach could simply divide the total Hamiltonian into several small clusters, and evolve each cluster separately in a quantum field. Based on the result for each cluster, the classical device can calculate the total cost function, and suggest a new configuration for the full problem.

These advantages compared to an adiabatic quantum annealer may motivate the additional efforts required for the development of hybrid annealer. In particular, for its efficiency it will be crucial to have fast ways of preparing quantum states in the computational basis and of

performing measurements projecting the states onto the computational basis.

### Acknowledgments

We thank Alireza Seif, Bruno Juliá-Díaz, and Stephen Jordan for fruitful discussions. We acknowledge financial support from AFOSR-MURI, from Fundació Cellex, from the European Union (ERC-2013-AdG Grant No. 339106 OSYRIS, FP7-ICT-2011-9 No. 600645 SIQS, H2020-FETPROACT-2014 No. 641122 QUIC, FP7/2007-2013 Grant No. 323714 Equam), from Spanish MINECO (FIS2013-46768-P FOQUS, SEV-2015-0522 Severo Ochoa), from the Generalitat de Catalunya (2014 SGR 874). TG acknowledges a JQI Postdoctoral Fellowship.

- 
- [1] 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, and G. Rose. Quantum annealing with manufactured spins. *Nature*, 473(7346):194–198, may 2011.
  - [2] Jerry M. Chow, Jay M. Gambetta, A. D. Córcoles, Seth T. Merkel, John A. Smolin, Chad Rigetti, S. Poletto, George A. Keefe, Mary B. Rothwell, J. R. Rozen, Mark B. Ketchen, and M. Steffen. Universal Quantum Gate Set Approaching Fault-Tolerant Thresholds with Superconducting Qubits. *Phys. Rev. Lett.*, 109:060501, Aug 2012.
  - [3] IBM. *IBM Quantum Experience*, page <http://www.research.ibm.com/quantum>, 2016.
  - [4] R. Barends *et al.* Digitized adiabatic quantum computing with a superconducting circuit. *Nature*, 534(7606):222–226, jun 2016.
  - [5] B. P. Lanyon, C. Hempel, D. Nigg, M. Müller, R. Gertsma, F. Zähringer, P. Schindler, J. T. Barreiro, M. Rambach, G. Kirchmair, M. Hennrich, P. Zoller, R. Blatt, and C. F. Roos. Universal Digital Quantum Simulation with Trapped Ions. *Science*, 334(6052):57–61, 2011.
  - [6] S. Debnath, N. M. Linke, C. Figgatt, K. A. Landsman, K. Wright, and C. Monroe. Demonstration of a small programmable quantum computer with atomic qubits. *Nature*, 536(7614):63–66, aug 2016.
  - [7] L. Isenhower, E. Urban, X. L. Zhang, A. T. Gill, T. Henage, T. A. Johnson, T. G. Walker, and M. Saffman. Demonstration of a Neutral Atom Controlled-NOT Quantum Gate. *Phys. Rev. Lett.*, 104:010503, Jan 2010.
  - [8] Stephan Ritter, Christian Nolleke, Carolin Hahn, Andreas Reiserer, Andreas Neuzner, Manuel Uphoff, Martin Mücke, Eden Figueroa, Joerg Bochmann, and Gerhard Rempe. An elementary quantum network of single atoms in optical cavities. *Nature*, 484(7393):195–200, apr 2012.
  - [9] J. S. Douglas, H. Habibian, C.-L. Hung, A. V. Gorshkov, J. H. Kimble, and D. E. Chang. Quantum many-body models with cold atoms coupled to photonic crystals. *Nat. Photon.*, 9(5):326–331, may 2015.
  - [10] W. B. Gao, A. Imamoglu, H. Bernien, and R. Hanson. Coherent manipulation, measurement and entanglement of individual solid-state spins using optical fields. *Nat. Photon.*, 9(6):363–373, jun 2015.
  - [11] N.Y. Yao, L. Jiang, A.V. Gorshkov, P.C. Maurer, G. Giedke, J.I. Cirac, and M.D. Lukin. Scalable architecture for a room temperature solid-state quantum information processor. *Nat. Commun.*, 3:800, apr 2012.
  - [12] Simon J Devitt, William J Munro, and Kae Nemoto. Quantum error correction for beginners. *Rep. Prog. Phys.*, 76(7):076001, 2013.
  - [13] Tadashi Kadowaki and Hidetoshi Nishimori. Quantum annealing in the transverse Ising model. *Phys. Rev. E*, 58:5355–5363, Nov 1998.
  - [14] Edward Farhi, Jeffrey Goldstone, Sam Gutmann, Joshua Lapan, Andrew Lundgren, and Daniel Preda. A Quantum Adiabatic Evolution Algorithm Applied to Random Instances of an NP-Complete Problem. *Science*, 292(5516):472–475, 2001.
  - [15] F Barahona. On the computational complexity of Ising spin glass models. *J. Phys. A*, 15(10):3241, 1982.
  - [16] Andrew Lucas. Ising formulations of many NP problems. *Front. Phys.*, 2(5), 2014.
  - [17] Thomas Jörg, Florent Krzakala, Jorge Kurchan, and A. C. Maggs. Simple Glass Models and Their Quantum Annealing. *Phys. Rev. Lett.*, 101:147204, Oct 2008.
  - [18] Boris Altshuler, Hari Krovi, and Jérémie Roland. Anderson localization makes adiabatic quantum optimization fail. *Proc. Nat. Acad. Sci. USA*, 107(28):12446–12450, 2010.
  - [19] S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi. Optimization by Simulated Annealing. *Science*, 220(4598):671–680, 1983.
  - [20] Philip Richerme, Zhe-Xuan Gong, Aaron Lee, Crystal Senko, Jacob Smith, Michael Foss-Feig, Spyridon

- Michalakis, Alexey V. Gorshkov, and Christopher Monroe. Non-local propagation of correlations in quantum systems with long-range interactions. *Nature*, 511(7508):198–201, jul 2014.
- [21] P. Jurcevic, B. P. Lanyon, P. Hauke, C. Hempel, P. Zoller, R. Blatt, and C. F. Roos. Quasiparticle engineering and entanglement propagation in a quantum many-body system. *Nature*, 511(7508):202–205, jul 2014.
- [22] Philipp Hauke, Lars Bonnes, Markus Heyl, and Wolfgang Lechner. Probing entanglement in adiabatic quantum optimization with trapped ions. *Front. Phys.*, 3:21, 2015.
- [23] Tobias Graß, David Raventós, Bruno Juliá-Díaz, Christian Gogolin, and Maciej Lewenstein. Quantum annealing for the number-partitioning problem using a tunable spin glass of ions. *Nat. Commun.*, 7:11524, may 2016.
- [24] B. Derrida. Random-Energy Model: Limit of a Family of Disordered Models. *Phys. Rev. Lett.*, 45:79–82, Jul 1980.