Eur. Phys. J. Special Topics **224**, 1–4 (2015)

DOI: 10.1140/epjst/e2015-02336-2

THE EUROPEAN
PHYSICAL JOURNAL
SPECIAL TOPICS

Editorial

Quantum annealing: The fastest route to quantum computation?

Sei Suzuki[1,a] and Arnab Das[2,b]

[1] Department of Liberal Arts, Saitama Medical University 38 Moro-hongo, Moroyama, Saitama 350-0495, Japan
[2] Theoretical Physics Department, Indian Association for the Cultivation of Science 2A & 2B Raja S. C. Mullick Road, Kolkata 700032, India

Received 9 December 2014 / Received in final form 16 December 2014

What is Quantum Computation. The world at the microscopic scale, i.e. at the scale of an atom or an electron, looks very different from the world of our daily life experiences. The very idea of the "state" of a microscopic (quantum) object is radically different from that of a macroscopic (classical) object. According to our daily experiences, at any given instant a system (say a switch or a particle) remains in a *definite* state characterized by some *definite set of values* of its physical properties. For example, a switch will either be in the state "off" state or in the state "on" but never in a state of simultaneously being on *and* off. Similarly, a particle moving on a surface (say) must be located at some definite point $\mathbf{R}$ on the surface at a given instant, and can't be simultaneously present at different points. But in the quantum world, a switch is allowed to be in a state which is a simultaneous "superposition" of both on and off states, and a particle can remain in a state that is "delocalized" over the entire surface at the same instant. Thus a quantum object can be in a state which has simultaneous access to all different configurations which are mutually exclusive classically. This property, often referred to as "quantum parallelism" endows quantum systems with enormous potential to store and process information with an efficiency that no classical machine can ever match.

To illustrate quantum parallelism a bit further, let us consider a device which takes input through n input switches. In order to obtain complete information about all the outputs of such a system one would need either to realize 2^n input configurations in sequence (implying an evaluation time that grows exponentially with n) or 2^n systems running in parallel, each realizing one input configuration (meaning exponential resources). But if one has an analogous system made of quantum switches, then one could put the input state to be a (say, with equal weight) superposition of all 2^n configurations, and the *system* would be able to access information about all those mutually exclusive configurations simultaneously. This apparently seems to be an exponential advantage to begin with. However, since in such a superposed input state the contribution of each configuration will be very (typicaly exponentially) small, retrieving the necessary information from the resulting output is far from trivial.

---

[a] e-mail: `sei01@saitama-med.ac.jp`
[b] e-mail: `arnab.das.physics@gmail.com`

Nevertheless, quantum parallelism can be employed cleverly to design quantum algorithms that would outperform any classical computer in solving certain hard computational problems, and even exponential quantum speed-up can be achieved. For example, Peter Shor has devised a quantum algorithm for factorizing an integer N within a time which is polynomial in the number of digits in N, while no classical algorithm is known to date that can solve the problem in a time which is faster than sub-exponential in the number of digits in N. Another interesting case is that of Grover's algorithm. This algorithm searches an unstructured database consisting of N entries in $\mathrm{O}(\sqrt{N})$ time, while no classical algorithm can manage to do it any faster than $\mathrm{O}(N)$ time, since knowledge of one failed attempt provides absolutely no information regarding the possibility of finding the target entry among the remaining ones. This is analogous to finding a hole (potential dip) on a flat landscape in the configuration space. Since there is no gradient, a classical algorithm will always take time proportional to the area of the landscape (*i.e.*, the size of the search space) to locate the hole. But suitably designed version of Grover's algorithm can do the job within a time that scales as $\sqrt{area}$ of the landscape.

***Quantum Annealing – The Analog version of Quantum Computation*:** Drawing the idea of Grover's search algorithm on a flat landscape a bit further, one can introduce the idea of quantum annealing, which is a very natural and so far the most widely accepted way of realizing analog version of quantum computation. The framework is particularly suitable for solving optimization problems where a cost function is to be minimized in a complex cost-configuration landscape. To illustrate the idea let us consider the simplest scenario of flat landscape with a single potential dip (Grover's hole search problem). In order to ensure that the initial state does not miss out the target state (hole), one starts with an equal superposition of all classical configurations (all position eigenstates in this case) with equal weight. That will be a momentum eigenstate with very high kinetic energy. Now if one reduces the kinetic energy adiabatically slowly, so that the state of the particle is always close to the ground state of the instantaneous Hamiltonian, then in the limit at which the kinetic energy is reduced to zero, the particle is expected to be at the ground state of the potential energy, i.e., within the hole. The condition of convergence to the actual ground state at the end of annealing is given by quantum adiabatic theorem of quantum mechanics, and hence such quantum annealing schemes are also known as quantum adiabatic algorithms. For most problems of interest, the landscapes are of course not flat but are highly rugged, plagued with local cost minima and barriers. Quantum fluctuations can have the edge over classical ones even in such cases under special circumstances as discussed in the following.

Though quantum parallelism (i.e. the ability of a single particle to delocalize over the entire space) is a crucial element of this quantum search mechanism, the idea of employing kinetic energy provided by thermal fluctuations to search global energy minimum of complex energy function was much older, and is known as *simulated annealing* (the dynamics is simulated using a suitable algorithm like Metropolis). There, instead of quantum kinetic energy term, the temperature is reduced from a very high value to zero slowly in order to reach the minimum energy state of a given classical Hamiltonian/energy function. The origin of neither the energy function nor the temperature needs to be physical of course. The technique can be used to optimize a multivariate cost function of several variables employing an artificial temperature that induces a thermal dynamics that allows the system to explore the configuration space, often jumping out of shallow local minima.

Why then quantum version is expected to do better, at least under certain circumstances? Though a major impetus behind such optimism (in retrospect) is of course quantum parallelism, but historically the motivation was quantum tunneling, which was thought to be a potentially superior mechanism to maintain ergodicity in rugged energy landscapes than thermal jumps. Crossing the barriers surrounding a local minimum using thermal fluctuations would require thermal jumps over the barriers, probability of which is suppressed exponentially with the height of the barrier, regardless of the barrier width. On contrary, quantum mechanical tunneling allows crossing any energy barrier, however high, if it is narrow enough. Hence for optimization problems where the cost barriers are generically high but narrow, annealing using quantum fluctuations can be expected to excel over purely thermal annealing. Historically speaking, at the formative stage the annealing process was not envisaged to be a purely quantum (Schrodinger) dynamics, and hence neither role of quantum parallelism nor quantum adiabatic convergence was explicitly identified. However, it is interesting to note that the early version of finite (but low) temperature quantum annealing is of no less potential than its later adiabatic version, particularly where loss of ergodicity due to localization can be restored by thermal hopping.

***The First Quantum Computer in the Making?*:** In recent days, a Canada based company named D-Wave has claimed to have realized the first quantum computer of a reasonable size which can do calculations of practical importance. The working of the device (DW1) is based on the principle of quantum annealing. This, quite expectedly, created a huge wave of enthusiasm and speculation reflected in several NEWS blogs and popular scientific media (see A. Ghosh, and S. Mukherjee, arXiv:1310.1339 for an interesting compilation). Several reputed Institutions bought the machine: The first two machines were bought by Lockheed and installed at USC (the DW1 in 2011 and the DW2 in 2013). Google followed suit and installed at NASA (the DW2 in 2013, 5 months after Lockheed-USC's DW2). Scientists are trying to build models from the output of the quantum annealer to understand its mechanism. Several questions regarding the "quantumness" and effectiveness of the machine are still looming large. Here we provide an authentic account of this development depicted by almost all major scientific groups directly involved with this great advancement at its different stages – starting from the early theoretical developments till the state-of-art activities. The issue has the following structure.

In the introductory article by *Das and Suzuki* a brief outline of the ideas of quantum annealing is illustrated. Broad failure modes of quantum annealing are discussed and certain generic areas where quantum annealing would outperform classical computers are identified. In the following review by *Nishimori* the present status of quantum annealing is summarized and a comparison is drawn between quantum annealing and simulated annealing. *Mukherjee and Chakrabarti* discuss potential advantages of quantum annealing in the context of multivariate optimization problems. *Silevitch, Rosenbaum and Aeppli* review experimental evidences of superiority of annealing dynamics using quantum fluctuations over those using purely thermal fluctuations. *Boixo, Ortiz and Somma* discuss the scope of quantum optimization algorithms broadened beyond the domain of adiabatic quantum annealing. They discuss different strategies (including adiabatic quantum annealing) for preparing low temperature states of complex systems. *Suzuki* presents a general review on performance of quantum annealing in multivariate optimization problems and reviews certain rigorous proofs of faster convergence of quantum annealing compared to classical annealing due to Morita and Nishimori. *Hen* and *Young* review the performance of quantum annealing in constraint satisfaction problems and spin glass problems. They suggest modifications of adiabatic quantum annealing to speed up quantum algorithms in solving those hard problems where quantum annealing fails generically. *Laumann et al.* analyze two different failure modes of quantum annealing, namely, due to occurrence of small gap at the phase transition point and due to occurrence of small gap within the phase, in the context of decision problems. In the first part of their article, *Cohen and Tamir* review varieties of computational methods for optimization. Then they analyze the performance of the D-Wave machine which is claimed to be a quantum annealer by its maker. *Albash et al.* presents a thorough analysis of the D-Wave machine (DW1) data comparing it with those produced by theoretical models of Simulated Quantum Annealers as well as classical rotor models. They conclude that neither of these models can satisfactorily capture all the features of the D-Wave machine. *Smelyanskiy et al.* report their approach to the problem of diagnosing and detecting multiple faults in graph based systems using the D-Wave machine of 509 qubits. They achieved a new level of success in using quantum annealer to solve advance diagnostic problems surpassing previously reported ones limited to smaller system-sizes. *Inoue* reviews information processing, computing and inference using quantum fluctuations in the setting of infinite range Ising models in the transverse field. *O'Gorman et al.* introduce a method of learning the structure of a Bayesian network using adiabatic quantum annealing. In their article *del-Campo and Sengupta* review different protocols (not necessarily adiabatic) for suppressing defect generation while annealing through a quantum critical point. As a concluding chapter, we have compiled a debate and discussion section based on the response we received to a questionnaire containing few key questions that we circulated among the contributors with our commentaries.

Given the implication of realization of a usable quantum computer, even a small achievement in the right direction would be marked as a colossal milestone in the history of computation. We hope this small issue will serve as an early chronicle of the beginning of this great revolution.

Both the editors thank all the contributors for their valuable contributions and cooperation. Special thanks to Christian Caron, Isabelle Houlbert and Sabine Lehr (Springer) for their enthusiasm, encouragement and prompt responses. We also thank Bikas K. Chakrabarti for several useful suggestions. Finally we thank Asmi Haldar for creating the cover page picture for the issue.

Sei Suzuki
(Saitama, Japan)

Arnab Das
(Kolkata, India)
December 16, 2014