Eur. Phys. J. Special Topics **224**, 205–212 (2015)

DOI: 10.1140/epjst/e2015-02351-3

THE EUROPEAN PHYSICAL JOURNAL SPECIAL TOPICS

Discussion & Debate

Debate and discussion: Quo Vadis quantum annealing?

A. Das[1,a] and S. Suzuki[2,b]

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

Received 14 December 2014 / Received in final form 16 December 2014
Published online 5 February 2015

Here we compile the debate and discussion section based on the response we received to a questionnaire containing few key questions that we circulated among the contributors. Following each question we summarize the responses with a commentary, followed by the original comments of the participating contributors. A reference in bold face (e.g. Albash or Cohen and Tamir) appearing in a commentary following a given question refers to the answer to the said question provided by the referred participant/s (appended immediately below the commentary).

Q1. Do you think quantum annealing (QA) should work better than classical annealing for optimization in most cases? Which are the scenarios you think where QA might outperform classical annealing?

It seems that the question of superiority of quantum annealing over classical annealing in solving a problem of practical interest has no general answer. There are certain quantum mechanical aspects, e.g. efficient tunneling through narrow but high barrier, quantum parallelism etc. which might help quantum annealing to excel over classical annealing under certain circumstances, but such circumstances are highly system-specific. This is reflected in the comments by Albash, Das, Cohen and Tamir, Mukherjee and Chakrabarti, and Sengupta below.

In spite of this, there are quite a few areas of general interest where adiabatic quantum annealing would perhaps be the one most convenient means to achieve the goal. For example, this will be true for finding the ground state(s) of quantum frustrated systems where quantum Monte Carlo method suffers from the sign problem (see Das below and the article by Das and Suzuki in this issue).

Moreover, within an extended definition of quantum annealing, which includes any process where quantum fluctuations are used in the annealing, and the dynamics is neither restricted to be adiabatic, nor presence of thermal fluctuations are ruled out (as pointed out by Cohen and Tamir; and Somma and Ortiz), scope of quantum annealing might increase.

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

It seems possible that controlled thermal fluctuations might in fact help finding the ground state, particularly when adiabatic quantum annealing is plagued with localization (see article by Das and Suzuki, Laumann et al. in this issue). In those cases, ergodicity might be resorted by thermal fluctuations (in somewhat the same way finite conductivity is achieved by phonon-assisted hopping). In that case, the algorithm might benefit both from effectiveness of quantum tunneling through high but narrow barriers, and also acquire the desired mobility where localization would paralyze the system at $T = 0$.

Finally, interesting general theorems have been proved by H. Nishimori and his collaborators demonstrating superior convergence rate of adiabatic quantum annealing over the thermal one for very generic Ising models (see A. Das and B.K. Chakrabarti, Rev. Mod. Phys. **80**, 1061 and the article by S. Suzuki in this issue for a review). These theorems however do not necessarily imply superiority of quantum annealing over thermal annealing in solving the worst case instances of a problem or in reducing the complexity class of a given problem.

Responses:

Albash: As several articles in this review have pointed out, the general intuition about tunneling through large but narrow energy barriers strongly suggests that QA should have an advantage over classical annealing. However, given that we cannot predict the type of energy landscape that QA experiences for a general problem, it seems we have been forced into a trial-and-error sort of approach. An approach like in (B. Yucesoy, J. Machta, and H.G. Katzgraber, Phys. Rev. E **87**, 012104 (2013)) where some features of the classical landscape might be observed might lead to the development of tailor made problems that we can hope can reveal a quantum advantage for QA over classical annealing.

Das: It is difficult to answer in general based on our present state of understanding of many-body quantum mechanics. It seems (see Das and Suzuki, and Laumann et. al. in this issue) presence of strong disorder will force quantum annealing to fail generically, *i.e.*, makes it exponentially/sub-exponentially slow in system-size due to localization (Anderson localization or Many-body localization, as the case may be). This can happen even in cases where there is no frustration and finding the ground state is trivial.

On the other hand in certain cases, where the cost/energy function (not necessarily of physical origin) is somehow dominated by high but narrow enough barriers, quantum annealing can perform much better than classical annealing. Quantum annealing can also generically lead to non-classical square-root speed-up when the landscape is essentially flat (see the articles by A. Das and S. Suzuki in this issue).

Simulation of quantum systems with frustration without disorder, which are classically intractable (say, due to sign problems) is also a field where quantum annealer can grossly outperform a classical simulator.

Cohen and Tamir: Quantum annealing was suggested as an improvement of the simulated annealing technique which suffers a severe setback in cases where the system is "non-ergodic" (e.g. systems described by the spin glass model). Quantum annealing excels in tunneling through narrow (possibly cuspidal) barriers. Classical simulated annealing schedules might still have an advantage where the barrier is wide and shallow. In fact, for some specific problems the advantage of quantum annealing over simulated annealing is much clearer. In (R. Martonak, G.E. Santoro, E. Tosatti, [arXiv:cond-mat/0402330] (2004)) for instance, quantum annealing based on path-integral MC, showed better results for the Traveling Salesman Problem for

1002 cities. QA was shown to anneal more efficiently, and to decrease the solution's residual error at a much steeper rate than SA. The authors in (E. Farhi, J. Goldstone, S. Gutmann, [arXiv:quant-ph/0201031v1] (2002)) construct an example where the width between local minima is small and therefore the tunneling effect is strong. The simulated annealing counterpart of the example shows an exponential complexity. Quantum annealers are usually adiabatic computers, however the QA should be thought of being more general. One should differentiate between the benefits of the adiabatic computation and the benefits that are rooted in the tunneling effects (R.D. Somma, C.D. Batista, G. Ortiz, Phys. Rev. Lett. **99**, 030603 (2007)).

Mukherjee and Chakrabarti: It is not clear if QA works better than SA in most cases. As discussed above using the tunneling picture, if the "free energy" landscape consists only very "high" but extremely "thin" barriers, then QA should work (and SA cannot work). However, neither the landscape is known, nor the QA is seen to be superior in the all the cases studied (see articles by E. Cohen & B. Tamir and by S. Suzuki in this issue).

Sengupta: This is expected to occur for a certain class of problem for which one can use the annealing technique efficiently. For a generic problem, this question is currently being tested. One can define a relative speedup coefficient $\boldsymbol{S}$ = t_{classical}/t_{quantum} where t_{classical} and t_{quantum} are the times taken to solve the problem using classical CPU and quantum annealer (D-wave machines) and try to design algorithms which maximizes $\boldsymbol{S}$.

Somma and Ortiz: As exposed in our present manuscript, our definition of quantum annealing (QA) is one where the goal is to prepare the ground state of a generic quantum system by evolving a simple initial state according to the Schrodinger equation. Our results show that classical annealing (SA) can be interpreted as only one instance (or path) of QA, suggesting QA as a more powerful heuristics for optimization since additional arbitrary paths are available. Moreover, using our spectral gap amplification technique, we were able to prove a quantum speedup for QA with respect to SA in a quantum device. This speedup is more significant in scenarios where the optimization problems are computationally hard.

Q2. Do you think QA format of quantum computation is more robust and convenient to realize than the gate-based quantum computation?

The question has two facets, namely, robustness (which means its tolerance against faults) and simplicity in implementation.

On the issue of robustness most of the participants, namely, Albash, Das, Somma and Ortiz seems to believe that annealing based quantum computation still lacks the robustness of the circuit-based quantum computation, since so far there exists no systematic error-correction scheme for the quantum annealing format, as it does for circuit based quantum computation. However, Mukherjee and Chakrabarti strongly believe the contrary. Albash also notes that the quantum annealing version is more stable to some special kind of decoherence.

So far as simplicity is concerned, quantum annealing format has both its strengths and weaknesses compared to the circuit format. For optimization problems like finding ground state of a cost/energy function, quantum annealing seems to be the most natural and easily implementable (see Das, Somma & Ortiz and the article by A. Das and S. Suzuki in this issue). However, there are other kinds of problems, e.g. Shor's algorithm for factorizing integers, quantum annealing version seems to be quite

difficult (actually has not been realized to date). Hence the answer to this question seems to be largely problem specific. Sengupta also concurs.

Responses:

Albash: We know that QA can be more robust to some forms of decoherence (A.M. Childs, E. Farhi, J. Preskill, Phys. Rev. A **65**, 012322 (2001); M.S. Sarandy, D.A. Lidar, Phys. Rev. Lett. **95**, 250503 (2005); M.H.S. Amin, D.V. Averin, J.A. Nesteroff, Phys. Rev. A **79**, 022107 (2009)), which might make it more easily implementable in the near term, but there still remains a major theoretical problem for QA: the absence of an accuracy threshold theorem (E. Knill, R. Laflamme, W.H. Zurek, Science **279**, 342 (1998); J. Preskill, Quant. Inf. Comput. **13**, 181 (2013); D. Gottesman, [arXiv:1310.2984] (2013)). Without such a theorem, it is not clear that any large scale QA implementation will be possible.

Das: Circuit based quantum computation owes its robustness to the quantum error correction schemes, which can efficiently eliminate a large class of errors. Lack of its equal in the quantum annealing format renders it less robust than the circuit format. However, suitable dynamical decoupling schemes can perhaps fill this gap for this analog version of quantum computation in near future. As far convenience, both the formats have (dis)advantages depending on the nature of the problem. Quantum annealing format seems more directly implementable for solving combinatorial optimization problem; in particular those can be cast in form of Boolean variables. But there are other kind of problem, for example Shor's factorization problem, though circuit version is straightforward, quantum annealing version is yet to be formulated.

Cohen and Tamir: Regarding the coherence time of the qubits, there is currently an inherent problem with flux qubits, although it was shown that in some sense, the D-Wave computer is more robust against thermal noise (N.G. Dickson, et al., Nature **4**, 1903 (2013)). We feel this result should be better understood. The D-Wave computers were also capable of incorporating the largest number of qubits known today. Regarding algorithmics, it is nowadays unclear how to translate in the general case a gate-based quantum computation to a QA computation. This makes (only in practice) quantum annealers less universal than gate-based quantum computers, but we expect that the pool of QA will increase over time.

Mukherjee and Chakrabarti: Yes, we do.

Sengupta: I am not sure if this statement can be made independent of the class of problem that needs to be studied. For example, it is not clear at the moment if the D-Wave machine can achieve the expected speedups of standard problems such Schor's alogorithm.

Somma and Ortiz: We do not believe that adiabatic (QA format) quantum computation is more robust or convenient than the standard network or gate model of quantum computation. We do believe that adiabatic quantum simulators, without the requirement of computational universality, may constitute convenient routes to perform specific quantum computations at small scales, but quantum error correction will eventually be necessary. For physicists the QA paradigm of computation is appealing because it provides an intuitive way of developing new algorithmic tools

or specific algorithms that can eventually be written in any model of computation, including the network model.

Q3. Do you think the D-wave (DW1) machine is an effective quantum annealer?

It seems D-Wave machine demonstrated some quantum characteristics, but the scenario is still far from clear. This is because a convincing physical model that explains all the outcomes of the machine consistently is yet to be found. In some cases it seemingly outperforms its classical competitors, while in other cases it is much slower than those. However, in judging the "quantumness" of an annealer it should be kept in mind that inferiority of performance of the annealer in question compared to a standard classical annealer does not necessarily imply that the annealer under test is essentially a classical (and inferior) one. Depending on the problem, even a perfect quantum annealer might perform much worse than a classical annealer in principle.

Responses:

Albash: Given that we have no good model that effectively captures most features of the DW1 device at large system sizes, the cautious answer is, we don't know yet what the DW1 device is. Is an effective (finite temperature) quantum annealer one that can operate in a regime where the dominant dephasing is dephasing in the instantaneous energy eigenbasis, an innocuous form of dephasing for quantum annealing? If so, we can't really answer that question since we do not have a model with such a dephasing mechanism that captures all features of the device at large problem sizes. At least for small system sizes, the DW1 device appears to be well modeled by an adiabatic Markovian master equation, so for small system sizes, we can say that it might be an effective quantum annealer. For larger system sizes, the master equation becomes numerically prohibitive to simulate, so we are no longer able to test whether the DW1 continues to agree with the master equation predictions.

If an effective quantum annealer is one that gives us any quantum advantage over a thermal annealer, then again, we have no evidence that the DW1 has such an advantage. One thing to keep in mind in this regard is that the D-Wave devices are possibly severely handicapped by control errors on the programmable couplings and local fields. Especially as problem sizes grow, these errors can dramatically hurt the performance of the annealer.

Cohen and Tamir: The DW1 computer is manifesting an interesting regime which is in-between adiabatic and thermic. It was already shown (S. Boixo, et al., [arXiv:1304.4595] (2013)) that it excels in some problems, but fails on others. We believe there is a need for more conclusive results; however it already seems that DW1 has made a major leap in the field of quantum annealing.

Mukherjee and Chakrabarti: Yes, indeed (see articles by Albash et al. and Perdomo-Ortiz et al. in this EPJ ST issue).

Sengupta: It is not clear at the moment if DW1 or DW 2 has better speedup than other traditional annealers for arbitrary problems. Thus this question cannot be answered independent of the problem which one aims to solve.

Somma and Ortiz: There are mixed results in the literature regarding the computational model that best describes DW1. While some quantum effects seem to be part of the computation, thermal effects also play an important role and, under our definition of QA, the DW1 machine does not seem to behave as an effective quantum annealer.

Q4. What do you think are the major hurdles in the way to realizing a perfect quantum annealer?

From the discussion below it seems decoherence and scalability are the two main hindrances in the way of realizing a *practical* quantum annealer (perfect annealer is neither possible, nor seems necessary). These two problems are actually intertwined on physical ground. Moreover, there seems to be a curious and unavoidable trade-off between scalability and achievable coherence among the systems realizable in laboratory. Whether D-Wave has already made a serious progress in striking a subtle balance, is still an open question to the community (see Das below).

Responses:

Albash: Besides the technological hurdles, a critical theoretical aspect that remains missing is how to do scalable error correction for quantum annealing.

Das: Perfect quantum annealer is an idealization, and most possibly, a superfluous one. What is perhaps more reasonable to look for is a practical quantum annealer. Achieving scalability in presence of decoherence – a fundamental requirement for a practical quantum annealer – seems (unfortunately) to be the most difficult goal. The difficulty is multi-pronged: with increasing system-size, the gaps reduce, which means the system becomes more prone to decoherence due to ambient thermal fluctuations. Moreover, smaller gaps mean longer annealing time, which again implies more decoherence. In systems like NMR, where decoherence is well controlled, system size cannot be increased beyond a handful of qubits. Taming decoherence and/or formulating a new strategy to attain scalability seems to be the major challenge.

Cohen and Tamir: (1) *Scalability*: This question is deeply connected to the lack of fault tolerant gate theory for adiabatic computation. It is possible that the complexity to construct such a computer with all its inter-couplings, grows itself exponentially. This will mean that the possible gain in algorithmic complexity is paid out in building a coherent circuit (see also G. Kalai, [arXiv:1106.0485] (2011)). (2) *Coherence*: Increasing the coherence time of the qubits is needed for meeting the required computational time.

Sengupta: I think achieving better decoherence times, ensuring that one reaches the final state with a good degree of certainty, and applicability to a major class of problems are the chief hurdles.

Somma and Ortiz: Strictly speaking, the laws of physics prevent the design of a perfect quantum annealer: All operations in a computation are performed within a certain precision and, especially in quantum computational models, operation errors accumulate eventually destroying the desired-state preparation. One way to overcome these complications is via quantum error correction. However, a fault-tolerant model of quantum annealing is yet to be developed and it is not clear whether such a model can ever exist. Also, notice that in order to achieve a given computational task there could be issues related to the locality of the physical probes used to address the physical quantum annealer.

Q5. What do you think is the most promising setup (ultra-cold atoms in optical lattice/quantum optics/interacting superconducting qubits... etc.) for realizing a perfect quantum annealer?

Given the fact that D-Wave has already launched a quantum annealer (consisting of about 500 qubits) which seems to have shown certain quantum characteristics,

majority of the participants gave their verdict in favor of the D-wave building blocks, namely, the superconducting qubits. However, it is also pointed out by Das, Somma and Ortiz, and Sengupta that an ultra-cold atom in optical lattice is also a very promising candidate for quantum annealing format of quantum computation where it is applied to problems of condensed mater physics.

Responses:

Albash: Since all setups have their particular advantages and disadvantages, a parallel effort on all fronts will probably establish a different "perfect" quantum annealer for each task. Some may be better suited for solving optimization problems of a particular type, others might be better suited for studying quantum simulations.

Das: Interacting superconducting qubits have recently been used successfully in D-wave one quantum annealer (see articles by Albash, et al., Pedromo-Ortiz, et al., and Cohen and Tamir in this issue), and hence have made a place as a bona-fide candidate for hardware of quantum annealers. Other potential candidates for quantum annealing (or quantum simulation in general) of physical systems are cold atoms in optical lattice, where both realizing interaction between the atoms and tuning them with time can be achieved with great accuracy for almost macroscopic time-scales (see, e.g., I. Bloch, J. Dalibard, S. Nascimbene, Nat. Phys. **8**, 267 (2012)). These systems are likely to evolve as experimentalist's quantum annealers, settling age-old questions and controversies in condensed matter physics, chemistry and material science.

Cohen-Tamir: The flux qubits used by the D-Wave group have the advantage that a large set of them could be easily constructed and manipulated using well established techniques. It is currently the only known way which allows concatenating such a large set of qubits. Moreover, a promising progress was achieved recently regarding the decoherence time of flux qubit (M. Stern, et al., Phys. Rev. Lett. **113**, 123601 (2014)). It is too early to tell, but anyons (A.Y. Kitaev, Ann. Phys. **303**, 2 (2003)), being topologically protected against decoherence, seem (at least theoretically) to be a promising alternative.

Mukherjee-Chakrabarti: Interacting superconduting junction qubits (M.W. Johnson, et al., Nature **473**, 194 (2011)).

Sengupta: I would think that given scalability and relatively easy experimental implementation, superconducting qubits would be the most promising setup.

Somma-Ortiz: Perfect quantum annealers do not exist, as explained above. Nevertheless, realistic quantum annealers may be more efficient quantum simulators than quantum computers, as they are designed for special purposes. Therefore, the most promising setup will depend on the specific task/problem one is trying to optimize. For example, if one wants to determine the ground state energy of the Bose-Hubbard model, a typical optimization problem, the optical lattice setup may be the most convenient one.

List of Participants:

(1) **Albash:** Tameem Albash (University of Southern California, Los Angeles, USA)
(2) **Chakrabarti:** Bikas K. Chakrabarti (Saha Institute of Nuclear Physics, Kolkata, India)
(3) **Cohen:** Eliahu Cohen (Tel Aviv University, Tel Aviv, Israel)

(4) **Das:** Arnab Das (Indian Association for the Cultivation of Science, Kolkata, India)
(5) **Mukherjee:** Sudip Mukherjee (Saha Institute of Nuclear Physics, Kolkata, India)
(6) **Ortiz:** Gerardo Ortiz (Indiana University, Bloomington, USA)
(7) **Sengupta:** K. Sengupta (Indian Association for the Cultivation of Science, Kolkata, India)
(8) **Somma:** Rolando Somma (Los Alamos Natl. Lab., Los Alamos, USA)
(9) **Tamir:** Boaz Tamir (Bar-Ilan University, Ramat-Gan, Israel)