

Stronger benchmarks needed to fully quantify quantum speedup, physicist says

February 5 2014, by John German



Credit: Texas Advanced Computing Center.

(Phys.org) —Texas A&M University physicist Helmut G. Katzgraber's research takes him to the crossroads of physics, computer science, quantum information theory and statistical mechanics. For more than a decade, he has been pushing the frontiers of computational physics to study hard optimization problems and disordered materials, applying his expertise to problems in the fascinating and fast-evolving field of quantum computing.

This past week, his work caught the attention of the global research community because of a study related to a particular commercial quantum computing device, the USD 10M D-Wave Two—more specifically, its [documented failure](#) to outperform traditional computers in head-to-head speed tests by Ronnow et al.

Not so fast, says Katzgraber, whose own National Science Foundation-funded research points to an intriguing possible explanation: Benchmarks used by D-Wave and research teams alike to detect the elusive quantum speedup might not be the best to do so and, therefore, not up to the test.

In a paper submitted earlier this month, Katzgraber details his team's innovative results on quantum speedup. Among other findings, he proposes potentially hard benchmark problems to detect and quantify such a mysterious target as quantum speedup, which Katzgraber says is highly dependent on the combination of the chosen benchmark and optimization algorithm. In particular, his results suggest that the current benchmarks might not be best suited to truly showcase the potential of the [quantum annealing](#) algorithm, a quantum version of thermal simulated annealing and the technology upon which the D-Wave machine is based.

Simulated annealing borrows its name from a type of heat treatment that involves altering a material's properties by heating it to above its critical temperature, maintaining the temperature and then cooling it slowly with the hope of improving its ductility, Katzgraber explains. In using simulated annealing as an optimization method, the system is heated to a high temperature and then gradually cooled in the hope of finding the optimal solution to a particular problem. Similarly, in quantum annealing, quantum fluctuations are applied to a problem and then slowly quenched again in the hope of finding the optimum of the problem.

Katzgraber's work, primarily done by simulating spin-glass-like systems (disordered magnetic systems) on D-Wave's chip topology using the facilities at the Texas A&M Supercomputing Facility and the Stampede Cluster at the Texas Advanced Computing Center (TACC), shows that the energy landscape of these particular benchmark instances might often be simple, with one dominant basin of attraction. Optimization algorithms such as simulated annealing excel in studying these type of problems. Not surprisingly, he advocates for additional testing and better benchmark design prior to proclaiming either defeat or victory for the D-Wave Two machine.

"Simulated annealing works well when the system has one large basin in the energy landscape," Katzgraber said. "Think of a beach ball on a golf course with only one sand pit. You let it go, and it will just roll downhill to the lowest part of the pit without really getting stuck on the way. But if you have something with one dominant pit embedded in a landscape with many other hills and valleys, then the ball might get stuck on its way to the deepest pit and therefore miss the true minimum of the problem.

"My results seem to indicate that the current benchmarks might not have the complex landscape needed for quantum annealing to clearly excel over simulated annealing; i.e., a landscape with deep valleys and large barriers where the quantum effects can help the system tunnel through these barriers to find the optimum (i.e., the deepest pit) efficiently. This, of course, does not mean that quantum annealing does not perform well in the current benchmarks, but the signal over simulated annealing could be stronger by using better benchmarks. I am merely proposing benchmark problems that have an energy landscape more reminiscent of the Texas Hill Country versus the comparatively flat terrain in the College Station area—benchmarks where we know that simulated annealing will fail quickly."

The D-Wave machine currently in use by Google and NASA was

benchmarked by a team of scientists from the University of Southern California, ETH Zurich, Google, the University of California at Santa Barbara and Microsoft Research in [work that was independent](#) of Katzgraber's but submitted near-simultaneously. The two teams do agree on one important point: The jury's still out because better benchmarks need to be developed.

"While on the one hand, D-Wave wants to dismiss the tests, and on the other, scientists have shown the machine is only faster for certain instances, I am proposing potentially harder tests before any concrete conclusions are drawn," Katzgraber said.

More information: Read the paper online:
arxiv.org/pdf/1401.1546.pdf

Provided by Texas A&M University

Citation: Stronger benchmarks needed to fully quantify quantum speedup, physicist says (2014, February 5) retrieved 26 June 2024 from <https://phys.org/news/2014-02-stronger-benchmarks-fully-quantify-quantum.html>

<p>This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.</p>
--