What problems can you solve on a quantum annealer?

Enlarge / The D-Wave hardware is, quite literally, a black box.

reader comments

37 with 31 posters participating

Share this story

The D-Wave quantum annealer isn’t a general-purpose computer, in that it can only solve a set of problems that can be structured as energy minimizations. And even on those problems, D-Wave employees will acknowledge that the current generation of D-Wave hardware generally can’t outperform algorithms implemented on standard computers (though they’re optimistic that the next generation of machines may change that in some cases).

But a key reason the company has been selling time on their machines before they have a clear advantage is to give developers the chance to identify the sorts of problems where quantum annealing will prove to be effective. Last week, Ars attended a D-Wave user’s group meeting, where we got a chance to talk to the people who are developing software to run on these systems. We also spoke with D-Wave’s staff. What emerged was a sense of the sorts of things that people are hoping will be able to demonstrate a clear speedup when run on a sufficiently advanced quantum annealer.

Quantum annealing

D-Wave's systems work through a process called quantum annealing. This begins by placing a system's qubits in an absolute energy minimum. From there, the hardware gently alters the configuration of the system so that its energy landscape reflects the problem that needs to be solved. If everything goes well, all the qubits will end up with the lowest possible energy in the new landscape. Viewed literally, this will end up identifying the lowest energy state of that landscape. But if the energy landscape represents something else—some other problem restructured so that it looks like an energy minimization—the final state will represent a solution to that problem.

That's how it's supposed to work, at least. Energy landscapes tend to look like a collection of peaks and valleys, and there's a chance that the system will get stuck in a valley that's not the absolute minimum. While the system should be able to take advantage of quantum effects to tunnel out of these valleys, the process isn't deterministic; sometimes, it will simply remain stuck. In some cases, this won't matter much, as a good minimum can be nearly as valuable as the best. In cases where it does matter, the same problem can be run multiple times, with the best solution identified by its frequency.

(It's also possible to do what's called "reverse annealing," in which the system is started out in a known minimum and then set loose to see if it tunnels to a better one.)

The annealing itself is a very fast process; it's possible to do multiple runs to sample solutions in a fraction of a second. But that's not the only thing at issue here. To begin with, there is some overhead with setting up the problem and transferring it from traditional computers to the hardware that performs the quantum annealing. And as we described when

we covered D-Wave's new hardware, the complexity of the problem that can be solved depends in part on the number of qubits and in part on the connections among them. It's again helpful to think of this as an energy landscape—the more qubits you have, the larger the energy landscape that you can explore at once.

The other challenge annealing faces is that classical algorithms are constantly improving. Fengqi You of Cornell University estimated that thanks to the combination of computing power and better algorithms, a computation that would have taken 124 years in 1988 would now take one second. And D-Wave's Cathy McGeoch said that in response to some of the early results produced on D-Wave hardware, people started revisiting some classical algorithms and made large strides in performance.

What's likely to work?

Given the performance of many of these algorithms on classical computers, one of the main areas of work has been finding situations where quantum annealing will provide an obvious performance advantage. Some of these efforts are quite straightforward. Since the quantum annealer performs an energy minimization, other contexts where energy needs to be minimized should map directly to the sorts of algorithms that run well on this hardware. There are a lot of optimization problems—traveling salesman problems, workflow optimizations, and more—that also map relatively directly to energy minimization.

The challenge at the moment is that classical algorithms often work well on smaller optimization problems but choke on larger ones. While a quantum annealer might have advantages for these complex problems, the current generation of hardware simply doesn't have the qubits needed to handle the large problems. In response, a number of people are working on "hybrid

algorithms," programs where some of the work is done on a classical computer, and then the parts that are most likely to see a benefit from the D-Wave hardware are run there.

This can involve breaking up a large energy minimization problem into a series of smaller ones and solving the smaller ones on the D-Wave hardware. Alternately, sets of solutions can be found using classical computations and then be tested for optimality using quantum annealing. Cornell's You found that to be the case with a scheduling optimization problem, where optimizing the schedules of 18 items ground to a halt on an all-classical system but could still be performed with a hybrid approach.

For hybrid algorithms, the back and forth between classical and annealing worlds means regular communications with the D-Wave hardware, which helped motivate the company to try to reduce the latency of those communications. In addition, the company is working on software that would ease the process of identifying how best to break up a problem into smaller parts—D-Wave's Alan Baratz told Ars that there's a "D-Wave hybrid" framework in the works.

Another approach people are taking is focusing on the fact that even the best classical algorithms may have a subset of cases where they either slow down or fail entirely. This often occurs when the algorithms have to explore an extremely large problem sequentially, which either takes a lot of time or exceeds the memory available. Sridhar Tayur of Carnegie-Mellon found cases like this in non-linear optimizations, where a hybrid algorithm could find an optimization after two to three calls to the D-Wave machine, while a classical algorithm took roughly eight hours to grind through the problem.

The challenge, however, is figuring out if your problem falls into the group of cases where classical algorithms would choke. The amount of time it takes to do that would be added to the amount of time used to solve the problem with a hybrid

algorithm and could potentially negate the benefits of quantum annealing.

Close to the qubits

Other people at the meeting discussed alternative approaches to getting more out of the D-Wave hardware, focusing on how the problems are structured to operate on a quantum annealer. One option is structuring problems as something called an Ising Hamiltonian. University of Maryland Baltimore County's Ramin Ayanzadeh said that it's possible to represent a single problem using different Ising Hamiltonians, and it wasn't clear in advance which of these would perform best on D-Wave hardware. He's using machine learning to see if there's a way of identifying the optimal formulation for a quantum annealer.

Ising Hamiltonians can also be converted to what are called Quadratic unconstrained binary optimization problems, or QUBOs. George Hahn, who is presently at Harvard, talked about some work he did at Los Alamos on examining the structure of QUBOs. He found that it's possible to identify cases where certain terms in the QUBO can be replaced by a constant, and others where two terms have a consistent relationship. This allows the QUBO to be simplified down, potentially allowing larger problems to fit on D-Wave's hardware.

Two government researchers were also experimenting with the D-Wave hardware itself to understand the annealing process better. Scott Pakin, who works at Los Alamos, described how researchers were essentially starting an optimization and then bringing it to a halt so they could examine what the annealer looked like at intermediate states in the process. Similarly, NASA's Shon Grabbe has found that if you temporarily halt the system at specific points in the annealing process, it's possible to get a five-fold improvement in the solution provided in a given sampling.

In the end, some of these approaches are undoubtedly going to

come up short of providing a practical solution. But it’s clear there are a lot of people looking into ways of finding problems where the quantum annealer could have an edge over traditional ones. Given that this is happening in parallel with hardware improvements, it’s easy to see why a number of people at the meeting felt that we’ll be seeing some clear advantages before the end of 2020.

In a future article, we’ll take a look at some of the specific problems that people think are worth solving with quantum annealing.

https://arstechnica.com/?p=1577435