IBM compensates for errors, gets usable results out of quantum processor

Enlarge / IBM’s Eagle processor has reached Rev3, which means lower noise qubits.

reader comments

30 with

Today’s quantum processors are error-prone. While the probabilities are small—less than 1 percent in many cases—each operation we perform on each qubit, including basic things like reading its state, has a significant error rate. If we try an operation that needs a lot of qubits, or a lot of operations on a smaller number of qubits, then errors become inevitable.

Long term, the plan is to solve that using error-corrected qubits. But these will require multiple high-quality qubits for every bit of information, meaning we’ll need thousands of qubits that are better than anything we can currently make. Given that we probably won’t reach that point until the next decade at the earliest, it raises the question of whether quantum computers can do anything interesting in the meantime.

In a publication in today’s Nature, IBM researchers make a strong case for the answer to that being yes. Using a technique termed “error mitigation,” they managed to overcome the problems with today’s qubits and produce an accurate result despite the noise in the system. And they did so in a way that clearly outperformed similar calculations on

classical computers.

Living with noise

If we think of quantum error correction as a way to avoid the noise that keeps qubits from accurately performing operations, error mitigation can be viewed as accepting that the noise is inevitable. It’s a means of measuring the typical errors, compensating for them after the fact and producing an estimate of the real result that’s hidden within the noise.

An early method of performing error mitigation (termed probabilistic error cancellation) involved sampling the behavior of the quantum processor to develop a model of the typical noise and then subtracting the noise from the measured output of an actual calculation. But as the number of qubits involved in the calculation goes up, this method gets a bit impractical to use—you have to do too much sampling.

Advertisement
So instead, the researchers turned to a method where they intentionally amplified and then measured the processor’s noise at different levels. These measurements are used to estimate a function that produces similar output as the actual measurements. That function can then have its noise set to zero to produce an estimate of what the processor would do without any noise at all.

To test this system, the researchers turned to what’s called an Ising model, which is easiest to think of as a grid of electrons where each electron’s spin influences that of its neighbors. As you step forward in time, each step sees the spins change in response to the influence of their neighbors, which alters the overall state of the grid. The new configuration of spins will then influence each other, and the process will repeat as time progresses.

While an Ising model involves simplified, idealized behavior,

it has features that show up in a variety of physical systems, so they've been studied pretty widely. (D-Wave, which makes quantum annealers, recently published a paper in which its hardware identified the ground state of these systems, so they may sound familiar.) And as the number of objects in the model increases, their behavior quickly becomes complex enough that classical computers struggle to calculate their state.

It's possible to perform the calculations on a quantum computer by performing operations on pairs of qubits. To simplify matters, IBM used an Ising model where the grid was configured in a way that coincides with the physical arrangement of qubits on its processor. But this wasn't a case where the processor was simply being used to model its own behavior; as mentioned above, Ising models existed independently of quantum hardware.

https://arstechnica.com/?p=1947845