Quantum computing’s also-rans and their fatal flaws

Enlarge / IBM’s 16-qubit quantum computer from 2017.

reader comments

20 with 12 posters participating

Share this story

Last month, Google claimed to have achieved quantum supremacy—the overblown name given to the step of proving quantum computers can deliver something that a classical computer can’t. That claim is still a bit controversial, so it may yet turn out that we need a better demonstration.

Independently of the claim, it’s notable that both Google and its critics at IBM have chosen the same type of hardware as the basis of their quantum computing efforts. So has a smaller competitor called Rigetti. All of which indicates that the quantum-computing landscape has sort of stabilized over the last decade. We are now in the position where we can pick some likely winners and some definite losers.

Why are you a loser?

But why did the winners win and the losers lose?

In the end, the story comes down to engineering. A practical quantum computer requires that we can create many quantum bits (qubits). Those qubits have to stay in a quantum state for multiple gate operations. Gate operations require that we are able to manipulate qubits on both an individual basis and in

groups (or at least pairs). And, of course, you have to be able to read out the result of a computation.

Many of these individual features have been demonstrated to work using qubits in liquids, in Rydberg atoms, in Bose Einstein condensates (BECs), in solid-state systems, nitrogen vacancies in diamond, defects in silicon, trapped ions, light, and of course, superconducting rings. That list is incomplete, but, most of those options are pretty much dead in the water, and for very good reasons. While qubit behavior is dictated by physics at the individual qubit level, once you think about scaling, engineering really matters, and a lot of these options aren't very amenable to scaling.

Randomness is bad

Early in the decade, nitrogen-vacancy centers, silicon vacancies (which took a lot longer), and solid-state materials were among the front runners. These materials all operate on similar principles: a small percentage of a contaminant material is introduced to a crystal. Nitrogen is put in diamond, phosphorous in silicon, and ytterbium in yttrium-aluminum-garnet crystals.

The qubits in each material are formed by similar physics. The contaminant material can't match the bonding requirements of its neighboring atoms, which will leave an isolated electron or positively charged nucleus (an ion). The states of these isolated objects can be used as a qubit, and its states can last for a very long time—often longer than their more successful counterparts.

But there are fundamental disadvantages to these technologies as well. A good example of many of these can be seen in nitrogen-vacancy centers in diamond. Each qubit consists of a single electron left hanging by nitrogen's inability to bond with a fourth carbon atom. This electron is addressed (set and read) optically. Hence, the first problem is to search through

a crystal for the few isolated vacancies that can be individually addressed. Optically addressing the qubits means that the vacancies are too far apart to directly couple, so qubit operations and entanglement between qubits has to be done via optical and microwave photons. Unfortunately, microwave emission will couple to *all qubits*, reducing the precision with which the qubit scan be controlled.

Even worse, each vacancy is different. The quantum properties of the vacancy are determined by the precise arrangement and type of atoms that surround it. For instance, in diamond, the two common isotopes of carbon provide enough difference that the presence of carbon-13 changes the performance of nearby qubits. To make the qubits identical, local magnetic fields need to be applied, which shifts the energy levels of the qubit states. That needs to be done by running relatively strong currents through nearby wires but simultaneously isolating the effects so that they don’t influence other qubits.

Essentially, every chip of diamond will produce a different computer, with a different arrangement of qubits that have different properties. The wire routing to ensure that the local magnetic fields are truly local to the target qubit seems insanely difficult. Then you have to engineer tiny lens arrays (milled directly onto the diamond surface) to couple all the qubits to the outside world.
The tiny, suppressed part of my brain that understands engineering screams at the very thought.

These issues apply to almost all vacancy-based qubit systems, which is why we’re hearing less and less about them.

https://arstechnica.com/?p=1594833