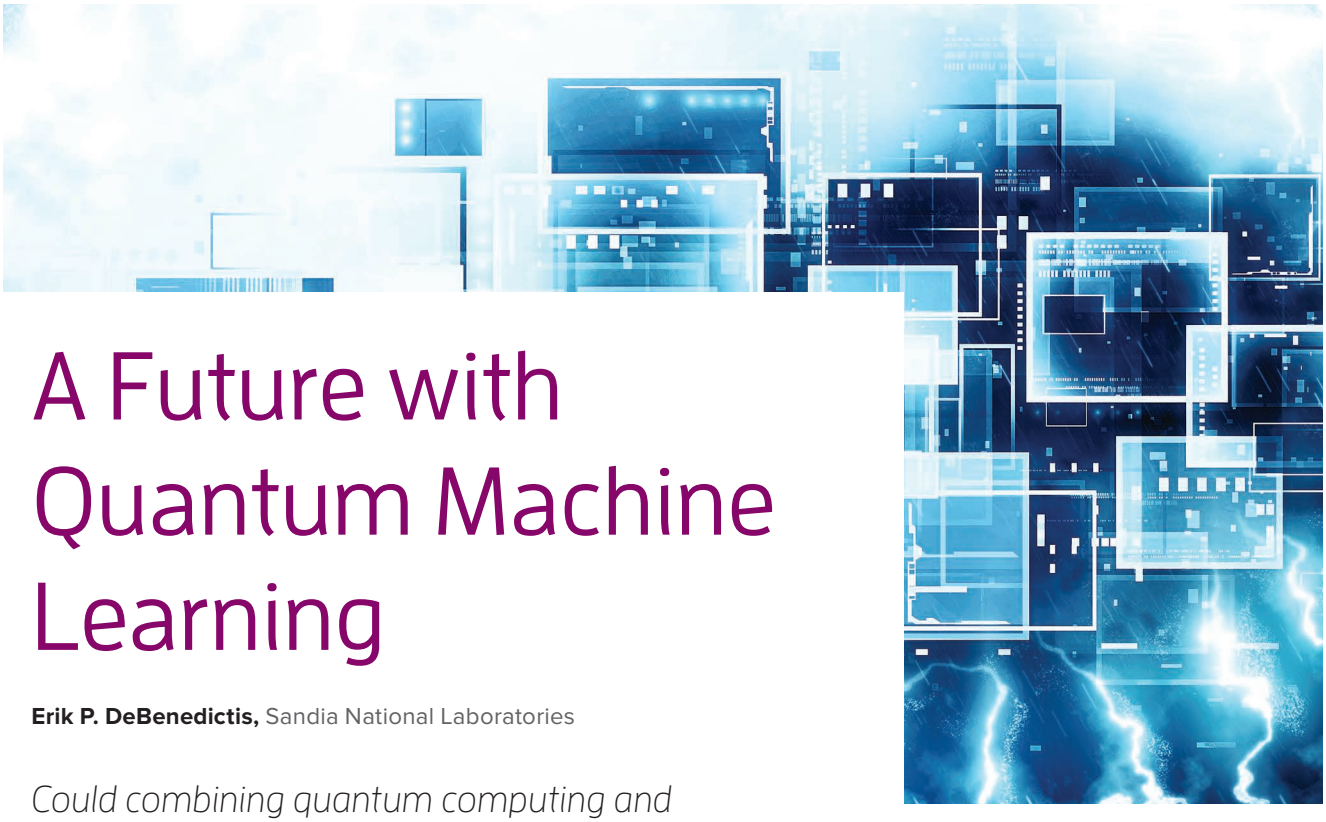


A Future with Quantum Machine Learning

Erik P. DeBenedictis, Sandia National Laboratories

Could combining quantum computing and machine learning with Moore's law produce a true "rebooted computer"? This article posits that a three-technology hybrid-computing approach might yield sufficiently improved answers to a broad class of problems such that energy efficiency will no longer be the dominant concern.

Over the past two years, this column has highlighted the technologies being considered as candidates to reboot computing. Yet none of the individual technologies has been very exciting. For example, in December 2016 I wrote about neuromorphic crossbars' potential as machine learning accelerators,[1] concluding that they face the same thermodynamic limit as today's microprocessors.

However, in the last year, the combination of quantum computing, machine learning, and Moore's law has taken form with more potential than anything seen to date.[2] While it's too early to tell whether this three-technology hybrid will survive the test of time, the field is attracting both venture capital and substantial investment by big companies and the government. My goal here is to show how this combination of technologies has a synergy that could affect people outside "the club" creating it.

MACHINE LEARNING CHANGES THE QUANTUM GAME

To explain the combination of quantum computing and machine learning, we must express quantum computing in terms that don't unnecessarily hide the role of programming, because this would obfuscate machine learning's value in shifting some portion of human-developed programming to computers.

Quantum computers' popular success story is the factoring of large numbers. Historically, numbers were

0018-9162/18/$33.00 © 2018 IEEE

EDITOR **ERIK P. DEBENEDICTIS**
Sandia National Laboratories;
epdeben@sandia.gov



factored using trial division, which requires a three-line program. However, the three lines iterate an exponential number of times ($2^{n/2}$) when factoring an n-bit number. This leads to an exponential expenditure of energy, given the thermodynamic kT model of minimum energy per binary operation. The community has explored two options to reduce the cost of factoring numbers:

1. The subexponential number field sieve algorithm was developed using perhaps 100 person-years effort by mathematicians, computer scientists, and programmers.
2. Shor's polynomial-time quantum algorithm was developed, albeit requiring a quantum computer that has yet to be built.

Both improvements are important.

The discovery of quantum algorithms occurred in parallel with improvements to the equivalent classical algorithms, leading to competition between the 100 person-years research and the special properties of quantum information. However, computational complexity theory seeks the best algorithm without reference to algorithm development and programming effort, unfairly disadvantaging quantum computers. This retelling of the quantum computer story opens the door for machine learning to contribute by making programming more efficient.

Machine learning moves the dividing line between humans and computers. In a typical machine-learning scenario, a server farm consumes a dollar's worth of energy learning how to recognize your pet in images or how to target advertisements by scanning your emails. This learning might compute neuron weights for a recognition circuit, which is essentially a program. In many cases, each person's pet and mailbox have an underlying structure similar to the number field that enabled improved factoring algorithms. It might be possible to improve the computational efficiency of the neural network that cost a dollar to synthesize in the first place through 100 person-years of research. However, there's no way to recoup 100 years' salary, given that the learned behavior is applicable to only one person.

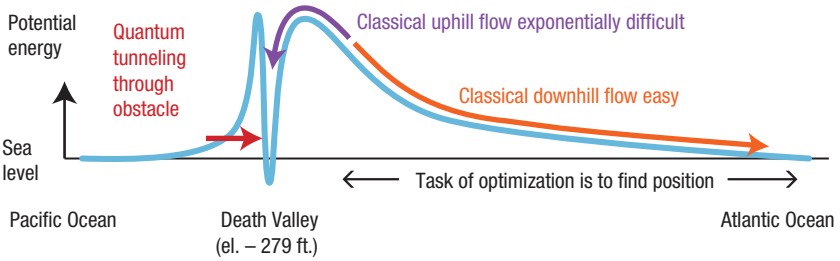


Figure 1. Optimization involves finding the lowest point on a potential energy curve (blue), which is Death Valley, even though most water flows to the oceans. Classical optimization (orange) works like raindrops flowing downhill, but simulated annealing allows limited uphill movement (purple). However, quantum computer optimization can use a quantum physics principle called "tunneling" to go through a high energy barrier (red). The text also describes how this type of optimization could apply to organizing the slides in a slide deck to make a compelling presentation.

The opportunity for quantum machine learning will be in learning lots of simple lessons—concepts that will make society more efficient, not just the hard problems currently attracting geniuses and armies of researchers. I suggest that quantum machine learning be benchmarked on learning a completely original behavior and performing it as few as, say, 10 times. The cost metric would include both the learning and running times.

CHIP LAYOUT TO SLIDE DECKS

How can a quantum computer's computational advantage in optimization,[3] factoring numbers, and other algorithms be repurposed to machine learning? While classical computers can perfectly optimize small systems, they only find incremental improvements for large systems such as transportation routes and product pricing. This is due to their rapidly rising running time as a function of problem size.

Placement of logic gates on an integrated circuit is an example. Chip design tools have optimizers that place logic gates on a chip's surface with just enough space between to hold the wiring that defines the chip's function. Better placement reduces chip area—and hence cost—while simultaneously increasing the chip's speed because the shorter wires convey information in less time. However, a chip might be profitable even if it's a few percent larger than necessary, so perfect optimization isn't essential.

Classical placement algorithms such as simulated annealing follow the same principle as raindrops trying to find the lowest elevation by flowing downhill. Figure 1 shows an energy landscape for water by position across the US. Water dropped almost anywhere will flow to an ocean. Oceans are low, but not as low as Death Valley. However, Death Valley has a small rainfall basin surrounded by high mountains, so a random raindrop would be unlikely to fall into its basin.

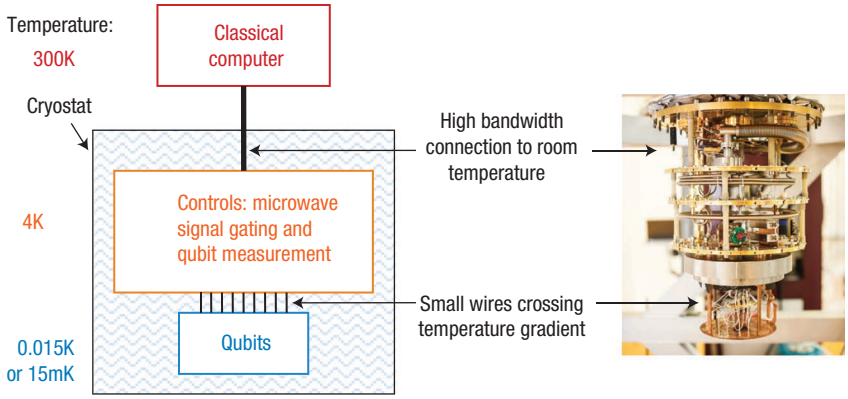


Figure 2. General structure of a quantum computer system. The user interacts with a classical computer. If the problem requires optimization, the classical computer translates the user's problem into a standard form for a quantum computer, such as QUBO, or into a different form if another quantum algorithm is required. The classical computer then creates control signals for qubits (quantum bits) located in a cryogenic environment, receiving data from measurements of the qubits. Some classical electronics are placed in the cold environment to minimize heat flow through wiring across the cryogenic-to-room-temperature gradient. Photo source: A. Hellemans, "Europe Bets €1 Billion on Quantum Tech," *IEEE Spectrum*, 22 Jun. 2016; spectrum.ieee.org computing/hardware/europe-will-spend-1-billion-to-turn-quantum-physics-into-quantum-technology.

Mathematicians, computer scientists, and programmers have improved simulated annealing so that potential solutions can jump over an obstacle, but the probability of this occurring decreases exponentially with the height of the jump. Human effort has also created heuristics, such as chip design tools that handle memories, busses, and clock lines in special ways.

One form of quantum machine learning uses "quantum tunneling"[3] to go through the peak in Figure 1, with the probability of this occurring declining exponentially with the width of the peak. The tunneling approach may or may not be better than simulated annealing, but applying both techniques might give a better answer than either alone. Other quantum algorithms work quite differently, such as not using potential energy at all.

Optimization can be applied to development of a slide deck for a presentation, such as tuning the ordering of the slides to meet the expected interests of a particular audience. Hypothetically, there is a "potential energy landscape" for every audience based on the listeners' background knowledge and receptiveness to new ideas. For example, one audience might prefer an emotional appeal first and facts later. If Figure 1's horizontal axis represents slide interchanges, optimization just needs to find the sequence of interchanges that yields the most compelling presentation. There are exponentially many orderings—too many to test exhaustively—so the classical approach is to follow downhill paths as shown in Figure 1, or slide interchanges that each make each potential presentation a little better than the previous. However, a quantum computer's unique ability to tunnel through high potential barriers might let it find the most compelling slide deck when simulated annealing cannot.

Presentations can be optimized through use of human labor, such as mock juries in criminal trials. However, the effort required is too high for everyday situations.

A VISION FOR FUTURE APPLICATIONS

I've painted a picture in which today's corporate applications, such as optimizing transportation routes, are improved and then applied to everyday personal situations. But are there enough such applications to bother with? Computers assist people with numerical calculations countless times a day, such as when a smartphone computes how far you jogged. But there are also occasions when you need to say or do something that requires nonnumerical judgement—such as preparing a compelling slide-deck presentation, as in my previous example, or answering a question in a way that impresses your boss. With today's knowledge and technology, there should be as many ways for computers to address these nonnumerical activities as the numerical ones.

MOORE'S LAW AND SUPERCONDUCTING ELECTRONICS

In the early 1940s, IBM president Thomas J. Watson reputedly said, "I think there is a world market for maybe five computers." If quantum machine learning meets the expectations of the venture capitalists who are funding start-ups, Watson's statement won't hold for these three-technology hybrid computers either, and we'll need a path to produce them in large volume.

The quantum effect frequently, although not always, requires components operating near absolute zero, making just about every aspect of the design exotic. Quantum computer components operated at room temperature inevitably acquire error from the thermal motion of the atoms in the computer's structure. The errors must be removed by quantum error correction, yet the error accumulation rate is too high for practical removal unless the components are cooled to millikelvins, or thousandths of a degree above absolute zero—273.15 °C or 0 K.

The architecture of these quantum–classical hybrid computers is zeroing

in on the structure shown in Figure 2. The qubits (quantum bits) must be kept at a temperature of approximately 15 mK. They need support from classical superconducting electronics based on Josephson junctions operating at temperatures around helium's boiling point, or 4 K.[4] The electronics must have extremely low energy dissipation, because the external refrigeration must dissipate at least the temperature ratio (300 K/4 K = 75× or 300 K/15 mK = 20,000×) times as much energy to remove the heat to room temperature (300 K)—and, in practice, several times this amount. Logic-gate circuits based on Josephson junctions are available that perform the logic functions for error correction as well as the gate microwave signals required to control qubits.

As Moore's law demonstrated, industry knows how to take control of a technology and work relentlessly to improve it to the limits of physics. So if the ideas in this column pan out, industry will need to do the same for superconducting electronics.

SOCIETAL IMPLICATIONS

So far I've discussed future computers as though they were standalone, yet we use computers today as agents for many of our business transactions and information handling. Some day we might have the option to upgrade our computerized agents to more advanced versions that use quantum machine learning internally. However, we must be prepared for competitors and bad actors that upgrade early to take advantage of us.

If you are a defense lawyer trying to defend a client, you will be at a disadvantage if your presentation is less well tuned to the jury than the prosecution's is. Similarly, web traffic is monitored by machine learning software purportedly to send us advertisements, but bad actors can use the same technology to better target phishing emails that can cause us harm. Computers can use machine learning to find phishing emails, but this will lead the opposing sides to mount an arms' race for better quantum computers.

My December 2016 article comparing the energy efficiency of analog memristor-style crossbars for learning showed that this new technology was subject to the same thermodynamic limits as digital chips. Each approach might beat the other in some portion of a parameter space, but the common limit implied that the best win would be an order of magnitude or two. Companies could live or die based on a couple orders of magnitude in product performance, but changing the world typically requires a bigger difference. The triad of quantum computing, machine learning, and a continuation of Moore's law could possibly address a broad class of problems, with only distant competitors. So what are the practical challenges that quantum machine learning must overcome to survive the test of time?

There will be technical challenges beyond just building hybrid quantum–classical hardware. We haven't systematically looked for applications that depend on exorbitant amounts of machine learning or optimization, nor have we applied quantum computing to general problem solving.

The computer industry has been producing chips intended to operate at room temperature, which was convenient. A quantum–classical computer, however, has unique capabilities that require a cryogenic environment. Materials, devices, and circuits for this environment are known but haven't been refined to the same level of manufacturability as semiconductors.

Classical computers' rapid emergence has stretched society's ability to assimilate their capabilities, creating concerns regarding cybersecurity, robots and AI, social media, and so on. Rolling out quantum machine learning products could introduce similar issues, but they should be seen as challenges to overcome, not reasons to hold back progress or ignore the uncomfortable questions they present.

ACKNOWLEDGMENTS

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International, Inc., for the US Department of Energy's National Nuclear Security Administration under contract DE-NA-0003525.

REFERENCES

1. E. P. DeBenedictis, "Computational Complexity and New Computing Approaches," *IEEE Computer*, vol. 49, no. 12, 2016, pp. 76–79.
2. J. Biamonte et al., "Quantum Machine Learning," *Nature*, vol. 549, no. 7671, 2017, pp. 195–202.
3. V.S. Denchev et al., "What is the Computational Value of Finite-Range Tunneling?," *Physical Review X*, vol. 6 no. 3, 2016, p. 031015.
4. M. Veldhorst et al., "Silicon CMOS Architecture for a Spin-Based Quantum Computer," arXiv, 30 Sept. 2016; arxiv.org/abs/1609.09700.

ERIK P. DEBENEDICTIS is a technical staff member at Sandia National Laboratories Center for Computing Research. He is a member of IEEE, ACM, and APS. Contact him at epdeben@sandia.gov.

myCS Read your subscriptions through the myCS publications portal at **http://mycs.computer.org**