Quantum Realism

Travis S. Humble, Oak Ridge National Laboratory
Erik P. DeBenedictis, Sandia National Laboratories

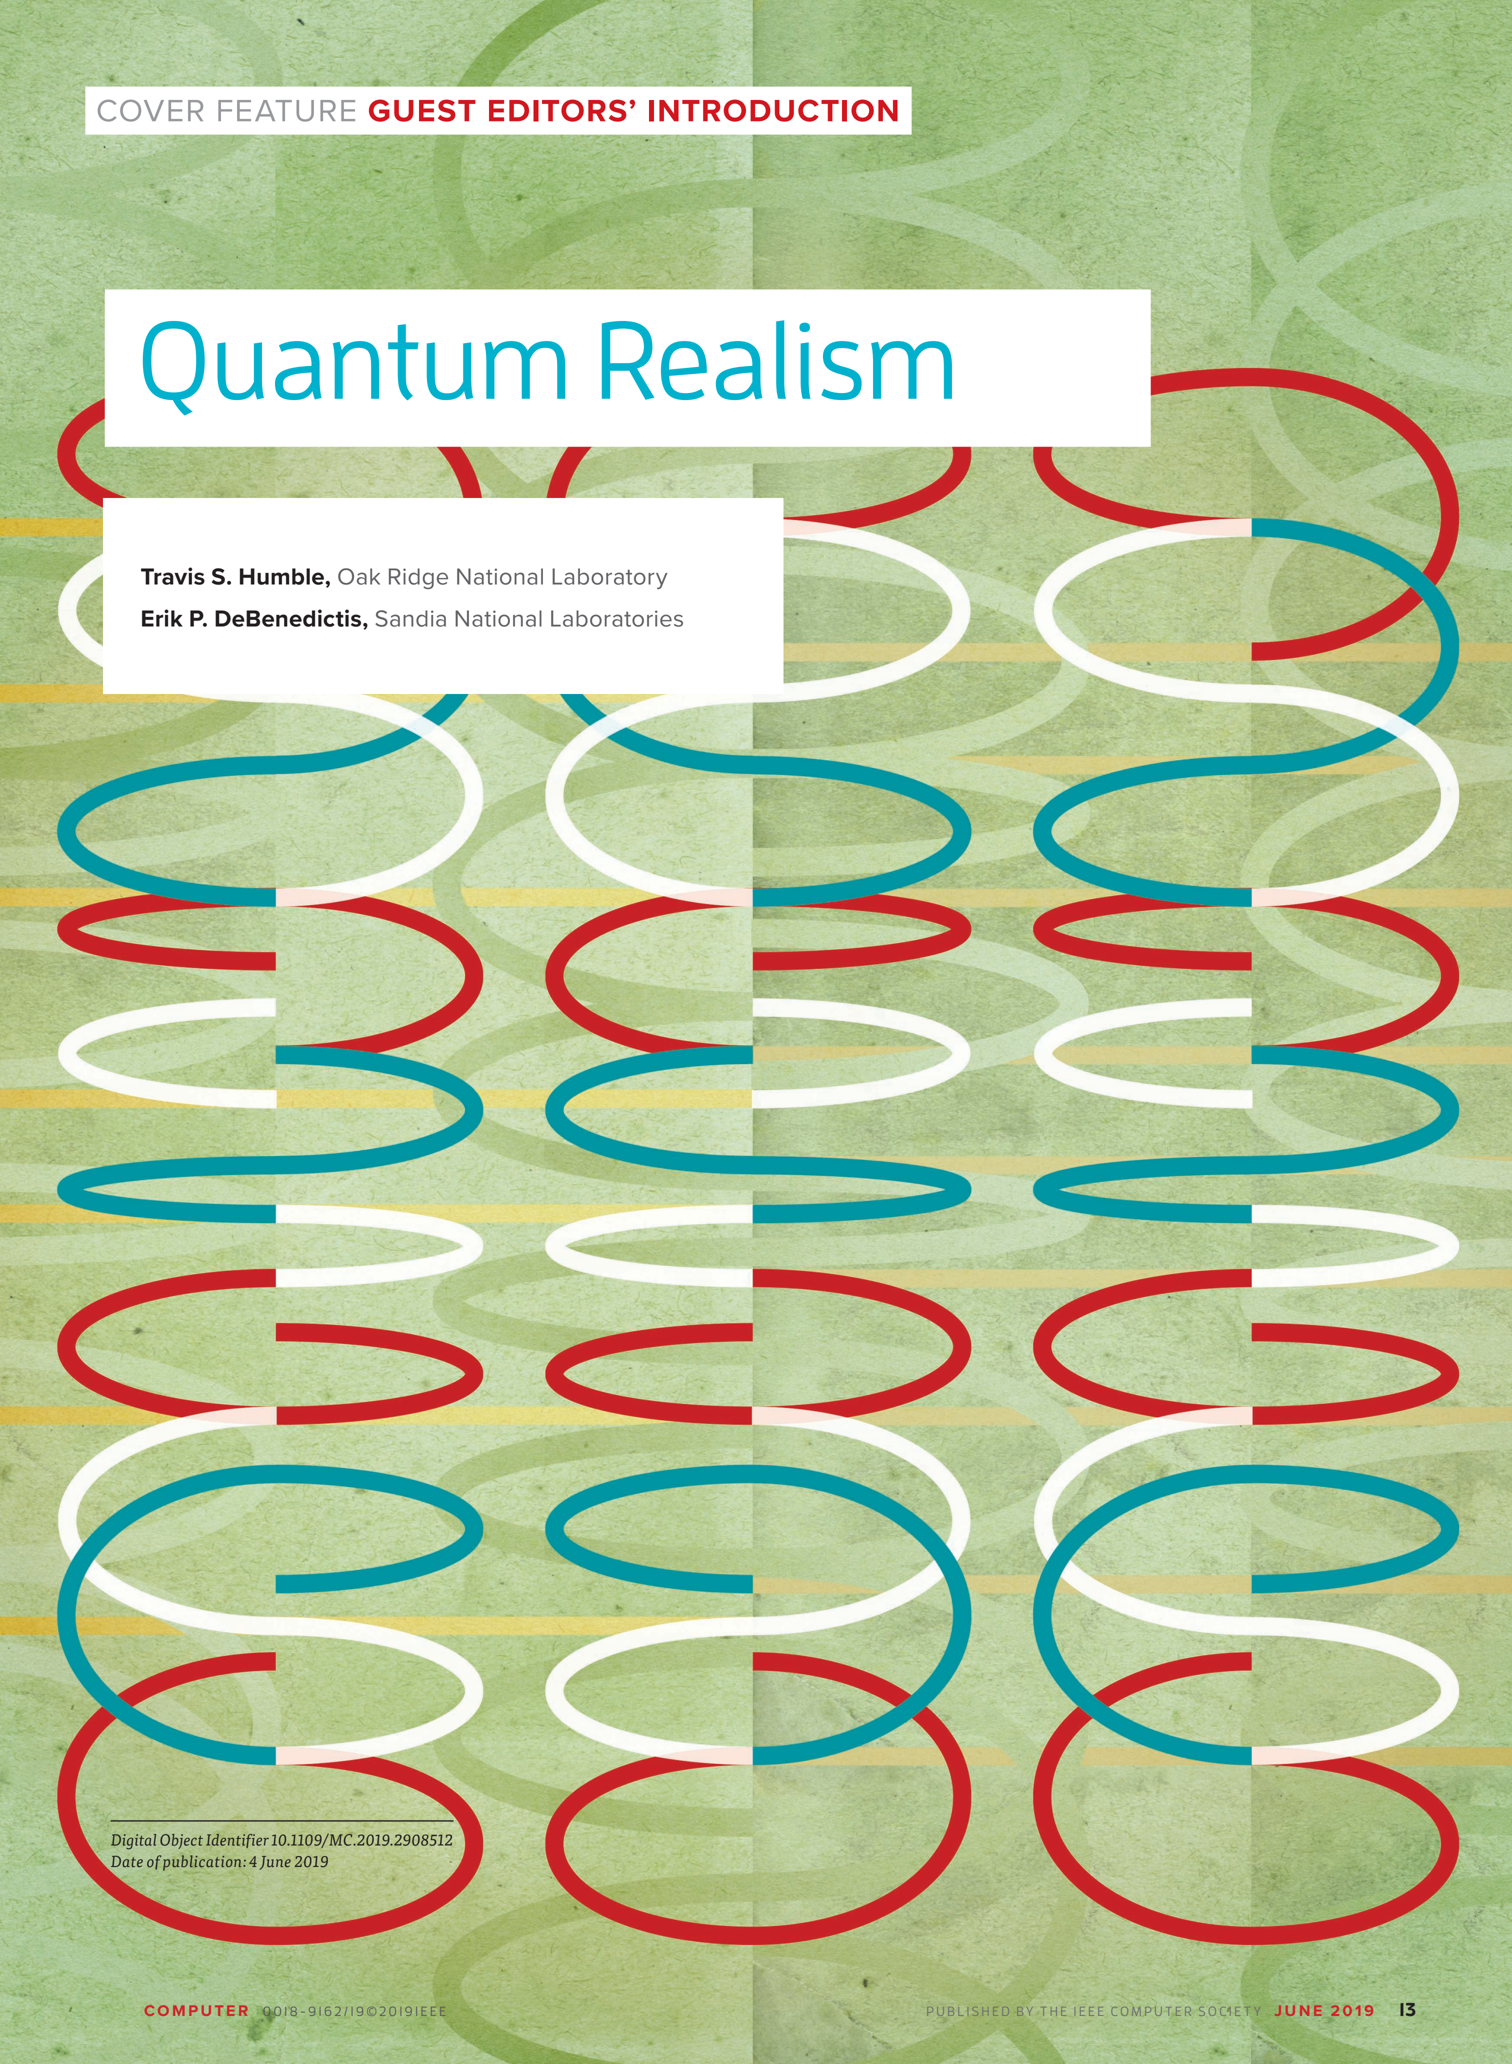

Digital Object Identifier 10.1109/MC.2019.2908512
Date of publication: 4 June 2019

In this issue, we take a look at quantum computing: what has been accomplished and what is its future. We present a collection of views about the realistic challenges facing the future of quantum computing.

The field of quantum computing has quickly become a topic du jour for computer scientists in industry, academia, and government. Even a politician or two has added to the conversation. Their excitement is motivated by the palpable expectation that quantum computers are coming and these new computers will bolster national security, accelerate scientific innovation, and boost computational power. Such remarkable hopes come from the equally remarkable theories that have harnessed the laws of quantum mechanics for the power of computing. But what is realistic for one of the most rapidly rising fields in computer science?

Computational complexity theorists have long hinted that quantum computers could surpass conventional expectations of what is possible with classical computers. These results have suggested roles for quantum computers in, for example, cracking public-key encryption by trivializing the problem of integer factorization and dramatically reducing the time needed to perform database searches. Nobel Laureate Richard Feynman went further. He made the case for why a quantum computer will be the key to unraveling the mysteries of nature. Using quantum computers to study chemical reactions and materials science seems almost obvious for advancing drug discovery and nanotechnology.

These extraordinary expectations are now fueling a global effort to perfect quantum computing. Multibillion-dollar national initiatives in North America, Europe, Asia, and Australia are bringing the theory of quantum computing closer to reality, and they are being matched by investments from industry and venture capitalists. These investments are as much about economic security as national security with the technological basis of the 21st century at stake. Scientists are moving to meet this demand. Professional societies, including the IEEE, are focusing on education, workforce development, and communication to enable the desired transformative change. For example, the IEEE Future Directions Committee has recently launched an initiative to define how the field of engineering can shape the future of quantum computing.

[USING QUANTUM COMPUTERS TO STUDY CHEMICAL REACTIONS AND MATERIALS SCIENCE SEEMS ALMOST OBVIOUS FOR ADVANCING DRUG DISCOVERY AND NANOTECHNOLOGY.]

In light of this enthusiasm, a realistic examination of the technical challenges facing quantum computing is needed. Current quantum computing devices are available both commercially and in private laboratories, but they are far from useful. As a field of practice, quantum computing is in its infancy with first-in-kind advances of experimental hardware still routine. These nascent devices are barely able to perform the simplest computations and only for cases in which the answers are already known. These experiments provide clear hints of progress, and maybe even promise, but too little is known about how scientists and engineers may overcome the challenges needed to operate devices at scales that are computationally useful. As a case in point, many theoretical estimates for the resources required to factor numbers or search through databases exceed the current state of the art by a millionfold. It is certainly not clear if or how the current technology could scale up to such numbers.

Nonetheless, these estimates may be too conservative. Unequivocal demonstrations of quantum computational utility may be much closer. If experimental hardware can be increased by even a factor of 10 in size and quality, then their rudimentary calculations of quantum mechanical systems could even surpass those of our largest supercomputers. But this quantum computing milestone assumes that conventional methods, including algorithms and architectures, do not achieve similar progress. While this may seem to augur competition between quantum and classical computing, it appears far more realistic to expect that these two approaches will become codependent in form and function.

What problems will quantum computers help to solve better than classical-only computers, and how much impact can these advances have on our global society? And if a quantum computer can help overcome a known but previously intractable problem, will this sustain the promise of a new technological base? What products will quantum computing companies sell, and how

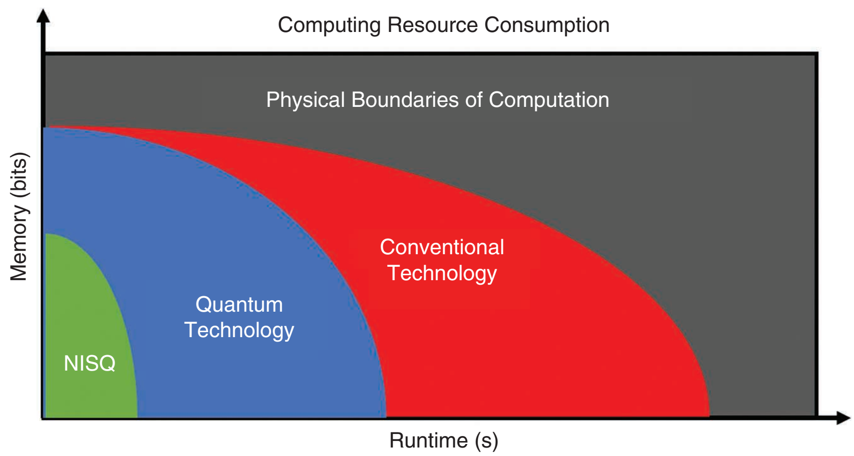

FIGURE 1. A notional representation of the resources available from different computing technologies and how they relate to each other. While conventional and quantum technologies may be able to solve similar problems, quantum approaches are expected to require significantly less runtime. However, it remains an open question as to whether NISQ devices could offer such advantages.

will the rest of the world respond? What are the lessons a quantum computer might learn, which a classical computer cannot, that would advance machine learning and artificial intelligence?

We frame these articles in terms of Figure 1, which provides a notional representation for how the scale of computing resources typically defines solvable problems. We focus on runtime and memory, where the latter may be in bits because even a quantum computer must generate a classical output. For example, designs for future high-performance computing systems provide a boundary for modern computing, such that any computer program whose resource requirements fall within the red area should be feasible. Similarly, mean-time-between-failure estimates how long a conventional computer can run before a hard fault occurs. Meanwhile, quantum technology is expected to provide computational advantages for certain classes of problems once systems of sufficient size and quality are constructed.

These problems define the blue region in Figure 1, which is bounded by the size of the problem and the fewest number of steps required to compute the solution. As the list of problems that quantum computing might help address continues to grow, vigorous debate goes on about specific challenges. Finally, we may speculate what should be the ultimate boundaries for computing, shown as the black area, based on the size and remaining lifespan of the universe. Given a log–log scale for our axes, this black area appears just a bit further out. But where in Figure 1 lies the field of quantum computing today, and where will it go in the future?

Today's demonstrations of quantum computing lie in the green area in Figure 1. These early experimental systems have relatively small memory along with relatively brief runtimes for performing reliable computations. Progress over the next several years is anticipated to favor noisy, intermediate-scale quantum (NISQ) devices, a term coined only recently. But the NISQ regime may itself outperform classical methods. In fact, to maintain speculative investments in the field, the overall quantum computing enterprise must demonstrate a meaningful milestone. For example, if a problem too large for a supercomputer could be solved with the help of quantum computing technology, then a "quantum advantage" could be established. Of course, this problem would not need to run entirely on a quantum computer. An equally convincing demonstration could use a smaller NISQ device alongside a conventional computer. Such an approach would split the computational workload across paradigms, enabling each type of computer to work on the portion of the problem it is best suited to address.

[CURRENT QUANTUM COMPUTING DEVICES ARE AVAILABLE BOTH COMMERCIALLY AND IN PRIVATE LABORATORIES, BUT THEY ARE FAR FROM USEFUL.]

A LOOK AT THE ARTICLES

In this special issue, we present a collection of views about the realistic challenges facing the future of quantum computing. We sought the opinion of renowned experts in the emerging field of quantum computer science to identify and explain the balance between the possibilities and the reality, as well as the limitations, of various aspects of quantum computing. While these contributions may not solve the open problems of quantum computing, they offer the reader a look into the current state of the field, the challenges that are anticipated, and some of the best efforts underway to get past them. The articles address these different challenges and offer insights into how the future of quantum computing may unfold. All of them focus on forthcoming changes in how we use quantum computers, and they give indications for how the ideas of quantum computing are migrating from physics experiments to bona fide computational platforms.

In "A Hybrid Approach for Solving Optimization Problems on Small Quantum Computers" by Shaydulin et al., a variation of the hybrid computing model shows how a problem too large for current hardware can be solved using a series of quantum classical programs. The authors adopt a programming paradigm that decomposes an input problem into a sequence of smaller quantum subroutines suitable for the available quantum processing unit (QPU). Using community detection within a large network as a motivating example, they show how to offload only the relevant computation to a QPU while allocating the remainder of the processing to a classical supervising program. Although this method does not yet demonstrate a quantum advantage on available hardware, the authors argue that a hybrid computing model may be able to achieve this first major milestone for solving a real-world problem.

ABOUT THE AUTHORS

TRAVIS S. HUMBLE is with the Oak Ridge National Laboratory, Tennessee, where he studies the application of quantum information technologies to scientific discovery and innovation. Humble received a Ph.D. from the University of Oregon in 2005. He is a Senior Member of the IEEE. Contact him at humblets@ornl.gov.

ERIK P. DeBENEDICTIS is with Sandia National Laboratories, Albuquerque, New Mexico. DeBenedictis received a Ph.D. in computer science from the California Institute of Technology in 1982. He is a Senior Member of the IEEE and a member of ACM and APS. Contact him at epdeben@sandia.gov.

The crucial issue of estimating resources for a quantum computer is addressed in "Really Small Shoe Boxes: On Realistic Quantum Resource Estimation" by Paler et al., in which the individual components contributing to such calculations are examined. This article characterizes these resources using the concept of a 3D volume with an area that represents the size of the calculation and a depth that measures how long it takes. This space–time representation of the computing resources is then analyzed with respect to the different contributions made by hardware, software, and application. Paler et al. argue for realistic estimates of a fault-tolerant quantum computer, in which they use the example of a surface code design to capture the essential steps needed for resource costing of future quantum computing systems.

The contribution from McGeoch et al., "Practical Annealing-Based Quantum Computing," explains quantum annealing, one of the prevailing paradigms for how to build and operate a quantum computer. This article provides an overview of the quantum annealing hardware offered by D-Wave Systems as well as a survey of the practical applications possible using this approach. These applications provide convenient baselines for evaluating the performance of the D-Wave system to other computational methods. McGeoch et al. argue that computationally meaningful performance measures are essential for evaluating the practical relevance of any quantum computing system. They articulate several of the open questions for evaluating the performance of quantum annealing and offer insights into the future of programming this integrated hardware and software system.

The contribution from Hu et al., "Reduction-Based Problem Mapping for Quantum Computing," acknowledges that the adoption of quantum computing faces conceptual barriers. Differences in the underlying hardware logic defy everyday intuition for how to program in the quantum domain. Hu et al. recognize that software can mitigate some of these challenges by automating the reduction of an input problem into a program tailored to run directly on a quantum computer. They focus broadly on the class of NP-complete problems, which are generally hard for conventional computers, and provide an explicit implementation for the case of satisfiability (3-SAT). Using an example of map coloring reduced to 3-SAT and executed on an IBM quantum processor, they show how quantum computing can be more tightly integrated into modern computing workflows.

Compiling efficient code is an essential step in programming, even for quantum computers. Liu et al., in "Stochastic Optimization of Quantum Programs," describe a new approach to optimizing programs using quantum computing while guaranteeing logical correctness. Instead of techniques based on heuristic rules, Liu et al. use an automated approach derived from a stochastic feedback loop to optimize the score of a program implementation. The score accounts for program correctness as well as program efficiency, where measuring the latter depends on the program feature being optimized, e.g., number of instructions. They show how improvements in code compilation may be obtained by selecting from a set of predetermined mutation operators, and their empirical data for the case of a quantum program to solve 3-SAT (compare with the contribution from Hu et al.) hint at new opportunities for quantum compilation techniques.

The prospects of quantum computing are exciting, and we look forward to seeing this technology used in the future, for example, to advance our understanding of materials and medicine. But there are tremendous challenges along the way of harnessing quantum computers to solve real-world problems, and we hope the articles in this issue provide a balanced assessment of this reality. ■

Digital Object Identifier 10.1109/MC.2019.2915413