Quantum Computing Enters 2018 Like It Is 1968

nextplatform.com/2018/01/10/quantum-computing-enters-2018-like-1968/

January 11, 2018

January 10, 2018 Paul Teich

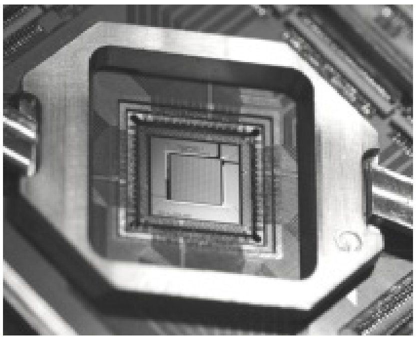

The quantum computing competitive landscape continues to heat up in early 2018. But today's quantum computing landscape looks a lot like the semiconductor landscape 50 years ago.

The silicon-based integrated circuit (IC) entered its "medium-scale" integration phase in 1968. Transistor counts ballooned from ten transistors on a chip to hundreds of transistors on a chip within a few short years. After a while, there were thousands of transistors on a chip, then tens of thousands, and now we have, fifty years later, tens of billions.

Quantum computing is a practical application of quantum physics using individual subatomic particles chilled to millikelvin temperatures as compute elements. These subatomic compute elements are called "qubits." Qubits can be manufactured with CMOS technologies, like standard ICs. But the interconnect, control, and sensor circuits needed to manipulate and coordinate an increasing number of qubits in the intensely cold operating environment of a quantum computer require new science and technology development.

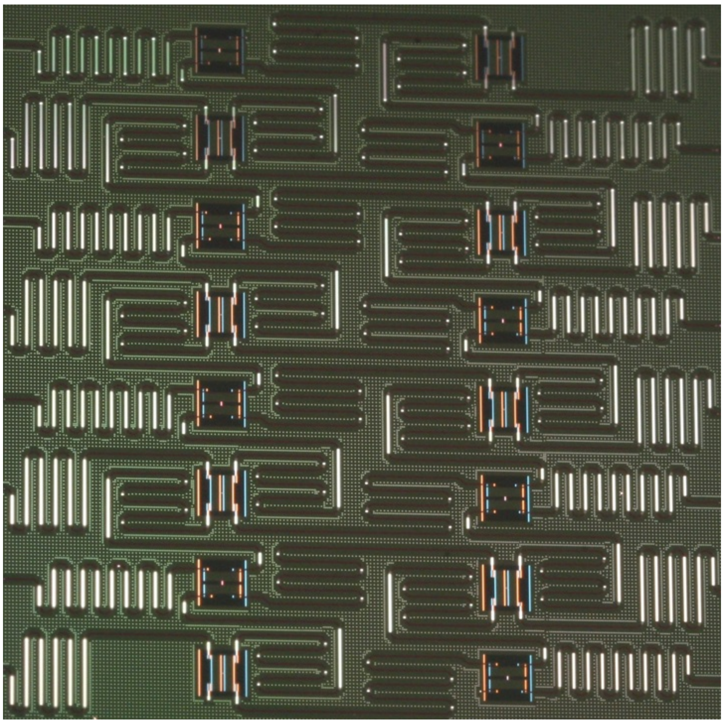

IBM 16-qubit processor (Source: IBM)

Quantum computing is currently enjoying its double-digit qubit era. In 2017, we saw a chip with 20 general-purpose physical qubits, and in 2018 I think we'll see chips with over 50 general-purpose qubits. But the first mass-market general-purpose quantum computers will be built from thousands of logical qubits. Logical qubits are fault-tolerant, with error detection and eventually error correction. Thousands of logical qubits translates into at least tens of thousands of physical qubits, depending on the physical qubit architecture perhaps orders of magnitude more.

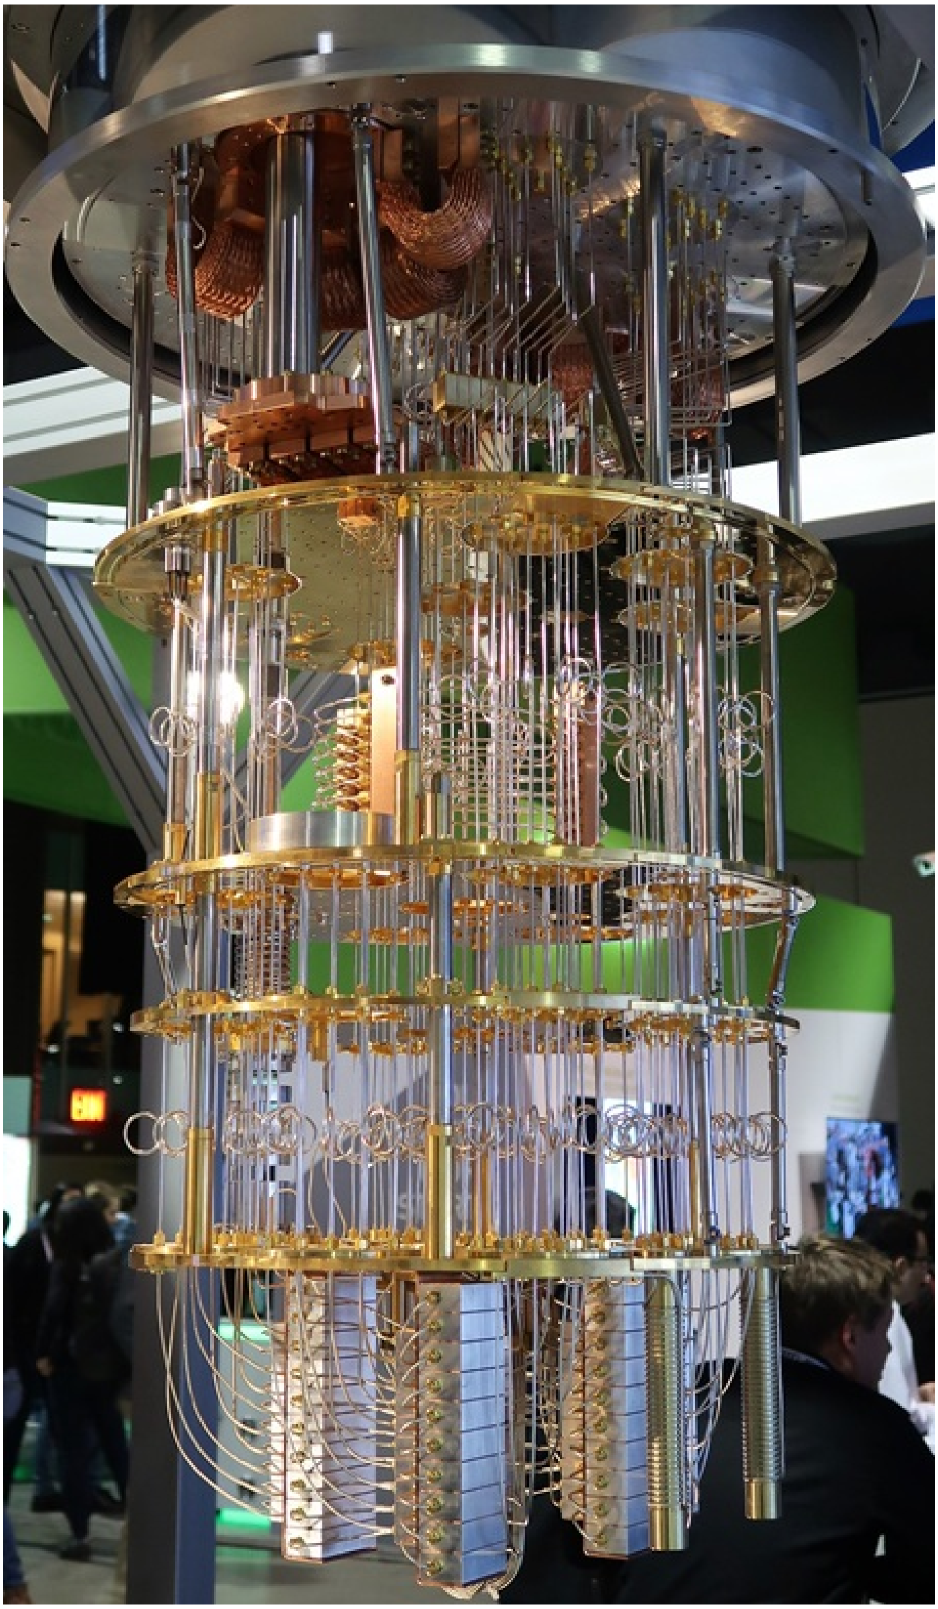

IBM's quantum computer I/O subsystem for getting electrical signals in and out of a millikelvin liquid helium bath (Source: TIRIAS Research)

As it is, evolving from tens to hundreds of physical qubits could take a while; from hundreds to thousands will take longer. Experts agree that a commercially deployable quantum computer with thousands of logical qubits is over a decade away, and possibly over two decades away. "Quantum supremacy" will not be a thing for a long while yet. Meanwhile, many vendors are making compelling progress.

Here's a scorecard for quantum computing in early 2018. It's in a rapid expansion phase, but when the numbers are small, rapid growth is easy.

Under The Hood Of Quantum Systems

IBM and Rigetti both introduced capable general-purpose cloud-based quantum computers for public and limited access use (20- and 19-qubit systems, respectively), each with a full-stack software development toolkit (devkit). NTT introduced its cloud-based quantum dot and photonics-based architecture, with its full-stack devkit. Microsoft and Google introduced their general-purpose quantum computing R&D programs along with full-stack devkits and simulators but have not yet publicly demonstrated hardware. Intel is showing prototype chips, but has not demonstrated them yet. IonQ, Quantum Circuits, and RIKEN are investing in hardware development but have not shown their work yet. And yet, only two companies are selling dedicated systems for customer on premise use that can arguably be called quantum computers: D-Wave's quantum annealing architecture and Atos's dedicated quantum simulator.

D-Wave and NTT implement 2,048 physical qubits, though they use completely different technologies to do so, and their systems do not show full general-purpose quantum computing capabilities. Their architectures are suitable for solving some problems in optimization, molecular dynamics, and even deep learning training and inference tasks.

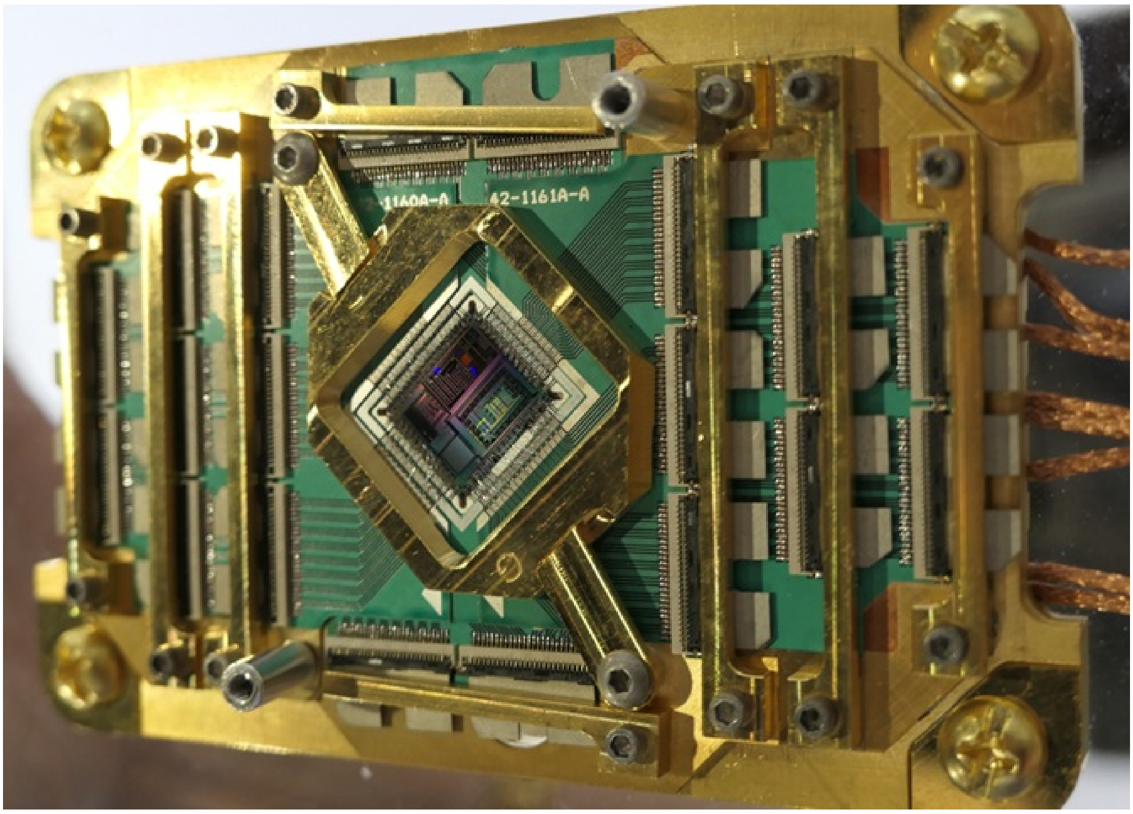

D-Wave quantum computing chip (Source: TIRIAS Research)

Simulating Quantum Iron

Simulating tens of physical qubits takes a lot of "classical" compute power, meaning today's most advanced IC-based compute, memory, storage, and networking architectures. These software simulations may run orders of magnitude slower than the quantum computers they simulate, could researchers actually build real systems as large as they are currently simulating.

Last week, a team of European researchers from Jülich Supercomputing Centre, Wuhan

University, and the University of Groningen successfully simulated a 46-qubit general-purpose quantum computer. This simulation broke the 45-quibit record set in April by the US Department of Energy's Lawrence Berkeley National Laboratory. In July, a US-based team from the Harvard-MIT Center for Ultracold Atoms and the California Institute of Technology simulated a 51-qubit quantum computer, but it was built to solve one specific equation and was not a general-purpose simulation. In November, a team from University of Maryland and the US National Institute of Standards and Technology (NIST) published a paper on a 53-qubit simulator, also intended to solve a specific problem.

Meanwhile, in the cloud, IBM has internally simulated a 56-qubit general-purpose system on a classical supercomputer. However, with 16-qubit systems available publicly, it looks like IBM has deemphasized end-user simulations outside of its new Q Network program. Microsoft's new Quantum Development Kit supports simulations of "over 40 qubits" in its Azure cloud, and its local PC-based simulation can scale to about 30 qubits in 16 GB of memory. I have to wonder if Microsoft's Azure quantum computing simulations are related to its recent partnership with Cray. Rigetti's cloud-based Forrest simulator can simulate up to 36 qubits. Google's Quantum Playground can simulate up to 22 qubits.

General Purpose Chips

IBM introduced its 20-qubit chip in late 2017, it was a cornerstone of IBM's Q Network announcement (read more on the Q Network announcement below). IBM says it has built and internally tested a 50-qubit chip. IBM Q Network participants have access to the new 20-qubit systems and will have early access to the 50-qubit chip as it progresses. Intel delivered a 17-qubit test chip to its partner QuTech in October and showed a 49-qubit chip in early 2018 at the Consumer Electronics Show (CES). Rigetti announced this week that its 19-qubit chip is available for cloud access (access requires approval by Rigetti). Rigetti's chip is a 20-qubit architecture where one of the qubits has a fabrication defect; it is close behind IBM. Google has internally tested six-, nine-, and 20-qubit chips, and is working on a 49-qubit chip that was to ship by the end of 2017, but had not as of this post's publication time.

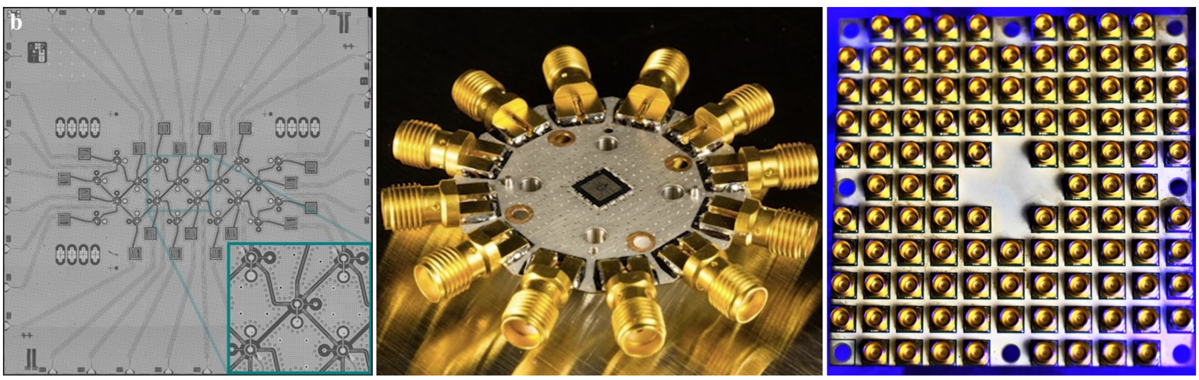

Rigetti 20-qubit silicon (left), Google 6-qubit silicon plus carrier (middle), Intel 49-qubit carrier (right) (Sources: respective manufacturers)

Atos says its 40-qubit simulator is based on Intel Xeon processors but that dedicated hardware accelerators are "coming soon." This is not surprising, as IBM is internally using its Power Systems to simulate quantum computers during development.

Quantum Software Development

On the software front, open sourcing key pieces of devkits is mandatory to attract academic researchers to specific architectures, because these researchers have been open sourcing their in-house quantum computing development environments for the past couple of decades.

This year, IBM open sourced its QASM (Quantum ASseMbler), a key component of IBM's QISKit (Quantum Information Software Kit). XACC (EXtreme scale ACCelerator) interfaces to Rigetti's simulator and prototype chips, and to D-Wave's production systems. QuTiP (Quantum Toolbox in Python) is an open source quantum computing simulator in use across a wide swath of the quantum computing hardware community (logos for Alibaba, Amazon, Google, Honeywell, IBM, Intel, Microsoft, Northrup Grumman, Rigetti, and RIKEN appear on its site). Presumably QuTiP is being used to simulate hardware architectures under development. Google partnered with Rigetti to open source OpenFermion, a software package for compiling and analyzing quantum chemistry problems. And Microsoft introduced its Q# ("Q-sharp") quantum computing language (read more about Microsoft below). There was much more activity here, but too much to convey in this post.

What About China?

Chinese companies have been conspicuous by their absence from quantum computing press and announcements. This year China announced a $10 billion National Laboratory for Quantum Information Sciences, it is scheduled to open in 2020. Alibaba, Baidu, and Tencent are all heavily invested in AI and deep learning, so I would expect to hear more about their interest in quantum computing over the next year.

Recent Big Announcements

Microsoft Announces Quantum Development Kit

Microsoft started working on quantum computing almost two decades ago, in 2000. In September, at its Ignite event, Microsoft announced it is basing its quantum computing program on Majorana Fermions, which were found in 2012. If Microsoft can harness Majorana Fermions, scaling logical qubit count may be much more economical than alternative qubit technologies, with only 10 or so physical qubits to one logical qubit, instead of thousands or more.

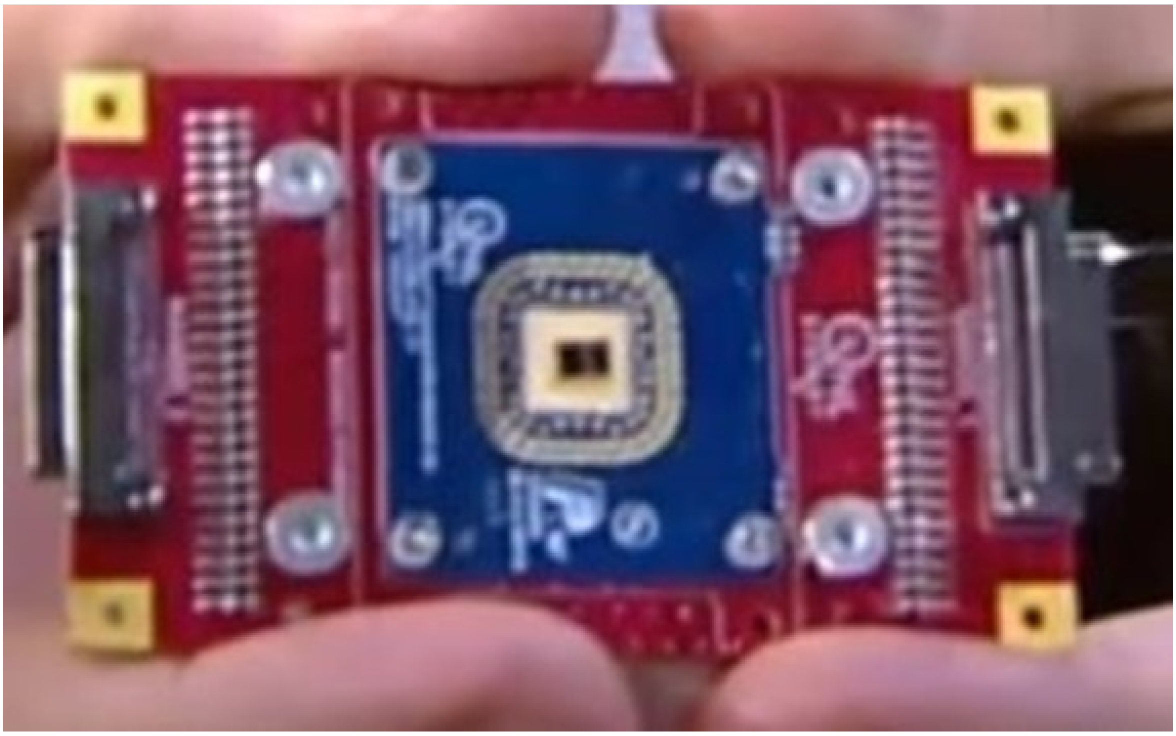

Microsoft's prototype Majorana Fermion quantum computing chip (Source: TIRIAS Research)

But, after its big Majorana Fermion unveil, Microsoft has been silent on its hardware progress. Instead, Microsoft is focusing on quantum simulation its new Q# language, tightly integrated into its Visual Studio Integrated Development Environment (IDE), and quantum computer simulation tools, including a trace simulator to analyze resource utilization, plus lots of libraries, code samples, and comprehensive documentation.

Microsoft's quantum simulator uses Intel's AVX extensions, supported in Intel processors since its 2011 "Sandy Bridge" processor generation. Microsoft is also announced its project "Brainwave" FPGA-based AI accelerator this year and has hinted that it is running "quantum inspired optimizations" on Brainwave. My guess is that Microsoft is optimizing Brainwave's FPGA deep neural network (DNN) logic to improve deep learning model accuracy, speed, or both.

Microsoft's announcement is important because of the sheer number of enterprise software developers using its Visual Studio IDE. It is a mature, highly productive toolkit. Integrating quantum computing into Visual Studio may attract a new generation of academic researchers away from open source IDEs, much as Nvidia has been able to do with GPU programming through its CUDA application programming interface (API) and toolkit.

IBM Announces Q Network

IBM had already launched its QISKit API and devkit for developers to access IBM's cloud-based Quantum Experience and local simulators. In December, IBM launched its Q Network ecosystem development initiative. IBM hasn't talked about membership levels, but membership seems to scale based on both ability to pay and likelihood of potential contributions to IBM's quantum ecosystem.

Access to IBM's quantum computing resources is simple; general access will always be a generation or two behind Q Network participants for both access to hardware and access to the latest devkit resources. There are three types of memberships, the announced participants are:

- Hubs (regional centers of education, research, development, and commercialization): Keio University, University of Melbourne, Oak Ridge National Laboratory (ORNL), University of Oxford, and IBM Research
- Partners (pioneers in a specific industry or academic field): Daimler, JPMorgan Chase & Co, JSR, and Samsung
- Members (developing a strategy to get quantum ready): Barclays, Honda, Materials Magic (Hitachi Metals Group), and Nagase

Target users in IBM's Q Network and in the larger IBM Q experience user base are graduate students, academic researchers, and commercial researchers. Quantum computing is still at an experimental stage, both in providing infrastructure and in understanding how to program quantum computers to solve useful problems. Quantum computing is today dominated by discovery and heuristics.

IBM says its Q Experience tools are used by over 1,500 universities, 300 private education institutions, and 300 high schools as part of their physics curriculum. This, also, is like Nvidia's successful strategy of educational outreach with CUDA tools. IBM claims that 35 third-party research publications used Q Experience tools, an impressive number that highlights the intense competition for researcher mindshare underway in this early phase of quantum computing.

Where To From Here?

We still have a long road ahead of us to commercialize quantum computing. There will be some temporary advantages along the way. But with the amount of investment pouring into quantum computing R&D, short-term quantum advantage by any one competitor will be fleeting without sustained, long-term R&D and commercialization strategies.

We will be surprised if we don't see systems with 50 or more general-purpose qubits in 2018. We think we will also see some of the more specialized systems that already have over 2,000 physical qubits show significant quantum advantage within constrained problem domains. We are planning to attend The International Conference on Quantum Communication, Measurement and Computing (QCMC) in March to stay up-to-date on the research side of quantum computing.

Paul Teich is an incorrigible technologist and a principal analyst at TIRIAS Research, covering clouds, data analysis, the Internet of Things and at-scale user experience. He is also a contributor to Forbes/Tech. Teich was previously CTO and senior analyst for Moor Insights & Strategy. For three decade, Teich immersed himself in IT design, development and marketing, including two decades at AMD in product marketing and management roles, finishing as a Marketing Fellow. Paul holds 12 US patents and earned a BSCS from Texas A&M and an MS in Technology Commercialization from the University of Texas McCombs School.