

Comments on IEEE Framework for Metrics and Benchmarks of Quantum Computing

Mark B. Ritter, Chair, Council of Physical Sciences, IBM T.J. Watson Research Center

General Comments:

It is very early, in my view too early, to develop metrics and benchmarks for quantum computing technologies. This field is very young, and nobody has yet built any quantum computing system that can outperform a laptop computer on any task. Every aspect of these young quantum systems is changing, from the qubit itself to how to best control the qubits to the environment of the qubits all the way up the stack to the quantum-specific languages. In short, we do not yet know how to make a quantum computer which demonstrates the holy grail: quantum advantage (performing better, in some dimension, than conventional computing).

There is another potential concern that I have-- that of allowing traditional ideas of computing force themselves into terminology and “layers” of the quantum system stack which do not properly describe quantum systems. Further, it is hard, and perhaps even misleading, to define and compare technology “layers” across different quantum computational models. As an example: what is called a qubit for a quantum annealer is so different from other models of quantum computation that an “annealer qubit” would be useless in a circuit or digital quantum computer. Therefore, they should not be compared and maybe should be called by a different name to avoid confusion. I would also recommend, if IEEE publishes any documents, that any table of technologies be made only for systems that can be compared across the dimensions considered in the table. For example, all quantum computation systems could be compared when performing an Ising model simulation (of suitable size and connectivity), but annealers could not perform a Heisenberg model or other more complex Hamiltonian simulations.

One final comment that I will make emphasizes the nascent technologies that comprise the field of quantum computing. In traditional computing, a “separation of concerns” developed allowing clean definitions of memory, CPU, and I/O interfaces and functions that have served computer scientists very well over the decades, allowing separate development of memory, CPU, and I/O. In present-day NISQ quantum systems, there is no separation between memory and CPU, and the “circuits” or interactions among qubits are orchestrated by a conventional computer that sends electromagnetic pulses (laser or RF) of the proper amplitude and phase and in the right sequence! Where do you define the separation of concerns here? Further, in the future, to add error correction codes to noisy qubits to create a fault-tolerant machine, the state and manipulations of the Logical qubit (constructed from many physical qubits by sequences of electromagnetic pulses) will reside totally in a conventional computer (at least as now envisioned). The logical code and algorithms at a logical level will be translated into a complex orchestral score of

electromagnetic pulses to individual physical qubits. On the other hand, there are researchers seeking qubits protected by physics (topological qubits); although they have not yet been demonstrated, if they can be realized, then a fault-tolerant quantum computer would look much different.

Details: Table 1: Quantum Computational Models

Because of the lack of consensus in terms, I believe that Table 1 is not complete and mixes up different types of quantum computational models. Here is my proposal for a more correct table; however, the very fact I must make this proposal is another indication of the nascent state of quantum computing and the caution that the IEEE and other organizations should not move too quickly when there does not yet seem to be a scientific consensus. Even this proposed table and definitions should be discussed by quantum experts.

Universal Quantum Computing	Technology which implements a universal set of quantum logic gates on qubits.
Circuit Quantum Simulator	Technology which implements gates (potentially universal) that can entangle qubits and flexibly model Hamiltonians.
Adiabatic Quantum Computer	Technology which can implement a Hamiltonian that can be adiabatically evolved from an initial Hamiltonian of known solution to a final Hamiltonian whose solution is desired. If adiabatically evolved, there will be no tunneling.
Analog Quantum Simulator	Technology that employs qubits with fixed connections intended to mimic a particular Hamiltonian.
Quantum Assisted Annealer	Technology that employs qubits with a fixed connection topology intended to solve Ising Hamiltonian annealing problems assisted by quantum effects (tunneling).

Details: Table 2: Families of Systems:

I also feel Table 2 descriptions should be clarified; my suggestions are below.

Fault-tolerant Quantum Error Corrected Systems (FTQS?)	Systems that suppress errors with codes to detect/correct qubit errors or with new qubit technologies.
Noisy Intermediate Scale Quantum Systems (NISQ)	A system comprised of few to hundreds of qubits without full error correction- uncorrected errors affect the calculation.
Quantum Simulators	Systems that use discrete (circuit) or analog (continuous-time) qubit interactions to solve the static or dynamic properties of a Hamiltonian expressing a physical system or mathematical problem.
Quantum Annealers	Systems employing quantum effects (like tunneling) to assist in finding the lowest energy state of a Hamiltonian that describes a particular mathematical or physical model. (Present systems only solve the Ising model.)

Technology Layers

The IEEE should appreciate that not all of the layers of conventional computers are applicable to quantum technology. Further, the layer naming proposal given in this table is inscrutable to people who are experts in the field (see, for example, the layers in Van Meter 2013). If anything, the IEEE should solicit input from experts in different quantum computing technology types to begin to achieve a consensus on what the layers should be and how to name them. Since nobody has yet made a fault-tolerant quantum computer or an adiabatic quantum computer, layer naming cannot even begin! I would contend that we have ideas of how to create these systems, but there is too much research to be done to reliably predict what the layers will finally be.

Now for feedback on the technology layers in Table 3. In general, it is my current view that these layers will differ from one implementation of quantum computing technology to another. I think it is very good to give this example table to make people start thinking of how to “abstract” layers that are invariant among various quantum systems and still make sense, but I think that is work best done after we have actually demonstrated systems which show advantage and not at this stage which could be described as a “research prototype” phase.

Examples of areas that I consider missing from the table: cryogenic system (cools quantum device(s)), cryogenic packaging, cryogenic amplification chain (for qubit readout), RF subsystem, optical subsystem (for ions, NV centers, etc), vacuum subsystem (for ions), control electronics, readout electronics, ion trap subsystem (ions), atom trap subsystem (neutral atoms). (Note on the last two: ion traps use RF electrical trapping fields, neutral atoms use optical fields. The technologies are so different that they have very different requirements.)

Examples of areas that I consider incorrect:

- Physical register, physical device → replace both with “quantum device.” A collection of qubits holds the current quantum state which is manipulated by electromagnetic pulses.
- Physical circuit → This is very ambiguous. In some systems (an annealer, for example) the physical connections (circuits?) are fixed with programmable strength. In other systems, electromagnetic pulses (laser or RF) control which qubits talk to each other, but there is no fixed connection or circuit.
- Logical Register → This is an abstract concept in quantum computing; a set of qubits are controlled to create a logical qubit, and the information about the logical qubit is contained in a classical computer memory. I am not sure what should be done with this term.
- Logical Device → Again, this would simply be a classical computer or state machine supervising the coding to create a logical qubit. Please do not use “device”- there is no physical meaning to saying “a logical qubit device” – it is a construct created from a group of physical qubits being controlled, measured, and possibly corrected by a conventional computing system.
- Integrated Device → Same as above. There is a “logical layer” but not “logical device.” I would not use “Integrated Device” at all in a layered description.
- System Architecture → I think most people knowledgeable in the field (and there aren’t many of us who have full quantum system stacks!) would say that the system architecture of a fault-tolerant quantum computer will end up being a distributed computing system which will have a conventional system controlling coding to create logical qubits and another conventional computing system above that to manipulate the logical qubits. The system manipulating logical qubits could look very similar (in language and operations) to full stacks that exist today, like IBM’s qiskit. Nobody has yet made such a machine; invariably, engineering tradeoffs will be learned and the final architecture evolved from research and engineering work that has yet to be done.

Benchmarks and Metrics

Finally, the field of quantum computing is beginning to create metrics relevant to a particular system’s performance. There is much research (and entire sessions at conferences) devoted to benchmarking quantum annealer performance measures—but no consensus. For NISQ machines there are proposed single and two-qubit gate fidelity measurements, including Quantum Process Tomography (QPT), Gate Set Tomography (GST), Randomized Benchmarking, and Interleaved Randomized Benchmarking.

Though one of the more developed benchmarks for qubit gate performance, GST is a field of research which is still in flux and being debated at physics conferences [Blume-Kohout 2017]. If I can explain just one example of why this is so, consider that the GST benchmark requires a number of measurements which grows exponentially with the number of qubits being measured, and so can be used only for a handful of qubits. Further, the measurements

themselves can take longer than typical heuristic NISQ algorithms. This is important because present-day NISQ technology is not yet stable in time—the coherence time of the qubits can fluctuate, calibrations of equipment controlling the qubits can drift. Many things are being debated—but one of them is how to measure quantum device performance over a time period that is meaningful, and how to disentangle system changes from device changes.

It is my view that the IEEE could begin working toward metrics and benchmarks for individual classical components that are required in quantum systems. For example, for atom-based systems, lasers of a particular wavelength and stability are required to accurately control the ions or atoms. I believe the frequencies and frequency stability required for some operations is known and could be defined by metrics and benchmarks. Similarly, quantum limited amplifiers and cryogenic RF components like attenuators, isolators and circulators could have defined performance metrics and benchmarks. Right now, vendors do not realize that some attenuators use resistive elements that become superconducting and are useless in a cryogenic environment. It is up to the researcher to “buy and try” components which were not designed for or tested in the cryogenic environment. Further, the resistance of all materials changes over temperature—what temperature range(s) and what variability with temperature would be adequate for cryogenic applications? Though these examples may seem trivial, I believe they exemplify “low-hanging fruit” that the IEEE or other organizations could profitably address to the benefit of the entire field and component suppliers.

Summary

I commend the IEEE for opening discussions about metrics and benchmarks for quantum computing. Just starting the discussions was a necessary and good step, but the quantum systems and technologies have not yet matured to the point to have stable definitions of systems, let alone the metrics and benchmarks. I believe that it is profitable at this time to begin discussing and defining metrics and benchmarks for traditional technologies (lasers, RF devices) that have different requirements (stability, temperature range) when used in quantum systems. I will end with a question posed in Van Meter: “When will a quantum computer do science, rather than be science?” The question remains unanswered.

I would be glad to give further clarifying input, if the IEEE requests.

References:

Blume-Kohout 2017:

“Demonstration of qubit operations below a rigorous fault tolerance threshold with gate set tomography,” Blume-Kohout, Robin and Gamble, John King and Nielsen, Erik and Rudinger, Kenneth and Mizrahi, Jonathan and Fortier, Kevin and Maunz, Peter, *Nature Communications* **8**, pg. 14485 (2017).

Van Meter 2013:

“A Blueprint for Building a Quantum Computer,” Rodney Van Meter, Clare Horsman, *Communications of the ACM*, **56**, pp. 84-93 (2013).