

Comments on “An IEEE Framework for Metrics and Benchmarks of Quantum Computing”  
Eleanor Rieffel  
Feb 8, 2019

My first reaction is that it generally seems too early to be coming up with benchmarks and metrics for quantum computing, at least for most aspects of quantum computing, because the field is at such an early stage of research. At the hardware level, there are many different approaches (superconducting, ion trap, anyon, optical, etc.), with different tradeoffs. All approaches will need breakthroughs before even having a design for a large-scale quantum computer that is fully scalable. At the algorithmic end, we only have a few dozen algorithms with a clear advantage over classical algorithms, and even fewer when one restricts to algorithms having a practical use. As small-scale quantum computing hardware becomes available, and experimentation well beyond the range at which classical supercomputers can simulate quantum computers, we can begin to evaluate quantum heuristics which is expected to broaden the class of quantum algorithms substantially (most classical algorithms run on supercomputers are heuristics, i.e. they have not been mathematically proven superior to previous algorithms, but rather have been evaluated numerically), but we do not know yet which algorithms will be effective, and for what applications. The uncertainties at all levels are so great, and so much still research needs to be done, not only would setting standard benchmarking and metrics be difficult, but it could also be detrimental. I am particularly worried that novel and useful ideas, in directions no one has imagined, could be suppressed (e.g. denied funding, rejected for publication) simply because they do not meet IEEE standards.

The main point I want to get across is that great care must be taken that this effort does not in any way impede good research.

The opening paragraph mentions using the benchmarks and metrics to “track technical progress,” to “to guide the decisions of policy makers,” and to “monitor the overall growth.” In many cases it is too early to do any of these things. When one has a clear destination, one can measure progress toward that destination. When the destination is unclear, one should be wary of guides that are too definitive.

One area where metrics and benchmarks could be useful is in classifying different types of quantum advantages, and standardizing the meaning and terminology. Some work has already been done in this direction (e.g. S. Mandra et al. [arxiv:1604.01746](#), Rønnow et al. [arXiv:1401.2910](#)).

The call for open partnership and collaboration in the second paragraph of the proposed IEEE framework highlights an issue of critical importance at this time. Up to this point, the community has been quite open, sharing most knowledge freely. With the increasing commercial presence with pressures to make money in the near-term, international rivalries, and the current level of hype and competition, there is a danger of losing the openness that has provided such fertile ground for research of the past couple decades. If this IEEE framework can

help maintain an open ethos, that will be a valuable contribution to the future of quantum computing research.

The second section of the document, “Use Cases,” is confusingly written. The concepts here should be untangled and made clear. Further, the framework currently contains a very odd description of what a “use case” is; we should not be equating “distinct designs [for quantum computers] as uses cases.” Please use other terminology here (e.g. “quantum computational designs” or “quantum computational models”). Use cases are generally specific applications, with all sorts of practical parameters filled in. Here, it appears to be referring to broad classes of models for quantum computation. The first two are universal quantum computers, so can be applied in any use case where quantum computing turns out to be relevant. The third narrows to simulation, but that is still quite broad, with potential applications to a wide variety of use cases.

I agree that it is worth distinguishing universal and special purpose quantum hardware, and that there are strong reasons to pursue each. (I trust that the readers and writers of this document are fully aware that there are many other models of quantum computation, so that Table 1 is just providing a few such examples.) I appreciate highlighting that different benchmarks and metrics will be needed for different models, and different uses/applications.

We should be careful to make sure that specifying the layers of technology exclude reasonable approaches to quantum computing. As one example, the definition for physical register is that it stores qubits. But we should not rule out continuous variable, qutrit, or higher qudit approaches.

The framework currently contains no examples of Benchmarks or Metrics. I understand that the what is proposed is a framework, and that the purpose was not to propose benchmarks or metrics yet, but rather provide a framework for doing so. Nevertheless, naively one might have expected to see some examples, preferably some uncontroversial ones. The early research stage in which the field currently sits means that finding solid and uncontroversial benchmarks and metrics is not trivial.

Let’s approach this effort humbly and carefully, recognizing that for many aspects of quantum computing we may be years, even decades, away from a point at which we can reasonably, and without impeding research, establish benchmarks and metrics.