QUANTUM COMPUTING

Quantum advantage deferred

A type of optics experiment called a boson sampler could be among the easiest routes to demonstrating the power of quantum computers. But recent work shows that super-classical boson sampling may be a long way off.

Andrew M. Childs

Quantum computers have the potential to dramatically outperform classical computers at solving certain problems. However, despite impressive experimental progress, practical quantum computers remain a long-term goal. A simpler objective is to build a quantum device that merely outperforms classical computers, if only for a toy problem. Boson sampling is perhaps the best-known framework for establishing such a result, yet it remains unclear at what size it becomes classically intractable. Previous estimates have placed this threshold around 30 bosons. Writing in *Nature Physics*, Alex Neville and co-authors[1] give strong evidence that the threshold for classically hard boson sampling is at least this high — and that boson-sampling experiments will need to exceed approximately 50 photons to push quantum computation beyond the reach of our present classical abilities.

In boson sampling, single photons are prepared, sent through a linear optical network consisting of beam splitters and phase shifters, and finally measured with detectors that can resolve individual photons. Although such a linear optics experiment cannot perform universal quantum computation — it cannot run Shor's celebrated factoring algorithm, for example — it has been argued that sampling from its output distribution (even approximately) is hard for classical computers[2]. Since realizing linear optics is easier than performing a general quantum computation, boson sampling offers a gentler path to demonstrating a quantum computational advantage.

So far, experiments have performed boson sampling with up to five photons[3]. Extending them to more photons is challenging for many reasons, not least of which is the difficulty of reliably creating several single-photon states that arrive at the detectors simultaneously. This raises a natural question: just how large an instance of boson sampling must one consider such that a classical computer cannot solve it in a reasonable amount of time?

Unfortunately, even the asymptotic hardness of boson sampling is unclear. Arguments for its intractability rely on unproven conjectures in both linear algebra (about the distribution of permanents of matrices with Gaussian-distributed entries) and computational complexity theory. But for concrete instances of boson sampling, we are interested in a more mundane — yet slipperier — question: how long would the best possible classical algorithm take to sample from the output distribution? This is challenging to answer, as there are many possible approaches to simulating boson sampling with a classical computer.

Considering the difficulty of scaling up single-photon experiments, the results of Neville *et al*. suggest that it could be quite some time before boson sampling is used to demonstrate a definitive quantum advantage.

The algorithm presented by Neville *et al.* is based on a technique called Metropolised independence sampling, which apparently reproduces the boson sampling distribution much more efficiently than brute-force calculation of the full distribution. Roughly speaking, the algorithm runs a Markov chain whose stationary distribution corresponds to boson sampling. Random steps in the Markov chain are first proposed by drawing from the distribution corresponding to classical distinguishable particles — which can be sampled efficiently — and then accepted with a probability that depends on the chance of making such a transition with indistinguishable bosons. Empirically, this Markov chain appears to converge to the boson-sampling distribution after a number of steps that grows slowly with the number of bosons, suggesting that the overall cost should be dominated by the cost of determining whether to accept each step. While that is a computationally demanding task, it can be done with an algorithm of Ryser's[4] that takes time growing only exponentially with the number of bosons. Using this approach, the authors generate samples from the 20-photon distribution using a laptop and from the 30-photon distribution using a cluster of four off-the-shelf computers. Furthermore, they estimate that sampling from the 50-photon distribution should be within reach of a supercomputer.

One shortcoming of the sampling algorithm of Neville *et al.* is that it does not come with a guarantee of correctness. The Markov chain should be run long enough that it provides a good approximation of the boson-sampling distribution. However, it is challenging to determine whether this is the case, precisely because it is hard to sample from the desired distribution. The authors address this issue in several ways. They show that the output of their Metropolised independence-sampling procedure agrees with a brute-force calculation for very small sizes, and they find that it agrees with a different, less-efficient Markov chain method for somewhat larger sizes. They also show that the distribution they sample from is far from the distribution of distinguishable particles — a necessary, but far from sufficient, condition. But because boson-sampling experiments also suffer from the same challenge of verifying correctness of the output distribution, it seems reasonable to accept this limitation.

Considering the difficulty of scaling up single-photon experiments, the results of Neville *et al.* suggest that it could be quite some time before boson sampling is used to demonstrate a definitive quantum advantage. However, it is possible that another way of asserting quantum computational advantage will be even easier to realize. Numerous alternative sampling problems have been proposed, including the so-called instantaneous

quantum polynomial-time model, sampling the output of random quantum circuits, and analogues of boson sampling in spin systems[5]. Other tasks such as quantum simulation[6] and approximate optimization[7] are also promising candidates for early implementation. Progress towards building quantum processors continues apace, so demonstrations of super-classical computation may be viable in the relatively near term. However it is eventually achieved, quantum computation that exceeds the reach of classical computers will be an exciting development that begins a new era of quantum science. ❐

Andrew M. Childs is in the Department of Computer Science, at the Institute for Advanced Computer Studies, and at the Joint Center for Quantum Information and Computer Science, University of Maryland, Maryland 20742, USA.
e-mail: amchilds@umd.edu

References
1. Neville, A. *et al. Nat. Phys.* http://dx.doi.org/10.1038/nphys4270 (2017).
2. Aaronson, S. & Arkhipov, A. *Theory Comput.* **9,** 143–252 (2013).
3. Wang, H. *et al. Nat. Photon.* **11,** 361–365 (2017).
4. Ryser, H. J. *Combinatorial Mathematics* (MAA, 1963).
5. Lund, A. P., Bremner, M. J. & Ralph, T. C. *npj Quant. Inf.* **3,** 15 (2017).
6. Feynman, R. P. *Int. J. Theor. Phys.* **21,** 467–488 (1982).
7. Farhi, E., Goldstone, J. & Gutmann, S. Preprint at http://arxiv.org/abs/1411.4028 (2014).

Published online: 2 October 2017