

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-authors¹ 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 computers². 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 photons³. 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's⁴ 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 systems⁵. Other tasks such as quantum simulation⁶ and approximate optimization⁷ 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.* **13**, 1153–1157 (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

SELF-ASSEMBLY

Misfits unite

The spontaneous assembly of particulate or molecular ‘building blocks’ into larger architectures underlies structure formation in many biological and synthetic materials. Shape frustration of ill-fitting blocks holds a surprising key to more regular assemblies.

Gregory M. Grason

In the canonical picture, self-assembly results in an arrangement where identical building blocks fit together, like a jigsaw puzzle, into a complementary pattern stabilized by non-covalent attractions between neighbours (Fig. 1a). The uniform distribution of cohesive energy in the interior, together with the drive to minimize cohesive loss at the surface, leads to a single morphology in the condensed state: the macroscopic aggregate. But a minimal model, reported in *Nature Physics* by Martin Lenz and Thomas Witten¹, shows that assembly of ill-fitting, or geometrically frustrated, elements breaks this basic paradigm — yielding multiple distinct morphologies, including assemblies with finite and well-defined multi-element dimensions.

Geometric frustration refers to the incompatibility of local interactions with global geometric constraints², making the propagation of uniform and strain-free order impossible. Although it's most commonly associated with bulk materials³ such as antiferromagnets and hard-sphere glasses, geometric frustration has important and lesser-known implications for self-assembling structures whose dimensions are variable and possibly finite.

As an illustration, consider the assembly of the misshapen puzzle pieces shown in Fig. 1b. The warp of the piece does not prevent pairwise fit of any neighbours, which we take as a proxy for cohesive interactions. But in this example the tapered shape favours bent 1D rows of the puzzle that lead to geometric conflicts when trying to extend the

structure in 2D. It is intuitive, for example, to imagine that one will have to deform the pieces in the 2×2 unit shown in Fig. 1b — not unlike a determined puzzle enthusiast forcing the wrong pieces together. In the context of the assembly problem, the energetic cost of deformation required to extend the cohesive fit over an arbitrary size distinguishes geometrically frustrated assembly from the canonical process.

The resulting structure, thermodynamics and kinetics of geometrically frustrated assemblies are far less understood than the standard unfrustrated cases⁴. To date, geometric frustration has been implicated as an a posteriori mechanism to explain structures found in a diverse range of systems, including phase-separated lipid vesicles⁵, particle coated droplets⁶, viral capsids⁷, chiral membranes^{8,9} and protein filament bundles¹⁰. A primary motivation for Lenz and Witten's study is the assembly of protein building blocks, whose complex and sometimes irregular tertiary folded structures often require some measure of shape strain to form multi-protein architectures.

Lenz and Witten simulated the cohesive 2D assembly of flexible polygons whose shapes cannot tile the plane: pentagons, heptagons and irregular hexagons. Their model considered the sequential addition of polygonal blocks to an existing cluster in order to maximize the net energetic of gain of assembly: the increased cohesion, which is proportional to the number of newly contacting edges, minus the elastic cost to deform block edges into contact.

In their model, unlike the unfrustrated case, in which block shapes tile space without deformation, incompatible blocks exhibited three distinct morphologies (Fig. 1c). In the limiting case of rigid blocks, they found that assemblies were tree- or branched polymer-like with a large fraction of unbound edges. In the opposite extreme of very floppy blocks, they found that cohesive gain was sufficient to deform blocks into a compatible shape that assembles without bound in any dimension, albeit with a slightly reduced energetic gain per block. At the boundary between these two regimes, where deformation cost roughly balances cohesive gain, they found a surprising intermediate state of fibre-like assemblies, where blocks assemble into fairly regular finite-width domains, and continue to assemble unabated in a quasi-1D fashion by growth along the short dimension.

What is so remarkable about this fibre formation is that the assembly proceeds to a particular finite width that is larger than a single block (in the range ~4–8 depending on the conditions) — and yet it is less than infinite. It is as if the process has a mechanism to ‘measure’ the finite width of the assembly, despite the lack of any long-range interactions in the model.

What, then, accounts for the size selection of the assembly? This effect seems to derive from the dual roles played by the free boundaries of the assembly. As in canonical assemblies, boundaries lead to a deficit of cohesive contacts. But for frustrated assemblies, they also lead