Quantum computing

A fast quantum route to random numbers

Mohan Sarovar

Using a quantum computer to speed up one step in a textbook approach to generating random numbers proves to be a savvy strategy, and one that could make good use of quantum computers that will be available in the near future. **See p.282**

Generating a random number might seem like an easy task, but it can be surprisingly difficult – especially if the probability distribution from which the number is drawn is complex. This is often the case in scientific research, for example when training a neural network. In such cases, researchers can use a technique that was one of the first uses of general-purpose computing: the Metropolis algorithm. This was first run on the groundbreaking MANIAC computer in 1953 (refs 1, 2). Its modern generalization is known as the Markov chain Monte Carlo (MCMC) algorithm. On page 282, Layden *et al.*[3] report an even more modern twist in this algorithmic tale, by using a quantum computer to accelerate the program's performance.

The MCMC algorithm is a framework for generating random numbers according to specified probability distributions, a task known as sampling. The framework encompasses several variations, all of which involve iterating through samples that, after enough cycles, are guaranteed to be distributed according to the desired target distribution. Each iteration in this process has two components: a proposal step, in which a sample is suggested on the basis of the current sample; and an accept-or-reject step, in which the new sample is either accepted as the next sample in the iteration, or rejected in favour of repeating the process (Fig. 1a).

Variants of the MCMC algorithm are distinguished by the different strategies used for each of these steps. Crucially, both steps must be constructed in such a way as to guarantee that repeating them eventually results in samples that are distributed according to the desired distribution. Exactly how long this takes is a key property of any MCMC variant. Does the process need to be repeated 1,000 times before the samples are distributed according to the target distribution? Or one million times?

The number of iterations required is known as the convergence time, and it is dependent on the dimension of the random variable – the number of bits needed to describe the sampled variable. The larger the dimension, the longer the convergence time. Unfortunately, for most of the MCMC variants used today, the exact mathematical dependence of convergence time on the variable dimension is not known rigorously[4]. However, this has not stopped people from using the MCMC algorithm. Rather than insisting on a rigorous understanding of convergence time, users tend to fall back on empirical and statistical characterizations of convergence.

Layden *et al.* devised an MCMC variant that uses a quantum computer to produce a sample in the proposal step (Fig. 1b). In any iteration, the random sample is encoded as a quantum state, and a series of quantum operations are applied to it to produce an output state that can be measured to generate the new sample. This in itself is not particularly impressive – almost any procedure could be used to generate a new sample in the proposal step of an MCMC algorithm, including simply applying noise to the current sample. However, to have confidence in the process, researchers must be able to prove that the steps converge to the target distribution, which is not possible for arbitrary proposal procedures. This brings us to Layden and colleagues' key innovation: they designed a set of quantum operations that allow convergence to be verified when the quantum proposal step is coupled with a standard accept-or-reject step.

The authors demonstrated their quantum-enhanced MCMC algorithm through a combination of numerical simulations and experiments on early-stage quantum-computing hardware. Their findings show the predicted convergence of the iterations to the target distribution. More importantly, they also demonstrate that the convergence is faster than several classical alternatives that have previously been devised for the proposal step.

The actual rate of convergence is difficult to measure, and the authors managed to do so only for processes of limited complexity: those with target distributions over variables that can be described by up to ten bits. They also approximated the convergence rate for target distributions over 20-bit variables. In all cases, they found convincing evidence that their quantum version of the MCMC algorithm converged faster than did its classical counterpart. They established empirically that this speed-up is polynomial, with the convergence time for the quantum-enhanced strategy being about the cube root of the convergence time for the conventional strategies.

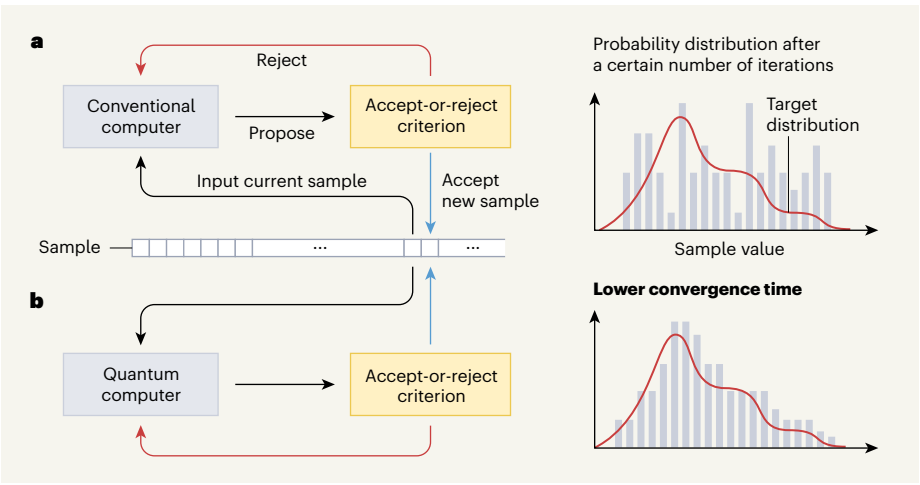

Figure 1 | Accelerating random number generation with a quantum computer. a, Markov chain Monte Carlo (MCMC) algorithms for sampling random numbers from a target probability distribution involve two steps: proposal of a sample and then either its acceptance as the next sample in the iteration, or rejection, which triggers the process to begin again. Both steps must be designed so that sufficient iterations will eventually result in samples that are distributed according to the target distribution, and the number of iterations required to achieve this is known as the convergence time. **b,** Layden *et al.*[3] devised a variation of this algorithm that uses a quantum computer to propose a sample, and offers a marked improvement in speed over existing classical algorithms. Specifically, the samples converge to the target distribution in fewer iterations than they do using conventional MCMC algorithms.

Where does this speed-up come from? As with most quantum enhancements, it is difficult to ascribe it to any one feature of the quantum system. Layden *et al.* offer numerical evidence that their chosen quantum operations strike a delicate balance between generating proposals that are diverse with ones that satisfy the constraints imposed by the target probability distribution — a trade-off that classical proposal strategies struggle to achieve.

Although Layden and colleagues' work is comprehensive, there are some limitations. First, the proof of convergence of the quantum-enhanced algorithm is valid only if the quantum operations are executed perfectly — in the absence of any noise arising from the hardware. However, their experimental results suggest that the rate of convergence is somewhat robust to noise, especially if the hardware noise can be randomized. Second, the accelerated convergence was observed only for small-scale problems, and could disappear at larger scales, especially in the presence of noise. If the authors' explanation for the reason for the speed-up is valid, and if hardware noise can be suppressed at larger scales, it seems likely that the speed-up would persist, but this is far from certain at this stage.

Finally, although Layden *et al.* have demonstrated that their quantum-enhanced algorithm shows faster convergence than do some common classical proposal strategies, there are many MCMC variants that they haven't tested. It is therefore possible that this gap could be closed by other classical proposal strategies that exist or could be devised — perhaps even some that are inspired by this work. Despite these limitations, Layden and colleagues' research forges an important and exciting application of early-stage, noisy quantum computers to generate useful solutions and, in doing so, it defines many directions for fruitful future research.

Mohan Sarovar is in the Quantum Information Science Department, Sandia National Laboratories, Livermore, California 94550, USA.
e-mail: mnsarov@sandia.gov

1. Metropolis, N. *et al. J. Chem. Phys.* **21**, 1087–1092 (1953).
2. Gubernatis, J. E. *Phys. Plasmas* **12**, 057303 (2005).
3. Layden, D. *et al. Nature* **619**, 282–287 (2023).
4. Cowles, M. K. & Carlin, B. P. *J. Am. Stat. Assoc.* **91**, 883–904 (1996).

The author declares no competing interests.

Cell division

A lack of commitment to proliferation

Alexis R. Barr

It turns out that commitment to cell division is not an irreversible switch. In the absence of sustained stimulation by growth factor proteins during DNA replication, cells can quit the cell cycle before cell division occurs. **See p.363**

When cells proliferate, they commit to replicating their DNA and then dividing the duplicated genome and cellular contents into two new cells. This commitment to proliferation is dependent on proteins called growth factors (also known as mitogens) and has been likened to an irreversible switch, termed the restriction point, that occurs before DNA replication starts. According to that model, if growth factors are withdrawn before this molecular switch is flipped, cells will return to a non-proliferative state called quiescence (also known as G0). But if they are removed after this switch has been flipped (Fig. 1a), cells will complete a round of DNA replication and cell division before re-entering quiescence[1,2]. Or so we thought. On page 363, Cornwell *et al.*[3] present data that challenge this model.

The authors show that cells that were thought to be irreversibly committed to proliferation by the flipping of this switch do not necessarily complete cell division. Instead, Cornwell and colleagues show that if growth factors are withdrawn after the proposed switch has been flipped, the cells sometimes just replicate their DNA without dividing. Intriguingly, whether a cell completes cell division or withdraws from the cell cycle can be attributed to the amount of a single protein — cyclin A2.

This work stems from the authors' initial observation that if growth-factor signalling was disrupted after human cells had flipped the switch, a small population of the cells (up to 15%, depending on the cell type) did not complete cell division and only replicated their DNA. Therefore, about 15% of cells were not committed to cell proliferation. Using single-cell time-lapse imaging, the authors were able to

From the archive

Concerns about preventable cases of infant death, and praise for a museum guide book about fossils.

50 years ago

[A] meticulous study of birth records in New York City ... involved examining the records of all the births which took place ... in 1968 ... If all the women ... had received adequate prenatal care, infant mortality could have been cut by one third, the study suggests ... Small wonder, therefore, that Dr Robert Coles of Harvard University said in a preface to the report that "we do things wrong, we are indifferent to the needs of others — and here, right here is the proof."

From *Nature* 13 July 1973

100 years ago

British Museum (Natural History). Guide to the Exhibition Galleries of Geology and Palaeontology — The Keeper of Geology, in his preface to this small book, says, "It is merely a guide, not an introduction to the study of fossils." Those familiar with official scientific publications may appreciate the modesty and wisdom of this statement. But intelligent members of the general public ... will soon find that the statement errs on the side of diffidence; they will say, "This is not merely a guide, but a remarkably good guide" ... The casual visitor to these magnificent geological collections is often bewildered by the multitude of objects and oppressed by the strangeness of nomenclature. With this guide ... the systematic names are explained in everyday terms and the essential characters of the fossils are made clear, while no opportunity is lost of showing how the forms of these extinct creatures throw light upon their habits and phylogeny. Thus a great deal of sound information is woven into a readable story, which does not neglect human interest but links up the fossils with their discoverers or with some apt reference to literature or history. Who will not be tempted after reading of Thomas Hawkins to look up his descriptions of the hunt for Ichthyosauri, or to renew an acquaintance with "The Chambered Nautilus" of Oliver Wendell Holmes?

From *Nature* 14 July 1923