nature

NEWS & VIEWS

QUANTUM INFORMATION

To compute or not to compute?

Jonathan P. Dowling

Quantum physics aims another blow at common sense: a simple quantum computer gives the right answer, even when it is not run. (Traditionalists be comforted: the computer must be turned on.)

Is it possible to get a sensible answer from a computer without even running a program? Yes indeed, report Hosten *et al.* on page 949 of this issue[1] — and they have the experimental proof.

The thread that leads thus far has its origin in the early 1990s, with the introduction of a quantum paradox known as the Elitzur–Vaidman bomb thought-experiment[2]. They showed that by sending a photon into a simple optical interferometer (Box 1), one can sometimes detect the presence of a light-triggered bomb without setting it off — in other words, without light interacting with it at all. 'Sometimes' — there's the rub. At other times, you set the bomb off; at still other times, you get no information about the existence of the bomb at all. But then that's the probabilistic nature of quantum mechanics. It might seem to defy common sense, but the result is a straightforward application of the quantum theory of light.

A few years ago, it was pointed out[3] that an all-optical quantum computer could replace Elitzur and Vaidman's bomb. Rather than exploding or not, this computer either runs a program or doesn't, depending on whether or not a single photon impinges on its 'run' switch[3]. And just as we can obtain the answer to the question 'is the bomb there?' without exploding it, so we can obtain the right answer from the computer without running it. (The computer must, at least, still be turned on and properly programmed in order for the scheme to work.) However, although intriguing, this original 'counterfactual computation' scheme suffered from the same quantum-probabilistic issues as the bomb business: sometimes the device would actually run the quantum computer to get the answer; sometimes it gave no answer at all. The probability that it would give the correct answer while not running the computer was so small that one could always do better by guessing the answer randomly. Not much of a computer, that.

Another spin-off of the bomb idea was Anton Zeilinger and colleagues' elegant and real (as opposed to thought) experiment called 'quantum seeing in the dark'[4,5]. That experiment allowed safe, non-explosive objects such as a human hair to be imaged, sometimes using no photons at all. Again, the key word is 'sometimes': just as with the bomb, sometimes no photons were needed to image the hair; sometimes some photons were needed to image the hair; and sometimes the hair was not imaged at all. By sorting the individual imaging events carefully, however, an image of the hair could be constructed using only the events in which no photon had interacted with the hair. Better still, the probability of these photon-free imaging events could be increased using yet another paradox of quantum mechanics, the quantum Zeno effect[6].

This effect is named in honour of Zeno of Elea, the philosopher who gave us paradoxes such as that of Achilles and the tortoise (you can run, but you can't catch up). It is the quantum-physical equivalent of the saying "a watched pot never boils". If you measure the state of a quantum system often enough and fast enough, the system will remain in that state and never evolve to another state — even if it would evolve if you were not measuring it. Zeilinger and colleagues applied this idea by using a sequence of polarization rotators as 'Zeno boosters' that forced their interferometer system to remain far more often in the state of 'imaging the hair with no photons', and far less often in the states of 'imaging the hair

Box 1 How to stop worrying and learn to love the bomb

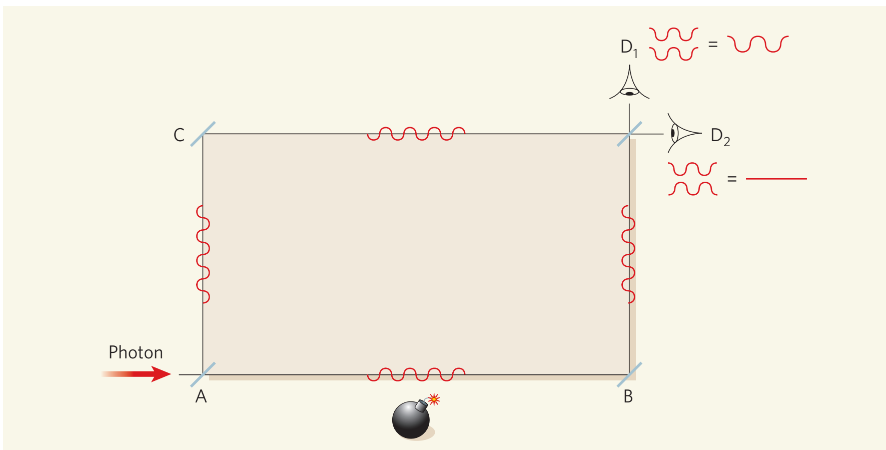

Avshalom Elitzur and Lev Vaidman's bomb thought-experiment[2] uses a simple wave-splitting device known as an optical interferometer and the quantum phenomenon of wave–particle duality. If just one photon enters their interferometer at A, its associated wave splits, going with a certain probability by way of either B or C. The geometry of the interferometer is set up such that, at one of two detectors at its far end (D_1), the waves from the two paths interfere constructively, so light is detected. At the other detector (D_2), on the other hand, the waves interfere destructively, so it is never triggered — unless one of the pathways, say B, is blocked off, preventing the passage of a photon. In this case, there can be no interference, so the photon will be detected with equal probability by D_1 or D_2.

Now suppose that B is — possibly — blocked by a photon-triggered bomb. A single photon enters the interferometer. Three outcomes are possible. First, the bomb goes off: so there is a bomb. Second, D_1 triggers. This outcome does not help us, as it could mean one of two things: there is a bomb, but the photon passed by C, or there isn't a bomb, and waves from the two paths are interfering constructively. Third, D_2 triggers. In this case, there must be a bomb blocking B, but our photon passed by C. In other words, we have 'seen' the bomb without a photon ever touching it. **J.P.D.**

with photons' or 'not imaging the hair at all'[4,5].

One might think that the same Zeno trick could be used to improve counterfactual quantum computers, so that they could actually be useful for something — or at least, that the probability of their finding the right answer should be better than that achieved by flipping coins or tossing dice. But this hope proved misguided: the straightforward application of the Zeno idea to counterfactual quantum computing did not seem to beat the random-guessing limit, and things were at a standstill.

Enter Hosten *et al.*[1]. These authors have discovered that chaining Zeno boosters together produces a super Zeno booster that is indeed capable of beating the random-guessing limit in counterfactual quantum computing. They also demonstrate that the idea works with a quantum-optical implementation of the Grover search algorithm. This algorithm proves that a quantum computer can always perform better than any classical computer in searching an unsorted database for a target datum[7]. It has a particularly simple implementation in quantum optical interferometry, as does the super-sized quantum Zeno effect. Hosten and colleagues show how the two can be married to locate the target element in a database — without running the quantum computer that searches for that element in the first place.

So what is counterfactual quantum computing good for, other than to illustrate the conclusion of a rather obscure quantum paradox? The computer must still be programmed and turned on, even if it is not run, so the approach will not save on electricity bills or labour costs. In fact, the value of the experiment lies simply in furthering our understanding of quantum mechanics and its interface with computation. The entire field of quantum information arose from physicists trying to understand the implications of paradoxes in quantum theory, and already the field is beginning to evolve on its own. The first commercial applications of quantum information, such as real-world quantum cryptography, are now being deployed.

A century ago, we saw the start of the first quantum revolution — the discovery of the quantum rules that underpin our world. We are now on the verge of a second revolution[8], in which these rules spawn technological applications. Results such as those of Hosten and colleagues[1] are significant markers on the road to that revolution. ■

Jonathan P. Dowling is at the Hearne Institute for Theoretical Physics, Department of Physics and Astronomy, Louisiana State University, Baton Rouge, Louisiana 70803, USA, and the Institute for Quantum Studies, Texas A&M University.
e-mail: jdowling@lsu.edu

1. Hosten, O., Rakher, M. T., Barreiro, J. T., Peters, N. A. & Kwiat, P. G. *Nature* **439,** 949–952 (2006).
2. Elitzur, A. C. & Vaidman, L. *Found. Phys.* **23,** 987–997 (1993).
3. Mitchison, G. & Jozsa, R. *Proc. R. Soc. Lond. A* **457,** 1175–1193 (2001).
4. Kwiat, P. *et al. Phys. Rev. Lett.* **74,** 4763–4766 (1995).
5. Kwiat, P., Weinfurter, H. & Zeilinger, A. *Sci. Am.* **275**(5), 72–78 (1996).
6. Knight, P. *Nature* **344,** 493–494 (1990).
7. Grover, L. K. *Phys. Rev. Lett.* **79,** 325–328 (1997).
8. Dowling, J. P. & Milburn, G. J. *Phil. Trans. R. Soc. Lond. A* **361,** 1655–1674 (2003).

NEUROBIOLOGY

Efficiency measures

Michael R. DeWeese and Anthony Zador

The nervous system translates sensory information into electrical impulses. The neural 'code' involved seems to represent natural sounds and images efficiently, using the smallest number of impulses.

Our perception of the outside world relies on the transformation of physical signals (such as light and sound) into a pattern of neural impulses, or spikes. These spikes are then transmitted to higher brain regions, where they are further transformed into other patterns of sensory spikes, and ultimately into the motor spikes that mediate behaviour. What is the relationship (the 'neural code') between these neural responses and the sensory signals they represent? Are there general principles underlying the neural code?

The 'efficient-coding hypothesis'[1] proposes that sensory neurons are adapted to the statistical properties of sensory signals to which neurons are exposed. Two papers in this issue invoke this principle to predict how neurons encode natural auditory and visual stimuli, as opposed to the artificial stimuli often used in experiments. Smith and Lewicki (page 978)[2] develop an algorithm to find an efficient representation of natural sounds and speech, and show that this theoretically predicted representation matches that observed experimentally in the auditory nerve of cats. Sharpee and colleagues (page 936)[3] show that cortical neurons adapt over seconds or minutes during the course of an experiment to maximize the information they provide about the stimulus. Together, the two papers show how the efficient-coding hypothesis can help to make sense of properties of the neural code on both evolutionary and behavioural timescales.

In the cochlea, sound is encoded into spikes, which are transmitted along the auditory nerve to higher stations in the auditory system. Auditory nerve fibres each respond to a narrow range of sound frequencies, with the range generally increasing with the median frequency. The response of each auditory nerve fibre can therefore be modelled as a (nonlinear) 'filter' that removes frequencies outside a particular range. Why do the auditory nerve filters have the particular form they do? Smith and Lewicki reasoned that if the auditory code is indeed 'efficient', then they should be able to predict the form of the auditory filter bank by finding the sparsest code; that is, the one that requires the least activity.

To obtain this prediction, Smith and Lewicki first expressed the efficient-coding hypothesis as an algorithm whose input is an ensemble of sounds, and whose output is a sparse encoding for transmitting or representing this ensemble. The algorithm discovers that the sparsest encoding of sounds is into brief events suggestive of spikes, the precise timing of which conveys much of the information. The sparsest code depends on the ensemble of sounds to be encoded; a code that is most efficient for one set of sounds is not necessarily most efficient for another.

Why should the most efficient code depend on the stimulus ensemble? The basic intuition is straightforward. Suppose I ask you to describe individual sounds produced by different musical instruments, but I limit your vocabulary to only four words (of your choosing). If you know that the instruments are used in a rock band (that is, they are chosen from the rock ensemble), you might choose a code consisting of the words 'guitar', 'bass', 'drums', 'keyboard'; but if the instruments are used in a classical orchestra (the classical ensemble), you might choose instead 'woodwind', 'brass', 'percussion', 'string'. So, the choice of the most efficient code depends on what is being described.

The dependence of the sparsest code on the ensemble raises the question of exactly which ensemble should be considered the evolutionarily relevant 'natural' ensemble. Smith and Lewicki therefore tested three subsets of sounds: animal vocalizations, transient environmental sounds such as crunching leaves and cracking twigs, and ambient environmental sounds like rain. They found that if the ensemble consisted of any single sub-ensemble, then the predicted sparse code failed to match previous experimental observations of cat auditory filters. However, the code predicted from a carefully balanced mixture of these natural sounds provided a good match to the experimental data, supporting the idea that these filters evolved to encode natural sounds efficiently. Interestingly, filters calculated from a human speech ensemble also fit the cat data well; because (as any cat owner will attest) the cat auditory system did not evolve to