It from Qubit

David Deutsch

Centre for Quantum Computation, The Clarendon Laboratory, University of Oxford

September 2002

To Appear in *Science & Ultimate Reality*, John Barrow, Paul Davies, Charles Harper, Eds. (Cambridge, UK: Cambridge University Press, 2003)

Introduction

Of John Wheeler's 'Really Big Questions', the one on which the most progress has been made is *It From Bit?* – does information play a significant role at the foundations of physics? It is perhaps less ambitious than some of the other Questions, such as *How Come Existence?*, because it does not necessarily require a metaphysical answer. And unlike, say, *Why The Quantum?*, it does not require the discovery of new laws of nature: there was room for hope that it might be answered through a better understanding of the laws as we currently know them, particularly those of quantum physics. And this is what has happened: the better understanding is the quantum theory of information and computation.

How might our conception of the quantum physical world have been different if *It From Bit* had been a motivation from the outset? No one knows how to derive *it* (the nature of the physical world) from *bit* (the idea that information plays a significant role at the foundations of physics), and I shall argue that this will never be possible. But we can do the next best thing: we can start from the qubit.

Qubits

To a classical information theorist, a *bit* is an abstraction: a certain amount of information. To a programmer, a bit is a Boolean variable. To an engineer, a bit is a 'flip-flop' – a piece of hardware that is stable in either of two physical states. And to a physicist? Quantum information theory differs in many ways from its classical predecessor. One reason is that quantum theory provides a new answer to the ancient dispute, dating back to the Stoics and the Epicureans and even earlier, about whether the world is discrete or continuous.

Logic is discrete: it forbids any 'middle' between true and false. Yet in classical physics, discrete information processing is a derivative and rather awkward concept. The fundamental classical observables vary continuously with time and, if they are fields, with space too, and they obey differential equations. When classical physicists spoke of discrete observable quantities, such as how many moons a planet had, they were referring to an idealisation, for in reality there would have been a continuum of possible states of affairs between a particular moon's being 'in orbit' around the planet and 'just passing by', each designated by a different real number or numbers. Any two such sets of real numbers, however close, would refer to physically different states which would evolve differently over time and have different physical effects. (Indeed the differences between them would typically grow exponentially with time because of the instability of classical dynamics known as 'chaos'.) Thus, since even one real variable is equivalent to an infinity of independent discrete variables – say, the infinite sequence of zeros and ones in its binary expansion – an infinite amount of in-principle-observable information would be present in any classical object.

Despite this ontological extravagance, the continuum is a very natural idea. But then, so is the idea (which is the essence of information processing and therefore of *It From Bit*) that complicated processes can be analysed as combinations of simple ones. These two ideas have not been easy to reconcile. With the benefit of hindsight, I think that this is what Zeno's paradox of the impossibility of motion was really about. Had he been familiar with classical physics and the concept of information processing, he might have put it like this: Consider the flight of an arrow as described in classical physics. To understand what happens during the flight, we could try to regard the real-valued position coordinates of the arrow as pieces of information, and the flight as a computation that processes that information, and we could try to analyse that computation as a sequence of elementary computations. But in that case, what is the 'elementary' operation in question? If we regard the flight as consisting of a *finite* number of shorter flights, then each of them is, by any straightforward measure, exactly as complicated as the whole: it comprises exactly as many sub-steps, and the positions that the arrow takes during it are in one-one correspondence with those of the whole flight. Yet if, alternatively, we regard the flight as consisting of a literally infinite number of infinitesimal steps, what exactly is the effect of such a step? Since there is no such thing as a real number infinitesimally

greater than another, we cannot characterise the effect of this infinitesimal operation as the transformation of one real number into another, and so we cannot characterise it as an elementary computation performed on what we are trying to regard as information.

For this sort of reason, *It From Bit* would be a non-starter in classical physics. It is noteworthy that the black body problem, which drove Max Planck unwillingly to formulate the first quantum theory, was also a consequence of the infinite information-carrying capacity of the classical continuum.

In quantum theory, it is continuous observables that do not fit naturally into the formalism (hence the name *quantum* theory). And that raises another paradox – in a sense the converse of Zeno's: if the spectrum of an observable quantity (the set of possible outcomes of measuring it) is not a continuous range but a discrete set of values, how does the system ever make the transition from one of those values to another? The remarkable answer given by quantum theory is that it makes it continuously. It can do that because a quantum observable – the basic descriptor of quantum reality – is neither a real variable, like a classical degree of freedom, nor a discrete variable like a classical bit, but a more complicated object that has both discrete and continuous aspects.

When investigating the foundations of quantum theory, and especially the role of information, it is best to use the Heisenberg picture, in which quantum observables (which I shall mark with a caret, as in $\hat{X}(t)$) change with time, and the quantum state $|\Box\rangle$ is constant. Though the Schrödinger picture is equivalent for all predictive purposes, and more efficient for most calculations, it is very bad at representing information flow and has given rise to widespread misconceptions (see Deutsch and Hayden 2000).

Apart from the trivial observables that are multiples of the unit observable $\hat{1}$, and hence have only one eigenvalue, the simplest type of quantum observable is a *Boolean observable* – defined as one with exactly two eigenvalues. This is the closest thing that quantum physics has to the classical programmer's idea of a Boolean variable. But the engineer's flip-flop is not just an observable: it is a whole physical system. The simplest quantum system that contains a Boolean observable is a *qubit.* Equivalently, a qubit can be defined as any system all of whose non-trivial

observables are Boolean. Qubits are also known as 'quantum two-state systems' (though this is a rather misleading term because, like all quantum systems, a qubit has a continuum of physical states available to it). The spin of a spin-$\frac{1}{2}$ particle, such as an electron, is an example. The fact that a qubit is a type of physical system, rather than a pure abstraction, is another important conceptual difference between the classical and quantum theories of information.

We can describe a qubit Q at time t elegantly in the Heisenberg picture (Gottesman 1999) using a triple $\hat{\mathbf{q}}(t) = \left(\hat{q}_x(t), \hat{q}_y(t), \hat{q}_z(t)\right)$ of Boolean observables of Q, satisfying

$$\begin{matrix} \hat{q}_x(t)\,\hat{q}_y(t) = i\hat{q}_z(t) \\ \hat{q}_x(t)^2 = \hat{1} \end{matrix} \quad \left(\text{and cyclic permutations over } (x, y, z)\right). \tag{1}$$

All observables of Q are linear combinations, with constant coefficients, of the unit observable $\hat{1}$ and the three components of $\hat{\mathbf{q}}(t)$. Each Boolean observable of Q changes continuously with time, and yet, because of (1), retains its fixed pair of eigenvalues which are the only two possible outcomes of measuring it.

Although this means that the classical information storage capacity of a qubit is exactly one bit, there is no elementary entity in nature corresponding to a bit. It is qubits that occur in nature. Bits, Boolean variables, and classical computation are all emergent or approximate properties of qubits, manifested mainly when they undergo decoherence (see Deutsch 2002a).

The standard model of quantum computation is the *quantum computational network* (Deutsch 1989). This contains some fixed number N of qubits

$$Q_a \quad (1 \leq a \leq N), \text{ with } \left[\hat{\mathbf{q}}_a(t), \hat{\mathbf{q}}_b(t)\right] = 0 \quad (a \neq b), \tag{2}$$

where $\hat{\mathbf{q}}_a(t) = \left(\hat{q}_{ax}(t), \hat{q}_{ay}(t), \hat{q}_{az}(t)\right)$.

In physical implementations, qubits are always subsystems of other quantum systems – such as photons or electrons – which are themselves manipulated via a larger apparatus in order to give the quantum computational network its defining properties. However, one of those properties is that the network is *causally autonomous*: that is to say, the law of motion of each qubit depends only on its own

observables and those of other qubits of the network, and the motion required of the external apparatus is independent of that of the qubits. Hence, all the external paraphernalia can be abstracted away when we study the properties of quantum computational networks.

Furthermore, we restrict our attention to networks that perform their computations in a sequence of *computational steps*, and we measure the time in units of these steps. The *computational state* of the network at integer times t is completely specified by all the observables $\hat{\mathbf{q}}_a(t)$. Although any real network would interpolate smoothly between computational states during the computational step, we are not interested in the computational state at non-integer times. The network at integer times is itself a causally autonomous system, and so, just as we abstract away the external apparatus, we also abstract away the network itself at non-integer times.

The computational state is not to be confused with the Heisenberg state $|\Box\rangle$ of the network, which is constant, and can always taken to be the state in which

$$\langle\Box|\hat{q}_{az}(0)|\Box\rangle = 1, \tag{3}$$

so that all the $\hat{q}_{az}$ observables are initially sharp with values +1. (In this convention, the network starts in a standard, 'blank' state at t=0, and we regard the process of providing the computation with its input as being a preliminary computation performed by the network itself.)

During any one step, the qubits of the network are separated (dynamically, not necessarily spatially) into non-overlapping subsets such that the qubits of each subset interact with each other, but with no other qubits, during that step. We call this process 'passing through a quantum gate' – a gate being any means of isolating a set of qubits and causing them to interact with each other for a fixed period. Because we are interested only in integer times, the relevant effect of a gate is its net effect over the whole computational step. The effect of an n-qubit quantum gate may be characterised by a set of $3n$ functions, each expressing one of the observables in the set $\{\hat{\mathbf{q}}_a(t+1)\}$ (where a now ranges over the indices of qubits passing through the gate between times t and t+1) in terms of the $3n$ observables $\{\hat{\mathbf{q}}_a(t)\}$, subject to the constraint that the relations (1) and (2) are preserved. Every such set of functions describes a possible quantum gate. For examples see Deutsch and Hayden (2000).

Between these interactions, the qubits are computationally inert (none of their observables change); they merely move (logically, not necessarily spatially) from the output of one gate to the input of the next. Thus the dynamics of a quantum computational network can be defined by specifying a network of gates linked by 'wires'.

It might seem from this description that the study of quantum computational networks is a narrow sub-speciality of physics. Qubits are special physical systems, and are often realised as subsystems of what are normally considered 'elementary' systems (such as elementary particles). In quantum gates, qubits interact in a rather unusual way: they strongly affect each other while remaining isolated from the environment; their periods of interaction are synchronised, alternating with periods of inertness; and so on. We even assume that all the qubits of the network start out with their spins pointing in the $+z$-direction (or whatever the initial condition (3) means for qubits that are not spin-$\frac{1}{2}$ systems). None of these attributes is common in nature, and none can ever be realised perfectly in the laboratory. At the present state of technology, realising them well enough to perform any useful computation is still a tremendously challenging, unattained target.

Yet quantum computational networks have another property which makes them far more worthy of both scientific and philosophical study than this way of describing them might suggest. The property is *computational universality*.

Universality

Universality has several interrelated aspects, including:

- the fact that a single, standard type of quantum gate suffices to build quantum computational networks with arbitrary functionality;
- the fact that quantum computational networks are a universal model for computation;
- the fact that a universal quantum computer can simulate, with arbitrary accuracy, the behaviour of arbitrary physical systems;
- the fact (not yet verified) that such computers can be constructed in practice.

The first of those concerns *universal gates*. One of the ways in which the theory of quantum computation lives up to the *It-From-Bit* intuition is that in the most natural sense, the computation performed by the component gates of a network can indeed be simpler than that performed by the network as a whole. The possible motions of one or two qubits through a gate, though continuous, are not isomorphic to the possible motions of a larger network; but by composing multiple instances of only a single type of gate that performs a fixed, elementary operation, it is possible to construct networks performing arbitrary quantum computations. Any gate with this property is known as a *universal quantum gate*. It turns out that not only do there exist universal gates operating on only two qubits, but in the manifold of all possible two-qubit gates, only a set of measure zero are *not* universal (Deutsch, Barenco and Ekert 1995).

Thus, computational universality is a generic property of the simplest type of gate, which itself involves interactions between just two instances of the simplest type of quantum system. There are also other ways of expressing gate-universality: for instance, the set of all single-qubit gates, together with the controlled-not operation (measurement of one qubit by another) also suffice to perform arbitrary computations. Alternatively, so do single-qubit gates together with the uniquely quantum operation of 'teleportation' (Gottesman and Chuang 1999). All this constitutes a strikingly close connection between quantum computation and quantum physics – of which there were only hints in classical computation and classical physics. Models of classical computation based on idealised classical systems such as 'billiard balls' have been constructed in theory (Fredkin and Toffoli, 1982), but they are unrealistic in several ways, and unstable because of 'chaos', and no approximation to such a model could ever be a practical computer. Constructing a universal classical computer (such as Babbage's analytical engine) from 'elementary' components that are well described in a classical approximation (such as cogs and levers) requires those components to be highly composite, precision-engineered objects which would fail in their function if they had an even slightly different shape.

The same is true of the individual transistors on the microchips that are used to build today's classical computers. But it is not true, for instance, of the ions in an ion trap (Cirac and Zoller 1995, Steane 1997) – one of many quantum systems that are currently being investigated for possible use as quantum computers. In an ion trap, a

group of ions is held in place in a straight line by an ingeniously shaped oscillating electric field. In each ion, one electron forms a two-state system (the states being its ground state and one of its excited states) which constitutes a qubit. The ions interact with each other via a combination of the Coulomb force and an external electromagnetic field in the form of laser light – which is capable of causing the observables of any pair of the qubits to change continuously when the laser is on. *The engineering problem ends there.* Once an arrangement of that general description is realised, the specific form of the interaction does not matter. Because of the generic universality of quantum gates, there is bound to exist *some* sequence of laser pulses – each pulse constituting a gate affecting two of the qubits – that will cause an N-ion trap to perform any desired N-qubit quantum computation.

The same sort of thing applies in all the other physical systems – nuclear spins, superconducting loops, trapped electrons and many more exotic possibilities – that serve, or might one day serve, as the elementary components of quantum computers. Lloyd (1995) has summed this up in the aphorism: 'Almost any physical system becomes a quantum computer if you shine the right sort of light on it'. There is no classical analogue of this aphorism.

Quantum computers are far harder to engineer than classical computers, of course, but not for the same reason. Indeed the problem is almost the opposite: it is not to engineer precisely-defined composite systems for use as components, but rather, to isolate the physically simplest systems that already exist in nature, from the complex systems in their environment. That done, we have to find a way of allowing arbitrary pairs of them to interact – in some way – with each other. But once that is achieved in a given type of physical system, no shaping or machining is necessary, because the interactions that quantum systems undergo as a matter of course are already computationally universal.

The second aspect of universality is that quantum networks are a universal model for computation. That is to say, consider any technology that could, one day, be used to perform computations – whether quantum or classical, and whether based on gates or anything else. For any computer C built using that technology, there exists a quantum computational network, composed entirely of simple gates (such as instances of a single two-qubit universal gate), that has at least the same repertoire of computations as C. Here we mean 'the same repertoire' in quite a strong sense:

- Given a computational task (say, factorisation) and an input (say, an integer), the network could produce the same output as C does (say, the factors of the integer).

- The resources (number of gates, time, energy, mass of raw materials, or whatever) required by the network to perform a given computation would be bounded by a low power of those required by C. I conjecture that this power can be 1. That is to say, there exists a technology for implementing quantum computational networks under which they can emulate computers built under any other technology, using only a constant multiple of the resources required under that technology.

- The network could emulate more than just the relationship between the output of C and its input. It could produce the output *by the same method* – using the same quantum algorithm – as C.

The upshot is that the abstract study of quantum computations (as distinct from the study of how to implement them technologically) is effectively the same as the study of one particular class of quantum computational networks (which need only contain one type of universal quantum gate). This universality is the quantum generalisation of that which exists in classical computation, where the study of all computations is effectively the same as the study of any one universal model, such as logic networks built of NAND gates or Toffoli gates, or the universal Turing machine.

However, quantum universality has a further aspect which was only guessed at – and turned out to be lacking – in the case of classical computation: quantum computational networks can simulate, with arbitrary accuracy, the behaviour of arbitrary physical systems; and they can do so using resources that are at most polynomial in the complexity of the system being simulated. The most general way of describing quantum systems (of which we are at all confident) is as *quantum fields*. For instance, a scalar quantum field $\hat{\phi}(\mathbf{x},t)$ consists of an observable for every point $(\mathbf{x},t)$ of spacetime, satisfying a differential equation of motion. There are many possible approximation schemes for computing the behaviour of such a system by approximating the continuous spacetime fields with continuous spectra as finite sets of observables with finite spectra, on a spacetime lattice. Such approximation

schemes would be suitable for quantum computation too, where, for instance, a finite number of qubits would simulate the behaviour of the field $\hat{\phi}$ in the vicinity of each of a set of spatial grid points.

However, suppose that we had come upon quantum field theory from the other direction, convinced from the outset that 'it' (a quantum field) is made of qubits. A quantum field can certainly be expressed in terms of fields of Boolean observables. For instance, the set of all Boolean observables 'whether the average value of the field over a spacetime region R exceeds a given value ϕ', as R ranges over all regions of non-zero volume and duration, and ϕ ranges over all real numbers, contains the same information as the quantum field $\hat{\phi}(\mathbf{x},t)$ itself (albeit redundantly). For each of these Boolean observables, we can construct a 'simplest' quantum system containing it, and that will be a qubit.

Local interactions could be simulated using gates in which qubits interact with close neighbours only. In this way, quantum networks could simulate arbitrary physical systems not merely in the bottom-line sense of being able to reproduce the same output (observable behaviour), but again, in the strong sense of mimicking the physical events, locally and in arbitrary detail, that bring the outcome about.

In most practical computations, we should only be interested in the output for a given input and not (unless we are the programmer) in how it was brought about. But there are exceptions. An amusing example is given in the science fiction novel *Permutation City* by Greg Egan. In it, technology has reached the point where the computational states of human brains can be uploaded into a computer, and simulations of those brains, starting from those states, interact there with each other and with a virtual-reality environment – a self-contained world of the clients' choice. Because these computations are expensive, the people who run the service are continually seeking ways to optimise the program that performs this simulation. They run an optimisation algorithm which systematically examines the program, replacing pieces of code or data with other pieces that achieve the identical effect in fewer steps. The simulated people cannot of course perceive the effect of such optimisations – and yet … eventually the optimisation program halts, having deleted the entire simulation with all its data, and reports '*this program generates no output*'.

By the way, there is no reason to believe that a universal *quantum* computer would be required for such simulations (see Tegmark 2000). There is every reason to believe that the brain is a universal classical computer. Nevertheless this strong form of universality of quantum computation assures us that such a technology, and artificial intelligence in general, must be possible, and tractable, regardless of how the brain works.

Provided, that is, that *universal quantum computers can be built in practice*. This is yet another aspect of universality, perhaps the most significant for the *It From Qubit?* question. Indeed, universality itself may not be considered quite as significant by many physicists and philosophers if it turns out that qubits cannot, in reality, be composed into networks with universal simulating capabilities.

The world is not 'made of information'

Let us suppose that universality does hold in all four of the above senses. Then, since every physical system can be fully described as a collection of qubits, it is natural to wonder whether this can be taken further. Might it have been possible to *start* with such qubit fields and to interpret traditional quantum fields as emergent properties of them? The fact that all quantum systems that are known to occur in nature obey equations that look fairly simple in the language of fields on spacetime, is perhaps evidence against such a naive 'qubits-are-fundamental' view of reality. On the other hand, we have some evidence in its favour too. One of the few things that we think we know about the quantum theory of gravity is expressed in the so-called Bekenstein bound: the entropy of any region of space cannot exceed a fixed constant times the surface area of the region (Bekenstein 1981). This strongly suggests that the complete state space of any spatially finite quantum system is finite, so that, in fact, it would contain only a finite number of independent qubits.

But even if this most optimistic quantum-computation-centred view of physics turned out to be true, it would not support the most ambitious ideas that have been suggested about the role that information might play at the foundations of physics. The most straightforward such idea, and also the most extreme, is that the whole of what we usually think of as reality is merely a program running on a gigantic computer – a Great Simulator. On the face of it, this might seem a promising approach to explaining the connections between physics and computation: perhaps the reason why the laws of physics are expressible in terms of computer programs is

that they are in fact computer programs; perhaps the existence of computers in nature is a special case of the ability of computers (in this case the Great Simulator) to emulate other computers; the locality of the laws of physics is natural because complex computations are composed of elementary computations – perhaps the Great Simulator is a (quantum?) cellular automaton – and so on. But in fact this whole line of speculation is a chimera.

It entails giving up on explanation in science. It is in the very nature of computational universality that if we and our world were composed of software, we should have no means of understanding the real physics – the physics underlying the hardware of the Great Simulator itself. Of course, no one can prove that we are not software. Like all conspiracy theories, this one is untestable. But if we are to adopt the methodology of believing such theories, we may as well save ourselves the trouble of all that algebra and all those experiments, and go back to explaining the world in terns the sex lives of Greek gods.

An apparently very different way of putting computation at the heart of physics is to postulate that 'all possible laws of physics' (in some sense) are realised in nature, and then to try to explain the ones that we see, entirely as a selection effect (see e.g. Smolin 1997). But selection effects, by their very nature, can never be the whole explanation for the apparent regularities in the world. That is because making predictions about an ensemble of worlds (say, with different laws of physics, or different initial conditions) depends on the existence of a measure on the ensemble, making it meaningful to say things like 'admittedly, most of them do not have property X, but most of the ones in which anyone exists to ask the question, do'. But there can be no *a priori* measure over 'all possible laws'. Tegmark (1997) and others have proposed that the complexity of the law, when it is expressed as a computer program, might be this elusive measure. But that merely raises the question: *complexity according to which theory of computation?* Classical and quantum computation, for instance, have very different complexity theories. Indeed, the very notion of 'complexity' is irretrievably rooted in physics, so in this sense physics is necessarily prior to any concept of computation. 'It' cannot possibly come from 'bit', or from qubit, by this route. (See also my criticism of Wheeler's *Law Without Law* idea – Deutsch 1986.)

Both these approaches fail because they attempt to reverse the direction of the explanations that the real connections between physics and computation provide. They seem plausible only because they rely on a common misconception about the status of computation *within mathematics*. The misconception is that the set of computable functions (or the set of quantum-computational tasks) has some a priori privileged status within mathematics. But it does not. The only thing that privileges that set of operations is that it is instantiated in the computationally universal laws of physics. It is only through our knowledge of physics that we know of the distinction between computable and non-computable (see Deutsch, Ekert and Luppaccini 2000), or between simple and complex.

The world is made of qubits

So, what does that leave us with? Not 'something for nothing': information does not create the world ex nihilo. Nor a world whose laws are really just fiction, so that physics is just a form of literary criticism. But a world in which the stuff we call information, and the processes we call computations, really do have a special status. The world contains – or at least, is ready to contain – universal computers. This idea is illuminating in a way that its mirror-image – that a universal computer contains the world – could never be.

The world is made of qubits. Every answer to a question about whether something that could be observed in nature is so or not, is in reality a Boolean observable. Each Boolean observable is part of an entity, the qubit, that is fundamental to physical reality but very alien to our everyday experience. It is the simplest possible quantum system and yet, like all quantum systems, it is literally not of this universe. If we prepare it carefully so that one of its Boolean observables is sharp – has the same value in all the universes in which we prepare it – then according to the uncertainty principle, its other Boolean observables cease to be sharp: there is no way we can make the qubit as a whole homogeneous across universes. Qubits are unequivocally multiversal objects. This is how they are able to undergo continuous changes even though the outcome of measuring – or being – them is only ever one of a discrete set of possibilities.

What we perceive to some degree of approximation as a world of single-valued variables is actually part of a larger reality in which the full answer to a yes-no question is never just yes or no, nor even both yes and no in parallel, but a quantum

observable — something that can be represented as a large Hermitian matrix. Is it really possible to conceive of the world, including ourselves, as being 'made of matrices' in this sense? Zeno was in effect asking the same question about real numbers in classical physics: how can we be made of real numbers? To answer that question we have to do as Zeno did, and analyse the *flow of information* – the information processing – that would occur if this conception of reality were true. Whether we could be 'made of matrices' comes down to this: what sort of experiences would an observer composed entirely of matrices, living in a world of matrices, have? The theories of decoherence (Zurek 1981) and consistent histories (Hartle 1991) have answered that question in some detail (see also Deutsch 2002a): at a coarse-grained level the world looks as though classical physics is true; and as though the classical theories of information and computation were true too. But where coherent quantum processes are under way – particularly quantum computations – there is no such appearance, and an exponentially richer structure comes into play.

As Karl Popper noted, the outcome of solving a problem is never just a new theory but always a new problem as well. In fundamental science this means, paradoxically, that new discoveries are always disappointing for those who hope for a final answer. But it also means that they are doubly exhilarating for those who seek ever more, and ever deeper, knowledge.

The argument that I used above to rule out Great-Simulator-type explanations has implications for genuine physics too: Although in one sense the quantum theory of computation contains the whole of physics (with the possible exception of quantum gravity), the very power of the principle of the universality of computation inherently limits the theory's scope. Universality means that computations, and the laws of computation, are independent of the underlying hardware. And therefore, the quantum theory of computation cannot explain hardware. It cannot, by itself, explain why some things are technologically possible and others are not. For example, steam engines are, perpetual motion machines are not, and yet the quantum theory of computation knows nothing of the second law of thermodynamics: if a physical process can be simulated by a universal quantum computer, then so can its time reverse. An example closer to home is that of quantum computers themselves: the last aspect of universality that I mentioned above – that universal quantum computers can be built in practice – has not yet been

verified. Indeed, there are physicists who doubt that it is true. At the present state of physics, this controversy, which is a very fundamental one from the *It From Qubit* point of view, cannot be addressed from first principles. But if there is any truth in the *It From Qubit* conception of physics that I have sketched here, then the quantum theory of computation as we know it must be a special case of a wider theory.

Quantum constructor theory (Deutsch 2002b) is the theory that predicts which objects can (or cannot) be constructed, and using what resources. It is currently in its infancy: we have only fragmentary knowledge of this type – such as the laws of thermodynamics, which can be interpreted as saying that certain types of machine (perpetual motion machines of the first and second kind) cannot be constructed, while others – heat engines with efficiencies approaching that of the Carnot cycle arbitrarily closely – can. One day, quantum constructor theory will likewise embody principles of nature which express the fact that certain types of information processing (say, the computation of non-Turing-computable functions of integers) cannot be realised in any technology while others (the construction of universal quantum computers with arbitrary accuracy) can. Just as the quantum theory of computation is now *the* theory of computation – the previous theory developed by Turing and others being merely a limiting case – so the present theory of computation will one day be understood as a special case of quantum constructor theory, valid in the limit where we ignore all issues of hardware practicability. As Einstein (1920) said, "There could be no fairer destiny for any physical theory than that it should point the way to a more comprehensive theory in which it lives on as a limiting case".

References

Bekenstein, J.D., 1981, *Phys. Rev.* **D23**(2), 287-98.
Cirac, J. I., and Zoller, P., 1995, *Phys. Rev. Lett.* **74** 4091-4
Deutsch, D., 1986, *Found. Phys.* **16**(6), 565-72.
Deutsch, D. 1989, *Proc. R. Soc. Lond.* **A425** 1868.
Deutsch, D., 2002a, 'The Structure of the Multiverse' *Proc R Soc Lond.* (to appear).
Deutsch, D., 2002b, in *Proceedings of the Sixth International Conference on Quantum Communication, Measurement and Computing*, Shapiro, J.H. and Hirota, O., eds, Rinton Press, Princeton, NJ.
Deutsch, D. Barenco, A. and Ekert, A., 1995, *Proc. R. Soc. Lond.* **A449** 669-77
Deutsch, D., Ekert, A. and Luppachini, R., 2000, Bull. Symb. Logic **3**, 3.
Deutsch, D., and Hayden, P., 2000, *Proc. R. Soc. Lond.* **A456** 1759-74.
Einstein, A., 1920, *Relativity: The Special and General Theory*. Ch. 22. (*Über die spezielle und die allgemeine Relativitätstheorie,* 1917.)
Fredkin, E. and Toffoli , T. , 1982, *Int. J. Theor. Physics*, **21** 219-53.

Gottesman, D. 1999, in *Group22: Proceedings of the XXII International Colloquium on Group Theoretical Methods in Physics*, Corney, S. P., Delbourgo, R. and Jarvis, P. D., eds. 32-43 International Press, Cambridge, MA.
Gottesman, D. and Chuang, I.L. 1999 *Nature* **402** 390-93.
Hartle J.B., 1991, *Phys. Rev.* **D44** 10, 3173.
Lloyd, S., 1995, Remark made at the Workshop on Quantum Computation, Villa Gualino, Torino, Italy.
Smolin, L., 1997, *The Life of the Cosmos,* Oxford University Press.
Steane, A., 1997, *Applied Physics* **B64**, 623.
Tegmark, M., 1997, Preprint gr-qc/9704009.
Tegmark, M., 2000, *Phys. Rev.* **E61** 4194-206.
Zurek, W. H., 1981, *Phys. Rev.* **D24** 1516-25.