

On Demons and Oracles

Alastair A Abbott, Cristian S Calude and Karl Svozil

1. Turing's Barrier

The concept of digital computation which emerged from the works of Church, Turing and Gödel is an important achievement of the last century. A large variety of mathematical models of computers and computations have been developed. Turing machines, lambda-calculus, combinatory logic, recursive functions, Markov algorithms, register machines are among the best known classical models. Newer models range from programming-oriented models including concurrent models like actor model and process calculi to quantum Turing machines, DNA computers, molecular computers, wetware computers and many others. A remarkable result was gradually proved: in spite of the apparent diversity, the computational capability of every model of computation is the same. All models are computationally equivalent. This strong mathematical evidence motivated a more general belief: the Turing model of computation is the right and most general concept for digital computation.

Turing proved that Hilbert's Entscheidungsproblem — the decision problem for the predicate calculus^a — is unsolvable by any Turing machine. Independently, Church obtained the same negative result by using his lambda-calculus, so by proposing to

define the notion ... of an effectively calculable function of positive integers by identifying it with the notion of a recursive function of positive integers (or of a lambda-definable function of positive integers)

he argued that Hilbert's Entscheidungsproblem is *unsolvable* (not only unsolvable by any lambda-definable function). Motivated by a similar identification proposed by Turing, and the (mathematical) equivalence between the sets of functions computed by Turing machines, lambda-definable

functions and recursive functions, Kleene introduced the Church–Turing Thesis:

A function of positive integers is effectively calculable only if recursive.

In the quest to give meaning to their negative solutions to Hilbert's Entscheidungsproblem, Turing and Church were interested in describing what humans could "in principle" compute, so originally the Church–Turing Thesis' main scope was purely mathematical. However, the Thesis itself is not a mathematical statement as one of the terms involved — the notion of an effectively calculable function — is not mathematically defined. In particular, the Thesis cannot be (mathematically) proved: it needs empirical verification, which can continue for ever, or a plausible refutation. This fact was recognised almost immediately by Post, who objected to its presentation as a definition because it "blinds us to the need of its continual verification".

In time the scope of the Church–Turing Thesis shifted towards a more general goal, the ultimate limits of digital computation. In this new context the Church–Turing Thesis can be stated as^b:

Every function of positive integers which can be computed in a physical system is recursive, or equivalently, it can be computed by a Turing machine.

In this form the Church–Turing Thesis is a statement about what can be computed in a system of physical laws. Two components are involved: (i) the mathematical component determines the dynamics of evolution of physical states into others and the relation between inputs and outputs, and (ii) the physical component determines which dynamics can be performed in the given system of physical laws. Once we fix a system of physical laws described mathematically, the corresponding Church–Turing Thesis becomes a well-posed mathematical question which can be

^aFind an effective method to determine whether an arbitrary formula of the predicate calculus system is provable in the system.

^bSometimes this is called the Physical Church–Turing Thesis.

mathematically investigated: it can be proved, disproved, or proved undecidable. Gödel, who was initially unconvinced by Church's argumentation, changed his mind after reading Turing's paper, suggested the idea of an axiomatic approach for the notion of "effective calculability" meant to capture its generally accepted properties. In this spirit Gandy proposed a programme where physical laws (like bounded velocity and finite density of information) are used to "prove" the Church–Turing Thesis; this approach was extended to quantum theory by Arrighi and Dowek. The Church–Turing Thesis has morphed into a class of Church–Turing Theses, each depending on the underlying system of physical laws; in some cases it is true, in some false.

The physical determination of the Church–Turing Thesis was rightly pointed out by Deutsch [1]:

The reason why we find it possible to construct, say, electronic calculators, and indeed why we can perform mental arithmetic, cannot be found in mathematics or logic. The reason is that the laws of physics "happen" to permit the existence of physical models for the operations of arithmetic such as addition, subtraction and multiplication.

However, omitting the mathematical component

Computers are physical objects, and computations are physical processes. What computers can or cannot compute is determined by the laws of physics and not by pure mathematics.

is wrong. The laws of physics can determine what dynamics can be performed in a given system of physical laws, but not what such dynamics "compute": this a mathematical issue. According to Timpson

We must recognise that their [mathematical determinants] place is *prior* to that of physical determinants.

Here is an example. Consider an undecidable problem, say the halting problem^c. In every system of physical laws the halting problem will

^cDoes there exist a Turing machine capable of deciding whether an arbitrary Turing machine halts on a given input?

be undecidable because of Turing's proof. This proof and its conclusion — the undecidability of the halting problem — tell us nothing about the system of physical laws, as no possible dynamics can be a solution to the problem: this is a mathematical fact true in any system of physical laws. Of course, the halting problem can be solved by other types of "machines": such a "solution" — obtained by mathematical or physical means — does not challenge the validity of Turing's result which concerns only the mathematical concept of Turing machine.

2. The Land of Hypercomputation

Hypercomputation studies models of computations in the hope of breaking the Turing barrier. By placing precise physical constraints on computations, hypercomputation contributes to the program of continuous verification of the Church–Turing Thesis suggested by Post.

The possibility of executing infinitely many "operations" in a finite amount of time is the core of many proposals. This idea is not new: Zeno's analysis of motion paved the way for the accelerated Turing machines featured in Sec. 3.

In 1939 Turing introduced the seminal notion of *oracle* Turing machine [2], a standard machine having access to an infinite sequence of bits — the oracle — coding answers to as many questions, and made this machine compute with finite approximations of the infinite oracle. If the oracle is computable the resulting computation is equivalent to a standard computation, but in case the oracle is incomputable the machine trespasses the Turing barrier. In the expert hands of recursion-theorists oracle Turing machines have been used to scrutinise the land of incomputable. The "crucial question", in the words of the mathematician M Davis [3], is:

Are there real physical processes that can be harnessed to do the work of Turing "oracles"?

Davis gives an unequivocally negative answer. This issue will be re-visited in Sec. 5.

Rewind to 1970: In a footnote to [4] (p. 143) the logician G Kreisel makes an astonishing suggestion: a collision problem related to the 3-body problem^d could be regarded as "an analogue com-

^dThe problem of predicting the motion of a group of celestial objects that interact with each other gravitationally.

putation of a non-recursive function”, so an instance of hypercomputation. This possibility gets a new dimension with Xia’s [5] construction of no-collisions singularities in small Newtonian systems. Harnessing the incomputability identified in different physical systems becomes a possible source of hypercomputation.

Hypercomputation models have been constructed using neural networks, quantum mechanics, relativity theory, inductive Turing machines and many other ideas.

3. A Case Study: Accelerated Turing Machines

Centuries ago the ancient greek philosophers worried about the implications of an infinite divisibility of space and time. Zeno of Elea pointed out that motion itself would unexist, since the slightest finite movement would require an infinity of actions. Subsequently, differential calculus suggests to formally overcome these issues by taking the finite differential quotient of spatial and temporal change.

The revival of these ancient ideas came with computation. Weyl noted the potentiality to “complete” infinite computations in finite proper physical time as follows ([6], pp. 41–42):

Yet, if the segment of length 1 really consists of infinitely many subsegments of lengths $1/2, 1/4, 1/8, \dots$, as of ‘chopped-off’ wholes, then it is incompatible with the character of the infinite as the ‘incompletable’ [...and] there is no reason why a machine should not be capable of completing an infinite sequence of distinct acts of decision within a finite amount of time; by supplying the first result after $1/2$ minute, the second after another $1/4$ minute, the third $1/8$ minute later than the second, etc. In this way it would be possible [...] to achieve a traversal of all natural numbers and thereby a sure yes-or-no decision regarding any existential question about natural numbers!

Such a device — termed an accelerated Turing machine, reflecting that the rate of computation accelerates over its computational period — goes beyond the Church–Turing barrier. Questions such as the halting problem can be solved as

the infinity of computational steps performed by a non-halting computation are performed within a finite period of time. This power comes at an interesting cost however: for such machines to hypercompute they must by necessity use an infinite amount of space.

As noted earlier, the mathematical model of an accelerated machine undergoing these infinite dynamics must be supplemented by a search for physical implementations. A growing body of research has been exploring this; several proposals exist to harness infinite divisibility of space and time, and to utilise physical processes for the construction of such infinity machines. Some of these physically inspired proposals involve, for example, investigating the state of a lamp with ever decreasing switching cycles, and several ultrarelativistic methods put observers in “fast orbits” to exploit relativistic differences in time or throw them toward black holes where space-time behaves strangely.

4. Statistical Physics

Hypercomputation is also relevant in statistical physics — where information plays a key role — and can add to the dialogue on Maxwell’s infamous paradoxical demon. Here, the ability to hypercompute yields the potential to improve, if not reverse, energy dissipation.

Maxwell’s demon is supposedly capable of separating lower-energy from higher-energy particles in an isolated box divided into two chambers. Traditionally, the demon accomplishes this by controlling a shutter in the dividing barrier, only opening it for particles with suitable velocity. In this way a temperature gradient might be produced which would, at least in the long run, contradict the second law of thermodynamics (in the form that no process exists whose sole purpose is the transformation of heat into work).

The contemporary “exorcism” of the demon and resolution of the paradox seems to presuppose that nature behaves reversibly — that is, the evolution of physical microstates is one-to-one. Thus, every separation or contraction of microstates in one region of phase space has to be accompanied by an equal compensating amount of mixing or expansion. In particular, any sorting action of the demon, associated with a decrease of entropy of the rest of the system, should be

compensated by an increase of information in the demon's "mind" memory which is at least as large as the entropy decrease caused by the demon.

But what if the demon's memory and computational capacity is, at least in principle, unbounded? This case corresponds to certain types of hypercomputability: here the contraction in physical configuration or phase space could go on forever at the price of consuming more and more memory of the demon, thereby realising a sort of *Hilbert's Hotel* scenario.

Furthermore, if the demon's capacity to compress information is unbounded, would it be possible to "compute" the (classical incomputable) algorithmic information content of the information acquired? Present literature postulates that only the optimal compression yields optimally small compensation on the demon's side; other less optimal compressions result in an overall increase in entropy. Thus, in order to attain optimal performance, hypercomputational abilities of the demon must be assumed.

5. Quantum Oracles

Progress in the natural sciences during the 20th century was marked by two distinct departures from classical thought: Einstein's theories of relativity and the theory of quantum mechanics. Quantum mechanics has primarily been explored in the realm of computation as a medium for the alternative computational model of quantum computing, but there are possibilities of attaining hypercomputational power through the more subtle approach of considering quantum oracles instead of quantum computational models.

The proposal for a quantum oracle which we present here is based around the heart of non-classicality in quantum mechanics: the measurement process. While there are strong results relating to the impossibility of a classical description of measurement which lie at the centre of our argument, there are several competing interpretations of the ontological structure of quantum systems which alter the way measurement is viewed. Making explicit the assumptions we rely on is critical as the oracle is not independent of these.

In quantum mechanical theory even simple systems can exist in states that are superpositions of other states, for which any attempt to measure the state will yield one of the possible

outcomes seemingly at random. Formally, the theory only describes the probability distribution of this process; the fact that, after measurement, a subsequent measurement of the state will yield the same result seems to indicate that the measurement process irreversibly changes the state of the system at random. The nature of this "state collapse" is outside the theory and in the realm of "interpretations", of which many exist [7]: the standard Copenhagen interpretation(s), the de Broglie-Bohm theory, the many-worlds interpretation of Everett, and many more exotic ones.

A natural interpretation of these state of affairs, and one argued early on by Einstein, Podolsky and Rosen (EPR) [8], is that

the description of reality as given by a [quantum mechanical] wave function is not complete

in that the result of a measurement is not probabilistic but in fact determined by some unknown, yet pre-existing, "element of physical reality". However, the failures of a classical, deterministic viewpoint to account for the predictions of quantum mechanics are exemplified by the "no-go" results of Bell [9] and Kochen and Specker [10]. Bell's results show the impossibility of any hidden variable theory to reproduce the statistical predictions of quantum mechanics under the assumption of locality. However, of more interest is the result of Kochen and Specker applicable to individual quanta. The Kochen-Specker Theorem proves that it is impossible to assign pre-existing values to the outcomes of measurements under the conditions of (i) *value indefiniteness*: all observables, even those which are not compatible (cannot be simultaneously measured) have definite values corresponding the result of a measurement of them; and (ii) *non-contextuality*: the value corresponding to the result of a measurement does not depend on which other compatible measurement are made alongside of it.

In a bid to maintain realism, Bell proposed that

the result of an observation may reasonably depend not only on the state of the system ...but also on the complete disposition of the apparatus.

Attempts to give complete, contextual, interpretations for quantum theory exist, such as the de Broglie-Bohm theory, and while such inter-

pretations reproduce certain quantum mechanical predictions, they must by necessity embrace non-locality and remain distinctly non-classical and counterintuitive.

If we interpret the Kochen–Specker Theorem as an evidence for Born’s proposal to “give up determinism in the world of atoms”, then we can construct a simple device acting as an incomputable oracle. The Kochen–Specker Theorem, however, does not give us a straight choice between contextual and indeterministic realities. Even if we choose to reject the notion of a contextual reality, the theorem does not exclude the possibility of partial-determinism; we may need to give it up only for some observables. However, the apparent co-ordinate independence of the proof makes such a situation rather implausible unless there is a fundamental asymmetry in the measurement process. Put bluntly, one can conceive of a *demon* possessing the observer and ensuring only those observables with definite values are ever measured; conversely, a *demon* could inhabit the state ensuring the observable we choose to measure is assigned a definite value, while those we do not measure are allowed to be indefinite. Such “super-deterministic” loopholes are known to exist in, and would invalidate tests of Bell inequalities. If we are to use a quantum system as an oracle, we must refuse to accept such demons from existing and conspiring to make our output be due to predetermined elements of reality. We hence choose to consider a complete departure from classical omniscience and only allow those observables in which the state was prepared as an eigenstate of to have definite values; elsewhere we have complete value indefiniteness.

Before we can construct our oracle we must make one final connection between value definiteness and computability by returning to EPR. Specifically, we ask what does it mean to be able to assign a definite value to a measurement outcome? According to EPR,

if, without in any way disturbing a system, we can predict with certainty the value of a physical quantity, then there exists an element of physical reality [hidden variable] corresponding to this physical quantity.

A definite value exists exactly when there is a value allowing us to predict exactly the result

of a measurement. Thus, if we repeat the state preparation and measurement process *ad infinitum* and the sequence produced by the concatenation of measurement outcomes is computable, then *every* measurement can be predicted with certainty and was thus of a value definite observable. This final assumption makes the important connection between computability and the classical notion of determinism that quantum mechanics appears to have abandoned.

From here it can be argued that, if each measurement is of a value indefinite observable (i.e. not of the observable which the state is in an eigenstate of) then the infinite sequence of results considered above must be incomputable. If it were computable, this would mean the measured observables were all value definite, contradicting the assertion of value indefiniteness everywhere.

Under these assumptions we are hence guaranteed that such a device would produce an incomputable sequence and act as an oracle. Such devices have in fact been considered for the use of random number generation, so perhaps Davis’ verdict on the existence of physical process that can be harnessed as Turing oracles was a little premature; one of the most plausible physical interpretation of the conundrums presented by quantum mechanics allows us to do just that.

It is useful to compare and contrast this envisaged quantum oracle to a hypothetical realisation of a probabilistic Turing machine — a Turing machine which chooses transitions probabilistically from some predefined computable probability distribution. If we consider a trivial such device which, regardless of the input, accepts with probability one-half, and rejects also with probability one-half, and consider the infinite sequence generated by running this machine on inputs $1, 2, \dots$ (where “accepting” on input i means the i th bit is 1), is this device different in any real respect to our quantum oracle? Naïvely it would seem not: both devices act as oracles where the i th bit is 1 with probability one-half. However, the sequences produced by this probabilistic oracle are only uniformly distributed — we cannot rule out the crucial probability-zero possibility of a computable sequence being produced. This probabilistic oracle would hence, in practice, be a Turing oracle with probability-one. While this may appear to contradict the well known Turing-equivalence of probabilistic Turing machines, we

note the crucial distinction that this device does not formally compute any sequence at all — we are simply envisaging a single output of an infinite run of it being used as an oracle. The existence of a physical realisation of such a probabilistic Turing machine is, of course, as difficult a problem to solve as that of a Turing oracle; we simply note that such a device is not the same as our proposed quantum oracle which is stronger in that it is unable to produce any computable output.

While our quantum oracles behave as oracles in the Turing sense, we know of no way of saying more about the set which they are an oracle for. The most important open question is: *what is the computational power of a Turing machine working with a quantum oracle?*

6. The Known, The Unknown, and The Unknowable

The body of every subject can be divided into three parts: the known, the unknown, and the unknowable. In time, the unknown shrinks, with some facts migrating to the known and the unknowable parts. The unknowable is the most problematic part as to prove that a condition is impossible one has to show that it implies a contradiction or an absurdity. A limit implies an impossibility, but the converse implication is false. Impossibilities are provable, hence objective; limits tend to be subjective and temporal. In contrast to mathematics where the triad is sharp and its poles — the known and the unknowable — are rather stable, the division fluctuates in science. Mathematical limits, like Gödel's

incompleteness theorem or incomputability results, cannot be automatically transferred to physics. Hypercomputation is a subject at the intersection of mathematics, computer science and various particular sciences, physics, chemistry, biology, so here impossibility and limits are difficult to obtain and tend to be temporal.

References

- [1] D. Deutsch, Quantum theory, the Church–Turing principle and the universal quantum computer, *Proceedings of the Royal Society of London. Series A, Mathematical and Physical Sciences (1934–1990)*, **400** (1985), pp. 97–117.
- [2] A. M. Turing, Systems of logic based on ordinals, *Proceedings of the London Mathematical Society, Series 2*, **45** (1939), pp. 161–228.
- [3] M. Davis, The myth of hypercomputation, in *Alan Turing: Life and Legacy of a Great Thinker*, C. Teuscher (ed.) (Springer, Berlin, 2004), pp. 195–212.
- [4] G. Kreisel, Church's thesis: a kind of reducibility axiom for constructive mathematics, in *Intuitionism and Proof Theory: Proceedings of the Summer Conference at Buffalo N.Y. 1968*, Studies in Logic and the Foundations of Mathematics, A. Kino, J. Myhill and R. E. Vesley, (eds.), **60** (North Holland, 1970), pp. 121–150.
- [5] Z. Xia, The existence of noncollision singularities in the n-body problem, *Annals of Mathematics*, **135** (1992) 411–468.
- [6] H. Weyl, *Philosophy of Mathematics and Natural Science* (Princeton University Press, Princeton, NJ, 1949).
- [7] J. A. Wheeler and W. H. Zurek, *Quantum Theory and Measurement* (Princeton University Press, Princeton, NJ, 1983).
- [8] A. Einstein, B. Podolsky and N. Rosen, Can quantum-mechanical description of physical reality be considered complete? *Physical Review*, **47** (1935) 777–780.
- [9] J. S. Bell, On the Einstein Podolsky Rosen paradox, *Physics*, **1** (1964) 195–200.
- [10] S. Kochen and E. P. Specker, The problem of hidden variables in quantum mechanics, *J. Mathematics and Mechanics*, **17** (1967) 59–87.



Alastair A Abbott

University of Auckland, New Zealand

Alastair Abbott is a Doctoral Candidate at the University of Auckland and the École Polytechnique, Paris. He has a background in theoretical computer science and quantum theory, and works in the interface between these two disciplines. His current research is looking at the role of algorithmic notions of randomness in physical processes which are intuitively seen as random such as quantum events and chaotic systems.



Cristian S Calude

University of Auckland, New Zealand
www.cs.auckland.ac.nz/~cristian

Cristian S Calude, is a mathematician and computer scientist based at the University of Auckland, New Zealand, where he is a chair professor and the founding director of the Centre for Discrete Mathematics and Theoretical Computer Science. His research includes theoretical and experimental work in computability, complexity, randomness and quantum theories. Calude is a member of the *Academia Europaea*.



Karl Svozil

Vienna University of Technology, Austria

After studying theoretical physics in Vienna and Heidelberg, Karl Svozil has held visiting positions at various academic organisations; among them UC Berkeley and the Lawrence Berkeley Lab in California, Moscow State University and the Lebedev Physical Institute, as well as the Centre for Discrete Mathematics and Theoretical Computer Science in Auckland, New Zealand. Svozil studies and teaches theoretical physics in Vienna.