

# Hypercomputation and the Physical Church-Turing Thesis

Paolo Cotogno

---

## ABSTRACT

A version of the Church-Turing Thesis states that every effectively realizable physical system can be defined by Turing Machines ('Thesis P'); in this formulation the Thesis appears an empirical, more than a logico-mathematical, proposition. We review the main approaches to computation beyond Turing definability ('hypercomputation'): supertask, non-well-founded, analog, quantum, and retrocausal computation. These models depend on infinite computation, explicitly or implicitly, and appear physically implausible; moreover, even if infinite computation were realizable, the Halting Problem would not be affected. Therefore, Thesis P is not essentially different from the standard Church-Turing Thesis.

- 1 *Introduction*
  - 2 *Computability and incomputability*
  - 3 *The physical interpretation of the Church-Turing Thesis*
  - 4 *Supertasks and infinite computation*
  - 5 *Computation on non-well-founded domains*
  - 6 *Analog computation*
  - 7 *Quantum computation*
  - 8 *Retrocausal computation*
  - 9 *Conclusions*
- 

## 1 Introduction

The Church-Turing Thesis identifies the effectively computable functions with the recursive functions, or equivalently with the functions computable by Turing Machines; it is usual to consider the Thesis as a fundamental law of mathematics, true though not demonstrable. Some authors (e.g., Gandy [1980], Copeland [1997]) stress that what Turing actually did was to analyze computation as a process performed with paper and pencil by a human clerk (a 'computer', in Soare [1996]). Therefore, the Church-Turing Thesis should be neatly distinguished from the claim that whatever physical system can compute just recursive functions (Physical Church's Thesis, or 'Thesis P').

The interesting aspect is that Thesis P could be false without affecting the standard version: this would be the case if some physical systems were capable of computing functions that are not computable by Turing Machines, in principle. Recent developments of non-conventional computation theories seem to support this idea.

In this paper we discuss the principal approaches to computation beyond Turing definability, or *hypercomputation*;<sup>1</sup> our focus is on the general decision problem, while complexity is touched on only as a side topic. We shall argue that although some approaches are, at least potentially, very efficient, no model of hypercomputation can deliver effective results. Since there is no unified theory that explains everything and is accepted by everyone, some think that the present limitations may be overcome, possibly in the near future; we shall see, however, that even if hypercomputation were realizable it would not compute recursively unsolvable objects in principle, and therefore would not falsify the Church-Turing Thesis.

Section 2 recalls the notions of computable and incomputable function, with some emphasis on the latter. Section 3 presents the pure form of the Church-Turing Thesis and its interpretation as a physical hypothesis; we shall review some possible counterexamples and give some comments on physical incomputability. Section 4 presents the notion of a supertask, a process consisting of an actual infinity of steps, executed in finite time; this is a key point, since the theories of hypercomputation rest, directly or indirectly, upon some form of infinite computation. We shall see that supertasks are subject both to criticism of a physical character and to objections of a logical nature. Section 5 deals with computation on non-well-founded domains, a technique used in systems science to model interacting and self-referencing objects; we shall see that its hypothetical hypercomputing power depends essentially on supertasks and non-effective oracles. Section 6 deals with analog computation, often held to be intrinsically more powerful than its digital counterpart: we shall see that analog hypercomputation is a form of non-effective oracle computation, possibly depending on the implausible assumption of infinite precision. Section 7 is dedicated to quantum computation, the most promising approach to non-conventional computing at the present; we shall recall the main approaches to quantum hypercomputation, and we shall see that they depend either on infinite precision or on supertask computation. Section 8 deals with the retrocausal form of quantum computation; we shall see that even this speculative case can produce no violations of the Church-Turing Thesis. Section 9 presents some concluding remarks.

<sup>1</sup> In some contexts, 'hypercomputation' refers just to particularly powerful devices, such as massively parallel computers. Here we use the term in the ambitious sense popularized by Copeland and Proudfoot ([1999]), i.e. computation of non-recursive functions, 'beyond the Turing limit', as Siegelmann ([1995]) has it.

## 2 Computability and incomputability

**2.1** A function  $f(x) \simeq y$  is (*effectively*) *computable* when there is some method to produce explicitly the value  $y$  and check that  $y$  is the correct value; computations must produce a result, possibly the undefined value  $\perp$ , within a finite lapse of proper time after the beginning of the process—time being measured by elementary operation steps.

A function  $f$  is *incomputable* when there is no method to compute its values; it is essential to distinguish the cases where there is no such method as a matter of fact (*practical* incomputability) from those cases when there is a proof that  $f$  cannot be computed (*logical* incomputability) (see DeLong [1970], §28 for this Humean perspective on incomputability).

A function is (*general*) *recursive* when it can be defined and evaluated within the formalism of Gödel-Herbrand recursion equations, or any other universal programming system, such as the well-known Turing Machines (TMs).

**2.2** Logical incomputability plays the central role in the following discussion, so we give here a brief reminder. Let  $\psi_1, \psi_2, \psi_3, \dots$  be any enumeration of the computable functions expressed in some system  $\{\psi_i\}$ , capable of indexing. One defines then a binary function like:

$$2.2.1 \ f(x, y): \quad \begin{cases} \text{if } \psi_x(y) \neq \perp & \text{then } 1 \\ \text{if } \psi_x(y) = \perp & \text{then } 0 \end{cases}$$

and the *diagonal* monadic function:

$$2.2.2 \ g(x): \quad \begin{cases} \text{if } f(x, x) = 0 & \text{then } 0 \\ \text{if } f(x, x) = 1 & \text{then } \perp \end{cases}$$

This is incomputable for all  $x \geq 0$ : otherwise, there would be an index  $i$  such that  $g = \psi_i$ , and we would have  $\psi_i(i) = 0$  if and only if  $g(i) = f(i, i) = 0$  if and only if  $\psi_i(i) = \perp$ , a contradiction. Since all universal programming systems are recursively isomorphic by Rogers' Theorem, an equivalent argument can be formulated in any system: in particular, the TM version establishes the unsolvability of the *Halting Problem* (see e.g. Machtey and Young ([1978]) for details).

**2.3** A function  $f(x) \simeq y$  is *recursive in* another function  $\gamma$  when the computing process depends on  $\gamma$  at some point; in terms of TMs this is represented by an additional tape, called an *oracle*, where one can look at tabulated values for  $\gamma$ . This has a double aspect: if  $\gamma$  is assumed to be a function computable by 'unspecified means', as Turing had it, then the oracle defines a notion of *relative computability*. For the sake of generality, however,  $\gamma$  should be an

arbitrary function, and thus can be an object like the  $g(x)$  of 2.2.2: in this case the oracle defines rather a notion of *relative incomputability*, as any function  $f(x)$  depending on  $g$  would inherit its incomputability. The position of Copeland ([1997], [1998a], [1998b]), who views Turing's oracles as projects for future 'nonclassical' computers, is evidently based on the positive aspect, while we shall constantly refer to the negative one.

Holding that  $g(x)$  is computable by itself would yield one of the following alternatives:

- (\*) one output is chosen by convention
- (\*\*) both outputs are accepted

The latter requires an inconsistent system; the former yields indifference to the result of the computation, a form of randomness that remains untouched by probabilistic attempts at solving practical incomputability.

### 3 The physical interpretation of the Church-Turing Thesis

**3.1** The *pure* form of Church's Thesis is the claim that the syntactic notion of a recursive function is coincident with the semantic notion of an effectively computable function of positive integers (Church [1936]):

3.1.1 The computable functions are all and only the recursive functions.

Turing ([1936]) presented the first working description of the computation process needed to evaluate recursive functions, so the Thesis is often reformulated as the *Church-Turing Thesis*:

3.1.2 The computable functions are all and only those definable by TMs.

These propositions are first of all a definition of computable function; the Thesis can be viewed also as an explanation for the notion of computability (see, e.g., Mendelson [1964]), as an axiom for constructive mathematics (see, e.g., Kreisel [1970a]), and as a theorem of computability theory (see, e.g., Soare [1996]). Any other interpretation, such as the widespread view of the Thesis as a statement of computationalistic philosophy, goes beyond the letter of the original proposition.

**3.2** When Church and Turing advanced their theories, computation was chiefly a process of a mental nature, but the subsequent development of computing machinery has laid ground for viewing computation itself as a physical process. This has turned the Thesis into an empirical proposition, 'more like a proposition of natural science than of mathematics, [. . .] even though its subject matter appears to be mathematical in character' (Galton [1996], p. 138). A first 'physical' version of the Church-Turing thesis is:

3.2.1 The computable functions are all and only those computable by a discrete deterministic mechanical device.

This is the ‘Thesis P’ of Gandy ([1980]); see also Pitowsky and Shagrir ([forthcoming]). Deutsch ([1985]) has another statement, called the Church-Turing *Principle*:

3.2.2 Every finitely realizable physical system can be perfectly simulated by a universal model computing machine operating by finite means.

*Perfect simulation* here means input-output equivalence in an extensional sense, rather than some form of intensional, step-to-step, equivalence; a claim of this tenor is presented by Smith ([1999a]) as a ‘profound claim about the physical laws of our universe’. The difference with Gandy’s version is that ‘universal models’ of computing machines are not required to be classical mechanical systems, but may be finite-dimensional quantum-mechanical systems. Deutsch ([1985], p. 100) says that 3.2.2 is ‘both better defined and more physical than Turing’s own way of expressing it’; apparently, 3.2.1–2 are also more philosophical than the pure Church’s Thesis, since they include a general mechanistic view of the universe—‘a version of Laplacian determinism’, as Cooper ([1999]) points out. The physical computation community is specially interested in a stronger claim, the *Quantitative Church’s Thesis*:

3.2.3 Any physical computing device can be simulated by a TM in a number of steps polynomial in the resources used by the computing device.

Aharonov ([1998]) calls this the ‘modern’ form of the Thesis; it is intended to apply just to *probabilistic* TMs.<sup>2</sup>

**3.3** The status of these propositions is not exempt from criticism. Rocha ([1995], p. 65) remarks:

There is a circular argument in all the extensions of the Church-Turing Thesis beyond mathematics; the very first assumption that physically realizable dynamics represent effective procedures is just what the thesis aims at hypothesizing.

The risk of circularity is already known to affect the pure Thesis; as observed by Péter ([1959]), if one assumes that the construction of recursion equations is a recursive process itself, then Church’s Thesis may turn to infinite regression. One should be able to give ‘effective computability’ a meaning that does not coincide exactly with its recursion-theoretic version, and this is

<sup>2</sup> A probabilistic TM has random states that choose a new configuration with probability  $p \leq 1$ , at some point of the process; the sum of all probabilities associated with the possible choices is equal to 1. This allows more efficient computations in some classes of problems (notably, in primality testing), but no extension of the set of computable functions.

not so obvious. The specific danger of the physical interpretation is confusion between ontological and epistemological levels: this may become evident if we consider not only computability but also the opposite notion of incomputability, as possibly endowed with an interpretation in physical terms (see e.g. Cooper [1999], ‘Occurrent incomputability in Nature’).

**3.4** Pour-El and Richards ([1981]) have described a scalar wave equation in Euclidean space, such that its unique solution is incomputable although initial conditions are computable. Since continuous wave propagation is a reliably predictable phenomenon in classical physics, some have inferred that ‘wave computation’ would falsify Thesis P.

Weihrauch and Zhong ([2002]) show that the absence of effective results at time 1 in the Pour-El/Richards counterexample depends on the necessity of using an infinite set of conditions at time 0, to prepare the initial state of the system. Computability can be restored if wave equations are represented in a *Sobolev space*, a functional space where the norm involves derivatives; the problem of finding the solution of a three-dimensional wave equation with computable initial conditions is then shown to be effectively computable, ‘without any loss of regularity’, by Weihrauch’s generalized TMs. Weihrauch and Zhong interpret these results rigorously, and prudently, as the proof that there is no ‘obvious way’ to construct a super-Turing computer based on wave systems; in epistemological terms, the point is that incomputability is just a property of the mathematical way of representing physical systems.

**3.5** A simpler example is the classic *N-body problem* of Newton’s mechanics: this consists in describing the time evolution of a system of  $N$  point-masses, for any integer  $N \geq 2$ , according to the laws of gravitation  $F = Gm_1m_2/r^2$  and acceleration  $F = ma$ , and assuming initial positions and velocities. The problem can be solved analytically only for  $N = 2$ , and for  $N = 3$  in the planar case, but not for larger systems, where approximation techniques should be used. This may be construed as an objection to the Church-Turing Thesis in the sense that  $N$ -body systems apparently are finite objects, yet they are not perfectly simulated by TMs, against 3.2.1–2. Now,  $N$ -body systems can be thought of as physical entities, but their Newtonian descriptions are just formal objects; incomputability arises because the involved functions deal with real numbers specified in infinite precision, rather than with rational numbers resulting from actual measures. Smith ([1993]) points out two ‘key facts’ that make the  $N$ -body problem incomputable: 1) point-masses can move arbitrarily close; 2) point-masses can arbitrarily increase their velocity. Both are just hypotheses: 1) may be replaced by different assumptions—e.g., that masses merge when they come close enough, and 2) is in contrast with the relativistic constancy of the speed of light. The classical formulation of

the  $N$ -body problem thus yields that finite physical systems can be mathematically described in incomputable ways; nothing prevents other formulations of the  $N$ -body problem, within more realistic laws of nature, from being computable by TMs in principle, in accordance with 3.2.1–2.

**3.6** Penrose ([2000]) describes another type of physical incomputability, not depending on the practical incomputability of chaotic behaviours. The idea is to use ‘toy universes’, where space is simply a bidimensional plane, time evolution is discrete, and each state at time  $n$  is described by a *polyomino* set—a generalization of domino configurations, collections of squares arranged in fixed shapes. The only action admitted in a Penrose toy universe is repeating polyomino configurations, possibly until the entire plane is tiled without overlapping or gaps:

The point is that, although its evolution is entirely deterministic [. . .] it is *non-computable*. It follows from a theorem of Robert Berger that there is no computer action which can simulate the evolution of the universe because there is no computational decision procedure for deciding when a polyomino set will tile the plane. (Penrose [2000], p. 119)

The chief difference from Newton’s  $N$ -body systems is that incomputability here is logical, rather than practical. Penrose’s argument, then, is that the universe suffers from incomputability *a fortiori*, being evidently more complicated than toy models, but the analogy is vague. Short of precise rules of correspondence between polyominoes and natural phenomena, however, the incomputability of tilings has no direct physical meaning; Berger’s undecidability theorem by itself is just a computation-theoretic result, stating that the polyomino formalism is powerful enough to be subject to the same restrictions as recursion equations or TMs.

**3.7** What these examples illustrate is that computability and incomputability are properties of the formal descriptions of physical systems, and that one should be very careful in separating the epistemological level from the ontological level. If this distinction is not given the role it deserves, the falsification of Church’s Thesis may become a rather ordinary affair, a consequence of the existence of physical processes that happen to possess no recursive description at the moment, though being evidently effective. One could obtain counterexamples at will, simply by considering the practical instructions of ‘mundane’, or ‘quotidian’ procedures (Cleland [1993], [2001]),<sup>3</sup> or unusual pieces such as the following ‘function’ (from Nelson [1987], p. 586):

<sup>3</sup> Curiously, the typical instances of ‘quotidian procedures’ are cooking recipes. Cleland ([2001]) argues that the recipe for—our dear—Caprese salad is literally as effective as numerical algorithms and TM-instructions.

$$3.7.1 f(x, r): \begin{cases} \mu y (y > x) \\ x = \text{'the number of hickory railway ties} \\ \text{laid in region } r \text{ in the United States between} \\ \text{January 1, 1887 and January 1, 1900'}/C \end{cases}$$

where  $r$  ranges on any enumeration of American regions, and  $C$  is—for some unexplained reason—Euler's constant. Instances of this genre bear little effect on the Church-Turing Thesis, since propositions 3.2.1–2 are not comparing physical objects to mathematical objects, but descriptions of physical systems to their computational definitions—that is, objects belonging to different formal systems.

The standard Thesis claims that all possible descriptions of computation are equivalent to TM definitions, hence that nothing can compute more functions than TMs: the physical versions raise the interesting question of whether physics can provide models of computation that are more powerful than TMs, in principle. Let us accept this perspective as a working hypothesis, and let us consider hypercomputation processes, which can effectively yield results in this universe—leaving complexity aspects aside—but cannot be defined and evaluated by TMs. We should thus ask: is hypercomputation realizable?

## 4 Supertasks and infinite computation

**4.1** The most direct approach to hypercomputation consists in renouncing the finiteness of computations; the basic notion is that of a *supertask*, an actually infinite sequence of operations carried out in a finite interval of time, without patent violations of physical laws. Supertask operations either happen in null time, or become infinitely fast by geometric compression; only the latter case is considered meaningful. Copeland ([1998c]) has a short history of supertask machines, first conceived at the beginning of the 20th century, even before TMs: these theoretical devices have been appropriately called 'Zeno Machines' after the ancient paradoxes of Zeno of Elea. Stewart ([1991]) christens them RAC, for 'Rapid Accelerating Computers'; Copeland ([1998a]) speaks instead of AUTM, for 'Accelerating Universal TM'; in a more prosaic fashion, Davies ([2001]) admits that they are rather a sublimation of 'brute force' computation. Svozil ([1998]) has a concise description: first, one distinguishes between two time scales  $\tau$  and  $t$ , where  $\tau$  is *proper time*, dense and measured within the observer's system, and  $t$  is *intrinsic time*, measured in discrete cycles of the given machine. Then one has to assume that any cycle of  $t$ , measured in terms of  $\tau$ , is squeezed by a factor  $k < 1$ , so that:



$$4.1.1 \quad \tau_t = \sum_{n=0}^t k^n - 1 = \frac{k(k^t - 1)}{k - 1}$$

Therefore, for  $t \rightarrow \infty$  the proper time  $\tau_\infty$  has the value  $k/(1 - k)$ ; the whole subject is based on this elementary mathematical fact, that some infinite series has a finite limit—and this, contrary to Zeno’s most famous paradox, allows Achilles to catch his Tortoise.

**4.2** Supertasks may be a ‘mere plaything of philosophers’, as Earman and Norton ([1998]) have it, but there is a more widespread interest in the perspective that these processes are realizable, in some appropriate physical form. A direct objection to the material realization of Zeno machines is that they would violate the ordinary stress-energy conditions; the usual reply (e.g., Copeland [1997], Copeland and Sylvan [1999]) is that Zeno machines are as physically unrealizable as standard TMs, on the grounds that TMs are also just idealized devices requiring infinite memories—trouble shared, trouble halved. One could thus think that, as TMs are implemented by digital computers after all, so Zeno machines could be realized by other computer models. The problem here is conflating the actual infinity necessary for *bona fide* supertasks with the fact that TMs have no fixed bounds on their tape length: in the former one considers results—hypothetically—obtainable only through a denumerable sequence of operations; in the latter the point is just that a TM with bounded memory might run out of tape, and thus fail to be universal.

**4.3** Classical physics offers a favourable environment for supertasks. For instance, Saari and Xia ([1995]) show that for  $N = 5$  the Newtonian  $N$ -body problem produces non-collision singularities where one point-mass is ejected infinitely far from the others in finite time—we may call this a *Newtonian supertask*. One might view results of this kind as a *reductio ad absurdum* for the immediate application of classical physics to reality, in the tradition of Zeno’s paradoxes, but others think differently: ‘it seems a reasonable conjecture that a suitable system of point masses, obeying Newtonian gravitation, can also simulate RAC,’ claims Stewart ([1991], p. 9); the same point appears also in Burgin ([1999], §5).

Another form of classical supertask is introduced by Davies ([2001]) through an infinite extension of Moore’s Law: the idea is considering machines capable of constructing smaller and faster replicas of themselves *ad infinitum*, within a universe where ‘matter may be subdivided more and more finely while retaining the same properties’—a Newtonian, noiseless universe. In more detail, for any  $i > 0$  a (mechanical, rather than electrical) device  $M_i$  working with internal clock  $c_i$  and a memory size  $m_i$  is assumed to be capable of building a new machine  $M_{i+1}$  working with clock  $c_{i/8}$  and memory size  $2m_{i-1}$ , which is ‘smaller, faster and more powerful than  $M_i$ ’. As for single

Zeno machines, the total computing time of an infinite array of these machines is the finite sum of a geometric series, and so is the total space they require. The key assumption is that  $M_i$  should start building its descendant  $M_{i+1}$  if it has produced no output on a given input  $P_i$  within the fixed span of  $2^i$  computation steps; the input is then passed unchanged to the descendant machine. The main advantage over traditional Zeno machines is that each  $M_i$  requires neither an infinite energy supply nor an infinitely resistant material, to counter objections like those hinted at above.

**4.5** General relativity theory has been considered as the framework of choice for physical supertasks by Pitowsky, Hogarth, and others (a full survey is Etesi and N emeti [2002]).<sup>4</sup> Briefly, this approach requires a division of labour between two observers A and B: A should carry out a computation in an infinite amount of proper time, while the second observer B should be in a position where his past light-cone contains A's entire world-line. B could thus access A's infinite computation steps in finite proper time: assuming that A can complete her infinite task, the central problem is modifying the Minkowski space-time so as to ensure that B can be informed of the result of A's infinite work. A solution is constituted by the metric of the Malament-Hogarth (MH) model, which admits observers whose past can contain entire future-directed, timelike half-curves of infinite proper length. Since MH observers are allowed to access infinitely long geodesics in a finite span of time, they should be in control of non-terminating computations, such as, say, the complete decimal expansion of  $\pi$ .

Earman and Norton ([1993]) have objections of physical character to the idea that this model allows the realization of supertasks. A first point is that receiving A's supertask output implicitly involves another supertask on B's part; actually, B should be capable of an infinite power of discrimination to tell A's signals from non-deterministic perturbations in the intended point  $p$ :

No matter how carefully and expansively we set up experiments—that is, no matter how large we choose our initial value hypersurface—we cannot prevent spurious signals from reaching  $p$  or coming arbitrarily close to  $p$ . (Earman and Norton [1993], p. 39)

Another point concerns the safety of observers: in the MH model, the space-time structure acts as an energy amplifier, so that the slightest amount

<sup>4</sup> The apotheosis of relativistic supertasks is Tipler's Omega-Point Theory, a speculation on the long-term evolution of the Universe towards the Big Crunch—the singularity hypothesized as final collapse. Tipler ([1995]) considers the whole Universe as becoming a single computer, and claims that the amount of information that it can process diverges to infinity on approaching 'Omega-Point', the God-like final state; the extreme deformation of space-time produced by the violence of shearing gravitational forces would provide the unlimited supply of energy necessary to power this cosmic Zeno machine. Lloyd ([2002]) considers the Universe as a single computer as well, but reaches completely different conclusions.

of thermal radiation can be indefinitely amplified, thus preventing the communication of information; moreover, as a very disturbing side-effect, this amplification could eventually vaporize the receiver.<sup>5</sup> Other objections are discussed by Pitowsky and Shagrir ([forthcoming], §2).

Etesi and Némethi ([2002]) claim that the Kerr-Newman space-time used for representing a slowly rotating black hole of small electric charge is actually MH, and develop the idea that a relativistic process capable of realizing infinite computations is traversing a Kerr black hole—such as the one possibly existing in the centre of our galaxy. In this no-return mission, observer A first prepares B for the computation, then plunges into the black hole; A's journey past the event horizon, in finite proper time, corresponds to infinite intrinsic time for B. Etési and Némethi accurately discuss conditions to ensure that A will not hit the singularity of the black hole, and will thus avoid being crushed by tidal waves 'in the stomach of the Kerr black hole'. Another problem concerns the transmission of information between A and B: A should receive B's signals before passing the Cauchy horizon. The assumption that time and space measurements have 'arbitrary accuracy' is necessary to counter the main objection of Earman and Norton, and this confirms that infinite precision is an ingredient of relativistic as well as of Newtonian supertasks.<sup>6</sup>

**4.6** Notwithstanding some support from classical and relativistic arguments, the plausibility of supertasks is spoiled, possibly in irredeemable ways, by objections coming from other branches of physics, in particular quantum mechanics and thermodynamics. A substantial body of studies indicates that any material system viewed as a computer is subject to quantitative limitations in terms of operations per second, memory space and communication speed rate. Lloyd ([2000]), building on previous results of Bremermann, Bekenstein, and others, shows that the laws of quantum mechanics limit the speed of any computing device to its energy. In particular, a system with average energy  $E$  can perform a maximum of  $2E/\pi\hbar$  operation

<sup>5</sup> In a classical universe, one could assume that the computer temperature converges towards zero as time increases to infinity, with energy proportionally decreasing (as pointed out by a referee). In a quantum and relativistic universe, however, computers would start being very dangerous as soon as they reach the energy threshold for producing electron-positron pairs, as positrons would annihilate with environmental electrons, with gamma rays as a by-product (Pakvasa *et al.* [1999]). Some authors are actually considering the evolution of computers towards the most extreme conditions known to physics, those of black holes (Lloyd ([2002]), Ng ([2001])).

<sup>6</sup> Another problem for black hole supertasks is that information could evaporate in the form of thermal radiation before it is communicated outside; the result of the computation would thus be lost for good in this universe, according to Hawking's theory (see Elitzur and Dolev [2000]). Otherwise, one could assume that information survives in a *remnant*, a stable point-like object capable of maintaining information in an infinite amount of quantum states after the black hole evaporation (see e.g. the *cornucopion* described by Banks [1995]). Even this hypothesis, however, does not ensure that data resulting from hypercomputation can be retrieved.

per second, where  $\hbar$  is Planck's constant; each operation should be realized by the most elementary physical event, the evolution of a single quantum state into an orthogonal state. In a similar way, the laws of thermodynamics limit the amount of the information, i.e. of the accessible physical states, that a computer can store and process. In this case the restrictive conclusion is that any material computing system can handle a maximum of  $S/k_B \ln 2$  bits, where  $S$  is the entropy of the system, and  $k_B$  is Boltzmann's constant. Lloyd ([2002]) pushes these bounds to the extreme, by considering the whole universe as a computer; his estimates are developed within the standard cosmological model, and produce huge, but in any case finite, values.

**4.7** Now let us concede, for argument's sake, that machines with infinite intrinsic time are realizable, in some physical context that we cannot yet foresee: apparently such machines would solve the cases of practical incomputability, but what could they do for logical incomputability? We know that diagonal functions cannot be properly computed on pain of inconsistency, but some are tempted by the ancient belief that infinity conquers the principle of non-contradiction (Nicholas of Cusa, *De Docta Ignorantia* I, 4): the conjecture is that a Zeno machine can solve the Halting Problem by executing infinitely many computation steps. This position appears, for instance, in Stewart ([1991], p. 8):

RAC can cram an infinite number of computations into a single second. It is indifferent to the algorithmic complexity of the task set to it; everything runs, not in exponential or polynomial time, but in *bounded* time. It can, therefore, solve the halting problem for Turing machines [. . .] by running a computation in accelerating time and throwing a particular switch if and only if the Turing machine halts.

The same claim is held also by Copeland ([1998a], [2000]), who assumes that in any enumeration  $\psi_1, \psi_2, \psi_3, \dots, \psi_n, \dots$  of all TMs, in a programming system  $\{\psi_i\}$ , the  $n$ th element corresponds either to a terminating TM or to a non-terminating TM. Then one should consider the characteristic sequence of the Halting Problem set for the given enumeration, a binary real number  $K_0$  that has a 1 at the  $n$ th place if the corresponding TM halts, and a 0 if it does not: a Zeno machine might thus solve the Halting Problem by scanning the whole  $K_0$  in finite time. In a comparable perspective, Davies ([2001]) can expect that uncomputable functions are computed by an actually infinite series of nano-machines: if the first machine  $M_1$  receives back no answer for a given input  $P$  after the time interval sufficient for generating infinitely many descendant machines, then one should be entitled to conclude that  $P$  is unsolvable.

**4.8** The critical point here is that the real  $K_0$  is assumed as given a priori, and this is just what the logical incomputability argument rules out: as Webb ([1983]) has it,  $g(x)$  and its numerous isomorphic formulations are the

‘guardian angels’ of Church’s Thesis.  $K_0$  obtains by running  $UTM$ , a universal Turing Machine, on input  $\ulcorner \psi_n \urcorner$  for all  $n > 0$ , but there is no way of deciding termination other than executing  $UTM$  ( $\ulcorner \psi_n \urcorner$ ) and waiting for the halt state, in the general case; therefore, one would still need a denumerable number of (non-accelerated) operations *before* any Zeno machine can start and do its infinitely fast job on  $K_0$  in finite time. Even if it is generated by an algorithmic process, the sequence  $K_0$  is not a computable number like  $\pi$  or  $e$ , because some  $n$ th digit of  $K_0$  represents an individual object with paradoxical properties—a diagonal function. Considering these as digits that ‘take infinitely long to compute’ (Stannett [2001], p. 20) begs the question:  $UTM$  processing of diagonal objects may be viewed as an endless loop, but is not recursive in geometric forms like 4.1.1. There is no convergence to a limit because step  $i + 1$  always destroys the result of step  $i$ , *ad infinitum*.

This might be a good chance for appreciating the intuitionistic caveat on the application of the excluded-middle law to infinite objects—in terms of Stewart’s description, one could say that the existence of the ‘halting switch’ is not certain everywhere—but the essential condition is classical consistency of the programming system  $\{\psi_i\}$ .

Actually, all forms of Zeno machines are no more immune to logical incomputability than standard TMs; this can be shown in a rigorous way by tailoring new unsolvable Halting Problems, through the usual techniques of diagonalization (for details, see Geroch and Hartle [1986], Svozil [1998]).

**4.9** A vantage point for seeing how deeply  $K_0$  is resistant to computational treatment is the *algorithmic information theory* developed by Kolmogorov, Solomonoff, Chaitin, and others. In that framework, one can show that although the information contained in  $K_0$  is redundant, since most instances of the Halting Problem are mutually recursive, the sequence can be transformed into a random number by considering the probability that each TM halts. Chaitin ([1975]) has introduced an incomputable real number defined as follows:

$$\Omega = \sum_{h \in A} 2^{-|h|}$$

where  $A$  is a denumerable set of binary strings  $h$  that encode TMs; for technical reasons these TMs must be *self-delimiting*, i.e. such that they include all their parameters and cannot be concatenated. Now,  $\Omega$  is an *incompressible* number, in the sense that the length of the shortest algorithm that generates  $\Omega$  is not smaller than  $|\Omega|$ . As Chaitin ([1990]) says,

each bit in base-two of this real number  $\Omega$  is an independent mathematical fact. To know whether that bit is 0 or 1 is an irreducible mathematical fact which cannot be compressed or reduced any further.

No method that can be expressed in standard mathematical theories—including, therefore, infinite computations—can by-pass the incompressibility of  $\Omega$ : Solovay ([2000]) has even devised a version of  $\Omega$  such that *ZFC* (Zermelo-Fraenkel set theory with Choice Axiom) cannot determine any single bit of the sequence, assuming a classical soundness theorem. This can be interpreted as a form of intrinsic randomness in the core of mathematics: yet, the halting probability  $\Omega$  is plainly recursive in the  $K_0$  defined for the same enumeration  $A$  (see Raatikainen [2000], §3), so this randomness turns out to be just a special way of looking at the unsolvability of the Halting Problem—as we should already know from 2.3(\*).

**4.10** The limitations induced by logical incomputability do not prevent supertasks from being very interesting subjects on their own, as they are in the theory of *infinite-time Turing Machines (ITTMs)* developed by Hamkins and Lewis ([2000]). ITTMs are the Zeno extension of ordinary TMs: they have input-output and work tapes divided into cells, and the sequences of computation steps are numbered by transfinite ordinals, to represent infinite intrinsic time. On reaching any ordinal  $\alpha$ , an ITTM continues the computation at stage  $\alpha + 1$ , even when  $\alpha$  is a limit ordinal; the idea is that in cases where an ordinary TM enters an endless loop, never reaching the halt state, an ITTM can pass through the loop, after the first  $\omega$  steps. This should solve the Halting Problem for TMs and make all arithmetical formulae ‘decidable’ in one strike:

4.10.1 The truth of any arithmetic statement is infinite-time decidable (Hamkins and Lewis ([2000], Theorem 2.1)

Since some true arithmetic statement happens to be built around diagonal functions, one should check this claim against the dilemma of §2: there is no reason to think of case (\*\*), inconsistency, so the oxymoron of infinite-time decidability must rest on case (\*), random choice. The relevant point is the behaviour of ITTMs at limit ordinal stages. The rules are as follows:

[I]f the values appearing in a cell have converged, that is, if they are either eventually 0 or eventually 1 before the limit stage, then the cell retains the limiting value at the limit stage. Otherwise, in the case that the cell values have alternated from 0 to 1 and back again unboundedly often, we make the limit cell value 1. (Hamkins and Lewis [2000], p. 570)

The third rule, the value 1 assigned by convention,<sup>7</sup> rather than transfinite ordinals by themselves, is what makes ITTMs capable of facing logical

<sup>7</sup> A historical precedent is the convergence of the Grandi series  $+1 - 1 + 1 - 1 + 1 - 1 \dots$ , with the dubious choice between 1 and 0 as limit values (see Rucker ([1983], p. 126). Guido Grandi (1671–1742) speculated that a series that can have both a null and a positive value could be the formula of divine creation *ex nihilo*. From the mathematical point of view, however, he reached the conclusion that the series converges to  $\frac{1}{2}$ , reasoning from  $1 - x + x^2 - x^3 + \dots = 1/(1+x)$ , almost an anticipation of fuzzy logic; both Leibniz and Euler agreed on this Solomonic value.

incomputability without becoming inconsistent. Proposition 4.10.1, therefore, does not expand our knowledge of arithmetical truth, unfortunately, but is instructive under other respects: it shows that the omniscience involved in infinite computation has a random seed, and exposes the supertask-like assumptions that may be implicit in claims of comparable strength, in higher recursion theory (see also Welch [2000]).

## 5 Computation on non-well-founded domains

**5.1** Another hypercomputation idea, found especially in some areas of systems science, is that the Church-Turing Thesis is falsified by computing systems interacting with a complex environment, as they would realize effective computations that are not represented by ordinary TMs. This position is held independently, and in various forms, by many authors, e.g., Bains *et al.* ([forthcoming]), Burgin ([1999], [2001]), Copeland and Sylvan ([1999]), Kamps ([1991], [1995]), Wegner ([1997], [1998]). Bains *et al.* ([forthcoming], §4) claim that in physical computation,

there are only two elements: the agent and the rest of the universe. [. . .]  
Nothing in the agent’s behaviour can be independent of its external environment, because the two systems are inextricably linked.

Therefore, the evolution of a physical computer, ‘taken in its entirety’, and dependent on the state of the universe, should be essentially superior to ordinary digital computation because it includes any phenomenon possibly affecting the system. This appeal to an undefined wholeness is certainly inspiring, but it is dubious that a comparison of computing power can be made with systems where one has no control on input-output flows.

Kamps has a more precise stand, holding that ‘mechanisms’ representable by TMs with a fixed set of states are overpowered by ‘*component systems*’, which ‘construct’, rather than compute, their states; the components these systems are made of acquire meaning only within the system itself, in a circular, self-referencing way. The argument can be thus formalized by considering functions that belong to their own domain: these objects were almost banned after Russell’s criticism of impredicative definitions as a source of inconsistency. In Von Neumann-Bernays-Gödel set theory (*NBG*), self-application is prevented by a dedicated postulate, the *restriction axiom*

$$(RA): \forall x(x \neq \emptyset \rightarrow \exists y(y \in x \wedge y \cap x = \emptyset))$$

(any non-empty class  $x$  contains an element  $y$  disjoint from  $x$ ). In *ZFC* there is an equivalent, somewhat more complicated, proposition called *foundation axiom* (*FA*), which forbids infinite descending  $\in$ -chains of the form:

$$5.1.1 \dots \in x_2 \in x_1 \in x_0$$

In function-theoretic terms, *FA* amounts to assuming that no function is contained in its own domain.

**5.2** The traditional view (see for instance Wittgenstein's *Tractatus*, 3.333) was that functions that do not respect *RA* are useless, but Løfgren ([1968]) pointed out that they can be of interest in mathematical biology, as formalizations of self-reproducing systems. Moreover, Løfgren noticed that since *RA* is independent from the other axioms of *NBG*, its negation  $\neg RA$  also is: one can thus give an axiomatization of self-reproduction including  $\neg RA$ , and admitting  $\in$ -chains of form 5.1.1. Kampis ([1995]) refers to this point as support for his claim that self-referencing component systems falsify the Church-Turing Thesis.

A thorough analysis of self-application has then been developed by *non-well-founded set theory*, a version of *ZFC* where *FA* is replaced by the *anti-foundation axiom* (*AFA*); this allows sets that are represented by graphs with cyclic subgraphs (Aczel [1988]). A non-well-founded set (also called a *hyperset*) can thus be a member of itself, as in the trivial case  $x = \{x\}$ . The main achievement of Aczel's approach is that self-reference, either in direct or in indirect form, does not cause the theory to be *ipso facto* inconsistent: *AFA* does not lead to Russell's paradoxical set  $\{x|x \notin x\}$ , as long as one maintains a suitable set/class distinction.

**5.3** Hypersets have an intrinsic interest as a non-standard approach to set theory; from the philosophical viewpoint, Rucker ([1977]) refers to non-well-foundedness as a concise formulation of Royce's concept that the Absolute is essentially self-representing, as in the map that contains an exact replica of itself (*The One, the Many, and the Infinite*). More concretely, hypersets can be considered as models for objects with all sorts of cyclic and recurrent behaviour: a nice example is the hall of mirrors, with a pair of facing mirrors  $x, y$  that reflect each other's reflection. If one describes the mirrors  $x, y$  by a pair of hypersets  $x = \{y\}$ ,  $y = \{x\}$ , then the mutual reflections (ignoring inversion) correspond to the infinite chain:

$$5.3.1 \quad y = \{x = \{y = \{x = \{y = \{ \dots \}}\}}\}$$

Wegner ([1997], [1998]) refers explicitly to non-well-foundedness as a formal way of describing interactive computing agents with infinite tape resources: the idea is still that TMs cannot model the behavior of a computing agent, no matter if human or mechanical, in dynamical interaction with an environment outside its control. Wegner's *persistent Turing Machines* (PTMs) aim to model interactive computation by extending usual TMs with a 'persistent' work tape—a tape that is not reinitialized after execution of sequential computations—and an infinite input tape to represent the continuing input streams. PTMs are intended as models of real-life processes



such as bank-account managers, airline reservation systems, even ‘driving home from work’. According to Wegner ([1997]), PTMs provide ‘well-defined mathematical and machine models that go beyond algorithms, allowing Church’s thesis to be extended’; the objection to Church’s Thesis is that it would not cover the interactive aspect of real-life computation.

**5.4** This stand is criticized by Goertzel, who points out that computations on *finite* hypersets can be simulated in principle by non-deterministic TMs: ‘After all, manipulations with finitely given hyperrelations are merely manipulations of finite graphs!’ (Goertzel [1994], §7.2.3). As long as PTMs operate in a discrete timescale, one can always establish finite time boundaries—as is evident in the ‘driving-home’ example. Therefore, any instance of interactive computation in a given time interval can be unfolded into a concatenation of a finite number of TM computations, each of which produces its output, halts, and returns control to some higher-level control function; similar considerations can be made for the ‘super-recursive algorithms’ of Burgin ([1999], [2001]), another description of non-deterministic interactive computation. A detailed criticism of Wegner’s stand is developed by Prasse and Rittgen ([1998]), who explain how interaction can be effectively reduced to a succession of subroutine calls.

In short, reflection 5.3.1 would escape recursive treatment only by our considering the sequence as an infinite set; non-well-foundedness could go beyond TM capabilities only by our considering an infinite chain 5.1.1. In fact, the ‘surprisingly easy’ proof (Wegner [1998]) of the irreducibility of interaction machines to TMs is but a consequence of the infinity of input tapes—allegedly necessary to model interaction between machine and environment. Now, if this infinity is viewed as a Gaussian *façon de parler*, we are still within TM definability, since standard TMs come with unbounded tapes. Otherwise, if one intends to view infinite tapes as actual Cantorian totalities, effective computability becomes problematic, and some form of oracle computation must be invoked.

We have already dealt with supertask oracles, but there can be stronger assumptions: in the admittedly ‘trivial’ argument of Copeland and Sylvan ([1999], §3) for *coupled* Turing Machines, i.e. TMs linked to their environment ‘via one or more input channels’, superiority over standard TMs would depend on the chance of receiving incomputable real numbers on their input tapes. The point overlaps with analog computation, as we shall now see.

## 6 Analog computation

**6.1** This approach is based on theoretical machines operating on continuous quantities, as opposed to the discrete data handled by ordinary digital

devices. Analog computers predate the digital age: a mechanical ‘differential analyzer’ was first described by Kelvin in 1876.<sup>8</sup> Today, they are of special interest in fields such as numerical analysis and computational geometry, where most methods are defined in terms of real numbers, and in the study of biological systems, where relevant dimensions are described in continuous terms. According to the ‘Manifesto’ of Blum *et al.* ([1996]), the digital computation performed by standard TMs is

fundamentally inadequate for giving [. . .] a foundation to the theory of modern scientific computation, where most of the algorithms—with origins in Newton, Euler, Gauss *et al.*—are *real number algorithms*.

In this field, at present, we have no result like Rogers’ Theorem for  $\Sigma_1$  computability, as there are several non-coincident approaches to computation on real numbers. For instance, Blum, Shub and Smale have developed a theory of *real* Turing Machines, now well known as the BSS model: variables and constants of these flowchart machines store real or complex numbers, rather than approximations; the operators are rational functions, ideally performed in constant time, even if the operands have no round-off. See Smale ([1990]) for a short introduction, and Meer and Michaux ([1997]) for a detailed survey.

**6.2** Siegelmann and Sontag ([1994], [1995]) present another renowned model for the analog computer, the *recurrent neural network* (ARNN) or *processor net*, a modern refinement of early studies by McCulloch and Pitts on artificial neurons. ARNNs are finite interconnections of  $n$  processors (‘neurons’), responding to  $m$  input lines, for some  $n, m > 0$ ; the connections can be cyclic, and feed values back into the net. Each neuron has an activation value, updated in discrete time  $t \geq 1$  according to a ‘sigmoid’ function of activation states and inputs at time  $t - 1$ , with a set of weights that can have integer, rational, or real values. Contrary to BSS machines, inputs and outputs of ARNNs in any case obey conditions of finiteness; according to Siegelmann ([1995]) this makes the model ‘more realistic’.

The so-called *hybrid* systems are constituted by a combination of intercommunicating discrete and continuous components: for instance, Alur *et al.* ([1995]) consider finite automata with discrete states but real—valued inputs—as is the case with common devices like thermostats and water-level monitors.

Other models derive from Shannon’s ‘General Purpose Analog Computer’ (GPAC), in turn a modern development of Kelvin’s machine, and have also a continuous time variable; in this case the net behaviour is described by

<sup>8</sup> Aharonov ([1998]) cites a beautiful example of the analog computer in architecture: the systems of ropes used by Antoni Gaudí to calculate and project the arcs of the neogothic cathedral *Sagrada Família* in Barcelona.

differential equations (see e.g. Campagnolo and Moore ([2000])). The GPAC model inspires also the theory of  $\mathfrak{R}$ -recursion presented by Moore ([1996]), which aims to extend the standard recursion-theoretic methods from the integers to a subset of real numbers. The elementary  $\mathfrak{R}$ -recursive functions are *composition*  $h(\vec{u}) = f(g(\vec{u}))$ , *integration*  $h = f + \int g$ , and  $\mu$ -*recursion*  $h(\vec{u}) = \mu y(f(\vec{u}, y))$ , for any real vector  $\vec{u}$  and any  $\mathfrak{R}$ -recursive functions  $f, g$ . These operators are defined as real-valued functions, and are computed in continuous time; integration takes the place of primitive recursion, and  $\mu$ -recursion is to find the infimum of the set  $\{y | f(\vec{u}, y) = 0\}$  on the real line.

**6.3** From the computational point of view, the relationship of digital and analog computers is symmetric. In one direction, analog computation is routinely simulated by digital means: although truly analog machines can be employed for special-purpose computations, BSS machines and other analog models are routinely implemented as software systems running on ordinary digital computers. The BSS model, in particular, is presented by Blum *et al.* ([1996]) as a mathematical theory of Fortran programming; although actual computations use just finite rationals, they can fill up a bounded set of reals ‘sufficiently densely that viewing the computer as manipulating real numbers is a reasonable idealization’.

In the other direction, analog computers of sufficient complexity do perform digital computation. For instance, Moore ([1990]) describes a model for the analog computer based on *mirror systems*—billiards with particle-like balls bouncing between a finite number of planar and parabolic surfaces according to the laws of reflection. This computer is shown to be universal, capable of simulating all TMs, and therefore subject to undecidability; Moore’s aim is also to show that the behavior of simple physical systems is undecidable ‘even if the initial conditions are known exactly’—in our terms, that classical (descriptions of) physical systems are subject to logical incomputability, not only to practical incomputability, as we know from Penrose’s tilings. Similar results hold for ARNNs: Siegelmann and Sontag ([1995]) show that an ARNN of sufficient dimension with rational weights can simulate all TMs, and is thus universal. These results, besides clarifying what analog computers can do in principle, give a basis for comparing the strength of different theories of analog computation, also in terms of their space and time efficiency.

**6.4** The most common view is that discrete approximations with any desired degree of accuracy can perfectly simulate computations on real-numbered quantities. For instance, Rosen explains that the differential equations describing any physical computing system can be treated by relaxation techniques and reduced to difference equations; numerical solutions can then be obtained by TM equivalents, so that ‘there will be no physically detectable

difference between the behavior of the system as predicted by the approximate solution and the actually observed behavior of the system' (Rosen [1962], p. 389; see also Rubel [1989] for a mathematical treatment, and Toffoli [1999] for an engineering perspective). Vergis *et al.* ([1986], p. 92) have the following claim as the analog version of Church's Thesis:

6.4.1 Any finite analog computation can be perfectly simulated by a Turing Machine.

More precisely, this is called 'Weak Church's Thesis', while 3.2.3 is dubbed 'Strong Church's Thesis'. Bains and Johnson ([2000]) object that perfect simulation requires foreknowledge of the necessary degree of precision  $\epsilon$ , and this 'is not practical for intelligent systems designed to function in an unforeseeable universe'. Practicality may be the first concern for robot engineering, but it does not affect thesis 6.4.1: 'If we need more bits of precision, we can just supply them', as Bains *et al.* ([forthcoming]) acknowledge. If one is interested just in decision problems, it is sufficient to assume that  $\epsilon$  matches the lower bounds of any observable phenomenon in the given reference frame.

**6.5** In the field of non-conventional computation, there is also a different position, that computers with analog components are definitely more powerful than their digital counterparts:

although the Church-Turing Thesis is indeed a fundamental observation for a large class of discrete computing devices, it may not provide the whole picture, and the analog computation thesis is thus required. (Siegelmann [1995], p. 546)

Some admit that this superiority is merely hypothetical, and 'almost certainly unphysical' (Moore [1996], p. 23), but others (e.g., Kreisel [1970b], Doyle [1982], Copeland [1998b], Copeland and Sylvan [1999], and references therein) have conjectured the existence of physical processes described in analog terms that could be used to actually transcend Turing computability—we have already seen an instance with the Pour-El/Richards wave systems. The idea is that even if one cannot directly follow and verify the correctness of each single step of an analog computation, contrary to the digital case, correctness is ensured for the whole process, by the validity of the 'best theories' that govern the relevant physical dynamics (Pitowsky and Shagrir ([forthcoming])). Therefore, looking at extant physical theories, which are based on continuous quantities and do not involve finite precision bounds, one can expect that there are systems capable of realizing effective and correct computations, without the limitations of their digital counterparts. In this perspective, an analog machine should be more powerful than a digital one by being able to perform computations at *any* level of its internal

organization, although it has no definite programmable primitives—as in Leibniz’s concept of ‘divine automaton’ (*Monadology*, §§64–5).

For some analog models, the claim of hypercomputing power is explicit. For instance, Siegelmann and Sontag ([1994], p. 359) hold that ARNNs with real-valued (instead of rational-valued) weights are ‘super-Turing’, as the long-term behaviour of these systems could not be perfectly simulated by digital approximations; the hypercomputing capability admittedly depends on the ‘implicit use of infinite precision for internal computations’. Likewise, Bournez and Cosnard ([1996]) hypothesize that their version of Moore’s mirror systems is capable of computing functions that are not recursive; the necessary assumption is that the trajectories of particles between mirrors are continuous, thus allowing infinitely precise spatial discriminations.

**6.6** Strictly speaking, these claims do not violate the Weak Church’s Thesis, which refers to analog computation as a finite process, but touch the closely related question: does infinite precision possess an operational meaning? Any object described in terms of Euclidean lines is a system endowed with an infinite density of information: the problem is whether one can actually retrieve and operate on this information—say as if we could obtain any digit of the decimal expansion of  $\sqrt{2}$  by measuring the diagonal of a perfectly square object. This perspective is ingrained in the classical space-time manifold, where quantities obey Leibniz’s *loi de continuité*; as we have seen in the context of Newtonian and relativistic supertasks, it can still be considered a reasonable principle (e.g., Copeland [1997], Stannett [2001]).

A general objection (e.g., Lokhorst [2000]) is the ‘indistinguishability argument’: even if classical physics permits this infinite precision in principle, one could never verify infinitely precise discriminations through empirical observations, which are necessarily finite. Stannett ([2001], §5) replies that the two aspects are separate: numerical representations of physical objects are always approximate, no matter whether irrational, rational, or even integer, numbers are used, but all digits of irrationals like  $\sqrt{2}$  should be physically real though we cannot measure them. This aims to posit that there is a purely physical way of doing what cannot be done algorithmically: the irony is that a purely mathematical argument, like the Pythagorean proof, is all one has to establish that some physical object realizes the infinite irrational sequence  $\sqrt{2}$ , independently of any measurement.

**6.7** Now, if the idea of extracting unbounded information from the continuous representation of a given material object is an empirical hypothesis, we should say that the development of physics is diminishing its plausibility in many ways. In particular, most approaches to the fusion of quantum mechanics and general relativity postulate some form of non-manifold space, and develop physical geometries with fuzzy or discrete

features, where the possibility of unbounded precision breaks down even in principle (see e.g. Butterfield and Isham [2001]).

From the specifically computational point of view, the chief objection is that measurement and computing operations cannot be physically performed without taking their energetic aspect into account. If all computations represent the time evolution of a physical system, then the quantitative limits of computers may be estimated in principle by considering the behaviour of the most fundamental space-time constructs: black holes. Bekenstein, in particular, has considered the entropy of the largest black hole fitting in any finite spherical region of space containing total energy  $E$ ; entropy is viewed as a measure of the maximum amount of bits that a system can hold, and information bits are materialized, as we have seen in §4.6, as distinguishable quantum states. The main conclusion is that the value of informational entropy cannot be larger than:

$$6.7.1 \quad 2\pi ER/\hbar c \ln 2$$

where  $R$  is the radius of the given sphere,  $\hbar$  is Planck's constant and  $c$  is the speed of light (Bekenstein [2001] has a narrative account).

Bound 6.7.1, so as other bounds we have recalled above, is fairly loose and depends on several hypotheses that can be discussed and criticized (see, e.g., Smith [1995], and Bains *et al.* [forthcoming]): adjustments and corrections are likely, but one cannot expect to remove the constraints altogether without making major changes in the laws of physics—for instance, without resorting to an arbitrarily variable velocity for light, or to some form of superluminal information transfer. The most important aspect for our topic is that these general bounds on information processing make no distinction between computer architectures, serial or parallel, digital or analog. In a sense, all computing processes are analog, from the physical point of view: digital computation is just a special process, where one selects distinguished values out of continuum, and takes them as Boolean values (McLennan [1993], Toffoli [1998]).

**6.8** One might notice that Siegelmann ([1995]), and Gavaldà and Siegelmann ([1997]) justify the claim that ARNNs can compute ‘beyond the Turing limit’ seemingly without recourse to infinite precision: hypercomputation would be ensured by the non-uniformity of real-valued analog computation. Briefly, a computation is *non-uniform* when it is performed by a TM augmented with an ‘advice’, that is an oracle assumed to evaluate a function  $\varphi(|x|)$  in constant or polynomial time, for any  $x$ . The notion of advice as a function of the size of the input  $x$  has been introduced by computation theorists for defining complexity classes like  $C/Poly$ ; the relevant point here is that the function  $\varphi$  is arbitrary, and might be not only a practically but also a logically

incomputable one. The astute move is therefore to use a complexity-based relation to introduce a non-effective oracle, as an undercover source of hypercomputation power. This is not a specific feature of analog machines, however; standard TMs with oracles for incomputable functions can turn into hypercomputers as well, as Copeland ([1998c]) painstakingly explains.

**6.9** For completeness, let us consider analog computation also from an idealized point of view, abstracting completely from thermodynamic bounds and round-off problems in measurement. The superiority over TMs would apparently consist in using mathematical methods that involve actual infinities of operations: for instance, the analog *X-machines* defined by Stannett ([1990]) on a continuous topological space are proven to ‘decide’ the Halting Problem—for ‘a Turing program that generates the same configuration infinitely often’—just through a denumerable conjunction of predicate instances.

In the model of  $\aleph$ -recursion, computing on the subset of recursive reals that represent integers is declared ‘far more powerful’ than computing directly on integers: the Halting Problem should be ‘solved’ analytically by encoding TMs into linear maps on the real line (as shown by Moore [1990]). The search for halting states is then translated into an application of  $\mu$ -recursion on the real line  $(-\infty, \infty)$ ; after that, Moore ([1996]) invokes the ‘compression trick’, a mapping  $g$  of the real line into a compact interval  $(g(-\infty), g(\infty))$ , which is finite and thus allegedly decidable. We already know that the finite interval  $(0, 1)$  contains intractable objects like Chaitin and Solovay’s  $\Omega$ , therefore decidability requires also a stronger ‘trick’, the patching of  $\mu$ -recursion into a *total* zero-finding function  $\eta$ :

$$\eta_Y(f(\vec{u}, y)) : \begin{cases} \text{if } \exists y f(\vec{u}, y) = 0 \text{ then } 1 \\ \text{if } \exists y f(\vec{u}, y) \neq 0 \text{ then } 0 \end{cases}$$

Contrary to standard  $\mu$ -recursion, this function is defined so that it can always find a value, even after actual infinities of unsuccessful attempts: an ‘additional zero’ is strategically placed at the end of the search interval, something comparable to the conventional value assignment we have seen above for ITTMs. The method is really powerful, so much so that it can be used to prove that  $\aleph$ -recursion encompasses the whole Arithmetical Hierarchy  $\Delta^0_\omega$  (Moore [1996], Proposition 16); this result shows that  $\aleph$ -recursive computability is equivalent to TM computability augmented with an unbounded number of non-recursive oracles—a generalized form of Siegelmann’s point.

**6.10** In summary, nothing suggests that analog models, even in idealized form, can go ‘beyond the Turing limit’ any better than digital computers. Analog hypercomputation arguments appear but a variation of discrete

supertasks, and face the same objections: if reals are treated as digital expansions, then one could not consistently determine all digits of  $K_0$  by the presence of the ‘guardian angels’ of Church’s Thesis—which are just syntactic constructions, indifferent to the mathematical properties of the enumerated functions. On the other hand, if reals are viewed as proper ‘quantities’, instead of digital sequences, then continuous computations do not actually generate the digits of irrationals like  $K_0$ , but just give conventional or random approximations.

The crucial issue in analog computation is the treatment of randomness in input and data manipulation—in one word, of *noise*. Randomness is the price to pay whenever one deals with real-valued objects:

It is essential to model the inherent accuracy limitations of physical sensors that must be employed to ‘read out’ the solution of each instance, by assuming that each analog computer has an associate *absolute precision*  $\varepsilon$ . (Vergis *et al.* [1986] p. 94)

Bains and Johnson ([2000]) admit that the argument ‘gets to the heart of the controversy concerning super-Turing computation’, but object that ‘no-one has, so far, proven that noise removes all *non*-Turing abilities’ from analog computers. There are some pretty close results, however: for instance, Maass and Orponen [1998] find out that even a small amount of noise has the effect of reducing the computational power in a broad class of discrete-time analog computers, including BSSs and ARNNs. Maass and Sontag ([2000]) show that ARNNs where each gate is subject to noise have strong limitations—for instance, they cannot recognize a regular language as elementary as:

$$\{w \in \{0, 1\}^* \mid w \text{ begins with } 0\}$$

This happens even when the noise distribution is ‘known and of a rather benign type’, such as, say, a Gaussian distribution. In cases of practical interest, noise is likely to be more malign: the so-called *Newcomb-Borel paradox* indicates that mathematically defined real numbers differ significantly from the finite sets of numbers obtained by measuring Nature. This makes us conclude that the true concern with analog computers should be to ensure that implementations are reliable enough to maintain, rather than to surpass, the power of TMs.

## 7 Quantum computation

**7.1** This approach consists in a strengthening of probabilistic TMs through the mathematical methods of quantum mechanics. Berthiaume ([1997]), Shor ([1998]), Steane ([1998]), and Aharonov ([1998]) have detailed introductions; see also Smith ([1999a]) for historical remarks, Lomonaco ([2000]) for the



physical background, and Deutsch *et al.* ([2000]) for a philosophical perspective.

The original idea was to build special-purpose computers to reduce the exponential cost of simulating quantum-mechanical systems on standard machines (Feynman [1986]); at first, the time-reversibility of the unitary evolution of quantum systems appeared as a possible computational restriction, but then the theory of reversible computation reached the conclusion that there is no loss of computing power by considering only operations that can be systematically undone. The project evolved into a new paradigm of computation theory when Deutsch ([1985]) introduced the *quantum Turing Machine (QTM)*. This is a hybrid version of a TM, with digital input and output and internal states described in quantum-mechanical terms; other models, such as *quantum circuits* and *quantum cellular automata*, have also been proposed. Shor ([1997]) gave a substantial boost by presenting a new method, based on discrete logarithms and Fourier transforms, for factoring integers in polynomial time on a quantum computer—a task of paramount interest, for its potential use in breaking public-key cryptography. Another optimal quantum algorithm, also of practical interest, has been presented by Grover, for searching an unstructured database of  $N$  items in time  $O(\sqrt{N})$ .

**7.2** The QTM processes information in units called *qubits* (quantum bits), i.e. two-state quantum systems of the form:

$$7.2.1 \quad |\varphi\rangle = \alpha|0\rangle + \beta|1\rangle$$

where  $\alpha$  and  $\beta$  are complex numbers such that  $|\alpha|^2 + |\beta|^2 = 1$ . Any quantum system with orthogonal, and thus distinguishable, states could in principle represent a qubit: for instance, a spin- $\frac{1}{2}$  particle, with  $|0\rangle$  and  $|1\rangle$  corresponding respectively to spin-down and spin-up eigenstates. This allows defining computing machines where head and tape cells can be in any coherent superposition of the logical states  $\{0,1\}$ ; a measurement operation should then reduce superposition, and give a classical value as output. The superposition of states is in a sense the engine of quantum computing: the total state of classical systems composed of  $n$  subsystems, for any  $n > 0$ , contains information that grows linearly with  $n$ , but the superposition components of quantum systems grow exponentially with  $n$ . Thus,  $n$  qubits can store in principle  $2^n$  arguments, and linearly bounded resources allow the performance of computations on exponentially many arguments at a time; contrary to any classical architecture, this parallelism does not require spatially separated subsystems, but is intrinsic in the two-dimensional Hilbert space where qubit states are defined. In a sketchy way, the QTM can be compared to a probabilistic TM with the capability of following

simultaneously all computation paths in a single operation, instead of choosing one at random; the main differences are that probability amplitudes replace classical probabilities, and that computations are reversible, as required by the unitary operators of quantum mechanics.<sup>9</sup>

**7.3** An inherent limitation of the model is that the higher the number of superposed input states, the smaller the probability of measuring a particular value as actual output. If the evolution of quantum systems is interpreted as information processing, then only a small fraction of the computed information can be actually retrieved. Jozsa ([1999]) explains why quantum computation is nevertheless so promising:

the small amounts of information that are possible to obtain about the identity of the final processed state do *not* coincide with the particularly meagre kinds of information processing that can be achieved by classical computation on the input running for a similar length of time.

Shor's factoring algorithm has first shown how to make the 'computation paths giving the wrong answers interfere to cancel each other out, leaving a high probability of obtaining the right answer' (Shor [1998], p. 468). The mathematical pattern appears to be finding a certain subgroup in an Abelian group, but it is not known to what extent the method can be generalized, and which complexity classes can be dealt with successfully. A reasonable conjecture is that quantum algorithms are efficient for computational problems that have few or unique solutions—such as, e.g., *travelling salesman* problem instances where there is a unique optimal route, or a small number of equivalent optimal routes. These instances should have a better chance of exploiting interference: 'in the interference phase, the single solution could allow for a higher contrast and thus a better detection efficiency than a situation which would allow for many solutions' (Calude *et al.* [2000]).

**7.4** Returning to our chief question, let us ask: can quantum computers realize hypercomputation? Lloyd ([1993]) *proves* the negative answer by establishing a formal correspondence between the time evolution of a QTM and a standard TM computation, so that the unsolvability of the Halting Problem can apply, as for any other programming system. This is not considered conclusive, apparently because the Halting Problem is—unduly—read as a physical fact, instead of a logical proposition: the idea is that unsolvability is established only for idealized classical machines, and not for

<sup>9</sup> A superior feature of quantum-mechanical computers could be the capability of dealing with truly random numbers, as opposed to the usual pseudo-random generation obtained by algorithmic methods (Stannett [2001]). Pseudo-random number generation can be made arbitrarily close to quantum randomness, however, and no truly random sequence can be recognized by inspecting any *finite* segment of it; Knill ([1996]) concludes that quantum randomness does not add computational power over classical coin-flips.

the QTM, as idealization of a non-classical system. According to Deutsch ([1985]), the QTM should be not only universal but also possibly more powerful than conventional TMs, since it ‘can simulate various physical systems, real and theoretical, which are beyond the scope of the universal Turing Machine’.

That quantum wave systems have aspects which cannot be simulated by classical means is revealed in particular by *entanglement*, a correlation of joint states that can encode information, but which has no counterpart in classical preparations, in consequence of Bell’s inequalities (see, e.g., Penrose [1989], [1994] for an introduction). Entanglement is considered as a physical resource different from anything used in conventional computers, and is held to be a necessary ingredient of quantum computation (see, e.g., Linden and Popescu [1999], and Ekert and Jozsa [1998]). A striking aspect of the non-classical nature of entangled systems is their capability of performing counterfactual operations. A *counterfactual* is some observable phenomenon that might have happened according to given theoretical conditions, but did not happen as a matter of fact: any good theory allows this kind of conditional prediction as part of its capability to describe the world. Strikingly, the notion is not purely epistemological in quantum mechanics: that some event in a quantum-entangled system *could* have happened is sufficient to obtain information about it (Mitchison and Jozsa [2001]).<sup>10</sup> In a sense, counterfactual operations would allow us to run a quantum computer ‘for free’, without ever turning it on—something almost meaningless in classical terms.<sup>11</sup>

7.5 As for the classical analog case, a central issue is the reliability of computations in the presence of noise, with the important difference that the output of quantum computations can also be affected. Linden and Popescu ([1999]) show that if the noise level in the system is high enough, even entanglement produces no substantial speed-up in exponentially complex computations, and performance degrades to that of ordinary TMs. Entanglement can rapidly link the states 7.2.1 to their non-prepared environment, thus dispersing information in the rest of the universe, and making the system classical—a phenomenon called *decoherence*, arguably the

<sup>10</sup> A dramatic example is the Elitzur-Vaidman *bomb testing problem*. One has a stock of bombs, equipped with detonators connected to super-sensitive mirrors—one photon hitting the mirror suffices to cause an explosion. All bombs have an identical external appearance, but some are duds, with mirrors not connected to the detonator: the problem is to tell the ‘good’ bombs from the duds without detonating them. The task is impossible in classical terms, since the lightest contact with the mirror can cause the explosion, but the Elitzur-Vaidman protocol does have a quantum solution, employing a Mach-Zender interferometer to detect duds through destructive interference effects (see details and diagrams in Penrose [1994], §5.2, or Jozsa [1999]).

<sup>11</sup> Entanglement has other non-classical effects, such as *quantum teleportation* of qubits. In the *Gedankenexperiment* described by Brukner *et al.* ([2001]), a Bell-state measurement of entangled qubits could give, with a certain non-null probability, the result of an arbitrarily long quantum computation in no wait time.

main physical obstacle to the realization of working quantum computers, including counterfactual ones (Mitchison and Jozsa [2001]). Successful error-correcting strategies are thus essential for the realization of quantum computation: one can extend the methods developed in classical information theory, but new, quantum-specific principles are also required (see Gottesman ([2000]) for a survey).

**7.6** Computation enhanced by non-conventional methods is not genuine hypercomputation, however: for this one should explicitly attack the Halting Problem.<sup>12</sup> Some attempts aim to exploit the inherently probabilistic aspects of quantum mechanics: Calude *et al.* ([1999]) start from a ‘halting function’

$$7.6.1 \ H_U(P): \quad \begin{cases} \text{if } U \text{ halts on input } P \text{ then } 1 \\ \text{if } U \text{ loops on input } P \text{ then } 0 \end{cases}$$

where  $U$  is a standard UTM, and consider its computation through an Elitzur-Vaidman counterfactual device. This should be equipped with a classical one-bit register  $i$ , such that the ‘value of  $i$  can toggle between 0 and 1, depending on whether the computation of  $U$  on  $P$  is progressing toward the halting state’. As instance of  $P$ , the authors then use Radó’s *busy beaver* function  $\Sigma(n)$ , which has the property of growing faster than any other computable function of  $n$ , and expect a result with probability greater than  $\frac{1}{2}$ .

Calude *et al.* ([2000]) use another strategy, and define first a ‘halting qubit’:

$$7.6.2 \ |halt\rangle = c_h(t) |h\rangle + c_n(t)|n\rangle$$

where  $|h\rangle$  represents the halting state and  $|n\rangle$  the non-halting state, with  $c_h(t)$  and  $c_n(t)$  being their time-dependent amplitudes. Radó’s  $\Sigma(n)$  is then used again to obtain a worst-case analysis, and to argue that the probability  $|c_h(t)|^2$ , though extremely small, gives ‘a non-vanishing chance to obtain a solution of the halting problem in finite time’.

Calude *et al.* ([2002]) try trespassing the ‘Turing barrier’ through a semi-quantum method, tailored for the *infinite Merchant Problem* (the task of sorting a stack of  $n$  false coins out of an infinite sequence of stacks by a single weighting operations—false coins being heavier than genuine ones of a fixed value  $\delta > 0$ ), assumed as equivalent of the Halting Problem. The mathematical method, which should detect a solution with a ‘tiny but non-zero probability’, involves searching on truly random vectors in an infinite dimensional Hilbert space, within  $T > 0$  units of time. Negative search results are interpreted either as non-termination or as provenience of the test vector from a set  $F_{\epsilon, T}$ , of indistinguishable vectors, for a given precision-bound  $\epsilon$ .

<sup>12</sup> We refer to the proposition outlined in §2. In the literature on quantum computing there is also a quite different ‘halting problem’: this consists in ensuring that the ‘halt qubit’, designated to signal the termination of processing by a given eigenstate, can be inspected without affecting the computation. Ozawa ([2000]) has a solution, but the topic is still under discussion.

The technique aims to remove the latter chance by reducing the set  $F_{e,T}$  as  $T$  tends to infinity; this should determine ‘with a pre-established precision whether an arbitrary program halts or not’.

**7.7** It is not easy to tell what could be effectively achieved through the technical sophistication of these approaches. Calude *et al.* ([2000]) admit that hypercomputation will remain unreachable ‘for all practical purposes’; the reason is not worked out in detail, but we may presume that the involved timescale becomes too narrow for the desired output measurement—the same problem we know for relativistic supertasks. The chief objection in principle, however, is that probabilistic results can affect only practical incomputability, and not logical incomputability, where one is interested exclusively in the limit values.

As we know from the classical analog case, an avenue to hypercomputation starts from the assumption of infinitely precise discriminations. Nielsen ([1997], p. 2915) defines a quantum observable:

$$7.7.1 \hat{h} = \sum_{x=0}^{\infty} h(x)|x\rangle\langle x|$$

from the ‘halting function’  $h(x)$ , a version of 7.6.1; the  $|x\rangle$  of 7.7.1 is an orthonormal basis for the state space of some physical system with a countably infinite dimensional state space. Then, assuming that one can prepare the system in state  $|x\rangle$  for any  $x \geq 0$  and can apply some measurement apparatus to 7.7.1, obtaining the outcome  $h(x)$  with probability 1, or at least with ‘arbitrarily high confidence’, would amount to computing a recursively incomputable function. The argument is presented as a dichotomy: either some observables, such as  $\hat{h}$ , are not measurable in principle, or Church’s Thesis is false.

Ozawa ([1998]) observes that the critical assumption in Nielsen’s argument is indeed infinite-precision measurement: it is sufficient to consider definite bounds for measurement, for making computability of the halting function vanish, so that Church’s Thesis remains unaffected.

**7.8** Otherwise, hypercomputation can be sought directly through some quantum form of supertask. For instance, Kieu ([2001], [2002]) proposes to use the energetic status of the vacuum to attack Hilbert’s Tenth Problem, i.e. the decision problem for Diophantine equations in  $n$  variables, for any  $n > 0$ , which is known to be equivalent to the Halting Problem after the Matijasevich reduction (see e.g. Manin [1977], ch. 6). Kieu’s approach is based on a *Fock space*, a special Hilbert space used in quantum-field theory as a model of creation and annihilation of multi-particle systems: in this case the integer variables  $x_i$  of Diophantine equations are not represented by qubits,

but by a set of *creation operators*  $a_x^\dagger a_x$ , which have positive integer eigenvalues. The problem of finding the zeroes of a Diophantine polynomial  $D(x_1, \dots, x_n) = 0$  is first converted into the problem of finding the absolute minimum of its square, which is assumed to exist and be finite in any case (Kieu [2001], p. 5). This is in turn converted into the problem of realizing the ground state of a quantum Hamiltonian, involving just finite energy levels, on the assumption that the probability distributions used as criteria for verifying ground states can always be computed to the required precision. Tsirelson ([2001]) criticizes the method, and points out that it is but a form of non-effective oracle computation. The crucial point is that *all* values of the  $x_i$  need to be examined: finding a solution requires performing ‘an infinite search through the integers in a finite amount of time’, as Kieu himself admits, and thus a true supertask; we may also presume that Newtonian supertasks are hidden in the continuous interactions of particles allowed in a Fock space of infinite dimensionality.

**7.9** Quantum supertasks have also been considered by Svozil ([1998]), who wonders whether the Halting Problem can be solved by a *quantum Zeno machine*.<sup>13</sup> The difference from the classical case is that qubits are not directly affected by the self-reference of function 2.2.2: if the two qubits  $|1\rangle$  and  $|0\rangle$  represent the halting and the non-halting state, then a quantum Zeno machine that receives its own code as input can have a coherent superposition  $|1/\sqrt{2}, 1/\sqrt{2}\rangle$  of states as the acceptable fixpoint of an infinite sequence of operations. Now, this allows us to deal with the Halting Problem without producing inconsistency, but it is not a computational solution: the fixpoint is expressed in terms of probability amplitudes, which do not amount to observable probabilities. In consequence, no definite value could be extracted from the output:

This qubit solution corresponds to the statement that it is impossible for the agent to control the outcome; [. . .] the result of the computation appears to occur entirely at random. (Svozil [1998])

We can thus see that the quantum-mechanical approach again meets the logical dilemma between inconsistency and randomness.

**7.10** To summarize, logically incomputable functions appear as unsolvable for quantum computers as they are for classical machines. Quantum methods

<sup>13</sup> This—purely theoretical—construct should not be confused with Misra and Sudarshan’s *quantum Zeno effect*, where measurements performed on a quantum observable alter the dynamics of the system, in the limit case completely blocking its free evolution. This effect promises to be the most powerful way to improve the efficiency of the Elitzur-Vaidman ‘bomb-testing’ schema, and thus of counterfactual computation (see Jozsa [1999]). The connection with Zeno’s paradoxes is very loose, as Ropolyi and Szegedi ([1999]) point out, so it may be more precise to talk about the *watched-pot effect*, from the old saying ‘a watched pot never boils’.

have nevertheless the potential of becoming specially efficient: they might be the tool for solving some classes of exponentially difficult problems, provided that one succeeds in making systems composed of a realistic number of qubits work for a sufficiently long time, without being degraded by decoherence. The full depth of this quantum speed-up is not known at present; after the exploit of Shor's factoring algorithm, some have even conjectured  $NP = QP$ . Castagnoli ([1998]), for instance, hypothesizes quantum networks to solve 'with any desired confidence level' the SAT problem, which is  $NP$ -complete, like the travelling-salesman problem. The problem with this proposal is that measuring output observables should be done through an—admittedly speculative—notation of 'continuous Von Neumann measurement'.

Subsequent discussions have reduced some exceedingly optimistic hopes: an analysis of existing quantum algorithms leads Calude *et al.* ([2000]) to conjecture that most  $NP$  problems will remain intractable. Nielsen ([2000]) claims that if we could prove that the amount of entanglement in some class of quantum systems exceeds a certain lower bound, then we would have the long-sought demonstration that the  $P = NP?$  question has a negative answer.

As far as Church's Thesis is concerned, we observe that even a wholly successful realization of quantum computers could affect just the Quantitative form 3.2.3. Bernstein and Vazirani ([1997]) supply the 'first formal evidence' that 3.2.3 may be false, by showing that some oracle-computations can be performed in polynomial time by the QTM, and only in exponential time by a classical probabilistic TM. Even in these terms, however, the superiority over ordinary TMs remains conjectural. A formal proof that the same results *cannot* be obtained by arbitrary non-quantum algorithms would ultimately amount to a demonstration that  $P \neq NP$ ; in the unlikely—but still conceivable—case that  $P = NP$ , the speed-up obtainable by quantum computers on  $NP$  problems could be simulated by classical algorithms.<sup>14</sup>

## 8 Retrocausal computation

**8.1** Finally, let us deal with the most far-fetched route to hypercomputing power: computation with *retrocausal*, or *time-reversed* operations. In all models considered thus far, including reversible TMs, any computation path starts from some input state  $q_0$  and follows through a forward-ordered sequence of operations  $c_1 c_2 c_3 \dots c_j \dots$  until the halt state is possibly reached, for some  $j > 0$ . A time evolution of this kind can be viewed as the

<sup>14</sup> Arguments pointing in this direction are already available: according to Pitowsky ([forthcoming]) the same performance of quantum computers on the SAT problem can be obtained by standard probabilistic TMs. Battacharya *et al.* ([2002]) show that the performance of Grover's algorithm obtains also in a classical optical setting, by analyzing the interference patterns of a short laser pulse bouncing between two mirrors.

propagation of a signal at finite velocity within the computing system, with a natural interpretation in causal terms: like any other ordinary phenomenon, the future of a computation is determined by its past. Retrocausal computation is based on a wholly different idea, and hypothesizes computer circuitry with ‘negative time delay elements—which output signals that predict what their inputs will receive some fixed time later’ (Moravec [1992]); in other words, the computation path should proceed back in time from the future.

**8.2** Gödel presented a first formal theorization of time travel within a rotating universe; the space-time geometry of his relativistic model allows *closed timelike lines* (CTLs)—i.e., world-lines of objects that loop back to intersect themselves in finite time, by a major tilting of their light cone. Subsequently, other models of the universe and other gravitational structures allowing CTLs have been considered, such as *wormholes* connecting different parts of space-time (see Arntzenius and Maudlin [2000], and Earman and Smeenk [forthcoming] for surveys of relevant results).

From the logical point of view, the very idea of time travel is more than simply counter-intuitive, it is also exposed to the risk of inconsistent retrocausation: in the popular ‘grandfather paradox’, a time traveller encounters his own grandfather as a young man and shoots him dead, thus preventing the birth of his own father, and his own coming into existence. A special assumption, the *chronology protection* principle, may be invoked to avoid inconsistency: briefly, the idea is that retrocausation paradoxes can be blocked by suitable low-probability events—in the grandfather paradox, the gun could malfunction, the bullet could miss the target, the wound could heal, and so on, almost as for Leibniz’s pre-established harmony. Carlini *et al.* ([1995]) speak of a ‘self-consistency’ principle, and show that it can be naturally interpreted as a consequence of the *principle of minimal action*, well-known to physics since the 17th century.

**8.3** At the quantum level, time travel could be allowed naturally by minor fluctuations of the metric field; objects small enough to be sensitive to such fluctuations could tolerate these local, ‘zigzagging’, violations of the causality principle. The compatibility of CTLs with quantum systems is under discussion (see Earman [1995]), but is often considered plausible, as long as quantum laws are formulated in a time-symmetric way. Some physicists think that the difference between past and future should be viewed as an artefact in the quantum world, since nothing distinguishes the final states from the initial ones.

In Cramer’s *transactional interpretation* of quantum mechanics in particular, time symmetry is fundamental: any quantum event is described as a sort of ‘handshake’ in space-time between retarded and advanced waves,



‘a two-way contract between the future and the past for the purpose of transferring energy, momentum, etc’ (Cramer [1988]).

**8.4** Svozil ([1995]) and Stenger ([2001]) describe a similar EPR-Bell type protocol to explain how time-travelling quantum objects can avoid the inconsistent consequences of the grandfather paradox: unpolarized photons are placed in a pure quantum state by a horizontal polarizer  $H$ , and then sent to a circular polarizer  $C$ . This splits the pure state and emits two beams  $L$  and  $R$ , with opposite polarization. A detection apparatus should then intercept left-polarized beam  $L$ , and send back in time a signal to insert an absorber that blocks the unpolarized photon before it can reach  $C$ —the quantum analogue of grandfather-killing. The point is that the photon to be absorbed is in a coherent superposition of states  $|L, R\rangle$ , and there is no way to block just the  $L$  state; if the absorber was inserted anyway, the result would be random. The time-reversed signal would thus have the same effect as a ‘random shot’, without paradoxical consequences; any measurement performed to individuate the target would reduce the entangled state, and destroy the wanted information by decoherence.

**8.5** If retrocausal computation was possible, it would succeed in eliminating the random aspect of non-deterministic choices: by the chronology protection principle, Nature should produce some event  $e_0$ , with an extremely low probability  $p_0 \ll 1$ , to counteract the retrocausal process triggered by time reversal (see Moravec [1992], and Koshelev and Kreinovich [1997]). Let us consider a computing agent that accepts an approximate solution of some problem as input, and calculates an improved approximation as output. For a sufficiently large problem, one should settle for some approximation that can be arbitrarily far from the fixed point, assuming there is any, of the given input. Now, if the computer had time-loop components, the result of each approximation could be brought back in time to serve as input: hence, with iterations performed in reverse time, the optimal fixed point would be found literally at once, as soon as the machine receives its input and the ‘Start’ command—even better than counterfactual and teleportation processing. From the computational point of view, the low-probability events  $e_0$  triggered by causality violations are just the solutions of the given problem. If we consider, for instance, a  $NP$ -complete problem like SAT with  $n$  variables, the probability of a single satisfying assignment of truth values is  $2^{-n}$ ; hence, provided  $p_0 < 2^{-n}$ , it is more probable that Nature, forced by the time-loop computer, will supply the correct solution. This leads some researchers to consider time reversal in quantum computers as the avenue to obtain non-exponential methods for solving  $NP$ -problems. The idea is taken seriously by Castagnoli ([1998]); Koshelev and Kreinovich ([1997]) emphatically consider

retrocausal computers as ‘generation omega’, the ultimate frontier of computer science.

**8.6** The first and immediate objection is that any transfer of information backward in time is thermodynamically impossible, as the Second Law yields that the future cannot have lower entropy than the past. Even in the transactional interpretation, quantum operators characterizing separate measurements do commute, and there is no chance of superluminal communication; in Suarez’s EPR-Bell type experiments with time-like separated impacts, care is taken to ensure that there is no ‘superluminal telegraphing’ from the future (Suarez [1998]). The picture might change by adopting non-standard extensions or reformulations of quantum mechanics. For instance, in the approach of Aharonov *et al.* ([1990]), the principle of superposition is applied to the time evolution of quantum states: under special gravitational conditions, the superposition of time-evolving systems would allow, at least as a probabilistic effect, a ‘time translation’ machine. One can hypothesize that the negative time displacements obtainable in these conditions could be used to realize backward-in-time communications.

**8.7** Now let us assume, for argument’s sake, that retrocausal computers can actually be constructed in some physical setting. Efficiency would be beyond doubt, but let us ask once more whether something changes for the Church-Turing Thesis: could time-reversed computation solve the Halting Problem? A quotation from Penrose goes in the affirmative direction:

in the space-time geometries with closed timelike lines, a Turing-machine operation can feed on to its own output, running rounds indefinitely, if necessary, so that the answer to the question ‘does that computation ever stop’ has an actual influence on the final result of the quantum computation. (Penrose [1994], p. 383)

Unfortunately, though not surprisingly, even a TM augmented with CTL-operations is in no position to compute logically incomputable functions like 2.2.2. The diagonal dilemma is still in force: either we obtain inconsistency, or we have to settle for a random result. In the first case, if the value being sent back in time is a 0, then the present output should be  $\perp$ , and if the output sent back is a  $\perp$ , then the present output should be a 0, inconsistently. In Stenger-Svozil’s protocol, ‘diagonalization’ corresponds to removing the polarizer  $H$  from the apparatus, so to send an incoherent mixture of  $L$  and  $R$  states: in that case, the retrocausal signal from  $C$  would produce the blocking of  $L$ -photons destined to be emitted by  $C$  in the  $L$  beam, inconsistently. Otherwise, if states are superposed, we would have just a random answer, pointing in no steady direction; in particular, assuming that outputs consist of coherent

beams of opposite phase, the computing circuitry would find itself ‘perpetually in a dark fringe of an interference pattern’ (Moravec [1992]).

## 9 Conclusions

We have seen that non-conventional approaches to computation are but projects to obtain ever-increasing degrees of efficiency, but do not affect computability *in principle*: the functions that can be effectively evaluated are nowhere external to Turing definability.

We started with the most explicit approach, infinite computation: we saw that classical and relativistic physics allow us to think of computers capable of actually infinite operations in principle, but only within idealized conditions, and with *ad hoc* exclusion of quantum and thermodynamical constraints. Moreover, we saw that even hypothetical machines capable of performing supertasks could not compute the logically incomputable functions that define the boundaries of the Church-Turing Thesis. We saw then that interactive computation on non-well-founded domains can be simulated by ordinary TM processing, and has no hypercomputing aspect unless viewed in a non-effective way, as infinite computation. We then dealt with analog computation; we saw that the alleged superiority over digital processing depends essentially on infinitely precise discriminations, or again on supertask-like processes. We also noted that analog computation cannot be described realistically without accounting for the limiting effects of thermodynamical bounds. Passing to quantum computation, we stressed that the native parallelism of qubits could realize extraordinary efficiency, and that entanglement could allow feats like counterfactual computation. We also saw that quantum approaches to hypercomputation are based either implicitly on infinite precision, or explicitly on supertasks, as in the analog case. Eventually, we saw that ultimate speed-up might be allowed by retrocausal computation, but even this model cannot escape the ‘guardian angels’ of Church’s Thesis.

As regards hypercomputation, the conclusion is therefore negative: one may expect some progress in efficiency, but none in the enlargement of the set of effectively computable functions; this applies to all theories of computation, as long as they share the same Boolean background and are thus subject to the effects of logical incomputability. In particular, we argued that all ways of bypassing the Halting Problem within physical models reproduce the dichotomy between inconsistency and randomness, and are thus equally pointless from the computational perspective.

As regards Thesis P, the conclusion is that this statement is but a duplication of the standard Church-Turing Thesis. We reviewed the principal approaches to physical hypercomputation, but the epistemological reasons

outlined in Section 3 could suffice: we may agree that Turing's model draws its basic assumptions from classical physics, but it is not less true that physical theories are based on the same corpus of logical principles, embodied in any mathematical theory. Intuitionistic logic and other constructive systems do not make a substantial difference, as they intend to be more faithful to real-life mathematics than classical logic itself; even *quantum logic*, where the commutative and distributive laws have no universal validity, or other calculi with built-in quantum properties (e.g., the *matrix logic* of Stern [1988]), do not change the picture, since they are still constructed around standard logic and algebraic notions.

The idea that physics can discover hypercomputers somewhere in Nature forgets that the computational aspect of any present physical theory depends unavoidably on branches and refinements of algebra and calculus; any computing power beyond Turing definability should already be contained therein. Quantum algorithms may be quite different from classical ones, but in any case they are formulated within theories of Hilbert spaces and Hamilton operators, which can be formalized and proven correct in *ZFC* or comparable theories (Ozawa [1998]). The 'physical' systems that Thesis P concerns actually amount to their mathematical formulations; therefore, rejecting Thesis P would mean describing some effective computation technique capable of falsifying the standard Church's Thesis as well.

These conclusions, we finally observe, should not be construed as a proof that the Thesis is absolute truth: sensible objections have been raised in the past—see e.g. Péter ([1959]), Kreisel ([1970a]), and Ross ([1974])—and new ones may be found, but only within the logico-mathematical domain where the Church-Turing Thesis belongs.

### Acknowledgements

I would like to thank the referees for helpful comments on the earlier version of this paper. I am also grateful to Gert-Jan Lokhorst and to Seth Lloyd for some bibliographical indications, and to Itamar Pitowsky and to Klaus Weihrauch for letting me see their forthcoming papers on hypercomputation.

*via Corsica 301  
I-25125 Brescia  
Italy  
pc-logic@cydon.com*

### References

- Aczel, P. [1988]: *Non-Well-Founded Sets*, Stanford, CA: Center for the Study of Language and Information.

- Aharonov, D. [1998]: ‘Quantum Computation’, in D. Stauffer (ed.), *Annual Reviews of Computational Physics VI*, Singapore: World Scientific.
- Aharonov, Y., Anandan, J., Popescu, S. and Vaidman, L. [1990]: ‘Superpositions of time evolutions and a quantum time-translation machine’, *Physical Review Letters*, **64**, p. 2965.
- Alur, R., Courcoubetis, C., Halbwachs, N., Henzinger, T. A., Ho, P. H., Nicollin, X., Olivero, A., Sifakis, J. and Yovine, S. [1995]: ‘The Algorithmic Analysis of Hybrid Systems’, *Theoretical Computer Science*, **138**, pp. 3–34.
- Arntzenius, F. and Maudlin, T. [2000]: ‘Time Travel and Modern Physics’, in E. Zalta (ed.), *Stanford Encyclopedia of Philosophy*, <<http://plato.stanford.edu/entries/time-travel-phys/>>.
- Bains, S. and Johnson, J. [2000]: *Noise, physics and non-Turing computation*, 4th International Joint Conference on Information Systems, Atlantic City, preprint.
- Bains, S., Johnson, J. and Hall T. [forthcoming]: ‘Physical Computation and “Big” Artificial Intelligence’.
- Banks, T. [1995]: ‘Lectures on Black Holes and Information Loss’, *Nuclear Physics Proceedings, Supplement*, **41**, pp. 21–65.
- Battacharya, N., van Linden van den Heuvell, H. B. and Spreeuw, R. J. C. [2002]: ‘Implementation of Quantum Search Algorithm using Classical Fourier Optics’, *Physical Review Letters*, **88**, 137901.
- Bekenstein, J. D. [2001]: *Buchi neri, comunicazione, energia*, Roma: Di Renzo.
- Bernstein, E. and Vazirani, U. [1997]: ‘Quantum complexity theory’, *SIAM Journal on Computing*, **26**, pp. 1411–73.
- Berthiaume, A. [1997]: ‘Quantum Computation’, in L. A. Hemaspaandra and A. L. Selman (eds), *Complexity Theory Retrospective II*, New York: Springer, pp. 23–51.
- Blum, L., Cucker, F., Shub, M. and Smale, S. [1996]: ‘Complexity and Real Computation: A Manifesto’, *International Journal of Bifurcation and Chaos*, **6**, pp. 3–26.
- Bournez, O. and Cosnard, M. [1996]: *On The Computational Power And Super-Turing Capabilities Of Dynamical Systems*, NeuroCOLT Technical Report NC-TR-96-005.
- Brukner, C., Pan, J.-W., Simon, C., Weihs, G. and Zeilinger, A. [2001]: *Probabilistic Instantaneous Quantum Computation*, arXiv:quant-ph/0109022, preprint.
- Burgin, M. [1999]: ‘Super-recursive algorithms as a tool for high-performance computing’, *Proceedings of the High Performance Computing Symposium*, San Diego, CA, pp. 224–8.
- Burgin, M. [2001]: *Mathematical Models for Artificial Intelligence*, UCLA Department of Mathematics, Los Angeles, preprint.
- Butterfield, J. and Isham, C. J. [2001]: ‘Spacetime and the Philosophical Challenge of Quantum Gravity’, in C. Callender and N. Huggett (eds), *Physics Meets Philosophy at the Planck Scale*, Cambridge: Cambridge University Press.
- Calude, C., Campbell, D. J., Svozil, K. and Stefanescu, D. [1995]: ‘Strong Determinism vs. Computability’, in W. Depauli-Schimanovich, E. Koehler and F. Stadler (eds), *The Foundational Debate, Complexity and Constructivity in Mathematics and Physics*, Dordrecht: Kluwer, pp. 115–31.

- Calude, C., Dinneen, M. J. and Svozil, K. [1999]: *Counterfactual Effect, the Halting Problem, and the Busy Beaver Function*, CDMTCS-107, University of Auckland, New Zealand.
- Calude, C., Dinneen, M. J. and Svozil, K. [2000]: 'Reflections on Quantum Computing', *Complexity*, **6**, pp. 35–7.
- Calude, C. S. and Pavlov, B. [2002]: *Coins, Quantum Measurements, and Turing's Barrier*, CDMTCS-170, University of Auckland, New Zealand.
- Campagnolo, M. L. and Moore, C. [2000]: 'Upper and lower bounds on continuous-time computation', *Proceedings of the 2nd International Conference on Unconventional Models of Computation*, pp. 135–53.
- Carlini, A., Frolov, V. P., Mensky, M. B., Novikov, I. D. and Soleng, H. H. [1995]: *Time Machines: the Principle of Self-Consistency as a Consequence of the Principle of Minimal Action*, NORDITA preprint 95/49.
- Castagnoli, G. [1998]: *An Inherently Quantum Computation Paradigm: NP-complete = P Under the Hypothetical Notion of Continuous Uncomplete von Neumann Measurement*, quant-ph 9810017, preprint.
- Chaitin, G. J. [1975]: 'A theory of program size formally identical to information theory', *Journal of the ACM*, **22**, pp. 329–40.
- Chaitin, G. J. [1990]: 'Undecidability and Randomness in Pure Mathematics', in *Information, Randomness & Incompleteness*, Singapore: World Scientific, pp. 307–13.
- Church, A. [1936]: 'An unsolvable problem of elementary number theory', *American Journal of Mathematics*, **58**, pp. 345–63.
- Cleland, C. [1993]: 'Is the Church-Turing Thesis True?', *Minds and Machines*, **3**, pp. 283–312.
- Cleland, C. [2001]: 'Recipes, Algorithms, and Programs', *Minds and Machines*, **11**, pp. 219–37.
- Cooper, S. B. [1999]: 'Clockwork or Turing U/universe? Remarks on Causal Determinism and Computability', in S. B. Cooper, and J. K. Truss (eds), *Models of Computability*, Cambridge: Cambridge University Press, pp. 63–116.
- Copeland, B. J. [1997]: 'The Broad Conception of Computation', *American Behavioral Scientist*, **40**, pp. 690–716.
- Copeland, B. J. [1998a]: 'Super Turing-Machines', *Complexity*, **4**, pp. 30–2.
- Copeland, B. J. [1998b]: 'Turing's O-machines, Searle, Penrose, and the Brain', *Analysis*, **58**, pp. 128–38.
- Copeland, B. J. [1998c]: 'Even Turing Machines Can Compute Uncomputable Functions', in C. S. Calude, J. Casti, M. J. Dinneen (eds), *Unconventional Models of Computation*, Singapore: Springer, pp. 150–64.
- Copeland, B. J. [2000]: 'Narrow Versus Wide Mechanism: Including a Re-examination of Turing's Views on the Mind-Machine Issue', *The Journal of Philosophy*, **96**, pp. 5–32.
- Copeland, B. J., Proudfoot, D. [1999]: 'Alan Turing's Forgotten Ideas in Computer Science', *Scientific American*, **280**, pp. 99–103.
- Copeland, B. J. and Sylvan, R. [1999]: 'Beyond the Universal Turing Machine', *Australasian Journal of Philosophy*, **77**, pp. 46–66.

- Cramer, J. G. [1988]: 'An Overview of the Transactional Interpretation', *International Journal of Theoretical Physics*, **27**, p. 227.
- Davies, E. B. [2001]: 'Building Infinite Machines', *British Journal for the Philosophy of Science*, **52**, pp. 671–82.
- De Long, H. [1970]: *A Profile of Mathematical Logic*, Reading, MA: Addison-Wesley.
- Deutsch, D. [1985]: 'Quantum theory, the Church-Turing principle and the universal quantum computer', *Proceedings of the Royal Society*, **A 400**, pp. 97–117.
- Deutsch, D., Ekert, A. and Lupacchini, R. [2000]: 'Machines, Logic and Quantum Physics', *Bulletin of Symbolic Logic*, **3**, pp. 265–83.
- Doyle, J. [1982]: *What is Church's Thesis?*, Laboratory of Computer Science, MIT, Cambridge, MA.
- Earman, J. [1995]: 'Outlawing Time Machines: Chronology Protection Theorems', *British Journal for the Philosophy of Science*, **45**, pp. 5–19.
- Earman, J. and Norton, J. D. [1993]: 'Forever is a Day: Supertasks in Pitowsky and Malament-Hogarth Spacetimes', *Philosophy of Science*, **60**, pp. 22–42.
- Earman, J. and Norton, J. D. [1998]: 'Comments on Laraudogoitia's "Classical Particle Dynamics, Indeterminism and a Supertask"', *British Journal for the Philosophy of Science*, **49**, pp. 123–33.
- Earman, J. and Smeenk, C. [forthcoming]: 'Take a Ride on a Time Machine', in R. Jones, P. Ehrlich (eds), *Reverberations of the Shaky Game*, Chicago: University of Chicago Press.
- Ekert, A. and Jozsa, R. [1998]: *Quantum Algorithms: Entanglement Enhanced Information Processing*, quant-ph 9803072, preprint.
- Elitzur, A. C. and Dolev, S. [2000]: *Black Hole Evaporation Entails an Objective Passage of Time*, arXiv:quant-ph/0012081, preprint.
- Etesi, G. and Némethi, I. [2002]: 'Non-Turing computations via Malament-Hogarth space-times', *International Journal of Theoretical Physics*, **41**, pp. 341–70.
- Feynman, R. P. [1986]: 'Quantum Mechanical Computers', *Foundations of Physics*, **16**, pp. 507–31.
- Galton, A. [1996]: 'The Church-Turing Thesis: Its Nature and Status', in P. J. R. Millican and A. Clark (eds), *Machines and Thought*, Oxford: Clarendon Press, pp. 137–64.
- Gandy, R. [1980]: 'Church's Thesis and Principles for Mechanism', in J. Barwise, H. J. Keisler and K. Kunen (eds), *The Kleene Symposium*, Amsterdam: North-Holland, pp. 123–48.
- Gavaldà, R. and Siegelmann, H. [1997]: *Discontinuities in Recurrent Neural Networks*, NeuroCOLT Technical Report NC-TR-97-037.
- Geroch, R. and Hartle, J. B. [1986]: 'Computability and Physical Theories', *Foundations of Physics*, **16**, pp. 533–50.
- Goertzel, B. [1994]: *Chaotic Logic*, New York: Plenum Press.
- Gottesman, D. [2000]: 'An Introduction to Quantum Error Correction', *Proceedings of the Symposia of Applied Mathematics*, Providence, RI: AMS.
- Hamkins, J. D. and Lewis, A. [2000]: 'Infinite Time Turing Machines', *Journal of Symbolic Logic*, **65**, pp. 567–604.

- Jozsa, R. [1999]: ‘Quantum effects in algorithms’, *Chaos, Solitons and Fractals*, **10**, pp. 1657–64.
- Kalmár, L. [1959]: ‘An argument against the plausibility of Church’s Thesis’, in A. Heyting (ed.), *Constructivity in Mathematics*, Amsterdam: North-Holland, pp. 72–80.
- Kampis, G. [1991]: *Self-Modifying Systems In Biology And Cognitive Science*, Oxford: Pergamon Press.
- Kampis, G. [1995]: ‘Computability, Self-Reference, and Self-Amendment’, *Communications and Cognition—Artificial Intelligence*, **12**, pp. 91–109.
- Kieu, T. D. [2001]: *Quantum Algorithm for the Hilbert’s Tenth Problem*, arXiv:quant-ph/0110136, preprint.
- Kieu, T. D. [2002]: *Computing the noncomputable*, arXiv:quant-ph/0203034, preprint.
- Knill, E. [1996]: *Quantum Randomness and Nondeterminism*, LANL report LAUR–96–2186.
- Koshelev, M. and Kreinovich, V. [1997]: ‘Towards Computers of Generation Omega—Non-Equilibrium Thermodynamics, Granularity, and Acausal Processes: A Brief Survey’, *Proceedings of the International Conference on Intelligent Systems and Semiotics (ISAS ’97)*, Gaithersburg, MD: National Institute for Standards and Technology Publications, pp. 383–8.
- Kreisel, G. [1970a]: ‘Church’s Thesis: A Kind of Reducibility Axiom for Constructive Mathematics’, in A. Kino, J. Myhill and R. E. Vesley (eds), *Intuitionism and Proof Theory*, Amsterdam: North-Holland, pp. 121–50.
- Kreisel, G. [1970b]: ‘Hilbert’s Programme and the Search for Automatic Proof Procedures’, *Proceedings of the IRIA Symposium on Automatic Demonstration*, Berlin: Springer, pp. 128–46.
- Linden, N. and Popescu, P. A. [1999]: *Good Dynamics Versus Bad Kynematics. Is Entanglement Needed For Quantum Computation?*, quant-ph 9906008 v2, preprint.
- Lloyd, S. [1993]: ‘Quantum Computers and Uncomputability’, *Physical Review Letters*, **71**, pp. 943–6.
- Lloyd, S. [2000]: ‘Ultimate physical limits to computation’, *Nature*, **406**, pp. 1047–54.
- Lloyd, S. [2002]: ‘Computational capacity of the universe’, *Physical Review Letters*, **88**, 237901.
- Lokhorst, G.-J. C. [2000], ‘Why I Am Not a Super Turing Machine’, Workshop on Hypercomputation, London, University College.
- Lomonaco, S. J. Jr [2000]: ‘A Rosetta stone for quantum mechanics with an introduction to quantum computation’, Washington DC: AMS Short Course Lecture Note.
- Löfgren, L. [1968]: ‘An Axiomatic Explanation of Complete Self-Reproduction’, *Bulletin of Mathematical Biophysics*, **30**, pp. 415–25.
- Maass, W. and Orponen, P. [1998]: ‘On the effect of noise in discrete-time analog computations’, *Neural Computation*, **10**, pp. 1071–95.
- Maass, W. and Sontag, E. D. [2000]: *A Precise Characterization of the Class of Languages Recognized by Neural Nets under Gaussian and other Common Noise Distributions*, NeuroCOLT2 Technical Report NC2–TR–2000–63.
- Machtey, M. and Young, P. [1978]: *An Introduction to the General Theory of Algorithms*, New York: Elsevier.



- Manin, Yu. [1977]: *A Course in Mathematical Logic*, New York: Springer.
- McLennan, B. J. [1993]: ‘Grounding Analog Computers’, *Think*, **2**, pp. 48–51.
- Mendelson, E. [1964]: *Introduction to Mathematical Logic*, Princeton: Van Nostrand.
- Meer, K. and Michaux, C. [1997]: ‘A survey on real structural complexity theory’, *Bulletin of the Belgian Mathematical Society*, **4**, pp. 113–48.
- Mitchison, G. and Jozsa, R. [2001]: ‘Counterfactual computation’, *Proceedings of the Royal Society*, **A 457**, pp. 1175–94.
- Moore, C. [1990]: ‘Unpredictability and Undecidability in Dynamical Systems’, *Physical Review Letters*, **64**, pp. 2354–7.
- Moore, C. [1996]: ‘Recursion Theory on the Reals and Continuous-Time Computation’, *Theoretical Computer Science*, **162**, pp. 23–44.
- Moravec, H. [1992]: ‘Time Travel and Computing’, *Extropy*, **9**, pp. 15–20.
- Nelson, R. J. [1987]: ‘Church’s Thesis and Cognitive Science’, *Notre Dame Journal of Formal Logic*, **28**, pp. 581–614.
- Ng, Y. J. [2001]: ‘From Computation to Black Holes and Space-Time Foam’, *Physical Review Letters*, **86**, p. 2946.
- Nielsen, M. A. [1997]: ‘Computable functions, quantum measurements, and quantum dynamics’, *Physical Review Letters*, **79**, pp. 2915–18.
- Nielsen, M. A. [2000]: *On The Units Of Bipartite Entanglement: Is Sixteen Ounces Of Entanglement Always Equal To One Pound?*, arXiv:quant-ph/0011063, preprint.
- Ozawa, M. [1998]: *Measurability and Computability*, quant-ph 9809048, preprint.
- Ozawa, M. [2000]: ‘Quantum Turing machines: local transition, preparation, measurement, and halting’, in P. Kumer *et al* (eds), *Quantum Communication, Computing, and Measurement 2*, New York: Kluwer, pp. 241–8.
- Pakvasa, S., Simmons, W. and Tata, X. [1999]: *Relativistic Quantum Computing*, arXiv:quant-ph/9911091, preprint.
- Penrose, R. [1989]: *The Emperor’s New Mind*, Oxford: Oxford University Press.
- Penrose, R. [1994]: *Shadows of the Mind*, Oxford: Oxford University Press.
- Penrose, R. [2000]: *The Large, the Small and the Human Mind*, Cambridge: Canto/Cambridge University Press.
- Péter, R. [1959]: ‘Rekursivität und Konstruktivität’, in A. Heyting (ed.), *Constructivity in Mathematics*, Amsterdam: North-Holland, pp. 226–33.
- Pitowsky, I. [2002]: ‘Quantum Speed-up of Computations’, *Philosophy of Science*, **69**, pp. 168–77.
- Pitowsky, I. and Shagrir, O. [2003]: ‘Physical Hypercomputation and The Church-Turing Thesis’, *Minds and Machines*.
- Pour-El, M. B. and Richards, J. I. [1981]: ‘The Wave Equation with Computable Initial Data such that its Unique Solution is not Computable’, *Advances in Mathematics*, **39**, pp. 215–39.
- Prasse, M. and Rittgen, P. [1998]: ‘Why Church’s Thesis Still Holds’, *Computer Journal*, **41**, pp. 357–62.
- Raatikainen, P. [2000]: ‘Algorithmic Information Theory and Undecidability’, *Synthese*, **123**, pp. 217–25.
- Rocha, L. M. [1995]: ‘Artificial Semantically Closed Objects’, *Communications and Cognition—Artificial Intelligence*, **12**, pp. 63–89.

- Ropolyi, L. and Szegedi, P. [1999]: 'Zenonian arguments in quantum mechanics', in D. Greenberger, W. L. Reiter and N. Zeilinger (eds), *Epistemological and Experimental Perspectives in Quantum Mechanics*, Dordrecht: Kluwer, pp. 253–6.
- Rosen, R. [1962]: 'Church's Thesis and its relation to the concept of realizability in biology and physics', *Bulletin of Mathematical Biophysics*, **24**, pp. 375–93.
- Ross, D. [1974]: 'Church's Thesis: what its difficulties are and what are not', *Journal of Philosophy*, **71**, pp. 515–25.
- Rubel, L. A. [1989]: 'Digital simulation of analog computation and Church's Thesis', *Journal of Symbolic Logic*, **54**, pp. 1011–17.
- Rucker, R. [1977]: 'The One/Many Problem in the Foundations of Set Theory', in R. Gandy and M. Hyland (eds), *Logic Colloquium '76*, Amsterdam: North-Holland, pp. 567–93.
- Rucker, R. [1983]: *Infinity and the Mind*, New York: Bantam Books.
- Saari, D. G. and Xia, Z. J. [1995]: 'Off to Infinity in Finite Time', *Notices of the American Mathematical Society*, **42**, pp. 538–46.
- Shor, P. W. [1997]: 'Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer', *SIAM Journal of Computing*, **26**, pp. 1484–509.
- Shor, P. W. [1998]: 'Quantum Computing', *Documenta Mathematica*, ICM-I, pp. 467–86.
- Siegelmann, H. [1995]: 'Computation Beyond the Turing Limit', *Science*, **268**, pp. 545–8.
- Siegelmann, H. and Sontag, E. D. [1994]: 'Analog Computation via Neural Networks', *Theoretical Computer Science*, **131**, pp. 331–60.
- Siegelmann, H. and Sontag, E. D. [1995]: 'On The Computational Power of Neural Nets', *Journal of Computer and Systems Science*, **50**, pp. 132–50.
- Smale, S. [1990]: 'Some remarks on the foundations of numerical analysis', *SIAM Review*, **32**, pp. 211–20.
- Smith, W. D. [1993]: *Church's Thesis meets the N-body problem*, NEC Research Technical Report, Princeton.
- Smith, W. D. [1995]: *Fundamental physical limits on computation*, NEC Research Technical Report, Princeton.
- Smith, W. D. [1999a]: *Church's Thesis meets quantum mechanics*, NEC Research Technical Report, Princeton.
- Smith, W. D. [1999b]: *History of 'Church's Theses' and a manifesto on converting physics into a rigorous algorithmic discipline*, NEC Research Technical Report, Princeton.
- Soare, R. L. [1996]: 'Computability and Recursion', *Bulletin of Symbolic Logic*, **2**, pp. 284–321.
- Solovay, R. [2000]: 'A Version of  $\Omega$  for which ZFC can not Predict a Single Bit', in C. S. Calude and G. Paun (eds), *Finite Versus Infinite: Contributions to an Eternal Dilemma*, London: Springer, pp. 323–34.
- Stannett, M. [1990]: 'X-Machines and the Halting Problem: Building a Super-Turing Machine', *Formal Aspects of Computing*, **2**, pp. 331–41.

- Stannett, M. [2001]: *Computation and Hypercomputation*, Department of Computer Science, University of Sheffield.
- Steane, A. M. [1998]: 'Quantum Computing', *Reports on Progress in Physics*, **61**, pp. 117–73.
- Stenger, V. J. [2001]: 'Time's Arrows Point Both Ways. The View From Nowhen', *Skeptic*, **8**, pp. 92–5.
- Stern, A. [1988]: *Matrix Logic*, Amsterdam: North-Holland.
- Stewart, I. [1991]: 'The Dynamics of Impossible Devices', *Nonlinear Science Today*, **1**, pp. 8–9.
- Suarez, A. [1998]: *Does Quantum Mechanics imply influences acting backward in time in impact series experiments?*, quant-ph 9801061, preprint.
- Svozil, K. [1995]: 'On the computational power of physical systems, undecidability, the consistency of phenomena and the practical use of paradoxa', in D. M. Greenberger and A. Zeilinger (eds), *Fundamentals Problems in Quantum Theory*, New York: New York Academy of Sciences, pp. 834–41.
- Svozil, K. [1998]: 'The Church-Turing thesis as a guiding principle for physics', in C. Calude, J. Casti and M. J. Dinneen (eds), *Unconventional Models of Computation*, Singapore: Springer, pp. 371–85.
- Tipler, F. J. [1995]: *The Physics of Immortality: Modern Cosmology, God and the Resurrection of the Dead*, London: Macmillan.
- Toffoli, T. [1998]: 'Quo vadimus?—Much hard work is still needed', *Physica D: Nonlinear Phenomena*, **120**, pp. 1–11.
- Toffoli, T. [1999]: 'Non-Conventional Computers', in J. G. Webster (ed.), *Encyclopedia of Electrical and Electronics Engineering*, New York: Wiley & Sons.
- Tsirelson, B. [2001]: *The quantum algorithm of Kieu does not solve the Hilbert's tenth problem*, arXiv:quant-ph/0111009, preprint.
- Turing, A. M. [1936]: 'On computable numbers, with an application to the Entscheidungsproblem', *Proceedings of the London Mathematical Society*, **42**, pp. 230–65.
- Vergis, A., Steiglitz and Dickinson, B. [1986]: 'The complexity of analog computation', *Mathematics and Computers in Simulation*, **28**, pp. 91–113.
- Webb, J. [1983]: 'Gödel's Theorems and Church's Thesis: A Prologue to Mechanism', in R. S. Cohen and M. W. Wartofsky (eds), *Language, Logic and Method*, Dordrecht: Reidel, pp. 309–53.
- Wegner, P. [1997]: 'Why Interaction is More Powerful than Algorithms', *Communications of the ACM*, **40**, pp. 80–91.
- Wegner, P. [1998]: 'Interactive Foundations of Computing', *Theoretical Computer Science*, **192**, pp. 315–51.
- Weihrauch, K. and Zhong, N. [2002]: 'Is Wave Propagation Computable or can Wave Computers Beat the Turing Machine?', *Proceedings of the London Mathematical Society*, **85**, pp. 312–32.
- Welch, P. D. [2000]: 'The Length of Infinite Time Turing Machine Computations', *Bulletin of the London Mathematical Society*, **32**, pp. 129–36.