Physical Computability Theses

B. Jack Copeland & Oron Shagrir

Abstract: The Church-Turing thesis asserts that every effectively computable function is Turing computable. On the other hand, the *physical* Church-Turing Thesis (PCTT) concerns the computational power of physical systems, regardless of whether these perform effective computations. We distinguish three variants of PCTT – modest, bold and super-bold – and examine some objections to each. We highlight Itamar Pitowsky's contributions to the formulation of these three variants of PCTT, and discuss his insightful remarks regarding their validity. The distinction between the modest and bold variants was originally advanced by Piccinini (2011). The modest variant concerns the behavior of physical computing systems, while the bold variant is about the behavior of physical systems more generally. Both say that this behavior, when formulated in terms of some mathematical function, is Turing computable. We distinguish these two variants from a third – the super-bold variant – concerning decidability questions about the behavior of physical systems. This says, roughly, that every physical aspect of the behavior of physical systems – e.g., stability, periodicity – is decidable (i.e. Turing computable). We then examine some potential challenges to these three variants, drawn from relativity theory, quantum mechanics, and elsewhere. We conclude that all three variants are best viewed as open empirical hypotheses.

1. Introduction

The physical Church-Turing Thesis (PCTT) limits the behavior of physical systems to Turing computability. We will distinguish several versions of PCTT, and will discuss some possible empirical considerations against these. We give special emphasis to Itamar Pitowsky's contributions to the formulation of the physical Church-Turing Thesis, and to his remarks concerning its validity.

The important distinction between 'modest' and 'bold' variants of PCTT was noted by Gualtiero Piccinini (2011). Modest variants concern only computing systems, while bold variants concern the behavior of physical systems without restriction. The literature contains numerous examples of both modest and bold formulations of PCTT; e.g., bold formulations appear in

Deutsch (1985) and Wolfram (1985), and modest formulations in Gandy (1980) and Copeland (2000).

We will distinguish the modest and bold variants from a third variant of PCTT, which we term "super-bold". This variant goes beyond the other two in including decidability questions within its scope, saying, roughly, that every physical aspect of the behavior of any physical system – e.g., stability, periodicity – is Turing computable.

Once the distinction between the modest, bold, and super-bold variants is drawn, we will give three different formulations of PCTT: a modest version PCTT-M, a bold version PCTT-B, and a super-bold version PCTT-S. We will then review some potential challenges to these three versions, drawn from relativity theory, quantum mechanics, and elsewhere. We will conclude that all three are to be viewed as open empirical hypotheses.

2. Three physicality theses: Modest, Bold and Super-Bold

The issue of whether every aspect of the physical world is Turing computable was raised by several authors in the 1960s and 1970s, and the topic rose to prominence in the mid-1980s. In 1985, Stephan Wolfram formulated a thesis that he described as "a physical form of the Church-Turing hypothesis": this says that the universal Turing machine can simulate any physical system (1985: 735, 738). Wolfram put it as follows: "[U]niversal computers are as powerful in their computational capacities as any physically realizable system can be, so that they can simulate any physical system" (Wolfram 1985: 735). In the same year David Deutsch (who laid the foundations of quantum computation) formulated a principle that he also called "the physical version of the Church-Turing principle" (Deutsch 1985: 99). Other formulations were advanced by Earman (1986), Pour-El and Richards (1989), Blum, Cucker, Shub and Smale (1998) and others.

Pitowsky also formulated a version of CTTP, in his paper "The Physical Church Thesis and Physical Computational Complexity" (Pitowsky 1990), based on his 1987 lecture in the Eighth

Jerusalem Philosophical Encounter Workshop.[1] He said: “Wolfram has recently proposed a thesis — ‘a physical form of the Church-Turing thesis’ — which maintains, among other things, that no non-recursive function is physically computable” (1990: 86). The “other things” pertain to computational complexity: Pitowsky interpreted Wolfram as also claiming that the universal Turing machine *efficiently* simulates physical processes, and Pitowsky challenged this further contention (see also Pitowsky 1996; 2002). We will not discuss issues of computational complexity here (but see Copeland and Shagrir 2019 for some relevant discussion of the so-called "Extended Church-Turing Thesis").

Many have confused PCTT with the original Church-Turing Thesis, formulated by Alonzo Church (1936) and Alan Turing (1936); see Copeland (2017) for discussion of misunderstandings of the original thesis. It is now becoming better understood that, by ‘computation’, both Church and Turing meant a certain *human* activity, numerical computation; in their day, computation was done by rote-workers called "computers", or, more rarely, "computors" (see e.g. Turing 1947: 387, 391).Pitowsky correctly emphasized that PCTT and the original form of the thesis are very different:

> It should be noted that Wolfram's contention has nothing to do with the original Church thesis. By ‘every computable function is recursive,’ Church meant that the best analysis of our pre-analytic notion of ‘computation’ is provided by the precise notion of recursiveness. Indeed one sometimes refers to Church's thesis as ‘empirical,’ but the meaning of that statement, too, has nothing to do with physics. (Pitowsky 1990: 86)

We will use the term *physical* to refer to systems whose operations are in accord with the actual laws of nature. These include not only actually existing systems, but also idealized physical systems (systems that operate in some idealized conditions), and physically possible systems that do not actually exist, but that *could* exist, or will exist, or did exist, e.g., in the

[1] Some papers from the Workshop were published in 1990, in a special volume of *Iyyun*. The volume also contains papers by Avishai Margalit, Charles Parsons, Warren Goldfarb, William Tait, and Mark Steiner.

universe's first moments. (Of course, there is no consensus about exactly what counts as an idealized or possible physical system, but this is not our concern here.)

We start by formulating a modest version of PCTT:

> **Modest Physical Church-Turing Thesis (PCTT-M**) *Every function computed by any physical computing system is Turing computable.*

The functions referred to need not necessarily be defined over discrete values (e.g., integers). Many physical systems presumably operate on real-valued magnitudes; and the same is true of physical computers, e.g., analog computers. The nervous system too might have analog computing components. All this requires consideration of computability over the reals. The extension of Turing computability to real-valued domains (or non-denumerable domains more generally) was initiated by Turing, who talked about real computable numbers in his (1936). Definitions of real-valued computable functions have been provided by Grzegorczyk (1955, 1957), Lacombe (1955), Mazur (1963), Pour-El (1974), Pour-El and Richards (1989), Blum et al. (1988), and others. The definitions are related to one another but are not equivalent. The central idea behind the definitions is that a universal Turing machine can *approximate* (or simulate) the values of a function over the reals, to any degree of accuracy. We describe one of the definitions in Section 4.

Bold theses, on the other hand, omit the restriction to computing systems: they concern all (finite) physical systems, whether computing systems or not. Piccinini emphasized, correctly, that the bold versions proposed by different writers are often "logically independent of one another", and exhibit "lack of confluence" (2011: 747-748). The following bold thesis is based on the theses put forward independently by Wolfram and Deutsch (Wolfram 1985, Deutsch 1985):

> **Bold Physical Church-Turing Thesis (PCTT-B**) *Every finite physical system can be simulated to any specified degree of accuracy by a universal Turing machine.*

Pitowsky in fact interpreted Wolfram as advancing a modest version of the thesis, namely “that no non-recursive function is physically computable” (1990: 86). However, this is because Pitowsky was treating every physical process as computation; he said, for example: “According to this rather simple picture, the planets in their orbits may be conceived as ‘performing computations’” (1990: 84). Under this assumption, according to which all physical processes are computing processes, there is no difference at all between modest and bold versions. Piccinini, on the other hand, sensibly distinguished between computational and non-computational physical processes; he took it that the planets in their orbits do not perform computations. Against the backdrop of this distinction, both Wolfram's formulation and Deutsch's formulation are bold: they concern physical systems in general and not just computing systems.

In a recent paper, we introduced a new, stronger, form of PCTT, the "super-bold" form, here named PCTT-S (Copeland, Shagrir and Sprevak 2018). (The entailments between PCTT-S and PCTT-B and PCTT-M are: PCTT-S entails PCTT-B, and, since PCTT-B entails PCTT-M, PCTT-S also entails PCTT-M.) Unlike bold versions, the super-bold form concerns not only the ability of the universal Turing machine to simulate the behavior of physical systems (to any required degree of precision), but additionally concerns decidability questions about this behavior, questions that go beyond the simulation (or prediction) of behavior. Pitowsky (1996) provided some instructive examples of yes/no questions that reach beyond the simulation or prediction of behavior:

> There are, however, questions about the future that do not involve any specific time but refer to all the future. For example: ‘Is the solar system *stable?’*, ‘Is the motion of a given system, in a known initial state, *periodic?’* These are typical questions asked by physicists and involve (unbounded) quantification over time. Thus, the question of periodicity is: ‘Does there exist some *T* such that for all times t, $x_i(T+ t) = x_i(t)$ for i = 1, 2, ... *n?* Similarly the question concerning stability is: ‘Does there exist some *D* such that for all times *t* the maximal distance between the particles does not exceed *D?*’. (1996: 163)

The physical processes involved in these scenarios – the motion and stability of physical systems – may (so far as we know at present) be Turing computable, in the sense that the motions of the planets may admit of simulation by a Turing machine, to any required degree of accuracy. (Another way to put this is that possibly a Turing machine could be used to predict the locations of the planets at every specific moment.) Yet the answers to certain physical questions about physical systems – e.g., whether (under ideal conditions) the system's motion eventually terminates – may nevertheless be uncomputable. The situation is similar in the case of the universal Turing machine itself: the machine's behavior (consisting of the physical actions of the read/write head) is always Turing computable in the sense under discussion, since the behavior is produced by the Turing machine's program; yet the answers to some yes/no questions about the behavior, such as whether or not the machine halts given certain inputs, are not Turing computable. Undecidable questions also arise concerning the dynamics of cellular automata and many other idealized physical systems.

We express this form of the physical thesis as follows:

> **Super-Bold Physical Church-Turing Thesis (PCTT-S)**: *Every physical aspect of the behavior of any physical system is Turing computable (decidable)*.

Are these three physical versions of the Church-Turing Thesis true, or even well-evidenced? We discuss the modest version first.

3. Challenging the modest thesis: Relativistic computation

There have been several attempts to cook up idealized physical machines able to compute functions that no Turing machine can compute. Perhaps the most interesting of these are "supertask" machines—machines that complete infinitely many computational steps in a finite span of time. Among such machines are accelerating machines (Copeland 1998, Copeland and Shagrir 2011), shrinking machines (Davies 2001), and relativistic machines (Pitowsky 1990, Hogarth 1994, Andréka et al. 2009). Pitowsky proposed a relativistic machine in the 1987 lecture mentioned earlier. At the same time, Istvan Németi also proposed a relativistic machine, which we outline below.

The fundamental idea behind relativistic machines is intriguing: these machines operate in spacetime structures with the property that the entire endless lifetime of one participant is included in the finite chronological past of a second participant—sometimes called “the observer”. Thus the first participant could carry out an endless computation, such as calculating each digit of π, in what is a finite timespan from the observer's point of view, say one hour. Pitowsky described a setup with extreme acceleration that nevertheless functions in accordance with Special Relativity. His example is of a mathematician "dying to know whether Fermat's conjecture is true or false" (the conjecture was unproved back then). The mathematician takes a trip in a satellite orbiting the earth, while his students (and then their students, and then their students...) "examine Fermat's conjecture one case after another, that is, they take quadruples of natural numbers (x,y,z,n), with n≥3, and check on a conventional computer whether $x^n + y^n = z^n$" (Pitowsky 1990: 83).

Pitowsky suggested that similar set-ups could be replicated by spacetime structures in General Relativity (now sometimes called Malament-Hogarth spacetimes). Mark Hogarth (1994) pointed out the non-recursive computational powers of devices operating in these spacetimes. More recently, Etesi and Németi (2002), Hogarth (2004), Welch (2008), Button (2009), and Barrett and Aitken (2010) have further explored the computational powers of such devices, within and beyond the arithmetical hierarchy. In what follows, we describe a relativistic machine *RM* that arguably computes the halting function (we follow Shagrir and Pitowsky (2003)).

RM consists of a pair of communicating Turing machines T_A and T_B: T_A, the observer, is in motion relative to T_B, a universal machine. When the input (*m*,*n*)—asking whether the m^{th} Turing machine (in some enumeration of the Turing machines) halts or not, when started on input *n*—enters T_A, T_A first prints 0 (meaning "never halts") in its designated output cell and then transmits (*m*,*n*) to T_B. T_B simulates the computation performed by the m^{th} Turing machine when started on input *n* and sends a signal back to T_A if and only if the simulation terminates. If T_A receives a signal from T_B, it deletes the 0 it previously wrote in its output cell and writes 1 there instead (meaning "halts"). After one hour, T_A's output cell shows 1 if the m^{th} Turing machine halts on input *n* and shows 0 if the m^{th} machine does not halt on *n*. Since *RM* is able to do this for any input pair (*m*,*n*), *RM* will deliver any desired value of the halting function. There is further discussion of *RM* in Copeland and Shagrir (2007).

Here we turn to the question of whether *RM* is a counterexample to PCTT-M. This depends on whether *RM* is *physical* and on whether it really *computes* the halting function. First, is *RM* physical? Németi and his colleagues provide the most physically realistic construction, locating machines like *RM* in setups that include huge slowly rotating Kerr-type black holes (Andréka et al. 2018). They emphasize that the computation is physical in the sense that “the principles of quantum mechanics are not violated” and *RM* is “not in conflict with presently accepted scientific principles”; and they suggest that humans might “even build” a relativistic computer “sometime in the future” (Andréka, Németi and Németi 2009: 501). Naturally, all this is controversial. John Earman and John Norton pointed out that communication between the participants is not trivially achieved, due to extreme blue-shift effects, including the possibility of the signal destroying the receiving participant (Earman and Norton 1993). Subsequently, several potential solutions to this signaling problem have been proposed; see Etesi and Németi (2002), Németi and Dávid (2006) and Andréka et al. (2009: 508–9). An additional potential objection is that infinitary computation requires infinite memory, and so requires infinite computation space (Pitowsky 1990: 84). Another way of putting the objection is that the infinitary computation requires an unbounded amount of matter-energy, which seems to violate the basic principles of quantum gravity (Aaronson 2005)—although Németi and Dávid (2006) offer a proposed solution to this problem. We return to the infinite memory problem in a later section.

Second, does *RM compute* the halting function? The answer depends on what is included under the heading “physical computation”. We cannot even summarize here the diverse array of competing accounts of physical computation found in the current literature. But we can say that *RM* computes in the senses of “compute” staked out by several of these accounts: the semantic account (Shagrir 2006, Sprevak 2010), the mechanistic account (Copeland 2000, Miłkowski 2013, Fresco 2014, Piccinini 2015), the causal account (Chalmers 2011), and the BCC (broad conception of computation) account (Copeland 1997). According to all these accounts, *RM* is a counterexample to the modest thesis if *RM* is physical; at the very least, *RM* seems to show that non-Turing physical computation is logically possible. However, if computation is construed as the execution of an algorithm in the classical sense, then *RM* does not compute, since this requires some form of Turing-machine determinism. The classical conception of an algorithm does not accommodate the limit stages found in relativistic computation.

4. Challenging the bold thesis

PCTT-B says that the behavior of every physical system can be *simulated* (to any required degree of precision) by a Turing machine. Speculation that there may be physical processes whose behavior cannot be calculated by the universal Turing machine stretches back over several decades (see Copeland 2002 for a survey). The focus has been not so much on idealized constructions such as *RM* (which, if physical, is a counterexample to PCTT-B, as well as PCTT-M, since PCTT-B entails PCTT-M); rather, the focus has been on whether the mathematical equations governing the dynamics of physical systems are or are not Turing computable. Early papers by Scarpellini (1963), Komar (1964), and Kreisel (1965, 1967) raised this question. Georg Kreisel stated "There is no evidence that even present day quantum theory is a mechanistic, i.e. recursive theory in the sense that a recursively described system has recursive behavior" (1967: 270). Roger Penrose (1989; 1994) conjectured that some mathematical insights are non-recursive. Assuming that this mathematical thinking is carried out by some physical processes in the brain, the bold thesis must then be false. But Penrose's conjecture is highly controversial.

Another challenge to the bold thesis derives from the postulate of genuine physical randomness (as opposed to quasi-randomness). Church showed in 1940 that any infinite, genuinely random sequence is uncomputable (Church 1940: 134-135). Some have argued that, under certain conditions relating to unboundedness, PCTT-B is false in a universe containing a random element (to use Turing's term from his 1950: 445; see also Turing 1948: 416). A random element is a system that generates random sequences of bits. It is argued that if physical systems include systems capable of producing unboundedly many digits of an infinite random binary sequence, then PCTT-B is false (Copeland 2004, 2000; Calude et al. 2008, 2010; Piccinini 2011). One of us, Copeland, also argues that (again under unboundedness conditions) a digital computer using a random element forms a counterexample to PCTT-M (Copeland 2002). However, the latter claim, unlike the corresponding claim concerning PCTT-B, depends crucially upon one's account of computation — Shagrir denies that a digital computer with a random element *computes* the generated uncomputable sequences. In any case, though, it is an open question whether genuine random elements, able to generate unboundedly many digits of random binary sequences, exist in the physical universe, or physically could exist.

A further challenge to the bold thesis was formulated by Piccinini (2011). One of his argument's premises is: "if our physical theories are correct, most transformations of the relevant physical properties are transformations of Turing-uncomputable quantities into one another" (2011: 748). Another premise is: "a transformation of one Turing-uncomputable value into another Turing-uncomputable value is certainly a Turing-uncomputable operation" (2011: 748-749). The observation that in a continuous physical world, not all arithmetical operations are Turing computable is certainly correct (Copeland 1997). Where x and y are uncomputable real numbers, $x + y$ is in general not Turing computable, since the inputs x and y cannot be inscribed on a Turing machine's tape (except in the special case where x and y have been given proper names, e.g., where the halting number is named "τ"—but since there are only countably many proper names, most Turing uncomputable real numbers must remain nameless). If, therefore, the bold thesis is simply that “Any physical process is Turing computable” (Piccinini 2011: 746), then the thesis is indeed false in a continuous universe; as Piccinini argued, a Turing machine can receive at most a denumerable number of different inputs, and so the falsity of the bold thesis results from the cardinality gap between the physical functions, defined over non-denumerable domains, and the Turing computable functions, defined over denumerable domains.

However, this simple argument shows merely that Piccinini's version of the bold thesis is of little interest if the physical world is assumed to be continuous. Our own version of the thesis is sensitive to these considerations and requires only that physical processes be simulable to any specified degree of accuracy. Our version of the thesis is responsive to an account of (Turing machine) computation over the reals according to which—contra Piccinini's second premiss—the transformation of one Turing uncomputable value into another Turing uncomputable value can be a (Turing-machine) computable operation. Relative to this account, the real-valued functions of **plus**, **identity**, and **inverse** are computable by Turing machine, even though these functions sometimes map Turing uncomputable inputs to Turing uncomputable outputs: the definitions of real computable functions impose a continuity constraint, enabling the approximation (simulation) of uncomputable arguments and values. There are several (non-equivalent) characterizations of this continuity constraint; for the purpose of illustration, we select the characterization given by Andrzej Grzegorczyk (1955, 1957), and we adapt the exposition of Earman (1986).

We start with numbers:

Definition 1: A sequence of rational numbers $\{x_n\}$ is said to be effectively computable if there exist three Turing computable functions (over N) a,b,c such that $x_n = (-1)^{c(n)}a(n)\div b(n)$.

Definition 2: A real number r is said to be effectively computable if there is an effectively computable sequence of rational numbers that converges effectively to r. ('Converges effectively' means that there is an effectively computable function d over N such that $|r - x_n| < 1\div 2^m$ whenever $n \geq d(m)$.)

Now to functions:

Definition 3: A function f is an effectively computable function over the reals if:

(i) f is *sequentially computable*, i.e. for each effectively computable sequence $\{r_n\}$ of reals $\{f(r_n)\}$ is also effectively computable;

(ii) f is *effectively uniformly continuous* on rational intervals, i.e. if $\{x_n\}$ is an effective enumeration of the rationals without repetitions then there is a three-place Turing computable function g such that $|f(r) - f(r')| < 1\div 2^k$ whenever $x_m < r, r' < x_n$ and $|r - r'| < 1\div g(m,n,k)$ for all $r,r' \in R$ and all $m,n,k \in N$.

(If we confine f to a closed and bounded interval with computable end points then the above definition simplifies: no enumeration is necessary and g is only a function of k.)

Importantly, given that **plus** is *effectively uniformly continuous* on rational intervals, **plus** is a computable function, even though **plus** maps Turing uncomputable reals to Turing uncomputable reals.

It is interesting that, by and large, known physical laws give rise to functions over the reals that are computable in the sense just defined. A well-known exception was discovered by Marian Pour-El and Ian Richards (1981), who showed that the wave equation produces non-computable sequences for some computable initial conditions (i.e., for some computable input sequences). In that respect, the wave equation violates condition (i) in the definition of an effectively computable function over the reals (Definition 3). Clearly, this exception forms a potential counterexample to (our version of) the bold thesis.

However, this result of Pour-El and Richards is at the purely mathematical level, not the physical level. In his discussion of Pour-El and Richards, Pitowsky (1990) argued that their result does not refute the bold thesis, for two reasons:

> Firstly, the function *f* in the initial condition, though a recursive real function, is an extremely complex function. One can hardly expect such an initial condition to arise 'naturally' in any real physical situation. Secondly, we deal with recursive real functions, and in physics we never get beyond a few decimal digits of accuracy anyway. (Pitowsky 1990: 86-87)

Nevertheless, the result does demonstrate that being recursive

> is not a natural physical property. Physical processes do not necessarily preserve it. (Pitowsky 1990: 87)

Even if, in our world, the initial conditions envisaged in the Pour-El & Richards example do not occur, nevertheless these conditions could occur in some other physically possible world in which the wave equation holds, so showing, as Pitowsky said, that physical processes do not necessarily preserve recursiveness.

5. Challenging the super-bold thesis

We will turn next to PCTT-S. It might be objected that PCTT-S is immediately false, as may be shown by considering a universal Turing machine implemented as a physical system: many interesting questions about such a system are undecidable. Pitowsky (1996) describes one such construction, due to Moore (1990), in which a universal machine is realized in a moving particle bouncing between parabolic and linear mirrors in a unit square. This system can certainly be simulated by a Turing machine (since it is a Turing machine). But there are nevertheless undecidable questions about its behavior:

> At this stage we can apply Turing's theorem on the undecidability of the halting problem. It says: *There is no algorithm to decide whether a universal Turing machine halts on a given input.* Translating this assertion into physical language, it means that there is no algorithm to predict whether the particle ever enters the subset of the square corresponding to the halting state. This assertion is valid when we know the exact initial conditions with unbounded accuracy, or even with actually infinite accuracy. Therefore, to answer the question: 'Is the particle ever going to reach this region of space?', Laplace's Demon needs computational powers exceeding any algorithm. In other words, he needs to consult an oracle. (Pitowsky 1996: 171)

Pitowsky further noted that many other yes/no questions about the system are computationally undecidable. In fact, it follows from Rice's theorem that "almost every conceivable question about the unbounded future of a universal Turing machine turns out to be computationally undecidable" (Pitowsky 1996: 171; see also Harel and Feldman 2004). Rice's theorem says: any nontrivial property about the language recognized by a Turing machine is undecidable, where a property is nontrivial if at least one Turing machine has the property and at least one does not.

However, the objection that PCTT-S is straight-out false can hardly be sustained. It would by no means be facile to implement a Turing machine, with its infinite tape, in a finite physical object. The assumption that the physical universe is able to supply an infinite memory tape is controversial. If the number of particles in the cosmos is finite, and each cell of the tape requires at least one particle for its implementation, then clearly the assumption of an infinite physical tape must be false. Against this, it might be pointed out that there are a number of well-known constructions for shoehorning an infinite amount of memory into a finite object. For example, a single strip of tape 1 metre long can be used: the first cell occupies the first half-metre, the second cell the next quarter-metre, the third cell the next eighth of a metre, and so on. However, this observation does not help matters. Obviously, this construction can be realized only in a universe whose physics allows for the infinite divisibility of the tape—again by no means an evidently true assumption.

Another way to implement the Turing machine with its infinite tape is in the continuous (or rational) values of a physical magnitude, as described in Moore's system. Assume that each (potential) configuration of the Turing machine is encoded in a (potentially infinite) sequence of 0s and 1s. The goal now is to efficiently realize each sequence in a unique location of the particle in the (finite) unit square. Much like the example of the 1-metre strip of tape above, this realization can be achieved if we use potentially infinitely many different locations (*x*,*y*) within a unit square, where *x* and *y* are continuous (or rational) values between 0 and 1. This realization, however, requires that the (idealized) mirrors have the ability to bounce the particle, accurately, into potentially infinitely many different positions within the unit square. Moreover, if Laplace's Demon wants to predict the future position of the particle after k rounds in the square unit, the Demon will have to be able to measure the differences between arbitrarily close positions. As Pitowsky notes: "[T]o predict the particle position with accuracy of 1/2 the demon has to measure initial conditions with accuracy $2^{-(k+1)}$. The ratio of final error to initial error is 2^k, and growing exponentially with time k (measured by the number of rounds)" (1996: 167). This means that implementing a Turing machine in Moore's system requires determining a particle's position with practically infinite precision (an accuracy of $2^{-(k+1)}$ for unbounded k), and it is questionable whether this implementation is physically feasible. At any rate, the claim that this is physically feasible is, again, far from obvious.

To summarize this discussion, PCTT-S hypothesizes in part that the universe is physically unable to supply an infinite amount of memory, since if PCTT-S is true, the resources for constructing a universal computing machine must be unavailable (the other necessary resources, aside from the infinite memory, being physically undemanding). This point helps illuminate the relationships between PCTT-S and PCTT-M. Returning to the discussion of the relativistic machine RM and the infinite memory problem raised in Section 3, it is clearly the case that, since RM requires infinite memory, PCTT-S rules out RM (and this is to be expected, since PCTT-M rules out RM, and PCTT-S entails PCTT-M). Nevertheless, the falsity of PCTT-S, and the availability of infinite memory, would be insufficient to falsify PCTT-M—a universe that consists of nothing but a universal Turing machine with its infinite tape does not falsify PCTT-M. Thus, counterexamples to PCTT-M must postulate not only the availability of infinite memory but also additional physical principles of some sort, such as gravitational time dilation or unbounded acceleration or unbounded shrinking of components. Relativistic, accelerating and shrinking machines arguably invoke these principles successfully, and, hence, provide counterexamples to PCTT-M.

Moving on to challenges to PCTT-S at the quantum level, there are undecidable questions concerning the behavior of quantum systems. In 1986, Robert Geroch and James Hartle argued that *un*decidable physical theories "should be no more unsettling to physics than has the existence of well-posed problems unsolvable by any algorithm have been to mathematics"; and they suggested such theories may be "forced upon us" in the quantum domain (Geroch and Hartle 1986: 534, 549). Arthur Komar raised "the issue of the macroscopic distinguishability of quantum states" in 1964, claiming there is no effective procedure "for determining whether two arbitrarily given physical states can be superposed to show interference effects" (Komar 1964: 543-544). More recently, Jens Eisert, Markus Müller and Christian Gogolin showed that "the very natural physical problem of determining whether certain outcome sequences cannot occur in repeated quantum measurements is undecidable, even though the same problem for classical measurements is readily decidable" (Eisert, Müller and Gogolin 2012: 260501-1). (This is an example of a problem that refers unboundedly to the future, but not to any specific time, as in Pitowsky's examples mentioned earlier.) Eisert, Müller and Gogolin went on to suggest that "a plethora of problems" in quantum many-body physics and quantum computing may be undecidable (2012: 260501-1 - 260501-4).

Dramatically, a 2015 *Nature* article by Toby Cubitt, David Perez-Garcia, and Michael Wolf outlined a proof that "the spectral gap problem is algorithmically undecidable: there cannot exist any algorithm that, given a description of the local interactions, determines whether the resultant model is gapped or gapless" (Cubitt et al. 2015: 207). Cubitt describes this as the "first undecidability result for a major physics problem that people would really try to solve" (in Castelvecchi 2015).

The spectral gap, an important determinant of a material's properties, refers to the energy spectrum immediately above the ground energy level of a quantum many-body system (assuming that a well-defined least energy level of the system exists); the system is said to be gapless if this spectrum is continuous and gapped if there is a well-defined next least energy level. The spectral gap problem for a quantum many-body system is the problem of determining whether the system is gapped or gapless, given the finite matrices describing the local interactions of the system.

In their proof, Cubitt et al. encode the halting problem in the spectral gap problem, so showing that the latter is at least as hard as the former. The proof involves an infinite family of 2-dimensional lattices of atoms; but they point out that their result also applies to finite systems whose size increases: "Not only can the lattice size at which the system switches from gapless to gapped be arbitrarily large, the threshold at which this transition occurs is uncomputable" (Cubitt et al. 2015: 210-211). Their proof offers an interesting countermodel to the super-bold thesis. The countermodel involves a physically relevant example of a finite system, of increasing size, that lacks a Turing computable procedure for extrapolating future behavior from (complete descriptions of) its current and past states.

It is debatable whether any of these quantum models matches the real quantum world. Cubitt et al. admit that the model used in their proof is highly artificial, saying "Whether the results can be extended to more natural models is yet to be determined" (Cubitt et al. 2015: 211). There is also the question of whether the spectral gap problem could become computable when only local Hilbert spaces of realistically low dimensionality are considered. Nevertheless, these results are certainly suggestive. The super-bold thesis cannot be taken for granted—even in a finite quantum universe.

6. Conclusion

We have distinguished three theses about physical computability, and have discussed some empirical evidence that might challenge these. Concerning the boldest versions, PCTT-B and PCTT-S, both are false if the physical universe permits infinite memory and genuine randomness. Even assuming that the physical universe is deterministic, the most that can be said for PCTT-B and PCTT-S is that, to date, there seems to be no decisive empirical evidence against them. PCTT-B and PCTT-S are both thoroughly empirical theses; but matters are more complex in the case of the modest thesis PCTT-M, since a conceptual issue also bears on the truth or falsity of this thesis (even in a universe containing genuine randomness)—namely, the difficult issue of what counts as physical computation. Our conclusion is that, at the present stage of physical enquiry, it is unknown whether any of the theses is true.

Acknowledgements: An early version of the paper was presented at the *Workshop on Physics and Computation* (UCNC2018) in Fontainebleau, at the *International Association for Computing and Philosophy meeting (*IACAP2018*)* in Warsaw, and at the research logic seminar in Tel Aviv University. Thanks to the audiences for helpful discussion. We are also grateful to Piccinini for his critical comments on a draft of this paper. Shagrir's research was supported by the Israel Science Foundation grant 830/18.

References

Aaronson, S. (2005). Guest column: NP-complete problems and physical reality. *ACM Sigact News* 36.1: 30-52.

Andréka, H., Németi, I. and Németi, P. (2009). General Relativistic Hypercomputing and Foundation of Mathematics. *Natural Computing* 8:499–516.

Barrett, J.A. and Aitken, W. (2010). A Note on the Physical Possibility of Ordinal Computation. *The British Journal for the Philosophy of Science* 61: 867-874.

Button, T. (2009). SAD Computers and Two Versions of the Church-Turing Thesis. *British Journal for the Philosophy of Science* 60: 765–792.

Blum, L., Cucker, F., Shub, M. and Smale, S. (1998). *Complexity and Real Computation*. Springer.

Calude, C. S., Dinneen, M. J., Dumitrescu, M., Svozil, K. (2010). Experimental evidence of quantum randomness incomputability. *Physical Review A* 82: 022102–1-022102–8.

Calude, C. S., Svozil, K. (2008). Quantum Randomness and Value Indefiniteness. *Advanced Science Letters* 1: 165-168.

Castelvecchi, D. (2015). Paradox at the Heart of Mathematics Makes Physics Problem Unanswerable. *Nature* 528: 207.

Chalmers, D. J. (2011). A Computational Foundation for the Study of Cognition. *Journal of Cognitive Science* 12: 323-357.

Church, A. (1936). An Unsolvable Problem of Elementary Number Theory. *American Journal of Mathematics* 58: 345-363.

Church, A. (1940). On the Concept of a Random Sequence. *American Mathematical Society Bulletin* 46: 130-135.

Copeland, B. J. (1997). The Broad Conception of Computation. *American Behavioral Scientist* 40: 690-716.

— (1998) Even Turing Machines Can Compute Uncomputable Functions. In C. Calude, J. Casti and M. Dinneen (eds.), *Unconventional Models of Computation*. Springer.

— (2000). Narrow Versus Wide Mechanism. *Journal of Philosophy* 96: 5-32. Reprinted in M. Scheutz (ed.) (2002), *Computationalism*. Cambridge, Mass.: MIT Press.

— (2002). Hypercomputation. *Minds and Machines* 12: 461-502.

— (2004). Hypercomputation: Philosophical Issues. *Theoretical Computer Science* 317: 251-267.

— (2017). The Church-Turing Thesis. In E. N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy*. <https://plato.stanford.edu/archives/win2017/entries/church-turing/>.

Copeland, B.J. and Shagrir, O. (2007). Physical Computation: How General are Gandy's Principles for Mechanisms? *Minds and Machines* 17: 217-231.

— (2011). Do Accelerating Turing Machines Compute the Uncomputable? *Minds and Machines* 21: 221–239.

— (2019). The Church–Turing Thesis—Logical Limit, or Breachable Barrier? *Communications of the ACM*, 62: 66-74.

Copeland, B.J., Shagrir, O. and Sprevak, M. (2018). Zuse's Thesis, Gandy's Thesis, and Penrose's Thesis. In M. Cuffano and S. Fletcher (eds.), *Physical Perspectives on Computation, Computational Perspectives on Physics*. Cambridge University Press: Cambridge, pp. 39–59.

Cubitt, T. S., D. Perez-Garcia, and M. M. Wolf (2015). Undecidability of the Spectral Gap. *Nature* 528.7581: 207-211.

Davies, B. E. (2001). Building Infinite Machines. *The British Journal for the Philosophy of Science*. 52:671–682.

Deutsch, D. (1985). Quantum Theory, the Church-Turing Principle and the Universal Quantum Computer. *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences* (Vol. 400, No. 1818, pp. 97-117). The Royal Society.

Earman, J. (1986). *A Primer on Determinism*. Reidel.

Earman, J. and Norton, J. D. (1993). Forever is a Day: Supertasks in Pitowsky and Malament-Hogarth Spacetimes. *Philosophy of Science* 60: 22-42.

Eisert, J., M. P. Müller, and C. Gogolin (2012). Quantum Measurement Occurrence is Undecidable. *Physical review letters* 108.26: 260501.

Etesi, G. and Németi, I. (2002). Non-Turing Computations via Malament-Hogarth Space-times. *International Journal of Theoretical Physics* 41: 341–370.

Fresco, N. (2014). *Physical Computation and Cognitive Science.* Springer.

Gandy, R.O. 1980: Church's thesis and principles of mechanisms. In S.C. Kleene, J. Barwise, H.J. Keisler and K. Kunen (Eds.), The Kleene Symposium. North-Holland.

Geroch, R. and Hartle, J.B. (1986). Computability and Physical Theories. *Foundations of Physics* 16: 533–550.

Grzegorczyk, A. (1955). Computable Functional. *Fundamenta Mathematicae* 42: 168-203.

— (1957). On the Definitions of Computable Real Continuous Functions. *Fundamenta Mathematicae* 44: 61–71.

Harel D. and Feldman Y.A. (1992). *Algorithmics: The Spirit of Computing* (3rd edition). Addison-Wesley.

Hogarth, M. (1994). Non-Turing Computers and Non-Turing Computability. *Proceedings of the Biennial Meeting of the Philosophy of Science Association* 1: 126–138.

— (2004). Deciding arithmetic using SAD computers. *The British Journal for the philosophy of Science*, *55*: 681-691.

Komar, A. (1964). Undecidability of Macroscopically Distinguishable States in Quantum Field Theory. *Physical Review*, second series, 133B, pp. 542–544.

Kreisel, G. (1965). Mathematical Logic. In T.L. Saaty, ed., Lectures on Modern Mathematics, Vol. 3, Wiley.

— (1967), Mathematical Logic: What has it Done for the Philosophy of Mathematics? In R. Schoenman, ed., Bertrand Russell: Philosopher of the Century, Allen and Unwin.

Lacombe, D. (1955). Extension de la Notion de Fonction Récursive aux Fonctions d'une ou Plusieurs Variables Réelles III, *Comptes Rendus Académie des Sciences Paris* (French), 241:151–153.

Mazur, S. (1963). *Computational Analysis*. Warsaw: Razprawy Matematyczne.

Miłkowski, M. (2013). *Explaining the Computational Mind*. MIT Press.

Moore, C. (1990). Unpredictability and Undecidability in Dynamical Systems. *Physical Review Letters*, *64*: 2354-2357.

Németi, I. and Dávid, G. (2006). Relativistic Computers and the Turing Barrier. *Journal of Applied Mathematics and Computation* 178: 118–142.

Piccinini, G. (2011). The Physical Church-Turing Thesis: Modest or bold? *The British Journal for the Philosophy of Science* 62:733-769.

— (2015). *Physical Computation: A Mechanistic Account*. Oxford University Press.

Pitowsky, I. (1990). The Physical Church Thesis and Physical Computational Complexity. *Iyyun* 39: 81–99.

— (1996) Laplace's Demon Consults an Oracle: The Computational Complexity of Prediction. *Studies in History and Philosophy of Science Part B* 27.2: 161-180.

— (2002). Quantum Speed-up of Computations. *Philosophy of Science 69*: S168-S177.

Pour-El, M. B. (1974). Abstract Computability and Its Relation to the General Purpose Analog

Computer (Some Connections between Logic, Differential Equations and Analog Computers). *Transactions of the American Mathematical Society* 199: 1–28.

Pour-El, M. B. and Richards, I. J. (1981). The Wave Equation with Computable Initial Data such that its Unique Solution is not Computable. *Advances in Mathematics* 39: 215-239.

— (1989). *Computability in Analysis and Physics*. Springer.

Scarpellini, B. (1963). Zwei Unentscheitbare Probleme der Analysis. *Zeitschrift für mathematische Logik und Grundlagen der Mathematik* 9: 265–289. English translation in *Minds and Machines* 12 (2002): 461-502.

Shagrir, O. (2006). Why We View the Brain as A Computer. *Synthese* 153:393-416.

Shagrir, O. and Pitowsky, I. (2003). Physical Hypercomputation and the Church–Turing Thesis. *Minds and Machines* 13: 87-101.

Sprevak, M. (2010). Computation, Individuation, and the Received View on Representation. *Studies in History and Philosophy of Science* 41:260–270.

Turing, A.M. (1936). On Computable Numbers, with an Application to the Entscheidungsproblem. *Proceedings of the London Mathematical Society, series 2*, 42: 230-265.

— (1947). Lecture on the Automatic Computing Engine. In Copeland, B. J. 2004. *The Essential Turing: Seminal Writings in Computing, Logic, Philosophy, Artificial Intelligence, and Artificial Life, Plus the Secrets of Enigma*. Oxford University Press.

— (1948). Intelligent Machinery. In *The Essential Turing*.

— (1950). Computing Machinery and Intelligence. In *The Essential Turing*.

Welch, P. D (2008). The Extent of Computation in Malament-Hogarth Spacetimes. *The British Journal for the Philosophy of Science* 59: 659-674.

Wolfram, S. (1985) Undecidability and Intractability in Theoretical Physics. *Physical Review Letters* 54: 735–738.