

Technical Report No. 2013-609
NONUNIVERSALITY IN COMPUTATION:
THIRTEEN MISCONCEPTIONS RECTIFIED

Selim G. Akl

School of Computing, Queen's University

Kingston, Ontario K7L 3N6, Canada

akl@cs.queensu.ca

August 28, 2013

Abstract

In January 2005 I showed that universality in computation cannot be achieved. Specifically, I exhibited a number of distinct counterexamples, each of which sufficing to demonstrate that no finite and fixed computer can be universal in the sense of being able to simulate successfully any computation that can be executed on any other computer. The number and diversity of the counterexamples attest to the general nature of the nonuniversality result. This not only put to rest the Church-Turing Thesis, by proving it to be a false conjecture, but also was seen to apply, in addition to the Turing Machine, to all computational models, past, present, and future, conventional and unconventional.

While over eight years have now passed since nonuniversality in computation was established, the result remains largely misunderstood. There appear to be at least two main reasons for this state of affairs. As often happens to new ideas, the nonuniversality result was confronted with a stubborn entrenchment in a preconceived, deeply held, and quasi-religious belief in the existence of a universal computer. This was exacerbated by a failure to read the literature that demonstrates why such a belief is unfounded. Behavior of this sort, sadly not uncommon in science, explains the enduring mirage of the universal computer.

The purpose of this paper is to rectify the most notorious misconceptions associated with nonuniversality in computation. These misconceptions were expressed to the author in personal communications, orally, by email, and in referee reports. Each misconception is quoted verbatim and a detailed response to it is provided. The paper concludes by inviting the reader to take a computational challenge.

Keywords and phrases: nonuniversality in computation; myth of universality; universal computer; simulation; models of computation; time unit; analysis of algorithms; computational complexity; theoretical computer science; inherently parallel computations; time-varying variables; time-varying computational complexity; rank-varying computational complexity; interacting variables; uncertain time constraints; mathematical constraints; Church-Turing Thesis; Gödel's incompleteness theorem; quantum computing; finite and fixed number of operations per time unit, general-purpose computer; finiteness condition; unbounded space; unbounded time; communication with the outside world; Turing Machine; Random Access Machine; parallel computer; sequential computer; speedup theorem; slowdown theorem; Amdahl's law.

1 Introduction

For a long time people believed that the Sun orbits our Earth. We now know better. And so it is with universality in computation. Consider the following statement:

“It can also be shown that any computation that can be performed on a modern-day digital computer can be described by means of a Turing machine. Thus if one ever found a procedure that fitted the intuitive notions, but could not be described by means of a Turing machine, it would indeed be of an unusual nature since it could not possibly be programmed for any existing computer.” [32], p. 80.

The first sentence in the above quote is clearly false, and well known counterexamples abound in the literature (see, for example, [33, 59], and for further references [6, 11]). The second sentence, on the other hand, is only half true. Indeed, it is perfectly justified to say that a computation that cannot be performed on a Turing Machine must be of an *unusual nature*. However, this does not mean that such a computation cannot be programmed on any existing computer. Rather, the existence of such a computation would instead mean that the Turing Machine is simply not universal and that the Church-Turing Thesis (whose validity is assumed implicitly in the above quote) is in fact invalid.

In January 2005 I showed that the concept of a Universal Computer cannot be realized [6]. Specifically, I exhibited instances of a *computable* function \mathcal{F} that cannot be computed on any machine \mathcal{U} that is capable of only a finite and fixed number of operations per step (or time unit). This remains true even if the machine \mathcal{U} is endowed with an infinite memory and the ability to communicate with the outside world while it is attempting to compute \mathcal{F} . It also holds if, in addition, \mathcal{U} is given an indefinite amount of time to compute \mathcal{F} .

Formally, my result is stated as follows:

Theorem 1. Nonuniversality in computation: Given n spatially and temporally connected physical variables, X_1, X_2, \dots, X_n , where n is a positive integer, there exists a function $F^n(X_1, X_2, \dots, X_n)$ of these variables, such that no computer can evaluate F^n for any arbitrary n , unless it is capable of an infinite number of operations per time unit. ■

A constructive proof of Theorem 1 seeks to define a function F^n that obeys the following property: F^n is readily computable by a machine M_n capable of exactly n operations per time unit; however, this machine cannot compute F^{n+1} when the number of variables is $n + 1$. While a second machine M_{n+1} capable of $n + 1$ operations per time unit can now compute the function F^{n+1} of $n + 1$ variables, M_{n+1} is in turn defeated by a function F^{n+2} of $n + 2$ variables. In principle, the escalation continues without end.

This point deserves emphasis. While the function $F^{n+1}(X_1, X_2, \dots, X_{n+1})$ is easily computed by M_{n+1} , it cannot be computed by M_n . Even if given infinite amounts of time and space, machine M_n is incapable of simulating the actions of M_{n+1} . Furthermore, machine M_{n+1} is in turn thwarted by $F^{n+2}(X_1, X_2, \dots, X_{n+2})$, a function computable by a third machine M_{n+2} . The process repeats indefinitely. Therefore no computer is universal if it is capable of exactly $T(i)$ operations during time unit i , where i is a positive integer, and $T(i)$ is finite and fixed once and for all, for it will be faced with a computation requiring $V(i)$ operations during time unit i , where $V(i) > T(i)$ for at least one i . (Note that this last statement is, in fact, a more general form of the result in Theorem 1, as it invokes neither the physical variables nor the function F^n .)

Therefore, in order to actually prove Theorem 1 constructively, it suffices to exhibit (at least) one concrete instance of the function F^n satisfying the above property. Indeed we

describe several examples of such a function occurring in:

1. *Computations with time-varying variables:* The variables, over which the function is to be computed, are themselves changing with time.
2. *Computations with time-varying computational complexity:* The computational complexity of the function to be computed is itself changing with time.
3. *Computations with rank-varying computational complexity:* Given several functions to be computed, and a schedule for computing them, the computational complexity of a function depends on its position in the schedule.
4. *Computations with interacting variables:* The variables of the function to be computed are parameters of a physical system that interact unpredictably when the system is disturbed.
5. *Computations with uncertain time constraints:* There is uncertainty with regards to the input (when and for how long are the input data available), the calculation (what to do and when to do it), and the output (the deadlines are undefined at the outset); furthermore, the function that resolves each of these uncertainties itself has demanding time requirements.
6. *Computations with global mathematical constraints:* The function to be computed is over a system whose variables must collectively obey a mathematical condition at all times.

For instance, suppose that the X_i are themselves functions that vary with time. It is therefore appropriate to write the n variables as $X_1(t), X_2(t), \dots, X_n(t)$, that is, as functions of the time variable t . Further assume that, while it is known that the X_i change with time, the actual functions that effect these changes are not known (for example, X_i may be a true random variable). Let

$$F^{n+1}(X_1, X_2, \dots, X_{n+1}) = F_1(X_1), F_2(X_2), \dots, F_{n+1}(X_{n+1}),$$

where each F_i is a simple function of one variable that takes one time unit to compute. The problem calls for computing $F_i(X_i(t))$, for $i = 1, 2, \dots, n$, at time $t = t_0$. Specifically, let $F_i(X_i(t))$ simply represent the reading of $X_i(t)$ from an external medium. The fact that $X_i(t)$ changes with the passage of time means that, for $k > 0$, not only is each value $X_i(t_0 + k)$ different from $X_i(t_0)$, but also the latter cannot be obtained from the former. No computer capable of fewer than n read operations per time unit can solve this problem.

We note in passing that the above example is deliberately simple in order to convey the idea. Reading a datum, that is, acquiring it from an external medium, is the most elementary form of information processing. Any computer must be able to perform such operation. This simplest of counterexamples suffices to establish nonuniversality in computation. Of course, if one wishes, the computation can be made more complex, at will. While our main conclusion remains unchanged, for some, a more complex argument may sound more convincing. Thus, for example, we may add arithmetic by requiring that $F_i(X_i(t))$ call for reading $X_i(t)$ and incrementing it by one, for $i = 1, 2, \dots, n$, at time $t = t_0$. Reading $X_i(t)$, incrementing it by one, and returning the result takes one time unit.




Figure 1: Incompleteness in mathematics and nonuniversality in computation.

In any case, a computer M_n capable of n operations per time unit (for example, one with n processors operating in parallel) can compute all the $F_i(X_i(t))$ at $t = t_0$ successfully. A computer capable of fewer than n operations per time unit fails to compute all the $F_i(X_i(t))$ at $t = t_0$. Indeed, suppose that a computer is capable of $n - 1$ operations per time unit. It would compute $n - 1$ of the $F_i(X_i(t))$ at $t = t_0$ correctly. Without loss of generality, assume that it computes $F_i(X_i(t))$ at $t = t_0$ for $i = 1, 2, \dots, n - 1$. Now one time unit would have passed, and when it attempts to compute $F_n(X_n(t_0))$, it would be forced to incorrectly compute $F_n(X_n(t_0 + 1))$.

Is computer M_n universal? Certainly not. For when the number of variables is $n + 1$, M_n fails to perform the above computation. As stated above, a succession of machines M_{n+1}, M_{n+2}, \dots , each succeeds at one level only to be foiled at the next. The implication to the theory of computation is akin to that of Gödel's incompleteness theorem to mathematics. In the same way as no finite set of axioms A_i can be complete, no computer C_i is universal that can perform a finite and fixed number of operations per time unit. This is illustrated in Fig. 1: For every set of axioms A_i there exists a statement G_{i+1} not provable in A_i , but provable in A_{i+1} ; similarly, for every machine C_i there is a problem P_{i+1} not solvable on C_i , but solvable on C_{i+1} .

This result applies not only to idealized models of computation, such as the Turing Machine, the Random Access Machine, and the like, but also to all known general-purpose computers, including existing conventional computers (both sequential and parallel), as well as contemplated unconventional ones such as biological and quantum computers. It is true for computers that interact with the outside world in order to read input and return output (unlike the Turing Machine, but like every realistic general-purpose computer). It is also valid for computers that are given unlimited amounts of time and space in order to perform their computations (like the Turing Machine, but unlike realistic computers). Even accelerating machines that increase their speed at every step (such as doubling it, or squaring it, or any such fixed acceleration) at a rate that is set in advance, cannot be universal!

The only constraint that we have placed on the computer (or model of computation) that claims to be universal is that the number of operations of which it is capable per time unit be finite and fixed once and for all. In this regard, it is important to note that:

1. The requirement that the number of operations per time unit, or step, be *finite* is

necessary for any “reasonable” model of computation; see, for example, [53], p. 141.

2. The requirement that this number be *fixed* once and for all is necessary for any model that purports to be “universal”; see, for example, [22], p. 210.

Without these two requirements, the theory of computation in general, and the theory of algorithms, in particular, would be totally irrelevant.

The consequences to theoretical and practical computing are significant. Thus the conjectured “Church-Turing Thesis” is false. It is no longer true that, given enough time and space, any single general-purpose computer, defined a priori, can perform all computations that are possible on all other computers. Not the Turing Machine, not your laptop, not the most powerful of supercomputers. In view of the computational problems mentioned above (and detailed in [6] – [18], [26], and [42] – [47]), the only possible universal computer would be one capable of an infinite number of operations per time unit.

In fact, this work has led to the discovery of computations that can be performed on a quantum computer but not on any classical machine (even one with infinite resources), thus showing for the first time that the class of problems solvable by classical means is a true subset of the class of problems solvable by quantum means [44]. Consequently, the only possible universal computer would have to be quantum (as well as being capable of an infinite number of operations per time unit).

The results described in the previous paragraphs were obtained and published between 2005 and 2012. They appeared in technical reports, were presented at conferences, and published as journal papers and book chapters, as documented in the bibliography section. Yet, over eight years after they were first announced, they remain generally misunderstood, and consequently controversial. This is not surprising: Old habits die hard. It is difficult to abandon ideas that one has believed all of one’s life. This applies to inherited scientific ideas, that people persist in defending, even if these ideas are demonstrably wrong. Disruptive innovations usually disturb and destabilize. To announce that the long-held belief in universality in computation is, in fact, in error was met with great skepticism, to say the least. Most of the animosity towards the author came from scientists who had not bothered to read his work. The title of a paper was sufficient to ignite their anger.

It started with hostile email. Then came the uninspired referee reports. And finally Wikipedia removed the entry on Nonuniversality in Computation, created by a student of the author, after one reader objected. It was clear that these people had not read any of my work on the subject. This is when I decided to write this paper, based on a web page [4] I had created to reply to unfounded criticisms.

The remainder of this paper is organized as follows. Section 2 provides definitions for many of the terms used in the discussion to follow (some of the more common of these terms were already introduced in this section without definition, in order not to interrupt the flow of the opening paragraphs). Section 3 summarizes the work in my previous papers, culminating in the fundamental result that the ‘Universal Computer’ is a myth. Section 4 is a dialog with an imaginary debater who embodies all of my critics and voices all their misconceptions. Section 5 offers a computational challenge to the readers.

2 Preliminaries

A *time unit* is the length of time required by a processor to perform a *step* of its computation, consisting of three elementary operations: a *read* operation in which it receives a constant number of fixed-size data as input, a *calculate* operation in which it performs a fixed number of constant-time *arithmetic* and *logical* calculations (such as adding two numbers, comparing two numbers, and so on), and a *write* operation in which it returns a constant number of fixed-size data as output.

A *sequential computer*, consists of a single processor. A *parallel computer* has n processors, numbered 1 to n , where $n \geq 2$. Both computers use the same type of processor and that processor is the fastest possible [2].

The concept of *computational universality* is one of the dogmas in Computer Science, stated as follows: Given enough time and space, any general-purpose computer can, through simulation, perform any computation that is possible on any other general-purpose computer. Statements such as this are commonplace in the computer science literature, and are served as standard fare in undergraduate and graduate courses alike. Sometimes the statement is restricted to the Turing Machine (TM), and is referred to as the *Church-Turing Thesis* (CTT), which stipulates that there exists a Universal Computer, namely, the TM; thus, the following quote is typical: “A Turing machine can do everything that a real computer can do.” [53], p. 125. Other times, the statement is made more generally about a *Universal Computer*, that is, a computer capable of executing any computation that can be performed by any other computer; thus, we are told: “It is possible to build a universal computer: a machine that can be programmed to perform any computation that any other physical object can perform.” [22], p. 134.

As its name indicates, the CTT is a *conjecture*, whose proof has remained elusive due to the difficulty in defining what it means *to compute*. A few points are worth making in this regard:

1. The *Universal Turing Machine* (UTM) defined by Alan Turing, is ‘universal’ in the sense of being able to simulate any computation performed by any other special-purpose TM. In other words, the UTM can be programmed to execute any TM computation. This is a provable property, and is not in question here or elsewhere.
2. However, the UTM is not known to be ‘universal’ in the more general sense of being able to simulate by its own means any computation performed by means of any other computational device (not necessarily a TM). It is this universality of the UTM, in the more general sense, that is expressed by the CTT.
3. My work has shown the CTT, and more generally computational universality, to be false. Specifically, I proved the stronger result that no ‘Universal Computer’ is possible, thus eliminating, not only the UTM, but also all other models of computation as contenders for the title.

3 Inherently Parallel Computations

The term *inherently parallel computation* refers to a computation that can be performed successfully only on a parallel computer with an appropriate number of processors. Examples

are computations that involve: time-varying variables, time-varying complexity, rank-varying complexity, interacting variables, uncertain time constraints, and variables obeying mathematical constraints. These computations also provide counterexamples to the existence of a universal computer [6]-[15], [26], [42]-[47].

3.1 Time-Varying Variables

This computation was described in Section 1. It is reproduced briefly here for completeness in the context of inherently parallel computations. For a positive integer n larger than 1, we are given n functions, each of one variable, namely, f_1, f_2, \dots, f_n , operating on the n physical variables x_1, x_2, \dots, x_n , respectively. Specifically, it is required to compute $F_i(x_i)$, for $i = 1, 2, \dots, n$. For example, $f_i(x_i)$ may be equal to x_i^2 . What is unconventional about this computation, is the fact that the x_i are themselves (unknown) functions $x_1(t), x_2(t), \dots, x_n(t)$, of the time variable t . It takes one time unit to evaluate $f_i(x_i(t))$. The problem calls for computing $f_i(x_i(t))$, $1 \leq i \leq n$, at time $t = t_0$. Because the function $x_i(t)$ is unknown, it cannot be inverted, and for $k > 0$, $x_i(t_0)$ cannot be recovered from $x_i(t_0 + k)$.

A sequential computer fails to compute all the f_i as desired. Indeed, suppose that $x_1(t_0)$ is initially operated upon. By the time $f_1(x_1(t_0))$ is computed, one time unit would have passed. At this point, the values of the $n - 1$ remaining variables would have changed. The same problem occurs if the sequential computer attempts to first read all the x_i , one by one, and store them before calculating the f_i .

By contrast, a parallel computer consisting of n independent processors may perform all the computations at once: For $1 \leq i \leq n$, and all processors working at the same time, processor i computes $f_i(x_i(t_0))$, leading to a successful computation.

3.2 Time-Varying Computational Complexity

Here, the computational complexity of the problems at hand depends on *time* (rather than being, as usual, a function of the problem *size*). Thus, for example, tracking a moving object (such as a spaceship racing towards Mars) becomes harder as it travels away from the observer.

Suppose that a certain computation requires that n independent functions, each of one variable, namely, $f_1(x_1), f_2(x_2), \dots, f_n(x_n)$, be computed. Computing $f_i(x_i)$ at time t requires $C(t) = 2^t$ operations, for $t \geq 0$ and $1 \leq i \leq n$. Further, there is a deadline for reporting the results of the computations: All n values $f_1(x_1), f_2(x_2), \dots, f_n(x_n)$ must be returned by the end of the third time unit, that is, when $t = 3$.

It should be easy to verify that no sequential computer, capable of exactly one constant-time operation per step (that is, per time unit), can perform this computation for $n \geq 3$. Indeed, $f_1(x_1)$ takes $C(0) = 2^0 = 1$ time unit, $f_2(x_2)$ takes another $C(1) = 2^1 = 2$ time units, by which time three time units would have elapsed. At this point none of $f_3(x_3), \dots, f_n(x_n)$ would have been computed. By contrast, an n -processor parallel computer solves the problem handily. With all processors operating simultaneously, processor i computes $f_i(x_i)$ at time $t = 0$, for $1 \leq i \leq n$. This consumes one time unit, and the deadline is met.

3.3 Rank-Varying Computational Complexity

Suppose that a computation consists of n stages. There may be a certain precedence among these stages, or the n stages may be totally independent, in which case the order of execution is of no consequence to the correctness of the computation. Let the *rank* of a stage be the order of execution of that stage. Thus, stage i is the i th stage to be executed. Here we focus on computations with the property that the number of operations required to execute stage i is $C(i)$, that is, a function of i only.

When does rank-varying computational complexity arise? Clearly, if the computational requirements grow with the rank, this type of complexity manifests itself in those circumstances where it is a disadvantage, whether avoidable or unavoidable, to being i th, for $i \geq 2$. For example, the precision and/or ease of measurement of variables involved in the computation in a stage s may decrease with each stage executed before s .

The same analysis as in the previous section applies by substituting the rank for the time.

3.4 Interacting Variables

A physical system has n variables, x_1, x_2, \dots, x_n , each of which is to be measured or set to a given value at regular intervals. One property of this system is that measuring or setting one of its variables modifies the values of any number of the system variables unpredictably.

A sequential computer measures *one* of the values (x_1 , for example) and by so doing it disturbs the equilibrium, thus losing all hope of recording the state of the system within the given time interval. Similarly, the sequential approach cannot update the variables of the system properly: Once x_1 has received its new value, setting x_2 disturbs x_1 unpredictably.

A parallel computer with n processors, by contrast, will measure *all* the variables x_1, x_2, \dots, x_n simultaneously (one value per processor), and therefore obtain an accurate reading of the state of the system within the given time frame. Consequently, new values x_1, x_2, \dots, x_n can be computed in parallel and applied to the system simultaneously (one value per processor).

3.5 Uncertain Time Constraints

In this paradigm, we are given a computation each of whose components, namely, the input phase, the calculation phase, and the output phase, needs to be computed by a certain deadline. However, unlike the standard situation in conventional computation, the deadlines here are not known at the outset. In fact, and this is what makes this paradigm truly unconventional, we do not know at the moment the computation is set to start, *what* needs to be done, and *when* it should be done. Certain physical parameters, from the external environment surrounding the computation, become spontaneously available. The values of these parameters, once received from the outside world, are then used to evaluate two functions, call them f_1 and f_2 , that tell us precisely *what* to do and *when* to do it, respectively.

The difficulty posed by this paradigm is that the evaluation of the two functions f_1 and f_2 is itself quite demanding computationally. Specifically, for a positive integer n , the two functions operate on n variables (the physical parameters). Only a parallel computer equipped with n processors can succeed in evaluating the two functions on time to meet the deadlines.

3.6 Computations Obeying Mathematical Constraints

There exists a family of computational problems where, given a mathematical object satisfying a certain property, we are asked to transform this object into another which also satisfies the same property. Furthermore, the property is to be maintained throughout the transformation, and be satisfied by every intermediate object, if any. More generally, the computations we consider here are such that every step of the computation must obey a certain predefined constraint. (Analogies from popular culture include picking up sticks from a heap one by one without moving the other sticks, drawing a geometric figure without lifting the pencil, and so on.)

3.6.1 Rewriting Systems

An example of such transformations is provided by *rewriting systems*. From an initial string ab , in some formal language consisting of the two symbols a and b , it is required to generate the string $(ab)^n$, for $n > 1$. Thus, for $n = 3$, the target string is $ababab$. The rewrite rules to be used are: $a \rightarrow ab$ and $b \rightarrow ab$. Throughout the computation, no intermediate string should have two adjacent identical characters. Such rewrite systems (also known as \mathcal{L} -systems) are used to draw fractals and model plant growth [48]. Here we note that applying any *one* of the two rules at a time causes the computation to fail (for example, if ab is changed to abb , by the first rewrite rule, or to aab by the second).

A sequential computer can change only one symbol at once, thereby causing the adjacency condition to be violated. By contrast, for a given n , a parallel computer with n processors can easily perform a transformation on all the inputs collectively. The required property is maintained leading to a successful computation. Thus, the string $(ab)^n$ is obtained in $\log n$ steps, with the two rewrite rules being applied simultaneously to all symbols in the current intermediate string, in the following manner: ab , $abab$, $abababab$, and so on. It is interesting to observe that a successful generation of $(ab)^n$ also provides an example of a rank-varying computational complexity (as described in Section 3.3). Indeed, each legal string (that is, each string generated by the rules and obeying the adjacency property) is twice as long as its predecessor (and hence requires twice as many operations to be generated).

3.6.2 Sorting Variant

A second example of computations obeying a mathematical constraint is provided by a variant to the problem of sorting. For a positive even integer n , where $n \geq 8$, let n distinct integers be stored in an array A with n locations $A[1]$, $A[2]$, \dots , $A[n]$, one integer per location. Thus $A[j]$, for all $1 \leq j \leq n$, represents the integer currently stored in the j th location of A . It is required to sort the n integers in place into increasing order, such that:

1. After step i of the sorting algorithm, for all $i \geq 1$, no three consecutive integers satisfy:

$$A[j] > A[j + 1] > A[j + 2]$$

for all $1 \leq j \leq n - 2$.

2. When the sort terminates we have:

$$A[1] < A[2] < \dots < A[n].$$

This is the standard sorting problem in computer science, but with a twist. In it, the journey is more important than the destination. While it is true that we are interested in the outcome of the computation (namely, the sorted array, this being the *destination*), in this particular variant we are more concerned with *how* the result is obtained (namely, there is a condition that must be satisfied throughout all steps of the algorithm, this being the *journey*). It is worth emphasizing here that the condition to be satisfied is germane to the problem itself; specifically, there are no restrictions whatsoever on the model of computation or the algorithm to be used. Our task is to find an algorithm for a chosen model of computation that solves the problem exactly as posed. One should also observe that computer science is replete with problems with an inherent condition on how the solution is to be obtained. Examples of such problems include: inverting a nonsingular matrix without ever dividing by zero, finding a shortest path in a graph without examining an edge more than once, sorting a sequence of numbers without reversing the order of equal inputs, and so on.

An *oblivious* (that is, input-independent) algorithm for an $n/2$ -processor parallel computer solves the aforementioned variant of the sorting problem handily in n steps, by means of predefined pairwise swaps applied to the input array A , during each of which $A[j]$ and $A[k]$ exchange positions (using an additional memory location for temporary storage) [2]. A sequential computer, and a parallel computer with fewer than $(n/2) - 1$ processors, both fail to solve the problem consistently, that is, they fail to sort all possible $n!$ permutations of the input while satisfying, at every step, the condition that no three consecutive integers are such that $A[j] > A[j + 1] > A[j + 2]$ for all j . In the particularly nasty case where the input is of the form

$$A[1] > A[2] > \dots > A[n],$$

any sequential algorithm and any algorithm for a parallel computer with fewer than $(n/2) - 1$ processors fail after the first swap.

It is interesting to note here that a Turing Machine with $n/2$ heads succeeds in solving the problem, yet its simulation by a standard (single-head) Turing Machine fails to satisfy the requirements of the problem. Indeed, suppose that the standard Turing Machine is presented with the input sequence $A[1] > A[2] > \dots > A[n]$:

1. It will either use the given representation of the input, and proceed to perform an operation (a swap, for example), in which case it would fail after one step of the algorithm,
2. Or it will transform the given representation into a different encoding (perhaps one intended to capture the behavior of the Turing Machine with $n/2$ heads) in preparation for the sort, in which case it would again fail since the transformation itself will consist of more than one algorithmic step.

This is a surprising result as it goes against the common belief that any computation by a variant of the Turing Machine can be effectively simulated by the standard model [37].

3.7 The Universal Computer Is A Myth

The Principle of Simulation is the cornerstone of computer science. It is at the heart of most theoretical results and practical implements of the field such as programming languages, operating systems, and so on. The principle states that *any* computation that can be performed

on *any* one general-purpose computer can be equally carried out through simulation on *any* other general-purpose computer [22, 29, 41]. At times, the imitated computation, running on the second computer, may be faster or slower depending on the computers involved. In order to avoid having to refer to different computers when conducting theoretical analyses, it is a generally accepted approach to define a model of computation that can simulate *all* computations by other computers. This model would be known as a Universal Computer \mathcal{U} . Thus, Universal Computation, which clearly rests on the Principle of Simulation, is also one of the foundational concepts in the field [21].

Our purpose here is to prove the following general statement: There does not exist a *finite* computational device that can be called a Universal Computer. Our reasoning proceeds as follows. Suppose there exists a Universal Computer capable of n elementary operations per step, where n is a finite and fixed integer. This computer will fail to perform a computation *requiring* n' operations per step, for any $n' > n$, and consequently lose its claim of universality. Naturally, for each $n' > n$, another computer capable of n' operations per step will succeed in performing the aforementioned computation. However, this new computer will in turn be defeated by a problem requiring $n'' > n'$ operations per step.

This reasoning is supported by each of the computational problems presented in Sections 3.1–3.6. As we have seen, these problems *can* easily be solved by a computer capable of executing n operations at every step. Specifically, an n -processor parallel computer led to a successful computation in each case. However, *none* of these problems is solvable by any computer capable of at most $n - 1$ operations per step, for any integer $n > 1$. Furthermore, the problem size n itself is a variable that changes with each problem instance. As a result, *no* parallel computer, regardless of how many processors it has available, can cope with a growing problem size, as long as the number of processors is finite and fixed. This holds even if the computer purporting to be universal is endowed with an unlimited memory and is allowed to compute for an indefinite amount of time.

The preceding reasoning applies to any computer that obeys the *finiteness condition*, that is, a computer capable of only a finite and fixed number of operations per step. It should be noted that computers obeying the finiteness condition include all “reasonable” models of computation, both theoretical and practical, such as the Turing Machine, the Random Access Machine, and other idealized models [50], as well as all of today’s general-purpose computers, including existing conventional computers (both sequential and parallel), as well as contemplated unconventional ones such as biological and quantum computers [8].

Therefore, the Universal Computer \mathcal{U} is clearly a myth. As a consequence, the Principle of Simulation itself (though it applies to most *conventional* computations) is, in general, a fallacy. In fact, the latter principle is responsible for many other myths in computing, such as the *Speedup Theorem*, the *Slowdown Theorem*, and *Amdahl’s Law*. Counterexamples for dispelling these and other myths are presented in [3, 19].

4 Misconceptions and Replies

“Almost all the other fellows do not look from the facts to the theory but from the theory to the facts; they cannot get out of the network of already accepted concepts; instead, comically, they only wriggle about inside.” [24]

What follows are thirteen misconceptions relating to my nonuniversality result, and responses to them. They are presented as a dialog with an interlocutor who, I assume (perhaps

wishfully), has read my papers listed in the bibliography on the myth of universal computation.

4.1 Misconception 1

“So, you describe a number of functions that are uncomputable. What’s new about that? Uncomputable functions have been known since the time of Turing.”

Response: This is incorrect. The functions that I describe are eminently computable. In my papers listed in the bibliography, every function F^n of n variables can be easily evaluated by a computer capable of at least n elementary operations per time unit (an n -processor parallel computer [13], for example). However, a computer capable of only $n - 1$ or fewer elementary operations per time unit cannot compute F^n . Nonuniversality in computation immediately follows by simple induction.

4.2 Misconception 2

“You are proposing computations that cannot be simulated efficiently. What’s new about that? Your own book [2] and your earlier papers presented such computations that can be performed in constant time by n processors, but require time exponential in n when executed on $n - 1$ or fewer processors.”

Response: The error in the statement above is in the phrase “cannot be simulated efficiently”. Indeed, my nonuniversality result does not follow from computations that cannot be simulated efficiently. It follows from computations that *cannot be simulated at all*. Thus, for each of the functions F^n of n variables that I describe in my papers listed in the bibliography, no computer capable of fewer than n elementary operations per time unit can simulate the actions of a computer capable of n elementary operations per time unit. The latter computer is capable of evaluating F^n successfully, the former is not capable of doing so, even if given an infinite amount of time and memory space. It is this impossibility of simulation that leads to nonuniversality.

4.3 Misconception 3

“What do you think about this way of attacking your time-varying variables problem:

First of all, we should separate the tasks into sensing and computation tasks. That means, I assume that there are sensors equipped with memory which read the values of the variables at time t and write them into appropriate memory locations. So, for every t we have memory locations $x_1(t)$, $x_2(t)$, \dots , $x_n(t)$. Furthermore, assume that the values of n and t are stored in some other variables. Then, a universal machine would read the value of n (and t), read the values from memory, and perform the requested computations on these values, which is a rather simple task.

So, the main argument of this is that we allow for this separation. The sensors are peripheral components, which do not perform any computations, but ‘just’ provide the input for the computation, that is, they sense their values simultaneously at time t and store them in their respective memory locations, and the computations can access these values later on.

I guess that you will argue that this requires n values to be stored in the memory of the machine at the same time, which would contradict a specification of a machine which is independent of n . But if we can reduce the problem you posed to this separation of

concerns, we would have the consistence with the traditional theory of computation except for the addition of these sensors, which are not considered to be a part of the universal computing device but of the input specification.”

Response: Surely, you cannot “separate” part of the universal computer (in this case the input unit) from the rest of the computer just to fit the problem. The universal computer is one unit, and a computational step is: [read, calculate, write].

The definition of ‘to compute’ as ‘the process of transforming information’ applies to all three phases of computation, namely,

1. The input phase, where data such as keystrokes, mouse clicks, or temperatures are reduced, for example, to binary strings;
2. The calculation phase, where data are manipulated through arithmetic and logic operations, as well as operations on strings, and so on; and
3. The output phase, where data are produced as numbers on a screen, or rendered as images and sounds, for instance.

Each of the three phases represents information transformation; each is an integral part of every computation, and no computation worthy of that name can exclude any of them. In particular, input and output are fundamental in the design and analysis of algorithms. Input-sensitive and output-sensitive computations often play an especially important role in deriving lower and upper bounds on algorithmic complexity.

One of the most dramatic illustrations of the interweaving of input and output with calculation is the well-known linear-time planarity testing algorithm [31]. In order to determine whether a graph with V vertices is planar, Step 1 of that algorithm avoids a potentially quadratic-time computation, by reading in the edges of the graph *one at a time* from the outside world; if there is one more edge than $3V - 6$, the graph is declared nonplanar.

But let’s assume you have set up n sensors and succeeded in solving the problem. What happens when you discover, the next morning, that there are now, not n , but $n + 1$ inputs? Do you think it is a fair solution for you to go to the sensor shop, buy one more sensor and rush back to attach it to the extra input source and to the computer? And even if you did, what if by the time you return, the time t_0 at which the result is needed would have passed?

4.4 Misconception 4

“Your argument is great. It is simple and effective. However, it doesn’t show that there is something *theoretically* flawed with the concept of a universal computer; it shows that a universal computer could never be physically realized. So the Church-Turing thesis is alright if we take it that the space/time needed is infinite. Having said that, I imagine that the distinction between the theoretical and the implementation claim is often overlooked, which makes the proof integral in making that distinction very sharp indeed.”

Response: Thank you for reading. But, sorry, I beg to differ. Even if given infinite space and infinite time (which we allow the Turing Machine anyway), no computer that you define once and for all (and are not allowed to change ever again) can solve the problems that I define.

The issue is not with infinite space and infinite time. The issue is: How many operations can you do in one slice of time? You have to define this for every theoretical (and of

course practical) computer. Otherwise, analysis of algorithms becomes meaningless, and the running time of a program a void concept. For example, the Turing Machine can do one, two, or three operations per slice of time, depending on your definition, but it is always a finite number. It can read a symbol, then change state, or write a symbol, or move left or right or stay put. Then a new iteration starts. It cannot do an arbitrary number of these fundamental operations per slice of time (you may call the latter a step, an iteration, a time unit, whatever). In passing, I should say that it is this finiteness of the number of operations per time unit that caused the great success of computer science, making the Computer the greatest invention of the 20th century: A machine designed once and for all that can simulate any operation by another machine. In theory, you should be able to buy a computer and use it for the rest of your life, for it will never be faced with a computable function that it cannot compute. Or so we thought . . .

Suppose you have defined your machine once and for all (it can be a Turing Machine or a Supercomputer, or the chip in your digital camera or that in your toaster). I will now give you a computable function that it fails to compute. One example of such a computation is the problem of operating on variables that change with time (as described in Section 3.1). Even if given all the time in the world from here to eternity, even if given all the space in the known Universe and beyond, you could not solve the problem, not even in theory with pen and paper. Why? Because you defined your machine and fixed it once and for all (as one should if one claims to have a Universal Computer that can simulate anything that another computer can compute). And herein lies the Achilles heel of the Universal Computer: I, as a designer of the problem, ask you for a definition of your purported Universal Computer; I then concoct a computation that thwarts it. This computation, of course, can be easily carried out by another “Universal Computer”, but then I’ll give you another problem, and so on.

There is nothing here about implementation. I put no limits on what your machine can do, except for one: It must do a finite number of operations per time unit. This number can be as big as you want, it can even be defined by a function of time, but it must be fixed the moment you define your model, and it cannot be infinite. Once you define it, you cannot change it. We can play the game with paper and pencil. I will always win.

I believe that the result is perfectly *theoretical* as well as being perfectly *practical*.

4.5 Misconception 5

“While I could not make much sense of your “proof” (right off the bat, I do not think ‘spacially and temporally connected variables’ is well defined, although I also think it suffers from many other flaws) I decided to take your challenge seriously, within the bounds and the scope of your restrictions, language and assumptions, in an attempt to define a reasonable sounding computational device that would solve the time-varying variables computation. I will assume that the variables appear each on a luminous display of some sort.

I am going to assume relativity (please let me know if spacially and temporally connected variables is supposed to mean something else?). There is space k between displays so that, when they update their values, the light from their screens arrive at my computational device at nearly, but not precisely the same time (determined by k and where my device is in the ‘room’). I will move my computational device close to the displays so that I can exploit the hypotenuse the signals from each display must travel first to reach the sensor of my device. Then, it is simply a matter of setting the processing speed to a finite value around ‘a constant

$\times k$ meters / the speed of light.’ In this way, my computational device is able to compute the function F^n as desired.

My device has one sensor. It only needs one sensor, since the signals from the displays come in sequentially. Remember that the displays are separated by a distance and the signal from each display, after they update, must travel this distance (at the speed of light). So my device uses its one sensor to read in a display and then it reuses that sensor some time later after the signal from the next display has reached it (some time later).”

Response: Right off the bat, as is clear from my papers, the spatial and temporal relationship among the variables is such that:

1. The variables occupy the same region in physical space; specifically, the distance separating any two variables is no larger than a small constant (whose magnitude depends on the general paradigm under consideration);
2. The variables are constrained by a unique parameter representing true physical time.

Now turning to your solution, I would say it is original; unfortunately, I am afraid you are missing the point of the exercise.

The problems I pose are to be treated independently of specific physical properties. For example, the time-varying variables could be anything one would want them to be (atmospheric pressures for instance, or humidity readings, etc.), and they should not necessarily be on display (they would need to be acquired, that is, measured, first). Light rays, the basis of your argument, should play no role in the solution (for they may not help in general).

However, let’s consider your one-sensor solution, assuming the variables are indeed displayed. The difficulty with such setups is that they inevitably break down at some point as the problem scales up. Specifically,

1. The dimensions of the sensor are fixed.
2. As the number of variables n , and hence displays, grows, the angle of the light rays from the furthest display approaches 0. There will not be enough real estate on the sensor for that light to impinge upon.
3. Thus, the problem poser can make n sufficiently large so as to render the sensor useless.

And one more thing: How does your sensor handle multiple simultaneous (or perhaps overlapping) inputs from displays all equidistant from it?

4.6 Misconception 6

“Your definition of computation is unconventional.”

Response: Absolutely, if one considers unconventional the passage of time, or the interactions among the constituents of the universe subject to the laws of nature, or for that matter any form of information processing. In any case, my definition of computation may be unconventional, but it is not unrealistic.

Besides, it is important to realize that defining “to compute” as “what the Turing Machine does”, which is quite commonly done (see, for example, [25]), leads to a logical reasoning that is viciously circuitous. If computation is what the Turing Machine does, then clearly

the Turing Machine can compute anything that is computable. Furthermore, the “Church-Turing Thesis” would move immediately from the realm of “conjecture” to that of “theorem”. The fact that it has not done so to this day, is witness to the uncertainty surrounding the widespread definition of “computation”. As stated above, my counterexamples, by contrast, *disprove* the “Church-Turing Thesis”.

4.7 Misconception 7

“But the Turing Machine was not meant for this kind of computation.”

Response: Precisely. Furthermore, the nonuniversality result covers all computers, not just the Turing Machine.

4.8 Misconception 8

“Abstract models of computation do not concern themselves with input and output.”

Response: This opinion is held by those who believe that computation is the process that goes from input to output, while concerning itself with neither input nor output (see, for example, [25] cited above). By analogy, one might say that eating is all about digestion, and a Moonlight Sonata interpretation is nothing but the hitting of piano keys. Perhaps.

Perhaps not. A model of computation is useful to the extent that it is a faithful reflection of reality, while being mathematically tractable. In computer science, input and output are not cumbersome details to be ignored; they are fundamental parts of the computation process, which must be viewed as consisting of three essential phases, namely, input, calculation, and output. A computer that does not interact with the outside world is useless. In that sense, the thermostat in your house is more powerful than the Turing Machine. Please remember that my result goes beyond the Turing Machine (which, by the way, is a very primitive model of computation). To ask computer scientists to stick to the Turing Machine and not to look beyond, would be as if physics stopped progressing beyond the knowledge of the Ancient Greeks.

In fact, looking ahead, if computers of the future are to be *quantum* in nature, their main challenge is expected to be, not calculation but, input and output.

4.9 Misconception 9

“Classical computability theory does not include variables that change with time.”

Response: This is a serious lacuna in classical theory. My work shows that there are in fact many such lacunae. Their result is to severely restrict the definition of computation. Indeed, to define computation merely as function evaluation, with fixed input and fixed output, is unrealistic and naive. It also trivializes the Church-Turing Thesis (turning it into a tautology), because it necessarily leads to the kind of sterile circular reasoning mentioned in the response to Misconception 6.

Having said that, the time-varying variables counterexample is only one of many such refutations of universality in computation. Other examples arise in the computations described in section 3 (and detailed in [6] – [18], [26], [42] – [47]).

4.10 Misconception 10

“The Church-Turing Thesis applies only to classical computations.”

Response This is certainly not the case. As the multitude of examples listed in [7] amply demonstrate, the commonly accepted statement of the Church-Turing Thesis [36] is essentially this: “There exists a Universal Computer capable of performing any computation that is possible on any other computer.” There are no restrictions, exceptions, or caveats whatsoever on the definition of computation. In fact, a typical textbook definition of computation is as follows: “A computation is a process that obeys finitely describable rules.” [49]. What’s more, it is suggested in every textbook on the subject that, thanks to the fundamental and complementary notions of simulation and universality, every general-purpose computer is universal: A Turing Machine, a Random-Access Machine, a Personal Computer, a Supercomputer, the processing chip in a cell phone, are all universal. (My result shows this claim to be *false*.)

Going a little further, many authors consider all processes taking place in the Universe as computations. Interested readers may consult [16, 20, 22, 23, 27, 35, 38, 39, 49, 51, 52, 55, 56, 58, 60, 61, 62, 63, 64].

I happen to agree with the above authors on the pervasive nature of computation. However, in order to reach my conclusion about nonuniversality in computation, I do not in fact need to go that far. My counterexamples use very simple computations to refute universality: computations whose variables change with time, computations whose variables affect one another, computations the complexity of whose steps changes with the passage of time, computations the complexity of whose steps depends (not on the time but) on the order of execution of each step, computations that involve variables obeying certain mathematical conditions, and so on. These are not uncommon computations. They may be considered unconventional only when contrasted with the standard view of computation as a rigid function evaluation, in which, given a variable x , it is required to compute $f(x)$, in isolation of the rest of the world.

4.11 Misconception 11

“Actually, Alan Turing knew that computations involving physical time would cause problems for his machine.”

Response: This is the back-pedaling argument *par excellence*. Perhaps Turing knew about the computations that I describe here, but I doubt it.

Indeed, Turing did propose the Oracle machine (o-machine) [57] as an (unconventional) extension to his a-machine (today known as the Turing Machine). However, the o-machine has nothing to do with the problems I propose to counter universality in computation. Furthermore, the o-machine is a fanciful creation that appeals to some form of divine intervention in order to solve the problems it faces, while the computations that I use to prove nonuniversality are eminently executable on everyday computers, provided the latter are capable of performing the proper number of basic operations per time unit.

The Church-Turing Thesis [36] is proof that both Turing and Church were convinced of the universality of the Turing Machine. In any case, I am not aware of any writings by Turing, von Neumann, or anybody else, that hint to nonuniversality in computation, prior to my January 2005 paper on the myth of the universal computer.

This type of argument, exemplified by Misconception 11, reminds me of the words of the American philosopher William James:

“First, you know, a new theory is attacked as absurd; then it is admitted to be true, but obvious and insignificant; finally it is seen to be so important that its adversaries claim that they themselves discovered it.” [34].

4.12 Misconception 12

“The problem I see with the nonuniversality claim is that it does not appear to be falsifiable. It might be improved by thinking further about how to make the claim more specific, clear, testable, well-defined and worked out into detail. The entire claim would ideally be described in a few very worked out, specific, sentences - needing little external exposition or justification.”

Response: Have you had a chance to read any of my papers? There is no lack of mathematical formalism there. For example in one of my papers, M. Nagy and I used quantum mechanics to illustrate the first five of the six aforementioned paradigms [45]. I am certain you will enjoy that one.

But let me try this since you would like something simple and crisp. How about the best understood problem in computer science, arguably that of sorting a sequence of numbers, but with a twist? See Section 3.6.2 for a description.

4.13 Misconception 13

“A fine student just presented your NonUniversality in Computation work in my graduate theory course. Going simply on her presentation, I am not impressed. You start by quoting Hopcroft and Ullman and saying that their statement is clearly false. You took this quote out of the context of all of theoretical computer science which clearly defines that a ‘computation’ is to start with the entire input presented and is given as much time as it wants to read this input and do its computation. It is true that since Turing, the nature of computation has changed to require real time interactions with the world. But you should not misrepresent past work. Having not studied your arguments at length, the only statement that I have gotten is that a machine is unable to keep up if you only allowed it to read in a fixed number of bits of information per time step and you throw at it in real time an arbitrarily large number of bits of information per time step. In itself, this statement is not very deep.”

Response: I am sure the student did a wonderful job, but somehow the message did not get across. I will take your remarks one by one.

Remark A: “You took this quote out of the context of all of theoretical computer science which clearly defines that a ‘computation’ is to start with the entire input presented and is given as much time as it wants to read this input and do its computation.”

Response to Remark A:

1. I did not find your definition of ‘computation’ anywhere. It is certainly not in Hopcroft and Ullman’s book [32].
2. Your definition of ‘computation’ applies narrowly to some primitive computational models, such as the Turing Machine, Cellular Automata, and so on.

3. Because of (2), your definition trivializes the Church-Turing Thesis, rendering it a tautology: By defining ‘computation’ as ‘what the Turing Machine does’, it obviously follows that ‘the Turing Machine can compute everything that is computable’. As mentioned earlier (see the response to Misconception 6), this is a typical example of circular reasoning.
4. Your definition is not sufficiently general to capture the richness and complexity of the notion of ‘computation’. Others have proposed more encompassing definitions of ‘computation’. Here are a few quotes:

“In a sense Nature has been continually computing the ‘next state’ of the universe for billions of years; all we have to do - and actually, all we can do - is ‘hitch a ride’ on this huge ongoing computation.” [56]

“Computation is physical; it is necessarily embodied in a device whose behaviour is guided by the laws of physics and cannot be completely captured by a closed mathematical model. This fact of embodiment is becoming ever more apparent as we push the bounds of those physical laws.” [54].

“Information sits at the core of physics ... every particle, every field of force, even the space-time continuum itself derives its function, its meaning, its very existence entirely ... from answers to yes-or-no questions ... all things physical are information-theoretic in origin ... ” [61].

“Think of all our knowledge-generating processes, our whole culture and civilization, and all the thought processes in the minds of every individual, and indeed the entire evolving biosphere as well, as being a gigantic computation. The whole thing is executing a self-motivated, self-generating computer program.” [22].

5. And one more thing about your point that a computation is to “start with the entire input presented”. Since *all* of my counterexamples indeed require the entire input to be “present”, I will assume that you mean “start with the entire input residing in the memory of the computer”. (Incidentally, some of my counterexamples assume the latter as well.) The relevant point here is this: Hopcroft himself has a well-known algorithm that makes no such assumption about the entire input residing in memory. As stated in response to Misconception 3 above, one of the most dramatic illustrations of the interweaving of input and output with calculation is the well-known linear-time planarity testing algorithm of Hopcroft and Tarjan [31]. In order to determine whether a graph with n vertices is planar, Step 1 of that algorithm avoids a computation that could potentially run in time at least quadratic in n , by reading in the edges of the graph one at a time from the outside world; if there is one more edge than the absolute maximum stipulated by Euler’s formula, namely $3n - 6$ (the paper actually uses the loose bound $3n - 3$ in order to allow for $n < 3$), the graph is declared nonplanar [31].

Remark B: “It is true that since Turing, the nature of computation has changed to require real time interactions with the world. But you should not misrepresent past work.”

Response to Remark B:

1. I am glad to see that you agree with me about the nature of computation. You should know, however, that my counterexamples to universality are not all about “real time

interaction with the world”. There is a list of such counterexamples in section 1, a brief description of some in section 3, and a list of references to my papers in the bibliography section. One counterexample involving mathematical constraints (it is a variant of sorting, in which the entire input is available in memory at the outset of the computation) is described in Section 3.6.2. Note also that my nonuniversality result applies to putative ‘universal computers’ capable of interaction with the outside world. These ‘universal computers’ are endowed with unlimited memories and are allowed an indefinite amount of time to solve the problems they are given. They all still fail the test of universality.

2. However, I do take exception to the claim that I misrepresented past work. Below are quotes from famous computer scientists asserting, without caveat, exception, or qualification that a universal computer is possible (often explicitly stating that such a universal computer is the Turing Machine, and essentially taking for granted the unproven Church-Turing Thesis):

“A Turing machine can do everything that a real computer can do.” [53], p. 125.

“Anything which can be computed can be computed by a Turing Machine.” [1], p. 123.

“It is theoretically possible, however, that Church’s Thesis could be overthrown at some future date, if someone were to propose an alternative model of computation that was publicly acceptable as fulfilling the requirement of ‘finite labor at each step’ and yet was provably capable of carrying out computations that cannot be carried out by any Turing machine. No one considers this likely.” [37], p. 223.

“...if we have shown that a problem can (or cannot) be solved by any TM, we can deduce that the same problem can (or cannot) be solved by existing mathematical computation models nor by any conceivable computing mechanism. The lesson is: Do not try to solve mechanically what cannot be solved by TMs!” [40], p. 152.

“...any algorithmic problem for which we can find an algorithm that can be programmed in some programming language, any language, running on some computer, any computer, even one that has not been built yet but can be built, and even one that will require unbounded amounts of time and memory space for ever-larger inputs, is also solvable by a Turing machine.” [29], p. 233.

“It is possible to build a universal computer: a machine that can be programmed to perform any computation that any other physical object can perform. ... any computation that can be performed by any physical computing device can be performed by any universal computer, as long as the latter has sufficient time and memory.” [30], pp. 63–64.

Hundreds of such statements can be found in the literature; for a sample see [7].

3. Finally, please note that to correct previous mistakes is not to misrepresent the past. This is how science advances. Newton, Darwin, and Einstein were giants who built great scientific edifices. Each edifice, magnificent yet incomplete.

Remark C: “Having not studied your arguments at length, the only statement that I have gotten is that a machine is unable to keep up if you only allowed it to read in a fixed

number of bits of information per time step and you throw at it in real time an arbitrarily large number of bits of information per time step. In itself, this statement is not very deep.”

Response to Remark C:

Here we go again. Please note:

1. As mentioned above, the “time-varying variables” computation is but one of many counterexamples to universality.
2. Surely you must know that a reasonable model of computation must be finite. The Turing Machine has a finite alphabet, a finite set of states, and a finite set of elementary operations per step. To assume otherwise would render the fields of complexity theory and algorithm design and analysis useless.
3. Your characterization of my result is erroneous. I describe *computable functions* that no machine that claims to be universal can compute.
4. Let’s try the following: *You* define a “universal computer” fulfilling the requirement of ‘finite labor at each step’ (to quote [37] again), and I will give you a computable function that it cannot compute.

5 Conclusion

“Facts do not go away when scientists debate rival theories to explain them.” [28]

In conclusion, I offer the reader the following challenge.

Computational Challenge: Anyone who still does not accept my result, namely, that universality in computation is a myth, has but one option: to prove it wrong. In order to do this, one must exhibit a universal computer capable of a finite and fixed number of operations per time unit, on which each one of the computations in the following classes, and described in [6] – [18], [26], [42] – [47], can be performed:

1. Computations with time-varying variables
2. Computations with time-varying computational complexity
3. Computations with rank-varying computational complexity
4. Computations with interacting variables
5. Computations with uncertain time constraints
6. Computations with global mathematical constraints.

It is important that the purported universal computer be able to execute successfully all aforementioned computations, since each one of them, by itself, is a counterexample to computational universality. For a simplified (and perhaps more colorful) version of the challenge, please see [5].

References

- [1] Abramsky, S. et al, *Handbook of Logic in Computer Science*, Clarendon Press, Oxford, 1992.
- [2] Akl, S.G., *Parallel Computation: Models and Methods*, Prentice Hall, Upper Saddle River, New Jersey, 1997.
- [3] Akl, S.G., Superlinear performance in real-time parallel computation, *The Journal of Supercomputing*, Vol. 29, No. 1, 2004, pp. 89–111.
- [4] Akl, S.G., Non-Universality in Computation: The Myth of the Universal Computer, School of Computing, Queen’s University, 2005–2013. <http://research.cs.queensu.ca/Parallel/projects.html>
- [5] Akl, S.G., A computational challenge, School of Computing, Queen’s University, 2005. http://www.cs.queensu.ca/home/akl/CHALLENGE/A_Computational_Challenge.htm
- [6] Akl, S.G., The myth of universal computation, Technical Report No. 2005-492, School of Computing, Queen’s University, Kingston, Ontario, January 2005, <http://www.cs.queensu.ca/home/akl/techreports/universal.pdf>; also Chapter 7 in: *Parallel Numerics*, Trobec, R., Zinterhof, P., Vajteršic, M., and Uhl, A., Eds., University of Salzburg, Salzburg, Austria and Jozef Stefan Institute, Ljubljana, Slovenia, 2005, pp. 211-236.
- [7] Akl, S.G., Universality in computation: Some quotes of interest, Technical Report No. 2006-511, School of Computing, Queen’s University, Kingston, Ontario, April 2006, 13 pages. <http://www.cs.queensu.ca/home/akl/techreports/quotes.pdf>
- [8] Akl, S.G., Three counterexamples to dispel the myth of the universal computer, *Parallel Processing Letters*, Vol. 16, No. 3, September 2006, pp. 381–403.
- [9] Akl, S.G., Conventional or Unconventional: Is Any Computer Universal?, Chapter 6 in: *From Utopian to Genuine Unconventional Computers*, A. Adamatzky and C. Teuscher, Eds., Luniver Press, Frome, United Kingdom, 2006, pp. 101–136.
- [10] Akl, S.G., Gödel’s incompleteness theorem and nonuniversality in computing, *Proceedings of the Workshop on Unconventional Computational Problems*, Sixth International Conference on Unconventional Computation, Kingston, Canada, August 2007, pp. 1–23.
- [11] Akl, S.G., Even accelerating machines are not universal, *International Journal of Unconventional Computing*, Vol. 3, No. 2, 2007, pp. 105–121.
- [12] Akl, S.G., Unconventional computational problems with consequences to universality, *International Journal of Unconventional Computing*, Vol. 4, No. 1, 2008, pp. 89-98.
- [13] Akl, S.G., Evolving Computational Systems, Chapter 1 in: *Parallel Computing: Models, Algorithms, and Applications*, Rajasekaran, S. and Reif, J.H., Eds., Taylor and Francis, CRC Press, Boca Raton, Florida, 2008, pp. 1–22.

- [14] Akl, S.G., Ubiquity and simultaneity: The science and philosophy of space and time in unconventional computation, Keynote address, *Conference on the Science and Philosophy of Unconventional Computing*, The University of Cambridge, Cambridge, United Kingdom, 2009.
- [15] Akl, S.G., Time travel: A new hypercomputational paradigm, *International Journal of Unconventional Computing*, Vol. 6, No. 5, 2010, pp. 329–351.
- [16] Akl, S.G., What is computation?, Technical Report No. 2013-609, School of Computing, Queen’s University, August 2013.
<http://www.cs.queensu.ca/home/akl/techreports/computation.pdf>
- [17] Akl, S.G. and Nagy, M., Introduction to Parallel Computation, Chapter 2 in: *Parallel Computing: Numerics, Applications, and Trends*, Trobec, R., Vajteršic, M., and Zinterhof, P., Eds., Springer-Verlag, London, England, 2009, pp. 43–80.
- [18] Akl, S.G. and Nagy, M., The Future of Parallel Computation, Chapter 15 in: *Parallel Computing: Numerics, Applications, and Trends*, Trobec, R., Vajteršic, M., and Zinterhof, P., Eds., Springer-Verlag, London, England, 2009, pp. 471–510.
- [19] Akl, S.G. and Yao, W., Parallel computation and measurement uncertainty in non-linear dynamical systems, *Journal of Mathematical Modelling and Algorithms*, Special Issue on Parallel and Scientific Computations with Applications, Vol. 4, 2005, pp. 5–15.
- [20] Davies, E.B., Building infinite machines, *British Journal for Philosophy of Science*, Vol. 52, 2001, pp. 671–682.
- [21] Davis, M., *The Universal Computer*, W.W. Norton, 2000.
- [22] Deutsch D., *The Fabric of Reality*, Penguin Books, London, England, 1997.
- [23] Durand-Lose, J., Abstract geometrical computation for black hole computation, Research Report No. 2004-15, Laboratoire de l’Informatique du Parallélisme, École Normale Supérieure de Lyon, Lyon, France, April 2004.
- [24] Einstein, A., in a letter to Erwin Schrödinger, August 8, 1935, quoted in: Gilder, L., *The Age of Entanglement*, Vintage Books, New York, New York, 2009, p. 170.
- [25] Fortnow, L., The enduring legacy of the Turing machine, <http://ubiquity.acm.org/article.cfm?id=1921573>
- [26] Fraser, R. and Akl, S.G., Accelerating machines: a review, *International Journal of Parallel Emergent and Distributed Systems*, Vol. 23, No. 1, February 2008, pp. 81–104.
- [27] Gleick, J., *The Information: A History, a Theory, a Flood*, HarperCollins, London, England, 2011.
- [28] Gould, S.J., Evolution as fact and theory, *Discover*, Vol. 2, May 1981, pp. 34–37.
- [29] Harel, D., *Algorithmics: The Spirit of Computing*, Addison-Wesley, Reading, Massachusetts, 1992.

- [30] Hillis, D., *The Pattern on the Stone*, Basic Books, New York, New York, 1998.
- [31] Hopcroft, J. and Tarjan R., Efficient planarity testing, *Journal of the ACM*, Vol. 21, No. 4, 1974, pp. 549–568.
- [32] Hopcroft, J.E. and Ullman, J.D., *Formal Languages and their Relations to Automata*, Addison-Wesley, Reading, Massachusetts, 1969.
- [33] Hypercomputation, <http://en.wikipedia.org/wiki/Hypercomputation>
- [34] James, W., *Pragmatism: A new name for some old ways of thinking*, Lecture 6: Pragmatism’s conception of truth, Longman Green and Co., New York, New York, 1907, pp. 76–91.
- [35] Kelly, K. God is the machine, *Wired*, Vol. 10, No. 12, 2002.
- [36] Kleene, S.C., *Introduction to Metamathematics*, North Holland, 1952.
- [37] Lewis, H.R. and Papadimitriou, C.H., *Elements of the Theory of Computation*, Prentice Hall, Englewood Cliffs, New Jersey, 1981.
- [38] Lloyd, S., *Programming the Universe*, Knopf, New York, New York, 2006.
- [39] Lloyd, S. and Ng, Y.J., Black hole computers, *Scientific American*, Vol. 291, No. 11, November 2004, pp. 53–61.
- [40] Mandrioli, D. and Ghezzi, C., *Theoretical Foundations of Computer Science*, John Wiley, New York, New York, 1987.
- [41] Minsky, M.L., *Computation: Finite and Infinite Machines*, Prentice-Hall, 1967.
- [42] Nagy, M. and Akl, S.G., On the importance of parallelism for quantum computation and the concept of a universal computer, *Proceedings of the Fourth International Conference on Unconventional Computation*, Sevilla, Spain, October 2005, LNCS 3699, pp. 176-190.
- [43] Nagy, M. and Akl, S.G., Quantum measurements and universal computation, *International Journal of Unconventional Computing*, Vol. 2, No. 1, 2006, pp. 73–88.
- [44] Nagy, M. and Akl, S.G., Quantum computing: Beyond the limits of conventional computation, *International Journal of Parallel, Emergent and Distributed Systems*, Special Issue on Emergent Computation, Vol. 22, No. 2, April 2007, pp. 123–135.
- [45] Nagy, M. and Akl, S.G., Parallelism in quantum information processing defeats the Universal Computer, *Proceedings of the Workshop on Unconventional Computational Problems*, Sixth International Conference on Unconventional Computation, Kingston, Canada, August 2007, pp. 25–52; also in: *Parallel Processing Letters*, Special Issue on Unconventional Computational Problems, Vol. 17, No. 3, September 2007, pp. 233 - 262.

- [46] Nagy, N. and Akl, S.G., Computations with uncertain time constraints: Effects on parallelism and universality, *Proceedings of the Tenth International Conference on Unconventional Computation*, Turku, Finland, June 2011, LNCS 6714, pp. 152–163.
- [47] Nagy, N. and Akl, S.G., Computing with uncertainty and its implications to universality, *International Journal of Parallel, Emergent and Distributed Systems*, Vol. 27, Issue 2, April 2012, pp. 169–192.
- [48] Prusinkiewicz, P. and Lindenmayer, A. *The Algorithmic Beauty of Plants*, Springer Verlag, 1990.
- [49] Rucker, R., *The Lifebox, the Seashell, and the Soul*, Thunder’s Mouth Press, New York, New York, 2005.
- [50] Savage, J.E., *Models of Computation*, Addison-Wesley, 1998.
- [51] Seife, C., *Decoding the Universe*, Viking Penguin, New York, New York, 2006.
- [52] Siegfried, T., *The Bit and the Pendulum*, John Wiley & Sons, Inc., New York, New York, 2000.
- [53] Sipser, M., *Introduction to the Theory of Computation*, PWS Publishing Company, Boston, Massachusetts, 1997.
- [54] Stepney, S., The neglected pillar of material computation, *Physica D*, Vol. 237, No. 9, 2004, pp. 1157–1164.
- [55] Tipler, F.J., *The Physics of Immortality: Modern Cosmology, God and the Resurrection of the Dead*, Macmillan, London, England, 1995.
- [56] Toffoli, T., Physics and Computation, *International Journal of Theoretical Physics*, Vol. 21, 1982, pp. 165–175.
- [57] Turing, A.M., Systems of logic based on ordinals, *Proceedings of the London Mathematical Society*, Series 2, Vol. 45, 1939, pp. 161–228.
- [58] Vedral, V., *Decoding Reality*, Oxford University Press, Oxford, England, 2010.
- [59] Wegner, P. and Goldin, D., Computation beyond Turing Machines, *Communications of the ACM*, Vol. 46, No. 4, May 1997, pp. 100–102.
- [60] Wheeler, J.A., Information, physics, quanta: The search for links, *Proceedings of the Third International Symposium on Foundations of Quantum Mechanics in Light of New Technology*, Tokyo, Japan, 1989, pp. 354–368.
- [61] Wheeler, J.A., Information, physics, quantum: The search for links, in: *Complexity, Entropy, and the Physics of Information*, Zurek, W., Ed., Addison-Wesley, Redwood City, California, 1990.
- [62] Wheeler, J.A., *At Home in the Universe*, American Institute of Physics Press, Woodbury, New York, 1994.

- [63] Wolfram, S., *A New Kind of Science*, Wolfram Media, Champaign, Illinois, 2002.
- [64] Zuse, K., *Calculating Space*, MIT Technical Translation AZT-70-164-GEMIT, Massachusetts Institute of Technology (Project MAC), Cambridge, Massachusetts, 1970.