

Hypercomputation is Experimentally Irrefutable

Mike Stannett

Verification and Testing Research Group

Department of Computer Science, Regent Court, University of Sheffield

211 Portobello Street, Sheffield S1 4DP, United Kingdom

Email. m.stannett@dcs.shef.ac.uk

Abstract

This is the first in a series of discussion papers concerning various apparent relationships between hypercomputation and physics. In this paper, we demonstrate two logical interpretations of the statement “every unstable particle must eventually decay,” and argue that no finitely resourced physical experiment can distinguish between them. We demonstrate that one of these interpretations supports the construction of a hypercomputational device, while the other is a physical tautology (it is always true in the Standard Model). It follows that any physical experiment proposed as a refutation of hypercomputation would also refute the Standard Model itself. Therefore, the falsity of hypercomputation is logically inconsistent with the Standard Model. Hypercomputation is either logically entailed by, or else logically independent of, the Standard Model. We conclude this paper with the question we address in the next – how can a system be hypercomputational if its observable behaviour is functionally equivalent to that of a finite recursive system?

Introduction

The principles of modern physics have allowed researchers to establish a strikingly elegant model of natural entities, called simply the *Standard Model* (see *e.g.* [Loez96]) in which objects are constructed from a small family of elementary particles whose interactions at the smallest scales play a role in determining the large-scale structure of cosmological space-time, a model whose experimental predictive power is a by-word for exactitude. While it is accepted that certain aspects of the model are incomplete, it is widely held to be “essentially correct”. Against this backdrop, it is a brave (or foolish) author who would propose that the Standard Model may, in fact, have little or nothing to say about hypercomputation and its feasibility, where “hypercomputation” refers to the construction, design, specification and analysis of physical devices and systems whose behaviours are non-recursive (*i.e.* incapable of simulation by Turing machine, *see* [Davi01, BaiJ00, ZhoW00, CopS99, Sieg99] for a range of recent viewpoints). Nonetheless, I shall attempt to show that “scientific method” precludes the refutation of hypercomputation, and that the standard model may indeed be logically independent of it.

At the heart of the Standard Model lies the axiomatic assumption that experimental observation is the only meaningful evidence of existence, where observation is itself meaningful only if it satisfies the constraints of “scientific method”. In particular, an experimental result is only considered truly valid if it can be replicated by independent parties. By definition, replication of an experiment requires that it be recursively defined – independent scientists should be able to repeat the same experimental steps in the same order, starting from the same carefully prepared initial conditions, and so generate the same results, a concept that is only meaningful if the experiment is viewed as a

recursive operation to be conducted on a recursively recreated environment. Since the constraints of scientific method amount to a requirement of recursive replicability, the basic framework of physics places implicit faith in the completeness of recursive methods for determining physical reality. The Standard Model is therefore an intrinsically recursive representation of Nature, and it is not surprising that physical hypercomputation should be so difficult to demonstrate – the very language of modern physics precludes direct analysis of hypercomputational processes. It does not follow, however, that physics has nothing to say about hypercomputation. It means simply that any such statement is likely to rely on subtle arguments, and will probably require a deal of serendipity and lateral thinking to establish.

At the Hypercomputation Workshop at University College London in May 2000, I described a thought-experiment (essentially that of [Stan91]) which uses principles of radioactive decay to implement an infinite random number generator (this is outlined below). As we explain now, any such system must be hypercomputational. At the heart of our system lay the assumption that, given any radioactive sample, there must eventually come a time when half of that sample has decayed – the average time required is, of course, just the half-life of the sample in question. Further reflection shows, however, that this assumption is itself extremely subtle. In trying to express the assumption logically we find that two radically different interpretations are possible, and comparing the two interpretations allows us to deduce that: *either* our original assumption was correct, and hypercomputation is feasible; *or else* our assumption was flawed, but in this case we can show that hypercomputation cannot be refuted by any experiment conducted according to the rules of the Standard Model (*i.e.*, the two concepts are logically independent).

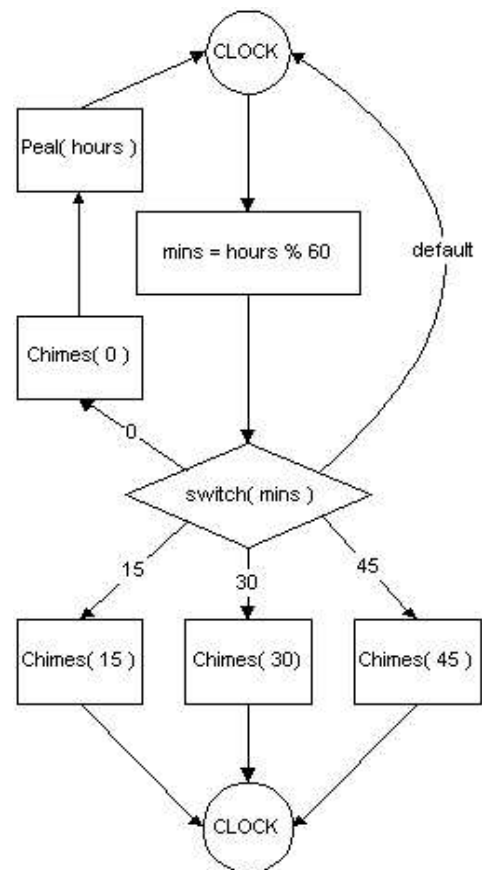
Random number generation and König’s Lemma

First, let us explain why true random number generators must be hypercomputational.

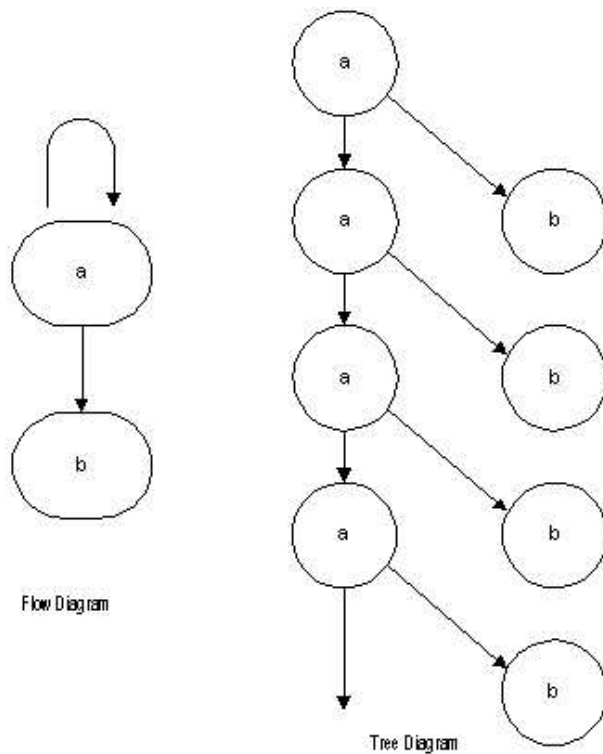
Any computer program written in a procedural language can be represented as a flow-diagram (this is the basis of “automaton-based” models of computation). Each flow-diagram (for example, a scheduling programme for Big Ben’s *Westminster Chimes*, right) contains a finite number of nodes (boxes) variously connected to one another by finitely many arrows. Any route that follows the arrows from box to box along the arrows is called a *path* through the program, and represents a “contiguous chunk” of programming.

Flow diagrams of this kind typically contain loops of one sort or another, and these can confuse the issue unnecessarily. To get round this problem we can imagine “straightening out” the diagram to obtain a tree diagram.

This will typically have infinitely many nodes, but each node will still be attached to at most finitely many arrows – technically, we call the new diagram a *finitely-branching tree*. For example, a simple two-node flow-chart might be “straightened out” as shown below. Notice that the behaviour



Westminster Chimes



represented is the same for both diagrams. At each point, the system is either in state b and can progress no further, or else is in state a and has the choice of returning to state a or moving to state b .

This example is particularly simple, but one can perform the same “straightening out” process on any flow chart to generate a finitely-branching tree exhibiting the same behaviour. [See e.g. [Miln89], where this process provides a foundation for formal concurrent semantics] In the general situation the “ b ” states correspond to (typically different) *terminal* states, those states in which the program can meaningfully be considered to have finished running, while the “ a ” states correspond to decisions and instructions in the main body of the program.

It is precisely because we can “straighten out” flow diagrams in this way that a mathematical result known as *König’s Lemma* is important to hypercomputation theory. König’s Lemma states that *any finitely-branching tree which*

contains infinitely many terminal nodes must also contain an infinite path. For example, the tree diagram shown left is finitely branching and contains infinitely many terminal nodes (the b s). According to König’s Lemma, it should also contain an infinite path – and it does. If we move down the a s one at a time, we follow a path comprising infinitely many arrows. Looking back at the original flow-chart, we see that this behaviour corresponds to cycling at state a , and represents a *non-terminating behaviour*.

This is quite general – any infinite path through any program’s flow diagram represents a non-terminating behaviour for the program in question. Its presence tells us that the program in question is capable of “running forever” without terminating. Since a recursive program’s flow-diagram is always finitely-branching, while the behaviours of flow diagrams and their corresponding tree diagrams are identical, König’s Lemma tells us that whenever a recursive program can terminate in any one of infinitely many different ways, then it is also possible for that program to run forever and never terminate.

In particular, then, suppose we manage to write an infinitely random number generator. This is a program which is guaranteed to generate a randomly selected value from the set $\equiv \{0, 1, 2, \dots\}$. We can regard the generation of the selected number as “successful termination” of a program, and by definition there are infinitely many different ways in which this program can terminate. If this program were recursive, König’s Lemma would entail the existence of at least one route through the program which fails to lead to termination, which would contradict the program’s specified behaviour, because it would subvert the program’s *guarantee* to generate an output. Consequently, any physical system which implements random integers must be hypercomputational. [More generally, any computation which is guaranteed to terminate, but which can do so in infinitely many different ways, is hypercomputational.]

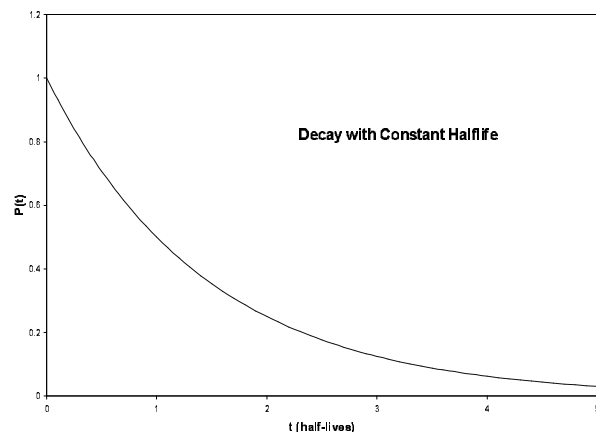
The semantics of particle decay

Although radioactive decay is a familiar concept in modern physics, we have found it to have remarkably subtle semantics. That particulate instability is important cannot be denied, as this popular description of the proton illustrates (but see [Loez96] for a more formal overview):

Protons are essential parts of ordinary matter and are stable over periods of billions and even trillions of years. Particle physicists are nevertheless interested in learning whether protons eventually decay, on a time scale of 10^{33} years or more. This interest derives from current attempts at grand unification theories that would combine all four fundamental interactions of matter in a single scheme. Many of these attempts entail the ultimate instability of the proton, so research groups at a number of accelerator facilities are conducting tests to detect such decays. No clear evidence has yet been found; possible indications thus far can be interpreted in other ways. ["Proton," *Microsoft® Encarta® 98 Encyclopedia*. ©1993-1997 Microsoft Corporation. All rights reserved.]

Given that protons are key components of every atomic nucleus, it is clear that some GUTs (grand unification theories, which are extensions of the Standard Model) would entail the instability of *all* atomic species. At the time of writing several hundred elementary particles are known, with mean lifespan ranging from the very long (protons $> 10^{33}$ years), through the everyday (free neutron ~ 15 minutes), to the vanishingly short ($Z^0 \sim 10^{-25}$ seconds). But what does “mean lifespan” actually mean in this context? The standard definition is straightforward. For example, if we start with one million free neutrons, we would expect that after roughly 15 minutes (the half-life of a free neutron) only about 500,000 would remain, the rest having decayed. Another 15 minutes later we would have only 250,000, and 15 minutes after that only 125,000.

The concept of “half-life” is so familiar that it is commonly accepted as an unquestioned basis for experimentation. For example, standard experimental support for the relativistic notion that time slows down for fast moving objects involves measuring the half-life of cosmic rays. These fast-moving particles rain constantly down upon the Earth, and it is possible to measure how their number changes with height. The lower we take our measurement, the longer the particles will have been in transit and undergoing decay, and the lower their apparent flux will become. Experimental evidence confirms that the flux



drops off much more slowly than would be expected given the amount of time involved (as measured by an observer on the Earth), because the fast-moving particles consider the journey to have taken considerably less time, whence considerably fewer of them decay.

Academic usage of the term suggests that physicists consider decay to be not merely possible but necessary, but this is nowhere expressed in the standard definitions which are firmly based in statistical analysis. We (and clearly physicists in general) would “expect” roughly half of our neutron sample to have decayed after 15 minutes, but we cannot guarantee it. The possibility always remains, however low the probability, that *none* of the neutrons will have decayed, even several days later.

This suggests two distinct semantics for the decay of unstable particles, one based on formal description, the other on pragmatic usage. Suppose that observation of an unstable particle p begins at time *zero*. Let us write $E(p,t)$ to denote the *a priori* probability that p still exists at time t , and has not yet decayed. The concept of half-life gives shape to the notion that $E(p,t)$ tends towards zero the longer we wait, *i.e.* $E(p,t) \rightarrow 0$ as $t \rightarrow \infty$. In the following statements, the value ε is assumed to be positive.

First Interpretation of $E(p,t) \rightarrow 0$ as $t \rightarrow \infty$

Any given particle p may or may not eventually decay, but the chances of its not having done so become arbitrarily close to zero as time passes. Formally, p satisfies

- $(\forall \varepsilon)(\exists T(\varepsilon))(\forall t) [(t > T(\varepsilon)) \Rightarrow (E(p,t) < \varepsilon)]$

Second Interpretation of $E(p,t) \rightarrow 0$ as $t \rightarrow \infty$

Any particle, p , must eventually decay after some finite time $T(p)$, but we do not know in advance what value $T(p)$ will take. We do know, however, that if we were to watch many particles the various values would be distributed as described. Formally, p satisfies

- $(\exists T(p))(\forall t) [(t > T(p)) \Rightarrow (E(p,t) = 0)]$

The symmetry between the two interpretations isn't immediately obvious, so we'll apply a little "trick" to help bring out the similarity. In the second interpretation, we're making the blunt assertion that $E(p,t)$ is zero, and consequently have no need of the test-parameter "ε" that appears in the first interpretation. We can re-introduce ε by observing that *zero* can be defined as "that unique non-negative number which is less than every positive number". Consequently,

$$(\exists T(p))(\forall t) [(t > T(p)) \Rightarrow (E(p,t) = 0)]$$

is the same statement as

$$(\forall \varepsilon)(\exists T(p))(\forall t) [(t > T(p)) \Rightarrow (E(p,t) < \varepsilon)]$$

provided $T(p)$ and E are independent of ε and ε is independent of p . Accordingly, we can write the two interpretations in strikingly similar terms, where ε is again understood to be positive:

- $(\forall \varepsilon)(\exists T(\varepsilon))(\forall t) [(t > T(\varepsilon)) \Rightarrow (E(p,t) < \varepsilon)]$
- $(\forall \varepsilon)(\exists T(p))(\forall t) [(t > T(p)) \Rightarrow (E(p,t) < \varepsilon)]$

Clearly, the only logical distinction between the two interpretations lies in the way T is chosen.

Experimental Indistinguishability of the Interpretations

I claim that there is no "experimental" way to distinguish between these two interpretations. The reason for this is intuitively obvious – the only type of particle which can distinguish them is one which never decays, and even then the experiments would have to be left running forever before the distinction could be verified.

The claim " $E(p,t) \rightarrow 0$ as $t \rightarrow \infty$ " can be viewed as a description of the experimental error we are prepared to accept when testing unstable particles against the assertion "every unstable particle will eventually decay". In order to test such an assertion we would build some apparatus, supply some particle p as input, wait for a time T and check to see whether or not p had decayed. The parameter

ϵ is our tolerance of experimental error – the experiment should produce observable evidence in favour of p 's decay all but ϵ of the time.

The two interpretative statements can be viewed as “experimental schemas”. They tell us that there are (at least) two distinct ways in which to relate the various constraints involved in this experiment. On the one hand we can envisage a situation in which (exact clones of) a single known particle will be repeatedly re-used as the input for our experiment, but the tolerances within which we have to work will be imposed by some outside agency and are not available *a priori*. We accordingly carry out research into the nature of our particle, and thereby come to understand its general nature. This yields enough information to allow a general rule of the form “if it hasn't decayed within T seconds, it's not going to”. Thereafter, no matter what tolerance we're told to work within, we apply the rule “run the experiment for T seconds and see what happens”. On the other hand, we can envisage a situation in we're told the tolerances we'll be expected to work to, but have no control over the particles that will be supplied for testing. Accordingly we conduct research into lots of different particles, and deduce a general rule of the form “if I run the experiment for T seconds, I can expect all but ϵ of the tests to be successful”.

We should also explain how E should be interpreted. The experimental schemas are in effect experimental systems which allow us to verify E by probing its structure, and they therefore assume at the outset that nothing is known about E beyond the standard assumption that it has the general shape of exponential decay with constant half-life. Each time a particle decays, however, the experiment generates information, and this information allows us to replace the assumed model of E with an observed version (a step-function equal to 1 before decay and 0 thereafter). This is why, when a particle decays, it is meaningful to say *of that particle* that $E(p,t)$ is zero thereafter. Until decay occurs, however, no extra information is forthcoming, and the representation of E remains abstract. Consequently, if a particle is believed to be unstable (but never in fact decays) we should continue to regard E as equal to the hypothetical exponential curve of standard decay. If a particle is known to be stable the situation is markedly different, because in this case we know that $E \equiv 1$ (no matter how long we wait, we remain certain that it hasn't yet decayed).

Is it possible to distinguish between these two experimental schemas experimentally? In other words, is there some apparatus we can construct which will generate particles p which can reliably distinguish the two schemas? Such a particle should have the property that when we supply it as an input, one of the experiments is certain to work (the associated logical expression evaluates to *true*), while the other is certain to fail (evaluates to *false*). We claim that no such apparatus can be constructed by finite means under the terms of the “scientific method” underpinning the Standard Model.

To see why, we need to examine the two statements in more details, and understand what sort of behaviour a “distinguishing particle” would need to exhibit. Recall the two interpretations, suitably instantiated in the context of this single-particle experiment.

First Interpretation - Our particle p may or may not eventually decay, but the chances of its not having done so become arbitrarily close to zero as time passes.

$$\bullet \quad (\forall \epsilon) (\exists T(\epsilon)) (\forall t) [(t > T(\epsilon)) \Rightarrow (E(p,t) < \epsilon)]$$

Second Interpretation - Our particle p must eventually decay after some finite time $T(p)$, but we do not know in advance what value $T(p)$ will take.

$$\bullet \quad (\exists T(p)) (\forall t) [(t > T(p)) \Rightarrow (E(p,t) = 0)]$$

Suppose p happens to decay after t_p seconds, where $t_p > 0$. [We know that our particle is properly defined when the experiment starts, so we can rule out the possibility $t_p = 0$.] Both statements evaluate to *true* if we take $T \equiv t_p$, since in both cases we know that $E(p,t)$ is experimentally measured to be zero for this particular particle p whenever $t > t_p$. Consequently, the only way a particle can distinguish between the two interpretations is if it *never* decays.

Pedants might argue that such a particle is, by definition, *stable*, and so cannot be used as a valid input to either experiment (according to this viewpoint we have already shown that there is no valid input which can distinguish the two interpretations). In fact, however, this distinction is academic. If we knowingly supply an *a priori* stable particle to the experiments, we need to reflect this prior knowledge by changing the definition of E . Because the particle never decays, we can be certain that $E \equiv 1$. Clearly, however, if E is always equal to *one*, it cannot be made arbitrarily small, and neither interpretation can ever be satisfied. Both interpretations evaluate to *false*, so they cannot be distinguished by *a priori* stable particles.

Suppose, however, we allow the use of *any* experimentally selected unstable particle, including those which are ostensibly unstable, but which never actually decay. Under this assumption the first interpretation evaluates to *true* (we have specifically allowed for unstable particles which fail to decay by our interpretation of E) while the second interpretation evaluates to *false* – there is no value of $T(p)$ for which the required implication holds.

It seems then, that unstable particles can distinguish the two schemas, *provided* they never actually decay. However, this requires waiting for ever, which violates “scientific method” on two distinct grounds. Firstly, an experimental result is only meaningful if it can be replicated, and clearly an experiment that runs for ever can never be repeated thereafter. Secondly, experiments must be conducted using “finite means”, and this includes the requirement that they run to completion in finite time. [If we restrict the experiments to run for no more than T seconds (say) they would clearly be unable to distinguish between particles that decay after T seconds from those which never decay.]

Experimentally, then, the two schemas should generate identical results. Any particle which, when supplied as an input to the first experiment, supports or undermines the assertion “all unstable particles must eventually decay” could also have been supplied to the second experiment and generated the same result.

Both interpretations entail experimental tautologies

We shall now demonstrate that *Interpretation One* yields a tautology (given the precepts of the Standard Model) and that *Interpretation Two* entails hypercomputation. Since they are experimentally indistinguishable, this will demonstrate that any experiment which purports to refute hypercomputation must also refute the Standard Model itself.

We observed above that both interpretations are true in the Standard Model for particles that decay and false for those which are *a priori* stable. For particles which are *a priori* unstable but never actually decay we found that the first interpretation was true, but the second false. Accordingly, if we restrict attention to *a priori* unstable particles, *Interpretation One* is a tautology relative to the Standard Model – it is true for every relevant particle p . That is, writing $Unstable(p)$ to indicate that p is *a priori* unstable,

$$StandardModel \models Unstable(p) \Rightarrow (\forall \epsilon)(\exists T(\epsilon))(\forall t) [(t > T(\epsilon)) \Rightarrow (E(p,t) < \epsilon)]$$

Given that both interpretations are uniformly false for *a priori stable* particles, we can extend our experimental indistinguishability claim to show – relative to the Standard Model – the following tautology also

$$\text{StandardModel} \models \text{Unstable}(p) \Rightarrow (\exists T(p)) (\forall t) [(t > T(p)) \Rightarrow (E(p,t) = 0)]$$

We claim, however, that this “experimental tautology” (essentially *Interpretation Two* restricted to *a priori* unstable particles) entails the feasibility of hypercomputation.

This is intuitively clear once we re-express the statement in everyday terms, for it simply expresses the statement “in a sample in which every particle is unstable, every particle must eventually decay”. The assertion of this statement allows us to re-construct the random number generator I described at the outset.

Interpretation Two entails hypercomputation

Our experimental strategy is straightforward. We start a clock running and wait until a radioactive species decays. The integer part of the number of minutes that have passed is the random number generated. The additional complexity of our proposed apparatus is there simply to overcome questions of observability and measurement.

We begin by replicating apparatus typically used to show that α -particles are in fact Helium nuclei. This comprises an evacuated flask placed inside a container which is filled with a suitable mass of α -radiating material. As the external material radiates, so particles are trapped in the inner flask, and this gradually gains mass as it fills with Helium (in the original experiment we would measure the spectrum of the material generated to determine its identity). The amount of material required is determined according to ones measuring equipment. Suppose we can reliably discern mass increments of size m or greater. If the radiating material is such that N grams of material will generate n grams of Helium nuclei in the evacuated flask, we choose N large enough to ensure that $n \gg 100m$.

We then automate the process by measuring the mass of the inner flask every minute (say), and stop when the perceived gain in mass is first observed to exceed $50m$ [any recursive procedure for resolving boundary disputes is acceptable]. If this occurs at the i 'th reading, the system is deemed to have generated the value i .

Given *Interpretation Two*, this system is a random number generator. Decay is axiomatically stochastic under the Standard Model, so our concern lies in verifying that the process will eventually terminate. However, this is a direct consequence of “guaranteed decay”, because we know that a time will eventually come when *all* of the finitely many atoms in the external container will have decayed, and by construction the corresponding mass of Helium is enough to ensure termination.

Summary and Further Questions

We have described two experimental schemas for affirming the statement “every unstable particle must eventually decay”, and have shown that at an experimental level these schemas are indistinguishable. Since one is an experimental tautology and the other entails hypercomputation, it follows that hypercomputation cannot be refuted without also refuting the Standard Model. Hypercomputation must, therefore, either be a consequence of the Standard Model, or failing that it is logically independent.

The question of “guaranteed decay” appears to be bound up with notions in the development of Grand Unified Theories, for which issues such as proton decay play a central role. It is possible that the feasibility of hypercomputation and the relevance of GUTs may be closely coupled. For example, suppose hypercomputation is found to be independent of the Standard Model, and that String Theory (which supports proton decay, see [Dani01]) is found to support “guaranteed decay”. Then String Theory would represent a hypercomputational model of nature in contrast to the recursion of the Standard Model. This raises the exciting possibility that GUTs may allow more freedom when discussing hypercomputational aspects of Nature.

It is perhaps worth observing that were a physicist to perform an experiment in which a highly radioactive species did *not* show signs of decay, they would not conclude “this is a fine example of Nature at work, showing us that decay really is probabilistic and need not occur over any finite period”, but rather “there is something wrong with my measuring equipment”. In other words, no matter how emphatically quantum physicists might insist that the equations of quantum theory are to be interpreted probabilistically, it is clear from their actions that this is *not* their pragmatic understanding of the concept. There is very clearly an unstated presumption that decay *must* eventually occur, something which is not implicit in the probabilistic interpretation of the standard model.

Let us conclude with a question. This will form the basis of our next discussion paper.

The random number generator described above can generate *any* positive integer. Consequently (our corollary of) König’s Lemma applies, and the system must be hypercomputational if it can be built. But is it possible to generate a hypercomputational system using a *finitely* random number generator, for example by tossing a coin? On the face of it, the answer should be “no”, because any physical experiment would have access to only a finite component of the bit-stream generated by the coin tossing, and any *finite* bit-stream could have been generated by a recursive program. Consequently, whatever the effect of including the coin-tossing, it cannot introduce more information into the computation (over the finite period during which the experiment takes place) than the equivalent recursive system would have done. However, a little reflection shows that this is also true of the *infinite* random number generator. During any finite period its output will be equivalent to that of some recursive simulation. How can we square this apparent recursiveness with its known hypercomputational nature?

References

- [BaiJ00] Sunny Bains and Jeffrey Johnson (2000) Noise, physics and non-Turing computation. *Joint Conference on Information Systems*, Atlantic City, 28 February – March 3, 2000.
- [CopS99] B. Jack Copeland and Richard Sylvan (1999) Beyond the Universal Turing Machine. *Australasian Journal of Philosophy* **77**, pp. 46-66.
- [Dani01] Ulf Danielsson (2001) Introduction to string theory. *Rep. Prog. Phys.* **64**, 51-96.
- [Davi01] E.B. Davies (2001, to appear) Building Infinite Machines. *Machines British Journal for Philosophy of Science*.
- [Loez96] Jorge L. Loez (1996) Supersymmetry: from the Fermi scale to the Planck scale. *Rep. Prog. Phys.* **59**, 819-65.
- [Miln89] Robin Milner (1989) *Communication and Concurrency*. Prentice-Hall International.

- [Sieg99] Hava T. Siegelmann (1999) Stochastic Analog Networks and Computational Complexity. *J. Complexity* **15**, 451-75.
- [Stan91] Mike Stannett (1991) *An Introduction to post-Newtonian and non-Turing Computation*. Technical Report CS-91-02, Dept of Computer Science, Sheffield University, United Kingdom. Available online at <http://noisefactory.co.uk/download/>.
- [ZhoW00] Ning Zhong and Klaus Weihrauch (2000, preprint) Computability Theory of Generalised Functions.

Mike Stannett
2 May 2001