

Universe creation on a computer

Gordon McCabe

August 17, 2004

Abstract

The purpose of this paper is to provide an account of the epistemology and metaphysics of universe creation on a computer. The paper begins with F.J.Tipler's argument that our experience is indistinguishable from the experience of someone embedded in a perfect computer simulation of our own universe, hence we cannot know whether or not we are part of such a computer program ourselves. Tipler's argument is treated as a special case of epistemological scepticism, in a similar vein to 'brain-in-a-vat' arguments. It is argued that the hypothesis that our universe is a program running on a digital computer in another universe generates empirical predictions, and is therefore a falsifiable hypothesis. The computer program hypothesis is also treated as a hypothesis about what exists beyond the physical world, and is compared with Kant's metaphysics of noumena. It is proposed that a theory about what exists beyond the physical world should be formulated with the precise concepts of mathematics, and should generate physical predictions. It is argued that if our universe is a program running on a *digital* computer, then our universe must have compact spatial topology, and the possibilities of observationally testing this prediction are considered. The possibility of testing the computer program hypothesis with the value of the density parameter Ω_0 is also analysed. The informational requirements for a computer to represent a universe exactly and completely are considered. Consequent doubt is thrown upon Tipler's claim that if a hierarchy of computer universes exists, we would not be able to know which 'level of implementation' our universe exists at. It is then argued that a digital computer simulation of a universe cannot exist as a universe. However, the paper concludes with the acknowledgement that an analog computer simulation can be objectively related to the thing it represents, hence an analog computer simulation of a universe could, in principle, exist as a universe.

Keywords: Computer Cosmology Creation Scepticism Information Entropy

1 The epistemology of universe creation on a computer

F.J.Tipler has suggested that our universe could be a computer program running on a computer in another universe, (see, for example, p240-244 of Tipler 1989, and p206-209 of Tipler 1995). Tipler imagines a perfect computer simulation of our universe, which precisely matches the evolution in time of our own universe, and precisely represents every property of every entity in our universe. Such a simulation would simulate all the people who exist in our own universe. Such simulated people, suggests Tipler, would reflect upon the fact that they think, would interact with their apparent environment, and would conclude that they exist. Their experience would be indistinguishable from our own experience, and Tipler infers from this that we ourselves cannot know that we are not part of such a computer program. *Ex hypothesi*, there is nothing in our experience which could be evidence that we are not part of such a program, hence, it might be argued, we cannot know that we are not part of a computer program.

This argument is a type of epistemological scepticism, similar to Descartes' dreaming argument. Descartes raised the possibility that one could experience a dream which is indistinguishable from the experience of a conscious, waking individual. The sceptical argument from this is that, *ex hypothesi*, there is nothing in one's experience which could be evidence that one is not dreaming, hence one cannot know that one is not dreaming.

A modern version of this is the 'brain in a vat' hypothesis. Jonathan Dancy characterises this sceptical hypothesis as follows: "You do not know that you are not a brain in a vat full of liquid in a laboratory, and wired to a computer which is feeding you your current experiences under the control of some ingenious technician/scientist...For if you were such a brain, then, provided that the scientist is successful, nothing in your experience could possibly reveal that you were; for your experience is *ex hypothesi* identical with that of something which is not a brain in a vat. Since you have only your own experience to appeal to, and that experience is the same in either situation, nothing can reveal to you which situation is the actual one," (Dancy 1985, p10).

One can identify two distinct premises in this argument:

- (a). It is possible for a brain in a vat to be fed experience of an illusional world.
- (b). It is possible for that experience to be indistinguishable from our own experience.

From these premises, the reasoning is as follows: Because the experience of the illusional world would be indistinguishable from one's own experience, it is not possible to know whether or not one's own experience is experience of a real world, or experience of an illusional world fed to a brain in a vat. Hence, it is not possible to know whether or not one is a brain in a vat.

There is, however, a vital ambiguity in the argument. There are two different senses in which real world experience could be indistinguishable from illusional world experience. One could claim either of the following two propositions:

1. The experience of the illusional world would be indistinguishable from the real world in terms of the detailed content of the experience.
2. The experience of the illusional world would be indistinguishable as experience from experience of the real world. In other words, the form of the illusory experience would be indistinguishable from the form of real-world experience.

It is not clear which of these claims Dancy is making. To illustrate the differences between these claims, consider the following scenarios:

Firstly, suppose that an individual is born in the real world, grows-up in the real world, and experiences the real world for 30 years, developing a range of cognitive skills, and accumulating a large collection of memories. Then, one night, whilst he lies asleep, the individual is unknowingly drugged and kidnapped by a scientist. As the victim lies unconscious in the scientist's laboratory, his brain is removed and wired up to a computer. When the individual is allowed to recover consciousness, he wakes up to experience an illusional world controlled by the computer. Suppose that the individual retains his memories of the real world. To prevent the individual from having a reason to believe that he is a brain in a vat, the experience of the illusional world must be indistinguishable from the individual's experience of the real world. Both the form and the detailed content of the individual's illusory experience must be indistinguishable from his experience of the real world. The illusional world must have the same spatial layout and the same apparent history as that part of the real world known to the victim, and the illusional world must evolve according to the same laws that operate in the real world. The victim must feel that he experiences his world, and influences events in his world, with the same body that he possessed before he fell asleep the previous night. The victim must not recognize any difference between the real world and the illusional world that is not explicable by the laws of the real world. The victim must appear to perceive the same world he perceived before he fell asleep the previous night.

If these conditions were satisfied, then the individual would have no justification for believing that he is a brain in a vat. In accordance with conventional definitions of knowledge, if the individual would not be justified in believing that he is a brain in a vat, then he could not know that he is a brain in a vat.

It is possible to imagine other sceptical scenarios which do not require the detailed nature of the illusional world to be indistinguishable from the detailed nature of the real world. If an individual's memories of the real world are deleted or suppressed, and apparent memories of an illusional world completely different from the real world are added, then experience of the illusional world would not give the individual reason to believe that he experiences an illusional world. The individual could experience an illusional world with a spatial layout and history totally different to the spatial layout and history of the real world. The illusional world could operate according to laws different to those that operate in the real world. Nevertheless, the experience of the illusional world would be indistinguishable, as experience, from experience of the real world. In other

words, the form of the illusory experience, if not the detailed content, would be indistinguishable from real-world experience.

To take another example, if an individual were fed illusory experiences from birth, that individual would have no memories of the real world. Hence, experience of an illusory world completely different from the real world in terms of detailed content, would not give the individual reason to believe that he experiences an illusional world.

It is not necessary to suppose that the individual who experiences an illusional world is an unwilling participant. It is possible, for example, that one's entire lifetime of experience upon 20th/21st century Earth, is part of a virtual reality game, played on a distant planet in the far-future. The technology of the far-future might enable game-players to play any role, in any factual or fictitious world. The game technology might suppress one's real-world memories, and supply the memories of the character one is playing. If the game technology suppressed one's real-world memories, one would be unaware of playing a virtual reality game. The game technology might even suppress one's real-world cognitive skills; one might experience birth, growth and mental development in the game world. Either way, one would have no memory of deciding to enter the game world. Once again, the sceptical argument is that one's own experience is indistinguishable from the experience of someone playing such a virtual reality game, hence one cannot know whether or not one is playing such a game.

Those sceptical arguments which require the detailed nature of the illusional world to be indistinguishable from the detailed nature of the real world, share a common point of vulnerability. It is possible for the hypothesis supporting such sceptical arguments to be false, and it is possible to know that it is false.

If the detailed content of the illusory experience is indistinguishable from the detailed content of real experience, then one can infer facts about the real world from one's experience, irrespective of whether one's experience is illusory or not. This allows one to determine, by scientific investigation, whether the hypothesis which supports the sceptical argument, is true or false.

For example, consider the brain in a vat argument. Recall that this sceptical argument is based upon the premise that it is possible for a brain in a vat to be fed experience of an illusional world. Because the illusional world would be indistinguishable, by hypothesis, from the real world, one's sensory systems and neurophysiology in the illusional world would be the same as one's sensory systems and neurophysiology in the real world. Hence, one could learn about one's real world physiology and neurology from one's experience, irrespective of whether one's experience is experience of the real world, or the illusional experience of a brain in a vat. One could not be led into forming false beliefs about the kind of entity one is without the violation of the indistinguishability condition.

Investigation of the human brain may reveal that it is impossible for it to be stimulated in a way which would produce experience indistinguishable from the experience of a person who is not a brain in a vat. Thus, the hypothesis upon which the sceptical argument is based, could be false. If one knew from neurophysiology that it is not possible for one to be a brain in a vat, then

one would know that one is not a brain in a vat. When Dancy characterises the sceptical argument he states that “you have only your own experience to appeal to,” (Dancy 1985, p10). This is false because one can also appeal to one’s scientific understanding, based upon both theory and empirical evidence.

The other sceptical scenarios share this vulnerability: neurophysiological investigation of the brain could reveal that it is not possible for dreams to be indistinguishable from the experiences of a waking individual; research in micro-electronics, computer science, and human physiology, might conclude that totally authentic virtual reality is not possible.

Those sceptical arguments which do not require the detailed nature of the illusional world to be indistinguishable from the detailed nature of the real world, are more robust. If the detailed nature of the illusional world is different from the detailed nature of the real world, then one cannot necessarily learn about real world physiology and neurology from illusory experience. However, these more robust sceptical scenarios are dependent upon the following premise:

- Either it is possible to delete or suppress an individual’s memories of the real world, and to replace them with apparent memories of an illusional world, or it is possible to feed an individual with illusory experience from birth.

If this premise is false, then all the sceptical arguments which concern illusional worlds might be refuted by empirical investigation. It is, however, difficult to establish whether this premise is true or false. If scientific investigation reveals that it is impossible in our world to feed an individual illusory experiences from birth, and that it is impossible in our world to delete or suppress an individual’s memories, and replace them with apparent memories of an illusional world, then this alone does not establish whether the premise is true or false. If our world is an illusional world, and if the detailed nature of the illusional world is different from the real world, then scientific discoveries about our world, the illusional world, do not tell us anything about the real world.

It has been assumed in this section that it is possible to make a distinction between the form and content of experience. If such a distinction is not possible, then the sceptical scenarios must be re-categorised as follows:

1. An individual in our world experiences an illusional world which is indistinguishable from experience of our world. The individual is unaware that his experience is illusional precisely because the illusional experience is indistinguishable from experience of our world.
2. An individual in our world experiences an illusional world which is distinguishable from experience of our world. The individual is unaware of the difference, either because his memories of our world have been deleted or suppressed, or because he has experienced the illusional world from birth.

In case 1 the sceptical argument is as before, with the reference to the content of experience omitted. In case 2, the sceptical argument is as follows: If an individual in our world could experience an illusional world which is distinguishable from experience of our world, and if that individual could be made unaware that what he experiences is illusional, then our own world experience could be illusional experience, distinguishable from the real world. We cannot know whether or not our experience is experience of the real world, or experience of an illusional world different from the real world.

The computer program hypothesis differs in one respect from the brain-in-a-vat type of hypothesis. The latter hypothesis suggests that an individual in a real world could be fed experiences of an illusional world, a world that does not objectively exist. The computer program hypothesis suggests that an entire universe could be created as a computer program, and that many individuals could be created as part of the program. This hypothesis does not merely suggest that there is a computer program which is feeding illusory experiences to individuals who exist in a real world. Instead, individuals capable of experience are themselves created by the program, and the world they experience is just as real relative to them, as our world is relative to us. It is not Tipler's claim that we cannot know whether or not our world is an illusional world. Instead, he claims that "we cannot know if the universe in which we find ourselves is actually ultimate reality," (Tipler 1995, p208). Tipler's claim is that we cannot know what level of reality we experience; that we cannot know whether or not the universe we experience has been created on a computer existing in another universe.

However, the hypothesis that our own universe is indistinguishable from a universe created on a computer, may be false. It will be demonstrated in this paper that physical predictions follow from the hypothesis that our universe is a program running on a *digital* computer. For example, it follows that the structure of the universe must be discrete, and that the spatial universe must be compact. If these predictions are found to be false, then it is impossible for our universe to be a program running on a digital computer. If the predictions are falsified, then our universe is distinguishable from a universe created on a digital computer. Alternatively, if these predictions are found to be true, then it remains possible for our universe to be a program running on a digital computer. Empirical investigation is necessary to determine if the computer program hypothesis is possible.

2 The metaphysics of universe creation on a computer

The hypothesis that our universe is a program running on a computer in another universe is not merely a sceptical epistemological hypothesis, but a metaphysical

hypothesis, in the sense defined below.

The term ‘metaphysics’ seems to have at least two different meanings. On the one hand, it is the study of that which possibly exists beyond the physical world. On the other hand, it is a whole group of philosophical subjects, such as the studies of time, causation, substance, and universals. These subjects seem to be united by the fact that they involve very general, foundational study of the nature of things.¹

For the purpose of this paper, metaphysics is defined to be the study of that which possibly exists beyond the empirically detectable world. In contrast, physics is defined to be the study of the empirically detectable world. The hypothesis that our universe is a program running on a computer in another universe, is clearly a metaphysical hypothesis, in the very specific sense defined here. The hypothesis is that the computer hardware on which the program is running cannot be empirically detected by the beings represented in the software, hence the hypothesis is metaphysical rather than physical.

It is important to distinguish Tipler’s hypothesis from a metaphysically distinct proposal made by J.D.Barrow. Barrow suggests that “If we were to regard the Universe as a vast computer...then we can readily envisage the laws of Nature as some form of software which runs upon the particular forms of matter that form the world of strings and elementary particles,” (Barrow 1991, p160). In Tipler’s computer program hypothesis, the computer hardware is inaccessible to the people represented in the computer program; the constituents of matter, elementary particles or not, are just as much a part of the program software as the laws of physics. Presumably, each different type of particle or field would correspond to a different data type in the program. Each individual particle or field would then correspond to an instance of the relevant data type. In programming parlance, an instance of a data type is called a data object. Hence, the constituents of matter would correspond to data objects defined in the program. The laws of physics would correspond to the algorithms which act upon the data objects defined in the program. In general, entities would correspond to data objects in a computer program, and processes would correspond to algorithms. For example, an individual electron would correspond to a data object, and the Dirac equation would correspond to an algorithm capable of acting upon any electron data object. To give another example, in the geometrodynamical formulation of general relativity, a 3-manifold Σ , and the tensor fields (γ_i, K_i, ϕ_i) representing the intrinsic geometry γ_i , extrinsic geometry K_i , and matter fields ϕ_i at time i , would all correspond to data objects. The geometrodynamical evolution process would correspond to an algorithm which calculates $(\gamma_{j+1}, K_{j+1}, \phi_{j+1})$ from (γ_j, K_j, ϕ_j) .

¹The historical reasons for the double-meaning can be traced to Aristotle, as Barry Smith explains: “The books of Aristotle’s *Physics* deal with material entities. His *Metaphysics* (literally ‘what comes after the *Physics*’), on the other hand, deals with what is beyond or behind the physical world - with immaterial entities - and thus contains theology as its most prominent part. At the same time, however, Aristotle conceives this ‘metaphysics’ as having as its subject matter all beings, or rather being as such. *Metaphysics* is accordingly identified also as ‘first philosophy’, since it deals with the most basic principles upon which all other sciences rest,” (Smith 1995, p373).

After suggesting that our universe could be a computer program running on a computer in another universe, Tipler goes one step further, and claims that there is no need for a computer to be running the program. The state of memory of a digital computer can be treated as a long string of binary digits, and this represents a natural number in binary notation. Given that a computer program maps an initial memory state to a final memory state, a computer program can be treated as a mapping on the set of natural numbers. Tipler duly treats a program as an abstract mapping $\mathbb{N} \rightarrow \mathbb{N}$, and claims that “if time were to exist globally, and if the most basic things in the physical universe and the time steps between one instant and the next were discrete,” (Tipler 1995, p208), then our universe could be in one-to-one correspondence with such an abstract object. Tipler acknowledges that the most basic things in the physical universe could be continuous, hence he proposes a further generalization of what a simulation is: “Let us say that a perfect simulation exists if the physical universe can be put into one-to-one correspondence with some mutually consistent subcollection of all mathematical concepts,” (ibid., p209).

This proposal does not merely suppose that mathematical Platonism is true, that mathematical objects exist independently of the physical universe, in an abstract realm. Nor does it merely suppose that physical objects possess intrinsic mathematical properties. Instead, it supposes that physical objects can be identified with mathematical objects. As Barrow puts it, “We exist in the Platonic realm,” (Barrow 1992, p282). Whilst this is a fascinating idea, I shall restrict the discussion in this paper to the hypothesis that our universe is a program running on a computer in another universe.

The notion that there is something which exists beyond the empirically detectable world has famous precedents in the history of philosophy. Various types of thing have been postulated to exist beyond the physical world: mental entities, theological entities, and mathematical entities. These types of metaphysical suggestion are of no relevance to this paper. Rather, the focus of attention is the metaphysical hypothesis that there is something non-mental, non-deistic, and non-mathematical, which exists beyond our physical world. For example, Kant proposed that there are things-in-themselves, so-called ‘noumena’, which exist beyond the empirically accessible world. The metaphysics of the computer program hypothesis can be compared with the metaphysics of Kant’s noumena.

To recall, Kant suggested that there is a distinction between noumena and phenomena. The noumena are things in themselves, and the phenomena are the appearances of things in sensory perception. There are three possible ways of defining noumena. The noumena could be things which exist independently of sensory perception, or they could be things which exist independently of empirical detectability, or they could be things which exist independently of cognition altogether. Obviously, things which exist independently of empirical detectability also exist independently of sensory perception, and things which exist independently of cognition also exist independently of empirical detectability.

If one merely stipulates that noumena are things which exist independently of sensory perception, then noumena could simply be things which are too small to see, like atoms and electrons. Things which are too small to see are still empir-

ically detectable. As a classic example, an electron leaves a luminescent trail in a Wilson Cloud Chamber. The electron is not directly perceivable, but it is nevertheless detectable. Kant seems to suggest that noumena exist independently of both sense perception and empirical detectability of any kind. Further, Kant seems to hold that noumena are beyond cognition altogether. The computer program hypothesis holds that the states and processes of the computer in another universe, exist beyond both sense perception and empirical detectability, but these states and processes are not beyond cognition. What exists beyond the physical world is conceivable, according to the computer program hypothesis. In contrast, Kant seems to hold that we cannot even conceive what things in themselves are like.

The computer program hypothesis is consistent with a threefold distinction between the phenomenal, the physical, and the metaphysical. This corresponds to the distinction between appearance, physical reality, and metaphysical reality. Appearances and phenomena consist of sensory experiences such as colours, sounds, and smells. Physical reality is the world described by physics, the world of atoms, electrons, and space-time. The hypothetical metaphysical reality consists of the states and processes of a computer in another universe. In this threefold distinction, space and time exist independently of sensory appearances, whereas Kant believed that space and time are merely the format into which sensory experience is arranged. Unlike Kant, the proposal in this paper will not relegate space-time to the merely phenomenal.

The computer program hypothesis is an interesting case because the global metaphysics is drawn from local physics. The nature of what lies beyond the entire physical universe (global metaphysics) is drawn from the nature of the computer, a part of the physical world (local physics).

This paper proposes that, in general, the relationship between metaphysics and physics should be similar to the relationship between topology and geometry. A particular topology constrains the possible geometry, but is, nevertheless, consistent with a range of different geometries. For example, given the topology of the n -sphere S^n , it is impossible for the geometry to be flat, but there are, nevertheless, many possible geometries on S^n with non-zero curvature. Topology has implications for geometry, but a particular topology does not entail a unique geometry. Similarly, a metaphysical theory should, at the very least, have implications for the physical world. A metaphysical model should constrain the possible physical models, but a metaphysical model should not entail a unique physical model. In general, a metaphysical model should be consistent with a range of different physical models.

Conversely, it is not possible to infer a unique topology from geometrical properties like sectional curvature. Similarly, it is not possible to infer a unique metaphysical model from a physical model. However, geometrical properties do have topological implications, and similarly, the theories of physics do have metaphysical implications.

The computer program hypothesis developed in this paper exemplifies these proposed standards for a metaphysical theory. Properly developed, the hypothesis makes predictions about the physical universe. The metaphysical hypothesis

that our universe is a program running on a digital computer entails that

- The universe is discrete
- The solutions to the fundamental evolution equations of physics must be computable functions
- The spatial universe has compact topology

These predictions are empirically testable, hence the metaphysical computer program hypothesis is empirically testable. It will be demonstrated that the computer program hypothesis might be verifiable or falsifiable by astronomical observation. None of the predictions above will be invalidated by the development of quantum computers. Although quantum computers might be able to perform certain calculations faster than computers based upon the notion of a Turing machine, the collection of uncomputable functions for a quantum computer is the same as the collection of uncomputable functions for a Turing machine. Like existing computers, quantum computers will possess a finite memory. And like existing digital computers, a quantum computer will only be able to represent discrete things.

Not only should a metaphysical hypothesis or theory have physical implications, but it should be formulated with the precise concepts provided by mathematics, just like a physical hypothesis or theory. There is no reason why natural language should be adequate in metaphysics when it is inadequate in physics. The meanings of the terms in mathematics are precisely established with stipulative definitions. In contrast, the meanings of the terms in natural language are often determined by use, and the use is inconsistent amongst the community of language users. If use determines meaning, and use is inconsistent, then it follows that meaning is ambiguous. Use varies from one member of the community to another, and use changes over time, hence if use determines meaning, then meanings vary from one member of the community to another, and meanings change over time.² This makes natural language inadequate if precision is sought.

If one defines metaphysics as the study of that which possibly exists beyond the empirically detectable world, there is nothing in that definition which entails that physics cannot have any metaphysical implications, or vice versa, and there is nothing in that definition which entails that mathematics cannot be used to formulate metaphysical theories. It would be strange indeed if a discipline were defined not by its subject matter, but by the methodological tools contingently used at one time to study that subject matter. If metaphysics were defined to be the natural language study of that which possibly exists beyond the empirically detectable world, it would be rather like an architect who refuses to use

²The meaning of a term in natural language should be given by a statistical distribution which specifies the relative frequency of the different uses of the term. The ambiguity of a term could be quantified with the variance of the statistical distribution. A highly peaked distribution, with a small variance, would indicate a term with low ambiguity. A flattened distribution, with a large variance, would indicate a term with high ambiguity.

Computer Aided Design tools because he has defined architectural design to be a matter of pen and paper alone. If the pursuit of knowledge is one's main objective, then one uses every conceptual resource at one's disposal.

Kant's metaphysics provides a good example of what is wrong with traditional metaphysics. His principal metaphysical assertion is that there are things-in-themselves, so-called 'noumena', which exist beyond the empirically accessible world. Although Kant seems to resist claiming that noumena cause phenomena, he does believe that phenomena are dependent in some way upon noumena. However, he also claims that the noumena are unknowable. If there are things-in-themselves, the noumena, and if they do have some relationship to phenomena, then any well-developed, detailed speculation about the nature of the noumena must have testable predictions about phenomena. Only if the noumena have no relationship to phenomena could they be unknowable. Kant, and his advocates in subsequent centuries, should have made some attempt to define noumena, using mathematical concepts, and should have derived testable predictions from their theoretical models. The only logical reason not to would be the dogmatic assertion that the noumena are beyond cognition; not just unknowable, but inconceivable as well.

Because Kant believed that a noumenal world beyond the empirical world must be unknowable, he attempted to re-define metaphysics. He believed that the faculties of cognition, the forms of experience, and the 'categories' of understanding exist independently of the content of our experience. Kant advocated the study of these things as metaphysics in the sense that they are independent of the content of experience. He urged that metaphysics should move away from speculation about the unknowable to study of the knowable. The false assumption in this reasoning is that a world beyond the empirical world must be unknowable.

Let us illustrate the claim that a metaphysical hypothesis should be formulated with the precise concepts of mathematics. One could hypothesize that both the transcendent, metaphysical world and the physical world are represented by mathematical structures, and they are related by a mathematical mapping. Let \mathcal{M} denote the mathematical structure which represents the metaphysical world, and let \mathcal{N} denote the mathematical structure which represents the physical world. One can postulate that both \mathcal{M} and \mathcal{N} must have

1. A cardinality
2. A topology
3. A dimension
4. A geometry

and one can postulate that they are related by a projection mapping $\mathcal{M} \rightarrow \mathcal{N}$ which is

1. Non-injective, and therefore, non-invertible

2. Continuous
3. Either dimension-lowering or dimension-preserving
4. Non-isometric

Note that condition 2 relates to the topological notion of continuity, and should not be confused with the cardinal notion of the continuum. A cardinality is not specified for either \mathcal{M} or \mathcal{N} . Even if both were discrete mathematical structures, they could each possess a topology, and the projection mapping could be continuous with respect to these topologies. Condition 3 means that the physical world, \mathcal{N} , is either of lower dimension than the metaphysical world, \mathcal{M} , or they are of equal dimension.

This particular metaphysical hypothesis is formally analogous to suggestions in some branches of physics that the universe has many more dimensions than we currently detect. For example, the original Kaluza-Klein theory predicted that the universe actually has 5 dimensions, and more recent versions predict many more dimensions. These theories have the empirically detectable 4-dimensional space-time diffeomorphically embedded in an n -dimensional manifold. The fields and geometry on the 4-dimensional submanifold are obtained from the n -dimensional manifold by what is either called a parallel or orthogonal projection. It must be conceded that these theories do not consider the n -dimensional reality to be metaphysical. In some cases, at least, this is because the n -dimensional reality is merely inaccessible at low-energies. Being detectable at higher energies, it is not a metaphysical reality. The fact that we cannot directly perceive the higher dimensions simply places them in the same category as the electron: they are empirically detectable, in principle, but not directly perceivable. The proposed relationship between the metaphysical world and the physical world is merely analogous to the relationship between the n -dimensional physical world and the perceptible world in Kaluza-Klein theories.

The relationship between the 3-dimensional physical world and phenomenal experience is also given by a projective transformation. The 3-dimensional physical world is projected onto the human retina, a 2-dimensional surface, by means of a perspective projection transformation. A perspective projection is simply a special case of the more general type of projective transformation which I defined above. The relationships between the 4-dimensional objects of general relativity and the measurements of an observer are also given by projection mappings.

To reiterate, the postulated projection mapping from the metaphysical world to the physical world is given to illustrate how a metaphysical hypothesis can be formulated with the precise concepts of mathematics. It is not intended to be a serious proposal. One can imagine other metaphysical hypotheses which do not satisfy this projective relationship. In fact, the computer program hypothesis might not satisfy this relationship. Notably, in the case of Platonic metaphysics, there appears to be a projective mapping in the opposite direction. If many objects in the physical world are instances of one Platonic form

in the metaphysical ‘Platonic realm’, then there appears to be a projection mapping from the physical world to the Platonic realm.

Mathematics provides a vast resource of precise concepts which are just as applicable to metaphysics as they are to physics.³ Resistance to the application of mathematics in metaphysics could be founded on two possible objections:

- Mathematics is a non-foundational construction; the fundamental concepts one is trying to elucidate in metaphysics are taken for granted in mathematics; the fundamental concepts of mathematics are expressed in natural language. Hence, it might be argued, natural language is the appropriate medium for analysing foundational concepts.
- Mathematics is merely formal and abstract.

The first objection underestimates the capacity of mathematics to engage in foundational work. Mathematics is able to take the foundational concepts found in natural language, and generalise. For example, the Boolean logic expressed in natural language corresponds to the special case of a distributive lattice in a more general range of mathematical structures. The second objection fails to appreciate that the models of mathematical structures, i.e. the things which instantiate them, need not be constructed from numbers. Hence, the models of mathematical structures need not be abstract themselves. Physical objects can be the instances of mathematical structures, and so can metaphysical objects.

3 Deriving empirical predictions from the metaphysical hypothesis

J.D.Barrow has claimed that if our universe is a computer program, then all the laws of physics must involve computable functions, (Barrow 1991, p205). A computable function is defined to be a function whose value can always be calculated by performing a finite sequence of well-defined steps, often called an ‘effective procedure’. Certainly, if a universe unfolds in time on a computer, evolution equations must be used to calculate each time-step from the preceding time-step, and a solution of those evolution equations implemented on a

³This paper urges the use of mathematics in metaphysics, defined as the study of that which possibly exists beyond the physical world. However, mathematics can also be applied to metaphysics, defined as the general, foundational study of the nature of things. For example, one could propose that to understand the mind-brain relationship, it will be necessary to represent the brain as a mathematical category and the mind as a different mathematical category. One could propose that the relationship between the mind and the brain will be given by a functor between these two distinct categories. One could propose that every brain state corresponds to an object in the brain-category, every brain process corresponds to a one-parameter family of morphisms in the brain-category, every mental state corresponds to an object in the mind-category, and every mental process corresponds to a one-parameter family of morphisms in the mind-category. A mental state is not isomorphic to a brain state, but instead there is a functor which maps each mental state to a brain state, and which maps each mental process to a brain process. The functorial relationship expresses precisely the sense in which the mind is ‘reducible’ to the brain.

computer must be a computable function. If the solutions of the fundamental evolution equations of physics were found to be non-computable functions, then the computer program hypothesis would be falsified. Whilst the computer program hypothesis therefore predicts that the solutions to the fundamental evolution equations of physics must be computable functions, computability would not be necessary to represent, at once, an entire space-time on a computer. Computability is only a requirement if the representation attempts to calculate one aspect of the universe from another aspect. As Tegmark remarks, “since we can choose to picture our Universe not as a 3D world where things happen, but as a 4D world that merely is, there is no need for the computer to compute anything at all - it could simply store all the 4D data,” (Tegmark 1998, p26).

Note also that algorithmic compressibility is not a necessary condition to represent a universe on a computer. A digital representation of something is defined to be algorithmically compressible if the length, in bits, of the shortest program capable of generating that digital representation, is shorter than the length, in bits, of the digital representation itself. Our universe might not be algorithmically compressible, but might still be digitally representable on a computer. What follows is an attempt to derive more specific empirical predictions from the computer program hypothesis.

To represent the entire universe on a computer one must use either:

- A unified theory of everything.

or

- A set of different theories, each with its own limited domain of applicability, such that the set of domains covers the entire universe.

We do not, at present, have a unified theory of everything, but we do have a set of different theories, which grow progressively closer to covering the entire universe, in all its detail. Of these, the only empirically verified theory which is capable of describing the universe as a whole is general relativity. However, although general relativity can represent the universe as a whole, when it does so, it is only concerned with the large scale structure of the universe. It cannot represent detail on all length scales, as a unified theory of everything could be expected to do. Nevertheless, because general relativity has been empirically verified, the predictions of a unified theory of everything would have to converge to the predictions of general relativity within general relativity’s domain of applicability.

The physical predictions derived from the metaphysical computer program hypothesis will be derived from an examination of how to represent a general relativistic universe on a computer. This is perhaps a weak point of the strategy. The universe may not be a 4-dimensional Lorentzian manifold, as it is represented to be in general relativity. We do not know what type of thing a unified theory of everything, incorporating a theory of quantum gravity, would represent the universe to be. It is, therefore, a provisional decision to consider a universe created on a computer to be a general relativistic universe.

In addition, the predictions derived assume that a digital computer is the only type of computer which has the potential to simulate an entire universe. Although it isn't proven to be impossible for an analog computer to simulate an entire universe, the current evidence suggests that an analog computer cannot have the representational capacity to do so. An analog computer uses concrete (and continuous) physical quantities of one sort, (e.g. electrical quantities or hydraulic quantities), to represent concrete (and continuous) physical quantities of another sort, (e.g. the varying height of tides). In other words, an analog computer uses the concrete physical quantities of its physical components to represent the physical quantities of the system to be simulated. Early analog computers were constructed from levers, cogs, cams, discs and gears, and used mechanical motions to perform calculations. Modern analog computers tend to use electrical quantities, such as voltage levels, to represent the quantities of a simulated system, and specially designed circuits are used to perform arithmetic upon these voltage levels. Whilst an analog computer might use voltage levels to represent the values of quantities on a simulated system, a digital computer uses voltage levels to represent bits, and then sequences of bits encode the values of quantities on a simulated system. Analog computers tend to rely upon a mathematical resemblance between the pattern of quantity-values possessed by the machine and the pattern of quantity-values possessed by the simulated system.⁴ Analog computers do not use the versatile, encoded, abstract representation of physical quantities that a digital computer uses, and this limits their representational capacity.⁵

It is often claimed that the variables of an analog computer are, in fact, continuously variable, but this claim can be disputed. Variables such as electrical voltage or fluid pressure are probably discrete when they are reduced to the quantum level. Even if there are other variables which are genuinely continuous, it would still not be possible to precisely control their value. Suppose for the sake of argument that voltage is continuously variable. It would be impossible to set a precise input voltage of, say, 5.34V. The best one could ever do is to set an input voltage within some interval, say $5.34V \pm 0.01$. This point is better illustrated for the case of irrational numbers. It is impossible to set an input voltage of π , and this is not because of the limitations of current technology, but because infinite precision is not attainable.

In general relativity, space is represented as a 3-dimensional differential manifold, and space-time is represented as a 4-dimensional differential manifold. A digital computer, as it is currently understood, can only deal with discrete items of data. The most crucial fact to recognize about a computer program is that the data objects defined in it are built from \mathbb{Z} , the set of integers. In contrast, the objects of analytic mathematics are built from \mathbb{R} , the set of real numbers. The memory of a classical, (i.e. non-quantum), digital computer consists of elec-

⁴A good example of a mechanical analog computer is an orrery, a clockwork device for simulating the solar system. The actual position and motion of the balls representing the planets, represents the actual position and motion of the planets.

⁵As an exception, Hava Siegelmann (1999) has proposed neural net analog computers which are abstract encoders, like a digital computer.

tronic circuits which have two possible voltage states. These voltage states are represented by binary digits, otherwise known as ‘bits’. An element of memory is therefore called a ‘bit’. Each bit of memory has two possible states, represented as 0 and 1. The set of possible states of a bit can be represented as $\mathbb{Z}_2 = \mathbb{Z}/2\mathbb{Z} = \{0, 1\}$, the additive group of integers, modulo 2. \mathbb{Z}_2 is a realisation of the cyclic group of two elements. Each byte of memory, a string of 8 bits, and the smallest addressable unit of memory, can be represented by

$$(\mathbb{Z}_2)^8 = \mathbb{Z}_2 \times \mathbb{Z}_2 \times \mathbb{Z}_2 \times \mathbb{Z}_2 \times \mathbb{Z}_2 \times \mathbb{Z}_2 \times \mathbb{Z}_2 \times \mathbb{Z}_2,$$

the 8-fold Cartesian product of \mathbb{Z}_2 . Thus, for a classical computer with n bytes of memory, the entire memory can be represented by $(\mathbb{Z}_2)^{8n}$, a discrete mathematical structure of $8n$ dimensions. All the data objects defined in a program correspond to regions of memory, hence the data objects defined in a program are built from subsets of the discrete mathematical structure $(\mathbb{Z}_2)^{8n}$.

The memory of a quantum computer consists of physical systems which possess a quantum state space isomorphic to the 2-dimensional complex Hilbert space \mathbb{C}^2 . Each such memory element is referred to as a ‘qubit’, or ‘Qbit’. A string of n qubits is represented by the n -fold tensor product of \mathbb{C}^2 . Hence, the state of 8 qubits is represented by a vector in

$$(\mathbb{C}^2)^8 = \mathbb{C}^2 \otimes \mathbb{C}^2 \otimes \mathbb{C}^2 \otimes \mathbb{C}^2 \otimes \mathbb{C}^2 \otimes \mathbb{C}^2 \otimes \mathbb{C}^2 \otimes \mathbb{C}^2$$

As a consequence, the state of the n qubits can be quantum mechanically entangled.

Each qubit is considered to have a fixed basis, $\{v_0, v_1\}$. Each vector in the n -fold tensor product consists of a complex linear combination of the 2^n basis vectors $\{v_{i_1} \otimes \cdots \otimes v_{i_n} : i_1 = 0, 1, \dots, i_n = 0, 1\}$. The algorithms of a quantum computer correspond to unitary operators upon these complex Hilbert spaces. Because \mathbb{C}^2 is built from the set of real numbers, and because each qubit \mathbb{C}^2 possesses a continuum of quantum states, it might appear that a quantum computer can store an infinite amount of information. This appearance, however, is deceptive. Whilst there are a continuum of possible unitary operators on a qubit Hilbert space, each quantum computer will only be engineered to implement a finite collection. Moreover, each quantum computation must cease with a measurement of the state of the n qubits, and this collapses the state from a linear combination of the basis vectors into one of the fixed basis vectors, $v_{i_1} \otimes \cdots \otimes v_{i_n}$. The initial state on which the unitary transformations can operate is also such a state, and as Mermin comments, “the entire role of the state of the Qbits at any stage of a succession of unitary transformations is to encapsulate the probability of the outcomes, should the final measurement be made at that stage of the process,” (2002, p16). Thus, a quantum computer, like a classical computer, possesses a finite number of *accessible* states. In fact, n qubits of memory possess exactly the same number of accessible states as n bits of memory, namely 2^n . The data objects defined in a program running on a quantum computer are discrete objects.

Because every manifold has the cardinality of the continuum, and because digital computers can only represent discrete objects, it is impossible to exactly represent a manifold on a digital computer. It is, therefore, impossible on a digital computer to exactly represent space and space-time as they are represented in general relativity.

If space and space-time actually are manifolds, and if a manifold cannot be exactly represented on a digital computer, then space and space-time cannot be exactly represented on a digital computer. If the space and space-time of our universe cannot be exactly represented on a digital computer, then our universe cannot be a computer program running on a digital computer in another universe.

However, as already mentioned, the space and space-time of our universe may not actually be manifolds. Space and space-time may not exactly be as they are represented to be in general relativity. Perhaps space and space-time are discrete, and perhaps the manifolds of general relativity only provide an idealisation of a discrete reality. The space and space-time of our universe can only be exactly represented on a computer if space and space-time are discrete.

Loop quantum gravity offers, perhaps, a mathematically rigorous means to quantize general relativity, and loop quantum gravity suggests that space is discrete in some sense. Using Ashtekar's 'new variables' approach, canonical general relativity can be cast in the form of a canonical gauge theory, albeit a gauge theory with additional constraints to the Gauss constraint. The configuration space is a space of $SU(2)$ -connections on a principal fibre bundle over a 3-manifold Σ . In loop quantum gravity, each closed curve ('loop') in the 3-manifold defines a functional on the space of $SU(2)$ -connections. This functional is obtained by taking the holonomy of each connection around the loop, representing that group element as an operator on a vector space, and then taking the trace of that operator. Furthermore, each 'spin network' embedded in the 3-manifold defines a functional on the space of $SU(2)$ -connections. A spin network, treated in isolation, is a discrete mathematical object consisting of a graph, (a collection of vertices and edges), an irreducible representation of $SU(2)$ assigned to each edge, and an 'intertwining' operator between such representations assigned to each vertex. Such a graph embedded in the 3-manifold Σ defines a functional on the space of $SU(2)$ -connections by taking the holonomy of a connection along each edge, using the representations to obtain operators along each edge, forming the tensor product of all those operators, tensoring that with all the intertwining operators, and then contracting to obtain a number, the number assigned to the connection, (Baez 1995, p19). Such functionals turn out to be eigenvectors of operators which purportedly represent the area of surfaces in the 3-manifold and the volume of regions in the 3-manifold. Furthermore, these operators have discrete spectra.

If one accepts that quantum theory provides a complete description of a physical system, then, arguably, it is not the configurations of the classical system which exist, but the quantum state function(al). Hence, in the case of loop quantum gravity, the 3-manifold used to define the classical configuration space does not exist. Rather, it is the state functional defined by the spin

network which exists.

Many important questions remain. For example, the dynamics of loop quantum gravity remain intransigent, and there is no obvious classical limit to the theory. Whilst it is claimed that area and volume are discrete, what are they the area and volume of, if a 3-manifold does not exist? Are area and volume re-interpreted as properties of spin networks?

The established means of finding a discrete approximation to a manifold, is to find a cell complex which is homeomorphic to the manifold. In particular, one tries to find a simplicial complex which is homeomorphic to the manifold. The schema of the simplicial complex is a discrete mathematical object, which can be exactly represented on a computer. By representing the schema on a computer, one approximately represents the manifold.

If the schema of a simplicial complex is the natural discrete approximation to a manifold, then, conversely, the manifold can be said to be the natural continuum idealisation of the schema. If space and space-time are actually discrete, but if they can also be represented in a continuum idealisation as a 3-manifold and 4-manifold, respectively, then it is natural to suggest that space is actually a 3-dimensional schema, and space-time is actually a 4-dimensional schema. Regge calculus is generally considered to be the ‘discretized’ version of general relativity, and Regge calculus duly represents space and space-time as a simplicial complex.

Loop quantum gravity demonstrates that, although space and space-time might not be manifolds, they might not be the schema of simplicial complexes either. However, if space and space-time actually are discrete, it may be that they are best represented by loop quantum gravity on small scales, and best represented by the schema of simplicial complexes on large scales.

Some explanation of the mathematics is in order here. An n -cell is an object which is homeomorphic with the n -ball in n -dimensional Euclidean space, $\mathbb{D}^n = \{x \in \mathbb{R}^n : \|x\| = 1\}$. For example, a 2-ball is a disc, bounded by a circle, while a 3-ball is a solid ball bounded by a 2-sphere. Any polygon is homeomorphic with a 2-ball, and is therefore a 2-cell. Any solid polyhedron is homeomorphic with a 3-ball, and is therefore a 3-cell.

A cell-complex is obtained by pasting together any number of cells, so that the faces of the cells are either disjoint, or so that they coincide completely. A 3-dimensional cell-complex is obtained by pasting together 3-cells in such a way that the faces, edges and vertices of the cells are either disjoint, or they coincide completely.

The most interesting type of cell is a simplex. A 0-simplex is a point, or ‘vertex’, a 1-simplex is a line segment, or ‘edge’, a 2-simplex is a triangle, and a 3-simplex is a solid tetrahedron. By pasting together simplices, one obtains a simplicial complex, (see Stillwell 1992, p23-24). A 3-dimensional simplicial complex is obtained by pasting together solid tetrahedra. The schema of a 3-dimensional simplicial complex can be specified as follows. First, one declares all the vertices in the complex. Next, one can specify which subsets of the set of vertices correspond to simplexes. By specifying a pair of vertices, $\{P_i, P_j\}$, one indicates that those vertices are connected by an edge. One can then specify

which triples $\{P_i, P_j, P_k\}$ of vertices correspond to the faces, and finally one can list which quadruples $\{P_i, P_j, P_k, P_l\}$ of vertices correspond to the tetrahedra. One could alternatively give each edge a name, and then specify which triples of adjoining edges are connected by a face. One would then name each face, and specify which quadruples of adjoining faces are connected by a tetrahedron, (see Geroch and Hartle 1986, p546).

Although the manifold models of general relativity may be idealisations, one particular manifold model may eventually be verified by observation, to the exclusion of all others. To be specific, either a Friedmann-Roberston-Walker (FRW) model, a small perturbation of a FRW model, or an exact solution close to a FRW model, may be verified by astronomical observation. If the computer program hypothesis predicts that space or space-time is actually the schema of a simplicial complex on large scales, then the manifold model of the large-scale universe must be homeomorphic with a simplicial complex whose schema can be represented on a computer. It is therefore important to determine which manifold models of general relativity can be discretely represented on a digital computer by the schema of a simplicial complex. If a particular manifold model were to be verified by astronomical observation, but that model could not be represented by a schema on a digital computer, then the hypothesis that our universe is a computer program running on a digital computer would be falsified.

Suppose, then, that one tries to represent space-time on a computer with the schema of a 4-dimensional simplicial complex. Unfortunately, it is not known if every 4-manifold is homeomorphic to a simplicial complex. Hence, there may be 4-manifolds which cannot be discretely represented by the schema of a simplicial complex. If the space-time of the universe has a manifold idealisation which does not have a homeomorphic simplicial complex, then the space-time of the universe would not be representable on a computer by the schema of a simplicial complex. If there were no other means of discretely representing such a 4-manifold on a computer, then the space-time of the universe would not be representable on a digital computer.

More seriously, because a computer can only store a finite amount of data, it can only represent the schema of a finite simplicial complex, a simplicial complex which contains a finite number of simplexes. A finite simplicial complex can only be homeomorphic to a compact manifold, hence only a compact 4-manifold is discretely representable by a schema on a computer. Unfortunately, a compact 4-manifold cannot accept a Lorentzian metric. If the space-time of our universe is Lorentzian, then our universe can only be a non-compact Lorentzian 4-manifold. One possible conclusion to draw is that one cannot represent a universe like our own on a computer if one tries to represent the entire 4-dimensional history of the universe.

As an alternative, the geometrodynamical formulation of general relativity employs a so-called ‘3+1’ decomposition of space-time. One chooses a 3-manifold Σ , and one studies the time-evolution of the geometry and matter fields on Σ . As the geometry and matter fields evolve, a 4-dimensional space-time unfolds. Such a space-time will, of necessity, have the topology of $\mathbb{R}^1 \times \Sigma$.

The geometrodynamical formulation is advantageous because of Moise’s tri-

angulation theorem for 3-manifolds, (Stillwell 1992, p25 and p242). Moise demonstrated that every 3-manifold is homeomorphic with a simplicial complex; one says that every 3-manifold can be ‘triangulated’. Although it is true that every n -manifold can be triangulated for $n \leq 3$, it is, to reiterate, unknown whether all 4-manifolds can be triangulated.

Moise’s theorem means that any possible topology of the spatial universe can be discretely represented with the schema of a 3-dimensional simplicial complex. Once again, however, a digital computer can only represent the schema of a finite simplicial complex. Whilst a compact 3-manifold is homeomorphic with a finite simplicial complex, a non-compact 3-manifold can only be homeomorphic with an infinite simplicial complex, a complex which contains an infinite number of simplexes.

Only a compact 3-manifold can be homeomorphic with a 3-dimensional simplicial complex whose schema is representable on a digital computer. Hence, if our universe is a program running on a digital computer, then our spatial universe must have a compact spatial topology in a continuum idealisation. The hypothesis that our universe is a program on a digital computer, predicts that the spatial universe is discrete, and yields the potentially testable prediction that our universe has compact spatial topology in a continuum idealisation.

The prediction of compact spatial topology means that the Euclidean \mathbb{R}^3 and hyperbolic H^3 FRW universes are both inconsistent with the computer program hypothesis. The only FRW universe which has both simply connected and compact spatial topology, is the S^3 -universe. Hence, the only simply connected FRW universe which could be discretely represented on a computer, is the S^3 -universe. There are, however, a host of multiply connected compact FRW universes. The spatial geometry of each such universe is obtained as a quotient Σ/Γ of a simply connected Riemannian space form⁶ Σ , where Γ is a discrete, properly discontinuous, fixed-point free subgroup of the isometry group $I(\Sigma)$, (O’Neill 1983, p243 and Boothby 1986, p406, Theorem 6.5).

Compact FRW models exist for any value of sectional curvature k . Of the 18 flat, $k = 0$, 3-dimensional Riemannian space forms, 10 are compact. Given that one can only create compact FRW universes on a computer, it follows that one can only create 10 topologically different $k = 0$ FRW universes on a computer.

All of the 3-dimensional Riemannian space forms of constant positive curvature are compact, hence they could all be created on a computer.

Whilst there are compact and non-compact quotients H^3/Γ , there are an infinite number of such compact quotients. The work of Thurston demonstrates that ‘most’ compact and orientable 3-manifolds can be equipped with a complete Riemannian metric tensor of constant negative sectional curvature. This means that ‘most’ compact, orientable 3-manifolds can be obtained as a quotient H^3/Γ of hyperbolic 3-space.⁷ One can therefore create an infinite number of possible negative curvature FRW universes on a computer. However, there is no compact $k = -1$ space form which is globally homogeneous. H^3 itself is the

⁶A complete and connected Riemannian manifold of constant sectional curvature is called a Riemannian space form.

⁷The meaning of ‘most’ in this context involves Dehn surgery, (Besse 1987, p159-160).

only globally homogeneous 3-dimensional Riemannian space form of constant negative curvature, and H^3 is, of course, non-compact. Given that one can only create a compact universe on a computer, one cannot create a $k = -1$ FRW universe on a computer which is globally homogeneous. Thus, if our own universe is a globally homogeneous $k = -1$ FRW universe, it cannot exist on a computer. However, a locally homogeneous $k = -1$ FRW universe, with compact, multiply connected topology, could exist on a computer, and it is only *local* homogeneity which our astronomical observations are capable of detecting.

In practice it is difficult to test the prediction of compact spatial topology. Observational evidence currently indicates that our universe is a FRW universe, but there is no observable parameter in a FRW model which determines the spatial topology. Thus, there is no necessary link between the spatial topology of a FRW universe and the value of the density parameter Ω_0 ; one cannot infer the spatial topology of our universe from Ω_0 .

However, in a ‘small’, compact, multiply connected universe, it is possible to see around the entire universe. To understand this, begin by recalling that a Riemannian manifold (Σ, γ) has a natural metric space structure. The metric tensor γ determines a Riemannian distance $d(p, q)$ between any pair of points $p, q \in \Sigma$. The Riemannian distance $d(p, q)$ is dimensionless, in the sense that it lacks any physical units. In a FRW model, it is the scale factor $R(t)$ which introduces physical units of distance. The physical distance between p and q at time t is $R(t)d(p, q)$. Because $R(t)$ has physical units, so does $R(t)d(p, q)$.

For any FRW universe, one can calculate the maximum Riemannian distance, d_{max} , that light has travelled by a time t_0 , which is considered to be the present time. The relevant equation is

$$d_{max}(t_0) = \int_0^{t_0} \frac{c}{R(t)} dt$$

A civilization located at some point p in space, will, at time t_0 , be able to see no further, in any direction, than a Riemannian distance of $d_{max}(t_0)$. This distance can therefore be referred to as the Riemannian horizon distance. It is, of course, a dimensionless quantity.

Now, recalling that the diameter of a metric space is the supremum of the distances which can separate pairs of points, it is a fact that any compact Riemannian manifold is a metric space of finite diameter. If one created, on a computer, a FRW universe in which (Σ, γ) were a compact Riemannian manifold of sufficiently small diameter, $diam(\Sigma, \gamma)$, then the Riemannian horizon distance $d_{max}(t_0)$ could exceed $diam(\Sigma, \gamma)$ by the time $t_0 \sim 10^{10}$. If so, the horizon would have disappeared for the observers in that universe. They would be able to see their entire spatial universe. No point of their universe could be separated from them by a Riemannian distance greater than $diam(\Sigma, \gamma)$, so if $d_{max}(t_0) \geq diam(\Sigma, \gamma)$, then they would be able to receive light from all regions of their spatial universe.

In such universes, individual galaxies and clusters of galaxies would produce multiple images upon the celestial sphere of planet-bound observers, (see El-

lis 1971). Different compact spatial topologies and geometries would produce different patterns of ghost images and multiple images upon the celestial sphere.

However, although compact spatial topology is a necessary condition for the entire spatial universe to be visible, it is not a sufficient condition. Our universe might have compact spatial topology, but if it is a ‘large’ compact universe, then all of space will not be visible. For all of space to be visible when the universe is only $\sim 10^{10}$ yrs old, the Riemannian manifold (Σ, γ) which represents the spatial universe must be sufficiently small, as well as compact. Even if our spatial universe is small and compact, it would be extremely difficult to identify multiple images of galaxy clusters. Hence, although the presence of multiple images would verify the hypothesis of a small, compact universe, the fact that they have not been identified at the current time does not falsify the hypothesis. A better means of testing the hypothesis is to search for paired circles in the microwave background radiation. Recent research indicates that if such paired circles exist, then one could derive the spatial topology from the specific pattern of paired circles, (see Cornish, Spergel, Starkman 1998). The CMBR power spectrum can also be used to determine whether our spatial universe is a small compact universe. A small compact universe would affect the CMBR power spectrum on large angular scales. The WMAP satellite has revealed anomalies in the CMBR power spectrum on large angular scales. The quadrupole $l = 2$ mode was found to be about 1/7 the strength predicted for an infinite flat universe, while the octopole $l = 3$ mode was 72% of the strength predicted for such a non-compact $k = 0$ universe, (Luminet *et al* 2003, p3).

The presence of paired circles or specific anomalies in the CMBR power spectrum would verify that the universe is spatially compact, and would thereby verify the computer program hypothesis. Unfortunately, the absence of paired circles or anomalies in the power spectrum would not entail that the spatial universe is non-compact. Our universe could simply be a large compact universe. Hence, the absence of paired circles or anomalies in the power spectrum would not falsify the computer program hypothesis.

Predictions about the lifetime of our universe are easier to test than predictions about the spatial topology. The lifetime of our universe is determined by parameters such as the Hubble parameter H_0 and the density parameter Ω_0 , which can be inferred from observation. Hence, if the computer program hypothesis made predictions about the lifetime of our universe, it would be easier to test it. If a universe is represented by a Lorentzian manifold (M, g) , then the lifetime of the universe corresponds to the ‘timelike diameter’ of (M, g) . The timelike diameter of (M, g) is the supremum of the length of all past-directed timelike curves in (M, g) . As Beem and Ehrlich comment, “the timelike diameter represents the supremum of possible proper times any particle could possibly experience in the given space-time,” (Beem and Ehrlich 1980, p329).

If a Lorentzian manifold with an infinite timelike diameter were represented by a numerical solution of the Einstein geometrodynamical equations, and if the size of the time steps in the numerical solution were constant, then an infinite number of time steps would be necessary. An infinite amount of information would have to be processed. Alternatively, if the size of the time steps diverge

exponentially as $t \rightarrow \infty$, a numerical solution would only require a finite number of time steps. The ever-expanding $k \leq 0$ FRW universes are examples of universes with an infinite lifetime. If a computer in a universe with an infinite lifetime could process information at a constant rate, then it could process an infinite amount of information. However, an ever-expanding universe will suffer an entropy ‘heat death’, the amount of free energy available converging to zero as $t \rightarrow \infty$. Brillouin’s inequality entails that there is a minimum, positive amount of free energy which must be expended to process a bit of information. Where ΔI is the amount of information processed in bits,

$$\Delta I \leq \Delta E/k_B T \ln 2.$$

ΔE is the free energy expended, T is the absolute temperature in degrees K, and k_B is Boltzmann’s constant, (Barrow and Tipler 1986, p661). At first sight, this suggests that it is impossible to process an infinite amount of information in an ever-expanding universe because the amount of free energy converges to zero. However, the amount of energy which must be expended per bit of information processed is temperature dependent. From the inequality above, one can derive the following constraint on the rate at which information can be processed:

$$dI/dt \leq \frac{dE/dt}{k_B T \ln 2}.$$

In turn, this entails the following constraint on the total amount of information I which can be processed between the current time t_0 and some future time t_f , which might be ∞ :

$$I = \int_{t_0}^{t_f} \frac{dI}{dt} dt \leq (k_B \ln 2)^{-1} \int_{t_0}^{t_f} T^{-1} \frac{dE}{dt} dt.$$

If the temperature converges to zero, $T \rightarrow 0$, as it does in an ever-expanding $\Omega_0 \leq 1$ universe, then the amount of free energy which needs to be expended per bit converges to zero. Hence, although the amount of free energy converges to zero, so also does the amount of free energy which needs to be expended per bit. Thus, because the integral $\int_{t_0}^{\infty} T^{-1} (dE/dt) dt$ can diverge, it may still be possible to process an infinite amount of information in an ever-expanding universe, even if the total free energy expended $\int_{t_0}^{\infty} (dE/dt) dt$ is finite, (Barrow and Tipler 1986, p663).

In $\Omega_0 > 1$ universes, because the temperature diverges near the final singularity, the rate at which free energy is expended dE/dt , and therefore the total energy expended $\int_{t_0}^{t_f} (dE/dt) dt$, must diverge if the total information processed is to diverge.

If $\Omega_0 \leq 1$ in our universe, as current astronomical evidence indicates, then our universe has an infinite timelike diameter. Assuming that the simulation of such a universe would require an infinite amount of information to be processed, the possibility of the computer program hypothesis then rests upon whether it is physically possible for the integral $\int_{t_0}^{t_f} T^{-1} (dE/dt) dt$ to diverge in either a $\Omega_0 \leq 1$ universe or a $\Omega_0 > 1$ universe. In both cases this remains a matter of

debate. If it is not physically possible for the integral to diverge in either case, and if the observation that $\Omega_0 \leq 1$ in our universe is reliable, then could one conclude that our universe is not a program running on a computer in another universe? If it is impossible to process an infinite amount of information, then the only type of universe which could be *entirely* simulated on a computer would be a finite lifetime universe. However, it remains possible that a partial simulation of a $\Omega_0 \leq 1$ universe could be created on a computer in another universe, a finite lifetime subset of the entire $\Omega_0 \leq 1$ universe. Hence, even if our universe is a $\Omega_0 \leq 1$ universe, and even if it is impossible to process an infinite amount of information, our universe could be a finite lifetime simulation running on a computer in another universe.

Not only could the computer program hypothesis be falsified by empirical investigation, but there are logical constraints upon what it is possible to simulate on a computer.

A computer is a finite volume subsystem of a universe which is capable of representing the state of other systems. A system can represent, exactly and completely, the state of another system, if and only if the amount of information which can be coded in the first system is greater than or equal to the amount of information which can be coded in the other system. An entire universe is a special type of system. Hence, a subsystem of a universe \mathcal{A} can represent, exactly and completely, the state of a universe \mathcal{B} , if and only if the amount of information which can be coded in the subsystem of \mathcal{A} is greater than or equal to the amount of information which can be coded in universe \mathcal{B} .

As a special case, if the amount of information which can be coded in a subsystem of a universe \mathcal{A} is less than the amount of information which can be coded in the entire universe \mathcal{A} , then it is impossible for the subsystem of universe \mathcal{A} to represent, exactly and completely, the entire universe \mathcal{A} .

The amount of information which can be coded in a system is determined by the number of possible different states of the system. If N denotes the number of possible states, then the amount of information I which can be coded, in bits, is simply $I = \log_2 N$. Hence, if the number of possible states of a subsystem of a universe is less than the number of possible states of the entire universe, then it is impossible for that subsystem to represent, exactly and completely, the entire universe.

However, just because a system is a subsystem of a universe, it does not follow that the number of possible states of the system is less than the number of possible states of the universe. True, if the number of possible states of a subsystem is finite, then by virtue of being a subsystem, that finite number must be smaller than the number of possible states of the entire universe. For every state of a subsystem, there must be multiple states of the entire universe which induce the same state upon that subsystem, hence the number of possible states of the entire universe must be larger. However, if the number of possible states of a subsystem is not finite, then it is possible that it has the same number of possible states as the entire universe. *A priori*, it is quite possible that a

subsystem of a universe \mathcal{A} , and the entire universe \mathcal{A} , both possess an infinite number of states. If the state space of a subsystem has the same cardinality as the state space of the entire universe, then, by definition, there exists at least one bijective mapping between the two state spaces. Any such bijective mapping would enable the states of the entire universe to be represented by the states of the subsystem.

This argument can be presented in another way. If a subsystem \mathcal{S} of our universe represents the entire universe \mathcal{U} , it must also represent \mathcal{S} representing \mathcal{U} . If it does this, it must also represent \mathcal{S} representing \mathcal{S} representing \mathcal{U} . And so on, *ad infinitum*. This is possible only if the subsystem can store an infinite amount of information.

If the entire universe only has a finite number of possible states, then a subsystem will also have a finite number of states, and the number of subsystem states will be smaller than the number of universe states. However, if the entire universe has an infinite number of possible states, then a subsystem could have either a finite number or an infinite number of possible states.

If the entire universe has an infinite number of possible states, then it could conceivably possess either a continuous infinity of possible states, or a discrete infinity. A digital computer could only represent the universe exactly if the universe is discrete, hence the only case of interest is the case in which the universe has a discrete infinity of possible states. A digital computer could only represent the universe exactly and completely if the entire universe and the computer subsystem of the universe both possess a discrete infinity of possible states. In other words, a digital computer could only represent the universe exactly and completely if both the computer and the universe can code a discrete infinity of information.

A computer is a finite volume subsystem of the universe, hence to determine if a computer could code the same amount of information as the entire universe, it is necessary to determine if a finite volume subsystem can code a finite or infinite amount of information. To answer this question, it is necessary to determine what the physical structure of the universe is.

At present, it appears that there are discrete levels of physical structure in the universe. All macroscopic material objects in our universe are composed of chemical elements and chemical compounds. The latter are composed of atoms in different combinations and organizations. Atoms are composed of electrons and atomic nuclei. The nuclei of atoms are themselves composed of protons and neutrons, which are themselves composed of quarks. The parts of material objects do not appear to lie on a continuum.

Electrons and quarks are purported to be elementary particles, pieces of matter which have no parts. If elementary particles do exist, then our universe could be said to have a finite lower level of structure. There would be no levels of structure below the level of elementary particles.

I propose that a finite volume subsystem is limited to coding a finite amount of information if and only if the following three conditions are satisfied:

- The number of structure levels available in a finite volume of space is

finite.

- On each structure level, there is a finite set of parts in a finite volume of space.
- Each of the parts on each level of structure has a finite set of states.

A finite volume subsystem which satisfies these conditions has only a finite number of possible states, and therefore cannot code the same amount of information which can be coded in the entire universe.

To reiterate, a computer could only represent the universe exactly and completely if a finite volume subsystem can code a discrete infinity of information. It seems safe to assume that, on each level of structure, there is a finite set of parts in any finite volume of space. The Bekenstein bound⁸ and the so-called holographic bound of Susskind and 't Hooft, purportedly entail that the parts on each level of structure have a finite set of states, (Bekenstein 2003). Moreover, the existence of elementary particles would mean that there is a finite set of structure levels in each finite volume of space. It would appear, therefore, at first sight, that all three conditions are satisfied. It would appear that a finite volume subsystem cannot code a discrete infinity of information, and it would appear that a computer cannot represent the universe exactly and completely.

However, further thought raises some doubts. Both the Bekenstein bound and the holographic bound place an upper limit on the *entropy* within a finite volume of space. Given a finite quantity of weakly self-gravitating energy E in a spherical volume of radius R , which is isolated from other systems, (Bekenstein 2004), the entropy S is subject to the following upper bound:

$$S \leq 2\pi ER/\hbar c .$$

The holographic bound is independent of the quantity of energy, and places the following limit on the entropy of a spherical volume of radius R , which is isolated from other systems:

$$S \leq \pi c^3 R^2/\hbar G .$$

In both cases, it is then assumed that a finite upper limit to the entropy of a finite volume of space entails a finite upper limit to the information storage capacity of that volume. This might be inferred from the following relationship:

$$\text{Information} = \text{Maximum entropy} - \text{entropy} .$$

By implication, it is the *statistical states* or *macrostates* of a system which are the bearers of entropy and information here. The states which provide a complete, detailed description of a system are referred to as 'microstates'. A statistical state expresses only partial knowledge of the state of a system, and, in classical mechanics at least, corresponds to a probability distribution ρ defined upon the space of microstates Γ . A macrostate is a set of macroscopically

⁸Otherwise known as the universal entropy bound.

indistinguishable microstates Γ_M , and corresponds to a special type of statistical state in which the probability distribution is of a constant value $|\Gamma_M|^{-1}$ on Γ_M , and zero elsewhere. $|\Gamma_M|$ denotes the volume of Γ_M . The microstate of a system inherits the entropy and information of the macrostate to which it belongs. The entropy of an isolated system increases because the microstate of the system moves into macrostates of ever greater entropy. The equation above means that the information possessed by a system at a point in time is the difference between the maximum entropy of the system, and the entropy possessed by the system at that point in time. The maximum information which can be possessed by a system is that which it possesses when the system's entropy is zero. Hence, according to the relationship above, the maximum information equals the maximum entropy.

Whether this entails that a finite volume of space possesses a finite number of states is a different question. In classical mechanics, a system consisting of n particles has a $6n$ -dimensional *continuum* state space Γ , called the phase space. The entropy $S(\rho)$ of statistical state ρ in classical mechanics is defined to be

$$S(\rho) = -k_B \int_{\Gamma} \rho \log \rho \, d\mu ,$$

where k_B is Boltzmann's constant. In the case of a macrostate ρ_M , this reduces to

$$\begin{aligned} S(\rho_M) &= -k_B \int_{\Gamma_M} |\Gamma_M|^{-1} \log |\Gamma_M|^{-1} d\mu \\ &= -k_B |\Gamma_M|^{-1} \int_{\Gamma_M} \log 1 - \log |\Gamma_M| d\mu \\ &= k_B |\Gamma_M|^{-1} \int_{\Gamma_M} \log |\Gamma_M| d\mu \\ &= k_B \log |\Gamma_M| . \end{aligned}$$

Hence, although the entropy of a macrostate of such a system can be finite, it corresponds to a continuum of possible microstates. An upper limit to entropy does not entail a finite number of possible states. I propose that the link between entropy and information storage capacity is only valid for finite state-space systems. When a system has an infinite number of states, but a finite maximum entropy, I propose that it has an infinite information storage capacity. Ultimately, each different state of a system can represent different information, so a system with an infinite number of possible states, but a finite volume state space, and therefore a finite maximum entropy, nevertheless has an infinite information storage capacity.

To argue that a finite volume of space possesses a finite information storage capacity, one might alternatively start from loop quantum gravity, and try to argue that a finite volume of space only possesses a finite number of quantum states. A finite volume of space corresponds to a finite number of spin network nodes, and for a fixed finite number of nodes, there are a finite number of

spin network states. For a system with a finite number of microstates, each macrostate M corresponds to an equivalence class containing a finite number of microstates, $\text{Num}(M)$. The entropy of such a macrostate is simply

$$S = k_B \log \text{Num}(M) .$$

Hence, a system with a finite number of microstates possesses a finite maximum entropy, and therefore possesses an upper limit on its information storage capacity.

However, quantum theory may not be the definitive theory of the physical world. A quantum state may correspond to many, or an infinite number of actual states. Even though there may be only a finite number of quantum states for a finite volume of space, there may be an infinite number of actual states. It may be that quantum theory is only valid for certain levels of structure, and it might merely be that the amount of information which can be coded above a certain length scale, or the amount of information which can be coded in a certain way, is finite.

There is also no decisive evidence that elementary particles exist. If the current candidates for elementary particles, such as quarks, do have parts, then those parts might only be detectable at energies which are not currently available in particle accelerators.

One could also dispute the assumption that, on each level of structure, there is a finite set of parts in any finite volume of space. If each part has a non-zero spatial extension with a well-defined boundary, and if the parts cannot interpenetrate, then it does indeed follow that there can only be a finite set of parts packed into a finite volume of space. However, parts in quantum theory do seem able to interpenetrate each other to some degree. If there are levels of structure below the levels of the electron and quark, one might find very strange things, beyond even quantum theory, like an infinite number of parts interpenetrating each other in a finite volume of space.

Tipler claims that there could be a hierarchy of computer universes, just like the hierarchy of so-called ‘virtual machines’ which can exist on a computer, and he claims that we would not know which level of the hierarchy our own universe exists at. Whilst I have argued that the Bekenstein bound does not entail that a finite volume subsystem has only a finite number of possible states, Tipler accepts this implication. This, I propose, is inconsistent with the claim that we would not know which level of a universe hierarchy our own universe exists at.

When one computer is programmed so that it precisely mimics the input-output behaviour of another computer, the latter is said to be emulated by the former. The emulation program, running on the real computer, is said to be a virtual machine. A real machine T_1 can be programmed to emulate another, producing a virtual machine T_2 . The virtual machine T_2 can then be programmed to emulate another computer, producing a higher level virtual machine T_3 . These levels are referred to as levels of implementation.

A universe running on a computer could itself contain computers, upon which other universes are running. The universes would be running at different levels of implementation, and Tipler suggests, (1995, p208), that in this case, the levels should be thought of as levels of reality. Tipler seems to assume that there must be a lowest level of the hierarchy, and refers to this as ‘ultimate reality’. Tipler claims that “we cannot know if the universe in which we find ourselves is actually ultimate reality,” (ibid.).

However, whilst any one computer may be able to emulate the input-output behaviour of another, that does not entail that any one computer has the same representational capacity as another. An actual computer, with a finite memory, does not have the same representational capacity as every other computer. A computer with N bytes of memory does not have the same representational capacity as a computer with M bytes of memory if $M > N$. There may be data structures which the computer with M bytes of memory can represent, but which the computer with N bytes cannot.

It was argued above that a computer with a finite set of states, (and hence a finite memory), cannot perfectly represent the universe to which it belongs. This is because a computer with a finite memory cannot code the same amount of information as the universe to which it belongs. In general, a computer with a finite memory cannot perfectly represent any universe which can code a greater amount of information than the computer. Any universe which can code a greater amount of information than the universe to which the computer belongs, will code more information than the computer.

If one accepts Tipler’s claim that “complexity is appropriately measured by the number of possible alternative states a system can be in,” (1995, p118), then the complexity of a system can also be measured as the amount of information which that system can code.⁹ If one accepts that a finite volume subsystem has only a finite number of possible states, then a computer can only have a finite memory. If a computer can only have a finite memory, then a computer cannot perfectly represent a universe of the same complexity, or greater complexity, than the universe to which the computer belongs. The complexity of a universe is observable, hence, *contra* Tipler, the levels of implementation are distinguishable. If a finite volume subsystem has only a finite number of possible states, then each higher level of universe implementation is less complex than the level below. A computer with a finite memory cannot perfectly represent a universe unless that universe is simpler than the universe to which the computer belongs. The more complex the universe one belongs to, the lower down the hierarchy that universe is placed. A universe of maximal complexity, if there is such a thing, could be proven to be the universe of ultimate reality.

⁹This should not be confused with the *computational complexity* of an algorithm used to calculate the values of a function. This is a measure of the growth in computation time with the growth of the size of the input. Those functions which are computable are divided into those which are calculable by an algorithm in polynomial time, **P**, and those which are not, **NP**. Tipler’s notion of complexity is also distinct from the *Kolmogorov complexity* of an object, which is the length, in bits, of the shortest computer program capable of producing the digital representation of that object, as output.

If our universe is a computer program running on a computer in another universe, then that universe must have a higher level of complexity to our own. This greater complexity might take the form of a higher number of spatial dimensions. If the metaphysical universe has a higher number of dimensions than our own, this would be consistent with the proposal for a projective relationship between a metaphysical reality and a physical reality.

Of course, if a finite volume subsystem has a discrete infinity of possible states, then a computer might be able to perfectly represent a universe with the same complexity as the universe to which the computer belongs. If so, then the levels of universe implementation might all have the same level of complexity. The point is that, if the Bekenstein bound does entail that a finite volume subsystem has only a finite number of possible states, then the Bekenstein bound is inconsistent with the thesis that universes at different levels of implementation are indistinguishable.

4 Reductionism, identity, and universe creation on a computer

The suggestion that a physical system can be perfectly simulated on a computer appears to have anti-reductionistic implications. Suppose, for example, that a tornado could be perfectly simulated on a computer. A tornado is described by a solution of the Navier-Stokes equations.¹⁰ To simulate a tornado on a computer, one would define program variables to represent the air pressure, velocity, density etc. in a volume of space, and one would represent the tornado by calculating a solution of the Navier-Stokes equations for these variables. Whilst a ‘real’ tornado is realised upon a collection of air molecules, a simulated tornado is realised upon the components and circuitry of a computer. Hence, if a tornado could be perfectly simulated on a computer, then a tornado could be realised on more than one medium. Two completely different lower-level processes would correspond to the same higher-level process.

Epistemological reductionism asserts that what can be known or theoretically represented about the higher levels of a composite system can be reduced to what can be known or theoretically represented about the lower levels. Ontological reductionism asserts that what exists on the higher levels of a composite system can be reduced to what exists on the lower levels. I propose that ontological reductionism is the conjunction of the two following assertions:

1. The higher-level properties of a composite system uniquely determine the parts of the system and the way in which the parts are organized and interact. In other words, the higher-level properties of a composite system uniquely determine the properties of the subsystems and the relationships

¹⁰There is, for example, an exact solution of the Navier-Stokes equations called the Sullivan Vortex, which describes the flow in an intense tornado with a central downdraft.

between the subsystems.¹¹

2. The parts of a system, and the way in which the parts are organized and interact, uniquely determine the higher-level properties of the system. In other words, the properties of the subsystems in a composite system, and the relationships between the subsystems, uniquely determines the higher-level properties of the composite system.

The second assertion on its own is ontological *supervenience*. This is the idea that there can be no difference in the higher-level state of a composite system without a difference in the lower-level state, otherwise one would have a one-many correspondence between the lower-level states and higher-level states.

Some of the expressions used here require some explanation. A composite system is simply a system composed of multiple parts, whilst a higher-level property of a composite system is a property which can be possessed by the whole system. In some cases, such as shape and size, a higher-level property can also be possessed by the subsystems. In other cases, a higher-level property cannot be possessed by any individual part. For example, liquidity is a higher-level property of matter which can only be possessed by a collection of particles. Liquidity cannot be possessed by an individual particle.

If a system satisfies some law of physics, then that law of physics is a property of the system. If the system is a composite system, and it satisfies some law of physics, then that law of physics is a higher-level property of the system. In one sense, the laws of physics are properties of properties of physical systems, or relationships between the properties of a physical system. If energy, position, and velocity, for example, are properties of physical systems, then evolution equations govern the time evolution of these properties, and constraint equations govern the possible relationships between the different properties.

The claim that the laws of Nature are relationships among properties is part of the ‘N-relation’ account of the laws of Nature. Consider the example $F = Ma$. There are two levels on which this law expresses a relationship between properties. The first level relationship is between the general properties of acceleration, force and mass. This is a relationship between so-called ‘determinable’ properties. This relationship between determinable properties entails an infinite number of other relationships between the so-called ‘determinates’ of those properties. Determinates are more specific properties: If the property of mass is a determinable property, then a mass of 3.0kg, or a mass of 23.7kg, are examples of determinate properties. The general relationship $F = Ma$ entails an infinite number of relationships between the determinates. For example, an acceleration of 5ms^{-2} , a mass of 10kg, and a force of 50N, are so related. However, even a determinate property such as a mass of 3.0kg, is a universal, which can be possessed by different objects at different times and places. One can have many different instances of a mass of 3.0kg. Hence, on a third level, the general law expresses a relationship between all instances of the determinates.

¹¹Epistemological reductionism asserts that the higher-level states and properties are *definable* in terms of the lower-level states or properties.

The organization and interaction between the parts of a composite system includes both the spatial arrangement of the parts, and the mutual forces exerted between them. For example, if the parts are atoms or molecules, the properties of the composite system depend upon whether the parts are arranged randomly or as a crystal lattice, and depend upon whether the electromagnetic bonds between the parts are covalent bonds, or Van de Waals bonds, etc.

Statement 1 claims that the entire set of higher-level properties, in combination, determines uniquely the parts of the system and the way they are organized and interact. Individual higher-level properties can be possessed by many different systems, composed of different parts. For example, liquidity is a higher-level property which can be possessed by many different chemical substances. Liquidity is not a property unique to water. A collection of helium atoms, or mercury atoms, for example, can be in a liquid state. Hence, liquidity is a higher-level property of a system which does not uniquely determine the parts of the system. The parts of any liquid body do interact in a similar manner, whatever the types of particle involved, hence liquidity, as a higher-level property, does determine the type of the relationships between the parts, but it does not determine what the parts will be. Furthermore, the specific liquid properties, such as viscosity, density, or pressure, will be different for different chemicals. Liquids of different chemical composition will all share the property of liquidity, but they will differ in the complete range of higher-level properties they possess.

If a physical system could be realised on more than one medium, this would entail the falsity of statement 1. It would not, however, affect statement 2, the principle of supervenience. For example, the properties of a tornado might not determine a unique medium upon which it must be realised, but the properties of air molecules, and the relationships between air molecules entail that a tornado can be realised on a collection of air molecules. Similarly, if it were possible to realise a tornado on a computer, then it would be the properties of, and relationships between, the components and circuitry of a computer which would entail that a tornado can be realised upon a computer.

Whether or not a physical system can in fact be realised on more than one medium depends upon how one defines the identity of a system. In the case of a tornado there are two possible approaches:

- (a). A tornado is a physical system composed of atmospheric molecules, which has the property that it satisfies a tornado-solution of the Navier-Stokes equations. The identity of a tornado is inseparable from being a collection of atmospheric molecules. Under this approach, a tornado is not realised upon a collection of atmospheric molecules, it is composed of atmospheric molecules. Under this approach, a tornado cannot be realised on more than one medium because there is no sense in which a tornado is realised on any medium. It is only if the identity of a tornado can be defined in a formal, mathematical sense, that one can speak of a tornado being realised upon a medium.
- (b). If the identity of a tornado is defined by a solution of the Navier-Stokes

equations, and if the identity of a solution of the Navier-Stokes equations is independent of any particular medium, then the identity of a tornado is independent of any particular medium. Under this approach, the identity of a tornado is independent of its realisation upon a collection of atmospheric molecules. If the components and circuitry of a computer can realise a tornado-solution of the Navier-Stokes equations, then a tornado can be realised on the components and circuitry of a computer.

The identity of a solution to the Navier-Stokes equations is independent of any particular physical medium because a solution of a differential equation is merely a mathematical object. A solution to a differential equation is given physical meaning when the solution variables are given a physical reference i.e. physical units. The solution variables can refer to many different things. Consider the diverse referents of solutions to the wave equation and the diffusion equation.

When a solution of the Navier-Stokes equations is realised on a medium, the solution variables are given reference. When a solution of the Navier-Stokes equations is realised on the medium of atmospheric molecules, the solution variables refer to air pressure, velocity, density etc. If a tornado could be realised on an economic system, the solution variables would refer to economic quantities.

For a computer to be able to realise a tornado-solution of the Navier-Stokes equations, the computer must possess physical properties which can be the referents of the solution variables for a tornado-solution. These physical properties of the computer might well be compound or collective properties, (compared to the possibly fundamental properties of the simulated system), but they must be physical. If a tornado can be realised on a computer, the solution variables do not refer to properties of the atmosphere, such as pressure, velocity, density etc. Instead, they refer to properties of the computer components and circuitry, such as, perhaps, the voltage states of the bytes in computer memory. The medium upon which a solution is realised is defined by the referents assigned to the solution variables.

A computer cannot realise a solution of the Navier-Stokes equations in which the solution variables refer to atmospheric pressure, velocity, density etc. because a computer does not possess these quantities. A computer can only *represent* variables such as atmospheric pressure. There is, therefore, a subtle but important distinction between (i) the realisation of a tornado-solution of the Navier-Stokes equations on a computer, and (ii) the representation on a computer of a realisation of a tornado-solution upon the medium of atmospheric molecules. A computer simulation of a tornado cannot realise a tornado in the sense of realizing a system with atmospheric pressure, velocity, density etc on the computer.

It is possible to accept approach (b), that the identity of a tornado is independent of any particular medium, without accepting that a tornado can be realised on a computer. A computer does not possess physical properties which can be the referents of the solution variables for the Navier-Stokes equations. One reason is that the solution variables are continuous, whilst the logical

states of electronic circuits are discrete. The example of a tornado-solution to the Navier-Stokes equations is probably a bad one at this juncture because the Navier-Stokes equations, and fluid mechanics *in toto*, merely provide a phenomenological approximation. A tornado-solution of the Navier-Stokes equations is not exactly realised on the medium of air molecules either. However, even if one goes down to the level of fundamental physics, a computer cannot exactly realise solutions to the fundamental equations of physics either. The reason is twofold:

- There is a one-many correspondence between the logical states and the exact electronic states of circuits.
- The logical states of multiple-bits in computer memory only represent numbers because they are deemed to do so under a numeric-interpretation.

In current computers, each bit of memory corresponds to an electrical circuit, and the two possible logical states of the bit correspond with two possible voltages between fixed points of the circuit. A voltage-value lies on a continuum. The logical state of 1 is not defined by a single precise voltage value, but by a range of values. The logical state of 0 is defined by a different range of possible voltages. There is, therefore, a one-many correspondence between logical states and voltage levels. Successive runs of the same program will not produce exactly the same sequence of electronic states in computer memory. The exact voltage levels will be different on successive runs. This level of electrical noise prevents a contemporary computer from exactly realizing anything, even discrete objects.

This suggests that a tornado-solution of the Navier-Stokes equations can only be approximately realised on a digital computer. This is crucial to the reductionistic question of whether the same physical system can be realised on more than one physical medium. If there cannot be an exact realisation of a tornado on the medium provided by the components and circuitry of a computer, this is presumably because the properties of, and relationships between, the components and circuitry of a computer differ from the properties of, and relationships between, the air molecules in a region of the atmosphere.

The condition that a computer must possess *physical* properties which can be the referents of the solution variables for a tornado solution, means that the properties must be objective properties if a computer is to realise a tornado-solution of the Navier-Stokes equations. Given the many-one correspondence between exact electronic states and logical states, the exact electronic properties of a computer's components cannot be the referents of the solution variables. Moreover, it is the numeric interpretation of the logical states of multiple electrical circuits which are the candidate properties. It is the pattern of numbers represented by a computer which resembles the pattern of values for the physical quantities of a simulated system. As will be explained at length, the numbers represented by a computer are interpretation-dependent, hence the numbers represented by a computer cannot be objective properties of the computer. The referents of the solution variables must be objective, not interpretation-dependent,

hence a computer cannot realise a tornado-solution of the Navier-Stokes equations, or any other physical system for that matter. If the numbers represented by the computer are interpretation-dependent, then the pattern of numbers represented by the computer must be an interpretation-dependent pattern. Hence, the resemblance between the pattern of numbers represented by the computer and the pattern of values for the physical quantities of a simulated system must be an interpretation-dependent resemblance.

Change the interpretation of the logical states of multiple electrical circuits, and there is no resemblance, not even an approximate one. Even if there was no electrical noise, and even if the simulated system was discrete itself, (even if there was a bijective correspondence), it would still be an interpretation-dependent resemblance. The numbers represented by the computer are not compound physical properties of the computer.

5 A digital computer simulation of a universe cannot exist as a universe

A digital computer simulation of a physical system cannot exist as, (does not possess the properties and relationships of), anything else other than a physical process occurring upon the components of a computer. In the current case of an *electronic* digital computer, a simulation cannot exist as anything else other than an electronic physical process occurring upon the components and circuitry of a computer. The general structure of the argument which establishes this conclusion is as follows:

1. A digital computer simulation is a type of representation.
2. There are three types of representation.
3. A digital computer simulation is a special case of the type of representation in which there is no objective relationship between the represented thing and the thing which represents it.
4. If there is no objective relationship between a universe and a digital computer simulation of a universe, then a digital computer simulation of a universe cannot exist as a universe.

The reasoning that justifies claim 3, outlined at the end of the previous section, is basically as follows: In a computer simulation, the values of the physical quantities possessed by the simulated system are represented by the combined states of multiple bits¹² in computer memory. However, the combined states of multiple bits in computer memory only represent numbers because they are deemed to do so under a numeric interpretation. There are many different interpretations of the combined states of multiple bits in computer memory. If the numbers represented by a digital computer are interpretation-dependent,

¹²Qubits in the case of quantum computers.

they cannot be objective physical properties. Hence, there can be no objective relationship between the changing pattern of multiple bit-states in computer memory, and the changing pattern of quantity-values of a simulated physical system.

Because a digital computer simulation of a universe cannot exist as a universe, it is, *a fortiori*, impossible for anyone to be embedded in a digital computer simulation. It is impossible for our experience to be indistinguishable from the experience of someone embedded in a digital computer simulation because it is impossible for anyone to be embedded in a digital computer simulation.

Tipler assumes that if a universe is simulated on a computer, then the simulation exists as a universe, at a so-called ‘higher level of implementation’. This ontological assumption can be generalized to the following proposition: If a physical system of type \mathcal{S} is simulated on a computer, then the simulation exists as a system of type \mathcal{S} , at a higher level of implementation. For example, if a tornado is simulated on a computer, it could be claimed that the simulation exists as a tornado, at a higher level of implementation. In opposition, it can be argued that a digital computer simulation of a physical system, even a perfect simulation, cannot exist as the thing it represents. Note, this does not entail that there is no such thing as a simulation of a physical system. A simulation of a physical system does exist, but it exists only as a physical process occurring upon the hardware of the computer.

A computer simulation is a special type of representation. A current digital computer can electronically encode numbers, and because numbers can be used to represent physical systems, a computer can represent physical systems. A digital computer simulation of a physical system is an evolving, automated, quantitative, adjustable, and encoded description. It is most important to appreciate that a digital computer simulation represents a physical system by means of an encoded description. Mathematical physics is able to describe physical systems in terms of numbers, and a current digital computer simulation electronically encodes the numerical description provided by mathematical physics.

To speak of an encoded description, is not to refer to the code of a programming language; in terms of current computer technology, it means, rather, that the description provided by mathematical physics in terms of numbers, undergoes a transformation into the states of electronic circuits. This transformation is an encoding transformation. To relate a computer simulation to the physical world it is necessary to use a decoding transformation. This decoding transformation maps the states of electronic circuits back into numbers.

The description provided by mathematical physics is itself a type of representation of the physical world. A current digital computer simulation represents a physical system by electronically encoding the numerical representation provided by mathematical physics.

A representation is defined by a mapping f which specifies the correspondence between the represented thing and the thing which represents it. An object, or the state of an object, can be represented in two different ways:

1. If an object/state is a structured entity M , it can provide the entire domain of a mapping $f : M \rightarrow f(M)$ which defines the representation. The range of the mapping, $f(M)$, is also a structured entity. The mapping f is a homomorphism with respect to some level of structure possessed by M and $f(M)$.
2. An object/state can be an element $x \in M$ in the domain of a mapping $f : M \rightarrow f(M)$ which defines the representation.

The representation of a Formula One car by a wind-tunnel model is an example of type-1 representation. The representation of a number by the electronic state of a word¹³ of memory in a current digital computer, is an example of type-2 representation.

In the example of the wind-tunnel model, there is an approximate homothetic isomorphism¹⁴ from the exterior surface of the wind-tunnel model to the exterior surface of the actual Formula One car. This notion of structure preservation can be seen in other cases of representation. The notorious map of the London Underground does not preserve geometry, but it does preserve the topology of the network. Hence, there is a homeomorphic isomorphism involved.

There is no homomorphism between a number and the electronic state of a word of computer memory. Each number is merely an element in the domain of a mapping which maps numbers to the electronic states of a word of computer memory. There are many ways to represent a number by the state of a word of computer memory. Moreover, the same electronic states of a word in computer memory can represent things other than numbers; they can represent character symbols, or parts of images and sounds.

Type-2 representation has two sub-types. The mapping $f : M \rightarrow f(M)$ can be defined by either (2i) an objective, causal physical process, or by (2ii) the decisions of thinking-beings.

The primary example of type-2i representation is the representation of the external world by brain states. Taking the example of visual perception, there is no homomorphism between the spatial geometry of an individual's visual field, and the state of the neuronal network in that part of the brain which deals with vision. However, the correspondence between brain states and the external world is not an arbitrary mapping. It is a correspondence defined by a causal physical process involving photons of light, the human eye, the retina, and the human brain. The correspondence exists independently of human decision-making.

The different types of representation proposed above are similar to C.S. Peirce's tripartite division of representational 'signs' into 'icons', 'indices', and 'symbols'. Peirce held that icons resemble what they represent, indices are causally connected to what they represent, and symbols are arbitrary labels for what they represent, (see Schwartz 1995, p536-537).

¹³A word is four consecutive bytes.

¹⁴A transformation which changes only the scale factor.

Type-1 and type-2i representation both involve objective facts, but type-2ii representation does not. It is an objective fact about the wind-tunnel model that it is approximately homothetically isomorphic to the actual Formula One car. The relationship between brain states and the states of the external world, exists objectively because it is determined by an objective physical process. However, type-2ii representation does not involve objective facts because the correspondence is neither homomorphic, nor is it a causal correspondence.

If a thing is merely an element x in the domain of a representational mapping f which maps x to another thing $f(x) = y$, then x can be mapped to things with which it shares no characteristics. Typically, there exist other representational mappings $\{g_i\}$, for which x is again merely an element in their respective domains, and which map x to things $\{g_i(x) = z_i\}$ which either share none of the characteristics of y , or which possess characteristics mutually exclusive to those possessed by y .

In contrast, if a thing is the entire domain M of a representational mapping F , which maps M , at some level of homomorphy, to another thing, $F(M) = N$, then any other object to which M can be mapped at the same level of homomorphy, must also be homomorphic to N . If $G : M \rightarrow P$ is a homomorphic mapping, then $F \circ G^{-1} : P \rightarrow N$ must also be a homomorphic mapping. Despite this, M can still be represented by things with mutually exclusive characteristics. For example, a torus can be represented by a red coffee cup, and a torus can be represented by a blue coffee cup. A torus is isomorphic to both coffee cups at the level of a topological isomorphism, but a coffee cup cannot be both red and blue in colour. As another example, a triangle is homeomorphic to a circle and homeomorphic to a square, yet the problems of squaring a circle are well-known! The point, however, is that $F(M)$ and $G(M)$ must share some characteristics, the ones which are preserved by the homomorphic mappings. In contrast, $f(x)$ and $g_i(x)$ might share no characteristics at all.

Consider another example of type-2ii representation. The state of a light switch could be used to represent things other than itself. One could decide that the On-position of a light switch represents the number 1, and the Off-position represents the number 0. This relationship between the states of the light switch and the set $\{0, 1\}$ does not exist objectively. In other words, the relationship does not exist independently of the interpretative decisions made by human-beings. Someone else could decide that the On-position represents the number 0, and that the Off-position represents the number 1. One could even decide that the On-position of a light switch represents the colour black, and the Off-position represents the colour white. There is no homomorphism between the On-position of a light switch and either the number 1 or the colour black. The position of the light switch is merely being used as an element in the domain of a mapping which defines the representation. The state of the light switch shares no characteristics with either the number 1 or the colour black.

In the case of a digital computer simulation, the bytes of memory are used to represent numbers and numbers are used to represent the quantities of the simulated system. Hence, the representation of a tornado by the electronic states of a current digital computer is an example of type-2ii representation.

There is no homomorphism between the electronic states of a current digital computer and the things those states are chosen to represent. The electronic states of a computer can be mapped to many different things, but in each case an electronic state is merely an element in the domain of the mapping which defines the representation.¹⁵ The electronic state of a computer is not the domain of a homomorphic mapping, and human decisions, rather than causal processes, determine what things the electronic states of a digital computer represent. For these reasons, the states of a digital computer are not objectively related to that which they are deemed to represent.

Note also that the processes occurring within the CPU of a computer are not arithmetical or logical operations in any objective sense. The processes occurring within the CPU of a computer are only arithmetical and logical operations *under a specific interpretation*.

The electronic states of a current digital computer do possess quite intricate structure, but that structure is not used for the representational applications of a computer. The state of each bit in the memory of a computer is defined by the 1-dimensional graph topology of an electrical circuit, and by the voltage between specific points of the circuit. Hence, the memory-state of a computer is something with quite intricate structure. However, this electrical circuit and voltage structure does not resemble the things which the memory of a computer is chosen to represent.

Note carefully that the distinction between the types of representation does not entail that a wind-tunnel model objectively represents a Formula One car, nor does it entail that a brain state objectively represents the spatial geometry of an individual's visual field. There is no such thing as objective representation. Representation is dependent upon the interpretational decisions taken by thinking-beings. This is true for type-1, type-2i and type-2ii representation. Whether x represents anything at all, and what type of representation it is, is dependent upon the interpretational decisions taken by thinking beings. The wind-tunnel model and the Formula One car are objectively related, and brain states and visual fields are objectively related, but whether or not one *represents* the other is determined by the decisions taken by thinking beings. Type-1 and type-2i representation require objective relationships to exist. Resemblance is not a sufficient condition of type-1 representation, but it is a necessary condition. Causal connection is not a sufficient condition for type-2i representation, but it is a necessary condition. In contrast, type-2ii representation does not require an objective relationship of any type between x and what it represents, and a digital computer simulation is a type-2ii representation.

If a digital computer simulation of a universe is a type-2ii representation, then a digital computer simulation of a universe is not objectively related to

¹⁵To be strictly accurate, one should recall that there is a one-many correspondence between the logical states and the exact electronic states of computer memory. Although there are bijective mappings between numbers and the logical states of words in computer memory, there are no bijective mappings between numbers and the exact electronic states of memory.

that universe. This rules out the claim that a digital computer simulation could exist as a universe.

Although the states of a digital computer are not objectively related to the things they represent, it is possible that the states of an analog computer could be so related. It is conceivable that there could be an homomorphism between the states of an analog computer and the things those states represent. Whilst an analog computer does not necessarily resemble the system it represents in terms of geometry or topology, a homomorphism between physical objects is not necessarily a homomorphism of spatial geometry or topology. The examples of the wind-tunnel model and London Underground map are misleading in this respect. The homomorphism could be a non-visual homomorphism. An analog computer could possess objective physical properties which do change with the same pattern as the changing pattern of values for the physical quantities on a simulated system. Hence, an analog computer simulation might provide type-1 representation. If this is so, then a more general argument would be required to demonstrate that no type of computer simulation at all could exist as a universe.

In the examples of type-1 representation given above, although there is a physical resemblance in some respects between M and $f(M)$, there is not total resemblance. For example, although the parts of a wind-tunnel model subtend the same angles as the actual car, the wind-tunnel model is not the same size as the actual car. However, there is no reason why a type-1 representor cannot possess all the properties of the thing it represents. At least, there is no reason why a type-1 representor cannot possess all the ‘intrinsic’ properties of the thing it represents. The intrinsic properties of an object are the properties it possesses independently of its relationships to other objects. If one object is numerically distinct from another, as a representor must be from what it represents, then the two objects cannot share the same set of relationships with other objects. To be numerically distinct, they must occupy disjoint regions of space-time, and therefore cannot share the same set of spatio-temporal relationships with other objects.

If a type-1 representor possesses all the intrinsic properties of the thing it represents, then one might conclude that it exists as the same type of thing as the thing it represents. Accordingly, an analog computer simulation of a universe might exist as a universe. However, to reiterate, it remains to be proven that an analog computer can possess the representational capacity to represent an entire universe.

In type-2 representation, a representor need not exist as the thing it represents because a type-2 representor need not possess any of the properties of the thing it represents. In particular, a type-2ii representor can represent many things, with mutually exclusive properties. If a type-2ii representor were to always exist as that which it represents, it would exist as many different things, a contradiction. A digital computer simulation, even a perfect simulation, is a special case of type-2ii representation. A current digital computer simulation of a tornado exists as an electronic process the computer undergoes, not as a

tornado. Similarly, a current digital computer simulation of an entire universe exists as an electronic process the computer undergoes, not as a universe.

Recall that Tipler imagines a perfect computer simulation of our universe, which would simulate all the people who exist in our own universe. Such simulated people, suggests Tipler, would reflect upon the fact that they think, would interact with their apparent environment, and would conclude that they exist. The claim that a simulated universe would be real to the simulated people, presupposes that simulated people exist. Digital computer simulations of people exist only as physical processes on a computer, not as people. Hence, there are no people in a digital computer simulation to reflect upon the fact that they think, or to interact with their apparent environment.

If a perfect digital computer simulation of a universe cannot exist as a universe, then Tipler's sceptical hypothesis cannot be true. It is impossible that our own experience is indistinguishable from the experience of somebody embedded in a digital computer simulation because it is impossible for anybody to be embedded in a digital computer simulation. People cannot exist in digital computer simulations.

There is a similarity between the argument above, and one of John Searle's arguments against the claim of 'Strong' Artificial Intelligence (Strong AI), that minds are computer programs. Searle argues that the brain cannot be a computer in any objective sense because nothing can be a computer in an objective sense. Searle argues that a process cannot be a computational process in any objective sense. He claims that "A process is computational only relative to some observer or user who assigns a computational interpretation to it," (Searle 1995, p548). In terms of what exists and what happens independently of observers, Searle states that a computer "is an electronic circuit with state transitions between voltage levels," (ibid., p547). The devices we refer to as computers can only be said to undergo computational processes because that is the interpretation which the designers and users assign to them.

If the brain cannot be a computer in any objective sense, then the mind cannot be a computer program running on the brain, in any objective sense.

Analogously, one can argue that a computer simulation is only a simulation because it is deemed so by the designers and users of the simulation. Our universe cannot be a computer simulation (program) in any objective sense because nothing can be a computer simulation (program) in any objective sense. Nothing can be a computer simulation (program) independently of our beliefs about it. A simulation of a system does not exist as the simulated system in any objective sense; rather, it is the electronic states and processes of the computer which exist. As Searle points out, a computer can "simulate the formal features of any process," but "the fact that the programmer and the interpreter of the computer output use the symbols to stand for objects in the world is totally beyond the scope of the computer," (Searle 1982, p370).

The fact that a simulation does not exist in any objective sense follows from the fact that representation is not an objective relationship between objects.

However, as emphasised, objective relationships do exist between objects, and there may be objective relationships between analog computers and the physical systems they simulate which enables an analog computer simulation to exist as the type of system it represents.

The things represented in a digital computer simulation do not exist in the memory of the computer; nothing exists other than the processes occurring to the components of the computer.

References

- [1] Aho, A.V. and Ullman, J.D. (1992). *Foundations of Computer Science*, New York: W.H.Freeman.
- [2] Baez, J.C. (1995). Spin networks in nonperturbative quantum gravity. arXiv:gr-qc/9504036 v1 21 Apr 1995. Published in Louis Kauffman (ed.), *The Interface of Knots and Physics*, (1996). Providence: A.M.S., pp167-203.
- [3] Barrow, J.D., Tipler, F.J. (1986). *The Anthropic Cosmological Principle*, Oxford and NY: Oxford University Press.
- [4] Barrow, J.D. (1991). *Theories of Everything*. Oxford and NY: Oxford University Press.
- [5] Barrow, J.D. (1992). *Pi in the Sky*, Oxford: Clarendon Press.
- [6] Beem, J.K. and Ehrlich, P.E. (1981). *Global Lorentzian Geometry*, New York: Dekker.
- [7] Bekenstein, J.D. (2003). Information in the holographic universe, *Scientific American*, August, pp48-55.
- [8] Bekenstein, J.D. (2004). How does the entropy/information bound work? arXiv:quant-ph/0404042 v1 7 Apr 2004.
- [9] Besse, A.L. (1987). *Einstein manifolds*, Berlin and New York: Springer-Verlag.
- [10] Boothby, W.M. (1986). *An introduction to differentiable manifolds and Riemannian geometry*, 2nd Edition, London: Academic Press.
- [11] Cornish, N.J., Spergel, D.N., Starkman, G.D. (1998). Circles in the sky: Finding Topology with the Microwave Background Radiation, *Classical and Quantum Gravity*, Vol.15, No.9, September, pp2657-2670.
- [12] Dancy, J. (1985). *Introduction to Contemporary Epistemology*, Oxford: Blackwell.
- [13] Dancy, J., Sosa, E. (eds.) (1992). *A Companion to Epistemology*, Oxford: Blackwell.

- [14] Ellis, G.F.R. (1971). Topology and Cosmology, *Gen.Rel.Grav.* vol.2, No.1, pp7-21.
- [15] Friedman, J.L. (1991). Space-Time Topology and Quantum Gravity, pp539-569 in A.Ashtekar and J.Stachel (eds.), *Conceptual Problems of Quantum Gravity*. Boston: Birkhauser.
- [16] Geroch, R., Hartle, J.B. (1986). Computability and Physical Theories, *Foundations of Physics*, vol.16, No.6, pp533-550.
- [17] Luminet, J-P., Weeks, J.R., Riazuelo, A., Lehoucq, R., and Uzan, J-P. (2003). *Dodecahedral space topology as an explanation for weak wide-angle temperature correlations in the cosmic microwave background*, arXiv:astro-ph/0310253 v1 9 Oct 2003
- [18] O’Neill, B. (1983). *Semi-riemannian geometry: with applications to relativity*, London: Academic Press.
- [19] Mermin, N.David (2002). From Cbits to Qbits: Teaching Computer Scientists Quantum Mechanics. arXiv:quant-ph/0207118 v1 19 Jul 2002.
- [20] Schwartz, R. (1995). Representation, pp536-541 in S.Guttenplan (ed.), *A Companion to the Philosophy of Mind*, Oxford: Blackwell.
- [21] Schmidhuber, J. (1997). A computer scientist’s view of life, the universe, and everything, pp201-208 in C. Freksa (ed.), *Foundations of Computer Science: Potential-Theory-Cognition, Lecture Notes in Computer Science*. Berlin: Springer.
- [22] Searle, J.R. (1995). Searle, John R., pp544-550 in S.Guttenplan (ed.), *A Companion to the Philosophy of Mind*, Oxford: Blackwell.
- [23] Searle, J.R. (1981). Minds, Brains, and Programs, pp353-373 in Douglas R.Hofstadter and Daniel C.Dennett (eds.), *The Mind’s Eye*. London: Penguin books.
- [24] Siegelmann, H. (1999). *Neural Networks and Analog Computation*. Boston: Birkhauser.
- [25] Smith, B. (1995). Ontology, pp373-374 in J.Kim and E.Sosa (eds.), *A Companion to Metaphysics*, Oxford: Blackwell.
- [26] Stillwell, J. (1993). *Classical Topology and Combinatorial Group Theory*, 2nd Edition, New York: Springer Verlag.
- [27] Tegmark, M. (1998). Is ‘the theory of everything’ merely the ultimate ensemble theory? arXiv:gr-qc/9704009 v2 1 Dec 1998. Published in *Annals of Physics*, 270, pp1-51, (1998).
- [28] Thurston, W.P. and Weeks, J.R. (1984). The Mathematics of Three-dimensional Manifolds, *Scientific American*, July 1984, pp94-106.

- [29] Tipler, F.J. (1989). The Omega Point as Eschaton: Answers to Pannenberg's Questions for Scientists, *Zygon* 24, pp217-253.
- [30] Tipler, F.J. (1995). *The Physics of Immortality*, London: Macmillan.
- [31] Wolf, J.A. (1967). *Spaces of Constant Curvature*, London: McGraw-Hill.