Close this window

A NEW COSMOGONY

EDWARD FREDKIN

Department of Physics Boston University

Boston, MA, 02215, USA

Abstract

Digital Mechanics is a model of physics based on the Finite Nature assumption; that at some scale, space and time and all other quantities of physics are discrete. In this paper we will assume that Finite Nature is true and we will explore the consequences with regard to the nature and origin of the Universe. Contemporary physics has a lot to say about models of the early universe; down to the first tiny fraction of a second after the Big Bang. Digital Mechanics can tell us something about what might have occurred before the Big Bang. We show that any reasonable estimate for the unit of length leads to a puzzle; the computational capacity of space dwarfs any reasonable requirement for what we know about physics. This paper will attempt to lead the reader down a connected path of consequences that all follow from the single assumption; Finite Nature.

1992 IEEE. Used with permission. This article is scheduled to appear in the Proceedings of the Physics of Computation Workshop, October 2-4, 1992.

Permission to copy without fee all of part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the IEEE copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Institute of Electrical and Electronic Engineers. To copy otherwise, or to republish, requires a fee and specific permission.

Introduction

The greatest mystery is "Why is there anything at all?" This mystery is tied to the great cosmogonical question, "Where did the Universe come from?" These questions raise one's curiosity about other things such as "If the Universe came from something or from somewhere, then where is that something or that somewhere and what are things like there?" "How did that place come into being?" and on and on. These are subjects that have been dealt with in mythology and religion. We presume to use science to look for plausible answers to these questions. If we assume that Finite Nature is true, we discover that surprising progress can be made in looking beyond our own world.

Finite Nature

Finite Nature[1] is the hypothesis that ultimately every quantity of physics, including space and time, will turn out to be discrete and finite; that the amount of information in any small volume of space-time will be finite and equal to one of a small number of possibilities. We call models of physics that assume Finite Nature

"Digital Mechanics". In Digital Mechanics[2], the basic element of physics *is* the cell and the rules. Things like particles are the consequence of stable patterns over a volume of space-time. The Finite Nature hypothesis makes no assumption about the scale of the quantization of space and time. Digital Mechanics is too immature a concept to say more about the scale of length other than it is probably around Planck's length, 1.6x10^-33cm. Some considerations make it hard to understand why the unit of length should be less than about 10^-17cm. The question can be settled by experimentally determining the value of the scale of length.

We simply do not yet know whether Finite Nature is true or false. Today, nearly every scientist in the world believes that there is insufficient experimental evidence in hand to decide the issue in favor of Finite Nature. The author, on the other hand, has managed to convince himself that the odds are greatly in favor of the Finite Nature Hypothesis. What has been decided up to now is that many things of our world that were once thought of as possibly continuous are now known to be discrete. The most famous is the *atomic hypothesis*. Dalton[3] wrote his papers in the early 19th century but as recently as 1900, a famous physicist (Ernst Mach) said that while there was evidence for the atomic theory, since no one had seen an atom and since no would ever see an atom, he was not convinced that the atomic theory was true. Times have changed and now we can see atoms with the scanning tunneling microscope. Now, we all ardently believe in the atomic theory.

The next to fall into the realm of the discrete was electricity. Originally thought of as a fluid, Thompson discovered the electron in 1897 and with it came the discovery that charge was a discrete or quantized phenomena. Einstein proposed that Planck's quanta of action could determine the relation between the energy and frequency of particles of light, which he called photons. Planck thought that Einstein was a very smart person in spite of Einstein's belief that light was made up of discrete particles! Today we all believe that photons are real and that light, electro-magnetic and other kinds of forces are made up of discrete particles.

As the consequences of the Quantum Theory became better understood, it became clear that the angular momentum of particles can only exist in multiples of +1/2 units of spin. This has the amazing consequence that a flywheel cannot have a continuous range of angular momenta, rather it must only have multiples of +-1/2h. Angular momentum is now known to be discrete. The story goes on with phonons and vibrons as quantized units of sound and other forms of energy.

So far, there is no convincing argument based on experimental evidence that points to any quantity of physics as definitely continuous. What we can often say is "If it is discrete, then the quantization must take place below some level." It is difficult to even propose a test that could verify that some quantity of physics was indeed continuous.

Since we know of no verified continuous quantity in physics, and since there has been a steady historical progression of finding that more and more of the fundamental quantities of physics are discrete, it is perfectly reasonable to assume the possibility that all quantities of physics will prove to be discrete. What we shall reveal is the amazing consequences of such an assumption; consequences that are independent of the scale of the quantization!

The cellular automaton

In this paper, we take the position that Finite Nature implies that the basic substrate of physics operates in a manner similar to the workings of certain specialized computers called cellular automata[4][5][6]. The logic behind this assumption is that a cellular automaton is a system of cells where each cell is in one of a finite number of states. Each cell transitions from one state to the next according to a rule where the outcome for a particular cell depends on the states of cells in the neighborhood of that cell. The definition of cellular automata seems so broad as to encompass every kind of discrete cellular process.

Some cellular automata are universal; they have the property that they can be made to do any computation for

which they have enough volume and time. If Finite Nature is true it means that a volume of space has a certain amount of computational capability. If Finite Nature is not true, then it seems necessary to assume that infinite computational resources are required to model physics exactly. Conversely it is reasonable to equate the order of computational power of any system with the order of computational power necessary to exactly model that system. If Finite Nature is false, then any volume of space time, no matter how small, would represent an infinite capacity for computation. The author believes that computational capability is a quantitative resource, like area, or energy. It should be possible to relate the physical units of computation to ordinary physical units (e.g. Mass, Length, Time...). This would make it untenable for a finite volume of space time to require an infinite amount of computation in order to model physics exactly; for example, if it requires infinite resources that are equivalent to mass.

In this paper we will be working with large finite numbers, but any kind of infinity dwarfs them all. It is very hard to imagine what the purpose or necessity could be for any sort of infinity, since simple to express finite resources can clearly dwarf the needs of a universe like the one we live in. Consider the Ackerman function:

A(a, n, j);
A(a, n, 1) = a + n
A(a, n ,2) = a * n,
A(a, n, 3) = a^n,
A(a, n, 4) = a^a^a^a (n times)

In Mathematica, A can be defined as follows:

```
A[a_, n_, j_]:= If [ j==1, a + n, If [ n=1, a, A[a, A[a, n-l, j], j-1] ] ]
```

Assume a = A(9, 9, 9). C is a large but finite number that easily exceeds the value and the precision of any number that might be encountered in any calculation about our universe. General recursion allows us to define much larger finite numbers as necessary. Nothing to do with anything we can learn from experimental physics will ever have any need whatsoever for such large or precise numbers. The ideas of infinite and infinitesimal numbers and of continuous variables are handy for allowing the use of calculus and for creating approximate physical formulas, but we shouldn't confuse those ideas with what is happening in the real world.

The Assumption: Finite Nature

If we could look into a tiny region of space with a magic microscope, so that we could see what was there at the very bottom of the scale of length, we would find something like a seething bed of random appearing activity. This would not be locally generated randomness. While the state of a particular bit is the immediate, deterministic function of its neighborhood, it is also functionally related to approximately 1/2 of all of the bits in the entire space-time history of the entire universe. This means that if one could magically reach back into the past and change any one of those functionally related bits, it would change the state of that particular bit. Because quantum mechanics limits our knowledge of the exact state of microscopic events, the state of such a bit can be perfectly random with regard to any kind of test that can be administered from within the system.

Space would be divided into cells and at each instant of time each cell would be in one of a few states. A snapshot would reveal patterns of two (or three or four or some other small integer) kinds of distinguishable states. It would be either pluses and minuses, blacks and whites, seven shades of gray, ups and downs, pluses and neutrals and minuses, clockwises and anticlockwises or whatever. The point is that it would be equivalent to digits. If every cell was either a black or a white, then we could rename them "1" and "0" or "+" and "-". It wouldn't matter.

What we would discover is that there is a rule that governs the behavior of the cells. It is logical to suspect that the state of each cell is some kind of function of a neighborhood; for each cell, a set of neighbor cells with some particular space-time relationship to the cell. We don't yet know what the rule is, or even the exact nature of the rule, but we know many kinds of rules it could be. The fact that each cell is like a digit and that the overall behavior is a consequence of a rule where the next state of each cell depends on some function of the neighborhood cells is why the underlying mechanism must be some kind of cellular automaton.

The meaning of the digits in Finite Nature is that the information process that the digits and cells are engaged in define the space and time and matter and energy of our world. It is important to understand that while the cellular space may normally resemble the space of physics, the two kinds of space would become very different under certain conditions such as near a black hole. The same is true for the passage of time. An ordinary clock might, if unaccelerated and far from a strong gravitational field, keep reasonable time with respect to the goings on in the cells, However, strong gravitational fields or velocities near the speed of light would cause an ordinary clock to slow down with respect to the events in the cellular space.

How did the Universe get started?

This is the greatest of puzzles because it seems impossible to find an answer. We are not asking "Did the Big Bang Happen?" but rather "What caused the creation of this universe?" The question seems at odds with both science and common sense. Common sense tells us that something can't come from nothing. Science proclaims that the quantity of Mass-Energy is conserved or always the same and unchanging. Science has a fancy way of saying the same thing as common sense. Common sense says that the story of the Universe has to have a beginning. Science says that the Universe started with the Big Bang. Once again Science and common sense agree. Common sense tells us that there is something here; "Cogito ergo sum". A reasonable estimate of the mass of the Universe is something like 10^53kg (the number of tons of stuff is 1 followed by 50 zeroes). So far science and common sense are in complete agreement. The rub is: if the Universe had a beginning then before the beginning the Universe wasn't here; the Universe can't come from nothing and yet the Universe is here.

It just doesn't hang together. Either we have to believe the Universe has been here forever or some kind of magic took place in the beginning. In a sense, physics lets us down when we try to peer at what was it that created the Universe so that the Big Bang could happen. Nevertheless, Finite Nature gives us a way to look beyond the Big Bang. It gives us the possibility of thinking about how the Universe got here and even about why the Universe is here. Additionally it gives us an idea of how the laws of physics were created, and why the laws have certain characteristics.

The answer lies in the amazing consequence of the simple assumption of Finite Nature. As we have explained, Finite Nature means that what underlies physics is essentially a computer. Not the kind of computer that students use to do their homework on, but a close cousin; a cellular automaton. Not knowing the details of that computer doesn't matter because a great and tragic British mathematician, Alan Turing proved that we don't need to know the details!

What Turing did in the 1930s was to invent the Turing Machine. It was a way to formalize all the things that a mathematician could do with pencil and paper. The result proves that any ordinary computer, given the proper program and enough memory, can do what any other computer can do. It can also do what any mathematician can do; if we only knew how to write the program! Finite Nature implies that the process underlying physics is a kind of computer, therefore it is subject to Turing's proof. This means that there is not just one kind of underlying computer, but there are many possible equivalent computers. Of course some are simpler, some are more elegant, some use the least amount of various resources, some are faster... This means that once we have figured out that it's a computer at the bottom, we already know a lot even if we don't know what kind of computer would be most efficient at the task.

Where is the Ultimate Computer?

When my son Richard was a precocious 7 year old, we were walking back to our car from an outing. We had parked some blocks away and suddenly we were uncertain as to where the car was. Richard piped up "I know exactly where the car is!" I was surprised and impressed, "Where is it?" I asked. "The car is in the Universe." he answered. As to where the Ultimate Computer is, we can give an equally precise answer, it is not in the Universe - it is in an *other* place. If space and time and matter and energy are all a consequence of the informational process running on the Ultimate Computer then everything in our universe is represented by that informational process. The place where the computer is, the engine thai runs that process, we choose to call "*Other*".

Where did *Other* come from? This question is actually quite easy to fence with. The nature of systems of laws that can support computation is very much broader than the nature of systems that are limited to the physics of our universe. In other words, many of the properties of our world that are necessary for our world to take the form it has are not necessary for other kinds of worlds that can support universal computation. Universal computation, the kind that can simulate other general purpose computers, is even a property of all ordinary commercial computers.

There is no need for a space with three dimensions; computation can do just fine in spaces of any number of dimensions! The space does not have to be locally connected like our world is. Computation does not require conservation laws or symmetries. A world that supports computation does not have to have time as we know it, there is no need for beginnings and endings. Computation is compatible with worlds where something can come from nothing, where resources are finite, infinite or variable. It is clear that computation can exist in almost every kind of world that we can imagine, except for worlds that are sterile or static at every level. Universal computation is essentially synonymous with the idea that something interesting is going on. Worlds where nothing interesting will ever happen are the only kinds of places that cannot support universal computation. What is certain is that worlds that are qualitatively beyond our power to imagine are also capable of supporting computation. What all this means is that the questions that puzzle us about the origin of our universe are most unlikely to apply to *Other*. It's not that the problem has been put off or postponed, the problem of the origin of *Other* is tautologically null. *Other* is that place that has such structure and laws as to not raise the question of its origins, as origins are a concept peculiar to our world.

What Can We Know About *Other?*

Surprisingly there is a great deal that we can deduce about the probable characteristics of *Other*. This can be done by a number of different approaches.

1. We need to experimentally determine the unit of length and the size of our universe. This allows us to measure the power of the computational engine.
2. We need to carefully quantify the apparent excess computational power in this universe in order to estimate how much could be going on that we are currently unaware of.
3. We need to look at the resources that are required for the engine, as a measure of the resources dedicated to our universe in Other.
4. By quantifying the amount of computational resources devoted to our universe we can eliminate many problems as too small or too large. This can enable us to construct a list of possible reasons for the existence of this universe.
5. Given a proposed purpose for our universe, we can subject it to the following tests: It should just require all of the computational resources of our universe - no more and no less.
6. Given the problem, what can be said about the answer? It might be as complex as the entire history of the Universe, or as simple as one bit: "yes" or "no". For example "Will the expansion of the Universe stop?" Finally, the answer may be forever beyond our knowledge or beyond our understanding.

7. Is this universe perhaps an artifact of some larger process?
8. Is physics at the bottom, or is what we know of as physics merely an artifact of a different and more fundamental process?
9. We can consider the laws of physics in this universe and the implications on Other.
10. We need to consider the grand design of this universe in the context of alternative universes.

Other intelligence

If the Universe serves a purpose, does that mean that there is an intelligent creator? First of all, we must get used to the idea that in this context, humanlike intelligence is not relevant. Humans are the prototype thinking creatures on our planet. As such, our brains are very close in design and capability to the brains of other primates such as chimpanzees or gorillas. This is not meant to be demeaning to humans but simply a statement of the relative position of human intelligence; what else would it be close to? As the prototype thinking creatures we do our intellectual work with mechanisms that were really designed (by evolution) for the tasks of surviving in jungles and forests. We haven't had scientists around long enough for brains to have evolved that are better at doing science; besides, what could the selection process be? We can already see the possibilities of artificial intelligence appearing on Earth that will greatly exceed human intelligence by most objective measures of capability.

A heavy jet airliner weighs about 300,000 kg. We can imagine a cube of silicon, of about the same weight, made into a computer that contains 6x10^23 computing-memory elements operatic at cycle times of 10^-12seconds. That is 10 atoms per computing-memory element. One such machine could out-compute the combined powers of all of the computers built so far and out-think the combined thinking powers of allof the people who have ever lived.

This shouldn't frighten us, since no such machine will be better at being human that we are, it will just out-think us. Today we have machines that exceed humans in every physical measure; they are faster, more accurate, stronger, can fly, sail across oceans, explore the surface of Venus, look into distant galaxies, peer at atoms, etc. We are not intimidated by the physical prowess of machines, and we still engage in human sporting events, such as the marathon despite the fact that a motorcycle could do so much better. Similarly we shouldn't worry about machines doing arithmetic billions of times faster than humans, and eventually doing every kind of intellectual activity faster and more accurately than humans. Such machines will be no better at being human that a human can be better at being a mouse. Mice would rather live, work and play in the company of mice as opposed to humans.

The point is that machines with Artificial Intelligence (AI) that we may build in the future will be qualitatively different in intellectual capability than are humans. We can't speak of AI as "intelligent" in human terms. Something that may have created our universe in a purposeful fashion is very different than something with human-like intelligence or the kind of intelligence we will see in Earthbound AI. There might be no more in common than the concept of questions and the concept of finding answers.

When primitive man gained control over fire, it was a great step. However, the Universe is full of fire; consider the stars. Because we can think, and nothing else on our planet seems to be able to do so, it is natural for us to hold our intellectual prowess in great esteem. However, it may be that information processing, instead of being the sole province of us humans and our machines, may be a part of almost everything else in physics. Life itself, is clearly mediated by digital information; the genetic code. Digital Mechanics assumes that physics is also a process based on informational processes. We may need to rid ourselves of the prejudice that purposeful, thought related action is exclusively the domain of humans or perhaps aliens similar to humans. There are kinds of thinking that are qualitatively unimaginable to us though we can think about it quantitatively. We should not be afraid to consider intellectual activity as the driving force behind the creation of the Universe. By a close and quantitative examination of the possible parameters of Digital Mechanics, we

can arrive at reasonable guesses as to what might be the purpose behind the creation of a universe like ours. That, in turn, can lead us towards intelligent speculations about *Other,* the space that contains the engine of our world.

What's happening in *Other*

If we assume that the Ultimate Computer was purposefully constructed in *Other,* we can immediately answer the puzzle of the origin of the Universe. It's simply a matter of the following process taking place in *Other.* The initial conditions are set into the engine and the engine is set into motion; it starts to compute. Those two steps are outside the domain of physics. If we write a program to simulate the operation of the Space Shuttle, we don't have to worry about conservation of mass if we teil a programmer to change the amount of fuel in the simulated shuttle. It's just a number that is set by a programmer. However once the simulation starts, if it is programmed properly, then what happens during the simulation will be in accord with the laws of physics.

If the purpose of *Other* is to find something like an answer to something like a question, we still have the possibilities that:

1. What we see as our universe might be working towards that answer.
2. The Universe we know, in its entirety, might be an artifact.

In either case, our existence here on Earth might or might not be completely incidental to the purpose.

We can also ask: "If something in *Other* was so smart and capable as to be able to create our universe, why didn't it just think and figure out the answer in its head?" (Please pardon the anthropomorphic allusions.) One of the interesting results of computer science, that transcends the laws of physics, is that an answer obtained by running a computer for a certain number of steps, cannot in general be found by some shortcut. This is a consequence of Turing's famous halting problem. The name "halting problem" comes from the old idea that a computer at work should halt when it gets the answer. The question is, can some other computer look at what the first computer is trying to do and figure out a way to get the answer in fewer steps? Of course, if the first computer is inefficiently implemented, then some more efficient computer could speed things up. If the program is coded so that it runs inefficiently, then reprogramming it could speed things up. However, there is no way, in general to take programs and to reprogram them so that they now run faster. This limitation known as the Speed Up Theorem.

What the Speed Up Theorem tells us is that if Finite Nature is true, and if the thing in *Other* was competent, then there is no way for it to get the *answer* any faster than by letting the Universe run its course. Interestingly, answers that do not require computation are just those &at have analytic solutions. We know our universe is universal (in the computer sense) simply because we can build computers; this would not be possible in a universe that was not universal.

As to the question "Why didn't the thing in *Other* just do it in its head?" The answer is quite straight forward: Doing it on a computer is exactly the same thing as doing it in one's head. Both are examples of using an informational process to get to the answer. We are not referring to the thing in *Other* finding an analytical solution in its head (the speedup theorem forbids such solutions) but rather to it imagining each step of some cellular automaton in its head. Strangely enough, that's exactly the same as doing it on a computer. A common fable is that maybe God is dreaming all this, that we are characters in that dream, and that we should be careful not to do anything that might wake him up!

The problem of the missing workload

The enormous computational capability of space-time is totally squandered by what we see and know. If we

assume that the radius of the Universe is 20 billion light years, then the number of cells (at Planck's Distance) is about 10^184. Insofar as the computational model, almost nothing is being done. Most of space is empty, except for a little radiation (photons) passing by. The space that is not empty, such as the space occupied by an atom, is also mostly nothing. Given the likely computational resources of space, it is like coming upon the worlds largest super computer and finding it spending years on the task of doing nothing but trying to count from 1 to 10 just one time.

There are many other ways in which we can see that what we think of as the Universe may be totally incidental to the main purpose. For example, let's assume our universe is a sphere 20x10^9 light years in radius, so that its volume is about 3x10^124 cubic fermis. If we make the generous assumption that we will give an automaton 10^12 cells per cubic fermi to simulate the essential physics that the operation of our macroscopic universe seems to require, the simulation will need a total of 3x10^136 cells. If the unit of length in our world is at Planck's length, 1.6x10^-33cm, then the volume of space needed for 3x10^136 cells at Planck's length would be a sphere 30,000,000 kilometers in radius (the size of a large star). In 100 seconds, the sphere could faithfully simulate the entire macroscopic evolution of our universe from the big bang to the present. The space-time volume of the Universe is 10^63 times larger than that sphere for 100 seconds. It seems that a tiny fraction of the Universe may have all of the computational resources to run all of the essential macroscopic physics we know of, from the big bang to the present in a couple of minutes.

We can think of this as the case of the missing workload. Unlike the missing neutrinos from the sun (perhaps a factor of 3) and unlike the missing mass from the heavens (perhaps a factor of 50), the amount of missing workload seems to be a factor of 10^63 times the workload that appears necessary. The idea that we will find quantum mechanics to be continuously interesting as we look at finer and finer detail seems unlikely. The degree of quantization of our current understanding of the most microscopic things we know, does not allow for finding useful quantum mechanical gears and wheels at ever more microscopic levels. We certainly might find something else that's interesting and that lies under quantum mechanics. It is also obvious that the physics we have discovered in the last 20 years seems dimly related to what most of us would consider to be the main macroscopic workings of the Universe.

Either something else is going on (other than what we know about), the important action is what's happening down at Planck's length or something more subtle is running things not too far below a fermi. The only remaining alternative is to admit that God was incompetent on a scale that boggles the mind.

---

[1] E. Fredkin, "Finite Nature" Proceedings of the XXVIIth Rencontre de Moriond (1992)

[2] E. Fredkin, "Digital Mechanics", Physica D, (1990) 254-270 North Holland

[3] J. Dalton, "New System of Chemical Philosophy" (part I, 1808; part II, 1810)

[4] S. Ulam, Random Processes and Transformations, Proceedings of the International Congress on Mathematics, 1950, Vol. 2 (1952) 264-275

[5] J. von Neumann, Theory of Self-Reproducing Automata, edited and completed by A. Burks (University of Illinois Press, Champaign, IL, 1966).

[6] T. Toffoli and N Margolus, Invertible Cellular Automata: A Review, Physica D 45 (1990) 229-253

Close this window