

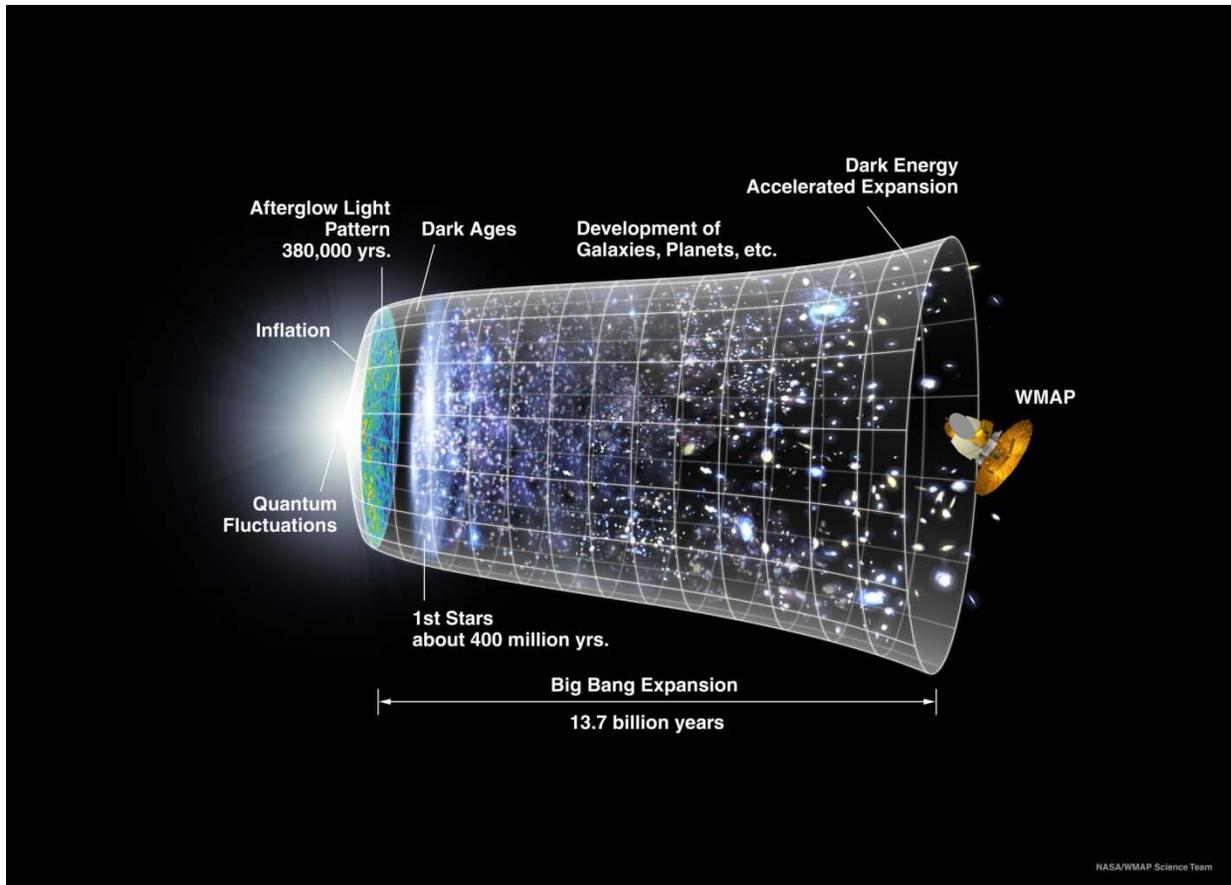
When Pierre de Fermat noted that he had discovered a ‘truly marvellous proof’ of an important conjecture, mathematicians looked forward to holding an elegant and succinct proof of the type that had once characterized their craft. What they got was a gargantuan proof constructed from newly invented techniques that initially took three days to present to an elite group of specialists who had the necessary competence to understand something so difficult. Most of us soon realized we were probably never going to grasp what Andrew Wiles and Richard Taylor had achieved.



Why is there something rather than nothing? Martin Heidegger considered this to be the most important question in philosophy. Symbolically, ‘nothing’ can be represented by ‘zero’, while ‘something’, in a very fundamental sense, can be represented by ‘unity’. Using just two symbols, ‘0’ and ‘1’, we can represent every possible number, and it is from this potential that we construct mathematics.

Numbers were originally invented to count things that exist (and are finite in number). They have since been used to count things which don’t exist, such as numbers themselves, leading to a recursive (and endless) hierarchy of infinities. If there were merely nothing rather than something, we would only have to account for ‘0’, which would be the end of the story (despite there being no one around to discuss it). It is ‘1’ that has got us thinking.

According to prevailing cosmological theory, our universe expanded exponentially from a microscopic region of the primordial quantum vacuum, an entity which has produced many other universes having any one of approximately 10^{500} different configurations; that’s quite a few gourmet jellybean recipes. Some of these configurations are viable, and go on to produce universes that contain incredible intellects like our own. Others go ‘poof’ and vanish only moments after they appear. Clearly the quantum vacuum is responsible for the production of a whole lot of stuff. The following diagram depicting the evolution of our particular universe is freely reproducible, courtesy of Goddard and Princeton, whose WMAP program was able to look back almost as far as the beginning.



The quantum vacuum has a somewhat circular heritage. The study of patterns in the behaviour of the physical world led to our invention of mathematics with which to model that behaviour, the most fundamental being general relativity and the wave function. Our extrapolation of these models back in time to a point before the universe came into existence, in turn gave them a life of their own (as the quantum vacuum), that is quite independent of this or any other universe. To put this conundrum simply, if you've already got yourself 'something', like three green apples, then you have a 'substrate' upon which to build the abstract notion of the number '3'. But is it meaningful (or possible) to speak of the number '3' before a universe, with at least three things in it to count, has come into existence?

Looking at the picture above, it is certainly meaningless to ask what the universe is expanding 'into'. It is space itself that is expanding, such that every point in space becomes the effective centre of this expansion. In the classical general relativity model, space is a continuum, a single entity that is expanding like the surface of an indefinitely inflating balloon. Loop quantum gravity theory, in contrast, suggests that space is not a single entity, like the thinning wall of a balloon, but is instead comprised of individual 'atoms' of space. A simple merging of general relativity and quantum theory suggests that these putative quanta of space (assuming they are spherical) have a diameter of about 10^{-35} metre. We cannot yet probe down to anywhere near this scale of length, and so our empirical tests of general relativity merely confirm that the assumption of a continuum remains a useful model at the macroscopic scales to which we have observational access (we have so far got down to about 10^{-18} metre).

If space is comprised of space 'atoms', and space is indeed expanding, then either the space atoms are expanding, as if each one of them were a classical continuum, or the volume of the space atom remains a constant of nature, and instead it is their number that is increasing. The latter can account for the isotropy (sameness in all directions) of our *observable* universe.

How do you create just one atom of space, let alone a whole universe full of them? Just that bit of our own universe which we can see would contain about 10^{185} of the blighters.

In modern computing, we increasingly rely upon 'virtualization', which is the practical realization of a principle that was identified early in the development of computing science. A real computer (such as an ordinary PC) can run a programme that precisely emulates (in software) the logic of its own hardware. The host PC can often run several instances of these 'virtual machines' or VMs, as they are known. Because the VMs are logically equivalent to the real PC, each VM can run any operating system, and any application, that can be run on the real PC. A VM can of course act as the host for yet another VM.

If we want to create a space 'atom' where none exists, we face the perennial problem of having to retrieve the raw ingredients from a storehouse that isn't even there. We can however employ the sort of device that certain climate scientists were reported to have described as a 'trick'. In scientific circles, a trick is never something dishonest, and often something of tremendous utility. Just suppose, for the sake of argument, that we have at our disposal a real computer, back there at the beginning of our universe. On that real computer, we execute a programme that *simulates* this same real computer. Then, on the simulated computer, we execute a programme that once again simulates the real computer. Finally, and this particular trick is known as a 'strange loop', we simply take away the real computer, and directly substitute the identical second simulated computer in its place.

Such a ploy is not possible in the physical world, because of well understood limitations imposed by entropy and information theory. Curiously, however, these machines have yet to enter into the physical universe and its jurisdiction, and the physical computer that was being used as a conceptual prop never actually existed. These simulated computers do not, of course, have anything like the complexity of an everyday PC. Rather, they are extremely simple and ideal universal machines, each comprised entirely of software – short strings of binary digits. In the first instance, the only capacity they require is that to simulate one another. Necessarily, one machine will parse through its string of digits, in the process of simulating the other. Once that first machine halts, the second machine then parses its (identical) string of digits, in the process of simulating the first, until it likewise halts, and the process repeats itself.

It is natural to think here that the strings of binary digits have become the persistent 'substrate' of these computers; that the strings are somehow streams of real '0's and '1's that feed back and forth into one another within some sort of platonic world. In fact, each is entirely *virtualized*. Each machine only enters into fleeting existence, and then merely as an abstraction, after it has been *simulated* by the other. After one machine has finished simulating the other, it vanishes entirely until it next comes to be simulated. There is no 'hardware' here; no persisting substrate, no length, and no mass. Such is the quantum vacuum, which as its name suggests, is about the *absence* of space.

The central idea, in the various string theories that orbit the mysterious body of M-theory, is that all particles (including those that mediate forces) are made up of vibrating 'strings'. Vibration implies periodicity, and in turn the most essential dimension of any fully functional physical universe. The pair of virtual computers described above represents, of course, a *clock*. Each computer parses its string of binary digits, from start to finish, in a finite period of time. Once again, the simple merging of general relativity and quantum theory indicates that this quantum of time is approximately 10^{-43} second. The value of this quantum derives from the absolute time it takes

this virtual computer to parse one of its virtual binary digits. There is no deeper level of abstraction providing this clock. Rather, it is fundamental that one computer – tick – computes the other computer – tock – and that it is this repeating cycle which *creates* the basic quantum of time. All we can surmise is that these computers execute their simulation of each other approximately 10^{43} times within the passing of what we perceive as one second in time. It is this fine granularity that gives our everyday experience of time its smoothness.

Once this first pair of universal computing machines has pulled itself up by its own bootstraps, and the clock of time has commenced, all manner of creativity becomes possible. In addition to simulating each other, these computers can execute functional applications (in the way a PC would execute a word processor), and the most basic of these applications is the simulation of space. Time and space thus become intimately and permanently related from the outset in the ‘time-space’ atom. In addition to any *hidden* variables, the processing of a *timespace* stipulates three visible dimensions of space, those that we detect empirically in the familiar geometry of reality. This simulated instance of space appears, then vanishes, only to appear again, according to which one of the pair of virtual computers happens to be extant and hosting the simulation.

The next fundamental application that a timespace can execute is the replication, or cloning, of its own computational routines. The first timespace produces a clone of itself. These two timespaces likewise reproduce, resulting in four timespaces, and so on. If it transpires that the timespace replication code can execute within one clock cycle, then after just one second, some $2^{10^{43}}$ timespaces will be produced. This results in a *very* rapid inflation in volume at the outset of the universe. The newly created timespaces are not being produced from a central point. Rather, each individual timespace becomes the centre of its own contribution to the ongoing doubling in the volume of space.

Each timespace unit can then proceed to host any of the material applications that are fundamental to the universe they engender, and this processing commences within several seconds of the start. As a universe matures, these fundamental applications of energy interact with each other, sometimes visibly. We proceed away from foundational issues and into familiar issues of cosmic engineering; condensation, transparency, accretion, synthesis, dispersal, conglomeration, geology, biology, consciousness, art, computation, the number Ω . Interestingly, *Omega*, which is the probability that any randomly selected computation taken from the ensemble of all possible computations will *finish* (and a number between ‘0’ and ‘1’), is algorithmically *incompressible*. This implies that the computation we find ourselves ‘inside’, the computation being executed by all the timespaces that comprise our universe, is also our most extensive calculation of the digits of Omega. It is algorithmic information theory, from which Omega arose, that encourages physicists to come up with simple explanations for complex phenomena. For if a theory is just as complex as the phenomenon it describes, the theorist may as well take up stamp collecting. Luckily for physicists, the conclusion to this universal computation will not be the end of world – the computation will simply output a completed theory of quantum gravity.

▣

The timespaces that comprise reality, were they spherical, could be arranged like stacked oranges in a market stall. However, the latest configuration of Euclidean quantum gravity – causal dynamical triangulations – approximates the timespaces as if they were tetrahedral pyramids conjoined into a ‘mosaic’. By insisting that timespaces only process time in the forward direction, the four familiar dimensions of reality emerge naturally from this analysis,

where earlier analyses that allowed time to move backwards as well as forwards produced an infinite (and unrealistic) number of dimensions.

The idea that reality is a mosaic (or lattice) of timespaces is particularly elegant, because the limiting speed of light is an emergent property of the configuration of the lattice, rather than an arbitrary empirical condition. Each timespace can be thought of as a cellular automaton, a machine that is able to store a particular computational state, *and* pass that state on to its neighbours, through an 'interface' that connects adjacent timespaces together. As with people, the 'mouth' of one timespace can pass on a message into the 'ear' of the timespace next to it, and so on.

To see how this works, we first carefully line up 10^{43} timespaces in a straight line, and connect them all together like a daisy chain. Because each timespace is about 10^{-35} metre wide, this line will stretch out for about 3×10^8 metres. Let's say the 'message' is a photon of light. A considerable conglomeration of timespaces actually participate in the definition of a photon, but for now let's assume that only one timespace is required. The photon message is passed on from one timespace to its neighbour in the course of one 'tick' of each timespace clock (10^{-43} second). This bucket brigade continues on down the line, each timespace passing on the photon message to its neighbour. After the patient participation and cooperation of some 10^{43} timespace individuals, the photon 'message' is finally delivered one second later to the other end of the line, some three hundred thousand kilometres away.

It requires *at least* one clock cycle to pass on a message (computational state) from one timespace to the next, and so this sets the limit for conventional propagation of a signal through the lattice. Various computational states (applications such as *mass*) can of course be transmitted through the lattice at much more leisurely rates, but *everything* apart from space and time is in motion *relative* to the timespace lattice.

Galaxies can be thought of as extremely stable islands of timespace latticework – absolute reference frames. However, timespace replication continues in the regions between galaxies, acting to push them apart (or indeed towards each other). Galaxies are like tectonic plates on the surface of the Earth being pushed apart by the volcanic upsurges at the plate boundaries.

The universal timespace lattice can be compared to a 'gas' with a 'temperature'. In the intergalactic regions, where timespaces are still replicating, the lattice is hot and turbulent, and the translation of information between the timespaces is 'noisy' and imperfect. Within the galaxies however, where timespace replication has effectively ceased, the 'temperature' of the lattice has cooled close to an absolute minimum, and it has become rigid – the lattice that comprises a galaxy is able to translate information between timespaces with practically no loss in fidelity (even where vast particle colliders have been assembled!).

How can it be that you and I are 'moving' through 'solid' superconducting space? Each one of us maps onto a very large number of timespaces, typically about 10^{103} of them. We displace this *volume* of timespaces, but we do not displace the timespaces themselves. Rather, we *are* what the computational states of these 10^{103} timespaces present us to the world as. But we do not consist in the same group of timespaces for very long. Every one of the timespaces that defines us at this instant, in the very next instant transfers its computational state over to its neighbour, in the direction of our net translation relative to the absolute timespace lattice. Indeed, we move on to 'inhabit' different groupings of timespaces in the lattice some 10^{43} times every second. So does everything else in the world. The extreme fidelity of the translation creates the illusion that the world is made of solid matter, which to be fair, is a very ancient and powerful illusion.

Although at any one instant our bodies might displace a *volume* of 10^{103} rigidly positioned timespaces within the galactic lattice, only a fraction of those timespaces are actually executing any particular material application. As we translate through the lattice, the vast majority of the 10^{103} timespaces we each displace are in an idle (zero energy) application state, waiting for just that one 10^{-43} second instant when they get to shine as the holder of the relay torch, only to return in the very next instant back to a resting state.

Other regions of the galaxy have higher energy densities, the sun being an obvious example. However, *all* material applications of the timespaces within our galaxy are understood, from observation, to be translating about a region of extremely high energy density located at the centre of the galaxy. Indeed this region, a super-massive ‘black hole’, has reached the maximum energy density possible. *Every* single timespace in this region is executing a material application, and transmitting vast numbers of gravitational messages out from the surface of its saturated event horizon, to any other active timespace in the galaxy that might receive its messages. Each one of us eventually receives a portion of these ancient messages, and their constant stream keeps us orbiting the centre. Not one timespace in the central region of the galaxy is idle. The lattice in this region has simply reached the limit of its computational capacity.

□

There is of course a whole family of ‘elephants’ inside this theoretical space. For a start, where do all these timespaces actually reside? At the outset, we established that space does not even exist until the timespaces begin to *simulate* it. Hence the machines themselves (during the part of the cycle when they have virtual existence) do not have any volume. They certainly do *not* reside within the space of their own making. So while it makes sense for there to be an ‘exclusion principle’ at the emergent physical level, where no instance of simulated space can encroach upon the territory of any other instance of simulated space, it is meaningless to apply this principle to the machines that are actually *hosting* this composite space. All these machines are literally ‘inhabitants’ of a singularity, a region *without* volume. In a very fundamental sense, every one of these machines is in a ‘superposition’ of computational states.

There is not an infinite number of timespace automata, but there are considerably more than the 10^{185} which we can see of the particular model that has been employed in our instance of universe. Then there are all the other instances of “universe” that have successfully employed this same model of automaton. Then there are all the instances of universe that have opted for one of the 10^{500} different models of vibrating strings that have the potential to construct members of the multiverse. And *all* these different automata collectives logically share the ‘superposition’ – they are all in the same place! Finally, of course, there may be realities which consist in something other than number, time, space and energy, but these are a bit more difficult to imagine.

The timespace automata have an external interface between each other that is defined according to their physical location in the lattice array. As discussed earlier, communication between different locations across the array (in space), cannot occur any faster than the speed of light, as ultimately set by the clock of the automata themselves. However, the ‘superposition’ of all these automata is *outside* space – it is quite literally *nowhere*. Within the superposition, any timespace can communicate directly to any and every other ‘entangled’ timespace through an internal interface. Thus, a couple of timespaces might be at opposite ends of our universe according to their physical location in the lattice, and yet directly pass a message to each other within just one tick of the ‘superposition’ clock.

The holy grail of quantum computing research is, of course, to develop an interface to this vast computation that is going down at the superposition; to develop a 'hyperlink' that would allow us to access these data directly. We could then sample these data, and render quite realistic (albeit approximate, or *uncertain*) facsimiles of stuff that is happening elsewhere in the multiverse, all without ever having to actually step outside (our local universe). Unlike the *virtual* worlds that we have imagined on physical computational hosts like the Internet, the myriad other worlds that we could visit through the superposition are *real* worlds just like ours.

We humans of course are relative latecomers to an understanding of this superposition. Our universe, for example, as noted in the WMAP image, has been creating its own space for about 14 billion years (or 10^{63} 'ticks' of those perfectly synchronized timespace 'clocks'). There has been plenty of time for other thinkers, even within just our local universe, to have progressed significantly beyond where we have so far managed to reach. When we learn how to sample the superposition data, we will be able to index and browse the multiverse's libraries, and merely read, for example, any one the various available proofs (some quite elegant) of a thing that Henri Poincaré once conjectured about simple things.

One of the joys of living is to discover such things by ourselves, rather than reading them in a book, or having them handed to us on a platter. Indeed for some, the life of discovery far transcends any concern for their own physical wellbeing. Luckily for mathematicians, it is the number Omega that guarantees the system of the multiverse has an infinite source of unprovable truths.

Between us, we pretty much know what is required to manage the planet sustainably, and thus indefinitely. What we desperately need to find out is how on earth we can take wing and pull out of this dive we are making towards oblivion. Assuming many others have faced a similar crisis of inertia, it would be very helpful to draw on the experience of those who have successfully managed to get through it. Indeed, we are seeking the best transition programme that the universe has available on offer. Sure, it would be nice to develop a programme in-house, but that's one luxury we can no longer afford.

Finally, we needn't be too concerned about how these strings of digits first came to be assembled. Clever though they, and all of their applications, might seem, algorithmic information theory has shown that a very short programme is capable of deterministically seeding all of these possible universal computations, including ours. Because time is not defined until the first timespace begins to oscillate, the quantum vacuum (which does not exist for half the time, and then neither for the other half) has an eternity in which to make this happen. Hence there is a probability of 1 that the quantum vacuum will eventually fluctuate sufficiently (as indicated in the diagram above) for it to successfully assemble this first simple universal computing machine. The rest is history.

To paint a picture, one needs a palette. Most of the ideas in this essay were borrowed from articles written for a general readership by leading specialists in their fields. Many thanks to John Barrow for all his Cosmic Imagery, and to Mariette DiChristina for orchestrating such a superb resource.

The Limits of Reason; March 2006; Scientific American Magazine; by Gregory Chaitin
Alle berechenbaren Universen; Spezial März 2007; Spektrum der Wissenschaft; von Jürgen Schmidhuber
The Limits of Quantum Computers; March 2008; Scientific American Magazine; by Scott Aaronson
The Cosmic Origins of Time's Arrow; June 2008; Scientific American Magazine; by Sean M. Carroll
The Self-Organizing Quantum; July 2008; Scientific American Magazine; by Jan Ambjørn, Jerzy Jurkiewicz and Renate Loll
Follow the Bouncing Universe; October 2008; Scientific American Magazine; by Martin Bojowald
Naked Singularities; February 2009; Scientific American Magazine; by Pankaj S. Joshi
A Quantum Threat to Special Relativity; March 2009; Scientific American Magazine; by David Z Albert and Rivka Galchen
Does Dark Energy Really Exist?; April 2009; Scientific American Magazine; by Timothy Clifton and Pedro G. Ferreira
Black Stars, not Holes; October 2009; Scientific American Magazine; by Carlos Barceló, Stefano Liberati, Sebastiano Sonego and Matt Visser
Portrait of a Black Hole; December 2009; Scientific American Magazine; by Avery E Broderick and Abraham Loeb
Looking for Life in the Multiverse; January 2010; Scientific American Magazine; by Alejandro Jenkins and Gilad Perez
Boundaries for a Healthy Planet; April 2010; Scientific American Magazine; by Jonathan Foley