Universal Simulation: From Feasability to Insights into our Universe

Thomas Kent

Introduction

From the beginning of human history, we as a species have desired to learn more about our surroundings. Prior to the early 20th century, almost all of human discovery has involved direct observation of our surroundings. This direct observation has led to vast human advancement, in a wide variety of fields from biology to physics. During the late 20th century, however, following the invention of the transistor and subsequently the computer, a new door of human discovery was opened. With the aid of computers, our ability to perform complex calculations has increased exponentially, and with that our ability to study the intangible has improved by orders of magnitude.

When testing hypotheses on large scales, it has become commonplace to perform mathematical simulations of the real world in massively parallel processes. Examples of the use of simulations include the simulation of the early universe following the Big Bang and galaxy formation (Adamek *et al.*, 2016; Lotz *et al.*, 2008). Even with our somewhat primitive computational capabilities, we are currently able to simulate complex systems.

Simulations are most frequently seen in use in videogames, where small self contained environments are simulated with fixed rulesets. Over the last 30 years, the fidelity of simulations within videogames have increased immeasurably, from simple arcade games such as space invaders, to modern videogames featuring physics engines and large 3D game worlds. Therefore, with ever advancing computational capabilities it stands to reason that our ability to simulate ever larger systems will improve.

Obviously, it is not currently possible with human technology to simulate something as complex as an entire universe. If we assume, however, the existence of other intelligent life outside of Earth and then subsequently consider that such a civilisation may have even a small technological “head start” over us, cosmologically speaking, there may exist civilisations that have the computational ability to simulate unfathomably large systems in respect to what we can perform in the present day. Therefore, we consider the possibility of simulation on a universal scale, the consequences this could have on our own reality, and the plausibility of our universe itself being a simulation.

Feasibility of Universal Simulation

When considering the feasibility of simulating something as complex as an entire universe it is important to define exactly what would need to be simulated to form a realistic functional *in-silico* universe, which we will refer to as U_n moving forward. The most obvious assumption would be that a simulated universe would need to simulate at least all of the matter contained within. In our oberservable universe this is estimated to be somewhere between 10^{78} and 10^{82} atoms (Planck Collaboration, 2016). Obviously, this implies there is no universe beyond our observable universe, so that number could be and likely is far greater, addtionally it implies a simulated universe would exist at similar scales to our own universe. For each of these atoms the simulation would need to account for atomic position in three-dimensional space, as well as other properties including electron state and velocity at the very least. The simulated universe would also have to simulate and track all forms of radiation emitted by any particle or cosmological entity as well as any non-matter. It is this non-matter that would likely make up the majority of all data in the simulation, with dark matter alone making up a considerable portion of information within the simulated universe (Bachall *et al.*, 1995).

Therefore:

$I_{Un} = EVM$

and; $A_{Un} = I_{Un} / D$

Where; I_{Un} describes all information contained within a simulated universe and is calculated by combining the total energy (E) within the simulated universe with vectoring information (V) of all constituent matter and non-matter entities and any additional metadata (M) for those entities. A_{Un} describes the total number of atoms required to store in the information in U_n and is dependent on the maximum bit density (D) possible in the parent universe.

Naturally, after defining such an astronomically large amount of information to store the question is - how exactly would a civilisation simulate such vast quantities of information within their own universe, hitherto referred to as U_0 for simplicity.

Any such U_0 civilisation would need exist in a universe that is cosmologically and quantifiably larger than the simulated universe, as through our current understanding no data structure could exist for storing the information contained in U_n using any smaller volume of information than is described in U_n as it would not be possible to simplify any information from U_n without subsequently losing that information (Susskind & Lindesay, 2004). This holds true even for subatomic particles such as the neutrino, which one might assume could be represented by a vector equation and starting position, due to their lack of mass an low interactivity with matter within a universe. However, to preserve the integrity of a simulated neutrino, the neutrino would have to be simulated in full to account for neutrinos that spontaneously change into other lepton variants.

Quantum computing in U_0 could aid in data storage for U_n by allowing for a reduction in the physical number of bits of data to a smaller number of qubits. However, it is important to recognise quite how much larger a traditional storage system would be than the information it contained from a traditional standpoint. For example for a single particle in U_n the equivalent storage structure in U_0 would have to store: a reference to a particle type, perhaps represented by an integer data type of length greater than or equal to all possible particle types; the particles location, which would need to be large enough in scope to allow accurate representation of position on a near continuous scale across massive cosmological distances (with the accuracy of femtometers or less across a range of possibly hundreds of billions of light years); as well as vectoring information and any additional metadata such as electron spin and cloud density. This presents a considerable issue for storage, as storing this much information at current storage densities would require many hundreds or thousands of bits, each requiring silicone transistors for storage, leading to a large information inflation.

A classical example in our current universe of this information inflation problem is seen in the Large Hadron Collider (LHC) at CERN. In this real world example, despite having relatively narrow focus in its measurement instrumentation, the LHC generates around 1 petabye (1 trillion bytes) of data per second when measuring only a few billion collisions. Extrapolating this up to cosmological scales would require an unfathomably large data storage structures to allow computation in any logical sense. The underlying problem rests in that within any universe while particles may exist, they don't necessarily have to be observed and their locations known at any given point in time. When converting this existence-only reality to a measured reality, the volume of information stays the same, but physical representations of this *in-silico* require considerably larger amounts of information. Consequently a single bit transistor has a finite mass that is certainly larger than a single particle, thus inflating the amount of mass needed to store a single bit. As previously described, the amount of raw storage data needed to even describe a single particle in sufficient detail to render its simulation accurate would require large amounts of storage in itself, inflating the information in U_n by orders of magnitude.

This inflation is a problem not only for processing but also for long term storage. Any U_0 civilisation using the U_n simulation for research purposes would likely wish to review historical data, inflating the amount of information storage needed even further. It therefore stands to reason that the parent universe would both have to be massive in relation to U_n but also dedicate vast regions of space to data storage alone. The inflation problem would also present a problem for processing data. As the volume of even transient storage medium increased to perhaps galactic scales, the volume of machinery required to process this basic stored data would rise in a likely exponential fashion.

Without some form of quantum computing, it is also not entirely clear how computation on such massive scales would occur. Even in present day computing, keeping circuitry small is an effective way of both reducing power requirements and increasing computational speeds. On modern processing units, distances of a few millimeters across take fractions of a second for electrons to flow through silicone circuits at the speed of light. In a computing unit on a galactic scale however, the massive distances would present considerable time delays during processing. This processing delay would render simulations on a universal scale extremely slow, which could present significant issues for scientists of U_0 when attempting to analyse any data.

While some of these issues could be alleviated with quantum computing in U_0 it is probable that any processing system, quantum or traditional, would have massive energy requirements. Without a firm understanding of exactly how large scale quantum computing works it is not possible to calculate with any certainty how much energy such a system would require, but it is clear that the amount will be non-trivial. Such a civilisation in U_0 would likely have to exist as a or beyond a Type III civilisation on the Kardashev scale, being able to control and harvest energy on a galactic scale perhaps through the use of a galactic network of Dyson spheres, superstructures theorised to encompass entire stars, harvesting all energy outputted by said stars (Cirkovic, 2015).

To compound these problems, if the civilisation in U_0 was performing these U_n simulations for research purposes, it is likely that they would simulate multiple universes simultaneously, perhaps experimenting with universal parameters to deduce some unknown constants in their own universe, similar to how we perform mathematical modelling to test theories in our universe. This would further increase the energy, processing, and storage requirements of such a large scale simulation.

To the modern scientist this task, while theoretically plausible, seems beyond anything we could envisage humans being able to do. Nevertheless, it is important to consider the possibility that the hypothetical U_0 civilisation could have considerable, currently unimaginable, technological advancements over modern humans.

Reducing computational burden in U_0

Given the vast quantity of information needed to fully simulate on a universal level, it stands to reason that any civilisation would aim to reduce computation burden as much as possible. In fact, we can apply many of the techniques used in modern video games to this hypothetical simulation scenario to vastly reduce the amount of information described previously, and in that regard reduce many of the seemingly impossible requirements of a U_0 civilisation.

The first technique, and possibly most poignant to our own universe, is the idea of "Sharding" the universe simulation. In our universe, we exist within the Milky Way, which itself exists within our Local Group, and it within the Laniakea Supercluster (Tully *et al.*, 2014). Due to cosmic expansion, however, only the entities within our Local Group are gravitationally bound to us, and other groups are instead accelerating away from us, eventually faster than the speed of light. Therefore, even at the cosmological speed limit, the speed of light, it is impossible for particles or otherwise to travel between such groups.

In modern video games, Sharding is used to describe the splitting of large game worlds into many smaller ones - called Shards, which exist within the same world (or universe) but with a reduced number of users or entities in a given shard. This massively reduces the amount of information that needs to be shared between clients by in turn reducing the number of non-static entities visible to each client.

In a universal simulation context, a U_0 civilisation could take this a step further, simulating only a single Group or supercluster of galaxies, ignoring the vast majority of the rest of the universe. Obviously, this would not be possible from the beginning of the simulation, when the majority of the universal mass is in close proximity. Instead the U_0 civilisation could broadly simulate the early universe with minimal resolution, similar to how recent simulations by Adamek *et al.*, have simulated the early universe. Such simplistic simulations would only have to be sufficiently high in resolution to rapidly simulate the early gravitational capture of pockets of the universe, and the formation of galaxies in a simplistic fashion. From there, using a large but restricted dataset, the properties of galaxies could be measured to generate models of how galaxies behave generally, and those models applied to galaxies in other Groups. Simulation of the main "simulated Group" could then continue in high fidelity, with properties such as radiation emission from distant groups being generalised without the full group and its constituent galaxies, stars, planets etc also being simulated in high fidelity.

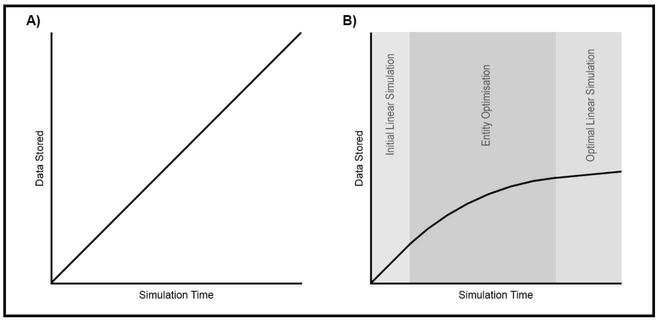

Figure 1: Representation of data accumulation when simulating over time in a linear non-sharding fashion (A) and with the use of sharding to discard unnecessary data (B). Initial simulations would accumulate data in a linear fashion as the early universe is generated. Following gravitational capture, entity optimisation could reduce the rate of data accumulation until complete segregation of the active group from extraneous groups allowing for optimal linear simulation.

One potential argument for Sharding is actually that of human dreams. When dreaming, we create small pockets of "reality", even using the limited processing capabilities of the human brain. These dreams, while limited in scope, provide evidence for vastly reducing the processing cost of simulating an entire reality, simply by truncating the vast majority of peripheral information.

One caveat of Sharding the simulated universe, however, may be that interactions between more distally close groups may be missed during the filtering and discarding process of non-gravitationally bound Groups. During later timepoints, when Groups are sufficiently separated, no interaction between groups will be seen. In a young universe, however, such Groups maybe close enough to contribute radiation to their neighbours, ultimately propagating an altered the state of nearby Groups. Due to the inherent flaws of using low resolution simulation these alterations will be ignored, leading to potentially unrealistic information later in the simulation once many of the precursor particles have already been discarded. Nevertheless, using Sharding the total computational burden of the U_n universe could be reduced by many orders of magnitude, to allow for considerably lower information storage and processing requirements.

Another method of reducing computational burden would be through "observer-only processing" - another technique borrowed from modern video games. This technique ultimately relies on two foundational principles, that the universe itself is deterministic, and that "consciousness" is a cosmic phenomena. In the case of video games, this is frequently used to reduce draw calls in graphical processing. When building a 3D world it will often take up many times more bytes of memory than a graphical processing unit (GPU) contains. However, to combat this, video game designers make use of observer only processing to only render the portions of the world that are currently being observed by a client. In some cases, developers take this a step further and additionally suspend physics and logic processing outside of view.

From the point of view of a simulated universe, this technique would have to be more restrained, as one could not simply stop processing radiation emission from distant stars simply because no conscious being could observe such radiation releases until they have travelled many lightyears. In this sense, stars themselves would be a paradox - as no star would be in view of a conscious being until the light had already travelled vast distances, but that light would never have been emitted due to a lack of an observer.

To combat this, insead an emitting pulsar could be simulated to release radiation, the vector of that radiation be calculated and the intersect timepoint with a fully simulated, conscious, being or particle calculated. At this point the path of the radiation could largely be ignored until such a time when the radiation encounters an observer or particle that is actively being observed and thus exists within that reality.

Naturally, such a system would be potentially lossy, ignoring potential interactions within the time between emission and interaction with an observer, however such lost information could be justified to decrease the complexity of the simulation.

An additional benefit of observer-only processing would be the potential for reducing the precision needed for spatial information for matter. In the previous description of the amount of data needed to store the position and related data for a single atom we suggested that positional data would have to have an incredible level of precision to allow for the vast range of positions available in a large 3D space such as a simulated universe. However, using oberserver-only processing, and simulated universe could hypothetically maintain spatial positions of matter in two formats, a large scale but low resolution format for basic position, and a large scale high resolution format for specific positioning. The prior, having low resolution, could be used to maintain the position of matter that is distal from other fully simulated matter, such as in the near emptiness of space, where interactions with other matter or radiation is less likely. Upon predicted interaction, the vectoring information of the matter could then be used to calculate an actual high resolution position, which could then be used from then onwards for calculating complex interactions. This would reduce the storage complexity of the entire system, by allowing the system to offload many of the non-interacting, partially simulated entities to a less memory intensive location system.

The final mechanism for reducing computational burden is not a technique *per se*, but an important distinction in relation to our perception of time. As with modern computing, simulations do not neccessarily have to take place in real time. Simulated entities would experience time from their own U_n viewpoint. That is to say when simulating computationally expensive timepoints in U_n the amount of time that passes in U_0 may be orders of magnitude longer, without impacting U_n in any way. This is common in modern videogames, where individual points in time are simulated faster or slower than real time, allowing for the perception of sped up or slowed down time for the player, but with no change in the internal time that passes from the point of view of an in-game entity. In this regard, the rules that govern physics within the game operate on the in-game passage of time, not the time that has passed outside the game.

This principle could also be applied to both U_0 and U_n, allowing a U_0 civilisation to rewind, slow down, or speed up passing time within the U_n simulation to observe cosmic phenomena or accelerate the formation of galaxies and stars. In this respect, a civilisation in U_0 would not need to be able to simulate U_n in real time, with the trade off that when simulating in slower than real time (i.e. where one second in U_n requires more than one second in U_0), observing the effects of perhaps billion-year events like galaxy formation would require more time than this in U_0, requiring long or infinite lifespans of researchers in U_0. Therefore, it is likely that a processing system in U_0 would have to be able to fully simulate a universe in Un at many thousandfold the speed of real time.

Consequences for our Universe

When considering the feasibility of simulation on a universal scale, it brings to mind a number of questions, the principle question being - is our universe a simulation? Obviously, this is likely an intractable question, as unless our U_0 creators indicated to us that we were in a simulated universe we would likely never know. Our current universe, however, does provide us with some evidence for the simulated universe theory, and highlights some striking consequences for our existence.

One such piece of evidence in favour of the simulated universe theory relates to the previously designed Sharding concept, and additionally adds a caveat to the simulation parameters chosen by our hypothetical U_0 civilisation. Should we assume that the U_0 civilisation does not intend to simulate a universe that is like their own, and is instead performing a simulation for entertainment or a thought experiment, we should consider the possibility that many of the fundamental laws of physics we experience were set that way to reduce computational load, and do not reflect the variables seen in U_0. One such variable is the speed of light itself. No plausible explanation has yet been proposed as to why the speed of light is fixed at 300,000km per second other than “that is the way it is”. If we consider the ramifications of such a speed limit, an increased speed limit (or infinite speed limit) would lead to a breakdown of the Sharding technique - as any matter could accelerate, perhaps infinitely, and cross between galactic Groups, rendering the technique useless.

Similar arbitrary variables include the rate of expansion of the universe, the strength of gravity itself which led to the formation of Groups in the early universe. All of these variables have allowed our universe to form exactly as it is now, and yet a change in any of these variables could have led to a markedly different universe.

In this regard, should our universe be simulated in a Sharded manner, it is possible that many of the more distant cosmic bodies we observe and hypothesise simply do not, and will not ever exist. Instead, such bodies may be observable from afar but never actually be available to visit, even if we could break the speed limits of our universe.

This brings about a fundamental philosophical question for humans - if we cannot directly observe something, in an interactive manner, can we be sure such an entity actually exists?

Another possible consequence of our universe being simulated is the idea that any computational system always has the risks associated with software design. In this regard we could compare our universe to the reality seen in the 1999 science fiction movie The Matrix. In the Matrix, humans have been enslaved by machines and their minds uploaded into a digital reality. A critical plot device in the film revolves around the ability of minds within the Matrix to alter the simulated reality itself. Obviously, it would not make much sense for a U_0 civilisation to give the inhabitants of their simulated universe such powers, however, it brings to mind the possibility that, should our universe be simulated, our universe may contain “glitches”, or software bugs, in our own “matrix” that were unintended by our universe’s creators. In the Matrix, these are sometimes seen as physical glitches in the world, but could also describe cosmic features such as theorised “wormholes” that would allow for effectively faster than light travel across vast cosmological distances.

One of the more existential consequences of a simulated universe would be that, despite the billions of years our universe has existed thus far, there is always the possibility that our universe is simply a research project that will, inevitably, cease to be of use at some point. Our universe, therefore, could at any moment be terminated and no more timepoints simulated. The possibly most frightening consequence of this is that we would never be able to perceive this, and simply vanish from simulated existence momentarily.

Another consequence of a simulated universe, specifically one that is utilising the observer-only processing method described earlier, is that humans (or life on Earth) could be designated as the only set of “observers” in our entire reality. This could explain why, despite the astronomically large number of potential planets capable of supporting life as we know it, and the possibility that any of these planets could have had, by evolutionary timescales, huge developmental head starts over Earth, we have yet to meet any intelligent life outside of Earth. Indeed, in order to maintain a relatively small number of observers, our U_0 creators may have intentionally prevented the formation of intelligent life outside of Earth, by narrowing the area of direct simulation down to only a few light years around our solar system.

On the contrary, another provocative thought experiment is to consider if our universe itself may be U_0, and somewhere within the vastness that is our observable universe, among countless trillions of galaxies, there may be a civilisation that had sufficient time to develop the technology required to engage in universe simulation. Further, given sufficient scientific progress, could humans one day engage in such large scale simulations?

A final consequence of a simulated universe brings about the notion that if our universe is indeed simulated, there may be other parallel U_n universes being simulated en mass in U_0. This theory aligns with the multiverse theory of multiple universes, in that our universe may be one of many, perhaps linked or perhaps forever distinct. In the case of simulated universes, it may or may not be possible to communicate with either our parent universe or other simulated universes under U_0. Additionally, our parent universe, U_0, may itself be simulated, in a recursive line of simulated universes. The origins of our universe may be simply be attributed to a scientific simulation by a vastly superior intelligent life form beyond our reality. However, universe simulation cannot adequately explain the origins of the top-level, original universe.

Conclusion

In conclusion, the concept of a simulated universe is one that involves many assumptions, and one that seems technologically infeasible to modern humans. Indeed, simulation on a universal scale would require more computational power than humans may ever be able to achieve by many orders of magnitude. However, when considering the scale of even our own universe, whether it be simulated or not, it is not impossible that a civilisation has already reached this level of technology. Should our universe itself be simulated, this could help describe many of the universal constants we take for granted, seemingly arbitrary constants in a turbulent, evolving universe. And, should our universe be simulated, there is always the risk our creators suddenly decide - *we don't need to continue this simulation any further.*

References

Adamek J, Daverio D, Durrer R, Kunz M; General relativity and cosmic structure formation. Nature Physics; 12 (2016).

Bachall N, Lubin L, Dorman V; Where is the Dark Matter? The American Astronomical Society; 447 (1995).

Lotz J, Jonsson P, Cox T, Primack J; Galaxy merger morphologies and time-scales from simulations of equal-mass gas-rich disc mergers. Monthly Notices of the Royal Astronomical Society; 391 (2008).

Cirkovic M; Kardashev's Classification at 50+: A Fine Vehicle with Room for Improvement. Earth and Planetary Astrophysics; 191 (2015).

Planck Collaboration; Planck 2015 results. XIII. Cosmological parameters. Astronomy & Astrophysics; 594 (2016).

Susskind L, Lindesay J; An intorduction to black holes, information and the string theory revolution. (2004).

Tully R, Courtois H, Hoffman Y, Pomarede D; The Laniakea supercluster of galaxies. Nature; 513 (2014).