Biography

Tom Barbalet created the Noble Ape Simulation in 1996 and continues its development to this day. Noble Ape is used by Apple and INTEL as well as a number of universities to teach biodiversity, multi-media education, vector processing, real-time graphical interfaces and a number of other technologies. He has been the editor of Biota.org (since 2005), a leading community resource for artificial life developers. He is the host of the weekly Biota Live internet radio show where he discusses a variety of topics relating to artificial life, artificial intelligence and simulation philosophy with a variety of guests. He is also the co-chair of the International Game Developers' Association's Intellectual Property Rights Special Interest Group.

X
Welcome to the Simulation

TOM BARBALET

tom@nobleape.com

I was awakened from my dogmatic slumber by the Nick Bostrom paper, "Are We Living in a Computer Simulation?" (2003). Simulators may want to be pigeon fanciers waiting for Darwin (Secord, 1981). I believe simulators can be both pigeon fanciers and Darwin.

Field of Reference

When I lived in England I would sit in my study and look out on a field. It was a large field that was sometimes used as a fair ground, sometimes as a sporting field and at least once a year as the town's location for fireworks. Aside from a large grassy center that could easily accommodate a football game and a group of spectators, there was a good combination of wild plants and some cultivated trees on the fringes. In the far corner of the field was the town's garden lots, where townsfolk would grow their vegetables in small plots. Much to the discomfort of people who strayed from the surrounding pathways there was poison ivy growing close off the path.

In all seasons, but winter, the field would be a sea of rich green vegetation. There was a good quantity of insects and a vast slew of birds that would fly around the field showing their aerobatic prowess. Following frequent rain, these birds would spend hours pulling worms from the field. It was clearly a rich source of organic production to support such a healthy worm population.

In the winter, the field was covered with snow. I peered onto it through the ashy silhouettes and complex geometries of the exposed tree limbs guarding the field. Footprints and paw prints would be left in the snow as a memory of humans, dogs, foxes, and prints that could only be attached to wolves. The prints in the fresh snow served as a memory of the life that moved through the field.

Fifty years before I sat looking over the field, Alan Turing lived in the town (Hodges, 1983; Copeland, 2004). Although it was the latter part of his life and he had already broken the Nazi's Enigma codes (in the process inventing modern computing), I imagined him walking through the field. Perhaps he had worked in the garden lots or perhaps he had watched a local football game being played.

About five minutes walk down the street from the field was a chain of ponds and an English swamp, affectionately known as a bog by the locals. There a prehistoric man had been found and named after the bog - Lindow man (Robins & Ross, 1989). Strikingly, the man had been exceptionally well groomed for an Iron Age man. Lindow man had no-doubt walked through the Iron Age version of the field on the way to his near-eternal preservation. Perhaps he had been part of a large group, perhaps he walked alone. He could have died naturally and been buried there. Maybe the field had been the site of Lindow man's home. What had this man seen? What kind of conflict had he experienced? Were his genes flickering in the local population, even now, crossing the field?

The field and its surrounds represented a theater of life. Every aspect of the field provided some connection with life. So while life may have an ethereal descriptive quality, every aspect of the field and its surrounds could be modeled with a simulation derived from observation; the field's operation in relation to the town as a social gathering place, the rationale of the human placed tree cultivation and the vegetable plots, the growth and variety of the competing ecosystem of plants and insects through summer, the paced routes of humans and animals through winter, also the chemical and bacterial interchange that preserved Lindow man, not to mention the possibilities of the Lindow man Iron Age society, and, even Lindow man's genetic descendants through to Turing and the snow walking humans.

Surely replicating the observable qualities of the field would yield no additional information?

Such a process is analogous to an animated film or perhaps a historical re-enactment where you can move through the observations but get no more information about history except through observation. You can't say to the re-enactor, "If your side had lost the battle rather than won, what would the world look like today?"

Simulated life solves this problem in the movement from creating environments that have elements of realism into what-if experiments — exploring limiting and expanding parts of the simulated environment. Rather than merely recreating observable life, simulation provides an environment where what-if experiments can be enacted.

There may be a bias to presuppose simulation refers to something deterministic, designed by intelligence or in some way constructed to be based on a small set of rules. This isn't the way the word simulation is used here. In fact, the simulations described in the field can (and should) in most cases be chaotic.

The purpose of this text is to argue that simulation developers (simulators) aspire to perfect their simulations to the point that simulation is reality. This change of reference where reality is simulation enables a better understanding of contemporary and future simulation development and also provides additional analytical tools which are not available if simulation is considered as a created model of a small fraction of reality.

Simulation Metaphysics

The word simulation comes from the Latin *simili* which means like. A simulation is a system constrained to reproduce another system.

Suppose you wanted to build a time tracking simulation of the Mississippi River to look at shipping along the river in the late nineteenth century (Twain, 1883). You could take a piece of paper and draw the traversing Mississippi River along the paper identifying points where the boats would stop. You could then move counters along the Mississippi River route to show the shipping commerce. You created something like the Mississippi River but rather than using the waterway itself or collecting information directly from the waterway, you created a simulation on paper of the waterway and the routes taken. It may seem rather strange to say by adding the route of the Mississippi River to the piece of paper you were constraining the paper in some way, but you were adding constraints to the system that the paper created. You could, however, add a number of fictitious tributaries to the Mississippi River and speed up some sections of water commerce. This identifies how the constraints on the simulation can be expanded to allow for greater freedom.

A computer simulation is another good example. If you wanted to simulate the interaction of two marbles striking, you would write a computer program that provided mathematical constraints. The idea of constraint is critical to understanding the purpose of a simulation (Winsberg, 1999). The computer is just as programable to make the marbles sprout arms and legs and have a multi-round boxing match. The constraints are designed to make the simulation close to what is being simulated. You want the simulation "like" the simulated.

Freedom (in a simulation) relates to possibility rather than actuality. For example a computer has near limitless possibilities which makes it an ideal simulation tool. We shouldn't be predisposed to thinking of simulation solely as part of some mechanistic construction. Simulations don't have to just exist within a computer or through the imaginations of intelligent agents as the paper Mississippi River example shows. The most powerful simulations can exist in organic environments by placing constraints on the participants. These constraints can be physical, they can exist in social mores, written laws, best practices and perceived constraints.

With this broad definition of simulation, it is possible to look at road systems as simulations, legal systems as simulations, the financial systems — from buying a pint of milk to the stock market — as simulations. In fact we are the willing and unwilling participants in a number of simulations that shape every aspect of our lives.

Aside from the appeal to the Platonic form of simulation as reality, an appeal for our collective participation in simulations could come through a contemporary Cartesian Cogito, ergo sum;

I am constrained, therefore I am simulated.

Consider that our environment is far richer an environment than we could possibly utilize. In fact our environment is constantly being constrained not by itself but by our interactions with the environment. The simulation dichotomy of two environments — the primary environment and the simulation environment — requires the simulation environment to be far freer than the primary environment. The simulation environment has the addition of constraints to bring it to the primary environment.

If the real world is the primary environment, the simulation of the real world will also allow for possible worlds. But through the constraints of the real world, the simulated world is the primary environment. Simulation allows the exploration of possible worlds as well as the real world.

Alive without Intelligence

The word “alive” has a number of different meanings through contextual use. If we ask the question “Is simulated life alive?” the first problem that is posed is which meaning of alive to use in the question. It would be trivial to define alive as being something biological — requiring respiration — or something that required a physical presence. These kinds of definitions of alive would result in an immediate negative in terms of simulated life being alive.

Aside from our intimate connection with “alive”, we observe others around us that are “alive”. The observation is easiest to see in those around us in physical proximity, both those we know and have conversed with, and also the people we may see driving or walking down the street. This observation is important because we don’t need to interrogate other entities in order to establish that they are “alive”. We don’t need to talk or interact with people on the street in order to establish they too are “alive”. It is something we are very familiar with on an implicit level. This idea of “alive” is also applicable to people we see in films and on television. In fact, when we see human actors — whilst we understand they may be performing a play — they are “alive”.

This characteristic of “alive” becomes particularly interesting with animated cartoons. The actors that give the voices to the cartoons are “alive”. But are the cartoons “alive”? This is a question of suspension of disbelief. This is particularly prevalent with films or television programs that use computer generated actors that interact with real world actors. The better the computer generation, the easier it is to suspend the disbelief and assume for an instant (or longer) that all the participants — human and computer generated — are “alive”.

Through this analysis, “alive” can be independent of scripting. Even if we are watching a play, we know the actors are “alive”. We can also acknowledge that there can be computer generated entities that also have this “alive” characteristic. This is particularly important for simulated life because “alive” here can have one of two components. Either the entity needs to look good — look “alive” — or they need to act well — act “alive”. Now whilst all simulated life developers strive to create stunningly beautiful representations of their simulated agents, the characteristic simulation developers strive to create is something that acts “alive”. That is the aim of simulated life.

It would be fair to presume that making something that acted “alive” required intelligence, intelligent design or an intelligent designer.

There are two simple means of showing this isn’t the case and that intelligent design is not part of the simulated life developer’s process. The first comes through paleobiology and the second comes through a simple thought experiment.

The fossil record shows that there was life before there was intelligence. Intelligence is a cunning adaption that aids life immeasurably, but life clearly existed before there was any intelligence in the life. The paleobiologist Roy Plotnick (Barbalet, Daigle, Kerr, DeJong & Plotnick, 2006) narrates this beautifully by describing the pre-Cambrian floating fauna that slowly started adapting to floating between feeding grounds. It was this understanding of feeding grounds that allowed for the amazing variety of fauna that came through the Cambrian period. The most primitive marker for intelligence comes through changing movement to optimize for feeding grounds. This was the first glimmer of intelligence but life came first. If you were to simulate this period, you couldn't do it if you designed with intelligence. Intelligence is the outcome you would desire, but if you wrote it in explicitly you would defeat the experiment.

The thought experiment for this is relatively easy too. Consider what life is fundamentally, not over generations, but in an instant. Life is an ability to survive through a changing environment. Consider that the changing environment is chaotic — there are things that may be predictive in some regard — but fundamentally there is just chaos. So life is an ability to survive chaos.

This definition of life is removed from intelligence. Intelligence is very good at understanding the predictive part, but intelligence can't quantify and work through chaos. There needs to be something that is sub-intelligence that deals with chaos.

There are a number of solutions to this problem. The most trivial example comes through tuned dynamic equations — which are clearly not intelligent but sufficiently reactive to act "alive". From this, there are higher orders of response which act "alive" too. This emerges into patterns which look intelligent but the components that create this perceived intelligence are dynamic and not something that could begin to be quantified as intelligence.

In order to understand this with real world examples, consider a moving object on a fluid — a surf board with a surfer — or a moving object that needs adjustment for balance — a bicycle with a rider — these are fundamentally dynamic and the responses to maintain them upright have to be equally dynamic and not intelligent. It is this reflex which the simulation developer needs to master if their creations are to act "alive".

Consider a surfer or a cyclist that thought about their response to the dynamic environment. This is a starting surfer or beginner cyclist fundamentally. It is also a response that would not yield balance. If you start from intelligence without reactive dynamics, the agent may be able to act clumsily, but they won't act "alive".

Intelligence and the Game Hunter

I have offered the briefest definition of intelligence as being a continuation of the survival element of life, life being the ability to survive chaos. With this definition a number of interesting properties of intelligence can be found.

Whilst this is an implicit property of this definition, I think it is important to slay the dragon early in this exploration. The human has no primacy in terms of intelligence. As someone who has developed artificial intelligence (in a simulated environment) for a number of years, I've always found the popular obsession with human intelligence rather curious. In fact it is a good litmus test for hucksters in the artificial intelligence community when they start talking about how we are looking to create a digital version of human intelligence.

When I have been interviewed about this issue, I have returned to the following parable (Barbalet, Trumbule, VanNuys, 2007):

Two horses are standing in a field overlooking a freeway. One horse looks at the other horse and then looks out at the cars driving along the freeway and says,

"Sure they are fast, those cars, but they'll never be horses."

This also illustrates a secondary point that people developing artificial intelligence rarely look to make human intelligence when they have something faster, and through utility, superior.

Once you lose the pinnacle of human intelligence as the height of intelligence you start to realize something very powerful. There are a lot more intelligent systems out there.

Looking at the survival property of intelligence, you start to find systems which hold this survival property very strongly. A personal favorite of mine is the road system. If you have ever driven through an area of natural disaster, you'll realize that it is amazing how felled trees are moved or new disaster roads are created. The need for movement produces a survival in the road system which is quite fantastic. Taxes, motor vehicles and roadworks crews go into making the road system, that's true. There can be no bias against systems that require additional components. The human brain requires a huge conditional upkeep in order to keep it functioning and until a few paragraphs ago, it was the pinnacle of intelligence.

There are three additional systems I like to discuss as simulations of intelligence. The legal system, the financial system and the information system containing the Internet. The latter is a favorite point of discussion particularly considering how one would begin to rate the intelligence of the Internet.

How do you begin to rate the intelligence of these systems?

I have a slight bias in thinking that all these systems including the road system are vastly more intelligent than humans. They have certain elements of fragility without question but they exist and survive for far longer and achieve far more through their survival than any human.

A pragmatic quantity of intelligence may relate to the number of humans it would take to stop a particular system. This is the big game hunter calibre test (Roosevelt, 1909) for intelligence.

It took but ten humans to slay this intelligent system.

Because this is such a rough measure, I'd like to empower the logarithm base ten as the way to fairly weight this measure. So in the case of ten humans to slay the intelligent system, it would have the intelligence value of 1, a hundred humans then the value of 2, and so on.

Whilst these ideas may seem a little radical, I am not appealing to conspiracies. I understand the road system is probably best measured with regards to a functioning city. Living in Las Vegas, I have observed the road system can survive with 10 human caused obstructions (Las Vegas Metropolitan Police Department, 2008) but I doubt it could survive with 100. So the intelligence value of Las Vegas' road system is somewhere between 1 and 2.

Similarly, I wouldn't suppose that there is an international legal system. But within a small county legal system, I have recently seen hundreds of immediate cases can effectively shut the legal system down (Dougherty & Holusha, 2008). This would indicate an intelligence value of a small county legal system as being somewhere between 2 and 3.

Take a large city's legal system – the effective division of labor – and this number would no-doubt be a lot larger.

When you consider a national or the global financial system, you begin to realize that it takes millions of humans to have but the tiniest effect. Now it could be argued that times of interpreted catastrophe involving a relatively small number of people can cause fluctuations in the global financial system but the metric is based on slaying the system, not leaving it fluttering.

How did these systems become so powerful?

There appears to be a combination of factors — a lot of time, divergent and competing engineering principles. Fundamentally a "survival first" pragmatism. This pragmatism isn't intelligent design, it's life.

Computational Power

When I started developing Noble Ape, I used the computers that were readily available to me as a nineteen year old university student of limited means. Not the university's computer resources, but the kinds of machines I would keep running near me for regular updates when I had a moment to spare. Machines that I could find for next to nothing (Barbalet, 2005b). I started developing Noble Ape on technology that was already seven to ten years older than the state-of-the-art of the time. For this reason, I have seen the same fundamental simulation run on more than two decades of advancing hardware. In addition to this, I have also had the privilege of working with some of the most brilliant hardware and software engineers at Apple Inc. (since 2003) and INTEL (since 2005) with Noble Ape (Barbalet, 2005a). Looking at how INTEL uses

Noble Ape in particular I have a deep respect for the advances in contemporary computer hardware.

It is difficult to quantify the advances in computational power over the past twenty years because all applicable popular cliches fail to grasp the increase in power but also the kinds of computation that can currently be achieved which would never have been considered possible even a decade ago. In this context, the greatest advances of recent years have also provided the revolution that simulation (as discussed in this text) must embrace. This is the idea of atomic computation – not of quantum computing – but of taking computer algorithms and translating them into something that can be run in a vast parallel way over a number of different processing cores and potentially over a distributed network. In addition to this, the advances in vector processing means that single processor cores can now be thought of as multi-stream processing pipelines.

This kind of processing power is inconceivable in a popular context. Contemporary computing hardware has the power to simulate anything we would wish to throw at it. The simulation problems merely need to be optimized for the current processor architectures. This is a non-trivial exercise.

The problem with translating contemporary computing power to something which is meaningful for simulation is a software problem. Software has always been the great lag in terms of translating the power available in computing to a popular audience. Ironically the simulation systems discussed – the financial system and the information system containing the Internet – utilize this power considerably greater than simulated environments that would translate well to human-scale understanding.

Consider the Newtonian model of the physical world. Mathematics was never the shortfall in the adoption the Newtonian model. No one ever said to Newton, "Your physics is great but unfortunately the mathematics isn't powerful enough yet to run your physics." The same is true with contemporary computing. The power is there to run vast simulations. The issue currently is there is a need for new mathematics to translate these simulation environments into a context that contemporary computing can embrace.

This is a challenge.

New Science

Simulation has existed in science prior to computation. The broad definition of simulation offered here includes most (if not all) of experimental science to-date. A confined hypothesis-testing experiment is a beautiful subset of the definition of simulation I've offered.

Computer simulation for testing specific scientific hypotheses isn't new either. To frame the problems discussed here, the simulation discussed in this text – with regards to the language, philosophy and mathematics of simulation – doesn't refer to serial computer simulation to find particular variables or test a simple hypothesis. Simulation in the context of this text relates to vast parallel simulation that has only been computational reality in recent years.

This text is arguing that the ground rules have fundamentally changed and contemporary computing can offer insights back to science which could never have been considered with the "simulation to test scientific hypotheses" interface (Barbalet, Damer, Gordon & Schafer, 2008).

The value of implementing existing science in simulations (as a starting framework) relates to a few basic problems. It is important to note the scope of science in this context because the use of mathematics in simulation is critical. The applied mathematical analysis that physics offers is also critical but one of the main problems with science in the context of simulation can be described through writing simulation that solely replicates physics. To-date this has been shown to be problematic outside simple physics simulation (Barbalet & Damer, 2007). There is nothing wrong with simple physics simulation but it isn't the same context of simulation discussed here.

The kinds of problems in physics relating to complex systems are the bread and butter of contemporary simulation. There is also a boundary condition problem with simulation of purely physics (or chemistry, or molecular biology, or another existing non-simulation derived scientific method). The equations and ideas from traditional science are created in a tight frame of reference and generally don't translate well to simulation.

In the face of these boundary conditions, simulators have used a wide variety of modeling methods that either under-approximate the effects of physics (for example) or even better allow the simulation to push the boundaries of contemporary physics. It is the ability for simulation to prompt and develop new science that is particularly exciting for many simulators and scientists alike but requires a new kind of thinking (Barbalet, Damer, Gordon & Schafer, 2008).

There are two major obstacles to this new science.

The first relates to the mathematical language of simulation. Whilst there are a number of traditional computational methods, contemporary computing creates new shortcuts which traditional computational methods underutilize. As much as this text is a call to change perspective on the philosophy of simulation, it is also a plea for improvement in the mathematics of simulation. As physics was based on mathematics, simulation is based on computation. In short, mathematics plus atomic computation – vast parallelism, intercommunicating mathematical processes with shared and local memory – is the new simulation paradigm that requires a new kind of mathematics.

The second obstacle relates to the corrosive nature of simulation on traditional methods. In short, science will need to have the capacity to change in order to embrace the potentially radical insight that contemporary simulation can offer.

The first and the second obstacles are fundamentally intertwined because there needs to be an education process that goes hand-in-hand with the new simulation methodology. The outline and methods of this education are outside the scope of this text. Except to say that the philosophical problems contemporary simulation offers to metaphysics are dwarfed by the education problems contemporary simulation poses.

There is a lead-by-example solution to these obstacles. If the new mathematics were in place, if the first obstacle has been overcome, the second obstacle melts away as the insight begins to flow. In this regard, contemporary simulators should focus on communicating and developing the new mathematics. The people who hold the best initial capacity for this are the super skilled employees of contemporary semiconductor manufacturers. My work with engineers at INTEL has identified the capacity for the new mathematics through these kinds of developers. Hopefully the new mathematics can move outside a proprietary context.

Moving the Discussion Forward

It may seem rather curious that a text of this nature would appear here in the context of science and belief. Within the simulation community there has been a lot of discussion about whether simulators are intelligent designers. My own view is that physicists and mathematicians don't have such a burden about their endeavors (Barbalet & Daigle, 2008). They just get on with doing what they are doing. In this regard simulators should follow suit and explore the amazing new scientific and philosophical landscape contemporary computer simulation offers.

This new landscape is about computation exploring aspects of the real world that may be discussed at a very high level but are still fundamentally unknown. Independent and inquiring minds should begin to push boundaries and explore new possibilities — simulation is the way forward.

References

Barbalet, T. S. (2005a). Apple's CHUD Tools, INTEL and Noble Ape. Noble Ape Website. Retrieved June 20, 2008, from http://www.nobleape.com/docs/on_apple.html

Barbalet, T. S. (2005b). *Original Manuals, Noble Ape 1996-1997*. San Mateo, CA: Cafe Press.

Barbalet, T. S., Daigle, J. P., Kerr, D., DeJong, G., & Plotnick, R. E. (2006). Can Artificial Life Explain the Cambrian Explosion?. *Biota Podcast*. Retrieved July 21, 2008, from http://www.biota.org/podcast/biota_020307.mp3

Barbalet, T. S., & Damer, B. (2007). Bruce Damer Chat. *Biota Podcast*. Retrieved July 21, 2008, from http://www.biota.org/podcast/biota_102307.mp3

Barbalet, T. S., Trumbule, J., & VanNuys, D. (2007). Shrink Rap Radio Live #3: Artificial Life and Artificial Intelligence. *Shrink Rap Radio*. Retrieved July 21, 2008, from http://www.nobleape.com/transcript/srr126.html

Barbalet, T. S., & Daigle, J. P. (2008). More Outstanding Questions. *Biota Podcast*. Retrieved July 21, 2008, from http://www.biota.org/podcast/biota_052308.mp3

Barbalet, T. S., Damer, B., Gordon, R. & Schafer, S. (2008). the EvoGrid (May Update). *Biota Podcast*. Retrieved July 21, 2008, from http://www.biota.org/podcast/biota_053008.mp3

Bostrom, N. (2003). Are We Living in a Computer Simulation? *The Philosophical Quarterly*, Vol. 53, No. 211, pp. 243-255.

Copland, B. J. (2004). *The Essential Turing*. Oxford University Press.

Dougherty, J., & Holusha, J. (2008). Chaotic Hearing in Polygamist Case. *New York Times*. Retrieved July 21, 2008, from http://www.nytimes.com/2008/04/18/us/17cnd-polygamy.html

Hodges, A. (1983). *Alan Turing: the Enigma of Intelligence*. London: Burnett Books.

Las Vegas Metropolitan Police Department (2008). *Current Traffic*. Retrieved July 21, 2008, from http://www.lvmpd.com/cgi-bin/traffic/traffic.pl

Robins, D., & Ross, A. (1989). *Life and Death of a Druid Prince*. New York: Simon & Schuster.

Roosevelt, T. (1909). *African Game Trails: An Account of the African Wanderings of an American Hunter-Naturalist*. New York: Syndicate Publishing.

Secord, J. A. (1981) Nature's Fancy: Charles Darwin and the Breeding of Pigeons, in *Isis*, Vol. 72, No. 2 (June, 1981), pp. 163-186.

Twain, M. (1883). *Life on the Mississippi*. London: Chatto & Windus.

Winsberg, E. (1999). Sanctioning Models: The epistemology of simulation, in Sismondo, S. & Gissis, S. (eds.) (1999) *Modeling and Simulation*. Special Issue of Science in Context 12.