Elves in the Machine

Tom Barbalet and Bruce Damer

Excerpted from Psychedelic Salon #283

September 28, 2011

Tom Barbalet: It's a pleasure to be on the Salon today and have the chance to rap with Bruce Damer about the strange and eclectic legacy of Terence McKenna in the artificial life community. It's probably a good place, Bruce, to introduce what the artificial life community has been historically and where we both fit in the puzzle.

Bruce Damer: I could probably give a bit of a summary, and one of the interesting things we should probably bring in, too, is that I recently came into a collection of letters written to and from Terence McKenna around his Time Wave Zero project, and the original Time Wave Zero software and manuals. I've got them running here, because I have old machines that run old software in the computer museum. Somehow that all leads into artificial life and machine novelty. . .

TB: We have slightly different views with regard to the artificial life community. I think the idea of artificial life predates the traditional discussion of computation. I read back into Plato's work a number of the concepts of artificial life associated with what-if scenarios, and it's interesting in terms of the early definition of artificial life given by Chris Langton as "life as it could be," because I think that humans have been thinking about life as it could be well prior to computation. Maybe you can pick up the story from there.

BD: At the beginning of research for my Phd on the EvoGrid project, I went to see Freeman Dyson at the Institute for Advanced Study at Princeton. At the same time as I was meeting with Freeman, I trundled across the lawn and met with the archivists for the Institute, because the Institute was the home of, one could argue, the first fully complete modern computer, designed by John von Neumann under the protection of J. Robert Oppenheimer, who was then the director of the Institute. Von Neumann and his team made this beautiful machine called the Electronic Computer Project which had internal registers.

It did not have patch cords on the outside. They convinced IBM to cut apart a punch card machine so that they could get programs into this thing, and it was the machine everybody copied, and it became known as the von Neumann architecture, which we're still with today.

One of the interesting things about this is that the third or so full-scale program written for this machine was an artificial life program by a fellow named Barricelli, who appeared at the Institute in early 1953. By the summer of '53 he had what he thought of as "numerical symbioorganisms" running around in this five kilobit memory made up of vacuum tubes and cathode ray tube storage. He had these "critters," he would call them, running around. I think it was a lot like Conway's Game of Life from the 1960s, or Chris Langton's cellular automata from the '80s. This thing was running on this first computer in the summer of '53. That's fascinating, because this interest in creating whole ecosystems of lifelike entities and studying them and looking for signs of authentic evolution, genomes, body plans and things like that has been a fascination of computer people ever since.

TB: The von Neumann point I just wanted to raise. Bruce and I have a series of divergent thoughts that will come out as we both narrate this history. Can you give a definition of what a von Neumann machine is, and we can move from there?

BD: Some of von Neumann's writings on this talk about it as the "contingency architecture." What they could get working was what they had in the late '40s, basically. They would load low-level machine language instructions in memory — they may have come off punch cards or a drum or something — and then they were put into registers, which were fast memory mechanisms, and then the machine goes at the instructions and goes and fetches the data from a short- or long-term cache somewhere, does its little operations and then puts everything neatly back in place and goes to the next instruction. It may be branching across this list of instructions or just going through them, but that's the basic idea of the von Neumann-type machine: serial processing of these instructions as fast as you can do it, and that's basically what we do today.

TB: There's an element of certainty in the von Neumann machine and that is critical in the localized idea of what a von Neumann machine is, but I think what fascinates me with modern computation, particularly if you look at a series of these things running in parallel, potentially accessing the same information, potentially networked; the boundaries associated with predictability disintegrate relatively rapidly, even if you have, for example, two of these processes in a dual core or a quad core machine and they're accessing the same memory. There are certain race conditions, which means that basically they are all trying to access, read or write off the memory space, which eliminates some of the certainty. Now, the interesting thing with regard to this idea of certainty is that you then — ideally, in most cases — program in the certainty to make sure you don't get into any of these race conditions, but there are some operations, particularly associated with network processing, where you do actually want to play into some of these elements of the race conditions. So it's interesting that you say that we're still dealing with von Neumann computing even to this day, because I would certainly argue against that. This is a nontrivial claim. It's been very curious even raising this in the circles associated with the artificial life community because there is a kind of traditional community that still thinks of modern-day computing as being very von Neumann, but I certainly don't consider modern-day computing, through the issues that I've described to do with parallel computing networks and these kinds of things, being von Neumann. Are you sympathetic to that, Bruce?

BD: Well, when we did the programming for the EvoGrid with little thousand-atom volumes, the horrifying news is that nature is running all thousand of those atoms at the same time and they're all slamming against each other, and this is known as a dissipative system. Stuff is trying to move around and if there's a hot area, the atoms are generally moving to a cooler area, or if there's a denser area, then they're moving to a less dense area, and reactions in the cell are happening because matter in the cell is just slamming into the other matter in cell, and binding sites are found by this stochastic, almost random walk process, and that's how nature works. Whereas when we try to simulate nature, the reason it takes so long is that we have to line up all those atom-to-atom interactions, and even if we have scalar pipelines, we have to cue them up and simulate them through this bottleneck. We don't have a computer per atom and we don't have a big computer that can control the whole volume, so we're doing our best but we're stuck in these bottlenecks. Otherwise, it wouldn't take hours or days or weeks to simulate these very small volumes. The folding of a protein may take weeks, whereas in nature it's happening in a few nanoseconds or less than a microsecond, so that's all the bottlenecks speaking there.

TB: This is again where you and I disagree, because my view is that we just don't have the correct model yet — that the computation is there, and certainly the atomic computation is there, associated with actually dividing this information into vastly parallelized computation — it's just a matter of finding the correct model to take this information and optimize it for processing. But returning to the idea of what artificial life is in the context of Terence McKenna and the emergences of McKenna through this discussion: in terms of the broad and written-down history of artificial life, it started with a fellow called Chris Langton, who basically gathered together a group of similar thinkers and formed a conference series, which is traditionally the way these things happen in academia. What happened through this conference series is that a series of books were written in parallel to this — some highly scientific, some philosophical and also some popular — and what you saw through this was almost a grass fire of interest in the artificial life idea in the late '80s and early '90s. Certainly, this is your and my story, and the story of a number of folks in the Biota community, getting their hands on one or many of these artificial life-related texts and starting to think about life as it could be. Through this there were around 100-200 folk that started creating relatively long-term simulations. Here we're talking almost exclusively about computational simulation, although some of them were robotic simulations as well. For example, Stephen Levy's book *Artificial Life* perfectly captures that. The idea is that you have these simulators that created their own particular stakes in computational simulation, and Bruce's original work in this was with regard to the simulation of plants — can you talk a little about that, Bruce?

BD: Yes, the first project of Biota.org was to "grow" plants from what are called Lindenmayer systems, and do it on the web through Java and VRML. This started in '96, and we did an exhibit at the Electric Garden at SIGGRAPH in '97; it was perfectly timed and titled. People would extrude these plants and then put them on a virtual island. That was the idea of using the procedural power of encapsulated "genomic" information to make a beautiful structure that everybody would recognize. Somewhere in the genes of real plants is something like a rewriting rule, something like an L-system, because you can create simple L-systems that will grow out the three-dimensional structure of a mustard plant, for example, with all the flowers in the right place and the right leaf shapes, etc. So, somewhere in nature there is code that does this, and that was our first project.

TB: In parallel to the formation of Biota, in the early '90s I wrote antiviral software, and at the time I was also writing a book called *Field of Chaos*, which I recently published. Through that antiviral simulation I developed a lot of extraneous code as well, associated with creating landscapes and ideas of early cognition and these kinds of things, and gathering these ideas together I created the Noble Ape simulation at about the same time that you were doing the first Biota project with L-system plants. Noble Ape has continued on to this day, it's something where various thinkers have come through and added their own little pepperings to it. The fundamental idea behind Noble Ape was to create a rich biological environment through simulation with plants, animals, birds, trees, insects, a wide variety of different things, and obviously the underlying landscape and these sentient ape-like creatures that wandered over the environment. Early on it was very primitive, it had black and white graphics and very simple visualization, but from that I was able to gather together a wide variety of folk that had interest in particular aspects of Noble Ape, and really it's been the legacy of Noble Ape going on that it's existed in software, philosophically and also with regard to these various intellectuals and others that have wandered through the simulation. In terms of Biota, how would you describe Biota to someone who listens to the Salon?

BD: Biota is a community center, an organization, a conference and now a podcast to make a discussion happen around, "Can we simulate nature in computers?" "Can we solve problems of emergence?" "Can we create systems that complexity science can use as an experimental laboratory?" "What are the philosophical and ethical implications of true artificial life should it ever emerge?"

Biota has been that since 1996, when it was founded, and we've had four conferences and hundreds of podcasts. Some from the old folks in the field from the time of Chris Langton and the artificial life conferences, and we also encourage hobbyists and new people who are entering the field, basically connecting them with the legacy and the history of the medium so that they can advance a lot more quickly and not reinvent the wheel that others have done before, and also connect them with academic and industry people who have the same fascination with these worlds.

TB: It is a fascinating idea that the hobbyist, someone who is just tinkering away at their own particular simulation, can actually have an impact on the academy. Developing Noble Ape early on, I always felt a sense that I wasn't doing anything that could contribute to a broader academic discussion, and it was only from my experiences with doing academic publishing that I realized that, having a project that has gone on for more than 15 years now, as Noble Ape has, actually the reverse is true — the academics wish they had been doing this kind of thing for such a long period of time! The beauty of the artificial life hobbyist is that they are independent of funding sources, so their passion and their enthusiasm and their drive associated with creating these very rich and interesting systems is independent of their ability to have a day job etc., so it has an intellectual longevity as well as an actual longevity. Linking the psychedelic community and the artificial life community, I was talking with Douglas Rushkoff a few months ago and he mentioned that he always saw the two communities as closely aligned, because it related to a familiarity and acceptance of alternative worlds and a sense that what is presented to us as the immediate world need not be, or should not be, the way that it is. This is very much the "life as it should be" of Langton. There are ebbs ands flows, there are people in the artificial life community, certainly Tom Ray and yourself, who also attend psychedelic conferences, and it has always struck me that there are skill sets that are needed from both communities in order to operate. This echoes your thinking, I guess.

BD: Terence and I sat down one evening at his house in Hawaii in February of 1999. As you know, Terence has this idea of the universe as a novelty-conserving engine. Terence talks about what is obvious to everybody, which is that complexity rolls forward and doesn't reverse, so you get life. Terence, in his psychedelic journeying, was interested in this concept, and then developed this in his waking hours into a whole set of interlocking theories that ended up with the concrescence of all this in 2012, that somehow we were accelerating through this novelty Time Wave and we were going to reach an infinite speed, almost like the inflationary model of the universe. I can comment about that a little, but what was interesting that evening in Hawaii is that I spent the better part of an hour with Terence trying to describe to him how the internet actually worked and why it wasn't a good environment for what he was talking about. The internet was too arid, it was too thin; there was no one actually working on the problem of creating an AI that would somehow rear its head in cyberspace.

The environment wasn't conducive for even a simple artificial life form, let alone a full-blown AI; this was all pretty much in the realm of fantasy. Terence opened the discussion with, "I don't suppose you're one of the people who believes that the AI will emerge and within twenty minutes we'll all be taken over." I said, "No, Terence, I do not, and let me explain why," and I believe that Terence and many other people are extremely naive: they see a virtual world, or a screen-saver with fish swimming on their screen, and they misinterpret that. They think that that is something that will one day be alive, but it's a chimera. It's truly artificial, it's truly limited. It's like shadow puppetry, there's no depth to any of this. If you start to study molecular biology and how things actually work in the cell, you realize that our computing systems are pathetic.

I recently joined the Astrobiology Institute at NASA Ames as an associate, and I had a conversation there with Andrew Pohorille, who is a well-known NASA researcher who has a lot of funded projects in astrobiology and the origin of life. He said, "What the people in the computer science side don't understand is that the palate they are working with to do emergent phenomena, self-assembly etc. is so limited in code compared to nature's palate. They think it's the opposite; they think that they have many more widgets and tools, but I posit that nature's toolset is much broader." There is this divide between those who work in the computer simulation side and those who work in chemistry. It's a divide that we saw at Digital Burgess in 1997, where we climbed up to the Burgess Shale in Canada. The paleontologists who were there described themselves as detectives after the greatest crime in history, the emergence of complex life through the Cambrian explosion, whereas the computer scientists were world makers, they were little microgods or nanogods. They could create their own perfect little worlds and watch the things that are happening, but no biochemist or paleontologist would really take them seriously, because they're staring into the full monty, the full, complex, seething, biochemical, multi-billion-year history. All the stuff coming out of computer science labs looks interesting, but it's ultimately not that useful for people dealing with the physical world. They would pat you on the head and say, "Very nice, maybe we can use this to figure out how trace fossils are made."

So, I was really trying to rein Terence in. Terence became more extreme over time; in the '80s he was talking about UFOs, in the '90s he was talking about AIs, and as the internet grew, as he got his Mac connected to the internet and he started seeing things, he really extrapolated way too far. He was reading too many science fantasy books, and he was talking about the AI rising and taking us into its folds within twenty minutes. Terence was fascinated by artificial life, and in fact the EvoGrid project is, in some ways, a novelty-conserving engine. It's one of the first to be built in software with a serious goal and measurable results. If Terence was here, he would be fascinated by it, and in fact it might bring him a bit more down to earth, because it shows how hard true novelty really is if you're trying to simulate nature in a computer.

TB: Here, Bruce and I disagree. Starting with the origins of Noble Ape and McKenna, at points in developing Noble Ape I had experiences with what he's describing, associated with primate psychology, primate evolution and the evolution of communities. So, independently having developed Noble Ape over 15 years, there have been a number of links between what McKenna has said and my own experiences developing Noble Ape. It's relatively difficult to describe the kind of intimacy that one needs to have with a simulation over such a long period of time, and particularly when you start looking at a simulation in terms of novelty and emergence. Emergence here means something that surprises you, but it also has a quality associated with how you can get the information from the simulation. Artificial life simulations are traditionally so complicated that it's almost impossible to get all the information from the simulation environment, so you have this kind of distancing of the kind of information that you can get from the simulation, and you have to find different ways of interrogating the simulation to try to get some of this information out. One of my favorite quotes from McKenna that resonated with Noble Ape — I think it was a quote of one of his thesis supervisors — he said that if you looked at the world as a world made by angels, it was very disappointing, but if you looked at it as a world made by monkeys, then it was truly amazing, and I think that this has been my experience developing Noble Ape. This idea of paranoid monkeys coming down from the trees with a hyper-sense of fear and long-term desire, charting a path through the environment; this is really the narrative of Noble Ape.

So from doing these simulations for long periods of time and looking at social evolution through the context of simulation, I started to realize that there was a gross lack of both social science and hard science data associated with things that were traditionally outside the boundaries of science. The thing that interested me in doing artificial life simulation was that you could construct simulation environments that weren't necessarily possible in the physical world or even in social science experiments. The idea of the interrogability of the simulation, and also the observability of the simulation, led me to start looking at things in the outside world in terms of them being potentially simulations as well. The idea of information that can exist both in deltas and in clouds rather than in actual interrogable points is a very interesting thing that came from this broader idea of artificial life simulation. For example, information in deltas means that you can't actually stop something and find something at a particular instance because it always exists in a flow, and "in the cloud" means that it can exist in multiple points rather than in a single point. Physics has a means of interrogating deltas, but understanding clouds is more difficult in terms of the spreading of information, and so the need for a new kind of science has really come through the simulation community, distilled in this idea of simulation science.

Returning to McKenna, the concept of AI as Bruce described it, this notion that artificial intelligence will in some way resemble human intelligence or needs to resemble human intelligence is one of the easy paradoxes to squash with artificial life. What you find with artificial life simulation is that there is a continuum of intelligence. In fact, Roy Plotnick, in his appearance on Biota, stated this best, the notion that intelligence is on a continuum that starts with survival. When the earliest organisms started eking their way toward feeding grounds they had a view of survival that emerged into intelligence. This notion of intelligence being on a continuum, or potentially a multidimensional space, and that we are merely one point, a small cloudy point in this environment. So Bruce and I differ very strongly on this idea of machine intelligence and what it will actually look like. In the context of the singularity movement in particular, I've often described my own views as being post-singular, in terms of this notion that what we have in computation is distinctly different than human intelligence and in many ways is vastly more powerful than human intelligence. In this regard I'm more of a McKennaphile than Bruce is.

The financial system or the legal system are these amazing environments that vast numbers of humans have created that are distinctly larger than human intelligence and also have a far greater temporal survivability — obviously, we all have finite lives that we live as humans — but also have the ability to self-maintain and to enact terrible things against humans. So my view is post-McKenna in that light and also reinforced by years of studying eclectic and diverse simulations. The Noble Ape simulation isn't just one simulation, it's a series of simulations that have been layered progressively. One of the curious things I found with McKenna is that he had an aversion to weather simulation; the weather simulation has actually been one of the most interesting aspects of Noble Ape in terms of the effects of communities and structures based on simulated meteorological effects. I guess that brings us up to AI in the context of human intelligence, which is the very narrow definition of intelligence defined by the singularity, versus the concept of intelligence where we are only one exhibit in the cloud of potential intelligences and where things like the internet and various other forms of machine intelligence are vastly more powerful and can outsurvive us. What's your view with regard to the Plotnick notion of survival intelligence and how we relate to that?

BD: I think that we vest far too much in our technology. If you really look at the internet as a whole, even the banking system, all this stuff is hand-built and hand-maintained by millions of people. Apart from very, very good Unix and Linux servers, if you left these systems without somebody poking and prodding them the whole internet would go down pretty quickly. It's almost as though in the 19th century people thought that the telegraph system was going to create a supermind, or that the steam engine rail system was bigger than we could comprehend and therefore you'd get Frankenstein coming out of the steam engine and electricity system.

TB: How do you walk away from the financial system in terms of its complete computerization? You've described these things as computational systems, but they're not just computational systems. They exist in a philosophical realm which is completely independent from the physical nature of the system. This is one of the things that has interested me about Rushkoff's more recent writing: that the ability to move away from the financial system is not just about no longer maintaining the computers that run the financial system, because it exists as something which is greater than the computers.

BD: Well, certainly the traders on the floor, the bankers, the policy-makers are constantly reacting to the dynamic that is rippling through the financial markets. It's very hard to predict what's going to happen. People poke and prod it, and then a wave is set up and currencies crash and all this stuff happens, but I think it's not an intelligence; it's just a great big hairball mess of stimulus-response. It's a huge finite-state machine, in a sense. We've built this huge finite-state machine and it comes and bites us all the time.

A good model for this is the International Space Station. Thirteen nations worked on that thing; I worked on some modeling and simulation of it back in the mid-2000s. There are 100+ types of connectors on it and hundreds of computers running it. The crew that's up there find that there's a dynamic to it. For instance, if there's a problem rotating the solar panels because there are metal shavings somewhere, it has this ripple effect that goes through the whole station and you have problems maintaining all kinds of equipment that will go out if power levels drop. It's just a huge complex dynamic that we've made. When we were inventing agriculture we learned about that, we learned about living in a complex, dynamic system: whether to flood the fields at a certain time, the spreading of seed and not weed, how to keep ergot mold from getting to your grain supply. Human beings have been living in these complex systems for a long time, we're just making them such that now we really can't drive them. We're victims of the system that we have made, and maybe we're doing this to the whole planet, but I don't think that there is an AI that will emerge, outside of a whole lot of reacting and maintaining and fire drills. I don't think that the system we've made is an autonomous standalone AI of any sort.

TB: My concern with your definitions is that you're defining AI, intelligence and the human brain in terms of what these things aren't, but the thing that I am interested in is what these things *are*, and how we can understand these things in terms of ourselves. If these systems are destructive, if they are polluting, if they are doing things which are causing us very fundamental problems, there is still no intellectual control, there is no understanding. There is no means of disassembling these systems or of extracting yourself from them. My criticism of Rushkoff's work and a lot of theorists that are coming through the dissection of the financial system is that there is no meaningful way to remove oneself from this environment. My discussion associated with being post-singular motivates intellectuals into starting to analyze this and starting to grapple with what they actually are, rather than abstract ideas of artificial intelligence that neither map onto the financial system or human intelligence. It seems to be a strange kind of argument to say that these things aren't artificial intelligence and they aren't intelligence. They may appear through particular analytical methods to be a kind of mess, but they are a vastly controlling mess which seems to be greater than can be described by individuals; they may be initially enacted by individuals, but I think that no one individual can intelligibly understand them, which is a property that is really critical. Not even large groups of individuals can understand them, so the whole notion of collective or single-intelligence humans being able to dissect these environments seems to be relatively flawed.

My concern is that I would like to see an emergence of intelligent folk who are able to start breaking apart these things in terms of what they actually are, rather than using negative terms against them. Without that kind of analysis we'll never be able to extract ourselves from these things. More importantly, these things are evolving — let's use that term — at such a vast rate that the previous terms, the previous means of explanation, even the ones that are emerging through the singularity and so on, are fundamentally flawed in terms of their analysis. I think there may be elements of McKenna that could offer some degree of insight into some of these things. What do you think Terence could teach us about how we can actually analyze these things? You talked in terms of Terence going too fast in a particular direction without connection to facts, but in terms of the inspirational nature of McKenna, which is motivating both of us actually having this conversation, what elements of McKenna may be useful in understanding these contemporary problems?

BD: I think that McKenna's concept of infinite novelty and the Time Wave and 2012 should be stripped out of this analysis and put aside. I feel that McKenna's own life was accelerating so quickly that he couldn't imagine his life in 2012 and it was really a personal eschaton that he was going through, but beyond that, he's no longer with us and he can't comment. I think that what he would say is that the psychedelic state that you enter, or that people do enter, allows you to see astoundingly complex realities, whether it be the elf machines or seething geometry or the entire planet from space, and that the human mind is actually quite capable in these elevated states of grokking a massive amount of interlocking complexity. At the same time — and this is very important, this comes from Leary and all these people who tripped in the '60s — you see the game. You see the mundane reality when you land from the elevated state, you see the game that people play and you realize that people are operating on very simple principles: stimulus-response, "me and I want to get mine," or, "I'm afraid of this."

Most people, even in institutions, governments, financial or service companies, act almost in a kindergarten-like way. There are a few devious ones among us who are gaining financial influence and really gaming the system, but for the most part you can almost predict what they're going to say at the next moment, you can predict what they're thinking and where they're going to go, because they're in this game which is relatively simplistic. The psychedelic experience reveals these games. Even up to the level of sitting in the Oval Office, these guys are operating on poor information, poor assumptions. Some of them have a vision, some of them have a fundamental understanding, but very few do; they're mostly just reacting. Within governmental or military organizations very few people are thinking outside of the box or coming up with new things. We're kind of carried along in this flotation device. The psychedelic space allows you to see this space, and with that information, the information that people are operating on relatively simple needs and simple scenarios, you realize that the system can collapse easily, but it can also be changed relatively easily, because it's not very sophisticated.

Steve Jobs talks about doing a trip in the '60s or early '70s and realizing that his mind was capable of changing the world and of making an entirely new world. He founded Apple Computer based on that, and he has continued to do that. The people who have had their minds opened up to the greater reality that they see are the power sources, they're the people who can make the future. They see the game being played. Some people who are opened up to this reality go back into what Terence would call the mundane reality and they get alienated, they check out; it's to their disadvantage because then they become financially disadvantaged and disempowered. Terence was one of the few people to ever make enough money to live on talking about mushrooms. As he himself said, this is no way to build a career. So, for the most of us, if we see these visions and we come back, we have to integrate, but we can do powerful things in the world with these insights. We can see the game for what it is, and we can not get panicked about what people say on TV and we can stay outside of panic-of-the-day discussions, because you see it as a game.

The psychedelic experience gives you a long view of history. Terence read widely in history, not just alchemical history but the history of technology, of the arts, so in the mode of Terence I would say: OK, we're worried about the growing, seemingly out of control complexity of the financial system: roll your clock back to the 1870s. Railroads have spread across various continents and there is no control mechanism, so you get these head-on collisions of trains, you get complete bedlam. They solved the problem: they invented telegraphy and a whole system of coding to keep these trains from slamming into each other. Roll your clock to the 1920s, when there were hundreds of telephone companies. It was complete bedlam trying to call between one town and the other when your phone company was only in your town. They moved a step up and by 1929-1930 they had direct dialing where you didn't have to wait for somebody to call you back and connect the line. So, we gradually mastered the complexity of a new medium that seemed to be out of control and would never be solved. I think we're doing this again. We connected computers together with money, with trading, with news and we're in this boom of complexity and there's a joy in it but there's also a panic over it overwhelming us, but I think that soon these systems will be subsumed again, they'll be part of the woodwork and they'll just function. The engineers will figure it out, the regulators will figure it out and we'll be into a new era, a new complex system on top of the old one, and we'll think back twenty years in the future and say, "Why did we worry so much about such and such?" because it was a problem that would get solved.

TB: An important point that you touched on is: I've done this relatively nihilistic, negative rap associated with this on our ability to understand our technology. The secret that we have is communication, and this is interesting in the psychedelic community because "find the others" is inhibited by various legal aspects. We're both students of computer history, we both love the rich narrative that computer history provides; you have a shrine set up in your backyard and I've maintained a lot of it in my head and on my bookshelf. The idea of

Steve Jobs is a good one, because Steve Jobs didn't create this stuff alone. He was able to communicate with groups of individuals and actually get a swell of engineering talent to come through and devote their 16 hours a day, 7 days a week to this dream as well. The power that we still hold as humans, versus machines, is the ability to put ideas out there, to communicate, in dissecting and solving these issues. I think when you talk about "the game," and the element of a relative simplicity, the game is still maintained by an environment which is very protective, you look at how the financial system can work against individuals that want to step out of the financial system. The way that we can solve that is through communication of ideas, and getting together in groups and having discourse and actually creating something which is independent of the environments that we see.

BD: Yes, absolutely. Another thing to touch on is the machine elves. If you talk about Terence McKenna you cannot *not* talk about machine elves. Terence described the mushroom experience as being extremely high-tech and futuristic, spaceships around distant worlds, these machine elves, etc. For him, at least, the mushroom experience was super high-tech. So, does that experience give us any insight into how technology is raveling itself around each of us? Are we all becoming machine elves?

TB: The idea about this that I've had is that the machine elves exist as the simulated Others. Certainly in terms of the Noble Apes communicating, I've thought about machine elf communication in that light. The idea that you have entities that are fundamentally internal, but represented as external entities, communicating in a way which seems to be unintelligible is ultimately a computational problem. My own view is that the appearance of the machine elves is analogous to the language of artificial life: the language of communication of these entities, but also an understanding that perhaps this is another form of language that we hold, that we can represent in another computational form... the need for these entities to communicate in a way that is potentially unintelligible to us but certainly is intelligible to these externalized entities. That's been my view from an artificial life perspective, very much embodied in my own machine apes and their existence externally.

BD: These worlds are so fantastic, there's that whole rap that this is communication from beings from some other world and that we can learn from it, that we can build technology based on it, so you could build a super- ultra-Noble Ape from something you saw while taking mushrooms. Or maybe not. Scientists have reported that they have seen the structure of chemical compounds that they could never have otherwise seen. This is certainly not in my experience, this is not something that I seek out, but the question is: can you use the psychedelic state as a tool to create wilder and woolier technologies, or to understand the fuzz of technology that is enveloping the earth and to see the threat or benefit of having this fuzz of technology all around us? Is this going to clothe us and keep us warm and help us survive, or is this going to suffocate us psychologically? It's hard to say, but perhaps the psychedelic state is the place to ask the question.