Francis Heylighen with David Sloan Wilson, Part One

David Sloan Wilson: So Francis, I often reflect upon how much has happened in the history of ideas since the second half of the 20th century which couldn't have happened before. I think that's true for evolutionary biology, my specialty, and also for complexity which is your specialty, and which among other things requires just the advent of widespread computing before we had the tools to even think about it or model complexity.

Because you go way back on this topic, I'd love for you to begin with just your own personal story as to how you became interested in complex systems, and against the background of really historic development in ideas. I think in some ways the enlightenment is still in progress. So, please just tell us your story against the background of these big ideas.

Francis Heylighen: Well, when I was in high school, actually I already started thinking about these things and I had even formulated what I called the generalized principle of natural selection which is that evolution to variation, selection is not limited to biological species but any system, physical, psychological, cultural, undergoes variation, selection.

So, I had this very broad train of interest but I didn't know very well what to study. I settled on studying physics basically for the reason that it was the most difficult discipline. So I thought, "If I can do this at the university, I can take all this on my own." That is more or less what happened. So I studied theoretical physics which is pretty hard but if afterwards, I was interested in psychology or in biology. I could do it on my own.

At the same time when I started studying physics, I did it with a kind of critical attitude in the sense that I was from the beginning skeptical about the traditional reductionist, deterministic, mechanistic culture. It's what I call the Newtonian worldview. At that moment in physics, there wasn't yet anything like complexity science. So, I started looking outside of physics.

What I discovered that was closest to the things I knew was systems theory and cybernetics. So in systems theory and cybernetics, there was an attempt to make a more mathematical theory using some methods from physics but without the assumption that everything is reducible to particles moving in space, the way the traditional physical worldview is.

It looked at systems and these systems could be biological systems, social systems, cultural systems that were connected by inputs and outputs, let's say systems theory point of view. And then cybernetics added to that, the idea that these systems could be goal-directed, that they were kind of like agencies. That they weren't just possibly receiving some input and transforming it into an output, that they were trying to reach certain goals. So that was a framework that to me sounded very natural, and which I recognized myself.

I also came into contact with a couple of colleagues working in that who actually in cybernetics were longer at that time than I was, and in particular came in contact with two guys, Valentin Turchin and Cliff Joslyn who had started something. It was in the very beginning stages, they called it the Principia Cybernetica Project.

So the idea is cybernetics and systems theory have all these great ideas. In principle, these great ideas allow us to unify scientific disciplines with the idea of general systems. It should be a theory that applies to any kinds of systems. In practice, the domain is anything but unified. It's kind of a mishmash of all kind of different approaches, these different kinds of theories.

It was of course very ambitious, was a little bit like the Principia Mathematica of Russell and Whitehead where they tried to create a foundation for mathematics. So, we wanted to create a foundation for cybernetics. We called this as Principia Cybernetica Project, and I assume we made some quite good progress in the sense that we made a website where we formulated the basic concept, the basic principles that we organized a little bit like a semantic network.

It was one of the first serious websites in the world. It's 1993 actually. So, it had a lot of influence. A lot of people learned about these things from our website, but that's not yet what is conventionally known as complexity theory. So systems theory and cybernetics in principle are about complex systems but in practice, without computers or the models who could make mathematically. They had to be pretty simple. So, it's early in the 1980s that with the Santa Fe Institute, people started to model complex systems with computers. That made me also see that the traditional cybernetics and systems paradigm lacked something.

The traditional cybernetics and systems paradigm is kind of you have a well-defined system and you have a control mechanism on top of that system. On top of that control mechanism, you can have another control mechanism. So, you build hierarchies of systems within systems. All very fine, but the assumption is that the system has a clear boundary. That there is this one system, and that system has particular subsystems.

What we now saw in this new paradigm which is best called complex adaptive systems. I think the guy who was most responsible for that is John Holland, who was also associated with the Santa Fe Institute, is that systems consist of agents and the agents themselves are pretty simple. If you put lots of agents together, they have all these kind of very complex behavior.

There are self-organizations. There are positive feedbacks. There are negative feedbacks. There are evolutionary processes going on and that led to domains like artificial life where you try to simulate a whole ecosystem of evolving organisms. It was a very interesting approach to complexity that I would say complemented the one of cybernetics and general systems theory.

DSW: Let me cut in here. This is all great. I just wanted to drill down on a certain point before you continue, which is the fact that former analytical models are both enabling and constraining. On the one hand of course, they're so much more precise than verbal modeling. So they were huge events over verbal modeling, but they have to make so many simplifying assumptions that they end up being a de facto denial of complexity.

When I talk about this, the difference between the two-body problem and the three-body problem in physics, you hit this complexity wall beyond which formal analytical models can't go. So this embracive mathematics, formal mathematical models, for example in the whole field of economics, which is if you can't model it formally, then it doesn't have much of an impact. That's the sense in which computer simulation models were able to go beyond the complexity wall, and yet they were definitely frowned upon.

James Gleick in Chaos, I think made this point is that there is such a prestige and hubris associated with formal proofs that computer simulation models initially were looked down upon as a kind of a second class citizen. So finally now we've realized that there's no alternative to that. So, that's my interpretation. I just wondered if you agree or if you'd like to elaborate upon the account that I just gave.

FH: Well, I agree but I want to elaborate which is that the computer simulations themselves are also reductionistic. One of the reasons I said why I started to do physics but with a critical stance was because I intuitively knew that these reductionistic, deterministic models, that they weren't right. They couldn't describe something like let's say living organisms or societies.

Even in physics itself, a large part of the history of physics in the 20th century is showing that all these deterministic assumptions which we thought we had at the end of the 19th century don't work. Quantum mechanics with the uncertainty principle, relativity theory saying that even space and time are not absolute. Chaos theory, of course showing that the slightest nonlinearity in your system and all kind of thing become, in practice, unpredictable like you said the three-body problem.

What physics basically has been doing in a way was to prove that it couldn't solve a lot of problems. That was what happened with other impossibility proofs like the famous Godel, or the halting problem in computer science. So the mathematical models themselves, they show that they were incomplete. So, computer simulation then was possibly an alternative. But for the computer simulations, you also have to make a model. At least you have to make a simplification.

Now, the great innovation for me of this domain of complex adaptive systems, I prefer to use that term rather than complexity science because in complexity science, different people attempt to look at the different aspects and some people will speak about chaos theory as complexity science where chaos theory is not really complexity science in my view.

So complex adaptive systems, the greatest insight for me was that we can model things by having these agents. Agents are kind of little bits of programs that are programmed with what is called the condition-action rule. In this condition, the agent does that. In this condition, it does something else. Now, you just throw a whole bunch of agents together. Each agent, perceives some condition, performs an action. That action changes the conditions not only for itself but for all the other agents. So you now get all this kind of direct and indirect interactions between the agents, and that leads you to all kinds of very interesting results.

I already gave the examples of ecosystems. In economics, of course where the agents that are people who buy and sell things depending on what one agent sells and other one may buy. The prices may increase and you may have all these kinds of effects on the stock exchange. That insight is a very interesting paradigm that illuminates a lot. Of course, it depends on what do you choose as your element. What are your agents and what do you program your agents to do? What people typically do is they have some kind of a simplistic view of what the system is.

They program their agents. They see that maybe the simulation doesn't do what they want, so they play a little with the rules until they get something that looks right. Then they say, "Well, we have proven that this and this will happen," but actually they have just proven that this simulation of these agents with these rules does that thing.

DSW: Let me take my turn, Francis. I'm really enjoying this conversation. So, here is my turn. I'd like to spend more time on the concept of complex adaptive systems and actually distinguish a number of varieties. Let's begin with Conway's famous Game of Life. In this case, they're not actually not agents. They're just positions on a grid which can exist in an on or an off state. Then we put in simple rules as to for that on and off state basically based on neighboring cells, and their on and off state. So you put on those simple rules. Then out comes this amazing diversity of outcomes.

So, there's your point about basically agents with simple rules producing this amazing diversity of outcomes. But there’s nothing adaptive about those rules. I mean, those aren't evolutionary rules. These agents aren't organisms in any sense of the word. So we can make the point about agents with simple rules of behavior and then that leads to all of these amazing emergent properties.

We should not use the word adaptive yet because there's nothing about those agents that's adaptive. That's just an arbitrary set of rules which are applied by the modeler.

FH: It depends. The Game of Life is of course the most deterministic kind of system. So a lot of simulations have been using so-called cellular automata of which the Game of Life is one variety. I am personally very skeptical about whether cellular automata can teach us much about the real world because they are so simple and deterministic, that really the only thing they prove is that even with simple deterministic goals, you can get very complicated things. That's was an important insight. Once you're there, you need to go further.

Once you start to speak about agents, you can have different types of agents. Of course, you can have adaptive agents. That means agent that for example undergo variation, selection. That's the classical work of John Holland. He has rules that are formalized as things of ones and zeroes. Then he randomly changes some of these ones and zeroes, and he has a selection criterion. Then, the agents with the right rules may survive and the other ones get eliminated. That's one way to make them adaptive.

Actually, you can have adaptive processes even when your agents themselves are not adaptive. That for me personally, that's the most inspiring idea. You put a number of agents in an environment. The agents can change the state of the environment. They can do things. For example, suppose you have an environment and there are different sources of food in there and an agent can either eat a piece of food, or transform it into a different piece of food, or produce a piece of food.

The adaptivity is in effect that each agent has its own rules deterministically, but it doesn't know what the other agents are doing. The other agents have their own deterministic rules. Because they are in the same environment, what the one does affects the others but you don't know a priori, how that will happen. I think the novelty comes from the fact that you have independent agents, each of which is deterministic but their interaction is not deterministic. The interaction has to adapt to what all the other agents are doing.

So, it's a kind of a model of coevolution. Each agent does whatever it's programmed to do, but to do that, it has to take into account what the other agents do. And that changes all the time. So, you actually get the very adaptive systems. That is why I think the term complex adaptive system is not badly chosen.

DSW: There's two meanings. One meaning is, is that the agents within the system are following adaptive strategies. The other meaning is, is that the system as a whole is adaptive in some sense. So, I'd like to get your opinion on that distinction. If we're going to use the word adaptive, basically adaptation can exist at two levels. The agents within the system or the system as a whole.

In any practical sense, we're trying to create systems that are adaptive as whole systems. That's any kind of policy objective is to create whole systems that are adaptive— economic systems, environmental systems, social systems. So, how do we get from systems composed of agents following adaptive strategies, their own adaptive strategies to systems that are adaptive as whole systems?

FH: As I said you don't even need to have adaptive agents to have an adaptive system because the adaptive system is the whole of all the agents evolving and mutually adapting. One strategy that helped this whole adaptation, and that is I think one of the works that I'm most proud of. I have written several papers about it. I didn't invent it, but it is the concept of stigmergy.

What is stigmergy? Stigmergy means that an agent does something that leaves a trace in the environment. That environment is shared with the other agents. These other agents can, if you wish, read the traces left by previous agents, and then they will build on that. They will react on that. In another conversation, we spoke about Wikipedia working according to this principle. The different agents here are people who each write a little bit of texts in Wikipedia. These traces remain in Wikipedia. They are public, and they incite other people to add something to that.

The same applies at all of these levels even with very simple stupid agents and the classic example are the ants that leave pheromone traces. These pheromone traces then tell other ants what to do. The even simpler case that was the one that gave the name of stigmergy are the termite hills. Termites, when they build a cathedral-like termite hill, they just tell by dropping a little bit of mud here or there.

The mud contains pheromones which attracts other termites and thus termites tend to add their mud where there is already the most termites. So, you get a positive feedback. The higher a part with mud is, the more termites it attracts. It's again, this principle of stigmergy. The trace left by the activity of one agent, one termite, incites other agents to add to that. Adding in this case is pretty obvious, it's just doing more of the same.

In the case of Wikipedia, the adding means changing, correcting, whatever it is, the trace in a shared environment, now provides a kind of a template that is constantly being changed by the other agents like that. We have a self-organizing system that can be very adaptive and very coordinated even though the agent themselves are stupid.

DSW: If we were to understand the mechanisms of development of a single organism in enough detail, we find something just like that. This brings up the concept of the organism is really a central concept and metaphor because with an organism, we have something which is manifestly adaptive as a whole system, in which everything that makes it up is of course contributing to that adaptive whole but in a way that is very distributed. As you've written, even the brain is very distributed. There's no homunculus within the brain. It's all distributed.

What we're working towards, especially with The Human Energy Project is the concept of the Noosphere and ultimately the whole earth is some kind of superorganism. Before we can talk about superorganisms, we really should get our idea straight about organisms. That's our master kind of a concept. So let's talk about organisms.

What you just said, you said this whole system can be adaptive without the agents being adaptive. Is that true for an organism? Can you make that statement for an organism? In what sense are the elements that make up an organism, the cells, or the organs, not adaptive in their contribution to the functioning of the whole organism?

FH: Of course in an organism, the way we know it by the logic of organisms, they are a result of many different evolutionary transitions where each time you started with a system that had some degree of adaptivity. So, I'm not going to claim that the multicellular organism is only adaptive at the multicellular level. The cells themselves obviously are also adaptive.

Let's make a caricature or simplification. Let's assume that each cell is purely programmed by its DNA. So it has a DNA program that tells it, "If these molecules enter the cell, then you produce these other molecules to deal with it." In principle, that should be sufficient for the whole of this cell to coordinate. In my view, the critical term here is coordination. Each agent is capable of adaptive action. The agent itself does not need to be adaptive in the sense, they changes its rules. The agent is adaptive in the sense that if something changes in the environment, the agent will perform some action that makes the action that the agent appropriates to the environment.

The problem is if you have lots of agents, that each perform his actions, how do you coordinate them? I distinguished two forms of coordination. The one that I already described, that's the stigmergy one. The stigmergy one is you have this common medium. You drop signals in that medium. Everybody can read them, and whoever wants to react to what you have put in the medium can do so.

The equivalent in the multicellular organism would be hormones. A cell has some kind of a problem. It is programmed in this form of stress to release a certain type of hormone. That hormone goes to the bloodstream. Other cells that are specialized in reading that kind of hormone now will react by maybe releasing some other chemicals or maybe they will change some of the activities. Maybe your heart rate will go up. Or you will start sweating, all because of this one hormone that was deposited in the bloodstream. That's one way of coordinating the activities.

The other way, and that's the most sophisticated way is the one we find in the brain. In the brain, one neuron if it has let's say some kind of a problem, instead of depositing just the neurotransmitter that everybody can read, it will send a signal to particular other neurons that is connected to it. Then you get the network, and in the network, of course you wanted the right signals, go to the right agents.

Then you get this difficult problem that people in your network in theory are trying to solve. How do you get that network to self-organize so that it becomes effective in the sense that all the neurons together collectively are solving these difficult problems that single neurons cannot solve? So we know a number of algorithms that do that of which the most basic one is reinforcing whichever connection works, meaning that it did something that you wanted.

Let's say these are roughly the two prototypes of coordination. Either stigmergy, you just leave a message in a medium that everybody can read, or neural network, you send a particular message to one or more particular agents. If that's the right agents, the connection is reinforced. The next time, they are more likely to get this message.