OUT OF THE BLUE

Can a thinking, remembering, decision-making, biologically accurate brain be built from a supercomputer?

seedmagazine.com

by Jonah Lehrer • March 3, 2008

In the basement of a university in Lausanne, Switzerland sit four black boxes, each about the size of a refrigerator, and filled with 2,000 IBM microchips stacked in repeating rows. Together they form the processing core of a machine that can handle 22.8 trillion operations per second. It contains no moving parts and is eerily silent. When the computer is turned on, the only thing you can hear is the continuous sigh of the massive air conditioner. This is Blue Brain.

The name of the supercomputer is literal: Each of its microchips has been programmed to act just like a real neuron in a real brain. The behavior of the computer replicates, with shocking precision, the cellular events unfolding inside a mind. "This is the first model of the brain that has been built from the bottom-up," says Henry Markram, a neuroscientist at Ecole Polytechnique Fédérale de Lausanne (EPFL) and the director of the Blue Brain project. "There are lots of models out there, but this is the only one that is totally biologically accurate. We began with the most basic facts about the brain and just worked from there."

Before the Blue Brain project launched, Markram had likened it to the Human Genome Project, a comparison that some found ridiculous and others dismissed as mere self-promotion. When he launched the project in the summer of 2005, as a joint venture with IBM, there was still no shortage of skepticism. Scientists criticized the project as an expensive pipedream, a blatant waste of money and talent. Neuroscience didn't need a supercomputer, they argued; it needed more molecular biologists. Terry Sejnowski, an eminent computational neuroscientist at the Salk Institute, declared that Blue Brain was "bound to fail," for the mind remained too mysterious to model. But Markram's attitude was very different. "I wanted to model the brain because we didn't understand it," he says. "The best way to figure out how something works is to try to build it from scratch."

The Blue Brain project is now at a crucial juncture. The first phase of the project—"the feasibility phase"—is coming to a close. The skeptics, for the most part, have been proven wrong. It took less than two years for the Blue Brain supercomputer to accurately simulate a neocortical column, which is a tiny slice of brain containing approximately 10,000 neurons, with about 30 million synaptic connections between them. "The column has been built and it runs," Markram says. "Now we just have to scale it up." Blue Brain scientists are confident that, at some point in the next few years, they will be able to start simulating an entire brain. "If we build this brain right, it will do everything," Markram says. I ask him if that includes selfconsciousness: Is it really possible to put a ghost into a machine? "When I say everything, I mean everything," he says, and a mischievous smile spreads across his face.

Henry Markram is tall and slim. He wears jeans and tailored shirts. He has an aquiline nose and a lustrous mop of dirty blond hair that he likes to run his hands through when contemplating a difficult problem. He has a talent for speaking in eloquent soundbites, so that the most grandiose conjectures ("In ten years, this computer will be talking to us.") are tossed off with a casual air. If it weren't for his bloodshot, blue eyes—"I don't sleep much," he admits—Markram could pass for a European playboy. But the playboy is actually a lab rat. Markram starts working around nine in the morning, and usually doesn't leave his office until the campus is deserted and the lab doors are locked. Before he began developing Blue Brain, Markram was best known for his painstaking studies of cellular connectivity, which one scientist described to me as "beautiful stuff...and yet it must have been experimental hell." He trained under Dr. Bert Sakmann, who won a Nobel Prize for pioneering the patch clamp technique, allowing scientists to monitor the flux of voltage within an individual brain cell, or neuron, for the first time. (This involves piercing the membrane of a neuron with an invisibly sharp glass pipette.) Markram's technical innovation was "patching" multiple neurons at the same time, so that he could eavesdrop on their interactions. This experimental breakthrough promised to shed light on one of the enduring mysteries of the brain, which is how billions of discrete cells weave themselves into functional networks. In a series of elegant papers published in the late 1990s, Markram was able to show that these electrical conversations were incredibly precise. If, for example, he delayed a neuron's natural firing time by just a few milliseconds, the entire sequence of events was disrupted. The connected cells became strangers to one another.

When Markram looked closer at the electrical language of neurons, he realized that he was staring at a code he couldn't break. "I would observe the cells and I would think, 'We are never going to understand the brain.' Here is the simplest possible circuit—just two neurons connected to each other—and I still couldn't make sense of it. It was still too complicated."

Markram has good reason to cite physics—neuroscience has almost no history of modeling. It's a thoroughly empirical discipline, rooted in the manual labor of molecular biology. If a discovery can't be parsed into something observable—like a line on a gel or a recording from a neuron—then, generally, it's dismissed. The sole exception is computational neuroscience, a relatively new field that also uses computers to model aspects of the mind. But Markram is dismissive of most computational neuroscience. "It's not interested enough in the biology," he says. "What they typically do is begin with a brain function they want to model"—like object detection or sentence recognition—"and then try to see if they can get a computer to replicate that function. The problem is that if you ask a hundred computational neuroscientists to build a functional model, you'll get a hundred different answers. These models might help us think about the brain, but they don't really help us understand it. If you want your model to represent reality, then you've got to model it on reality."

Of course, the hard part is deciphering that reality in the first place. You can't simulate a neuron until you know how a neuron is supposed to behave. Before the Blue Brain team could start constructing their model, they needed to aggregate a dizzying amount of data. The collected works of modern neuroscience had to be painstakingly programmed into the supercomputer, so that the software could simulate our hardware. The problem is that neuroscience is still woefully incomplete. Even the simple neuron, just a sheath of porous membrane, remains a mostly mysterious entity. How do you simulate what you can't understand?

Markram tried to get around "the mystery problem" by focusing on a specific section of a brain: a neocortical column in a two-week-old rat. A neocortical column is the basic computational unit of the cortex, a discrete circuit of flesh that's 2 mm long and 0.5 mm in diameter. The gelatinous cortex consists of thousands of these columns—each with a very precise purpose, like processing the color red or detecting pressure on a patch of skin, and a basic structure that remains the same, from mice to men. The virtue of simulating a circuit in a rodent brain is that the output of the model can be continually tested against the neural reality of the rat, a gruesome process that involves opening up the skull and plunging a needle into the brain. The point is to electronically replicate the performance of the circuit, to build a digital doppelganger of a biological machine.

Cables running from the Blue Gene/L supercomputer to the storage unit. The 2,000-microchip Blue Gene machine is capable of processing 22.8 trillion operations per second—just enough to model a 1-cubic-mm column of rat brain. Courtesy of Alain Herzog/EPFL

"From Copernicus to Einstein, the big breakthroughs always came from conceptual models. They are what integrated all the facts so that they made sense. You can have all the data in the world, but without a model the data will never be enough."

Felix Schürmann, the project manager of Blue Brain, oversees this daunting process. He's 30 years old but looks even younger, with a chiseled chin, lean frame, and close-cropped hair. His patient manner is that of someone used to explaining complex ideas in simple sentences. Before the Blue Brain project, Schürmann worked at the experimental fringes of computer science, developing simulations of quantum computing. Although he's since mastered the vocabulary of neuroscience, referencing obscure acronyms with ease, Schürmann remains most comfortable with programming. He shares a workspace with an impressively diverse group—the 20 or so scientists working full-time on Blue Brain's software originate from 14 different countries. When we enter the hushed room, the programmers are all glued to their monitors, fully absorbed in the hieroglyphs on the screen. Nobody even looks up. We sit down at an empty desk and Schürmann opens his laptop.

laboratory, state-of-the-art equipment allows for computer-

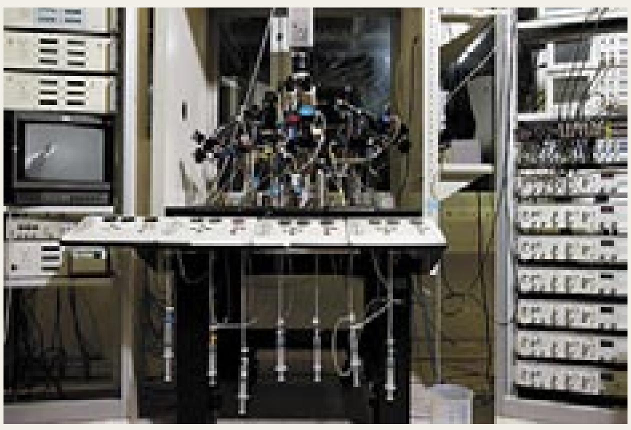

controlled, simultaneous recordings of the tiny electrical currents that form the basis of nerve impulses. Here, a technique known as "patch clamp" provides direct access to seven individual neurons and their chemical synaptic interactions. The patch clamp robot-at work 24 hours a day, seven days a week - helped the Blue Brain team speed through 30 years of research in six months. Inset, a system integrates a bright-field microscope with computer-assisted reconstruction of neuron structure.

The computer screen is filled with what look like digitally rendered tree branches. Schürmann zooms out so that the branches morph into a vast arbor, a canopy so dense it's practically opaque. "This," he proudly announces, "is a virtual neuron. What you're looking at are the thousands of synaptic connections it has made with other [virtual] neurons." When I look closely, I can see the faint lines where the virtual dendrites are subdivided into compartments. At any given moment, the supercomputer is modeling the chemical activity inside each of these sections so that a single simulated neuron is really the sum of 400 independent simulations. This is the level of precision required to accurately imitate just one of the 100 billion cells—each of them unique—inside the brain. When Markram talks about building a mind from the "bottom-up," these intracellular compartments are the bottom. They are the fundamental unit of the model.

But how do you get these simulated compartments to act in a realistic manner? The good news is that neurons are electrical processors: They represent information as ecstatic bursts of voltage, just like a silicon microchip. Neurons control the flow of electricity by opening and closing different ion channels, specialized proteins embedded in the cellular membrane. When the team began constructing their model, the first thing they did was program the existing ion channel data into the supercomputer. They wanted their virtual channels to act just like the real thing. However, they soon ran into serious problems. Many of the experiments used inconsistent methodologies and generated contradictory results, which were too irregular to model. After several frustrating failures—"The computer was just churning out crap," Markram says—the team realized that if they wanted to simulate ion channels, they needed to generate the data themselves.

That's when Schürmann leads me down the hall to Blue Brain's "wet lab." At first glance, the room looks like a generic neuroscience lab. The benches are cluttered with the usual salt solutions and biotech catalogs. There's the familiar odor of agar plates and astringent chemicals. But then I notice, tucked in the corner of the room, is a small robot. The machine is about the size of a microwave, and consists of a beige plastic tray filled with a variety of test tubes and a delicate metal claw holding a pipette. The claw is constantly moving back and forth across the tray, taking tiny sips from its buffet of different liquids. I ask Schürmann what the robot is doing. "Right now," he says, "it's recording from a cell. It does this 24 hours a day, seven days a week. It doesn't sleep and it never gets frustrated. It's the perfect postdoc."

The science behind the robotic experiments is straightforward. The Blue Brain team genetically engineers Chinese hamster ovary cells to express a single type of ion channel—the brain contains more than 30 different types of channels—then they subject the cells to a variety of physiological conditions. That's when the robot goes to work. It manages to "patch" a neuron about 50 percent of the time, which means that it can generate hundreds of data points a day, or about 10 times more than an efficient lab technician. Markram refers to the robot as "science on an industrial scale," and is convinced that it's the future of lab work. "So much of what we do in science isn't actually science," he says, "I say let robots do the mindless work so that we can spend more time thinking about our questions."

According to Markram, the patch clamp robot helped the Blue Brain team redo 30 years of research in six months. By analyzing the genetic expression of real rat neurons, the scientists could then start to integrate these details into the model. They were able to construct a precise map of ion channels, figuring out which cell types had which kind of ion channel and in what density. This new knowledge was then plugged into Blue Brain, allowing the supercomputer to accurately simulate any neuron anywhere in the neocortical column. "The simulation is getting to the point," Schürmann says, "where it gives us better results than an actual experiment. We get the same data, but with less noise and human error." The model, in other words, has exceeded its own inputs. The virtual neurons are more real than reality.

Every brain is made of the same basic parts. A sensory cell in a sea slug works just like a cortical neuron in a human brain. It relies on the same neurotransmitters and ion channels and enzymes. Evolution only innovates when it needs to, and the neuron is a perfect piece of design.

In theory, this meant that once the Blue Brain team created an accurate model of a single neuron, they could multiply it to get a three-dimensional slice of brain. But that was just theory. Nobody knew what would happen when the supercomputer began simulating thousands of brain cells at the same time. "We were all emotionally prepared for failure," Markram says. "But I wasn't so prepared for what actually happened."

A simulated neuron from a rat brain showing "spines"—tiny knobs protruding from the dendrites that will eventually form synapses with other neurons. Pyramidal cells such as these

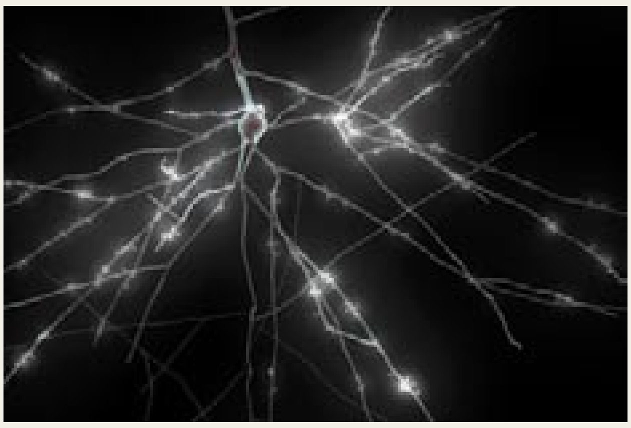

(so-called because of their triangular shape) comprise about 80 percent of cerebral cortex mass. Courtesy of BBP/EPFL

After assembling a three-dimensional model of 10,000 virtual neurons, the scientists began feeding the simulation electrical impulses, which were designed to replicate the currents constantly rippling through a real rat brain. Because the model focused on one particular kind of neural circuit—a neocortical column in the somatosensory cortex of a two-week-old rat—the scientists could feed the supercomputer the same sort of electrical stimulation that a newborn rat would actually experience.

It didn't take long before the model reacted. After only a few electrical jolts, the artificial neural circuit began to act just like a real neural circuit. Clusters of connected neurons began to fire in close synchrony: the cells were wiring themselves together. Different cell types obeyed their genetic instructions. The scientists could see the cellular looms flash and then fade as the cells wove themselves into meaningful patterns. Dendrites reached out to each other, like branches looking for light. "This all happened on its own," Markram says. "It was entirely spontaneous." For the Blue Brain team, it was a thrilling breakthrough. After years of hard work, they were finally able to watch their make-believe brain develop, synapse by synapse. The microchips were turning themselves into a mind.

But then came the hard work. The model was just a first draft. And so the team began a painstaking editing process. By comparing the behavior of the virtual circuit with experimental studies of the rat brain, the scientists could test out the verisimilitude of their simulation. They constantly fact-checked the supercomputer, tweaking the software to make it more realistic. "People complain that Blue Brain must have so many free parameters," Schürmann says. "They assume that we can just input whatever we want until the output looks good. But what they don't understand is that we are very constrained by these experiments." This is what makes the model so impressive: It manages to simulate a real neocortical column—a functional slice of mind—by simulating the particular details of our ion channels. Like a real brain, the behavior of Blue Brain naturally emerges from its molecular parts.

In fact, the model is so successful that its biggest restrictions are now technological. "We have already shown that the model can scale up," Markram says. "What is holding us back now are the computers." The numbers speak for themselves. Markram estimates that in order to accurately simulate the trillion synapses in the human brain, you'd need to be able to process about 500 petabytes of data (peta being a million billion, or 10 to the fifteenth power). That's about 200 times more information than is stored on all of Google's servers. (Given current technology, a machine capable of such power would be the size of several football fields.) Energy consumption is another huge problem. The human brain requires about 25 watts of electricity to operate. Markram estimates that simulating the brain on a supercomputer with existing microchips would generate an annual electrical bill of about $3 billion . But if computing speeds continue to develop at their current exponential pace, and energy efficiency improves, Markram believes that he'll be able to model a complete human brain on a single machine in ten years or less.

For now, however, the mind is still the ideal machine. Those intimidating black boxes from IBM in the basement are barely sufficient to model a thin slice of rat brain. The nervous system of an invertebrate exceeds the capabilities of the fastest supercomputer in the world. "If you're interested in computing," Schürmann says, "then I don't see how you can't be interested in the brain. We have so much to learn from natural selection. It's really the ultimate engineer."

Neuroscience describes the brain from the outside. It sees us through the prism of the third person, so that we are nothing but three pounds of electrical flesh. The paradox, of course, is that we don't experience our matter. Self-consciousness, at least when felt from the inside, feels like more than the sum of its cells. "We've got all these tools for studying the cortex," Markram says. "But none of these methods allows us to see what makes the cortex so interesting, which is that it generates worlds. No matter how much I know about your brain, I still won't be able to see what you see."

Some philosophers, like Thomas Nagel, have argued that this divide between the physical facts of neuroscience and the reality of subjective experience represents an epistemological dead end. No matter how much we know about our neurons, we still won't be able to explain how a twitch of ions in the frontal cortex becomes the Technicolor cinema of consciousness.

Markram takes these criticisms seriously. Nevertheless, he believes that Blue Brain is uniquely capable of transcending the limits of "conventional neuroscience," breaking through the mind-body problem. According to Markram, the power of Blue Brain is that it can transform a metaphysical paradox into a technological problem. "There's no reason why you can't get inside Blue Brain," Markram says. "Once we can model a brain, we should be able to model what every brain makes. We should be able to experience the experiences of another mind."

When listening to Markram speculate, it's easy to forget that the Blue Brain simulation is still just a single circuit, confined within a silent supercomputer. The machine is not yet alive. And yet Markram can be persuasive when he talks about his future plans. His ambitions are grounded in concrete steps. Once the team is able to model a complete rat brain - that should happen in the next two years - Markram will download the simulation into a robotic rat, so that the brain has a body. He's already talking to a Japanese company about constructing the mechanical animal. "The only way to really know what the model is capable of is to give it legs," he says. "If the robotic rat just bumps into walls, then we've got a problem."

Installing Blue Brain in a robot will also allow it to develop like a real rat. The simulated cells will be shaped by their own sensations, constantly revising their connections based upon the rat's experiences. "What you ultimately want," Markram says, "is a robot that's a little bit unpredictable, that doesn't just do what we tell it to do." His goal is to build a virtual animal—a rodent robot—with a mind of its own.

But the question remains: How do you know what the rat knows? How do you get inside its simulated cortex? This is where visualization becomes key. Markram wants to simulate what that brain experiences. It's a typically audacious goal, a grand attempt to get around an ancient paradox. But if he can really find a way to see the brain from the inside, to traverse our inner space, then he will have given neuroscience an unprecedented window into the invisible. He will have taken the self and turned it into something we can see.

Schürmann leads me across the campus to a large room tucked away in the engineering school. The windows are hermetically sealed; the air is warm and heavy with dust. A lone Silicon Graphics supercomputer, about the size of a large armoire, hums loudly in the center of the room. Schürmann opens the back of the computer to reveal a tangle of wires and cables, the knotted guts of the machine. This computer doesn't simulate the brain, rather it translates the simulation into visual form. The vast data sets generated by the IBM supercomputer are rendered as short films, hallucinatory voyages into the deep spaces of the mind. Schürmann hands me a pair of 3-D glasses, dims the lights, and starts the digital projector. The music starts first, "The Blue Danube" by Strauss. The classical waltz is soon accompanied by the vivid image of an interneuron, its spindly limbs reaching through the air. The imaginary camera pans around the brain cell, revealing the subtle complexities of its form. "This is a random neuron plucked from the model," Schürmann says. He then hits a few keys and the screen begins to fill with thousands of colorful cells. After a few seconds, the colors start to pulse across the network, as the virtual ions pass from neuron to neuron. I'm watching the supercomputer think.

Rendering cells is easy, at least for the supercomputer. It's the transformation of those cells into experience that's so hard. Still, Markram insists that it's not impossible. The first step, he says, will be to decipher the connection between the sensations entering the robotic rat and the flickering voltages of its brain cells. Once that problem is solved—and that's just a matter of massive correlation—the supercomputer should be able to reverse the process. It should be able to take its map of the cortex and generate a movie of experience, a first person view of reality rooted in the details of the brain. As the philosopher David Chalmers likes to say, "Experience is information from the inside; physics is information from the outside." By shuttling between these poles of being, the Blue Brain scientists hope to show that these different perspectives aren't so different at all. With the right supercomputer, our lucid reality can be faked.

"There is nothing inherently mysterious about the mind or anything it makes," Markram says. "Consciousness is just a massive amount of information being exchanged by trillions of brain cells. If you can precisely model that information, then I don't know why you wouldn't be able to generate a conscious mind." At moments like this, Markram takes on the deflating air of a magician exposing his own magic tricks. He seems to relish the idea of "debunking consciousness," showing that it's no more metaphysical than any other property of the mind. Consciousness is a binary code; the self is a loop of electricity. A ghost will emerge from the machine once the machine is built right.

And yet, Markram is candid about the possibility of failure. He knows that he has no idea what will happen once the Blue Brain is scaled up. "I think it will be just as interesting, perhaps even more interesting, if we can't create a conscious computer," Markram says. "Then the question will be: 'What are we missing? Why is this not enough?'"

Niels Bohr once declared that the opposite of a profound truth is also a profound truth. This is the charmed predicament of the Blue Brain project. If the simulation is successful, if it can turn a stack of silicon microchips into a sentient being, then the epic problem of consciousness will have been solved. The soul will be stripped of its secrets; the mind will lose its mystery. However, if the project fails - if the software never generates a sense of self, or manages to solve the paradox of experience - then neuroscience may be forced to confront its stark limitations. Knowing everything about the brain will not be enough. The supercomputer will still be a mere machine. Nothing will have emerged from all of the information. We will remain what can't be known.