

EdgeThe Third Culture

[Home](#)[About Edge](#)[Features](#)[Edge Editions](#)[Press](#)[Edge Search](#)[Subscribe](#)

My visit to Google? Despite the whimsical furniture and other toys, I felt I was entering a 14th-century cathedral — not in the 14th century but in the 12th century, while it was being built. Everyone was busy carving one stone here and another stone there, with some invisible architect getting everything to fit. The mood was playful, yet there was a palpable reverence in the air. "We are not scanning all those books to be read by people," explained one of my hosts after my talk. "We are scanning them to be read by an AI."

When I returned to highway 101, I found myself recollecting the words of Alan Turing, in his seminal paper *Computing Machinery and Intelligence*, a founding document in the quest for true AI. "In attempting to construct such machines we should not be irreverently usurping His power of creating souls, any more than we are in the procreation of children," Turing had advised. "Rather we are, in either case, instruments of His will providing mansions for the souls that He creates."

TURING'S CATHEDRAL [10.24.05]

A visit to Google on the occasion of the 60th anniversary of John von Neumann's proposal for a digital computer

by George Dyson



Introduction

Sorry, but the big news at the Frankfurt Book Fair this year was not about the international sales of your new book. It was about the activities of a new (second year) exhibitor: Google.

What is Google doing at the Frankfurt Book Fair? And why has a consortium of publishers filed a lawsuit against them? On the other hand, why do the "digerati" love Google Print and Google Print Library? How does Google's definition of "fair use" as it pertains to the digital domain, square with the notion that as a writer, own my own words? Clearly, we need to redefine "fair use" in the digital age as a "different use" with its own new set of benchmarks.

Whether we're talking about John Cage's idea of "the mind we all share" or

H.G. Well's "World Brain", Google has its act together and are at the precipice of astonishing changes in human communication...and ultimately, in our sense of who or what we are. And like nearly all science-driven, technological developments, governments can only play catch-up as no one is going to get to vote for Google's changes, and the current laws, written in a pre-digital age, don't address the new situation.

Some sincerely believe we are entering a golden age of wonder and Google is leading the way. And I am pleased to add from personal experience that the leading players, [Eric Schmidt](#), [Sergey Brin and Larry Page](#), are fine individuals: very serious, highly intelligent, principled. They don't come any better. Still, others believe there are reasons for legitimate fear of a (very near) future world in which the world's knowledge is privatized by one corporation. This could be a problem, a very big problem.

George Dyson visited Google last week at the invitation of some Google engineers. The occasion was the 60th anniversary of John von Neumann's proposal for a digital computer. After the visit, Dyson recalled H.G. Wells' prophecy, written in 1938:

"The whole human memory can be, and probably in a short time will be, made accessible to every individual," wrote H. G. Wells in his 1938 prophecy World Brain. "This new all-human cerebrum need not be concentrated in any one single place. It can be reproduced exactly and fully, in Peru, China, Iceland, Central Africa, or wherever else seems to afford an insurance against danger and interruption. It can have at once, the concentration of a craniate animal and the diffused vitality of an amoeba." Wells foresaw not only the distributed intelligence of the World Wide Web, but the inevitability that this intelligence would coalesce, and that power, as well as knowledge, would fall under its domain. "In a universal organization and clarification of knowledge and ideas... in the evocation, that is, of what I have here called a World Brain... in that and in that alone, it is maintained, is there any clear hope of a really Competent Receiver for world affairs... We do not want dictators, we do not want oligarchic parties or class rule, we want a widespread world intelligence conscious of itself."

—[JB](#)

GEORGE DYSON, a historian among futurists, is the author of *Darwin Among the Machines*; and *Project Orion: The True Story of the Atomic Spaceship*.

[George Dyson's Edge Bio page](#)

TURING'S CATHEDRAL [10.24.05]

A visit to Google on the occasion of the 60th anniversary of John von Neumann's proposal for a digital computer

by George Dyson



GEORGE DYSON, a historian among futurists, is the author of *Baidarka*; *Project Orion*; and *Darwin Among the Machines*.

[George Dyson's Edge Bio page](#)

[GEORGE DYSON:] In the digital universe, there are two kinds of bits: bits that represent structure (differences in space) and bits that represent sequence (differences in time). Digital computers — as formalized by Alan Turing, and delivered by John von Neumann — are devices that translate between these two species of bits according to definite rules.

Exactly sixty years ago, at the Institute for Advanced Study in Princeton, New Jersey, mathematician John von Neumann began seeking funding to build a machine that would do this at electronic speeds. "I am sure that the projected device, or rather the species of devices of which it is to be the first representative, is so radically new that many of its uses will become clear only after it has been put into operation," he wrote to Lewis Strauss on 24 October 1945. "Uses which are likely to be the most important are by definition those which we do not recognize at present because they are farthest removed from our present sphere."

Von Neumann received immediate support from the Army, the Navy, and the Air Force, but the main sponsor soon became the United States Atomic Energy Commission, or AEC. This deal with the devil was hard to resist. "The Army contract provides for general supervision by the Ballistic Research Laboratory of the Army, whereas the AEC provides for supervision by von Neumann," the Institute administration explained in 1949. When the machine finally became operational in 1951, it had 5 kilobytes of random-access memory: a 32 x 32 x 40 matrix of binary digits, stored as a flickering pattern of electrical charge, shifting from millisecond to millisecond on the surface of 40 cathode-ray tubes.

The codes that inoculated this empty universe were based upon the architectural principal that a pair of 5-bit coordinates could uniquely identify a memory location containing a string of 40 bits. These 40 bits could include not only data (numbers that mean things) but executable instructions (numbers that do things) — including instructions to transfer control to another location and do something else.

By breaking the distinction between numbers that *mean* things and numbers that *do* things, von Neumann unleashed the power of the stored-program computer, and our universe would never be the same. It was no coincidence that the chain reaction of addresses and instructions within the core of the computer resembled a chain reaction within the core of an atomic bomb. The driving force behind the von Neumann project was the

push to run large-scale Monte Carlo simulations of how the implosion of a sub-critical mass of fissionable material could lead the resulting critical assembly to explode.

The success of Monte Carlo led to compact, predictable fission explosives, and this, coupled with more Monte Carlo (and more Stan Ulam) led to the "Super" or hydrogen bomb. But the actual explosion of digital computing has overshadowed the threatened explosion of the bombs. From an initial nucleus of 4×10^4 bits changing state at kilocycle speed, the von Neumann's archetype has proliferated to individual matrices of more than 10^9 bits, running at speeds of more than 10^9 cycles per second, interconnected by an extended address matrix encompassing up to 10^9 remote hosts. This growth continues to speed up. Over 10^{10} transistors are now produced each second, and many of these are being incorporated into devices — no longer just computers — that have an IP (Internet Protocol) address. The current 32-bit IP address space will be exhausted within 10 years or less.

In the early 1950s, when mean time between memory failure was measured in minutes, no one imagined that a system depending on every bit being in exactly the right place at exactly the right time could be scaled up by a factor of 10^{13} in size, and down by a factor of 10^6 in time. Von Neumann, who died prematurely in 1957, became increasingly interested in understanding how biology has managed (and how technology might manage) to construct reliable organisms out of unreliable parts. He believed the von Neumann architecture would soon be replaced by something else. Even if codes could be completely debugged, million-cell memories could never be counted upon, digitally, to behave consistently from one kilocycle to the next.

Fifty years later, thanks to solid state micro-electronics, the von Neumann matrix is going strong. The problem has shifted from how to achieve reliable results using sloppy hardware, to how to achieve reliable results using sloppy code. The von Neumann architecture is here to stay. But new forms of architecture, built upon the underlying layers of Turing-von Neumann machines, are starting to grow. What's next? Where was von Neumann heading when his program came to a halt?

As organisms, we possess two outstanding repositories of information: the information conveyed by our genes, and the information stored in our brains. Both of these are based upon non-von-Neumann architectures, and it is no surprise that Von Neumann became fascinated with these examples as he left his chairmanship of the AEC (where he had succeeded Lewis Strauss) and began to lay out the research agenda that cancer prevented him from following up. He considered the second example in his posthumously-published *The Computer and the Brain*.

"The message-system used in the nervous system... is of an essentially statistical character," he explained. "In other words, what matters are not the precise positions of definite markers, digits, but the statistical characteristics of their occurrence... a radically different system of notation from the ones we are familiar with in ordinary arithmetics and mathematics... Clearly, other traits of the (statistical) message could also be used: indeed, the frequency referred to is a property of a single train of pulses whereas every one of the relevant nerves consists of a large number of fibers, each of which transmits numerous trains of pulses. It is,

therefore, perfectly plausible that certain (statistical) relationships between such trains of pulses should also transmit information.... Whatever language the central nervous system is using, it is characterized by less logical and arithmetical depth than what we are normally used to [and] must structurally be essentially different from those languages to which our common experience refers."

Or, as his friend Stan Ulam put it, "What makes you so sure that mathematical logic corresponds to the way we think?"

Pulse-frequency coding, whether in a nervous system or a probabilistic search-engine, is based on statistical accounting for what connects where, and how frequently connections are made between given points. As von Neumann explained in 1948: "A new, essentially logical, theory is called for in order to understand high-complication automata and, in particular, the central nervous system. It may be, however, that in this process logic will have to undergo a pseudomorphosis to neurology to a much greater extent than the reverse."

Von Neumann died just as the revolution in molecular biology, sparked by the elucidation of the structure of DNA in 1953, began to unfold. Life as we know it is based on digitally-coded instructions, translating between sequence and structure (from nucleotides to proteins) exactly as Turing prescribed. Ribosomes and other cellular machinery play the role of processors: reading, duplicating, and interpreting the sequences on the tape. But this uncanny resemblance has distracted us from the completely different method of addressing by which the instructions are carried out.

In a digital computer, the instructions are in the form of COMMAND (ADDRESS) where the address is an exact (either absolute or relative) memory location, a process that translates informally into "DO THIS with what you find HERE and go THERE with the result." Everything depends not only on precise instructions, but on HERE, THERE, and WHEN being exactly defined. It is almost incomprehensible that programs amounting to millions of lines of code, written by teams of hundreds of people, are able to go out into the computational universe and function as well as they do given that one bit in the wrong place (or the wrong time) can bring the process to a halt.

Biology has taken a completely different approach. There is no von Neumann address matrix, just a molecular soup, and the instructions say simply "DO THIS with the next copy of THAT which comes along." The results are far more robust. There is no unforgiving central address authority, and no unforgiving central clock. This ability to take general, organized advantage of local, haphazard processes is exactly the ability that (so far) has distinguished information processing in living organisms from information processing by digital computers.

Of course, dreams of Object-Oriented Programming Languages and asynchronous processing have been around almost as long as digital computing, and content-addressable memory was one of the alternative architectures that Julian Bigelow, von Neumann's original chief architect, and many others have long had in mind.

"For man-made electronic computers, a practice adopted, whereby events are represented with serial dependence in time, has resulted in computing apparatus that must be built of elements that are, to a large extent,

strictly independent across space-dimensions," Bigelow explained in 1965. "Accomplishment of the desired time-sequential process on a given computing apparatus turns out to be largely a matter of specifying sequences of addresses of items which are to interact... With regard to the explicit address nuisance, studies have been made of the possibility of causing various elementary pieces of information situated in the cells of a large array (say, of memory) to enter into a computation process without explicitly generating a coordinate address in 'machine-space' for selecting them out of the array."

Hardware-based content-addressable memory is used, on a local scale, in certain dedicated high-speed network routers, but template-based addressing did not catch on widely until Google (and brethren) came along. Google is building a new, content-addressable layer overlying the von Neumann matrix underneath. The details are mysterious but the principle is simple: it's a map. And, as Dutch (and other) merchants learned in the sixteenth century, great wealth can be amassed by Keepers of the Map.

We call this a "search engine" — a content-addressable layer that makes it easier for *us* to find things, share ideas, and retrace our steps. That's a big leap forward, but it isn't a universe-shifting revolution equivalent to von Neumann's breaking the distinction between numbers that *mean* things and numbers that *do* things in 1945.

However, once the digital universe is thoroughly mapped, and initialized by us searching for meaningful things and following meaningful paths, it will inevitably be colonized by codes that will start *doing* things with the results. Once a system of template-based-addressing is in place, the door is opened to code that can interact directly with other code, free at last from a rigid bureaucracy requiring that every bit be assigned an exact address. You can (and a few people already are) write instructions that say "Do THIS with THAT" — without having to specify exactly Where or When. This revolution will start with simple, basic coded objects, on the level of nucleotides heading out on their own and bringing amino acids back to a collective nest. It is 1945 all over again.

And it is back to Turing, who in his 1948 report on intelligent machinery to the National Physical Laboratory advised that "intellectual activity consists mainly of various kinds of search." It was Turing, in 1936, who showed von Neumann that digital computers are able to solve most — but not all — problems that can be stated in finite, unambiguous terms. They may, however, take a very long time to produce an answer (in which case you build faster computers) or it may take a very long time to ask the question (in which case you hire more programmers). Computers have been getting better and better at providing answers — but only to questions that programmers are able to ask.

We can divide the computational universe into three sectors: computable problems; non-computable problems (that can be given a finite, exact description but have no effective procedure to deliver a definite result); and, finally, questions whose answers are, in principle, computable, but that, in practice, we are unable to ask in unambiguous language that computers can understand.

We do most of our computing in the first sector, but we do most of our living (and thinking) in the third. In the real world, most of the time, finding an answer is easier than defining the question. It's easier to draw

something that looks like a cat, for instance, than to describe what, exactly, makes something look like a cat. A child scribbles indiscriminately, and eventually something appears that resembles a cat. A solution finds the problem, not the other way around. The world starts making sense, and the meaningless scribbles (and a huge number of neurons) are left behind.

This is why Google works so well. All the answers in the known universe are there, and some very ingenious algorithms are in place to map them to questions that people ask.

"An argument in favor of building a machine with initial randomness is that, if it is large enough, it will contain every network that will ever be required," advised Turing's former assistant and cryptanalyst Irving J. Good, speaking at IBM in 1958. A network, whether of neurons, computers, words, or ideas, contains solutions, waiting to be discovered, to problems that need not be explicitly defined. It is much easier to find explicit answers than to ask explicit questions. And some will be answers to questions that programmers will never have to ask.

"The whole human memory can be, and probably in a short time will be, made accessible to every individual," wrote H. G. Wells in his 1938 prophecy *World Brain*. "This new all-human cerebrum need not be concentrated in any one single place. It can be reproduced exactly and fully, in Peru, China, Iceland, Central Africa, or wherever else seems to afford an insurance against danger and interruption. It can have at once, the concentration of a craniate animal and the diffused vitality of an amoeba." Wells foresaw not only the distributed intelligence of the World Wide Web, but the inevitability that this intelligence would coalesce, and that power, as well as knowledge, would fall under its domain. "In a universal organization and clarification of knowledge and ideas... in the evocation, that is, of what I have here called a World Brain... in that and in that alone, it is maintained, is there any clear hope of a really Competent Receiver for world affairs... We do not want dictators, we do not want oligarchic parties or class rule, we want a widespread world intelligence conscious of itself."

My visit to Google? Despite the whimsical furniture and other toys, I felt I was entering a 14th-century cathedral — not in the 14th century but in the 12th century, while it was being built. Everyone was busy carving one stone here and another stone there, with some invisible architect getting everything to fit. The mood was playful, yet there was a palpable reverence in the air. "We are not scanning all those books to be read by people," explained one of my hosts after my talk. "We are scanning them to be read by an AI."

When I returned to highway 101, I found myself recollecting the words of Alan Turing, in his seminal paper *Computing Machinery and Intelligence*, a founding document in the quest for true AI. "In attempting to construct such machines we should not be irreverently usurping His power of creating souls, any more than we are in the procreation of children," Turing had advised. "Rather we are, in either case, instruments of His will providing mansions for the souls that He creates."

Google is Turing's cathedral, awaiting its soul. We hope. In the words of an unusually perceptive friend: "When I was there, just before the IPO, I thought the coziness to be almost overwhelming. Happy Golden Retrievers

running in slow motion through water sprinklers on the lawn. People waving and smiling, toys everywhere. I immediately suspected that unimaginable evil was happening somewhere in the dark corners. If the devil would come to earth, what place would be better to hide?"

For 30 years I have been wondering, what indication of its existence might we expect from a true AI? Certainly not any explicit revelation, which might spark a movement to pull the plug. Anomalous accumulation or creation of wealth might be a sign, or an unquenchable thirst for raw information, storage space, and processing cycles, or a concerted attempt to secure an uninterrupted, autonomous power supply. But the real sign, I suspect, would be a circle of cheerful, contented, intellectually and physically well-nourished people surrounding the AI. There wouldn't be any need for True Believers, or the downloading of human brains or anything sinister like that: just a gradual, gentle, pervasive and mutually beneficial contact between us and a growing something else. This remains a non-testable hypothesis, for now. The best description comes from science fiction writer Simon Ings:

"When our machines overtook us, too complex and efficient for us to control, they did it so fast and so smoothly and so *usefully*, only a fool or a prophet would have dared complain."

[John Brockman](#), Editor and Publisher
[Russell Weinberger](#), Associate Publisher
contact: editor@edge.org
Copyright © 2005 by [Edge Foundation, Inc](#)
All Rights Reserved.

|[Top](#)|