Lecture 6: God and Computer Science (8 December 1999)

Thanks to all of you for sticking with me until the end here. I have to confess that when I made up the table of contents for these lectures during the summer, I had almost no concept of what I might say in the last one. But I figured that I had two months' time at MIT to come up with something, and if I'd only keep my antennas open and tuned to the right frequencies, I might pick up some appropriate ideas. Well, sure enough, there's no shortage of stimulating ideas that relate God and computer science. I hope I'll be able to explain today some of the ones I think are most interesting and important. In fact, I thought of more than enough things to say in one lecture, and I've had to rank order them; so I might not have time to talk about everything that I promised you last week.

Lots of people have, of course, written about God and science in general. Shortly after my arrival here, I went to a dinner hosted by a thriving local organization called the Faith and Science Exchange, celebrating its tenth anniversary. Also, ever since the early 1980s, the University of California at Berkeley has had an active Center for Theology and the Natural Sciences.

But computer science is an *unnatural* science. Computer science deals with artificial things, not bound by the constraints of nature. When I chose the title of this lecture, I had a gut feeling that computer scientists could shed some new light on the subject—in addition to the fine contributions already made by biologists, physicists, and other scientists as well as theologians—because I think computer science gives us new analogies and theories that can help us to understand God.

For example, when Arthur Peacocke wrote twenty years ago about *Creation and the World of Science*, he compared God to the composer of a symphony, writing music that obeys fixed laws but contains random elements. Peacocke saw God as an "improvisor of unsurpassed ingenuity." Music is indeed a useful analogy, because it's a flexible form that moves through time. But certainly computer programs are much richer in this respect, because programs not only move through time, they also interact with people and they can even modify themselves.

When I talk about computer science as a possible basis for insights about God, of course I'm not thinking about God as a super-smart intellect surrounded by large clusters of ultrafast Linux workstations and great search engines. That's the user's point of view. I'm thinking about the *science* part of computer science, the abstract notions of processes, the theories that computer scientists have been developing about how to deal with large quantities of nonuniform data in dynamic ways. Such things are much better understood now than they ever have been before, and they clearly have intimate connections with God's role as creator and sustainer of the universe. Furthermore, computational models are better able to describe many aspects of the universe better than any other models we know. All scientific theories can, for example, be modeled by programs.

Years ago, I was pondering the difference between science and art. People had been asking me why my books were called *The Art of Computer Programming* instead of *The Science of Computer Programming*, and I realized that there's a simple explanation: Science is what we understand well enough to explain to a computer; art is everything else. Every time science advances, part of an art becomes a science, so art loses a little bit. Yet, mysteriously, art always seems to register a net gain, because as we understand more we invent more new things that we can't explain to computers.

During the panel discussion a few weeks ago, it developed that a lot of people in the audience — maybe a majority — thought that science would soon catch up with art and there wouldn't be any art left. According to that opinion, we'd know how to program a robot that would be able to do everything the people on the panel could do, and also the people in the audience I guess.

Now I personally see no signs that any such thing is on the horizon, but I do want to mention a point that Brian Hayes brought up earlier this year in *American Scientist*. Hayes observed that if we can create such robots, those robots would also be able to create such robots, and so on. And you see this would increase the likelihood that we ourselves were created by some designer.

In any case, I think people who write programs do have at least a glimmer of extra insight into the nature of God for that very reason, because creating a program often means that you have to create a small universe. For example, I spent most of this year writing programs to simulate a new computer called MMIX. One of those

programs was probably the most difficult that I've ever written, and I thought hard about how it would be to live with such a machine and with the new tools that I was creating — sort of like living in a new subculture.

I think it's fair to say that many of today's large computer programs rank among the most complex intellectual achievements of all time. They're absolutely trivial by comparison with any of the works of God, but still they're somehow closer to those works than anything else we know.

My main point, though, is not to debate the merits of computer science. Rather, I want to discuss some of the things that our experiences with computers during the past fifty years have taught us.

First I want to mention the fact that computing gives us great appreciation for the size of finite numbers. I spoke about this topic in Boston many years ago, at a big meeting of the American Association for the Advancement of Science held in 1976. My paper was called "Coping with finiteness," and it was reprinted later in several books. I hate to repeat myself, but I suppose some of you here today were unable to hear my talk in 1976; and maybe you never read the paper either. And even if you did, you've probably forgotten about it. So I'm going to spend the next few minutes repeating some simple but eye-opening facts about finite numbers.

I want to start small and remind you that the product x times n means x plus x plus x plus x and so on, the sum of n copies of the number x:

$$x\,n = \overbrace{x + x + \cdots + x}^{n \text{ copies of } x}\,.$$

Similarly we can talk about $x \uparrow n$, or x to the nth power, which is the *product* of n copies of x:

$$x \uparrow n = \overbrace{x\,x \ldots x}^{n \text{ copies of } x}\,.$$

For example,

$$\begin{aligned} 10 \uparrow 10 &= 10 \cdot 10 \cdot 10 \cdot 10 \cdot 10 \cdot 10 \cdot 10 \cdot 10 \cdot 10 \cdot 10 \\ &= 10{,}000{,}000{,}000\,; \end{aligned}$$

it's ten billion. People usually write that number as 10^{10}, but I like to write it with an arrow because of the next step up, which uses *two* arrows. I mean, if you have two arrows as in $x \uparrow\uparrow n$, you are computing x to the x to the x to the x and so on, n times:

$$x \uparrow\uparrow n = \overbrace{x \uparrow (x \uparrow (\cdots \uparrow x) \cdots)}^{n \text{ copies of } x}.$$

Now this arrow notation is my big claim to fame, because it's what got me into the *Guinness Book of World Records*. In particular,

$$10 \uparrow\uparrow 10 = 10^{10^{10^{10^{10^{10^{10^{10^{10^{10}}}}}}}}}$$

$$= 1 \text{ followed by } 10^{10^{10^{10^{10^{10^{10^{10^{10}}}}}}}} \text{ zeros.}$$

That's a fairly large number. If you put a monkey down at a typewriter and wait until he types out the entire text of Hamlet, with no errors, the number of trials is only 1 followed by about 40,000 zeros. Our number $10 \uparrow\uparrow 10$ is 1 followed by quite a few more zeros than that. Now you understand ten double-arrow ten.

The general rule, of course, is that if you have k arrows, you just define it as the operation with $k-1$ arrows, over and over again:

$$x \overbrace{\uparrow\uparrow \ldots \uparrow}^{k \text{ arrows}} n = \underbrace{x \overbrace{\uparrow \ldots \uparrow}^{k-1} (x \overbrace{\uparrow \ldots \uparrow}^{k-1} (\cdots \overbrace{\uparrow \ldots \uparrow}^{k-1} x) \cdots)}_{n \text{ occurrences of } x}.$$

I want to give you a small example of the arrow functions so that you can better understand these finite numbers. Let's look at ten quadruple-arrow three. By definition that means ten triple-arrow ten triple-arrow ten:

$$10 \uparrow\uparrow\uparrow\uparrow 3 \;=\; 10 \uparrow\uparrow\uparrow (10 \uparrow\uparrow\uparrow 10);$$

so we first have to evaluate ten triple-arrow ten. Well, of course, ten triple-arrow ten is

$$10 \uparrow\uparrow (10 \uparrow\uparrow (10 \uparrow\uparrow (10 \uparrow\uparrow (10 \uparrow\uparrow (10 \uparrow\uparrow (10 \uparrow\uparrow (10 \uparrow\uparrow (10 \uparrow\uparrow 10))))))))$$

and that is

$$10 \uparrow\uparrow (10 \uparrow\uparrow (10 \uparrow\uparrow (10 \uparrow\uparrow (10 \uparrow\uparrow (10 \uparrow\uparrow (10 \uparrow\uparrow (10 \uparrow\uparrow 10^{10^{10^{10^{10^{10^{10^{10^{10^{10}}}}}}}}})))))))$$

$$= 10 \uparrow\uparrow (10 \uparrow\uparrow (10 \uparrow\uparrow (10 \uparrow\uparrow (10 \uparrow\uparrow (10 \uparrow\uparrow (10 \uparrow\uparrow 10^{10^{\cdot^{\cdot^{\cdot^{10}}}}}))))))$$

where the stack of tens is $10 \uparrow\uparrow 10$ levels tall. I can't even tell you how big that stack of tens is without using double-arrow notation.

It's such a huge number, I can't even write it down, but then we repeat the double-arrow operation again, getting an even huger number; and so on, until finally we get this thing evaluated. We can save space by calling the final result $\mathfrak{K}$. (If you look with a magnifying glass you'll see that I have worked hard to make a very fancy K here; I couldn't just use an ordinary letter for such an immense quantity.)

Of course we're not done yet. $\mathfrak{K}$ is just ten triple-arrow ten. In order to get the final number that I started with — ten quadruple-arrow three — I have to take ten triple-arrow to this fancy K, namely

$$10 \uparrow\uparrow\uparrow\uparrow 3 = 10 \uparrow\uparrow\uparrow \mathfrak{K} = \underbrace{10 \uparrow\uparrow (10 \uparrow\uparrow (10 \uparrow\uparrow \cdots \uparrow\uparrow (10 \uparrow\uparrow 10) \cdots))}_{\mathfrak{K} \text{ times}}.$$

If you stop to think about this, you'll have to admit that, if anything is mind boggling, this is; it's *incredibly* huge. In my paper I noted that the three dots "$\cdots$" in this formula suppress a lot of detail — so much that I should probably have used four dots.

Now that we have evaluated $10 \uparrow\uparrow\uparrow\uparrow 3$, let's call it Super K. If you don't agree that Super K is so large as to be beyond human comprehension, I can at least prove conclusively that if you consider all the numbers less than or equal to Super K, almost all of them are impossible to describe in any way in the universe, because of a theorem by which computer scientists know that most number descriptions are incompressible. On the other hand, Super K is very small as finite numbers go; almost all finite numbers are a lot bigger than this. I don't care. Super K is big enough for me.

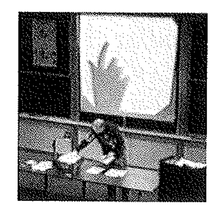

As soon as you begin to understand the immensity of Super K, you will realize that just being finite isn't much of a limitation, and you will see how pointless are the philosophers' discussions about finite versus infinite. Infinity is a red herring. I would be perfectly happy to give up immortality if I could only live Super K years before dying. In fact, Super K nanoseconds would be enough.

Many years ago, I learned about Cantor's famous theory about higher orders of infinity. I learned the beautifully simple fact that the set of all ways to label the elements of any given set with zeros and ones always has strictly more elements than the given set itself has. And I once thought, if I ever had to preach a sermon in church, I would try to explain Cantor's theorem to my non-mathematical friends so that they could understand something about the infinite.

But now I realize that infinity is not necessarily even one of God's attributes. I'm quite willing to grant that God might indeed be infinite, and that God might have the power to examine infinitely many possibilities in an instant. But even the ability to deal with finitely many numbers, on the order of Super K, is much more than enough to inspire awe.

Moreover, I don't think theologians can legitimately disagree with me on this. To say that God's abilities are not infinite, but limited by quantities like Super K, is not a realistic limitation at all. Such a limitation cannot contradict the Bible or any other sacred text, because natural language has no words to distinguish meaningfully between such unimaginably large magnitudes.

For example, the word "infinite" itself occurs only three times in the King James Bible. The first time is when one of Job's comrades says, "Job, your iniquities are infinite." In the second place, God's understanding is said to be infinite. And in the third place, the power of Egypt is said to be infinite. You see, only the second of these three applies this attribute to God; and the Hebrew words in that verse can be translated more accurately by the phrase "too big to count."

The amount of academic hair splitting about finite versus infinite in the literature is itself too much to count; but it misses the point. The real point, I think, is made rather well in Psalm 139:

(1) O Lord, you have searched me and known me . . .
(2) . . . You discern my thoughts from afar . . .
(4) . . . Even before a word is on my tongue,
lo, O Lord, you know it altogether . . .
(17) . . . How vast is the sum of [your thoughts] —
(18) If I try to count them, they are more than the sand.

In other words, God knows incredibly more than we can understand.

"Lord, you discern my thoughts from afar." I grew up with the idea that God constantly reads my mind, and I've always been comfortable with that invasion of my privacy. As a result, I haven't been especially successful in cryptographic research about keeping secrets. Of course I don't understand how it's possible for God to read my mind or to penetrate anybody's consciousness, especially because every individual brain probably has its own code for information processing. But that doesn't make me disbelieve that God can do it, even if we limit God to having finite capacity of size Super K. Peter Gomes has wisely described the Bible as an effort "to cram into the human imagination the unimaginable immensity of God." But he doesn't mean to imply that the Bible makes technical distinctions about subtle mathematical details.

Let's look at another example, based on a rather deep result of computer science that came out of an MIT Ph.D. thesis — Larry Stockmeyer's thesis in 1974, working jointly with Albert Meyer. I don't want to talk too much about it but there's a theorem of Büchi according to which we know that certain statements about the positive numbers can always be proved or disproved in a finite amount of time. (Technical people in the audience will understand that I'm talking about "weak second-order arithmetic.") Any statement in weak second-order arithmetic can be expressed in terms of 64 symbols, including a blank space. The statements might involve so-called quantifiers: For example, the statement "For all x there exists y such that y is less than $x+1$" is written `∀x∃y(y<x+1)`.

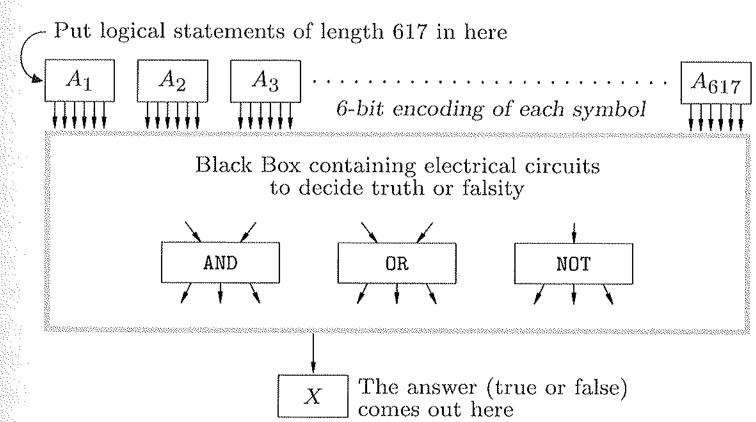

FIGURE 1.

According to Büchi's theorem, we can build an electrical circuit that will decide the truth or falsity of any statement in weak second-order arithmetic, if that statement has a bounded length, say 617 characters or less. We start out by adding blanks to make the statement exactly 617 characters long; then we use a 6-bit encoding for each of the 617 symbols, giving 6×617 inputs to the circuit shown in Figure 1. Then if appropriate combinations of "and," "or," and "not" have been put into the Black Box, out comes the answer: The statement is true, or false, as the case may be.

This circuit is finite. But Meyer and Stockmeyer proved that every such circuit must use at least 10^{125} components. And 10^{125}, which is of course really puny when compared to Super K, is still plenty big; it is larger than the number of protons and neutrons in the entire known universe. Furthermore, if we change 617 to 618, the problem gets harder yet. So you can see that there are absolute, fundamental limitations of complexity in certain computational problems.

I think it's fair to say that God may well be bound by the laws of computational complexity, even if we grant (as I do) that the Bible is God's inspired word. The Bible doesn't deal with Büchi's theorem or any other such fine points of detail, nor was it ever intended to.

Computer scientists today know many things about all kinds of levels of incredible difficulty that are inherent in the solution of different kinds of computational tasks, even when those tasks are solvable in finite time. These theoretical results could be used in academic discussions to restrict God's ability to be all-knowing and/or all-powerful in certain ways, if we assume that God has finite resources of a certain size. But I don't recommend that theologians undertake a deep study of computational complexity (unless, of course, they really enjoy it). Because the fact is, God can know much more than enough, and can be plenty powerful enough, to do *anything* relevant to the universe, *without* being strictly all-knowing or all-powerful. Finiteness is *not* a limitation in practice.

When I say that the question "finite or infinite?" is a red herring, I don't mean simply that philosophers and theologians have often been arguing about an unimportant issue. I also mean that physicists and other scientists fail to realize this. For example, take the literature of chaos theory: Hundreds of papers have been written about the behavior of solutions to unstable recurrences, by people who assume that *real* numbers are *real*.

(Let me explain, to non-mathematicians in the audience. When mathematicians talk about real numbers, they mean decimal numbers that have infinite accuracy — infinitely many decimal places.)

Well, the fact is, real numbers are an abstraction, an idealization. I grant you that they're an immensely useful abstraction: The concept of real numbers allows us to apply calculus and other tools of mathematics to solve all kinds of important problems.

But it's a tremendous leap of faith to assume that real numbers apply perfectly to the real world — to assume, for example, that two physically realizable objects could be in different places, even though their positions in space agree up to Super K decimal places. I can understand why people unconsciously make this assumption: All textbooks of mathematics start with real numbers. The concept is familiar and easy to work with.

In a similar way, I thought I "knew" what parallel lines were after I learned geometry in high school. Given any line and a point off the line, that line and that point determine a plane, and in the plane there's exactly one line through the point that's parallel to the line you started with. Parallel means that it never intersects that line. This statement is called Euclid's fifth postulate, Euclid's parallel postulate, and it seems obviously true.

Years later I learned about non-Euclidean geometries, which satisfy the other axioms stated by Euclid but not this parallel postulate. In some geometries there can be two or more lines parallel to the given one; in other geometries there aren't any at all.

I thought this was an amusing curiosity, but I never believed for a moment that non-Euclidean geometries had anything to do with reality. The possibility didn't even cross my mind, since I *knew* that the universe was Euclidean. Maybe twenty years went by before I was shocked to realize that I had no grounds for that hypothesis at all. Euclid's law was convenient for the practical calculations I needed, but it wasn't good enough for astronomers who were faced with the actual properties of the real world. And most scientists today believe that the geometry of the universe is *not* Euclidean.

Some years ago I wrote a book about so-called *surreal numbers*, which are much richer than real numbers because they include not only the real numbers but also infinitesimally small quantities, as well as numbers like $\omega^{\sqrt{\omega}}$ (infinity raised to the power of square root of infinity). You can add and subtract, multiply and divide surreal numbers, and they're algebraically closed.

In a sense, surreal numbers are actually simpler than real numbers, although they weren't discovered until fairly recently. For example, you can define surreal numbers with only two very simple axioms. I suspect that if physicists had been trained since childhood to work with surreal numbers they would implicitly imagine that surreal numbers describe the actual universe we live in. Certainly surreal numbers are no less likely than real numbers are for this purpose.

It seems to me that a new branch of physics is needed, called maybe "discrete physics" or something like that, to study the effects of the assumption that parameters can be infinitely precise and to consider instead that the universe probably has only a finite-but-extremely-large number of states. I've heard of a few people who are working on this; it seems to me that such ideas deserve to get into the mainstream.

Plato once said, "I have never known a mathematician who was able to reason." I think he was referring to the fact that mathematicians tend to believe that their abstractions apply perfectly to a world that is more complex than they can imagine. On the other hand, according to James Jeans, "Nature abhors accuracy and precision above all things."

Well, I can't dwell any longer on this subject, because I also want to cover several other ideas today. The next topic I want to discuss is John Conway's "Game of Life," as an example of an artificially created universe. It's an idea that John came up with about the same time that he invented surreal numbers in the late 1960s, and it's probably the simplest example of a cellular automaton that has really interesting properties.

We can imagine a grid that consists of square cells extending arbitrarily far in any direction. At every instant of time each cell is either off or on. In Figure 2, for example, the black ones are on and everything else is off. Exactly 197 cells are on.

There's a simple rule for determining the state of each cell at the next instant of time, based on the current state of that cell and its eight neighbors. Suppose k of the neighbors are on at a given time instant t. Then at time $t+1$ (the next instant):

- the cell becomes off, if k is less than 2.
  (If you don't have enough neighbors, you go off.)
- the cell becomes off, if k is greater than 3.
  (If you have too many neighbors, you go off.)

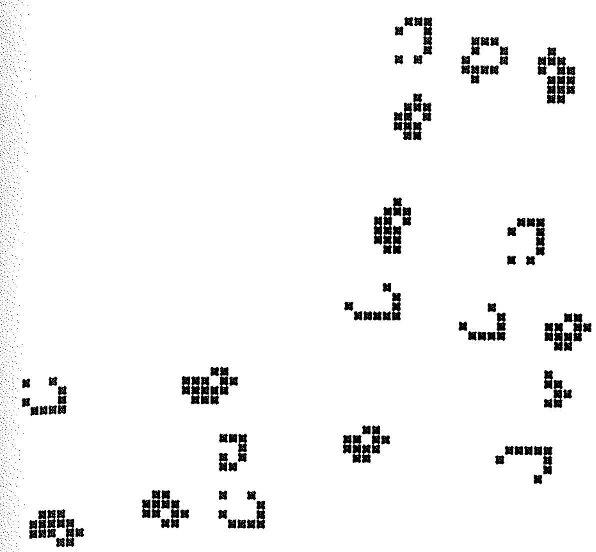

FIGURE 2.

- the cell becomes on, if k is exactly 3.
  (Three neighbors will turn you on.)
- the cell remains in the same state, if k is exactly 2.
  (Two neighbors do not affect you.)

The configuration in Figure 2 is an interesting one that Bill Gosper came up with on November 19, 1997; it's called a "totally aperiodic glider wave." Let's look, for example, at what happens after one unit of time. Several black cells have no neighbors, so they're going to go away. Some of them have too many neighbors; they're going to go away too. But the ones that have exactly two neighbors are going to stay as they were. The cells that have three neighbors are going to go on if they were off, and after one step you get the pattern in Figure 3. Now 202 cells are on.

The Web has wonderful resources for exploring this Game of Life. Just go to your favorite search engine and say "Conway Life," and you'll find an abundance of material, including some terrific Java applets. You can watch the cell patterns to your heart's content. I've heard people say that during the 70s more computer time was spent simulating this game than anything else; many companies banned it from their computers because it was chewing up so many machine cycles. It's quite fascinating to watch what happens.

FIGURE 3.

After 100 steps in this particular case, the cells have begun to send out so-called gliders, and the gliders interact in a really interesting pattern. (See Figure 4, which shows only 176 of the 394 cells that now are on.) The reason Gosper's construction is called "totally aperiodic" is that if you look at any cell whatsoever, its sequence of states is not periodic. No cell by itself could be described by a finite automaton.

I could give several lectures concerning lessons about real life that can be learned by studying examples of artificial life like this. But I have time to mention only a few of the main points. First, it's abundantly clear that a programmer can create something and be totally aware of the laws that are obeyed by the program, and yet be almost totally unaware of the consequences of those laws. Running the program with different starting configurations often leads to really surprising new behavior.

Secondly, the Game of Life illustrates the power of evolutionary mechanisms. Stable configurations arise out of random soup, usually very quickly; and many of those configurations have properties analogous to biological organisms. You can find a glossary of hundreds of names for such things on the Web: Besides the gliders in Figure 4, Lifenthusiasts are familiar with ants, beacons, bookends,

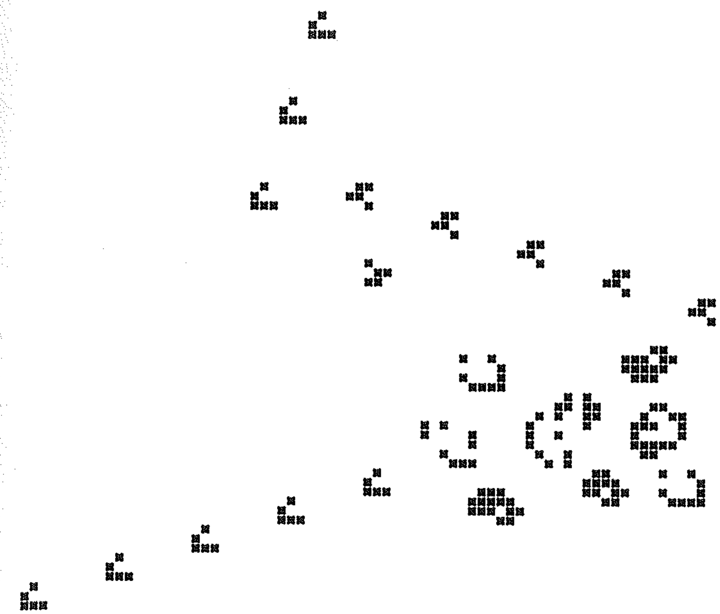

FIGURE 4.

bunnies, caterers, clocks, eaters, lightbulbs, and many other flora and fauna that tend to materialize almost spontaneously.

But the thing that strikes me most about this game is the fact that it is obviously *deterministic*, and I think it sheds light on the age-old question of free will versus a deterministic universe. At least it has helped me a bit with this issue. Somebody, I don't remember who it was (maybe Conway himself), told me in the 1970s that the computer-simulated behavior of patterns in this game tend to be so lifelike that it actually gave him pangs of conscience whenever he would shut the computer off or set it up to work on something else. He was killing off the poor creatures before those creatures had fulfilled their potential.

But Conway's Game of Life is completely deterministic; it needs only simple rules about having two and three neighbors. Thus all the future generations of every pattern must exist whether we simulate them or not. They can't die. Pulling the plug before a computer counts up to a million doesn't harm the number one million.

Now let's imagine that our universe — the real universe — is totally deterministic and finite, but of course extremely large. Conway and Gosper have proved that the Game of Life is *universal*, in the sense that this game can simulate anything that is computable by deterministic laws. Therefore, in principle, we could set up a gigantic on-off pattern that would perfectly describe the future of the entire universe, starting at any given state, if we simply followed Conway's rules. In fact, we could do this on a finite two-dimensional game board whose sides would not be much larger than the number of entities in the universe itself. The simulation wouldn't run in real time; it would in fact be mighty slow. But it wouldn't miss any detail.

That's what a deterministic universe is like. Such a universe exists without needing to be simulated, in the same sense as any number exists without needing to be named.

Raymond Smullyan's short story "Planet Without Laughter" ends with a discussion of free will. I actually like the early parts of that story better than the ending, but Smullyan is certainly entitled to his opinions, and here are some of the words that he put into the mouth of God in his story:

> Humans are like children. The only way you can get them to *do* anything is to make them think that it is *they* who are doing it. Their pride is so great that without having the illusion of free will they will never go forth and amount to anything.

In other words, Smullyan is saying that we don't have free will, but God wants us to think we do, because otherwise the world would get nowhere.

Apparently, then, the *illusion* of free will is good enough, because it's essentially indistinguishable from actually having free will. In a similar way, computer programmers have found that pseudorandom numbers turn out to be just as good as truly random numbers, for all practical purposes. Well, if that is true, the illusion of free will could even have *evolved*, by Darwinian principles.

Albert Einstein didn't believe in freedom of the will. He said, "This awareness ... preserves me from taking too seriously myself and my fellow men as acting and deciding individuals and from losing my temper." Here I cannot agree. Such thoughts are so totally different from my own, that when I first read them I thought they were self-contradictory. If lack of free will kept Einstein from losing his temper, he didn't have any temper to lose. You don't decide whether or not you have free will, if you don't have the power to make decisions.

But after thinking about it some more I had to admit that Einstein's viewpoint is indeed logically consistent, even though I don't subscribe to it myself. In my own view, people ought to take responsibility for the things they choose to do, and I think I can learn to control my temper about other people's choices without believing that they had no choice. Indeed I think it's impossible to be a parent and to observe one's own children *without* believing in free will.

The traditional argument against free will by some theologians is that God cannot know everything unless God knows what choices we're going to make. This debate goes on, but I think the evidence for that argument is fairly weak because of the inadequacies of natural language to deal with such technicalities in a precise way. It seems much more likely to me that God is interested in our decisions, and that he purposely held back from asserting total control. Why create something if you already know the outcome?

As a software designer I greatly enjoy watching surreptitiously what other people are doing with the programs I've written. (Well, sometimes I wince too.) It's similar to the feeling that parents have when their children are developing "attitude." Sometimes they're very happy about what their children come up with, and sometimes they're alarmed, but it's always interesting. Dorothy Sayers said that she enjoyed writing plays better than writing novels, because the actors and actresses would reveal deeper meanings that she hadn't specifically planned.

To carry this discussion further I need to talk a little bit about quantum mechanics. Several years ago, I chanced to open Paul Dirac's famous book on the subject and I was surprised to find out that Dirac was not only an extremely good writer but also that his book was not totally impossible to understand. The biggest surprise, however—actually a shock—was to learn that the things he talks about in that book were completely different from anything I had ever read in *Scientific American* or in any other popular account of the subject. Apparently when physicists talk to physicists, they talk about linear transformations of generalized Hilbert spaces over the complex numbers; observable quantities are eigenvalues and eigenfunctions of Hermitian linear operators. But when physicists talk to the general public they don't dare mention such esoteric things, so they speak instead about particles and spins and such, which are much less than half of the story. No wonder I could never really understand the popular articles.

The extra detail that gets suppressed when quantum physics gets popularized amounts to the fact that, according to quantum mechanics, the universe actually consists of much more data than could ever be observed. Dirac's preface states, for example, that

> [Nature's] fundamental laws do not govern the world as it appears in our mental picture in any direct way, but instead they control a substratum of which we cannot form a mental picture without introducing irrelevancies.

Quantum theories are almost exactly 100 years old now, and they seem to be holding up rather well even though Dirac's book first came out in 1930. James Jeans explained the need for quantum physics in this way:

> At the end of the nineteenth century it first became possible to study the behaviour of single molecules, atoms and electrons. The century had lasted just long enough for science to discover that certain phenomena, radiation and gravitation in particular, defied all attempts at a purely mechanical explanation. While philosophers were still debating whether a machine could be constructed to reproduce the thoughts of Newton, the emotions of Bach or the inspiration of Michelangelo, the average man of science was rapidly becoming convinced that no machine could be constructed to reproduce the light of a candle or the fall of an apple.

Well, I can't give you a tutorial about quantum mechanics today, but I do want to mention it because a lot of computer scientists have been working together with physicists for the past several years to develop something called quantum computing. These concepts have not yet been proved in practice, but steady progress is reported; and it's not impossible that quantum computing could turn out to be a truly revolutionary breakthrough, perhaps allowing us to deal with exponentially many possibilities in linear time. I'll try to describe the situation as simply as I can, without speaking of eigenvalues but still (I hope) conveying some of the essence of the ideas.

Every year at this time for the past six or seven years I've traditionally given a so-called Christmas tree lecture at Stanford, dealing with some aspects of tree structures. Tree structures such as the complete binary tree illustrated in Figure 5 rank among a computer scientist's most beloved concepts. Now please keep in mind that what I'm about to say is an idealization and oversimplification, but

we have to start somewhere. Imagine starting at the top of the diagram in Figure 5, which represents a branching point where some choice has to be made. Also imagine that going to the left or going to the right will affect the entire future history of the universe. First the effect will be very small but eventually the effects will build up. After the first choice (say we go to the left), another decision has to be made—maybe go to the right and then right and then left. The number of possible destinies for the world goes from two to four to eight to sixteen to thirty-two and so on. All of the paths in this tree are consistent with the laws of quantum mechanics, regardless of which choices are taken.

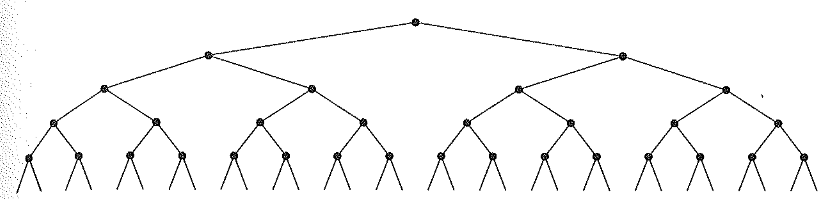

FIGURE 5.

If any of you saw the movie *Run Lola Run* that came from Berlin this year, you'll have a clearer idea of what I'm saying. There are three parts to that movie, and all three parts begin almost exactly the same except that Lola starts to run maybe one second later in the second part and two seconds later in the third. But then what happens in the three parts turns out to be quite different as the stories start to diverge, leading to three completely different endings. In one case the heroine dies; in one case her boyfriend dies; and in the other case they both live happily ever after. It's a fun movie, but it can also give you some idea about what I'm saying here about profound changes in the entire state of the universe, based on very small individual changes that each are consistent with quantum mechanical laws.

The most popular way to account for what we can see in physical experiments is to imagine flipping a coin at each branch point. Einstein made a famous comment that "God doesn't play dice with the universe," and he held to that position for the last thirty years of his life. But nowadays that's definitely a minority opinion. For example, Rustum Roy gave a major lecture in London twenty years ago called "Living with the dice-playing God." David Bartholomew's book *God of Chance*, published in 1984, said that God *uses* chance, because chance offers "many advantages which it is difficult to envisage being obtained in any other way," for example in genetic evolution.

Indeed, computer scientists have proved that certain important computational tasks can be done much more efficiently with random numbers than they could possibly ever be done by any deterministic procedure. Many of today's best computational algorithms, like methods for searching the Internet, are based on randomization. If Einstein's assertion were true, God would be prohibited from using the most powerful methods.

On the other hand, earlier this year Stephen Hawking said, "All the evidence points to [God] being an inveterate gambler who throws the dice on every possible occasion."

Some people, of course, are very suspicious of random choices. For example, Hugh Montefiore in 1985 said, "Chance and necessity may produce creativity but they cannot produce purpose." On this point I believe he was dead wrong: I use random numbers all the time with a very definite purpose, namely to help me discover something. Arthur Peacocke's opinion is that God has perfect knowledge of the *probabilities* of events like radioactive decay, but he doesn't have knowledge about the *outcome* of those events. And Peacocke also says that God created such a universe intentionally.

The picture isn't quite as simple as you might think, however, because quantum theory also implies that the probabilities aren't necessarily independent of each other. They're said to be *entangled*. In fact, quantum theory insists that certain observations have to agree, even though they are individually random, and even though they're being made simultaneously in two completely different parts of the universe. Mind boggling as it seems, quantum mechanics *requires* action at a distance, as if there were instantaneous communication much faster than the speed of light—in spite of what you may have been taught about the theory of relativity.

Yet quantum mechanics doesn't contradict the theory of relativity, because these widely separated events are individually random. Professor Abner Shimony of Boston University says that there is peaceful coexistence between those two theories, because he says there's not really "action at a distance" but rather "passion at a distance"—because of constraints on the way that entangled choices reveal themselves.

Exploiting the counterintuitive properties of entanglement is the basis for hopes about quantum computing, because entangled choices can perhaps be coerced to run through many, many possibilities almost simultaneously and to sort out the good ones based on some kind of resonance. I can't claim to understand much at all about entangled bits; but for me the significance of the probabilistic model for quantum theory is that it clearly makes room for free will, and it allows God to exert dynamic control over the world without violating any laws of physics.

In other words, using the simplified model of Figure 5, we can think of God as a tree pruner, occasionally influencing the outcome of various branches while simultaneously adjusting the nonobservable information behind the scenes so that all observations remain consistent with quantum mechanics. And we ourselves—even us, our spirits or souls or minds or whatever you want to call this part of our being—we might be little tree pruners too, with much more limited and local powers, of course, but still able to exercise free will in this way. Who knows, such a thing might even be *easy*, if God has set us up behind the scenes with some sort of useful hardware, off in the hidden dimensions. The word "spirit," meaning "breath of life," seems curiously appropriate in the context of tree pruning.

James Jeans said in 1930, "For ought we know, or for ought that the new science can say to the contrary, the gods which play the part of fate to the atoms of our brains may be our own minds." Last Sunday I heard another relevant quotation, this one from Martin Luther King, Jr.: "We through our deeds and words, our silence and speech, are constantly writing in the Book of Life." >

Now in closing I want to say a few words about *consciousness*. Stuart Sutherland, in the 1996 edition of the *International Dictionary of Psychology*, gave a great "definition" of that word. Here's what he said.

> Consciousness: The having of perceptions, thoughts and feelings; awareness. The term is impossible to define except in terms that are unintelligible without a grasp of what consciousness means. . . . Consciousness is a fascinating but elusive phenomenon: it is impossible to specify what it is, what it does, or why it evolved. Nothing worth reading has ever been written on it.

Well, during my visit here at MIT I did find one book about consciousness that was at least partly worth reading, namely *Mind and Matter* by the physicist Erwin Schrödinger. Schrödinger's main insight was to equate consciousness with learning; after you've learned something you do it unconsciously.

Yet consciousness remains the largest major question about which science has so far made little or no real progress. Computer scientists studying artificial intelligence may well have the best chance of unraveling this mystery, if it ever can be demystified. The most promising approach that I've heard of is the notion that consciousness might be a kind of genetic algorithm in which a large pool of ideas is constantly competing for attention in our brains. These ideas fertilize each other and the fittest survivors continue the process in Darwinian fashion; as I'm talking to you now, there's all kinds of survival going on inside our heads. Maybe some theory like that is going to work out. It seems to have many of the right properties.

But maybe the reins will still have to be controlled by free spirits acting outside of the observable portion of quantum mechanics. As Peter Gomes has said, we experience "close encounters of the transcendent kind that suggest relationships beyond the power of our experience to reckon, but which we know in some fundamental way to be true."

(Please excuse me for giving so many quotations. As I was preparing these lectures, I ran across lots of things that were said better than I could say them myself, and I couldn't resist giving you the benefit of these other people's wisdom.)

I would like to conclude by quoting one more thing from the end of a talk that James Jeans gave many years ago about a similar topic. He said,

> Every conclusion that has been tentatively put forward [in this lecture] is quite frankly speculative and uncertain. We have tried to discuss whether present-day science has anything to say on certain difficult questions, which are perhaps set for ever beyond the reach of human understanding. We cannot claim to have discovered more than a very faint glimmer of light at the best. . . . [Thus] our main contention can hardly be that the science of today has a pronouncement to make, perhaps it ought rather to be that science should leave off making pronouncements: the river of knowledge has too often turned back on itself.

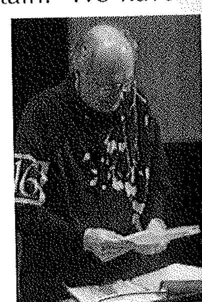

I hasten to add that, yes, we should stop making dogmatic pronouncements, but we should certainly not stop trying to learn more.

Thank you for listening. Once again I'm ready for questions.

Q (Guy Steele): I've got a comment. Twenty-two years ago I took a course here at MIT and the title of the course was "Digital Physics," taught by Ed Fredkin.

A: He was one of the people I had in mind when I mentioned that subject, but I didn't know he had actually taught such a course.

Steele: I distinctively remember a discussion in class about the possibility of modeling the universe as a cellular automaton. Fredkin spoke rapturously, I think, of electrons possibly being represented by billions of cells in some 3-D framework. One of the students in a very distressed voice said, "Wouldn't that take an awfully long time. Wouldn't it be awfully slow?" Fredkin said, "How fast do you want it to be?" And another student piped up and said, "One second per second."

A: The universe could very well be made out of discrete components — very, very tiny — and Fredkin has thought a lot about that. Besides this, a man named Poston in Warwick came up with some extremely interesting ideas that I read awhile ago.

When I study computer algorithms I have to keep track of the difference between continuous approximations and the discrete truth all the time. There's always the ideal thing, given by infinite calculus, to be compared to the reality of the computer program with its finite basis. So I have to say that X is equal to Y plus big-O of Z, and the big-O is what's missing in the books on theoretical physics that I consulted. In my own work *The Art of Computer Programming,* whenever I discuss an algorithm for which the original author has given only an asymptotic analysis in the style of a physicist, I always have to rework the argument to keep track of the errors being made. I think it would be extremely valuable to develop all the theories of physics with the big-O in there, to make rigorous estimates of the magnitudes of errors in your data and in your answer.

In some mathematical arguments, I have not been able to carry out the extra analysis of an appropriate big-O-type estimate. When things like the Borel–Cantelli lemma and other nonconstructive principles like subadditive functions are employed, I have no way to compute bounds on how fast a limiting value is approached. But I believe that the main theorems, the things that apply to reality, almost always allow you to understand the tremendous difference between the assumption that something is infinitely accurate and the assumption that its accuracy is finite but extremely precise (like good to one part in Super K).

Steele: Perhaps another answer to the student's question would be, "Yes, this simulation is slow, and it's slow down to Super K, but that's very small as finite numbers go."

A: Right. Or, my answer might have been, "Yes, it's slow, because it makes the universe operate at a student's speed."

Q: Suppose we make an assumption that at some point in the distant future, everything that our nervous system is capable of perceiving will be computable as a function of the observable state of the universe at that time, using some mechanism that we may not be able to conceive of right now. It's easy to say that computers are not capable of doing such things at the moment, but I find that argument very weak. But if we reach a state when deterministic computability is feasible, do you think that could then in some way negate the possible existence of God?

A: The question, if I understand you, is: If the world is deterministic, does that disprove the existence of God? No, it just disproves the existence of free will. Those are different questions.

But in fact, the idea of recording the entire state of everything in the world is quite incompatible with the theory of quantum mechanics, because modern physics uses complex numbers and all kind of unobservable things that can only be observable (converted from complex to real) if you destroy the information that's hidden.

The physicists have a nice theory about the way things work when they're undisturbed, and then they have this nice theory of what happens when you make a measurement. But I haven't seen any treatment about what happens when you close a loop, so that the physicist who takes a measurement and decides what to do next is also part of the picture, I mean part of the universe that obeys the laws of quantum mechanics. There should be a way to make the picture close on itself, the way it does when we write loops

in computer languages. "While measurement X gives a result less than Y, perform measurement Z." The theory about what happens in such cases wasn't addressed in any of the books I looked at. And I didn't have time to look further or to explore it myself, so I decided to continue writing Volume 4 instead.

Q: What do all these issues of computational complexity, infinity, and the positing of God have to do with good, bad, and purpose?

A: That's a terrific question. I didn't focus on aspects of purpose in my lecture, except for saying that I use randomness for my own purposes, and for talking about some of the reasons why I like to write computer programs and why Dorothy Sayers liked to write plays. I guess I sort of assumed, mistakenly, that God would have certain purposes as if God were a human being.

I know that it's not valid to try to second-guess God with human concepts, but human analogies are the only way I have of trying to understand transcendental things. Certainly there are big questions about such things: Why did the firemen die last week? Shouldn't God have done something so that it didn't happen? We have no idea what would have happened if other branches of the tree had been taken, in my simplified model of choice points for destiny.

I'm afraid that I haven't got any good answers to add to what people have said over the years about such questions. Leibniz said that we live in the best of all possible worlds, with an emphasis on *possible*. In other words, Leibniz argued that God has looked ahead and seen all the worlds that are possible, and there aren't very many that are possible—meaning consistent with the laws of physics. But of the ones that are possible, God chose the best, even though God's choice doesn't necessarily agree with our own idea of best. We just don't know what's possible and what's impossible. That was Leibniz's answer, and there have been many other people thinking about such things over the years. I don't know if computer science has anything to add.

Well, wait a minute, maybe our experience with software design does shed a little more light on the subject. I can design a program that never crashes if I don't give the user any options. And if I allow the user to choose from only a small number of options, limited to things that appear on a menu, I can be sure that nothing anomalous will happen, because each option can be foreseen in advance and its effects can be checked. But if I give the user the ability to write *programs* that will combine with my own program, all hell might break loose. (In this sense the users of Emacs have much more free will than the users of Microsoft Word.)

For example, my TEX system allows users to construct macros that can compute any computable function. This flexibility accounts for the fact that TEX continues to be used in new ways by new groups of people on new generations of printing equipment, even though TEX itself does not change. But the ability for users to add their own arbitrary programs to those of TEX is also dangerous; experience has shown that it is best to restrict TEX so that it cannot change the contents of system files under any circumstances. Otherwise viruses could be constructed to wreak great havoc.

From this standpoint we can see that the symmetry in Figure 5 is misleading, because it suggests that the choice points are fixed in advance and that the tree has a more-or-less uniform structure. I believe God has given us free will not only to choose between a given set of options in a menu, but also to change the options and to substitute new subtrees. My term "tree pruner" is inadequate to describe this vastly more significant aspect of our free will, unless we remember that the left branch of a tree might be totally different from the right branch. The second choice on the left and the second choice on the right might typically govern completely different aspects of the universe. One branch might trigger a nuclear holocaust.

I suppose we could even regard Figure 5 as the Tree of the Knowledge of Good and Evil. As we learn more about science, we find that the fruits of this tree include not only the ability to choose between given options, to choose either "the lady or the tiger," but also the far more dangerous ability to change the basic ground rules of the environment that we live in.

Q: Can you comment on Martin Luther's view of free will?

A: Luther did not believe in a deterministic universe, but he said that the free will of human beings is inherently limited to relatively low-level decisions—what to eat, whether to have children, whether to be Democratic or Republican, whether to give lectures or write books. Above all, he stressed that we cannot will ourselves to have faith; our own reason and strength are insufficient for that, just as the people in Smullyan's story were unable to will themselves to have a sense of humor. Faith is a gift of God, not a product of human free will. We can't control our own thoughts. In my tree-pruning metaphor based on Figure 5, the human tree-pruners have limited scope compared to the tree-pruning ability of God.

Luther was particularly concerned that nobody should be able to boast of having earned their own salvation by means of a choice they had made with their own free will. I agree strongly with that view, but I also think people have an ability to choose to follow God's wishes. Luther was uncomfortable with the latter idea, because he thought people would then be entitled to claim merit for having made the right choice. On the contrary, I think the proper attitude is that people deserve blame for messing up, but they don't deserve praise for just doing their job. The parable of the servant in Luke 17:7–10 brings out this point nicely. My view is essentially that God has provided us with some kind of hardware by which we can exercise our limited free will. We are doing the job expected of us if we control it well, but we are forgiven if we don't.

Q: What makes you think that the Bible is the word of God? Or is that an axiom chosen at random?

A: The assumption that the Bible is God's word is an unprovable axiom that I tend to find confirmed as I look at it. I don't treat the words of the Bible literally as axioms though. I don't say, for example, that the Book of Revelation is a set of axioms from which I could deduce theorems as Euclid did in geometry.

I do think that the Bible reflects God's messages. As time goes by, we understand more and more about the way in which it was written and the historical process it has gone through, the difficulties of transmission over thousands of years. I think the Bible holds some of the best clues other than the universe itself as to what God wants us to do; I spoke about that last week. But I would never expect to be able to persuade anyone else about this hypothesis by using a mathematical argument.

I also have great respect for several other religions, and I believe God is speaking in their scriptures as well.

Q: It's hard to express this question because I'm not sure I can find the right terms. Let's take "material" to mean those things describable by physics or understandable by physics. If that's the case, do you see God as standing outside the material universe or coexistent with it?

A: I don't think of God as standing outside the universe letting it run, I see God as active all the way though.

In fact, the way I view God myself has a lot in common with pantheism, where God is everywhere. Some people think monotheism and pantheism are diametrically opposed concepts, but we can see from the standpoint of computer programming that they aren't essentially different. I mean, we can think of distributed computing, when there are lots of cooperating processes but really only one program — in the same way as an ant colony can be thought of as a single organism, or in the same way as the individual neurons of our brains are part of a single person.

Q: So God is "material" in that sense?

A: Well, I guess not; it's in another dimension, I think, orthogonal to physics. You do understand that if we lived in a two-dimensional universe there would be 3-D things that we couldn't perceive.

Q: Then it will never be describable by physicists?

A: That's true. But I don't think that's bad. If you could describe it, it would be kind of boring after a week.

Q: You have this discrete deterministic model, which is a kind of computer, and we know the speedup theorem that there's no way to know what's going to happen without its happening. So doesn't that mean that the determinism you have is really quite a bit different than the kind people think of where you can calculate, say, the next position of the planets any time in the future?

A: That's an interesting point. Since classical physics deals with infinite-precision real numbers, their deterministic models are not realizable on actual computers.

There's also a classical question, "If a tree falls in the forest and nobody can witness it (because maybe the forest is in a black hole), did the tree really fall?" In the setup I talked about, there's a similar but maybe more subtle question: "If nobody has ever named the numbers up to Super K, do they exist?"

I'm quite willing to believe that the number five exists, and six and seven and so on, and all the numbers I can count up to. But does that mean that all finite numbers have an independent existence? I implied as much when I said that the deterministic results of the Game of Life need not be simulated because they have

always been present. It's something like if all the jokes in the world were numbered, and a person would come up to you and say "37!" and you would roar with laughter. Our universe would be number such-and-such in an appropriate encoding system, if it were finite, yet nobody could ever know that number.

I may therefore be guilty of the same sort of error that I've ascribed to physicists and philosophers. Why should I consider a sharp distinction between determinism and nondeterminism, or between existence and nonexistence, to be any more relevant than a distinction between finite and infinite? There might well be degrees of existence, when we consider the computational complexity of naming things.

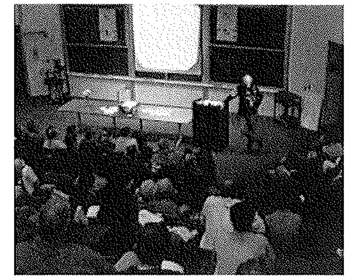

Suppose, for example, there exists a million-digit number K_w that contains the key to wisdom. But suppose further that it will take at least 100 years to compute K_w, even if we use all of the world's fastest computers. Should the United States Patent Office grant a patent on the numerical value of that integer, once it has been computed, even though purely mathematical things are not patentable, and even though one surely cannot patent a simple number like 2009?

Your question also reminds me of a completely different thing, the so-called "many-worlds hypothesis." I tend to think that it's too wild to be true, but some people have used such an assumption to account for entanglement, where they say that all of the branches in Figure 5 actually are present. In other words, as we're talking right now the universe is cloning itself into many almost-identical copies of itself, and we all exist in all those branches. And all three of the things that happened to Lola happened. Because all branches of the trees are solutions to the equations of quantum mechanics. This is the way a computer scientist looks at nondeterministic computation, as if all of them are going on. And we think we have free will in each thread.

Now we have no way to test this hypothesis; in some of those worlds we just die earlier. And then perhaps you have to ask the question "What is consciousness?" in a new way, because each of our individual consciousnesses is different on each branch of the tree. Lots of speculation of this abstract kind has been going on, but as I say it does seem quite far out to me. I probably didn't answer your question, but those are some of the complexities that are associated with it.

Q (David Rosenberg): This is sort of a follow-on to the last two questions. Rather than state it in general form, I'm going to pose only one specific example. Are you part of God?

A: Am I part of God? I'm not sure exactly what you're getting at, but it's Christian dogma that God is in us and we're in God. That is why, for example, the Bible verse 1 Thessalonians 3:16 says, "Refrain from sexual immorality because you're part of God." In other words, that's part of my motivation. I'm not an antagonist of God, because we're part of the same team. But I'm not entirely a subset of God in the mathematical sense, because I do think I have a free will that's independent of God.

Q: Is there any sense to which it can be true then that each of us is all of God?

A: Do you mean each of us individually?

Q: Yes.

A: Well I suppose there exists the sense, but I certainly don't feel like God.

Q: In your mathematical understanding of things, do you find a demarcation between running an experiment, as the question was asked back there, and a Gedanken experiment which you merely think through what might happen as Einstein did?

A: The distinction between running an experiment and thinking about the experiment? Are you assuming that my thought processes are accurate? Mine aren't, but maybe a perfect computer or perfect creator would be able to go through and wouldn't have to run anything.

My personal view is that God is dynamically involved and actively interested in the choices that we make. But, as I say, there is no way to prove or disprove any of these things.

I see it though like . . . it's a weak analogy, but still about as good as I can come up with from the standpoint of a software user . . . like when Joel Moses created the MACSYMA system at MIT twenty years ago. I used to log in from California to use MACSYMA, and people at MIT were often watching me as I was using it. And I was glad to know that they were following what I was doing, because every once in awhile I'd get stuck, but then my terminal would magically type out a message like "Try the `TELLRAT` function." Now *that's* the kind of attitude I'm thinking God is using with me in life.

Q: My question was somewhat different, tied in with the paradox of Schrödinger's cat.

A: Oh. I've heard that paradox a couple of times, but there's something about a cat dying and I hate to think about such things.

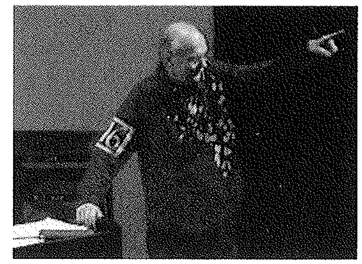

Q: What are your thoughts on prayer? Is prayer an effective way to communicate?

A: My thoughts on prayer are complicated and hard to verbalize, but I'll try to explain them. I don't believe in prayer for selfish things, where I would pray in order to get an advantage of some kind. I think of prayer as a conversation with God. Even if I didn't know that prayer was effective I would do it anyway. It's just something that feels natural.

When I was a child I thought about prayer in a very selfish way. I remember quite vividly being told that, if you ask God for something, and if you truly believe in Him, then He will grant your wish. And I wanted a Ferris wheel. (Really, it's true: I was visiting relatives in Cleveland, and we had gone to an amusement park, and I fell in love with the Ferris wheel.) So that night when I went to sleep and said my prayers, I said, "Please God, put a Ferris wheel on the front lawn tomorrow morning." I'm serious. And the next morning I ran to the window, fully expecting to see it there. I was devastated to see that the lawn was still empty.

That, I say, is a completely wrong attitude to have about prayer. But now I understand that prayer is a conversation. Maybe I hope that God will do something, but I don't pray in order to get an extra bonus. It's just something I feel like doing. When I give a prayer of thanks, it's essentially like the way I feel when I'm hugging my wife. And when I pray, I might also express disappointment or anger, like when I'm arguing with my wife. Communication is much better than keeping your feelings bottled up.

Q: In a nondeterministic world, let's say there are two brothers who are twins. Their behavior cannot be predicted, but if you look at the correlation between them they have a perfect correlation. So if you observe one's behavior you kind of know what the other is doing at the same time. Or let's suppose they have 90% correlation in their behavior with each other. In such a case, do you consider that each of them have free will individually, or do you consider that together they have a free will?

A: Well, of course I've never known twins whose behaviors are correlated in this way. I know that when identical twins are separated at birth and raised in completely different environments, they still turn out to have mannerisms in common, probably because of their genetic makeup.

But I guess I'm interpreting your question too literally. What you are saying is really a metaphor for entanglement, but with respect to human beings rather than to abstract bits. Then I'd say it is certainly reasonable to imagine that either their souls are correlated, or that consciousness without a soul is governed by very similar mechanisms. I haven't thought about this enough to say more.

Q (Doug Ross): Sometimes our free will gets tangled up with the real world. And the simplest example I know of is the fact that you so often get documents that say, "This page intentionally left blank." Do you have a comment about how to solve that problem in a proper way?

A: What you have to say is, "The *previous* page was intentionally left blank."

Q: I want to sort of try to tie together a lot of the comments you've made today. On the one hand, the sense seemed to be along the lines of peeling away more or less logical objections that some people would have to the existence of God. On the other hand, you've spoken of resonant images or metaphors, or whatever, that help you conceive of God. The gap I'm experiencing a little bit is that at least my own way of coming to a religious sensibility

was much more experiential; I had to get rid of some intellectual objections, while probably not as erudite as the kind of things you've been mentioning. So far, your comments don't seem to connect with the relationship aspect of religious belief experience. This is not so much a question as maybe a point of departure.

A: I focused my talk on things where computer science might contribute new notions to the mix of ideas that I had seen put forward by scientists in other disciplines. When I had read comments by physicists, biologists, and chemists, I would think of other points that I imagine a computer scientist would be more likely to bring up.

One of the things I explored while I was preparing these lectures might be closer to what you want to hear, although I had decided not to bring it up today because it didn't fit directly with the other topics. Earlier this year, I had been somewhat surprised to learn that Plato had once discussed something called "divine madness"—the Greek words are *theía manía*. He had Socrates speak about a God-given state of being beside oneself, which is manifested in four different ways: prophecy, catharsis, poesy, and eros. I learned a little bit about such things from the reading that I've been doing these past few weeks, and I think they add another significant dimension to the whole mystery of consciousness. But I didn't think computer science helped me understand *theía manía* in any special way.

Q: What was the thing you promised to talk about last week but didn't have time to discuss today?

A: That's a sneaky question, but I'm glad you asked. I had originally thought that I would discuss some experiences that I had when I was designing a computer chip. But this week I decided to leave that story out, because it relates more to economics than to God. I'll tell you anyway, because you asked for it, and because it does have possible connections to the organization of the universe.

So here's the story. As I was designing a chip for a very simple RISC computer, I was surprised to find that the easiest and somehow the best way to design this chip was to have it doing all kinds of things that would never be needed afterwards. I mean, two binary numbers were input to the chip at each clock cycle, and the adder would add them and the subtracter would simultaneously subtract them, and the multiplier would multiply them. These things were all going on at once inside the chip, but only one of those results would survive and actually be used in the computation in the next step. In this way the chips were operating quite differently from the computer programs I had been writing before.

The alternative would have been to design the chip so that every circuit inside the multiplier had extra inhibitors on it saying, "Don't multiply unless I tell you to." That would add an awful lot to the hardware.

I started thinking about this as an interesting metaphor for society and the world in general. It might help us to reformulate our notions of "purpose." There is good reason for a thousand people to work on a problem even though only one of them is going to solve it, and even though the people know in advance that only one of them is going to influence the final answer. Everybody can take pride in what they do even if it doesn't show up in the next generation, because otherwise less would get done and people would be idle.

You might find it interesting to muse about that a little bit.

Q: It's easier to let people make mistakes than to have such a tightly controlled environment that no one can make mistakes.

A: Yes. I remember something Dijkstra told me once. On his first visit to Stanford he saw the linear accelerator, and he just couldn't believe so many experiments were going on simultaneously. He thought you should run the accelerator for an hour and then sit and think for a year about what you learned; then it would be time to turn the machine on again. The other attitude was the American way instead of the Dutch way.

Q: Since there is time for only one more question, I want to ask you to mention one more thing that you had to cut from the main part of your lecture for reasons of time.

A: Well, I noticed debates in the faith and science community about "top-down causality" versus "bottom-up causality." I thought I might point out that my research on so-called *attribute grammars* shows that top-down and bottom-up causality can coexist quite nicely. But I never had time to explore that in depth.

In conclusion I want once again to thank the MIT administration in general and Anne Foerst in particular for inviting me to give these lectures. I thank you all for coming and asking such excellent questions. My fondest hope is that you continue to seek answers, even though the questions may be unanswerable.

Notes on Lecture 6

Page 167, improvisor of unsurpassed ingenuity: Arthur Peacocke, *Creation and the World of Science* (Oxford: Clarendon Press, 1979), Chapter 3. See also Arthur Peacocke, "Chance and law in irreversible thermodynamics, theoretical biology, and theology," in *Chaos and Complexity: Scientific Perspectives on Divine Action*, edited by Robert John Russell, Nancey Murphy, and Arthur R. Peacocke (Vatican City State: Vatican Observatory Publications, 1995), 123–143, especially page 140.

Page 168, art is everything else: See the reference to my paper "Computer programming as an art" in the notes to Lecture 4.

Page 168, panel discussion: See below.

Page 168, I personally see no signs: This in spite of recent books by people I respect, such as *Robot: Mere Machine to Transcendent Mind* by Hans P. Moravec (New York: Oxford University Press, 1998); *The Age of Spiritual Machines* by Ray Kurzweil (New York: Viking Press, 1999).

Page 168, Brian Hayes: "Computational creationism," *American Scientist* **87** (September–October 1999), 392–396.

Page 169, probably the most difficult: Donald E. Knuth, *MMIXware* (Heidelberg: Springer, 1999), especially pages 150–331.

Page 169, reprinted in several books: Donald E. Knuth, "Coping with finiteness," *Science* **194** (17 December 1976), 1235–1242. Reprinted with corrections in *Electronics, the Continuing Revolution*, edited by Philip H. Abelson and Allen L. Hammond, AAAS publication **77-4** (Washington, D.C.: American Association for the Advancement of Science, 1977), 189–196; and in *Mathematics: People, Problems, Results*, edited by Douglas M. Campbell and John C. Higgins, volume 2 (Belmont, California: Wadsworth, 1984), 209–222. Bulgarian translation by G. Chobanov and Z. Dokova in *Fiziko-Matematichesko*

Spisanie **21** (Sofia, 1978), 58–74. German translation by Arthur Engel in *Der Mathematik-Unterricht* **25**, 6 (1979), 5–26. Reprinted with corrections as Chapter 2 of *Selected Papers on Computer Science* (Stanford, California: Center for the Study of Language and Information, 1996), 31–57.

Page 171, because of a theorem: See, for example, Ming Li and Paul Vitányi, *An Introduction to Kolmogorov Complexity and Its Applications*, second edition (New York: Springer, 1997), Theorem 2.2.1.

Page 172, Cantor's famous theory: Georg Cantor, "Über eine Eigenschaft des Inbegriffes aller reellen algebraischen Zahlen," *Journal für die reine und angewandte Mathematik* **77** (1874), 258–262. See also Joseph W. Dauben, *Georg Cantor: His Mathematics and Philosophy of the Infinite* (Cambridge, Massachusetts: Harvard University Press, 1979).

Page 172, three times in the King James Bible: Job 33:5, Psalm 147:5, Nahum 3:9.

Page 173, Peter Gomes: *The Good Book* (see Lecture 3), page 313.

Page 173, a rather deep result: Larry Joseph Stockmeyer, *The Complexity of Decision Problems in Automata Theory and Logic*, report MAC TR-133 (Ph.D. thesis, Massachusetts Institute of Technology, 1974), Chapter 6. See also the exposition in my paper "Coping with finiteness" cited earlier.

Page 173, theorem of Büchi: J. R. Büchi, "Weak second-order arithmetic and finite automata," *Zeitschrift für Mathematische Logik und Grundlagen der Mathematik* **6** (1960), 66–92.

Page 175, tremendous leap of faith: Coincidentally, I happened to see a car with the bumper sticker

Question Reality

when I returned to California.

Page 175, *surreal numbers*: See Donald E. Knuth, *Surreal Numbers* (Reading, Massachusetts: Addison–Wesley, 1974).

Page 176, Plato: *The Republic*, vii:531e.

Page 176, James Jeans: *The Mysterious Universe* (Cambridge University Press, 1932), page 30.

Page 176, Game of Life: Martin Gardner's original column "Mathematical Games: The fantastic combinations of John Conway's new solitaire game 'life,'" *Scientific American* **223**, 4 (October 1970), 120–123, in which Conway's game was introduced to the world, has been reprinted with extensive additions in Gardner's book *Wheels, Life and Other Mathematical Amusements* (New York: W. H. Freeman, 1983), Chapters 20–22. See also Conway's own description in *Winning Ways* by Elwyn R. Berlekamp, John H. Conway, and Richard K. Guy (London: Academic Press, 1982), Chapter 25.

Page 176, probably the simplest example: An even simpler example of a universal cellular automaton was discovered by E. R. Banks in his Ph.D. thesis at MIT in 1971. The transition rules in his system involve only the four nearest neighbors of a cell: If three or four of those neighbors are on, the cell goes on; if exactly two are on and they are adjacent, the cell goes off; otherwise the cell's state doesn't change. However, this scheme apparently needs to have infinitely many cells in each of the two states, while Conway's scheme requires only finitely many cells to be on. See Edwin Roger Banks, "Universality in cellular automata," *Symposium on Switching and Automata Theory* **11** (1970), 194–215.

Page 177, an abundance of material: For example, a good starting point is "Paul's Page of Conway's Life Miscellany," by Paul Callahan, `www.cs.jhu.edu/~callahan/lifepage.html`.

Page 178, artificial life: For a stimulating introduction to this active area of research, see *Artificial Life* by Steven Levy (New York: Pantheon, 1992).

Page 180, Planet Without Laughter: See the end of Lecture 4. This quotation begins on page 181 (slightly paraphrased).

Page 180, Albert Einstein: The quotation is from a speech, "My credo," to the German League of Human Rights, Berlin, in 1932. See *Einstein: A Life in Science* by Michael White and John Gribbin (New York: Dutton, 1993), page 262.

Page 181, she enjoyed writing plays better: Dorothy L. Sayers, *The Mind of the Maker* (London: Methuen, 1941), page 64.

Page 181, Paul Dirac's famous book: P. A. M. Dirac, *The Principles of Quantum Mechanics* (Oxford: Clarendon Press, 1930).

Page 182, Jeans explained the need: James Jeans, *The Mysterious Universe* (Cambridge University Press, 1932), page 22.

Page 183, *Run Lola Run*: A movie written and directed by Tom Tykwer, 1999; see the webpage `www.spe.sony.com/classics/runlolarun/runlolarun.html`. Or, in the original German, *Lola Rennt*, `www.lolarennt.de`.

Page 183, God doesn't play dice: Albert Einstein, letter to Max Born, 4 December 1926: "Jedenfalls bin ich überzeugt, dass *der* nicht würfelt," in *Albert Einstein, Hedwig und Max Born, Briefwechsel* (Munich: Nymphenburger, 1969), page 130. English translation by Irene Born, *The Born–Einstein Letters* (London: MacMillan, 1971).

Page 183, Rustum Roy: *Experimenting with Truth: The Fusion of Religion with Theology Needed for Humanity's Survival* (Oxford: Pergamon, 1981), page 188.

Page 183, David Bartholomew's book: David J. Bartholomew, *God of Chance* (London: SCM Press, 1984).

Page 184, more efficiently with random numbers: See the first note on Lecture 2. I must admit, however, that this statement is true only in the limit as our horizon becomes infinite. Deterministic rules of finite complexity Super K would be quite adequate for our universe.

Page 184, methods for searching the Internet: See Andrei Z. Broder, Moses Charikar, Alan M. Frieze, and Michael Mitzenmacher, "Min-wise independent permutations," *Journal of Computer and System Sciences* **60** (2000), 630–659.

Page 184, Stephen Hawking: "Does God play dice?" `www.damtp.cam.ac.uk/user/hawking/dice.html`, 12 July 1999.

Page 184, Hugh Montefiore: *The Probability of God* (London: SCM Press, 1985), page 98.

Page 184, Arthur Peacocke's opinion: See "God's interaction with the world: The implication of deterministic 'chaos' and of interconnected and interdependent complexity," in *Chaos and Complexity* (as cited above), 263–287, especially page 280.

Page 184, Abner Shimony: I found this in "Why God plays dice" by Mark Buchanan, *New Scientist* **159**, 2148 (22 August 1998), 26–30.

Page 185, James Jeans: *The Mysterious Universe* (as cited above), page 36.

Page 185, Martin Luther King, Jr.: I have not been able to trace the source of this quotation, which was part of last Sunday's sermon at the Methodist Church near Harvard Square. But King did say, "There is an invisible book of life that faithfully records our vigilance or our neglect," in an address given at Riverside Church in 1967; see *A Testament of Hope*, edited by James Melvin Washington (San Francisco: Harper & Row, 1986), pages 243 and 633.

Page 185, Stuart Sutherland: *The International Dictionary of Psychology* by N. S. Sutherland (New York: Crossroad, 1996).

Page 185, Erwin Schrödinger: *Mind and Matter* (Cambridge University Press, 1959).

Page 186, competing for attention: See, for example, William H. Calvin, *The Cerebral Code* (Cambridge, Massachusetts: MIT Press, 1996); William H. Calvin and Derek Bickerton, *Lingua ex Machina: Reconciling Darwin and Chomsky with the Human Brain* (Cambridge Massachusetts: MIT Press, 2000).

Page 186, Peter Gomes: *The Good Book* (as cited above), page 214. "There is in Celtic mythology the notion of 'thin places' in the universe, where the visible and the invisible world come into their closest proximity."

Page 186, James Jeans: *The Mysterious Universe* (as cited above), pages 187–188.

Page 187, Digital Physics: See Edward Fredkin, "Digital mechanics," *Physica* **D45** (1990), 254–270; "A physicist's model of computation," in *Massive Neutrinos — Tests of fundamental symmetries*, 26th Rencontre de Moriond (1991), 285–197; "Finite nature," in *Progress in Atomic Physics, Neutrinos and Gravitation*, 27th Rencontre de Moriond (1992), 345–354; "A new cosmogony," *PhysComp '92: Proceedings of the Workshop on Physics and Computation* (Los Alamitos, California: IEEE Computer Society Press, 1992), 116–121. A website for discussion of further developments related to Fredkin's cellular-automata-based approach has been set up by Joel C. Dobrzelewski at `cvm.msu.edu/~dobrzele/dp`.

Page 187, Fredkin has thought a lot about that: To my surprise I was greeted after this lecture by Fredkin himself, who I hadn't seen for more than 20 years; although he was in the audience, he chose to remain incognito and silent during the discussion.

Page 187, Poston: See Tim Poston, *Fuzzy Geometry* (Ph.D. thesis, Warwick University, 1971), 177 pages. An informal introduction to this work appeared in *Manifold* **10** (Autumn 1971), 25–33.

Page 189, the firemen: The big news story at the time of this lecture concerned the tragic deaths of six firemen in Worcester, Massachusetts on 3 December 1999.

Page 189, many other people: The best discussion I know on this question appears in Chapters 15 and 16 of Martin Gardner's book *The Whys of a Philosophical Scrivener* (New York: William Morrow, 1983); a second edition, published by St. Martin's [sic] Press in 1999, adds 32 pages of new notes.

Page 189, wait a minute: These comments did not arise during the public lecture but in a conversation with members of the audience immediately afterwards.

Page 190, Emacs versus Microsoft Word: The "E" in "Emacs" stands for "extensible," meaning that users can readily add new features of their own design. Microsoft Word, by contrast, allows only features that have been provided in advance, so that users who want more features must wait to purchase a new release of the program. (My analogy would therefore suggest that Microsoft Word should be considerably more reliable than Emacs. This, unfortunately, is not the case, but for entirely different reasons having to do with the concept of open source software.)

Page 190, the Tree of the Knowledge of Good and Evil: Genesis 2:17. This tree is "in the midst of the garden" (Genesis 3:3).

Page 190, Martin Luther's view: This question, likewise, came up after the lecture. I thank Anne Foerst and Jon Allen for presenting me with a copy of Luther's *The Bondage of the Will*, translated by J. I. Packer and O. R. Johnston (London: James Clarke, 1957); in the original Latin it was *De Servo Arbitrio* (Wittenberg: Johann Lufft, 1525).

Page 191, parable of the servant: See the hymn in the notes following Lecture 3.

Page 192, the speedup theorem: I'm not sure exactly which speedup theorem the questioner had in mind. But it is easy to prove from the incompressibility theorem stated earlier that, for example, Fredkin's simulation of the universe by cellular automata operating in the universe could not run at a rate of 1.001 seconds per second.

Page 193, degrees of existence: For example, Shannon's information theory has proved to be hopelessly inadequate to express principles of cryptography; therefore that theory needed to be reformulated by taking computational models into account. See, for example, Andrew C. Yao, "Theory and applications of trapdoor functions," *IEEE Symposium on Foundations of Computer Science* **23** (1982), 80–91; M. Blum and S. Micali, "How to generate cryptographically strong sequences of pseudo-random bits," *IEEE Symposium on Foundations of Computer Science* **23** (1982), 112–117.

Page 194, Christian dogma: See, for example, Isaiah 7:14; John 6:56, 14:20; 1 Corinthians 6:15, 19.

Page 194, 1 Thessalonians 3:16: Actually 1 Thessalonians 4:3.

Page 194, Gedanken experiment: The questioner was probably referring to the famous paper by A. Einstein, B. Podolsky, and N. Rosen, "Can quantum-mechanical description of physical reality be considered complete?" *Physical Review* **47** (1935), 777–780. But I have a poor knowledge of physics and I didn't understand the question.

Page 194, MACSYMA: The Mathlab Group, *MACSYMA Reference Manual*, version six (Cambridge, Massachusetts: Massachusetts Institute of Technology, Project MAC, 1974).

Page 195, Schrödinger's cat: E. Schrödinger, "Die gegenwärtige Situation in der Quantenmechanik," *Die Naturwissenschaften* **23** (1935), 807–812, 823–828, 844–849; English translation, "The present situation in quantum mechanics," in *Quantum Theory and Measurement* (cited in Lecture 5), pages 152–167. This was Erwin Schrödinger's response to the paper of Einstein, Podolsky, and Rosen. According to Fredkin's paper of 1990 in *Physica*, cited above, cellular automata allow us to view such questions in a new way.

Page 197, Plato: See, for example, Josef Pieper, *Divine Madness* (San Francisco, California: Ignatius Press, 1995). "Divine madness" is sort of the opposite of what Governor Festus meant when he shouted, "Paul, you are out of your mind; all that learning of yours is driving you mad!" (Acts 26:24).

Page 197, very simple RISC computer: See Donald E. Knuth, *The Stanford GraphBase* (New York: ACM Press, 1994), 238–259.

Page 198, Dijkstra: Edsger W. Dijkstra. As far as I know he never has expressed such sentiments in print, but his general approaches to computer science problems have had an enormously beneficial effect as counterbalances to "conventional wisdom." See, for example, his *Selected Writings on Computing: A Personal Perspective* (New York: Springer-Verlag, 1982).

Page 198, *attribute grammars*: See, for example, D. E. Knuth, "The genesis of attribute grammars," *Lecture Notes in Computer Science* **461** (1990), 1–12, and the other papers in that volume.

PANEL DISCUSSION: CREATIVITY, SPIRITUALITY, AND COMPUTER SCIENCE (17 NOVEMBER 1999)

Anne Foerst: As part of this year's lecture series on God & Computers, we decided to use the time when Don Knuth happens to be visiting our campus to invite other accomplished computer scientists to contribute their opinions on the whole interaction of religion, God, and computer science. We now have in front of us five very famous, very good, and very *interesting* (which is most important for me), distinguished computer scientists.

But one of them is an exception, because Harry Lewis from Harvard University is actually the moderator. Even though Harry is a well-accomplished Harvard professor of computer theory, he also happens to be Dean of Harvard College, and therefore very much accustomed to interacting with people, forcing them to tolerance, making them shut up when they talk too long, and so on. At the same time he's also competent enough to know when something is important.

Therefore I welcome Harry very very much, and I welcome the other four panelists very very much, thanking you for having the courage to be here to talk about an area of your life that is not part of your professional work and not often openly discussed. I thank everyone for coming, and I hope that we all will have a wonderful interaction.

Lewis: Thank you and welcome to this panel on the subject of Creativity, Spirituality, and Computer Science. I am very grateful for the four distinguished guests who have taken the time to share their afternoon with all of us.

Let me tell you what I would like the format of this to be. First I am going to introduce the four panelists to you. Then I will ask each of them, in the order in which they are seated to my right, to make a few introductory remarks on the subject of the panel. Then we will probably have some cross talk among them, exchanging