

Andrew Glassner's Notebook

<http://www.glassner.com>

Quantum Computing, Part 1

Andrew
Glassner

There's a revolution coming in the field of computing, and it's coming from the smallest of all places: the subatomic particles that form the basis of all matter.

Contrary to what some of us learned in high school, photons aren't billiard balls and electrons aren't little planets circling a nuclear sun. Those conventional models work well up to a point, but when you look more closely, they begin to break down and a new picture emerges. And from this picture, we obtain a new way to process information: the quantum computer.

Quantum computers will be useful for all kinds of computing. In addition to computer graphics they will have applications in cryptography, biology, and communications. All that makes them worth looking at, but even better, this is just incredibly cool stuff. The world of quantum physics is strange and beautiful. It's full of all sorts of interesting philosophical questions, including the nature of reality and consciousness. We won't get into all those lofty questions here, but we'll see some of the bizarre ways that reality behaves at the quantum level.

In this column, we'll get started with a look at the general principles of quantum mechanics as they apply to quantum computing. In the next issue, I'll present some of the math behind actually creating quantum algorithms.

You may wonder why I'm tackling this topic here, since it doesn't look like mainstream computer graphics. The reason is that quantum computing is an exciting, new way to think about algorithms of all sorts, including those for graphics. Someday we'll all be using quantum computers, and writing quantum algorithms. I'm just getting us off to an early start!

Before we get going, I need to point out that quantum computing is still in its infancy. The most powerful quantum computer built to date is a scant 3 bits wide. But these tiny machines have proven the basic theory's validity, from computation to error correction. The future is ripe with promise.

Quantum physics explained

The behavior of subatomic particles is the subject of quantum physics, so named because at this incredibly small scale, everything seems quantized, or broken up into tiny pieces. Energy, space, and perhaps even time aren't continuous but are made up of tiny discrete units, like beads on an abacus.

Quantum physics has a well-deserved reputation for being counterintuitive. The behavior of the subatomic world comes as a surprise to almost everyone who encounters it. Physics, after all, has developed like most of the other physical sciences: people study a phenomenon, hypothesize about how they think it's working, create a mathematical description of that mechanism, and then check to see how well that math matches (and predicts) reality. Human intuition and insight often drive this process. For example, if we take a ball to the park and roll it on the grass, it will roll along for a while before slowing down and stopping. Why doesn't it roll forever? Studying the problem, we can slowly build a picture of the world that includes things like how grass and dirt create friction and how gravity makes it hard for the ball to roll uphill. Our explanation matches our intuition of the world, and we've learned some physics.

In the subatomic world, our intuition often proves misleading. Stuff goes on that looks and sounds like magic: objects are in several places at once, a cat can be dead and alive at the same time, and you can instantly change the state of a particle that's halfway across the universe. This is just the tip of the weird iceberg that quantum mechanics describes.

Amidst all these oddities, we must remember two important truths. First, as weird as it is, quantum physics works. The math describes the reality, even if we can't explain why it does so in a satisfying way. Second, we can use these weird properties to make powerful computers.

Today we use *classical computers*. The coming machines based on quantum mechanics are called *quantum computers*. Here are three examples of what quantum computers will do for us one day:

- *Fast Z buffers*: The Z buffer is a great way to figure out visibility. Suppose you have an unordered list of N opaque objects at a pixel, and you need to find the closest one. On a classical computer you have to look through all N objects, which require at least N steps. On a quantum computer, you can find the nearest object in order \sqrt{N} steps.
- *Instant radiosity*: Radiosity is one way to compute a synthetic image. We describe the world as a collection of volumes and surfaces and follow the light in the scene as it propagates, hitting one surface and then another, around and around. We've seen lots of clever

algorithms for speeding this up, but high-quality radioactivity still requires a lot of computing time and energy. Quantum computers can compute a radioactivity solution in only one step, taking a tiny fraction of a second. As we'll see later, we'll typically take a few more steps to process the solution before it's useful, but if N is the number of surface elements in the scene, it will take no more than order \sqrt{N} steps to locate the complete solution. (That estimate uses today's theories; tomorrow we might be able to do it even faster.)

- **Double-speed transmission:** Suppose you've computed an image and you want to send it to a friend. (For simplicity, let's assume you don't want to compress it.) To send the image requires you to transmit N bits: the product of the width times the height times the number of bits per pixel. With a quantum computer, if you've shared $N/2$ quantum particles at any time in the past (even years ago), then to send the image now you'll only need to send $N/2$ bits.

Shopping for a quantum car

I'll forego a mathematical treatment to first get the ideas in place; we'll get to the details next time. Let's begin with a rough (and admittedly unrealistic) analogy: shopping for a quantum car.

Suppose that we want to buy a new car, and we've settled on a particular make and model. This particular model has three options: the number of doors, the body color, and whether it's a convertible or sedan. When we choose our options, we specify a particular configuration of the car, which I'll call a *state*.

Let's indicate each of these three choices with a variable. We'll set $f=0$ if we get a two-door car and $f=1$ if it's a four-door. We'll set $b=0$ if the color is red and $b=1$ if it's blue. Finally, $s=0$ if the car is a convertible and $s=1$ if it's a sedan with a hard top. Taken together, these three choices define the car's state. If we write the vector $[fbs]$, then we can name any of the eight possible states.

For example, if we set $f=1$, $b=0$, and $s=0$, then we have the state $[100]$ indicating a four-door red sedan. The important thing to note is that we need only three numbers to define a state.

Now suppose that on entering the showroom, you discover something strange. The dealer tells you that he has only one car but it's completely hidden under a blanket. The dealer tells you that nobody has ever looked under the blanket. In fact, he tells you that during the middle of the night, when nobody was around, robots in the showroom made the car. They were programmed to start themselves up in the middle of the night, select a state at random using an unpredictable random-number generator, build that particular version of the car, and then place the blanket over it. Hence, nobody has any idea what state the car is in.

I'll assume that the blanket is really opaque and is bulky enough that it doesn't reveal the car's shape. In other words, until we take the blanket off, we really can't tell what's under there, no matter how hard we try.

Suppose further that the dealer tells you that this has happened to 100 other dealers before, and they've shared with each other what they saw when they finally took off the blanket. For example, perhaps four of the

dealers got a car in state $[001]$ —that is, a two-door red hardtop.

We can summarize the odds of finding each kind of car in one expression:

$$.1[000] + .2[001] + .3[010] + .4[011] + .3[100] + .6[101] + .4[110] + .3[111] \quad (1)$$

Later, we'll want to distinguish between the *weight* associated with a given state and the *probability* of getting that state. It's easy to convert one to the other: the probability is the square of the weight. For example, the weight for state $[101]$ is 0.6, so the probability of getting that state when we remove the blanket is $0.6^2 = 0.36$. Notice that if we add up all the probabilities (not the weights) we get exactly 1.0. This just says that when we take the blanket off the car, we're certain to get something.




Until we look under the blanket, we don't know which of the eight states of the car is under there. If someone asks us, "Is that car in state $[110]$?", the best answer we can give is, "There's a 16 percent probability that when we look, that's what we'll find." The critical point here is that we need all eight probabilities to describe what's under the blanket, until we look. Comparing this to our earlier discussion, we can say that describing any given state requires only three numbers, but that when the car is still unseen, we need eight numbers to describe its total situation—one probability for each possible state.

Now suppose that before we look under the blanket, we push the car forward by a couple of feet. Whatever car is under there—whether it's $[010]$ or $[011]$ or anything else—it's been pushed forward, and its odometer has crept upward. So we can do things to the car without knowing just which one it is yet.

So far I hope you've found this thought experiment straightforward. Now I'm going to get weird and ask that you stick with me. It may sound like science fiction or like a silly philosophical distinction, but this way of thinking will prove important, because it will show us how to exploit the parallelism inherent in quantum computing.

When I first asked you to imagine walking into the dealer's showroom and looking at the blanket covering the car, you may have thought to yourself something like, "Okay, the car is obviously in one of the eight possible states for this model, but we just don't know which one yet. The probabilities found by the other dealers tell us what we're likely to find when we look under there, but that doesn't change what's actually under the blanket. Were I allowed to peek, I'd find that the car is in one of the states, and that's that."

Suppose instead I asked you to imagine that the car is in all eight states simultaneously. I don't mean that the car is partly in one state and partly in another, like the way a Gryphon is part lion and part eagle. I also don't mean that it's a blend of several states, like the way lemonade is partly water and partly lemon juice, but inseparably mixed. And I don't mean that the car is quickly shuttling between states somehow, rapidly switching between, say, a convertible and a sedan. Rather, imagine that the car is in all eight states at once. Somehow, while it's under the blanket, the car is in some



1 Several differently polarized photons striking a polarizer. Only photons aligned with the filter pass through. (a) A vertically oriented filter. (b) An obliquely oriented filter. (c) A horizontally oriented filter.

overlapping reality where it's a sedan and a convertible, red and blue, and two-door and four-door, all at once.

Naturally, you'd want to pull the blanket off and see what this crazy mixed-state car looks like. Perhaps it's some ghostly multiple-exposure 3D image of eight different cars all sharing the same space, with the states with higher probabilities looking a little sharper, or more opaque, than the unlikely states. But in this experiment, every time you pull off the blanket, you observe one car, solid as can be, in one state. Something magical happens when you pull off the blanket, causing the car to assume just one of its possible states. Which state will it be?

In Equation 1, I assigned a weight (and thus a probability) to each of the eight states. For example, 9 percent of the time we'll find that the car is in state [100] and 4

percent of the time we'll get a car in state [001]. We can't pick which state we'll get when we look; we can only predict what the probabilities are and then see what we get. Of course, if one of the probabilities is nearly one and the others are nearly zero, then we're likely to get a car in that high-probability state. If one probability is exactly one and all the others are zero (because they must all add up to one), then we'll get that state for sure.




The point here is that when the blanket is covering the car, the car is in all states simultaneously. When we pull off the blanket and observe or measure the car, it "snaps" into one of the eight states. The probabilities associated with the states determine the state the car goes into. Once we look at the car, the probability for that state becomes one and all the others become zero, so if we look away and then look back, the car will stay in the same state. Even if we cover it up again, it won't mysteriously go back to the simultaneous many-states condition. Once observed, the state remains. So if we see a hardtop, we keep a hardtop.

This is one of the key concepts behind quantum computing. Although this was a far-fetched description of a car salesroom, it's an accurate description of the quantum realm. Whatever it may mean philosophically, for all intents and purposes subatomic particles like photons and electrons, before they are observed, can really be in multiple incompatible physical states at the same time, in ways that larger objects (like cars) never are and seemingly never could be. Thousands of experiments have verified this. This idea (and some others we'll cover in the next section) is critical to creating quantum computers.

Quantum measurement

There's another interesting characteristic of quantum systems that's important to understanding how quantum computers work. I'll describe this effect using an example based on the *polarization* of light. Polarization comes from the wave nature of light. A detailed discussion of polarization here would take us far from our main subject, so let's use a rough but common analogy. Imagine a piece of rope: one end is in your hand; the other end is tied to a wall. If you shake the rope up and down, you'll create a vertical disturbance that will travel down the rope from your hand to the wall. If you shake it from left to right, you create a horizontal disturbance. Similarly, a photon that travels through air can be thought of as vibrating up and down (*vertically polarized*), left and right (*horizontally polarized*), or at any other angle with respect to the floor. The key point is that the photon wiggles as it travels and that these wiggles occur within a flat plane.


Now suppose that you make a flexible material that microscopically looks something like a jail-cell door, composed of many vertical parallel bars with thin gaps between them. For a photon to get through the material, it must wiggle in the plane that is parallel to the bars, as in Figure 1. Such *polarizing filters* (or film) are real—they're common in sunglasses. When sunlight bounces off a highway surface, it often becomes horizontally polarized in the process. Polarizing sunglasses are built so that when we wear them, the polarizing filters are



2 A laser and a clock face for specifying polarization. (a) A vertically polarized photon. (b) A horizontally polarized photon. (c) Many randomly polarized photons.

oriented vertically, and thus block the horizontally polarized glare.

Now let's carry out a little experiment. Our equipment list is short: a bright source of light (a laser will do nicely), a white screen (like a movie-theater screen), a light meter (like the one on a 35-mm camera), and three pieces of polarized film. We'll assume that the photons coming out of the laser have an equal probability of being polarized in every direction—that is, they're randomly polarized. If we look down the barrel of the laser toward the screen and imagine we're looking at a clock, as in



3 The general setup for the polarization experiment. A laser shines onto a screen. Near the laser sits a vertically polarizing filter, a diagonal polarizer, and a horizontal polarizer.

Figure 2, vertically polarized photons will travel in a plane containing 12 o'clock and 6 o'clock. (I'll just refer to directions from now on with just one of the two hours that define the plane.) Horizontally polarized photons lie in the 3 o'clock plane; the other photons will be polarized in all the other planes.

Figure 3 shows the final experimental setup: a laser shines onto a screen. (I'll assume that the screen is instrumented, so we can plug a meter into the back to read out how much light falls on the screen surface.) One by one, I'll insert three polarizing filters: a vertical one right after the laser, a horizontal one just before the screen, and a 45-degree filter in between them.






The experiment has five steps. In step 1, we point the laser at the screen and use the light meter to measure how much light is reflected, as in Figure 4a (next page). For simplicity, let's say that the meter in Figure 4a records A units. That's our reference value for how much light reaches the screen.

Now for step 2. Place a piece of polarizing film in front of the laser, orient the film vertically (that is, at 12 o'clock), and measure the light hitting the screen, as in Figure 4b. We'll find it's $A/2$, meaning that the filter has blocked about half of the light.

To anticipate what's to come, look at Figures 4c through 4e. When we have a single sheet of polarizing film between the laser and the screen, we get a reading of $A/2$. When we have two sheets one after the other, at 90-degree angles to each other, we find that the light meter reads 0. If we then place a third sheet of film between these two, oriented at 45 degrees to them, the light meter jumps up to read $A/8$. This definitely seems weird, so let's look at this experiment more closely.

In step 2, we had a single piece of vertically polarized film between the laser and the screen, and half the photons made it through. This almost makes sense, but remember that the laser generates randomly polarized photons, so there's no preference for photons to travel up or down as opposed to any other direction. If the polarizing material really only lets photons traveling parallel to the bars get through, then only a small number should pass through the filter. In other words, we'd expect a few

4 The polarization-filter experiment. (a) Step 1: The initial setup shines the laser onto the screen. The light meter reads A . (b) Step 2: We insert a vertical polarizer into step 1. The light meter reads $A/2$. (c) Step 3: We insert a horizontal polarizer into step 1. Again, the light meter reads $A/2$. (d) Step 4: We place a horizontal polarizer and then a vertical one into step 1. The light meter reads 0. (e) Step 5: We place a polarizer at time 10:30 between the two polarizers in Step 4. The light meter reads $A/8$.



photons to arrive in exactly the plane containing 12 o'clock but most of them will be at an angle. Because the filter works by blocking photons that aren't parallel to the bars, all these angled photons should have been stopped by the filter. Only a few photons should hit the screen—certainly far less than half. What's going on?

A natural interpretation is to think that the material lets through photons polarized in a direction close to the filter's orientation but not quite aligned with it. But if we measure the polarization of the photons coming out of the filter, we find that's not the case: each and every one of them is vertically polarized.

Because all the exiting photons are vertically polar-

ized, and only a few of the incoming photons arrive in that state, there are only two explanations. Either the material blocks most of the incoming photons and emits vertically polarized photons on its own, or the material changes the polarization of the incoming photons.

The second explanation is the right one. Our intuition might suggest that some of the incoming photons are adjusted somehow as they pass through the polarizer. Perhaps a photon is robbed of its side-to-side motion, leaving just the up-and-down part. In other words, we might imagine that the polarizer somehow turns the photons or that they turn themselves. Perhaps they arrive traveling at some angle to the bars, but to get through, they rotate

themselves to squirm between the gaps and then emerge traveling in that new direction. It turns out that there's no reasonable physical explanation for how photons or filters could do this, so we need to keep looking.

To find the explanation, let's look at it this way. The polarizer, as designed, can only do two things: pass photons aligned with (parallel to) the polarization direction and stop those that are perpendicular to it (either by reflecting or absorbing them). This description doesn't predict what will happen to photons that arrive at an angle, because no in-between answers exist. Such photons can't be partly passed through or partly blocked, because you can't cut a photon into pieces; it must be one or the other. What happens if a photon arrives so that it's polarized in the plane containing 12:30, as in Figure 5? It has to be categorized as either parallel or perpendicular, since passing or blocking the photons are the only choices.


Think of it this way: If the photon must be polarized one way or the other, what are the probabilities that it would end up in each state? Here's one place where our intuition works as we'd expect. The closer a photon is to vertical, the more likely it will be classified as vertical, and the same thing holds horizontally. Let's draw out the nearly vertical photon in Figure 5a in Figure 5b. At 12:30, the photon makes an angle of 15 degrees with the vertical bars of a vertically oriented polarizer. Because the polarization vector has length 1, the vertical component is $\sin 15 \approx 0.97$ and the horizontal component is $\cos 15 \approx 0.26$. We see from this that there's a probability of $(\sin 15)^2 \approx 0.93$ that this photon will be vertically polarized.

But it's not just a matter of interpretation. By virtue of being classified as one or the other, the photon is changed. This is called the *measurement postulate* of quantum mechanics. This isn't how our world tends to work, so let's use an analogy to nail this idea down.

Suppose that I ask you your favorite color and you tell me that it's yellow ochre. Then I lock you in a room with two doors (one red, one blue) and state that the only way out of the room is through the door painted with your favorite color. You protest that neither of these is your favorite color, since you prefer yellow ochre, but I'll hear none of it—you must pick red or blue to move on. You decide red is closer to yellow than blue and go through the red door. By choosing that door, you're transformed. By virtue of that choice, you have a new favorite color. You forget yellow ochre, as though you'd never heard of it. You selected the red door and went through it. Now whenever anyone asks, you say that red is your favorite color. If we bring a million people into that room, each with their own favorite color, after leaving the room they will all now like either red or blue, without exception.

Now suppose that after you've passed through the red door, I lock you in a new room with a violet door and a yellow ochre door and again ask you to go through the door with your favorite color. The fact that you once loved yellow ochre has disappeared as though it never existed; your favorite color has simply been redefined. You now pick the violet door because it's closer to your favorite color of red and from then on violet will be your favorite color.

There are two points to this weird experiment. First,



5 (a) A photon in the 12:30 plane. (b) Drawing the polarization of the photon in Figure 5a with respect to the axes of a vertical polarizer. The y component is parallel to the vertical bars, while the x component is perpendicular to them.

no matter what your favorite color was when you went into one of these rooms, when you passed through a door its color became your favorite color from then on. Second, your choices weren't up to you; they were imposed on you from the outside. In summary, the experiment gave you a limited number of choices for expressing something about yourself and by picking one of those choices it became true for you.

This analogy with the doors captures one of the realities of quantum measurement. When you measure an object, the act of measurement causes the measured object to take on one of the allowed outcomes and then it stays in that configuration. When a photon strikes the vertically oriented polarizing material, it becomes either vertically or horizontally polarized and is thus passed through or blocked. No in-between states exist. A photon that arrives vibrating in the 1 o'clock plane isn't adjusted to rob it of its side-to-side wiggle; it simply changes its nature, and most photons arriving in this state become vertically polarized. We say that the photon is *projected* into one of these two states: horizontally or vertically polarized.

Quantum measurement in practice

With these tools under our belts, now we can say what happens to any photon in Figure 4b when it strikes the polarizing filter. If the photon is precisely vertically polarized, then it's already in one of the two allowed states and passes through. Otherwise, we can break down its motion into a horizontal component and a vertical one, as in Figure 5. The relative sizes of these projections give the relative probabilities that the photon will go into that state. For example, at 1 o'clock the vertical projection is $\sqrt{3}$ as large as the horizontal one, so the photon is three times as likely to become vertically polarized as horizontally. At 45 degrees, the odds are 50/50.

Now we know why the meter reads $A/2$ in Figure 4b. When random photons hit the polarizing filter, half of them are more vertical than horizontal and therefore get projected into the vertical state and pass through.

When we measure a quantum object with respect to a preselected set of choices, the object is projected into one of those choices. The object actually changes. That's the measurement postulate of quantum mechanics.

The rest get projected to the horizontal state and the filter reflects or absorbs them.

Let's move on to step 3 of the experiment. We'll take out the vertical polarizer and put in another sheet of polarizing material that has the bars oriented horizontally, as in Figure 4c. As before, about half of the randomly polarized incoming photons will be projected parallel to the bars (in this case horizontally) and pass through, and the rest will be projected into a vertical state and blocked. When we run the experiment, the light meter again reads $A/2$.

In step 4, we'll put the vertical polarizer back in, just after the laser but before the horizontal polarizer, as in Figure 4d. The photons that make it through the first filter are vertically polarized. Thus, we'd expect that the second, horizontal polarizer would block every one of them. And that's just what happens: the light meter reads 0.

Now we get to the fifth, final, and weirdest stage of the experiment. I'll leave the two sheets of film where they are, so the meter reads 0, and I'll put a third sheet of polarizing film between them, rotated by 45 degrees, as in Figure 4e.

Under everyday conditions, we'd expect the light meter to read zero. After all, no light was getting through the two filters in Figure 4d. Putting another filter in between them shouldn't make a difference, right? Not so. The light meter actually reads $A/8$.

This makes sense if we interpret it according to our new rule that a measurement changes the state of the object being measured. All the photons that strike the 45-degree polarizer arrive in a vertically polarized state because they made it through the first filter. Now this new filter once again projects photons into one of the two states it allows: parallel or perpendicular to its bars. Like the experiment involving two doors and a favorite color, the previous polarization about the photon is lost when it's remeasured. In this case, every incoming photon is polarized vertically, and upon projection onto two axes that are each at a 45-degree angle to it, the photon has an equal probability of being projected into two new polarization directions. So half the arriving photons get projected into one state and get through, and the other half are projected into the other state and are blocked.

To summarize this final version of the experiment, randomly polarized photons leave the laser and hit the first piece of film. Upon striking the film, they're forever trans-

formed, with a 50/50 chance of becoming either vertically or horizontally polarized. Only the latter photons get through. Now these photons (half the original number emitted by the laser) reach the second filter and once again are transformed into a new diagonal state. Half are projected so that they make it through, so one-quarter of the original photons reach the third filter. Here again they're projected into either a horizontal or vertical state with equal probability. Half of them become horizontally polarized and get through this last filter, so one-eighth of the original photons strike the screen.

The essential point of this section, and the reason for presenting this experiment, was to make the following statement: When we measure a quantum object with respect to a preselected set of choices, the object is projected into one of those choices. The object actually changes. That's the measurement postulate of quantum mechanics.

Returning to our automobile example from the previous section, the act of measurement was the removal of the blanket. When we looked at the car, we measured it, and it was projected into one of the eight allowable states.

You might be wondering why polarizing material is special, or why we, as the observers of the car, are special. After all, aren't things being observed all the time? Perhaps the blanket was observing the car. I've never heard a blanket talk to me, but it's not impossible that the blanket was looking at the car. Was the blanket observing it? What if a dog peeked under the blanket? Did the car then get projected into one of its eight states? What if a snail looked under the blanket? What if someone else looked under the blanket, but I closed my eyes. Did the car get projected for that other person but not for me?

This line of questioning gets murky very quickly. This business of things changing themselves when they're observed presents all kinds of interesting philosophical questions like what it means to observe something and what it means to be an observer. Because it's so weird, physicists have sought to understand exactly what's going on here. The jury is still out on a satisfying explanation, but the reality of the situation can't be denied: of the thousands of experiments that have been run to investigate this behavior, every single one has confirmed this model. This property, too, is part of what makes quantum computers work.

Quantum radiosity

To bring this material home to graphics, let's look at an application by sketching out the broadest outlines of a quantum radiosity algorithm. To keep things from getting too messy I'll assume that we're making a black-and-white image (a color image wouldn't require any changes to the algorithm—just more bits in the computation). In this section, when I speak of light, I'll mean light at a single frequency.

Let's summarize the critical points I've previously mentioned:

- Before we examine it, a quantum particle can exist simultaneously in an infinite number of incompatible states (for example, our quantum car that was in eight states at once).

- We can operate on quantum particles in a superimposed state without knowing anything more about them (for example, we could push the car while the blanket was still over it).
- When we finally look at a quantum particle, it's projected into a particular state, depending on the probabilities of its allowed states (for example, when we looked at the car, it became one particular car and stayed that way).
- The measuring device defines the allowed states that a quantum particle can be projected into (for example, the two directions in which our photons got polarized depended on the orientation of the grating).


In my next column, I'll show that we can build quantum versions of classical computer gates (called *quantum gates*), which let us mimic any classical computer with a quantum computer. In fact, quantum gates are probably more powerful than classical gates.

Here's how our quantum radiosity computer will work. First, we set up a huge input register consisting of thousands of short chunks of bits, one after the other, and then one last chunk at the end. These bits are quantum bits, or *qubits*, which we'll discuss next time. For the moment, just think of a qubit as a single bit that can be 0 or 1, or some combination of 0 and 1 with an associated probability that it's one or the other.

Each chunk of bits except the last corresponds to a single surface in our scene. For simplicity I'll ignore volumes, although you can just think "surface or volume" every time you see "surface" and everything will work fine. As shown in Figure 6, each chunk holds five values: a pointer, three numbers, and a matrix. Respectively, these are the object information, the amount of light being emitted from the surface, how much incident light is absorbed, the total amount of incident light, and how much of the light not absorbed (that is, reflected light) falls onto every other surface in the scene (this is the famous matrix of radiosity coefficients). In the very last chunk of the register, we initially place a string of zeros for placing the free energy in the scene.

Now we'll turn some of those precious input bits into qubits. Every bit in the chunks that represent how much light is reflected and absorbed at each surface is turned into a state that simultaneously encodes a 0 and 1 with equal probabilities. That is, if that register is two bits wide, then it simultaneously holds the binary representations for 00, 01, 10, and 11 with equal probabilities, so that register holds all two-bit binary numbers at once. In fact, the entire system represents all possible incident and reflected values simultaneously. Every possible binary value for the light energy arriving at and leaving from every surface is represented with equal probability at the same time.

Next, we compute the amount of free energy bouncing around the scene—that is, the amount of light not accounted for. If we used a classical computer, we'd compute that by using standard radiosity techniques: we'd follow the light around the system and compare how much light is left to be shot or gathered from each surface. If that number wasn't zero (or below some threshold), then it would tell us to run another iteration of the



6 The structure of the monochrome quantum radiosity register. Each surface has five entries: a pointer to the surface data; the emitted, absorbed, and incident light; and the coupling matrix to the rest of the scene. At the end of the register is a space for placing the free energy in the scene.

Further Reading

Quantum computing is such a new field that it grew up with the World Wide Web. Much of the research in the field is available in electronic form on the publicly accessible Los Alamos Physics Preprint Archive at Los Alamos National Labs, located at <http://xxx.lanl.gov/abs/quant-ph>. The LANL server contains mostly Postscript papers but converts many of them into PDF and other output formats on demand. These papers are usually preprints of papers that later appeared in journals, although some are exclusive to the archive.

The LANL server has two excellent introductory articles. A good place to start, and my source for the experiment described in Figure 4, is *An Introduction To Quantum Computing for Non-Physicists* by Eleanor Rieffel and Wolfgang Polak (LANL 9809016). *Quantum Computation* by Dorit Aharonov (LANL 9812037) offers a good discussion of where quantum computing fits into the general context of computation and complexity theory. Both of these also describe Grover's quantum searching algorithm.

An entire book's worth of information is available in John Preskill's course notes, available at <http://www.theory.caltech.edu/~preskill/ph229>.

If you want to get a leg up on quantum mechanics in general, I highly recommend *The Structure and Interpretation of Quantum Mechanics* by R.I.G. Hughes (Harvard University Press, 1989). The book digs right into the math from the beginning. Although it's well presented, it can be rough sledding. If you're not in a rush, you might want to read my column in the next issue for an overview of the terminology and notation before plunging into the general theory.

radiosity computation. But in our quantum computer, we have all the possible values coexisting at once. So the free-energy register contains the free energy for every possible number in every coefficient, simultaneously. One of those superimposed states is our answer. Thus we've computed the solution in one step, as I promised earlier.

Now we search for the state (that is, the configuration of bits) with the smallest free energy. Because radiosity is deterministic, there's only one solution with zero free energy. (Technically, we'd want to make sure first that none of the coefficients are set to zero.) This search can be carried out with a variation on a searching algorithm invented by Grover, which lets us find this least-energy state in a number of steps proportional to the square root of the number of objects that need to be searched. For efficiency's sake, we'd probably want to limit our possible solutions to exclude crazy ones, like

IEEE Annals of the History of Computing


Did you know?

The April–June 2001 issue of *IEEE Annals* takes readers on a tour of digital paint programs from the late 1960s to the early 1980s. Alvy Ray Smith and Richard Shoup reminisce about the early days of painting digitally:

“Digital Paint Systems: An Anecdote and Historical Overview”
by Alvy Ray Smith

“Superpaint: An early Frame Buffer Graphics System”
by Richard Shoup

IEEE Annals covers a wide range of topics in computer history, including computer graphics, security, and networking.



Find more online at

<http://computer.org/annals/>

surfaces reflecting more energy than is hitting them or coffee cups glowing as brightly as supernovas. This algorithm changes the probabilities on the registers, so that we go from a state where every answer is as probable as every other to a state where the solution with zero free energy has a probability of almost one and all the others are almost zero.

Then we look at the output registers, which cause the registers to project from the simultaneous-possibilities state to one particular state. Because we changed the probabilities, the answer we want is almost certainly the one we observe. But if we measure the free energy and find that it isn't zero, then the system got projected to an unlikely state. We would just have to run through it again, knowing that the odds are with us that this time, when we observe the registers, they will be projected to the zero-energy state.

To summarize, we create a situation where we have every possible complete system of radiosity coefficients living simultaneously in the registers and then pick the one that has zero free energy.

If this description is a bit confusing, don't worry. First, I ran through it quickly; we'll see more detail next time. Second, this is weird stuff, and it can take a while to sink in.

I'm not suggesting that we can do quantum radiosity tomorrow. If we have 5,000 surfaces in a scene, then a brute-force encoding using 32-bit numbers would require about 800 million bits. As I mentioned before, the most powerful quantum computers built to date have 3 bits, so they're not powerful enough quite yet. The procedure that I just described is about as simple as a quantum algorithm can be and relies on brute force for storing and computing everything. We can certainly do much better. For example, we know that most radiosity matrices are sparse—that is, mostly zero—so we can store them much more efficiently than my crude layout here. Even so, quantum computing for a big radiosity scene is probably a ways off. But it will come someday.

Next time

The purpose of this column was to present the strange realities of the quantum world without getting bogged down in the notation and terminology of quantum mechanics. Now that we have the ideas in place, next time we'll dig into the tools and mathematics that let us describe and operate on quantum bits and registers. They will let us build up to the design of quantum algorithms, and we'll see how to turn general descriptions (such as the one I gave here for quantum radiosity) into real algorithms. ■

Acknowledgments

My thanks to Jules Bloomenthal, Matt Conway, Steven Drucker, Dave Forsey, Eric Haines, Don Mitchell, and Kirk Olynyk for their help and feedback. I created the figures in this column with 3DS Max 4.

Readers may contact Glassner by email at andrew_glassner@yahoo.com.