

## The Digital Ceraunoscope: Synthetic Thunder and Lightning, Part 1

---

The natural world is rich in complex and beautiful visual phenomena. Many of these are ripe for simulation by computer graphics. For me, developing models of the real world pays off in three different kinds of rewards. First, the process itself is stimulating and clear—it's fun to do, and you can easily see how well you're doing. Second, understanding the natural world gives me a richer appreciation of it, heightening my sensitivity to nuance and detail that I might otherwise miss. And third, I can direct the results—rather than waiting for nature to happen to act as I'd like for artistic reasons, I can direct my simulation as I would direct an actor.

In this issue I'll discuss digital lightning. I first put graphics together with lightning many years ago when one of the Siggraph conference receptions was held in a science center, where an enormous Van de Graaf generator created brilliant lightning chains for our amusement. It seemed irresistible to try to capture that lightning in the bottle of computer graphics.<sup>1</sup>

Of course, lightning is intimately bound up with thunder—one without the other is like Abbot without Costello. Because the same physical processes produce lightning and thunder, it seems reasonable to simulate the two together. In fact, we can go further and generate either one from the other. Given the shape of a lightning stroke and the position of an observer, we can compute just what thunder that observer would hear from that stroke. Similarly, we can start with a description of a thunderclap and an observer's position and create one of the many lightning strokes that would give rise to that sound.

The tremendous physical forces represented by thunder and lightning have fascinated people for millennia, and efforts to capture or simulate them probably date back at least as far. The ancient mysteries used a device called the *ceraunoscope* as an apparatus to imitate thunder and lightning and invoke their power in ritual ceremonies.

In this column, I'll talk about a digital form of the *ceraunoscope*, which not only creates the images and sounds of thunder and lightning, but allows the wielder to control and direct them both.

Before we get started, I should point out that my goal here is to accurately match the natural phenomena, and not simply make something that could pass as an example of it. Everyone knows that the basic structure of light-

ning is a twisty, downward-branching tree. A simulator that generates random lightning with no more physical basis than that can still produce a reasonable looking lightning stroke fairly often, though a designer usually needs to use discretion to select the random strokes that look best. Similarly, thunder can be simulated by a white-noise generator and a few moving bandpass filters. You don't even need a lot of technology. A sound effects pro can take a sheet of metal and, in radio theater tradition, shake it to produce convincing thunder.

In my next column I'll look at thunder. I'll show how to create the thunder that comes from a specific lightning stroke and how to generate the lightning corresponding to a given thunderclap.

### Lightning from the ground up

Here I'll use the term "lightning" to refer exclusively to the process that discharges negative energy from a cloud to the ground. I'll ignore other forms of lightning such as ground-to-cloud and cloud-to-cloud discharges and the more rare phenomena such as ball lightning.

Although a flash of lightning seems to occur in only an instant, the familiar jagged bolt is the visual climax of a process that develops in several stages.

Lightning begins when a large negative electrical charge builds up in a thunderhead. Normal air acts as an insulator, preventing this energy from leaping to the ground. When the negative charge builds up to a sufficient strength, it exceeds the ionization potential for air—essentially the air molecules break down and their resistance drops. The negative charge rushes forward into this ionized volume, like a stream filling a puddle. New charge flows down the channel, building up at its tip. When this charge reaches the air's ionization potential again, the air around the tip of the channel breaks down and the charge plunges forward, taking another step. In this way, the channel of negative charge marches through the air in short bursts.

This advancing pocket of charge is called the *step leader*. It takes about 1  $\mu$ s (1 microsecond) to punch through about 50 meters, then it pauses for about 50  $\mu$ s to build up enough charge to take another step. Although this channel is only a few centimeters in diameter, it's drawn towards the ground by a potential difference that averages about 300 million volts. The step leader's motion is faint and hard to see with the unaided eye.

Andrew  
Glassner

Microsoft  
Research

So far, we've assumed that the lightning is plunging downward into neutral regions of air. When a region of air is positively charged, it pulls the channel forward. On the other hand, when the advancing step leader strikes a pocket of negative charge, it will often split into two channels, each searching in different directions for the shortest electrical path to the ground. Note that this might be—and usually is—different from the shortest geometrical path, which is of course a straight line. The shortest electrical path takes into account the variation of electrical charge in the air through which the lightning channel passes. When the step leader gets close to the ground, often an upward-going positive leader forms at the ground and works its way skyward.

When the downward-stepping negative channel meets the upward-stepping positive channel, they unite to form a single continuous region of ionized air from cloud to ground. The enormous difference in potential now tries to settle to zero. In effect, the two ends of a giant atmospheric battery are short-circuited.

Enormous positive charge (as much as 20,000 amperes) now flows up the completed channel into the cloud. The tremendous flow of current rips apart the air molecules, causing a bright flash of light and a corresponding series of acoustic shock waves—that is, lightning and thunder. So although the step leader flows from the sky to the ground, it's the upward-flowing positive current that causes the actual bright illumination of the lightning channel. In other words, the visual flash of lightning goes up.

An old adage says "lightning never strikes in the same place twice." In fact, lightning often strikes the same place many times, because the ionization channel gets reused many times in quick succession. Although the upward-flowing positive charge neutralizes the region of the cloud where the process began, new negative charge rushes into that region of the cloud. Because the air below the cloud has not fully lost its ionization, it requires less energy to break it down again. After about 30 milliseconds, a new leader forms, and like a wheel in a rutted road, it follows the old ionization path to the ground.

Because the ionized channel has already been cut through the air, this leader can move in one fast plunge to the ground—this is called a *dart leader*. When the dart leader strikes the ground, an answering positive charge rushes back up the channel, creating another flash of light and explosion of sound, and the cycle begins again. Often this process results in four or five flashes along one lightning bolt. You can catch this yourself by moving your camera as you photograph a lightning strike, spreading out the images side-by-side over the film; one such photograph shows 26 flashes.<sup>2</sup> Sometimes the human eye perceives this as a flickering of the bolt. The ionization channel can move from flash to flash, particularly if there's a strong wind pushing it around during and between discharges.

The jaggedness of the lightning channel is caused by the many small and large diversions of the channel in response to localized regions of positive and negative charge in the air. The jagged structure is called *tortuosity*. More information on lightning can be found in Uman's book<sup>3</sup> and in the Uman and Krider paper.<sup>4</sup>

## Lightning geometry

Our first step is to create a 3D lightning bolt. As I mentioned earlier, virtually any algorithm that generates jagged, downward-branching patterns can create lightning bolts, though most still require a designer's eye to pick out the most realistic ones. Since my goal was to make naturally shaped lightning, I wanted a physically based approach, which begins with physical data.

Only a few quantitative investigations of the geometry of lightning strokes exist. Most were based on measuring photographs of a variety of strokes.

Hill gathered some of the first quantitative measurements on segment length and branching by measuring nine digitized strokes.<sup>5</sup> LeVine and Gilson carried out their analysis on eight photographs<sup>6</sup> and derived information on branching and segment lengths. Ribner and Roy tuned the parameters of a 3D random walk to get a rough statistical match to data derived from 10 photographs.<sup>7</sup> For this study, I started with a set of 40 digitized lightning channels collected for the Apollo space program. Let's look first at what the prior work revealed before moving on.


The major result of Hill's analysis was the observation that small-scale tortuosity doesn't seem to depend on height. That is, the amount of small wiggling is the same all along the length of the ionized channel. Hill found that the mean absolute value of the angle between adjacent short segments is about 16 degrees.

LeVine and Gilson modeled the large-scale structure of the lightning shape with a two-scale approach. This let them give each lightning bolt its own "personality," which in real life results from the electrical conditions in the air, the wind profile at different heights, humidity and temperature gradients, and other atmospheric effects. One lightning bolt might be predominantly horizontal with little branching, while another might be mostly vertical with many branches. They generated their lightning patterns using normally distributed independent random processes for the change in  $X$ ,  $Y$ , and  $Z$  of each segment endpoint relative to its start.

Ribner and Roy built the lightning channel step-by-step from the top by choosing two spherical angles around the last point generated. If the angles were both zero, the next point would lie along a vector that was the average direction of the last few segments. This "memory smoothing" prevented the stroke from becoming too erratic. As the angles increased, the points generated belonged to a cone whose main axis was this averaged vector. They reported a best experimental fit using an axis formed by averaging four strokes, a bias term that added in a downward vertical component with strength  $1/20$  of the randomly selected direction, and a mean of 30 degrees for the cone's half-angle.

The two-pass approach automatically handles branching because the segments of the first-pass skeleton are nothing more than straight lines that join one fork in the channel to the next. Thus there's no need to model the geometry and then the branching; when a first-pass segment ends, by definition that's where the channel branches.

For this study I fit skeletons to 40 digitized strokes and extracted parameters for branch length, branching frequency, and branching angle. I found that my measure-



1 (a) The measured data from 40 lightning skeletons. (b) The polynomial approximation. The vertical axis is the probability of choosing a segment length. The other axes indicate the length of the segment (as a percentage of the total stroke length) and the height of the segment above the ground.

ments agreed with Hill’s value of a 16-degree mean branching angle. The skeletal form was a useful representation for studying segment length and branching frequency, since they’re two sides of the same coin. In a skeleton, the segment only ends when the channel stops (either at the ground or in mid-air) or when it forks. I found that, as you might intuitively guess, segments became shorter as they approached the ground. Thus the skeleton automatically also becomes branchier.

I divided the atmosphere between ground and cloud into eight slabs and collected histograms of segment length in each slab. Figure 1 shows a graph of the data. I found that I got a good fit (within 7 percent maximum absolute relative error) to the data using a five-term polynomial for each slab. Table 1 shows the coefficients  $a_i$  through  $e_i$  for the polynomial

$$h_i = \frac{a_i}{y^2} + \frac{b_i}{y} + c_i + d_i y + e_i y^2$$

within a given slab  $i$ . After normalization, this polynomial serves as a probability distribution function (PDF) for the random variable that determines segment length. Basically that means that the height of the function tells how likely it is that a random variable will have that value. When generating a segment of the skeleton, I interpolate the PDFs linearly, using the height of the starting point to weight the nearest two PDFs and combine them.

### Lightning in a CRT bottle

To create a lightning stroke, I use a variation of the two-pass technique, modeling the large-scale structure first, then the small-scale details second. I select the general statistics in the first pass in order to generate a coherent and self-consistent coarse structure, or skeleton. The idea is to use the polynomials in the previous section to determine the length of each piece of the skeleton. The result is a low-resolution approximation of a lightning bolt. Then I pass through the rough skeleton, chopping up the long straight pieces into zig-zags of short strokes of about 3 to 5 meters each.

Table 1. Polynomial coefficients as a function of the height in kilometers.



Height	$a$	$b$	$c$	$d$	$e$
0	-538.8	819.6	-144.0	7.453	-0.113
0.86	-448.6	617.9	-89.85	4.119	-0.058
1.71	-282.7	354.4	-37.24	1.266	-0.014
2.57	-163.8	190.2	-13.00	0.163	0.002
3.43	-120.4	162.8	-11.40	0.171	0.001
4.29	-119.5	148.6	-10.46	0.208	-0.001
5.14	-52.83	76.12	-1.278	-0.193	0.005
6.00	-22.05	23.99	-0.010	-0.060	0.001

This choice of detail size is important because it sets a lower bound on the resolution of the stroke—that is, it sets the size of the finest level of detail. It’s important that this number be reasonable. If the shortest strokes are too short, then we’ll waste valuable modeling and rendering time on useless detail. If they’re too long, then the lightning won’t have the right shape, and the thunder we’ll generate from the stroke will not sound correct. I chose 3 to 5 meters as the finest detail for three reasons. Two of them have to do with thunder—when we use 3 to 5 meter segments, the resulting thunder statistically matches real thunder. I’ll discuss this in more detail in the next column.

The other justification is that this segment size is appropriate for rendering. Let’s imagine a rendering of a typical lightning bolt. At a distance of 3 km, a 3m segment subtends an angle of about 0.06 degree. A 5-km lightning bolt rendered with a pinhole camera at a distance of 3 km requires a half-angle of about 40 degrees. At a resolution of 1,024 pixels on a side, the smallest resolvable angle (at the center) is about 0.08 degree, which corresponds to a segment of about 4.8m. Thus a 3m segment is just comfortably smaller than a pixel. That means we avoid straight-line artifacts due to overly long segments, and we can use standard antialiasing techniques to avoid jaggies.


To summarize, the first pass creates the large-scale

2 The gray arrow represents the average of the last four segments. The cone at the end of the last endpoint contains the direction of the next segment; the gray arrow is its central axis.



3 Creating small-scale tortuosity. The line joining the two endpoints is the skeletal segment. The gray cylinder is a small-scale segment formed by joining points on two adjacent disks.

4 On the left, a drawn 2D skeleton has been promoted to 3D. On the right, the result of adding tortuosity to the designed stroke.



structure. Start at the top point of the lightning stroke. Use the altitude of that point to choose two of the  $h_i$  polynomials and interpolate their coefficients. Choose a random number according to the height of the normalized polynomial, and that's the length of that segment (the height of the curve above each value gives the relative probability of that value). Direct this first segment straight down, then push it a little off the vertical. If you're modeling wind shear, add some horizontal displacement to the endpoint to accommodate the effect of wind. Plunk this first segment down and move to the endpoint.

Use this point as the starting point for the next segment and compute its length the same way. When you have the new length, compute the average of the last few segments (I used four) and orient the new segment within a 30-degree cone centered on that average, as Figure 2 shows. Shear it in the wind direction as before, push it

towards the ground slightly, and plunk down the new segment. Next, make another segment starting from the same point as before and lay it down. Now move to the ends of both pieces and repeat. Stop when any chain reaches the ground. That completes the first pass.

To apply the second-pass detail, create a double cone along the length of each segment, as in Figure 3. Start by building a coordinate frame, using the segment itself as one of the axes. The remaining two axes span circular cross-sections of the cones. Now walk down the axis of the cone in randomly sized steps with a mean of 3m—basically we're marching along the length of the segment taking short and long steps. At each step, the selected point specifies a circular disk through the cone. Choose a random point on this disk to serve as the next point on the tortuous path. I found a good statistical match to my starting data when these points were computed by perturbing the previous point—this kept the segments roughly together, rather than bouncing erratically within the cone boundaries. The cone radius at the center is 25 percent of the skeletal segment's length.

Once the fine-structure lightning bolt has been built, thunder comes into play. We'll look at that more closely next time.

### Designer lightning

One of the pleasures of a digital model of lightning is that you can direct it. Suppose that you have a scene in a movie where a mother and daughter are standing in a field. The mother is pointing to the sky with a laser pointer, shining a red dot onto the cloud. She says "Now!" and a lightning bolt starts at the dot and strikes the top of the weather vane on the building next to them. Getting real lightning to behave this way, on cue, would be quite a trick. Doing it digitally is a snap.

To create a designer lightning bolt, we can start with as little as the desired 2D skeleton as seen by the observer. This can be drawn with any 2D drawing tool and imported into the simulator. From the 2D drawing, choose an angle out of the image's plane for each stroke to create a 3D skeleton. I try to choose angles that match the statistics given above. Then add fine structure to the skeleton using the double-cone technique described in the last section. The result is a lightning stroke whose small-scale statistics match the real world, with the added benefit of starting and ending right on target. Figure 4 shows an example of this process.

The designer then specifies how many flashes should occur and how long the interval between them should be. In my system, the default is four flashes, with a mean interflash interval of 30 ms.

Reed and Wyvill<sup>8</sup> presented techniques for rendering a lightning stroke with an appropriate amount of blooming and line thickness, as though film were actually being saturated. They also showed how to use a lightning stroke as an illumination element. Both of these techniques can be used to render these lightning strokes as well. Recently, Kruszewski<sup>9</sup> published a parametric model for lightning strokes. It may be possible to derive parameters for his algorithm to produce statistically natural strokes—I haven't tried that experiment.

## Results

Figure 5 shows three simulations created with this technique. In Figure 6, another simulation is composed over a real photograph.

Next time we'll look at how to compute the thunder that would be heard by a person on the ground observing these lightning strokes. We'll also see how to turn the process around and generate lightning that would give rise to a particular, orchestrated thunderclap. ■

## Acknowledgments

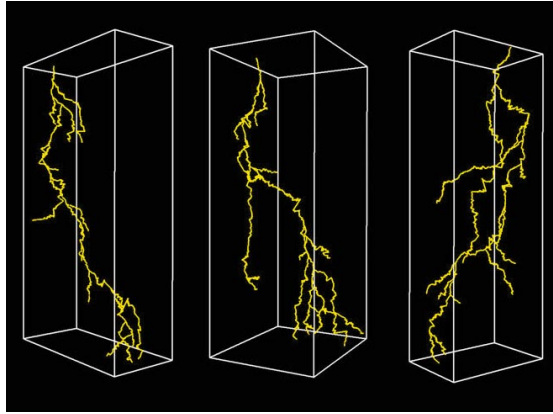
My grateful thanks to Philip Krider of the University of Arizona for his generosity in making available his original tapes of the 40 Apollo lightning strokes, which I used as my primary source data for this project. Thanks also to David Rogers of the US Naval Academy for reading the tapes and converting them to a plain text format.

---

## References

1. A. Glassner, *The Digital Ceraunoscope: Synthetic Thunder and Lightning*, Tech. Report MSR-TR-99-17, Microsoft Research, Redmond, Wash., 1999, <http://www.research.microsoft.com/pubs/>.
2. A.A. Few and T. Teer, "The Accuracy of Acoustic Reconstructions of Lightning Channels," *J. Geophysical Research*, Vol. 79, No. 33, 1974, pp. 5007-5011.
3. M.A. Uman, *Lightning*, Dover Publications, New York, 1969.
4. M.A. Uman and E.P. Krider, "A Review of Natural Lightning: Experimental Data and Modeling," *Trans. Electromagnetic Compatibility*, Vol. EMC-24, No. 2, May 1982, pp. 79-112.
5. R.D. Hill, "Analysis of Irregular Paths of Lightning Channels," *J. Geophysical Research*, Vol. 73, No. 6, March 1968, pp. 1897-1906.
6. D.M. LeVine and B. Gilson, *Tortuosity of Lightning Return Stroke Channels*, Tech. Report TM 86104, Goddard Space Flight Center, Nat'l Aeronautics and Space Administration (NASA), Greenbelt, Md., May 1984.
7. H.S. Ribner and D. Roy, "Acoustics of Thunder: A Quasi-linear Model for Tortuous Lightning," *J. Acoustical Society of America*, Vol. 72, No. 6, Dec. 1982, pp. 1911-1925.
8. T. Reed and B. Wyvill, "Visual Simulation of Lightning," *Proc. Siggraph 94*, ACM Press, New York, pp. 359-364.
9. P. Kruszewski, "A Probabilistic Technique for the Synthetic Imagery of Lightning," *Computers & Graphics*, Vol. 23, No. 2, April 1999, pp. 287-293.

Contact Glassner by e-mail at [glassner@microsoft.com](mailto:glassner@microsoft.com).



5 Three synthetic lightning bolts.



6 A synthetic lightning bolt composed over a photograph.