American Scientist

May-June 1996

Computing Science

Digital Diffraction

Brian Hayes

Note: This document is available in other formats.

Some years ago I visited M. F. Perutz, the Cambridge biochemist who deciphered the structure of the hemoglobin molecule. Professor Perutz showed me a series of artifacts from his 20-year struggle to unravel the twists and folds of the oxygen-carrying protein. There was an x-ray diffraction film whose symmetrical array of dots looked like a lace doily, a contour map of electron density drawn on stacked sheets of transparent plastic, and a huge molecular model supported by a forest of brass rods.

Looking at these objects, I could understand in a general way how an x-ray diffraction pattern reveals the structure of a crystallized molecule. The diffraction pattern is the Fourier transform of the crystal lattice, representing in "frequency space" the positions of the atoms in ordinary space. The lattice and the diffraction pattern have a reciprocal relationship: Widely separated dots on the x-ray film correspond to closely spaced planes of atoms in the crystal, and nearby dots on the film are generated by widely spaced atoms. That much I understood. What was lacking was more-specific knowledge of how to translate from crystal lattice to diffraction pattern and back again. I wanted to look at the x-ray film and see the geometry of the crystal. What I wanted most was a chance to experiment and explore. I wanted to nudge an atom in the crystal, and see how that displacement altered the diffraction pattern.

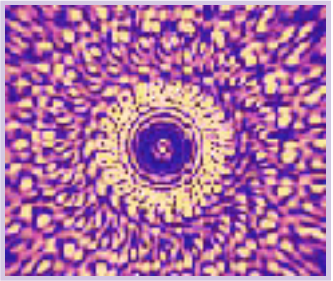

Figure 1. (106 kilobytes)

Only lately have I realized that such an experiment is not out of reach, provided I am willing to settle for a computer simulation. Growing real crystals and recording real x-ray diffraction patterns requires equipment and skills I don't have, but diffraction is a phenomenon the computer can readily imitate. And nudging an atom is a lot easier with virtual crystals than it is with real ones.

I am certainly not the first to think of simulating diffraction with a computer. My inspiration came from Marjorie Senechal's recent book *Quasicrystals and Geometry*, which includes many simulated diffraction patterns as well as a program (written by Stuart Levy) for generating them. Other programs that perform similar functions are listed at the SINCRIS site on the Internet. Rather than experiment with existing software,

however, I have written yet another diffraction program--not because I thought I could improve on the work of my predecessors but simply because writing the program is more edifying than running it.

Fraunhofer and Fresnel

In describing x-ray diffraction, the usual practice is to pretend that both the source of the radiation and the detector are infinitely far from the diffracting crystal, so that all waves move in parallel rays. This scheme, called the Fraunhofer approximation, simplifies analysis and also allows the industrial-strength mathematics of the Fourier transform to be brought to bear on the problem. But if you are trying to follow how individual diffracted waves meet and interfere with one another, infinite distances and parallel rays are not helpful. I have therefore taken a slightly different approach, called the Fresnel approximation, in which the source is still infinitely distant (so that the incoming rays are parallel) but the diffracted rays converge on a detector screen at a finite distance.

While opting for greater realism in this one aspect of the experiment, I have simplified it elsewhere. In particular I have replaced the three-dimensional crystal lattice with a two-dimensional *mask*--an opaque screen with an array of pointlike apertures. Diffraction by such a mask can be understood in terms of an experiment with visible light rather than x rays. The apparatus needed is a card with pinholes, illuminated from behind by a suitable light source (ideally a laser). An impressive collection of diffraction-pattern photographs made in just this way has been published by G. Harburn, C. A. Taylor and T. R. Welberry in their *Atlas of Optical Transforms*.

The key to understanding diffraction by a mask is Huygens's principle: Each pinhole opening acts as if it were an independent source of light, which radiates over the entire hemisphere on the "downstream" side of the mask. The bright spots and dark lanes of the diffraction pattern are created when the waves emitted by multiple pinhole sources interfere with one another. A point on the detector screen is bright if waves from the various apertures arrive there in phase. Otherwise the point is dark, as the out-of-phase waves cancel one another.

The core of an algorithm for simulating this process is based on straightforward geometric optics. First set up a detector screen at a specific distance, *z*, from the diffraction mask. Then for each pixel on the screen, calculate the amplitude and the phase of the waves reaching that pixel from each source in the mask. Both the amplitude and the phase are functions of the distance, *r*, from the source to the pixel. The amplitude falls off as 1/*r*, which implies that the pattern becomes fainter toward the edges. The phase depends on the relationship between *r* and the wavelength, lambda, of the radiation; specifically, the phase angle is equal to 2 pi *r* / lambda (or, if you dislike phase angles greater than 2 pi radians, 2 pi *r* modulo lambda).

To find the total brightness of a pixel, add up the waves reaching the pixel from all sources on the mask. How do you add waves? The state of a wave can be represented mathematically by a vector in the complex plane; the direction of the vector is given by the phase angle and the length of the vector is the wave amplitude. Summing the waves is therefore a matter of vector addition, done by decomposing the vector into its real and imaginary components and adding those components. Writing a routine to add complex numbers is not a major undertaking, but I have avoided even that minor bother by choosing a programming language (Scheme) that has complex arithmetic built in. (Source code for the program is available.)

In a physical diffraction experiment, neither the amplitude nor the phase of the waves is directly observable. What is measured at the detector screen is the *intensity* of the radiation, which is equal to the square of the amplitude. Accordingly, it is the intensity at each pixel that the simulation program calculates and displays.

Stripes and Spots

One of the frustrations of trying to understand diffraction through x-ray crystallography is that the simplest patterns are among the hardest to generate. There is no way to observe the diffraction of x rays by a single atom or by an isolated pair of atoms. But the analogous experiments are easy in a computer simulation.

The simplest experiment of all is diffraction by a blank mask--one with no openings. In this case, of course, the detector screen is uniformly black. With a single aperture, light passes through the mask but there is no interference of waves. Intensity is highest directly opposite the aperture, and it falls off elsewhere in proportion to $1 / r^2$.

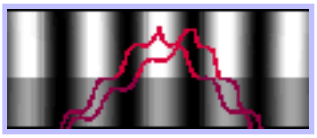

Figure 2. (13 kilobytes)

Things begin to get interesting with two apertures. Suppose pinholes are punched at two points aligned horizontally along the x axis of the mask, at $x = -1$ and $x = +1$. What diffraction pattern will be observed? The geometry of the problem is simplest when the detector screen is at a great distance from the mask--in other words, when the experimental arrangement is an approximation to the Fraunhofer approximation. At the origin ($x = 0$, $y = 0$) on the detector screen, the distance r from both sources is the same; thus the waves from the two pinholes are necessarily in phase at the origin, and the intensity must be at a maximum there. Indeed, all along the y axis the same equal-distance relation holds, and so there is a bright vertical stripe along the middle of the screen. Now move slightly away from the y axis. The path from one source becomes shorter while the other path becomes longer, and so the two waves are no longer in phase. Where the difference in path lengths is equal to one-half lambda, the two waves cancel exactly, and so there is a dark stripe parallel to the central bright one. Moving still farther, the difference in path length eventually becomes equal to lambda, and at this point the waves reinforce each other again. The alternation of bright and dark regions continues across the entire screen, producing a pattern of regularly spaced vertical stripes.

If the two apertures are arrayed vertically, the physics of the situation is the same, but the stripes in the diffraction pattern are horizontal. More generally, any two isolated apertures yield a pattern of stripes perpendicular to the line connecting the apertures.

What happens if the two apertures are moved father apart, say to $x = -2$ and $x = +2$? Now a given horizontal displacement along the detector screen induces twice the difference in path length, which means the vertical stripes are only half as wide; their spatial frequency is doubled. Moving the sources closer together, to $x = -1/2$ and $x = +1/2$, has the opposite effect: Path-length differences are halved, and so is the spatial frequency of the stripes. The closer the pinholes, the wider the stripes--a simple model of the reciprocal relation between the lattice spacing of a crystal and the dot spacing in an x-ray diffraction pattern.

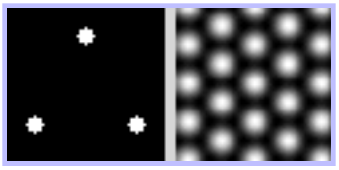

Figure 3. (11 kilobytes)

The striped pattern produced by a pair of pinholes is the key to understanding all diffraction patterns. Suppose you have a mask with three apertures, arranged at the vertices of an equilateral triangle. How would light diffract through these openings? If you were to delete any one vertex, the result would be easy to predict: a pattern of stripes perpendicular to the remaining side of the triangle, and with spatial frequency proportional to the length of that side. The three possible deletions yield patterns of stripes that differ in orientation by 120 degrees. The pattern created by the complete triangle is simply the sum of these three pairwise patterns. Wherever three bright stripes intersect, the pattern has a bright spot, and elsewhere it is dark. The result is an infinite array of dots with triangular symmetry.

More complicated masks can be analyzed in the same way, at least in principle. For any mask pierced by pointlike apertures, the diffraction pattern is the sum of the patterns associated with all the pairs of apertures. This way of understanding diffraction is a two-dimensional analogue of Fourier analysis, which breaks a waveform down into sinusoidal components of various frequencies. In this case the stripes are the sinusoids; the one distinctive complication is that they vary in orientation as well as frequency.

Seeing a diffraction pattern as a superposition of variously oriented parallel stripes works only in the Fraunhofer approximation. When the detector screen is brought closer to the mask, geometric distortions become apparent. In particular, the stripes generated by a pair of apertures are no longer straight lines. Each stripe is the locus of all points whose distances from the two sources differ by a fixed amount (namely lambda). Except along the central bisector, the form of this locus is not a straight line but a hyperbola. Thus a diffraction pattern is really a superposition of hyperbolic curves. In the simulation program, the curvature of the stripes can be made either subtle or obvious by adjusting the mask-to-screen distance *z*.

Crystals and Quasicrystals

Even with some basic principles in hand, predicting the features of a diffraction pattern by looking at the geometry of a mask is a tricky business. It's not easy to think like an x ray.

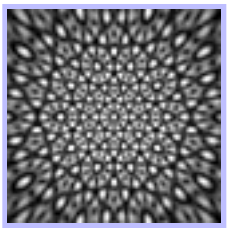
Figure 4. (19 kilobytes)

We have already seen that a triangular mask yields a diffraction pattern in which triangles tile the plane. Likewise, four apertures arranged at the corners of a square yield a square tiling of the plane--a pattern of bright spots at the vertices of a square grid. Does a pentagonal mask tile the plane with pentagons? Without knowing anything at all about diffraction, you could answer no to this question, because a tiling of the plane by regular pentagons is geometrically impossible. The actual diffraction pattern from a pentagonal mask has an elaborate fivefold symmetry but is not a tiling.

Some of the most visually appealing diffraction patterns in the *Atlas of Optical Transforms* come from polygonal masks and from spiral arrangements of apertures. These masks are two-dimensional analogues of fibrous molecular structures, such as microtubules and the helical protein collagen. The simulation program yields handsome patterns for such masks, particularly when the contrast is adjusted to bring out detail in the shadows, or when light intensity is encoded in color (*see Figure 1*).

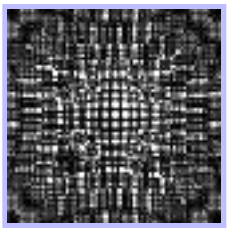
Figure 5. (18 kilobytes)

The masks that correspond most closely to a crystal structure are those with a periodically repeating array of apertures that in principle could cover the whole plane. The simplest example is a square lattice, with a pinhole at every vertex. In my first experiment with such lattices, I created a mask with a 10-by-10 array of openings, and chose some plausible-seeming values for lambda and *z*. By this stage in my work I thought I had a sense of how diffraction works, and so I thought I knew what the output would look like. Figure 5 was not what I expected. What went wrong? Crystallographers call it the finite-size effect. Because the mask does not go on forever, there is diffraction associated not only with individual apertures but also with the boundaries of the mask; the final pattern is the convolution of these two effects. In Figure 5 the small-scale features are artifacts of the boundary.

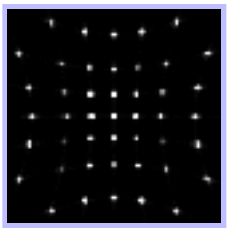
Figure 6. (8 kilobytes)

Finite-size effects cannot be eliminated altogether, but they can be suppressed by a more judicious choice of lambda and *z,* and by building a bigger array of apertures. The true pattern generated by a lattice is a *reciprocal lattice*: an array of dots with the same basic geometry as the mask, but with spacing inversely proportional to that of the original.

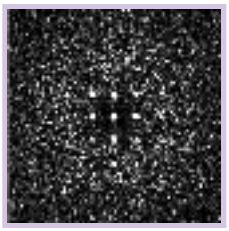
Figure 7. (16 kilobytes)

Irregularities or imperfections in a crystal can be modeled by randomly displacing the pinholes in a lattice mask. (This is where the nudging of atoms comes in.) In the resulting diffraction pattern the reciprocal lattice is still visible, but it is enveloped in a gray fog or dust cloud that indicates the essentially continuous spectrum of distances between apertures in the perturbed mask.


Figure 8. (14 kilobytes)

The limiting case of this process is a mask made up of apertures sprinkled totally at random, which produces a diffraction pattern consisting entirely of dust.

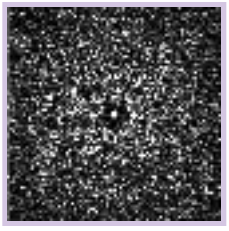
Figure 9. (17 kilobytes)

The random mask might be taken as a model of a glass--an amorphous solid--but it misses an important constraint on the distribution of atoms in a glass. In the random mask, apertures can fall anywhere, but in a glass no two atoms can approach any closer than the sum of their radii. This excluded-volume effect has an important bearing on the diffraction pattern of a glass: Although there are none of the discrete spots that characterize diffraction from crystalline solids, there are rings suggesting a tendency for atoms to cluster at the minimum interatomic distance and at multiples of that distance. Creating a mask with the statistical properties of the glassy state is not easy, but after some effort I have managed to generate diffraction patterns with at least a hint of annular structure.

quasicrystals and their two-dimensional analogues, the aperiodic tilings. These remarkable structures, which 25 years ago were not only undiscovered but unthinkable, have produced such an upheaval that crystallographers have been compelled to redefine the very subject matter of their discipline. Officially, a crystal is no longer a solid with a periodically repeating structure; instead it is a solid that has an "essentially discrete" diffraction pattern. The definition is calculated to include quasicrystals, which are nonperiodic but nevertheless produce a diffraction pattern of individual spots.

Figure 10. (33 kilobytes)

In two dimensions the best-known aperiodic tilings are the ones devised by Roger Penrose in the 1970s, especially the "kite-and-dart" construction of Figure 10. The tiling can be extended indefinitely--there is always room for another kite or dart--and yet it never becomes periodic. In other words, there is no unit cell that can be repeated in rubber-stamp fashion.

The kite-and-dart tiling is converted into an aperture mask by placing a pinhole at each vertex of each polygonal tile. Generating a diffraction pattern from this mask confirms the peculiar nature of the Penrose tiling: The diffraction pattern consists of discrete bright spots--the hallmark of a crystal--and yet it has fivefold symmetry, which is impossible for any periodic tiling.

Murky Patterns

Creating a program to simulate diffraction has answered some of the elementary questions that have nagged me ever since I first saw Perutz's x-ray film of the hemoglobin molecule. And yet, although I understand more about diffraction than I once did, I am puzzled by key aspects of my own simulation. In particular, I don't entirely understand how the pattern changes as a function of *z* and lambda. For a simple mask, such as a pair of apertures or a grid, the rules seem clear, and when I adjust *z* and lambda, the result is what I expect. With more complicated masks, however, the situation gets murky. The diffraction pattern for the Penrose tiling of Figure 10 is not nearly as sharp as it should be, but a long search failed to turn up values of *z* and lambda that yield a better result. (Of course the problem may be elsewhere. There could be an error in the conception or the implementation of the program. Or I might be seeing "numerical noise" as a result of exceeding the precision of numbers used in the calculation.)

Apart from matters of correctness there is the question of computational efficiency. In the present algorithm the number of operations needed varies as the product of the number of pixels in the detector screen and the number of apertures in the mask. By exploiting the symmetry of the diffraction pattern one might avoid calculating half the pixels. Is there an even better shortcut?

Bibliography

Harburn, G., C. A. Taylor and T. R. Welberry. 1975. *Atlas of Optical Transforms*. Ithaca, N.Y.: Cornell University Press.

Lipson, S. G., H. Lipson and D. S. Tannhauser. 1981. *Optical Physics*. Third Edition. Cambridge, England: Cambridge University Press.

Perutz, M. F. 1964. "The Hemoglobin Molecule." *Scientific American*, 211:5 (November): 64-76.

Senechal, Marjorie. 1995. *Quasicrystals and Geometry*. Cambridge, England: Cambridge University Press.

SINCRIS: Information Server for Crystallography. http://www.lmcp.jussieu.fr/sincris/.

Wagon, Stan. 1991. *Mathematica in Action.* New York: W. H. Freeman.

American Scientist Home Page