First, you find the simplest polynomial (with integer coefficients) that has your radical expression as a root. (A root of a polynomial is a number that makes the polynomial equal to 0.) For example, the simplest polynomial for $\sqrt[3]{2}$ is $x^3 - 2$. Second, you extend the set of rational numbers by tossing in all the roots of this polynomial, not just the original root, and then taking all possible sums and products, so that the extension is closed under addition and multiplication; this extension is called the splitting field because the polynomial "splits" into linear factors. For example, the splitting field of $\sqrt[3]{2}$ is formed not just by tossing $\sqrt[3]{2}$ in with the rational numbers, but by including $\sqrt{-3}$ as well, since $\sqrt[3]{2}(1 \pm \sqrt{-3})/2$ are the other two roots of the polynomial $x^3 - 2$.

Landau's algorithm requires the construction of a second algebraic structure, called the Galois group, associated to the splitting field. This is still familiar territory to algebraists—Galois theory was created for the purpose of understanding how and when the roots of a polynomial can be written using radicals. Unlike the field, which has infinitely many members, the Galois group is finite, and this makes the computer very happy.

The algorithm searches the Galois group for a sequence of nested subgroups satisfying certain technical conditions. (The subgroups sit inside each other like tightly fitting Russian dolls.) Once the shortest such sequence is found, the algorithm translates the subgroup nesting into a nested radical expression for the original number. The theory behind the algorithm guarantees that this translation produces the least nested version of the number.

Although it solves the problem, Landau's algorithm is not necessarily the last word in denesting radicals. One drawback is that the algorithm requires a potentially huge amount of computation—the splitting field and its associated Galois group can be extremely large. Inserting an extra radical sign in the initial expression can more than double the algorithm's work load.

It may be that inefficiency is the price to be paid for an all-purpose denesting algorithm. However, at the same time, Landau notes that "what is theoretically slow may be practically fast, and vice versa." In any event, it's nice to know that something can be done with those awkward algebraic expressions. Now if only the linguists could come up with some way of untangling bureaucratic officialese. ■ BARRY A. CIPRA

Barry Cipra is a mathematician and writer based in Northfield, Minnesota.

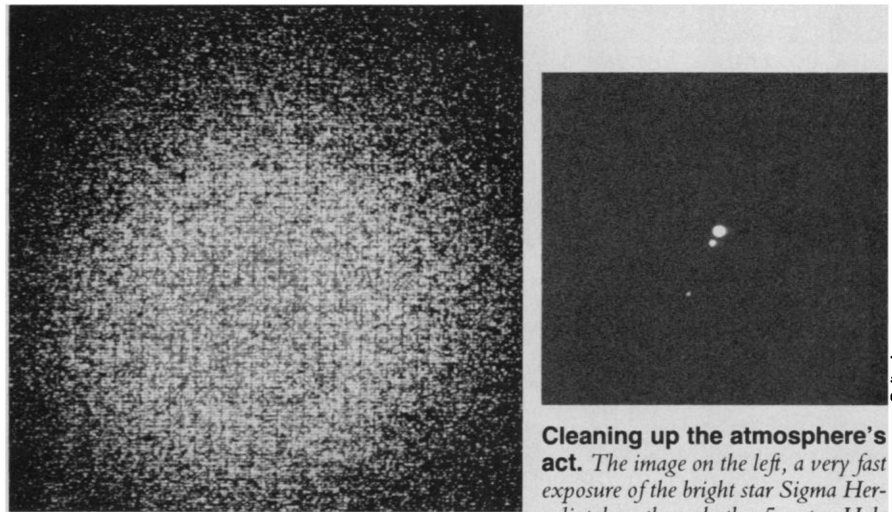

Cleaning up the atmosphere's act. *The image on the left, a very fast exposure of the bright star Sigma Herculi taken through the 5-meter Hale telescope, shows how the atmosphere breaks up a single shaft of starlight into a myriad of speckles. The image on the right shows the Caltech computer's reconstruction of the Sigma Herculi.*

Computer-Age Stargazing

When astronomers use the 5-meter Hale telescope at the Palomar Observatory to look at a star—in this case, the binary system Sigma Herculis—what they actually see is the image on the left: a shimmering, boiling blur caused by the incessant motion of the atmosphere. What they would like to see is the image on the right: a pair of crisp, well-defined spots blurred only by the unavoidable diffraction of light being focused by the huge mirror.

Now they can. Palomar director Gerry Neugebauer and seven other California Institute of Technology astronomers have developed and implemented two techniques that virtually eliminate the distorting effects of the atmosphere, thus allowing this telescope or any other telescope to approach its theoretical maximum resolution. The improvement in this case is a factor of 20, from roughly 1 arc second to about 50 milliarc seconds.

Both techniques rely on the fact that the granulated mess on the left contains precisely as much information about the source as the original starlight did—just scrambled. The trick is to use massive computer processing to unscramble it.

In Non-Redundant Masking, the method used to make the Sigma Herculis image, the Palomar team places an opaque screen pierced with five to seven tiny holes at the prime focus of the telescope. They then take the separate beams of starlight coming through the screen and recombine them into an undistorted image using mathematical algorithms developed for radio observations at the Very Large Array near Socorro, New Mexico.

In the Fully Filled Aperture technique, the group simply turns the computer loose on the whole blur, with nothing interrupting the light. The computer then recovers the image using a refinement of a long-established method known as speckle interferometry.

Neither method is easy. In both, the computational demands are horrendous, requiring the full parallel processing power of Caltech's 512-node "hypercube" supercomputer. Moreover, the techniques are cumbersome for the observers, demanding lots of 5- to 30-millisecond exposures to keep the motion of the atmosphere from smearing the image irretrievably. And they are limited to reconstructing relatively bright objects—in the case of nonredundant masking, no fainter than what can be seen with the naked eye.

Nonetheless, these techniques do show just how far ground-based astronomers have come: leaving aside such issues as faint-object sensitivity and wavelength coverage, where orbital observatories still have a decided advantage, the angular resolution demonstrated here is about as good as that expected from the Hubble Space Telescope. ■ M. MITCHELL WALDROP