Third, in a more general sense, the status of early modern human cannot be recognized by cultural products, because the Qafzeh people, and other equivalent populations in Africa, are associated with tool technologies that have been labelled "more primitive" than Upper Paleolithic. "We are going to have to look for more subtle differences in the tool technologies, such as regional variation," says Alison Brooks of the Smithsonian Institution in Washington.

With a suggested age of 92,000 years, the Qafzeh fossils are about as old as the oldest early modern humans in Africa, those from Klasies River Mouth cave in South Africa. Does this mean that modern *Homo sapiens* might have evolved in southwest Asia, and not in sub-Saharan Africa as many researchers now favor? Without the recent genetic evidence, this certainly would be a tenable position. However, data on both nuclear and mitochondrial DNA point strongly to a sub-Saharan origin, as Stringer and Andrews outline in their review.

Given the genetic data and the ages of the Qafzeh and Klasies River Mouth fossils, the origin of modern *Homo sapiens* must be put at substantially earlier than 100,000 years. The split between sub-Saharan populations and the rest of the Old World presumably has occurred by Qafzeh times, with the cave standing right in the corridor to the rest of the Old World.

Why early modern human populations took 50 millennia to penetrate further into the Old World is something of a puzzle. Perhaps the environment was unfavorable? Perhaps the established Neanderthal populations impeded migration? Perhaps most likely of all is that early modern human fossils dating from before 40,000 years ago exist but remain to be discovered in Eurasia.

Stringer and Andrews finish their review with the warning that "paleoanthropologists who ignore the increasing wealth of genetic data on human population relationships will do so at their peril." They are alluding to the prolonged confrontation during the 1960s and 1970s between molecular biologists and paleoanthropologists over the likely date and identity of the first member of the human family. In that case the molecular evidence was much closer to the mark than was the fossil evidence. This time around the geneticists' contribution is being welcomed by a few, considered cautiously by many, and flatly rejected by almost no one. A distinct improvement. ■ ROGER LEWIN

ADDITIONAL READING

H. Valladas *et al.*, "Thermoluminescence dating of Mousterian 'Proto-Cro-Magnon' remains from Israel and the origin of modern man," *Nature (London)* **331**, 614 (1988).

Y. Rak and B. Arensburg, "Kebara 2 Neanderthal pelvis," *Am. J. Phys. Anthrop.* **73**, 227 (1987).

Zeroing in on the Zeta Zeta Function

More evidence for the Riemann Hypothesis is provided by new method that speeds up the necessary calculations, also has applications to other large-scale computations

DAVID Hilbert—who in many ways set the course for 20th-century mathematics—was once asked what he would do if, as in the legend of Barbarossa, he found himself revived 500 years in the future. "I would ask," the great mathematician replied, "Has somebody proved the Riemann Hypothesis?"

The answer to Hilbert's question is still No, but mathematicians have amassed impressive amounts of numerical evidence in favor of the Riemann Hypothesis, which is a prediction about the mathematical behavior of a special function known as the Riemann zeta function. Recent work by Andrew Odlyzko at Bell Labs and Arnold Schönhage at the University of Tübingen has carried his numerical evidence to unprecedented levels, by means of algorithms that speedup the necessary calculations. Their techniques, moreover, have applications beyond the zeta function to other large-scale computations.

The Riemann zeta function has fascinated mathematicians for the past hundred years, ever since Bernhard Riemann showed that the properties of this complicated function have deep implications for the distribution of prime numbers—positive integers that are divisible only by themselves and 1—that is, 2, 3, 5, 7, 11, and so forth. In particular, Riemann established a program by which properties of the zeta function could be used to prove that the number of primes less than a given number x is approximately equal to x divided by the natural logarithm of x. The proof was completed some 30 years later, in the 1890s, independently by two French mathematicians Charles-Jean de la Vallée Poussin and Jacques Hadamard. The result is known as the Prime Number Theorem.

The Riemann Hypothesis is a prediction about the zeros of the zeta function—points where the function takes the value 0. Riemann showed that the zeta function—which was first studied by Leonhard Euler in the 18th century—could be defined for complex as well as real numbers. Complex numbers, which are obtained by adjoining the square root of minus one (a seemingly senseless, "imaginary" number) to the real line can be thought of as points in a plane. Riemann showed that the zeros of the zeta function (except for some "trivial" zeros that occur at the negative even integers) all lie within a thin strip of the complex plane—and this turns out to be enough to prove the Prime Number Theorem. The Riemann Hypothesis, however, goes much further: it asserts that these zeros lie exactly on the centerline of this strip.

One application is to check for bugs in computer operating systems and compilers.

If true, the Riemann Hypothesis has enormous implications for number theory. It would drastically improve the known estimates for the distribution of primes. (The Prime Number Theorem gives an *approximate* formula for the number of primes less than x; the *error* of this approximation is intimately related to the zeros of the zeta function.) Many other "theorems" in number theory are "conditioned" on the Riemann Hypothesis—that is, their proofs assume that the Riemann Hypothesis is true—and some other results are known to be equivalent to the Riemann Hypothesis. So there is no wonder why mathematicians keep looking at the zeta function.

The number and approximate location of zeros on any segment of the centerline can be determined by keeping track of sign changes of the zeta function—or more exactly, an adjusted version of the zeta function. (The zeta function can be adjusted so that it takes real values along the centerline; if it changes in sign from positive at one point to negative at another, then it must take the value 0 somewhere in between.) Another, more sophisticated technique makes it possible to count all the zeros within a rectangle (but does not specifically locate any of the zeros). If the rectangle contains a segment of the centerline and the two counts agree, then the Riemann Hypothesis has been verified to the extent of that rectangle. Said differently, if the rectangle count is ever *more* than the centerline count (it cannot possibly be less), then the Riemann Hypothesis is definitely *false*. This is one mild reason for all the computer searches: you could get famous by finding a counterexample.

Although counting zeros on the centerline is theoretically more straightforward than the rectangle count, it is computationally more difficult, in essence because it requires that each zero be located with some precision. (The rectangle count can be thought of as an *averaging* procedure which gives less information but is consequently less "expensive.") The search for zeros is a story in itself. J.-P. Gram published values of the first 15 zeros in 1903. This was actually the first solid evidence supporting the Riemann Hypothesis. There was even some wonderment at how Riemann had arrived at his prediction, since his paper on the zeta function contained no computations at all. It was generally believed that Riemann had based his hypothesis on aesthetics and intuition—two of the driving forces in mathematical research.

Aesthetics and intuition may have played a role, but in 1932, Carl Ludwig Siegel proved that there was more. Searching through Riemann's unpublished papers in the archives of the University Library at Göttingen, Siegel discovered that Riemann had indeed computed several zeros of the zeta function. Not only that, but he had done so by a method superior to those that Gram and others had used after him!

Siegel cleaned up Riemann's method, and the so-called Riemann-Siegel formula became the basis for computing zeros of the zeta function. The list of zeros has grown with the advent of high-speed computers. In 1952, Alan Turing identified the first 1054 zeros ("the first" meaning the zeros closest to the real axis—the purported centerline is perpendicular to the real axis at the point 1/2). The list grew to 25,000 in 1956, 3.5 million in 1968, and 81 million in 1979. In 1986, the list reached a staggering 1.5 billion zeros—every one of which lies on the predicted centerline. This last calculation occupied 2 months on a Cyber 205 supercomputer.

Odlyzko and Schönhage's new method is based on an analysis of the Riemann-Siegel formula, and is designed to locate quickly a large number of zeros in a segment of the centerline high above the real axis by computing the zeta function at equally spaced points along the segment. Odlyzko and Schönhage focused on the "costly part" of

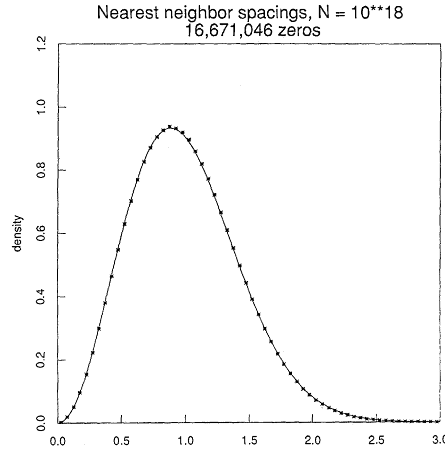

Nearest neighbors. *Probability density for normalized spacings between consecutive zeros of the zeta function. Solid line: theoretical prediction based on the Gaussian unitary ensemble. Scatter plot: empirical data based on 16 million zeros near the 10^{18}th zero of the zeta function.*

the Riemann-Siegel formula: an infinite sum of cosines. They determined that this sum could be approximated sufficiently well by a more manageable exponential sum. Then, using a technique called the Fast Fourier Transform, they reduced the problem to one of evaluating a rational function at equally spaced points—for which they also developed a fast algorithm. Their results are described in a paper which is to appear in the *Transactions of the American Mathematical Society*.

Their method is actually slower than the Riemann-Siegel formula at computing the zeta function at a single point, but by its use of the Fast Fourier Transform to compute the zeta function at equally spaced points, it is much faster at evaluating the zeta function at large numbers of nearby points. Odlyzko observes that a pair of computer runs show five orders of magnitude (10^5) difference in speed between the new and old methods. Using the old approach, a Cray X-MP spent 15 hours locating 10^5 zeros starting with the $(10^{12}+1)^{th}$ zero; the same computer spent 18 hours using the new method to locate 16 million zeros starting with the $(10^{18}+1)^{th}$. (The "cost per zero" of the Riemann-Siegel formula increases with the square root of the starting point, accounting for three orders of magnitude in the comparison; the other two orders of magnitude come from the hundredfold increase in the number of zeros located.) According to Odlyzko, more speedup would be possible if more memory were available.

The key ideas in Odlyzko and Schönhage's method are not dependent on the Riemann-Siegel formula. They can be applied to other functions of interest in number theory, and apparently much further afield. Odlyzko notes that their algorithm is essentially "interchangeable" with another recent algorithm developed at Yale for computing in discrete time steps for n-body simulations, such as in computer simulations studying the evolution of galaxies.

Why would anyone want so much data on zeros of the zeta function? Odlyzko says that some computations have been done "as jokes," but adds that there are serious mathematical reasons as well. One application is to check for bugs in computer operating systems and compilers. A small mistake is likely to exhibit itself as a counterexample to the Riemann Hypothesis. The program can also be run on overlapping segments; if the results agree on the overlap, that is evidence that both the program and the computer are operating correctly, because the computations are completely different for different segments. Odlyzko describes his program as "extremely good at detecting any kind of error" because the "final answers depend on absolutely everything that came before."

Another reason for the computations has to do with another hypothesis, stronger than the Riemann Hypothesis, concerning the spacing between consecutive zeros. This hypothesis asserts that the zeros are not only on the centerline, but they are also "randomly distributed" along the line in a very special way, which turns out to be familiar from mathematical models in high-energy physics: the zeros of the zeta function appear to be spaced much like the energy levels of many-particle systems. This observation has lead some physicists to propose the zeta function as a possible model for "quantum chaos."

By studying the distribution of zeros far away from the real axis, where one might expect them to have settled into a statistical pattern, Odlyzko has shown that the actual behavior of the zeros is in close agreement with their predicted behavior. The fit is fairly good near the 10^{12}th zero, but almost perfect when the data is taken near the 10^{18}th zero. More recently Odlyzko has moved out another two orders of magnitude, and is collecting data on the 31 million zeros near the 10^{20}th zero. For the time being he will not go any higher, he says, because the size of the numbers is already straining the double-precision accuracy of the supercomputer and higher precision computations would be too slow. Instead he plans to start looking for "special points" lower down, where the prime numbers, which got the whole subject started, conspire to make the zeta function "pretty wild." ■ **Barry A. Cipra**

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