Sowers and Bender monitor the same isotopic variations of seawater in the Vostok ice. These variations were recorded in the air bubbles trapped when snow was squeezed into ice. The air's oxygen had in turn been released from seawater by photosynthesis.

Sowers and his colleagues found that the ice did not begin to melt until roughly 3000 to 5000 years after carbon dioxide and methane began to increase and temperatures began to rise at the end of the penultimate ice age. Subtleties in the link between seawater and atmospheric oxygen create some uncertainties in the timing of the greenhouse and melting, but Bender believes that there is an 80 or 90% chance they have the right order of events. If so, it that the greenhouse warming is a cause, and not an effect, of the end of an ice age.

This finding from the Vostok core supports earlier work by Nicholas Shackleton of the University of Cambridge and Nicklas Pisias of Oregon State University. Analyzing sediments, they found that an indirect measure of atmospheric carbon dioxide increased before ice melting began, and that it followed a change in the shape of Earth's orbit. Many researchers consider such orbital variations, the so-called Milankovitch mechanism, to be the ultimate trigger for ice age initiation and termination.

Much work remains to be done on ice from Vostok drilling, which is now being analyzed at three U.S. institutions. For example, something else must have helped end the penultimate ice age, because an enhanced greenhouse can account for only about half of the observed warming.

■ RICHARD A. KERR

Big Number Breakdown

Armed with a powerful new method of factoring and assisted by hundreds of mathematicians and computer scientists around the world, two researchers have found the factors of a 155-digit number. This number, which is by far and away the largest ever factored, had spent the last several years at the top of a list of "most wanted" numbers that have not yet been factored but are known not to be primes.

Aside from its intellectual satisfactions, this feat has some significant implications for much more worldly matters—such as protecting bank accounts. Some of the most sophisticated current schemes for ensuring the accuracy of electronic fund transfers, for example, rely on the difficulty of factoring very large numbers. A truly efficient factoring method, such as an improved version of the one used to factor the 155-digit number, could cause commercial headaches.

That feat was accomplished on 15 June, after about 2 months of effort, by Mark Manasse of Digital Equipment Corporation and Arjen Lenstra of Bellcore. The big number, which can be written as $2^{512} + 1$ is known as F_9. When the dust settled, F_9 appeared on Manasse's computer screen as the product of a previously known seven-digit prime and two new primes having 49 and 99 digits, respectively.

F_9 derives its monicker from the 17th-century mathematician Pierre Fermat. Sometime around 1640 Fermat conjectured that numbers of the form $2n + 1$ were prime whenever n was a power of 2. And indeed, the first four such numbers—5, 17, 257, and 65,537—are primes. But after that Fermat was mistaken: the next four (make that five, now) have been factored and many more are known to be composite (nonprime).

The new algorithm that Mannasse and Lenstra used for factoring F_9 is called the number field sieve. The sieve is a child of the fertile brain of British mathematician John Pollard, who introduced two other factoring algorithms for large numbers during the 1970s. It is closely related to a method called the quadratic sieve, which Manasse and Lenstra have used in previous factoring work.

Both sieves break the task of factoring a large number into a huge set of smaller factoring problems. These can then be farmed out—hence the large number of collaborators in the F_9 work. When enough results are in, the master computer stitches them together to form a candidate factorization. If the first try fails, the computer simply tries another combination of factors.

In case of F_9, the stitching was itself a major computational operation. Manasse and Lenstra estimated that one key step would have taken 6 weeks of computing time on a VAX 11/780. They were ultimately able to do it in 3 hours on a Connection Machine supercomputer at the Supercomputer-Computational Research Institute at Florida State University.

Successful as it was in this case, the number field sieve is not yet practical for use on all the large numbers that are currently of interest. The key new feature of the number field sieve is that the smaller factoring problems are done not with ordinary integers but with a number system that includes the root of a carefully chosen algebraic expression called a polynomial. The advantage of working with such "algebraic integers" is that it makes the smaller factoring problems even smaller, increasing their yield through the sieve.

This speeds up the algorithm for numbers of the form $an + b$, such as F_9, where the algebra is relatively easy. A theoretical generalization of the method is known, but it doesn't outperform the quadratic sieve until the numbers get up into the 200-digit range—where neither algorithm is currently practical.

That's good news for at least one group: cryptographers who make a living using big numbers to protect things like transfers between Swiss bank accounts. Schemes for encoding data based on number theory are just coming into commercial use, and a generalized, efficient factoring algorithm could bring the robbers up to speed with the cops.

But that hasn't happened yet. So far the cryptographers have been able to stay out ahead of the factoring community. In fact, the advances that make it possible to factor 100-digit numbers also make it possible to use codes based on 200-digit numbers. "We have hardware technology on our side," says Burt Kaliski, a cryptographic systems scientist at RSA Data Security, Inc., in Redwood City, California. The factoring of F_9 "doesn't threaten our business," he adds. "It only confirms that our estimates of the difficulty of factoring are accurate."

And what is Manasse and Lenstra's next big-number target? The obvious successor at the top of the most wanted list is F_{10}. But that's out of the question for now. Manasse estimates that factoring F_{10} with existing methods would require half a million times the resources that went into breaking down F_9 into its components.

It might be possible to manage a tenfold increase in resources, Manasse says, but that still leaves a factor of 50,000 to be accounted for. And when you multiply 50,000 by the 2 months that it took to pick apart F_9, Manasse says, you're starting to "talk about some serious computing time." ■ BARRY CIPRA