

The Smart Money's on Numerical Analysts

I'd like to mention two of the biggest unsolved problems in numerical analysis. One of them, if solved affirmatively, would change the world. The other would have little practical effect but would settle a mystery that has been standing for fifty years.

The first problem is, can an $n \times n$ linear system of equations $Ax = b$ be solved in $O(n^{2+\epsilon})$ operations, where ϵ is arbitrarily small? In other words, could there be a "fast matrix inverse," like the fast Fourier transform?

The standard algorithms, like Gaussian elimination, take $O(n^3)$ operations, and you can see that cubic exponent reflected in the history of computing. In sixty years, computer speeds have increased by 15 orders of magnitude, from flops to petaflops. With a fast matrix inverse and plenty of memory, the dimension of matrix problems we could handle would have grown from $n = 10$ to $n = 10^{8.5}$. In fact, we are nowhere near that scale.

It has been known for more than 40 years that, in theory, the exponent of 3 can be beaten. Strassen reduced it to 2.81 in 1969, and Coppersmith and Winograd reduced it further in 1990, to 2.376. Despite much conceptual progress since then, however, especially in the past few years, the exponent has improved only to 2.373. And the algorithms implicit in this highly theoretical work are complicated, too complicated to have much promise in practice. It is an article of faith for some of us that if $O(n^{2+\epsilon})$ is ever achieved, the big idea that achieves it will correspond to an algorithm that is really practical.

FROM THE SIAM PRESIDENT

By *Nick Trefethen*

My own faith has been strong for years. Back in 1985, I made a \$100 bet with Peter Alfeld that a fast matrix inverse would be found by 1995. It wasn't, so I paid him \$100 and renewed the bet for another decade. It still hadn't been found by 2005, so I paid him a further \$100 and we renewed once more. One day I will win!—or my heirs?

If a fast matrix inverse is found, everything in computational science will be shaken up. That's what happened with the FFT after 1965, and the fast matrix inverse would be equally epochal. I believe it is in the same category of importance as the celebrated $P = NP?$

The second big unsolved problem is also in the area of linear algebra. Why is Gaussian elimination numerically stable in practice, even though it's unstable in the worst case? In other words, why does Gaussian elimination work?

For a brief period in the 1940s, it was believed that Gaussian elimination would not work, that it would be undone by rounding errors on a computer. Experiments soon showed, however, that it was utterly reliable in practice, and Jim Wilkinson became famous by developing a theory that went partway to explaining this. Wilkinson showed that an algorithm will have all the accuracy properties one could reasonably ask for if it is "backward stable," delivering the exact answer to a slightly perturbed problem. It was a beautiful and powerful observation that has been a centerpiece of numerical analysis ever since.

The trouble is, though most algorithms of numerical linear algebra are backward stable, Gaussian elimination is not. (Throughout this discussion I am speaking of the standard variant of Gaussian elimination, with row or "partial" pivoting.) Wilkinson showed that its stability can be quantified by a number known as the "growth factor," ρ , the ratio of the largest entry in the final upper-triangular matrix U to the largest entry in the original matrix A . Unfortunately, although ρ is usually small, it can be as large as 2^{n-1} if you pick A diabolically. All this was already known and published by Wilkinson in a landmark paper in 1961.

Given the existence of such matrices, why doesn't Gaussian elimination fail in practice? If you look in the textbooks, you will find that they do not confront this question mathematically but make empirical statements like "In practice such growth is very rare," or "This bound is much too pessimistic in most cases." My favourite analysis is the one attributed to Wilkinson himself:

"Anyone that unlucky has already been run over by a bus!"

So what is going on? It is clear from experiments that the fraction of matrices with dangerously large growth factors ρ is extraordinarily small. One might imagine various ways to make this idea precise, and I believe that a good approach is this: For random matrices with independent normally distributed entries, how fast does the probability of encountering a growth factor $>\rho$ decrease as a function of ρ ? Experiments show that it decreases exponentially, or at least faster than any inverse power of $\rho/n^{1/2}$ —see Fig. 22.2 in my textbook with Bau. But nobody has managed to prove this. Years ago I proved it for the bottom-right entry of U , but I was never able to extend the proof to the other entries, all of which enter into the definition of ρ .

Like the problem of the fast matrix inverse, the problem of the stability of Gaussian elimination has seen conceptual advances in recent years. Spielman, Teng, and Sankar have developed smoothed analysis, which gives a new probabilistic twist to algorithmic questions (*SIAM Journal on Matrix Analysis and Applications*, 2006; also Sankar's 2004 thesis at MIT). Spielman was recently named one of this year's MacArthur Fellows, partly for this work. Smoothed analysis hasn't yet solved the Gaussian elimination problem, however, and in fact, its best known results don't distinguish between pivoted elimination, whose striking stability we are trying to explain, and the unpivoted variant, which is unstable.

Precisely posed unsolved problems are one of the glories of mathematics. Even in applied mathematics, for all our close connections to practical applications, they can motivate and inspire us. I will give \$1000 for a proof that Gaussian elimination with partial pivoting is stable in the probabilistic sense specified above.