

**October 14, 1999**

**James Demmel, UCB**

## **Throwing away information for fun and profit**

Solving systems of linear equations, and finding eigenvalues of symmetric matrices are standard problems that arise in many mathematical modeling problems. Solutions have existed for a long time, but these solutions are no longer good enough, because the demands of applications keep growing to larger problems, and to more 'sensitive' problems that require higher accuracy. In addition, the larger problems in turn require large parallel computers where communication is much more expensive than computation.

We highlight several new algorithms under development, all of which have solved linear systems and eigenvalue problems many times larger, more sensitive or more quickly than before. Though these algorithms are all very different, one common mathematical approach distinguishes them from their predecessors: The new algorithms deliberately throw away just enough information about the problem to run both efficiently yet accurately. In contrast, their predecessors would either have to use much higher precision, send many more messages, or do many more operations to get the same answers.

Finally, we also describe one problem, computing reliable error bounds, where willingness to sacrifice information and accuracy provably does not help the complexity.