Correcting errors in a matrix inverse

Daniel S. Roche*

We consider the problem of computing the inverse of a matrix A, given a matrix B which is close to the inverse of A with a small number of erroneous entries. We show how to recover the true inverse of A in roughly $\tilde{\mathrm{O}}(n^2 + k^w)$ time, where n is the dimension of the matrix, k is the number of errors, and w is the exponent of fast matrix multiplication. With this running time, for sufficiently small number of errors k, the time to correct the inverse is linear in the size of the matrix itself, and is therefore worthwhile over the trivial solution of simply recomputing it. This can be seen as a continuation of work by Gasiniec et al on efficiently correcting errors in a matrix product, as well as an extension of recent papers at ISSAC and elsewhere on efficiently verifying the results of linear algebra computations. The main application is to distributed or outsourced computing, where errors can be introduced by a small number of faulty nodes or by network noise.

*U.S. Naval Academy, Annapolis, Maryland, USA; currently at Laboratoire Jean Kuntzmann, Université Grenoble Alpes, Grenoble, France