ARNE STORJOHANN, University of Waterloo
Faster algorithms for the Frobenius canonical form

Like the better known Jordan form, the Frobenius form of a square matrix A over a field K is a unique representative of the set of all matrices similar to A. The Frobenius form captures completely the geometric structure of a matrix, and clearly reveals invariants such as the minimal and characteristic polynomial.

The problem of computing the form has been very well studied. Let $2 < \theta \leq 3$ be such that $O(n^\theta)$ operations in K are sufficient to multiply together two matrices in $\mathsf{K}^{n\times n}$. Over a sufficiently large field, with $\#\mathsf{K} \geq n^2$, Giesbrecht's (1993) randomized algorithm computes the Frobenius form F, together with a similarity transformation matrix U such that $F = U^{-1}AU$, using an expected number of $O(n^\theta \log n)$ operations in K. More recently, Eberly (2000) describes a randomized algorithm, applicable over a field K of any size, that computes both the form and a transformation matrix using an expected number of $O(n^\theta \log n)$ operations in K.

In this talk I describe a new randomized algorithm for computing the Frobenius form. Over a sufficiently large field the new algorithm uses an expected number of $O(n^\theta)$ field operations, thus improving by a factor of $\log n$ on previously known results. Once the Frobenius form itself has been computed, a similarity transformation matrix can be constructed using an additional $O(n^\theta \log\log n)$ field operations.

Joint work with Clément Pernet.