22 Towards fill-in analysis

We finish our section on Cholesky factorisation by taking a more graph theoretic view of the occurrence of fill-in. This is a large topic, and we have only a lecture to devote to it, so we will just summarise some of the main ideas, and emphasise examples, rather than theoretical rigour. The underlying questions are:

(a) Given a sparse matrix A, can we predict where the non-zero entries L occur;

(b) Can we permute the rows and columns of A to reduce fill-in.

Question (a) is motivated by the fact that, in practice, when we compute L we first need to know its *structure*. This is because, on your compute, a sparse matrix are not stored as a rectangular array of numbers but in some other format such as *triplet* or *compressed column format*. (I'll give a brief explanation of these, but they are not of great mathematical interest).

Question (b) is motivated by the "arrow matrices" that we saw in Section 21.4; due to time constraints we won't dwell on it.

For more information, please read Secvions 1.2 and Chapter 4 of [Dav06].

. . . .

Surprisingly, most of the class was spent explaining and discussing these ideas. We looked at

- The triplet form of sparse matrix storage
- Revisited the arrow matrices
- Discussed why, rather than permuting just the rows of a matrix, we permute the rows *and* the column.
- Discussed why the Cholesky factor of the matrix A could be so different from PAP^T.