Abstract

In this thesis we consider sequential and parallel algorithms for the factorization of sparse, positive definite, linear systems. We especially focus on the computation of fill-reducing orderings for a general sparse matrix A.

Chapter 1 shows the importance of sparse direct methods in the field of Scientific Computing and compares Cholesky's method with iterative algorithms. The number of iterations performed by an iterative scheme highly depends on the numerical characteristics of A, whereas the performance of a direct algorithm is only determined by A's nonzero structure.

Chapter 2 introduces two different schemes for the computation of the factor matrix L, $A = LL^T$: the fan-in and the fan-out method. Both methods perform the same number of arithmetic operations – they only differ in the way the matrix elements are accessed. Furthermore, we show how the columnwise factorization of A can be modeled by a sequence of elimination graphs. Finally, we briefly describe three general classes of ordering algorithms: profile methods, bottom-up methods, and top-down methods.

In Chapter 3 we consider fill-reducing algorithms for Laplace matrices that correspond to $n \times h$ grid problems with a five point difference operator. If the $+$-shaped separators constructed by a classical nested dissection scheme are rotated to form $\times$-shape separators, the amount of fill and floating point operations will be reduced significantly. A detailed analysis of the modified nested dissection ordering shows that the shape of the elements formed during the elimination process is responsible for this improvement.

Chapter 4 provides us with a new fill-reducing ordering scheme for general sparse matrices that is based on the observations made in Chapter 3. Most state-of-the-art ordering schemes for general sparse matrices are a hybrid of a bottom-up method such as minimum degree and a top down scheme such as George's nested dissection. In this chapter we present an ordering algorithm that achieves a tighter coupling of bottom-up and top-down methods. In our methodology vertex separators are interpreted as the boundaries of the remaining elements in an unfinished bottom-up ordering. As a consequence, we are using bottom-up techniques such as quotient graphs and special node selection strategies for the construction of vertex separators. Once all separators have been found, we are using them as a skeleton for the computation of several bottom-up orderings. Experimental results show that the orderings obtained by our scheme are in general better than those obtained by other popular ordering codes.

In Chapter 5 we describe sequential and parallel algorithms for the symbolical and numerical factorization. We first introduce techniques that improve cache efficiency and register re-use of the numerical factorization process. To fully exploit the floating-point performance of modern computers we implement a multifrontal algorithm that is based on Level 3 BLAS operations. We then describe our parallel algorithm that uses a 2-dimensional mapping scheme of A to reduce the comunication overhead. We show that the orderings produced by the new scheme are also well suited for parallel factorization.

Chapter 6 summarizes the methods that are described in this thesis and discusses some open problems.