ELSEVIER

Parallel Computing 25 (1999) 1995–2014

PARALLEL COMPUTING

www.elsevier.com/locate/parco

Ordering strategies and related techniques to overcome the trade-off between parallelism and convergence in incomplete factorizations

Shun Doi[a,*], Takumi Washio[b]

[a] C&C Media Research Laboratories, NEC Corporation, 4-1-1 Miyazaki, Miyamae-ku, Kawasaki, Kanagawa 216-8111, Japan

[b] C&C Research Laboratories, NEC Europe, Rathausallee 10 D-53757, Sankt Augustin, Germany

Abstract

This paper is concerned with the parallel implementation of the incomplete factorization preconditioned iterative method. Although the use of such parallel ordering as multicolor ordering may increase parallelism in factorization, it often slows convergence when used in the preconditioned method, and thus may offset the gain in speed obtained with parallelization. Further, the higher the parallelism of an ordering, the slower the convergence; the lower the parallelism, the faster the convergence. This well-known *trade-off between parallelism and convergence* is well explained by the property of *compatibility*, the level of which can be clearly seen when ordering is presented in graph form (S. Doi, A. Lichnewsky, A graph-theory approach for analyzing the effects of ordering on ILU preconditioning, INRIA report 1452, 1991). In any given method, the fewer the *incompatible local graphs* in an ordering (i.e., the lower the parallelism), the faster the convergence (S. Doi, Appl. Numer. Math. 7 (1991) 417–436; S. Doi, in: T. Nodera (Ed.), Advances in Numerical Methods for Large Sparse Sets of Linear Systems, 7 Keio University, 1991). An ordering with no *incompatible* local graphs, for example, such as that implemented on vector multiprocessors by using the *nested dissection* technique, will have excellent convergence, but its parallelism will be limited (S. Doi, A. Lichnewsky, Int. J. High Speed Comput. 2 (1990) 143–179). To attain a better balance, a certain degree of *incompatibility* is necessary. In this regard, increasing the number of colors in *multicolor* ordering can be a useful approach (S. Fujino, S. Doi, in: R. Beauwens (Ed.), Proceeding of the IMACS Internation Symposium on Iterative Methods in Linear Algebra, March 1991; S. Doi, A. Hoshi, Int. J. Comput. Meth. 44 (1992) 143–152). Two related techniques also presented here are the overlapped multicolor ordering (T. Washio, K. Hayami, SIAM J. Sci. Comput. 16 (1995) 631–650), and a fill-in strategy selectively applied to *incompatible* local graphs. Experiments conducted with an SX-5/16A vector parallel supercomputer

* Corresponding author.

PII: S 0 1 6 7 - 8 1 9 1 (9 9) 0 0 0 6 4 - 2

show the relative effectiveness of increasing the number of colors and also of using this approach in combination with overlapping and with fill-ins.

Keywords: Sparse linear systems; Preconditioning; Iterative methods; Incomplete factorization; Convergence; Vector computer

1. Introduction

A class of the most effective methods for solving a system of large sparse linear equations is the preconditioned iterative method, which can be defined as a basic iterative method applied to a preconditioned system $M^{-1}Au = M^{-1}b$ (where M is the preconditioner) instead of being applied to the original system $Au = b$. The objective of the preconditioning is to reduce the condition number (or to cluster the eigenvalues) of the original system so as to reach an approximate solution with fewer iterations. This requires solving the system $Mv = g$. Hence, a good preconditioner needs to satisfy two requirements: it should be a good enough approximation of A to produce a reduced condition number (or better clusters of eigenvalues), and the system $Mv = g$ should be much easier to solve than the original system. Further, increases in the parallelism of recent computer architectures have led to a new requirement: the solution of $Mv = g$ must also have high enough parallelism to be mapped naturally on the computer to be used.

One type of commonly used preconditioning is represented by the incomplete factorization $M = LU$, where L and U are lower and upper triangular matrices of a sparsity structure similar to A, and are produced by neglecting certain fill-ins in the Gaussian elimination [19]. One way to attain parallelism in the solution of $Mv = LUv = g$ (in practice, this solution is obtained first by solving $Lw = g$ (forward substitution) and then by solving $Uv = w$ (backward substitution)) is to reorder the system $Au = b$ and to reconstruct an incomplete factorization for the reordered system. Actually, a number of techniques developed in the past do this [1,3,5,7,10,12–14,17,18,21]. In this regard, one possibility for this approach is to use the well-known *red–black* ordering. Solving $Lw = g$, where w has a red–black ordering can be performed simultaneously for half the unknowns in v; this is also true of the solving of $Uv = w$. In this case, however, a fundamental difficulty is presented by the *trade-off problem between parallelism and convergence*: in a system to which the incomplete factorization preconditioned iterative method applied, higher parallelism in the ordering will result in slower convergence, and higher convergence will require lower parallelism. This trade-off problem was first reported by Duff and Meurant [10] on the basis of intensive numerical tests using many orderings [10]. Since then, many attempts have been made to answer the following two fundamental questions: *why* there is a trade-off between parallelism and convergence in incomplete factorizations and *how* can the trade-off be overcome [2,3,6,10,11,16,20]? Before we give our own answers to these questions, it is worth repeating the remarks of Duff and Meurant [10] which are still valuable: