

Solution of Unassembled Linear Systems Using Block Stretching: Preliminary Experiments.

Michel J. Daydé¹, Jérôme P. Décamps¹, Nicholas I. M. Gould²
ENSEEIHT-IRIT Technical Report, **RT/APO/97/3**

December 3, 1997

Abstract

We consider so-called “matrix stretching” technique that make structured unassembled linear systems larger, but sparser. Our solution technique combines a direct factorization of the leading block diagonal submatrix of the stretched system, with a preconditioned conjugate gradient solution of the Schur complement system which results from the factorization of the diagonal blocks.

We show that matrix stretching is an effective technique, particularly for ill-conditioned systems. The Schur complement is often considerably better conditioned than the whole system. The main challenge is to find a suitable preconditioner for this matrix. We consider a range of preconditioners, including those proposed by Chan, and band approximations. We also study the use of some Element-by-Element preconditioners such as EBE and the recently introduced Subspace-by-Subspace preconditioner.

We report on experiments using structured problems and examples from the Harwell-Boeing sparse matrix collection. We also report on some preliminary parallel experiments that show how stretching improves the parallelisation of the linear solver.

1 Introduction

Stretching, a sparse matrix preprocessing technique that makes matrices sparser but, at the same time, larger was first introduced by Grcar (1990). He proposed two general stretching techniques, simple row and column stretching, and studied their efficacy on the factorization of arrow-head matrices. Alvarado (1997) showed that such techniques are an effective way of treating matrices with dense rows or columns before forming their LDU or QR factorizations. Similar ideas have been developed by Vanderbei (1991) and Andersen (1996) for solving the Schur complements arising from interior point methods in the context of linear programming.

In this paper, we apply matrix stretching to the solution of a linear systems of “finite-element” type without assembling the coefficient matrices from its “elements”. This requires that we introduce extra variables to recouple the “element” blocks of the matrix. The resulting enlarged system is a sparse “augmented system”, which can be solved using a range of direct and iterative methods.

To be precise, we consider the solution of the symmetric system of linear equations of the form

$$Bx = b, \tag{1.1}$$

¹ENSEEIHT-IRIT, 2 rue Camichel, 31071 Toulouse CEDEX, France

²Central Computing Department, Rutherford Appleton Laboratory, Oxfordshire, OX11 0QX, England.

where \mathbf{B} can be expressed as

$$\mathbf{B} = \sum_{i=1}^{n_e} \mathbf{B}_i, \quad (1.2)$$

and where each symmetric elementary matrix \mathbf{B}_i only involves a small subset of the variables \mathbf{x} . We let \mathbf{B}^i be the matrix obtained from \mathbf{B}_i by removing its zero rows and columns. Note that we do not assume that each \mathbf{B}_i is positive definite, merely that it is non-singular.

Linear systems of the form (1.1)–(1.2) arise when solving the Newton equations for the unconstrained minimization of a partially separable function (Griewank and Toint, 1982*a*, 1982*b*), and from finite-element methods for the solution of partial-differential equations (see, for instance, Zienkiewicz, 1977).

The stretched system we obtain is an augmented system of the form

$$\begin{pmatrix} \mathbf{B}^S & \mathbf{A} \\ \mathbf{A}^T & \mathbf{0} \end{pmatrix} \begin{pmatrix} \mathbf{x}^S \\ \lambda \end{pmatrix} = \begin{pmatrix} \mathbf{b}^S \\ \mathbf{0} \end{pmatrix} \quad (1.3)$$

where \mathbf{B}^S is block diagonal and \mathbf{A} is very sparse. The diagonal blocks of \mathbf{B}^S are the \mathbf{B}^i . The matrix \mathbf{A} is used to recouple the \mathbf{B}_i when they share variables. As we shall shortly show, the required solution \mathbf{x} may easily be recovered from \mathbf{x}^S .

The solution \mathbf{x}^S to (1.3) also solves the quadratic programming problem

$$\begin{aligned} \text{minimize} \quad & q(\mathbf{x}^S) = \frac{1}{2} \langle \mathbf{x}^S, \mathbf{B}^S \mathbf{x}^S \rangle - \langle \mathbf{b}^S, \mathbf{x}^S \rangle \quad \text{subject to} \quad \mathbf{A}^T \mathbf{x}^S = \mathbf{0}, \\ & \mathbf{x}^S \in \mathbb{R}^{n_b} \end{aligned} \quad (1.4)$$

where n_b is the dimension of \mathbf{B}^S and $\langle \cdot, \cdot \rangle$ denotes the Euclidean inner-product (see for instance Gill, Murray and Wright, 1981). The augmented system (1.3) or the quadratic program (1.4) may be solved in many ways (see for instance Gill et al., 1981), and we choose the well-known Schur complement approach. That is, we perform a block decomposition of the coefficient matrix

$$\begin{pmatrix} \mathbf{B}^S & \mathbf{A} \\ \mathbf{A}^T & \mathbf{0} \end{pmatrix} \quad (1.5)$$

by factorizing each of the diagonal blocks of \mathbf{B}^S , and use the method of preconditioned conjugate gradients (PCG) to solve the remaining linear system which results following elimination of these leading blocks.

2 Description of the stretching algorithm

2.1 Introduction: a simple example

We consider the solution of $\mathbf{B}\mathbf{x} = \mathbf{b}$ where \mathbf{B} is structured as

$$\begin{pmatrix} 8 & 1 & 1 & & \\ 1 & 8 & 1 & & \\ 1 & 1 & 8 & 1 & 1 \\ & & 1 & 8 & 1 \\ & & 1 & 1 & 8 \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \\ x_5 \end{pmatrix} = \begin{pmatrix} b_1 \\ b_2 \\ b_3 \\ b_4 \\ b_5 \end{pmatrix}. \quad (2.6)$$

Additionally, we assume that \mathbf{B} is composed of two elements \mathbf{B}_1 and \mathbf{B}_2 — \mathbf{B}_1 involves variables x_1, x_2, x_3 , while \mathbf{B}_2 involves x_3, x_4, x_5 . Thus

$$\mathbf{B}^1 = \begin{pmatrix} 8 & 1 & 1 \\ 1 & 8 & 1 \\ 1 & 1 & 4 \end{pmatrix} \quad \text{and} \quad \mathbf{B}^2 = \begin{pmatrix} 4 & 1 & 1 \\ 1 & 8 & 1 \\ 1 & 1 & 8 \end{pmatrix}$$

(Such an example might arise, for instance, if the linear system is the Newton equation arising from the unconstrained minimization of the partially separable cost function $f(x) = f^1(x) + f^2(x)$ with $f^1(x) = f^1(x_1, x_2, x_3)$ and $f^2(x) = f^2(x_3, x_4, x_5)$. In this case, \mathbf{b} is the negative of the gradient of f , \mathbf{B} is the Hessian of f , and \mathbf{B}^1 and \mathbf{B}^2 are respectively the projected Hessians of f^1 and f^2 onto the variables used by f^1 and f^2 .) The “overlap” in this simple example only involves the variable x_3 .

Since we do not want to assemble B , we expand—or stretch—the system $Bx = b$ as follows. Let \mathbf{B}^S be the matrix obtained by expanding \mathbf{B} without assembling the overlapping parts, ie,

$$\mathbf{B}^S = \begin{pmatrix} 8 & 1 & 1 & & & \\ 1 & 8 & 1 & & & \\ 1 & 1 & 4 & & & \\ & & & 4 & 1 & 1 \\ & & & 1 & 8 & 1 \\ & & & 1 & 1 & 8 \end{pmatrix}. \quad (2.7)$$

Our aim is to replace (2.6) by a system in which (2.7) is a submatrix. We look for a new solution vector $(x_1^S, \dots, x_6^S)^T$, where we identify

$$x_1 = x_1^S, \quad x_2 = x_2^S, \quad x_4 = x_5^S \quad \text{and} \quad x_5 = x_6^S \quad (2.8)$$

since these unknowns do not occur in the overlap. We also require that

$$x_3 = x_3^S = x_4^S. \quad (2.9)$$

In addition, if we let

$$x_1^S + x_2^S + 4x_3^S = \omega_1 \quad \text{and} \quad 4x_4^S + x_5^S + x_6^S = \lambda_1, \quad (2.10)$$

then the third equation in (2.6) requires that

$$\omega_1 + \lambda_1 = b_3. \quad (2.11)$$

Combining (2.6)–(2.10) and eliminating ω_1 using (2.11) then yields the symmetric augmented system

$$\begin{pmatrix} 8 & 1 & 1 & & & & & & & & \\ 1 & 8 & 1 & & & & & & & & \\ 1 & 1 & 4 & & & & & & & & \\ & & & 4 & 1 & 1 & & & & & \\ & & & 1 & 8 & 1 & & & & & \\ & & & 1 & 1 & 8 & & & & & \\ & & & +1 & -1 & & & & & & \end{pmatrix} \begin{pmatrix} x_1^S \\ x_2^S \\ x_3^S \\ x_4^S \\ x_5^S \\ x_6^S \\ \lambda_1 \end{pmatrix} = \begin{pmatrix} b_1 \\ b_2 \\ b_3 \\ 0 \\ b_4 \\ b_5 \\ 0 \end{pmatrix}. \quad (2.12)$$

2.2 The stretching algorithm

The general algorithm for building the required stretched system is simple. Suppose \mathbf{B} is of order n . We then construct a system of the form

$$\begin{pmatrix} \mathbf{B}_1^S & & & & \mathbf{A}_1 \\ & \mathbf{B}_2^S & & & \mathbf{A}_2 \\ & & \ddots & & \vdots \\ & & & \mathbf{B}_{n_e}^S & \mathbf{A}_{n_e} \\ \mathbf{A}_1^T & \mathbf{A}_2^T & \dots & \mathbf{A}_{n_e}^T & \mathbf{0} \end{pmatrix} \begin{pmatrix} \mathbf{x}_1^S \\ \mathbf{x}_2^S \\ \vdots \\ \mathbf{x}_{n_e}^S \\ \lambda \end{pmatrix} = \begin{pmatrix} \mathbf{b}_1^S \\ \mathbf{b}_2^S \\ \vdots \\ \mathbf{b}_{n_e}^S \\ \mathbf{0} \end{pmatrix} \quad (2.13)$$

using the algorithm described in Algorithm 2.1.

Algorithm 2.1: The stretching algorithm**Initialization of \mathbf{B}^S**

For each element $i = 1, \dots, n_e$

\mathbf{B}_i^S is the matrix \mathbf{B}^i

Initialization of \mathbf{A} and \mathbf{b}^S

For each variable $i = 1, \dots, n$

Compute $list$, the list of elements possessing i .

If i is shared by $deg(i)$ elements,

create $(deg(i)-1)$ columns for each \mathbf{A}_j for $j = 1, n_e$.

$list(1)$ is the first element possessing i

Let k be the row of $\mathbf{B}_{list(1)}^S$ associated with variable i .

Set the $(deg(i) - 1)$ columns of row k of $\mathbf{A}_{list(1)}$ to 1.

Set the corresponding value in \mathbf{b}^S to b_i .

For $j = 2, \dots, deg(i)$

Let k be the row of $\mathbf{B}_{list(j)}^S$ associated with variable i .

Set column $j - 1$ of row k of $\mathbf{A}_{list(j)}$ to -1

Set the corresponding value in \mathbf{b}^S to 0.

Having solved the stretched system, we simply recover \mathbf{x} from \mathbf{x}^S , by setting \mathbf{x}_i to its value in one of the blocks involving it. A number of comments are in order.

1. We start the stretching process for variable i with the first element in which it is involved. Other strategies may prove more effective, and further experiments are required.
2. It is easy to compute the number of variables involved in the augmented system. Let n_o be the number of variables shared by several elements, $deg(i)$ the number of elements in which the i^{th} variable appears, and $n_v(i)$ the number of variables involved in element i . Then the augmented system involves $n + \sum_{i=1}^{n_o} 2(deg(i) - 1)$ variables, of which $\sum_{i=1}^{n_o} (deg(i) - 1)$ make up $\boldsymbol{\lambda}$. The matrix \mathbf{A} has $n_s \stackrel{def}{=} \sum_{i_e=1}^{n_e} n_v(i_e)$ rows and $\sum_{i=1}^{n_o} (deg(i) - 1)$ columns, while \mathbf{B}^S is of order $\sum_{i_e=1}^{n_e} n_v(i_e)$; \mathbf{A} has $2 \sum_{i=1}^{n_o} (deg(i) - 1)$ nonzeros (two nonzeros per column).
3. If we are solving a sequence of problems of the same form (1.2), we only need to construct \mathbf{A} for the first such problem. For subsequent problems, we merely update the nonzeros of \mathbf{b}^S and \mathbf{B}^S .

2.3 A further example

In Figure 2.1, we consider a 6×6 matrix composed of 4 elements. Here, element 1 uses variables 1, 2, 3, 4, 5 and 6, element 2 uses 2, 3 and 4, element 3 uses 3, 4, 5 and 6, and element 4 uses variables 4 and 5.

When using the method described, we obtain the augmented system illustrated in Figure 2.2. Here, \mathbf{B}^S is a 15×15 matrix composed of the 4 unassembled elements, \mathbf{A} is a 15×9 matrix with 18 nonzeros, and $\boldsymbol{\lambda}$ is a vector of 9 elements.

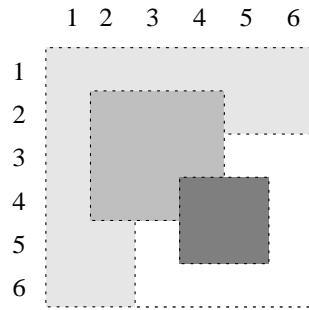


Figure 2.1: The initial matrix and its elements.

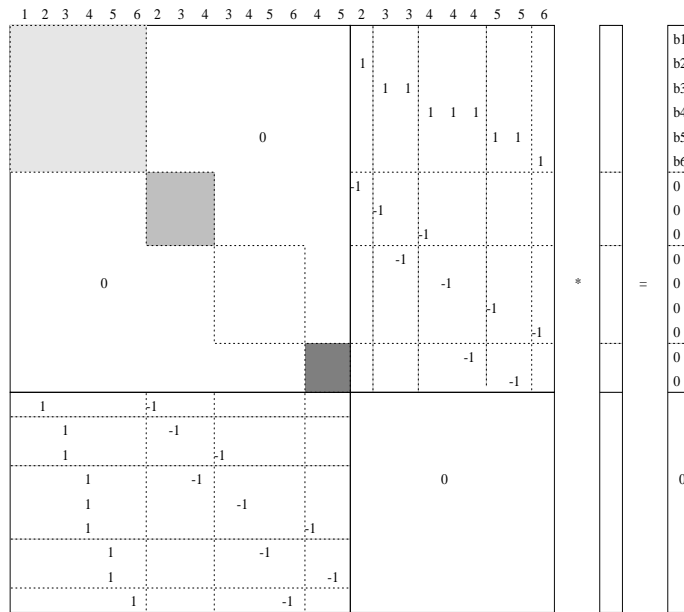


Figure 2.2: The stretched system.

3 The solution of stretched linear systems

The Schur complement method (see Haynsworth, 1968 and Axelsson, 1996) takes advantage of the structure of the system (2.13). The method is formalized as Algorithm 3.1.

Algorithm 3.1: The Schur complement algorithm on the stretched system

1. $\forall i \in 1, \dots, n_e$, solve $\mathbf{B}_i^S \mathbf{z}_i = \mathbf{b}_i^S$ (using the LDL^T factorization from LAPACK, see Anderson, Bai, Bischof, Demmel, Dongarra, Croz, Greenbaum, Hammarling, McKenney and Sorensen, 1990)
2. Use the PCG method to solve $\mathbf{S}\boldsymbol{\lambda} = \mathbf{s}$, where the n_s by n_s Schur complement $\mathbf{S} = \sum_{i=1}^{n_e} \mathbf{A}_i^T (\mathbf{B}_i^S)^{-1} \mathbf{A}_i$ and $\mathbf{s} = \sum_{i=1}^{n_e} \mathbf{A}_i^T \mathbf{z}_i$
3. Solve $\mathbf{B}_i^S \mathbf{x}_i^S = \mathbf{b}^S + \mathbf{A}_i \boldsymbol{\lambda}$ (using LAPACK)
4. $\forall k \in 1, \dots, n$, recover x_k from its value in \mathbf{x}_i^S , where (e.g.) \mathbf{B}_i^S is the last element involving x_k .

This approach offers great opportunities for parallelisation since Steps 1 and 3 can be computed element-wise in parallel. Moreover, one hopes, in Step 2, to exploit parallelism within the PCG iteration — conjugate gradients (CG) are chosen since we may not want to form \mathbf{S} . We use the LDL^T factorization of the elements formed during the step 1 to compute the Schur complement. We then obtain

$$\mathbf{S} = \sum_{i=1}^{n_e} \mathbf{A}_i^T (\mathbf{B}_i^S)^{-1} \mathbf{A}_i = \sum_{i=1}^{n_b} \mathbf{s}_i \mathbf{s}_i^T, \quad (3.14)$$

where n_b is the dimension of \mathbf{B}^S in (1.3). This decomposition proves to be useful during the construction of all but the probing-vector preconditioner we consider in the next section. The remaining ingredient is, of course, the choice of preconditioner.

4 Preconditioning the Schur complement

We have used a number of different preconditioners. These include

- probing-vector preconditioners — the Chan diagonal and band preconditioners,
- traditional band preconditioners, and
- Element-by-Element (EBE) and Subspace-by-Subspace (SBS) preconditioners.

We consider each in some detail.

4.1 Probing-vector preconditioners

The aim here is to estimate values of \mathbf{S} merely by forming products of \mathbf{S} with appropriate vectors. This avoids the need to explicitly form \mathbf{S} .

4.2 Traditional band preconditioners

We might equally assemble a band submatrix of the Schur complement, and factorize it using the LANCELOT band factorization. Unlike the probing vector preconditioners, the band submatrix computed in this way is the exact band restriction of the Schur complement \mathbf{S} .

4.3 Element-by-Element preconditioners

We can construct the elements of the Schur complement as the sum $\sum_{i=1}^{n_e} \mathbf{S}_i$ where $\mathbf{S}_i = \mathbf{A}_i^T \mathbf{B}_i^{-1} \mathbf{A}_i$ and then form an EBE preconditioner based upon the unassembled structure of the Schur complement. Alternatively, we can use the decomposition

$$\mathbf{S} = \sum_{i=1}^{n_e} \mathbf{S}_i = \sum_{i=1}^{n_b} \mathbf{s}_i \mathbf{s}_i^T, \quad (4.15)$$

where the $\mathbf{s}_i \mathbf{s}_i^T$ are rank-one contributions to the elements \mathbf{S}_i .

Element-By-Element (EBE) preconditioners were introduced by Hughes, Levit and Winget (1983) and Ortiz, Pinsky and Taylor (1983) and have been successfully applied in a number of applications in engineering and physics (see, for example, Hughes, Ferencz and Hallquits, 1987, and Erhel, Traynard and Vidrascu, 1991), and in the solution of partially separable linear systems (see Daydé, L'Excellent and Gould, 1997b). A detailed analysis of this technique for certain problems arising from partial-differential equations is given by Wathen (1989).

Let $\Delta(\mathbf{S})$ and $\Delta(\mathbf{S}_i)$ be the diagonal parts of \mathbf{S} and \mathbf{S}_i respectively. Then we obtain the EBE preconditioner

$$\mathbf{P}_{EBE} = \Delta(\mathbf{S})^{1/2} \left\{ \prod_{i=1}^{n_e} \mathbf{L}_i \prod_{i=1}^{n_e} \mathbf{D}_i \prod_{i=n_e}^1 \mathbf{L}_i^T \right\} \Delta(\mathbf{S})^{1/2},$$

where \mathbf{L}_i and \mathbf{D}_i are the LDL^T factors of the Winget terms

$$\mathbf{L}_i \mathbf{D}_i \mathbf{L}_i^T = \mathbf{W}_i \equiv \mathbf{I} + \Delta(\mathbf{S})^{-1/2} (\mathbf{S}_i - \Delta(\mathbf{S}_i)) \Delta(\mathbf{S})^{-1/2}$$

4.4 Subspace-by-Subspace preconditioners

More recently, Daydé, Décamps and Gould (1997a) have introduced a class of so-called Subspace-by-Subspace (SBS) preconditioners. These techniques consist of regrouping the rank-one terms in (4.15), and forming and applying a QR factorization of the each of the resulting low-rank matrices. If we use the decomposition of the Schur complement into n_b rank-one terms as indicated by (3.14), we may rewrite \mathbf{S} as

$$\mathbf{S} = \Delta(\mathbf{S})^{1/2} \left(\mathbf{I} + \sum_{i=1}^{n_b} \mathbf{E}_i \right) \Delta(\mathbf{S})^{1/2}, \quad (4.16)$$

where $\mathbf{E}_i = \mathbf{W}_i - \mathbf{I}$, $\mathbf{W}_i = \mathbf{D}_i + \mathbf{h}_i \mathbf{h}_i^T$, $\mathbf{D}_i = \mathbf{I} - \Delta(\mathbf{S})^{-1/2} \Delta(\mathbf{s}_i \mathbf{s}_i^T) \Delta(\mathbf{S})^{-1/2}$, and $\mathbf{h}_i = \Delta(\mathbf{S})^{-1/2} \mathbf{s}_i$. If, instead, we group together the rank-one terms into r_n sets, we obtain

$$\mathbf{S} = \Delta(\mathbf{S})^{1/2} \left(\mathbf{I} + \sum_{j=1}^{r_n} \mathbf{E}_j \right) \Delta(\mathbf{S})^{1/2} \quad (4.17)$$

where $\mathbf{E}_j = \mathbf{W}_j - \mathbf{I}$, $\mathbf{W}_j = \mathbf{D}_j + \mathbf{H}_j \mathbf{H}_j^T$, $\mathbf{H}_j = [\mathbf{h}_{j_1}, \dots, \mathbf{h}_{j_n}]$ and $\mathbf{D}_j = \mathbf{I} - \Delta(\mathbf{S})^{-1/2} \sum_{i=i_1}^{i_n} [\Delta(\mathbf{s}_i \mathbf{s}_i^T)] \Delta(\mathbf{S})^{-1/2}$. Replacing the summation in (4.17) by a product, an approximation

of \mathbf{S} is

$$\mathbf{P}_1 = \Delta(\mathbf{S})^{1/2} \prod_{j=1}^{r_n} \mathbf{W}_j \Delta(\mathbf{S})^{1/2}$$

To use this as a preconditioner, we merely need to be able to invert the matrices \mathbf{W}_j . To this end, we can use the factorization of the diagonal matrix \mathbf{D}_j to obtain

$$\mathbf{W}_j = \mathbf{D}_j^{1/2} \left\{ \mathbf{I} + \mathbf{C}_j \mathbf{C}_j^T \right\} \mathbf{D}_j^{1/2},$$

where $\mathbf{C}_j = \mathbf{D}_j^{-1/2} \mathbf{H}_j$. Thus we must factorize matrices – omitting the indices – of the form $\mathbf{I} + \mathbf{C}\mathbf{C}^T$, where \mathbf{I} is a $n_s \times n_s$ matrix and \mathbf{C} a $n_s \times l$ matrix (n_s is the dimension of the Schur complement \mathbf{S} , while l is the number of rank-one terms grouped together). One approach is to assemble the whole n_s by n_s matrix and use a LDL^T factorization — this is essentially what the EBE methods mentioned earlier do. However, since this ignores the inherent structure, we might instead use the QR factorization of the rank- l matrix

$$\mathbf{C} = \mathbf{Q} \begin{pmatrix} \mathbf{R} \\ \mathbf{0} \end{pmatrix}, \quad \text{where } \mathbf{Q} = \begin{pmatrix} \mathbf{Y} & \mathbf{Z} \end{pmatrix},$$

\mathbf{Q} and \mathbf{R} are, respectively, $n_s \times n_s$ orthonormal and $k \times k$ triangular matrices, and \mathbf{Y} is the leading $n_s \times k$ submatrix of \mathbf{Q} . We then have that

$$\mathbf{I} + \mathbf{C}\mathbf{C}^T = \mathbf{Q} \begin{pmatrix} \mathbf{L} & \mathbf{0} \\ \mathbf{0} & \mathbf{I} \end{pmatrix} \mathbf{Q}^T \mathbf{Q} \begin{pmatrix} \mathbf{L}^T & \mathbf{0} \\ \mathbf{0} & \mathbf{I} \end{pmatrix} \mathbf{Q}^T = \mathbf{M}\mathbf{M}^T$$

where $\mathbf{L}\mathbf{L}^T = \mathbf{I} + \mathbf{R}\mathbf{R}^T$ and

$$\mathbf{M} = \mathbf{Q} \begin{pmatrix} \mathbf{L} & \mathbf{0} \\ \mathbf{0} & \mathbf{I} \end{pmatrix} \mathbf{Q}^T$$

Using the orthonormality of the columns of \mathbf{Q} , it immediately follows that the action of the inverses of the matrices \mathbf{M} and \mathbf{M}^T on the vector \mathbf{v} can be expressed as

$$\begin{aligned} \mathbf{M}^{-1}\mathbf{v} &= \mathbf{Q} \begin{pmatrix} \mathbf{L}^{-1} & \mathbf{0} \\ \mathbf{0} & \mathbf{I} \end{pmatrix} \mathbf{Q}^T \mathbf{v} = \begin{pmatrix} \mathbf{Y} & \mathbf{Z} \end{pmatrix} \begin{pmatrix} \mathbf{L}^{-1} & \mathbf{0} \\ \mathbf{0} & \mathbf{I} \end{pmatrix} \begin{pmatrix} \mathbf{Y}^T \\ \mathbf{Z}^T \end{pmatrix} \mathbf{v} \\ &= \left(\mathbf{I} + \mathbf{Y}(\mathbf{L}^{-1} - \mathbf{I})\mathbf{Y}^T \right) \mathbf{v} \quad \text{and} \\ \mathbf{M}^{-T}\mathbf{v} &= \left(\mathbf{I} + \mathbf{Y}(\mathbf{L}^{-T} - \mathbf{I})\mathbf{Y}^T \right) \mathbf{v}. \end{aligned}$$

As $(\mathbf{I} + \mathbf{C}\mathbf{C}^T)^{-1} = \mathbf{M}^{-T}\mathbf{M}^{-1}$, there is no need to assemble the matrix \mathbf{M} to solve $\mathbf{M}^{-1}\mathbf{v}$. A possible preconditioner is then

$$\mathbf{P}_1 = \Delta(\mathbf{S})^{1/2} \prod_{j=1}^{r_n} \left[\mathbf{D}_j^{1/2} \mathbf{M}_j \mathbf{M}_j^T \mathbf{D}_j^{1/2} \right] \Delta(\mathbf{S})^{1/2}$$

but, to preserve symmetry, we prefer the Subspace-by-Subspace preconditioner

$$\mathbf{P}_{SBS} = \Delta(\mathbf{S})^{1/2} \prod_{j=1}^{r_n} \left[\mathbf{D}_j^{1/2} \mathbf{M}_j \right] \prod_{j=r_n}^1 \left[\mathbf{M}_j^T \mathbf{D}_j^{1/2} \right] \Delta(\mathbf{S})^{1/2}, \quad (4.18)$$

which results from a simple reordering of the product in \mathbf{P}_1 .

To find $\mathbf{M}^{-1}\mathbf{v}$ and $\mathbf{M}^{-T}\mathbf{v}$, it is sufficient to compute the \mathbf{Y} part of \mathbf{Q} in the QR factorization of \mathbf{C} . Thus we use a QR factorization based on the modified Gram-Schmidt algorithm (see for instance Golub and Loan, 1989), rather than a factorization based on Givens rotations or Householder transformations.

A simple calculation shows that the QR factorization is less expensive than assembling the total matrix $\mathbf{I} + \mathbf{C}\mathbf{C}^T$ and using a LDL^T factorization so long as k is less than $(\sqrt{5} - 2) \times n_s$.

5 Experiments with matrix stretching

In this section, we investigate whether matrix stretching is an effective means of solving systems of the form (1.1)–(1.2). We first consider an artificial problem, which allows us to investigate isolate different aspects of the technique. We then consider its effect on practical problems.

5.1 Experiments on structured matrices with specified values

We experiment on block diagonal matrices with overlapping between two consecutive blocks. The blocks are constructed using random values so that each block has eigenvalues between 10^{-1} and 10^3 . The percentage of overlap (number of variables shared by two elements over the total number of variables) is set to 10%, 20%, and 30%. Characteristics of the test examples are given in Table 5.1.

	10% overlap			20% overlap			30% overlap		
n_e	10	50	100	10	50	100	10	50	100
n	91	451	901	82	402	802	73	353	703
n_s	9	49	99	18	98	198	27	147	297

Table 5.1: Test problem characteristics. Key: n_e is the number of elements used, and n , the number of variables of the initial problem, n_s the number of variables in the Schur complement.

We compare the following methods:

1. CG on the initial system without preconditioning (**cg none**);
2. the same method with a diagonal preconditioner (**cg diag**);
3. the Schur complement method on the stretched system (**Schur none**); and
4. the Schur complement method with the Chan diagonal preconditioner (**Schur diag**).

The CG and its preconditioners are available with the **PAREBE** package (Daydé, L'Excellent and Gould, 1997c). The stopping criteria for all the methods is chosen so that $\|\mathbf{B}\mathbf{x} - \mathbf{b}\| \leq 10^{-9} \|\mathbf{b}\|$. For our experiments, the number of probing vectors used with the Chan diagonal preconditioner is 10% of the dimension of the Schur complement.

We give our results In Table 5.2. We report the number of iterations, the CPU time, the residual error ($\|\mathbf{B}\mathbf{x} - \mathbf{b}\|_\infty / \|\mathbf{B}\|_\infty \|\mathbf{b}\|_\infty$), and the component-wise error ($\max_{i=1,n} |\hat{\mathbf{x}}_i - \mathbf{x}_i| / |\hat{\mathbf{x}}_i|$ where $\hat{\mathbf{x}}$ is the theoretical solution and \mathbf{x} the computed one).

The time and number of iterations of the CG method increase as the number of elements increases. This is less the case when using the Schur complement method with and without preconditioning. The residual and component-wise errors are slightly better in the case of the CG

10% of overlap					
n_e	Method	#its	Time (s)	Residual	CW error
10	CG NONE	237	1.00	0.8×10^{-10}	0.7×10^{-6}
	Schur NONE	11	0.05	0.2×10^{-9}	0.3×10^{-7}
	CG DIAG	225	0.69	0.1×10^{-9}	0.7×10^{-6}
	Schur Chan	18	0.08	0.1×10^{-6}	0.5×10^{-4}
50	CG NONE	602	9.00	0.6×10^{-9}	0.3×10^{-5}
	Schur NONE	175	2.10	0.1×10^{-5}	0.4×10^{-4}
	CG DIAG	561	8.50	0.4×10^{-9}	0.1×10^{-5}
	Schur Chan	19	0.37	0.8×10^{-7}	0.2×10^{-5}
100	CG NONE	696	21.0	0.5×10^{-9}	0.2×10^{-5}
	Schur NONE	472	11.0	0.3×10^{-6}	0.5×10^{-4}
	CG DIAG	687	21.0	0.4×10^{-9}	0.3×10^{-5}
	Schur Chan	22	0.93	0.5×10^{-6}	0.5×10^{-5}
20% of overlap					
10	CG NONE	215	0.63	0.8×10^{-10}	0.1×10^{-5}
	Schur NONE	24	0.09	0.3×10^{-7}	0.8×10^{-6}
	CG DIAG	188	0.58	0.1×10^{-9}	0.9×10^{-7}
	Schur Chan	30	0.12	0.3×10^{-13}	0.6×10^{-12}
50	CG NONE	530	7.70	0.6×10^{-9}	0.4×10^{-5}
	Schur NONE	582	7.20	0.2×10^{-5}	0.7×10^{-3}
	CG DIAG	481	7.10	0.3×10^{-9}	0.8×10^{-5}
	Schur Chan	42	0.73	0.3×10^{-6}	0.6×10^{-4}
100	CG NONE	602	16.0	0.3×10^{-9}	0.5×10^{-5}
	Schur NONE	1464	36.0	0.5×10^{-6}	0.7×10^{-4}
	CG DIAG	517	15.0	0.4×10^{-9}	0.9×10^{-5}
	Schur Chan	53	1.90	0.2×10^{-6}	0.3×10^{-4}
30% of overlap					
10	CG NONE	179	0.56	0.1×10^{-9}	0.5×10^{-6}
	Schur NONE	50	0.14	0.6×10^{-8}	0.6×10^{-5}
	CG DIAG	153	0.44	0.1×10^{-9}	0.1×10^{-7}
	Schur Chan	32	0.13	0.2×10^{-8}	0.4×10^{-7}
50	CG NONE	428	5.70	0.3×10^{-9}	0.3×10^{-5}
	Schur NONE	1146	16.0	0.2×10^{-5}	0.8×10^{-4}
	CG DIAG	324	4.30	0.3×10^{-8}	0.3×10^{-5}
	Schur Chan	69	1.20	0.4×10^{-6}	0.1×10^{-3}
100	CG NONE	499	13.0	0.4×10^{-9}	0.7×10^{-5}
	Schur NONE	3384	90.0	0.2×10^{-5}	0.3×10^{-3}
	CG DIAG	401	11.0	0.4×10^{-9}	0.4×10^{-5}
	Schur Chan	75	2.80	0.1×10^{-5}	0.1×10^{-3}

Table 5.2: Number of iterations, execution time, residual, and component-wise error of CG applied to the initial system and on the stretched system relatively to the number of elements and the degree of overlap.

method. As the percentage of overlap increases to 30%, the execution times for the preconditioned Schur complement method and PCG applied to the initial system are very similar, while the Schur complement approach without preconditioning performs poorly. As one might expect, as the percentage of overlap increases, the Schur complement approach becomes less and less effective because of its increased size.

5.1.1 Preliminary parallel experiments

In Table 5.3, we study the parallel potential of the stretching method on our artificial examples without preconditioning. We give the execution time in seconds for solving the problem with n_e equal to 10 using one processor (1 Proc) and eight processors (8 Proc) of an Alliant FX/80. The **cg none** method is parallelised using a coloring algorithm (see Daydé et al. (1997c)). The speed-ups achieved using matrix stretching are almost double that achieved by CG on the initial system, and are close to 6 on the 8 processors of the ALLIANT FX/80, a very good result for such a computer. This clearly demonstrates the potential for parallelisation of matrix stretching techniques.

Solver	Percent. Over.	Execution time (s)		Speed Up
		1 Proc	8 Proc	
cg none on initial system	10 %	17.0	5.4	3.2
	20 %	14.0	4.1	3.4
	30 %	11.0	3.1	3.5
Schur none on stretched system	10 %	2.6	0.5	5.7
	20 %	5.8	1.0	5.9
	30 %	8.2	1.4	5.9

Table 5.3: Preliminary parallel experiments on the Alliant FX/80.

5.2 Experiments on matrices from the Harwell-Boeing collection

5.2.1 Description of the test problems

We have also experimented on some of the unassembled matrices available from the Harwell-Boeing collection (see Duff, Grimes and Lewis, 1989, 1992). Since only the structure of these problems is available, the numerical values of the elemental matrices have been set to random values, in the same way as in the previous section.

We have generated, for the three structures **CEGB2802**, **LOCK3491**, and **MAN5976**, four test problems with increasing condition numbers. These problems are termed **CEGB2802-1**, **CEGB2802-2**, **CEGB2802-3**, and **CEGB2802-4** for **CEGB2802**, and similarly for the other problems.

In order to have an efficient Schur complement method, the size of the Schur complement should ideally be much smaller than the size of the initial problem. For this reason, we have used element amalgamation in our test problems. This is achieved with the merge algorithm introduced by Daydé et al. (1997b, 1997c). In this algorithm, elements are amalgamated until a certain threshold is reached. The lower the threshold, the smaller is the Schur complement, as we illustrate in Figure 5.3 for our three classes of test problems. For our experiments, the threshold used by the amalgamation algorithm is -10^5 .

The characteristics of the problems used in our tests are given in Table 5.4.

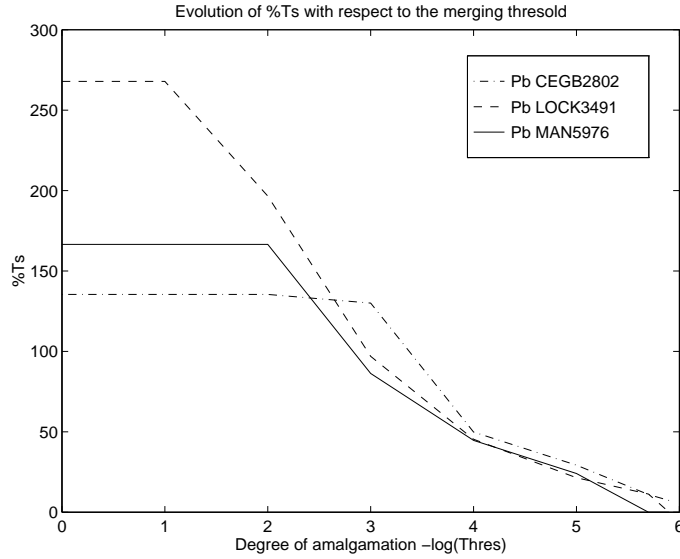


Figure 5.3: Size of the Schur complement with respect to the amalgamation threshold on the test problems (expressed as $100n_s/n$).

Problem	n	n_e	min	max	$aver$	κ	n_s	$\%T_s$	κ_S
CEGB2802-1	2694	12	267	432	290.3	2.1×10^2	789	29.3	9.4×10^1
CEGB2802-2	2694	12	267	432	290.3	1.1×10^4	789	29.3	2.7×10^3
CEGB2802-3	2694	12	267	432	290.3	5.2×10^5	789	29.3	8.6×10^4
CEGB2802-4	2694	12	267	432	290.3	2.3×10^7	789	29.3	3.4×10^6
LOCK3491-1	3416	11	281	468	377.1	3.8×10^2	732	21.4	1.2×10^2
LOCK3491-2	3416	11	281	468	377.1	1.2×10^4	732	21.4	1.2×10^3
LOCK3491-3	3416	11	281	468	377.1	4.3×10^5	732	21.4	5.4×10^4
LOCK3491-4	3416	11	281	468	377.1	1.5×10^7	732	21.4	1.3×10^6
MAN5976-1	5882	22	226	445	331.7	1.7×10^2	1416	24.1	8.0×10^1
MAN5976-2	5882	22	226	445	331.7	8.8×10^3	1416	24.1	1.9×10^3
MAN5976-3	5882	22	226	445	331.7	4.6×10^5	1416	24.1	5.7×10^4
MAN5976-4	5882	22	226	445	331.7	1.8×10^7	1416	24.1	1.9×10^6

Table 5.4: Characteristics of the test problems. Key: n is the number of variables used in the initial problem; n_e is the number of blocks in the initial system; min , max and $aver$ are respectively the minimum, maximum and average size of the blocks in the initial system; κ is the condition number of the problems; n_s is the number of variables used for the Schur complement; $\%T_s$ is the percentage $100 n_s / n$; and κ_S is the condition number of the Schur complement

5.2.2 Numerical experiments

We compare the following solution techniques:

- Solution of the initial system $\mathbf{B}\mathbf{x} = \mathbf{b}$:
 - CG Diag** — PCG with a diagonal preconditioner, and
 - CG EBE** — PCG with the Element-by-Element preconditioner EBE.
- Solution of the stretched system (1.3):
 - NONE** — unpreconditioned CG applied to $\mathbf{S}\boldsymbol{\lambda} = \mathbf{s}$,
 - CHAND** — PCG with the Chan diagonal preconditioner applied to $\mathbf{S}\boldsymbol{\lambda} = \mathbf{s}$,
 - CHANb** — PCG with the CHAN band preconditioner applied to $\mathbf{S}\boldsymbol{\lambda} = \mathbf{s}$,
 - EBE** — PCG with an EBE preconditioner applied to $\mathbf{S}\boldsymbol{\lambda} = \mathbf{s}$,
 - SBS** — PCG with a Subspace-by-Subspace preconditioner applied to $\mathbf{S}\boldsymbol{\lambda} = \mathbf{s}$,
 - Band** — PCG with a band preconditioner applied to $\mathbf{S}\boldsymbol{\lambda} = \mathbf{s}$, and
 - Dir** — a direct solver applied to $\mathbf{S}\boldsymbol{\lambda} = \mathbf{s}$.

When solving the initial system (1.1) by an iterative method, we stop as soon as

$$\|\mathbf{B}\mathbf{x} - \mathbf{b}\| \leq 10^{-9} \|\mathbf{b}\|;$$

when we solve the stretched system (1.3) using Algorithm 3.1, Step 2 is terminated as soon as

$$\|\mathbf{S}\boldsymbol{\lambda} - \mathbf{s}\| \leq 10^{-10} \|\mathbf{s}\|.$$

In our experiments, the number of probing vectors used for the Chan diagonal preconditioner is $0.1n_s$, and the semi-bandwidth of the approximation of the Schur complement for the Chan band preconditioner is $0.2n_s$. For the band preconditioner, we report results using a matrix of semi-bandwidth of $0.2n_s$, while the number of rank-one terms regrouped in the Subspace-by-Subspace preconditioner is also $0.2n_s$. Other values were tried for the various preconditioners but those reported here achieved the best compromise over all the problems considered in our experiments.

In the Tables 5.5, 5.6, and 5.7, we report the results of the tests performed on our test set. T_{total} gives the total time required to solve the linear system, T_{conv} gives the time to solve the system involving the Schur complement, T_{cons} is the construction time for the Schur complement preconditioner and $Iter$ gives the number of iterations required to solve the Schur complement system. The experiments were performed on a SUN workstation with a 125 MHZ HyperSPARC processor.

We make the following observations. Firstly, solving the systems using CG without preconditioning is the least effective method of all. Unfortunately, the Chan diagonal preconditioner also appears to be both rather costly and ineffective for these general matrices. Its banded sister is rather more effective, but is in general even more costly. The Band 20% method performs a comparable number of iterations, but is slightly cheaper than the Chan variant.

Secondly, the EBE preconditioner appears to be reliable and efficient with all our test problems, for both the initial and stretched systems. However, the cost of constructing the preconditioner may be high, the method only pays off overall for the more ill conditioned problems. (Daydé et al. 1997b) observed a similar behaviour under more general circumstances. The main drawback when using this preconditioner on the Schur complement is that its memory requirements are sometimes prohibitive, since it is necessary to compute the unassembled structure of the Schur complement

CEGB2802, Threshold = -1e-5, n = 2694, $n_s = 789$						
System	Preconditioner		CEGB2802-1	CEGB2802-2	CEGB2802-3	CEGB2802-4
Stretched	NONE	Ttotal	17.5	66.4	282.4	609.6
		Tconv	14.3	63.3	279.3	606.5
		Tcons	0.0	0.0	0.0	0.0
		Iter	92	408	1822	3945*
	CHANd	Ttotal	23.0	40.5	91.7	169.2
		Tconv	9.2	26.5	76.8	165.2
		Tcons	10.7	10.8	11.4	10.9
		Iter	59	169	446	969
	CHANb 20%	Ttotal	41.3	42.6	47.0	134.0
		Tconv	0.9	1.6	2.4	92.3
		Tcons	37.1	37.8	40.9	38.5
		Iter	4	7	11	448
	EBE	Ttotal	10.9	13.4	17.5	21.0
		Tconv	2.7	4.8	8.1	12.1
		Tcons	0.5	0.5	0.5	0.5
		Iter	14	24	38	62
	SBS 20%	Ttotal	17.0	20.6	28.5	39.9
		Tconv	5.5	8.9	16.0	28.0
		Tcons	4.8	4.8	5.0	4.9
		Iter	14	24	40	76
	Band 20%	Ttotal	13.2	13.4	14.8	107.6
		Tconv	1.0	1.3	1.6	94.7
		Tcons	5.2	5.3	5.6	5.6
		Iter	4	6	7	451
Dir	Ttotal	24.7	24.9	24.6	23.7	
	Tfact	6.2	6.5	6.3	6.4	
	Tcons	15.3	15.3	15.1	14.2	
	Iter	1	1	1	1	
Initial	CG	Ttotal	5.5	25.3	132.9	668.3
	Diag	Iter	47	234	1245	6150
	CG EBE	Ttotal	6.6	15.7	58.1	241.7
		Iter	13	48	211	915

Table 5.5: A comparison of the preconditioners on variants of CEGB2802.

from the rank-one contribution terms. An extreme example is when one element involves all the variables, as then the EBE method will require more memory than a direct solver.

When this occurs a Subspace-by-Subspace preconditioner may be particularly appropriate since its storage requirements will then be significantly lower. It is very close to EBE in terms of the number of iterations required for convergence, but, for the examples here, rather more expensive to use. We do not report here variations in the number of rank-one terms, n_g , which are grouped together for SBS — n_g is fixed to $0.2n_s$ in table 5.7. Empirically, this value has proved to be generally effective, but for **MAN5976**, the construction time crucially depends on n_g , and a value of $n_g = 0.05n_s$ proves to be better for the three first condition numbers ($Ttotal$ decreases from 178.2 to 34.6 seconds for the problem **MAN5976-1** with $n_g = 0.05n_s$). Since this preconditioner has only recently been proposed, it is clear that further experimentation is needed really to assess its effectiveness. For the cases considered here EBE appears to have an edge on SBS.

Thirdly, it is clear here that the direct factorization of the Schur complement is a very appealing alternative when it is feasible. In our experiments it is often the best option, but EBE proves to be almost as competitive. We would not expect that, for larger problems, the method will not be as attractive, as the Schur complement, while having considerable hidden structure, is quite

LOCK3491, Thres = -1e-5, n = 3416, $n_s = 732$						
System	Preconditioner		LOCK3491-1	LOCK3491-2	LOCK3491-3	LOCK3491-4
Stretched	NONE	Ttotal	27.1	74.0	204.6	582.0
		Tconv	22.7	69.6	200.2	577.1
		Tcons	0.0	0.0	0.0	0.0
		Iter	91	280	806	2105
	CHANd	Ttotal	30.5	46.5	104.3	351.9
		Tconv	11.4	26.2	83.8	331.4
		Tcons	14.6	15.6	15.9	15.7
		Iter	43	97	314	1287
	CHANb 20%	Ttotal	57.5	64.2	98.9	125.0
		Tconv	2.2	4.5	41.1	69.5
		Tcons	50.8	55.1	53.0	51.0
		Iter	7	14	133	232
	EBE	Ttotal	16.1	20.4	24.5	31.6
		Tconv	4.1	7.5	12.0	19.9
		Tcons	0.5	0.5	0.5	0.5
		Iter	14	24	40	65
	SBS 20%	Ttotal	24.3	33.2	40.0	56.5
		Tconv	7.1	14.0	22.5	39.7
		Tcons	7.0	7.9	7.0	6.7
		Iter	14	25	43	80
	Band 20%	Ttotal	18.5	20.9	57.4	86.5
		Tconv	2.2	3.8	40.1	69.0
		Tcons	6.2	6.3	6.4	6.2
		Iter	7	12	133	229
Dir	Ttotal	29.8	29.8	29.9	31.9	
	Tfact	5.0	5.0	5.1	5.4	
	Tcons	20.4	20.4	20.5	22.0	
	Iter	1	1	1	1	
Initial	CG	Ttotal	6.7	26.3	92.4	270.2
	Diag	Iter	40	161	569	1678
	CG	Ttotal	8.7	13.6	23.0	40.8
	Ebe	Iter	11	23	47	92

Table 5.6: A comparison of the preconditioners on variants of LOCK3491.

often dense. The limitation of these “direct” preconditioners lies in the fact that they do not take advantage of the structure of the Schur complement.

Finally, from a numerical point of view, the matrix stretching approach is highly interesting because of the observation we made in Table 5.4 that the Schur complement is, for the examples we considered, always better conditioned than the initial matrix. We noted that by amalgamating more and more of the elements from the problems CEG2802 and LOCK3491, the condition number of the Schur complement also improved. Once the condition number of the initial problem is large enough, matrix stretching appears to be more efficient than directly attacking the initial system.

6 Conclusions

In this paper, we have demonstrated the advantages of using a block stretching method in combination with a Schur complement solution approach. Such a method can be applied on the symmetric unassembled matrices arising from finite element problems or from partially separable optimization. The main benefit of matrix stretching appears to be that on ill-conditioned systems, the condition number of the Schur complement is lower than that of the initial system, and thus

MAN5976, Thres = -1e-5, n = 5882, $n_s = 1416$						
System	Preconditioner		MAN5976-1	MAN5976-2	MAN5976-3	MAN5976-4
Stretched	NONE	Ttotal	38.2	135.6	501.6	1695.8
		Tconv	32.7	130.1	496.1	1690.2
		Tcons	0.0	0.0	0.0	0.0
		Iter	84	337	1287	4368
	CHANd	Ttotal	72.6	121.7	264.5	1177.3
		Tconv	24.5	74.3	217.1	1132.0
		Tcons	42.1	41.5	41.5	39.7
		Iter	59	181	543	2913
	CHANb 20%	Ttotal	174.5	179.7	293.6	480.4
		Tconv	3.7	9.6	125.0	314.0
		Tcons	165.0	164.3	162.7	160.9
		Iter	6	17	231	588
	EBE	Ttotal	25.4	33.3	49.4	80.6
		Tconv	8.3	16.7	32.8	63.9
		Tcons	0.9	0.8	0.9	0.8
		Iter	17	36	73	136
	SBS 20%	Ttotal	178.2	181.4	212.7	279.1
		Tconv	19.7	38.9	80.8	112.2
		Tcons	143.9	127.6	118.0	152.7
		Iter	17	38	83	161
	Band 20%	Ttotal	40.0	42.6	45.2	368.9
		Tconv	3.7	5.9	9.5	331.0
		Tcons	22.2	22.4	21.7	22.8
		Iter	6	10	17	573
Dir	Ttotal	67.6	67.6	68.2	71.6	
	Tfact	31.6	31.7	32.3	34.4	
	Tcons	30.4	30.5	30.4	31.6	
	Iter	1	1	1	1	
Initial	CG Diag	Ttotal	12.2	51.0	196.4	632.0
		Iter	44	187	720	2482
	CG EBE	Ttotal	14.3	28.1	59.8	124.6
		Iter	13	36	86	190

Table 5.7: A comparison of the preconditioners on variants of MAN5976.

less sophisticated preconditioned may be required. The other main benefit of matrix stretching is obviously its potential for parallelisation, since the diagonal blocks within the stretched system can be factorized independently. The challenge is to make sure that the iterative solution step on the Schur complement is also effectively parallelised, and EBE and SBS preconditioners seem ideal in this respect.

Naturally, we do not expect such an approach to be efficient on all kinds of partially separable linear systems, but believe that we have demonstrated that, at least in some cases, there is considerable benefit to be gained from stretching. Amalgamation algorithms such as the ones described in (Daydé et al. 1997b) are of great importance since they may provide a decrease in size of the Schur complement. The most difficult problem is still to identify efficient preconditioners for the Schur complement. The Element-by-Element preconditioners, particularly the Subspace-by-Subspace one, offers some opportunity of taking advantage of the hidden structures within the Schur complement.

Further experiments on real problems from the Harwell-Boeing collection and from the CUTE collection of optimization problems (see Conn et al., 1992 and Bongartz, Conn, Gould and Toint, 1993) are currently being performed.

References

- Alvarado, F. L. (1997), ‘Matrix enlarging methods and their application’, *BIT* **37**(3), 473–505.
- Andersen, K. D. (1996), ‘A modified schur-complement method for handling dense columns in interior-point methods for linear programming’, *ACM TOMS* **22**(3), 348–356.
- Anderson, E., Bai, Z., Bischof, C., Demmel, J., Dongarra, J., Croz, J. D., Greenbaum, A., Hammarling, S., McKenney, A. and Sorensen, D. (1990), *Lapack : A portable linear algebra library for high-performance computers*, Technical Report Report CS-90-105, Computer Science Department, University of Tennessee. Technical Report LAPACK Working Note 20.
- Axelsson, O. (1996), *Iterative Solution Methods*, Cambridge University Press, Cambridge, USA.
- Bongartz, I., Conn, A. R., Gould, N. I. M. and Toint, P. L. (1993), *CUTE: Constrained and Unconstrained Testing Environment*, Technical Report TR/PA/93/10, CERFACS, Toulouse, France.
- Chan, T. and Goovaerts, D. (1988), A note on the efficiency of domain decomposed incomplete factorizations, Technical Report CAM 88–18, Department of Mathematics, UCLA, Los Angeles CA 90024–1555.
- Chan, T. and Mathew, T. (1992), ‘The interface probing technique in domain decomposition’, *SIAM J. Matrix Anal. Applics.* **13**, 212–238.
- Coleman, T. and Moré, J. (1984), ‘Estimation of sparse hessian matrices and graph coloring problems’, *Mathematical Programming* **28**, 243–270.
- Conn, A. R., Gould, N. I. M. and Toint, P. L. (1992), *LANCELOT: a Fortran package for large-scale nonlinear optimization (Release A)*, number 17 in ‘Springer Series in Computational Mathematics’, Springer Verlag, Heidelberg, Berlin, New York.
- Curtis, A. R., Powell, M. J. and Reid, J. K. (1974), ‘On the estimation of sparse jacobian matrices’, *J. Inst. Math. Appl.* **13**, pp. 117–120.
- Daydé, M., Décamps, J. P. and Gould, N. I. M. (1997a), Subspace-by-subspace preconditioners for structured linear systems, Technical report, ENSEEIHT-IRIT, Toulouse, France. to appear.
- Daydé, M. J., L’Excellent, J. Y. and Gould, N. I. M. (1997b), ‘On the use of element-by-element preconditioners to solve large scale partially separable optimization problems’, *SIAM Journal on Scientific Computing* **18**(6), 1767–1787.
- Daydé, M. J., L’Excellent, J. Y. and Gould, N. I. M. (1997c), PAREBE : parallel element-by-element preconditioners for the conjugate gradient algorithm, Technical report, ENSEEIHT-IRIT, Toulouse, France. To appear.
- Duff, I. S., Grimes, R. G. and Lewis, J. G. (1989), ‘Sparse matrix test problems’, *ACM Trans. Math. Softw.* **15**, 1–14.
- Duff, I. S., Grimes, R. G. and Lewis, J. G. (1992), Users’ guide for the Harwell-Boeing sparse matrix collection, Technical Report TR/PA/92/86, CERFACS, Toulouse, France.
- Erhel, J., Traynard, A. and Vidrascu, M. (1991), ‘An element-by-element preconditioned conjugate gradient method implemented on a vector computer’, *Parallel Computing* **17**, 1051–1065.

- Gill, P. E., Murray, W. and Wright, M. H. (1981), *Practical Optimization*, Academic Press, London and New York.
- Golub, G. H. and Loan, C. F. V. (1989), *Matrix Computations*, 2nd edn, The Johns Hopkins University Press, Baltimore, MD.
- Grcar, J. F. (1990), Matrix stretching for linear equations, Technical Report SAND90-8723, Sandia National Laboratories.
- Griewank, A. and Toint, P. L. (1982*a*), ‘Local convergence analysis for partitioned quasi-Newton updates’, *Numerische Mathematik* **39**, 119–137.
- Griewank, A. and Toint, P. L. (1982*b*), ‘Partitioned variable metric updates for large structured optimization problems’, *Numerische Mathematik* **39**, 429–448.
- Haynsworth, E. (1968), ‘On the schur complement’, *Basel Math. Notes* *20*.
- Hughes, T. J. R., Ferencz, R. M. and Hallquits, J. O. (1987), ‘Large-scale vectorized implicit calculations in solid mechanics on a CRAY X-MP/48 utilizing EBE preconditioned conjugate gradients’, *Computational Methods in Applied Mechanics and Engineering* **61**, 215–248.
- Hughes, T. J. R., Levit, I. and Winget, J. (1983), ‘An element-by-element solution algorithm for problems of structural and solid mechanics’, *Computational Methods in Applied Mechanics and Engineering* **36**, 241–254.
- Keyes, D. E. and Gropp, W. D. (1987), ‘A comparison of domain decomposition techniques for elliptic partial differential equations and their parallel implementation’, *SIAM J. Sci. Statist. Comput.* **8**, s166–s202.
- Keyes, D. E. and Gropp, W. D. (1989), Domain decomposition techniques for the parallel solution of nonsymmetric systems of elliptic BVPS, in T. Chan, R. Glowinski, J. Périaux and O. . Widlund, eds, ‘Domain Decomposition Methods’, Society for Industrial and Applied Mathematics.
- Ortiz, M., Pinsky, P. M. and Taylor, R. L. (1983), ‘Unconditionally stable element-by-element algorithms for dynamic problems’, *Computational Methods in Applied Mechanics and Engineering* **36**, 223–239.
- Powell, M. J. D. and Toint, P. L. (1979), ‘On the estimation of sparse hessian matrices’, *SIAM Journal on Numerical Analysis* **16**, 1060–1074.
- Vanderbei, R. J. (1991), ‘Splitting dense columns in sparse linear systems’, *Linear Algebra Appl.* **152**, 107–117.
- Wathen, A. J. (1989), ‘An analysis of some element-by-element techniques’, *Computational Methods in Applied Mechanics and Engineering* **74**, 271–287.
- Zienkiewicz, O. C. (1977), *The Finite Element Method*, McGraw-Hill, London.