INTERNATIONAL JOURNAL FOR NUMERICAL METHODS IN FLUIDS
Int. J. Numer. Meth. Fluids 2000; **32**: 1–27

Iterative solution of panel method discretizations for potential flow problems. The modal multipolar preconditioning

J. D'Elía*, M. Storti and S. Idelsohn

Grupo de Tecnología Mecánica del INTEC, Güemes 3450, (3000) Santa Fe, Argentina

SUMMARY

The iterative solution of linear systems arising from panel method discretization of three-dimensional (3D) exterior potential problems coming mainly from aero-hydrodynamic engineering problems, is discussed. An original preconditioning based on an approximate eigenspace decomposition is proposed, which corrects bad conditioning arising from a pair of surfaces that are very close to each other, which is a very common situation in slender wings and other aerodynamic profiles. This preconditioning has been tested with the standard Bi-conjugate gradient (Bi-CG) and conjugate gradient squared (CGS) iterative methods. Copyright © 2000 John Wiley & Sons, Ltd.

KEY WORDS: panel method; preconditioning; iterative solver; aerodynamics; potential flow; panel clustering

1. INTRODUCTION

The panel method is a well-established numerical technique for the solution of potential flow problems, especially in exterior aero- and hydrodynamics due to its ability to cope with complex geometries and the lack of artificial infinite boundaries [1–4]. This work is based on Morino's formulation [5–8] with plane low-order panels for the potential field. A characteristic of such formulation is that it gives a full non-symmetric matrix with relatively low condition numbers. To be more specific, the condition number does not degrade under refinement as it is common in the '*in volume*' discretization methods, like FDM, FEM, FVM or the relatively young 'meshless' methods. This advantage is somewhat compensated by the fact that the system matrix is full and a definitive assessment of the efficiency of the method, compared with the 'in volume' ones is rather involved. However, a clear advantage of the method, mainly from the practical point of view, is the lack of domain interior meshing, specially in 3D.

* Correspondence to: Grupo de Tecnología Mecánica del INTEC, Güemes 3450, (3000) Santa Fe, Argentina. Tel.: + 54 342 4550944; fax: + 54 342 550944; e-mail: mstorti@intec.unl.edu.ar

CCC 0271–2091/2000/010001–27$17.50

Received July 1997
Revised July 1998

Application of iterative solvers to panel/BEM (boundary element method) problems is described in many papers [9–14]. A rough list of parameter-free iterative solvers for non-symmetric systems of linear equations is given by Natchigal *et al.* [17], where three basic types are considered as follows. First, those methods based on the normal equations: CGN = CGNR (Hestnes and Stiefel [22]), CGNE (Craig, 1955), LSQR (Paige and Saunders [23]). Second, those ones based on orthogonalization: GCG (Concus and Golub, 1976, Widlund, 1978), ORTHOMIN (Vinsome, 1976), ORTHORES and ORTHODIR (Young and Jea, 1990), FOM (Saad, 1981), GCR (Elman, 1982, Eisenstat *et al.* [24]), GMRES (Saad and Schultz [25]). Third, those ones based on biorthogonalization: BCG = BIOMIN (Lanczos [18], Fletcher [19]), BIORES = BO (Lanczos, 1950, Jea and Young, 1983), BIODIR (Jea and Young, 1983), CGS = BIOMIN^2 (Sonneveld [20]), BIORES^2 and BIODIR^2 (Gutknecht [26]), Bi-CGSTAB (van der Vorst [27]), QMR (Freund [28]). We also have the USYMLQ and USYMQR methods (Saunders *et al.* [28]). This terminology approximately follows Nachtigal *et al.* and Gutknecht [26] and other references can be found in [29–31]. Iterative solvers are based on repetitive calculation of matrix–vector multiplication. As the matrices coming from panel discretizations are full, it is not possible to store the matrix coefficients in core memory, as it is usual in the 'in volume' methods, where the matrix is sparse. Then, the interaction coefficients have to be recomputed at each matrix–vector operation and the overall cost is roughly the number of matrix–vector operations times the cost of evaluating one of them. Global efficiency is controlled thus by: (a) a choice of the iterative solver and preconditioning in order to improve the convergence rate and reduce the number of matrix–vector operations, and (b) an efficient computation of the interaction coefficients.

Some preconditioners are purely algebraic, as those based on incomplete factorization (see [9] for instance), whereas others take into account the underlying physics. Most of the physically based preconditionings for the panel method are based on some kind of multi-polar expansion of the field produced by a '*cluster*' of panels [13–15]. The advantage of the algebraic preconditionings are that they can be applied to a broader range of problems, whereas physical based preconditionings yield better performances at the expense of being more specific.

Preconditioners in a general BEM context are extensively reviewed by Prasad *et al.* [10], where some success has been reported with the conjugate gradient and GMRES when they are used in conjunction with preconditioning approaches. Also, Hribersek *et al.* [11] have considered Jacobi, incomplete factorization and row–sum-type preconditioners for the BEM solution of viscous flow problems, showing improved convergence rates with the first two ones. Yan [21] obtained sparse preconditioners for dense system matrix in 2D BEM analysis through condensation by discrete Fourier transforms, whereas Vavasis [15] treats the panel/BEM case, which is rather near to the current case, i.e. solving the Laplace equation in a exterior 3D domain, so a brief account of his approach will be given. Vavasis considers three basic preconditioners: the *mesh neighbor*, the *matrix entries* and the *hierarchical clustering*, which have the following idea in common. From the system matrix $\mathbf{A}$, a small index list L is chosen, drawn from $\{1, 2, \ldots, N_i\}$ such that the variables in L have the most influence on the variable i. Next, a small system of equations $\bar{\mathbf{A}}^T\bar{\mathbf{p}}_i = \bar{\mathbf{e}}_i$ is solved, where the overbars denote that all the rows and the columns of $\mathbf{A}$ are deleted except for those in the index list L. Once this solution is known, it expands back to entries of the preconditioner and this procedure is done for all its rows. On one hand, in the *mesh neighbor* preconditioner, it is taken into account that the

matrix coefficient relating the control points i and j comes up roughly from a term like $1/|\mathbf{x}_i - \mathbf{x}_j|$. Therefore, the further apart two control points are, the less influence it would expect a change at one control point to have on the other. It is said that two control points are 'neighbors' if they border a common panel side. Since neighboring control points are the most interrelated, they are put in the mesh neighbor preconditioner. Also, it is noted that this strategy leads to a sparse preconditioner, where its sparsivity pattern will mirror the mesh connectivity. On the other hand, the mesh entries preconditioner constructs the index list L with the following criterion: if a_{ij}, a_{ji} satisfy $|a_{ij}a_{ji}| \geq t|a_{ii}a_{jj}|$, then the control point j is included in the index list L, where t denotes some user-specified tolerance, or 'magic number'. For instance, when $a_{ii} = 1/2$ and $|a_{ij}| < a_{ii}$, with $t = 1, 0$, a diagonal and full populated preconditioner are obtained respectively, whereas middle values of t do not offer any regular or predictable sparsivity pattern. In an algebraic context, Jemmings [21] proposed to adopt $0.01 \leq t \leq 0.10$. As it can be noted, this type of preconditioner does not depend on any panel/BEM formulation, hence it can be applied to an arbitrary system matrix. Finally, the *hierarchical clustering* preconditioner classifies the control points according to how far away they lie from the control point i. The first step in its construction is to make a hierarchy of clusters, next to the center and radii r for each cluster C are obtained. Once the clusters, centers and radii are computed, the hierarchy preconditioner is constructed with the aid of the index list L obtained as the acceptable clusters of each control point. It is said that cluster C is acceptable to the control point i when the distance from the control point i to the center of C is at least tr, where $t > 1$ is a user-specified number. Vavasis reports good improvements with the first two types for rather thick 3D geometries on industrial applications. Further details about these three types of preconditioners adapted for a dense and unsymmetric matrix system can be found in the Vavasis's work and the reference listed there.

However, for thin wings, the authors found that the performance of these kind of preconditioners is rather restricted. On one hand, the *mesh neighbor* preconditioner assumes that the importance of the influence matrix coefficient is only related with the mesh topology, i.e. neighbor panels have a strong interaction and remote ones a weaker interaction. However, this assumption is not always the case for *dipolar* matrices, which are proportional also to a view factor, so that neighbor panels (in a topological sense) that are nearly coplanar have weak interaction, which is actually the case of well-refined meshes on smooth surfaces. Moreover, remote panels (again in a topology sense) on opposite sides of thin wing geometries and facing each other, have a strong interaction. In other words, the mesh neighbor preconditioner strategy conduces to rather lower performances in such cases. On the other hand, the *matrix entries* preconditioner is a good alternative for thick geometries, with usually low rate in the filling of the preconditioner (i.e. its non-zero entries). But, for thin wing configurations (similar to the considered in the numerical examples) we had to use $t > 0.8$, with filling ratios greater than 80%. Finally, the *hierarchical clustering* preconditioner shares similar characteristics with the *mesh neighbor* one.

For most iterative solvers, the convergence rate is closely related to the condition number of the linear system, which highly depends on the geometry. To be more specific, high condition numbers arise whenever two surfaces are close together, so that the distance between them is smaller than the average size of the panels. Note that this implies that, in contrast with the 'in volume' methods, the condition number is reduced after refinement, since, eventually the

average size of the panels will get smaller than the distance between the surfaces. Now, suppose for a moment that the profile is symmetric and with a small thickness ϵ, and we are interested in $\epsilon \to 0$. Any distribution of double layer density μ given by $\mu_{1,2}$ on each side of the wing can be decomposed as the sum of a symmetric one μ^+ and a skew-symmetric one μ^- where $\mu^{\pm} = 1/2(\mu_1 \pm \mu_2)$. Now, in the limit of vanishing thickness, the field produced by a skew-symmetric density on both sides of the wing are added and the result is a flat surface at the plane of symmetry with a distribution of $2\mu^-$, whereas for the symmetric distribution, the fields tend to cancel each other and the result is also a flat surface but with a null double layer density distribution. This shows that the self-interaction coefficient for the skew-symmetric distribution behaves like $O(1)$, whereas those for the symmetric one behaves like $O(\epsilon)$ for $\epsilon \to 0$. This explains why the condition number degrades as $\epsilon \to 0$ and suggests that a physically based preconditioning, based on making the change of variables and scaling appropriately the symmetric part, will correct this degradation. The authors call this the 'modal preconditioning' and is, in their opinion, the main contribution of this paper.

One may argue that if the preconditioning is efficient for profiles that are too thin, then it would be better to handle those case specially, i.e. modeling them as plain flat surfaces. First, it will be shown in the examples that interesting gains are obtained for profiles of 8% and even 25% relative thickness. Second, even if the airfoil is not too thin, the thickness is smaller near the trailing edge and this is a cause of slower convergence. Third and last, using the preconditioning allows the solving of arbitrarily thin airfoils without worrying about the degree of the approximation of replacement by a zero thickness airfoil.

With respect to an efficient evaluation of the interaction coefficients, it is common practice to use approximated expressions based on far-field expansions, valid when both panels are separated by a distance that exceeds some threshold value scaled by the panel size. Typically, the cost of the far-field expression is faster than the exact evaluation by a ratio of 1:8. However, this introduces a consistency error, and in order to eliminate it, an outer loop is iterated, where a residual with the exact coefficients is computed and a correction is added. It is clear that if too many iterations of the outer loop (≥ 8) are performed, then it is cheaper to iterate directly on the exact coefficients, but it can be shown that this exterior problem is well-conditioned and, typically, two iterations are needed to reduce the error by a factor of 10^{-7}. This issue is discussed in depth elsewhere [16].

2. PANEL DISCRETIZATION OVERVIEW

Let Γ be a closed surface, and $\Omega^{i(e)}$ the corresponding interior (exterior) domain. The governing equations for potential flow are the Laplace equation in Ω^e with slip condition on Γ, which in terms of the perturbation potential can be written as:

$$\Delta\phi = 0, \quad \text{in } \Omega^e;$$

$$\frac{\partial\phi}{\partial n} = -\mathbf{u}_\infty \cdot \hat{\mathbf{n}}, \quad \text{at } \Gamma, \tag{1a,b}$$

where $\mathbf{u}_\infty$ is the undisturbed velocity and ϕ is the perturbation potential defined by:

$$\mathbf{u} = \mathbf{u}_\infty + \nabla\phi, \tag{2}$$

where $\mathbf{u}$ is the total velocity. This problem can be rewritten as a Fredholm integral equation in terms of the single and double layer densities σ and μ as:

$$\mu(\mathbf{x}) - \frac{1}{2\pi}\int_\Gamma \mu(\mathbf{x}')\frac{\partial}{\partial n}\left(\frac{2}{|\mathbf{x}-\mathbf{x}'|}\right)\mathrm{d}\Gamma' = \frac{1}{2\pi}\int_\Gamma \sigma(\mathbf{x}')\frac{1}{|\mathbf{x}-\mathbf{x}'|}\,\mathrm{d}\Gamma', \tag{3}$$

for $\mathbf{x}$ belonging to the surface Γ, and $\hat{\mathbf{n}}$ is the normal pointing into Ω^e. For the slip or inlet/outlet boundary condition as in (1b), it turns out to be that σ is simply $\sigma = -\hat{\mathbf{n}}\cdot\mathbf{u}_\infty$ and then, the right-hand-side of (3) is known. Moreover, μ is equal to the perturbation potential ϕ at the surface. The panel method is based on approximating Γ by a polyhedral surface, composed of a certain number N_{pan} of non-overlapping flat panels $\{\Gamma_i\}_{i=1}^{N_{\text{pan}}}$. Assuming that μ and σ are constant over each panel and imposing (3) by collocation at the centroids of the panels, a linear system of the form:

$$\mathbf{A}\boldsymbol{\mu} = \mathbf{C}\boldsymbol{\sigma} \tag{4}$$

is obtained, where $\boldsymbol{\mu} = [\mu_1\mu_2\ldots]^T$ is the vector of panel potentials, $\boldsymbol{\sigma}$ is the single layer densities per panel, computed as:

$$\sigma_i = -\mathbf{u}_\infty\cdot\hat{\mathbf{n}}_i, \tag{5}$$

and $\hat{\mathbf{n}}_i$ is the normal to panel i. The interaction coefficients are computed from:

$$\begin{aligned} C_{ij} &= \int_{\Gamma_j}\frac{1}{|\mathbf{x}-\mathbf{x}_i|}\,\mathrm{d}\Gamma, \\ A_{ij} &= \int_{\Gamma_j}\frac{\partial}{\partial n}\left\{\frac{1}{|\mathbf{x}-\mathbf{x}_i|}\right\}\mathrm{d}\Gamma, \end{aligned} \tag{6}$$

where $\mathbf{x}_i$ is the centroid of the i panel.

3. THE MODAL PRECONDITIONING

It will be shown that for thin airfoils, a symmetric/skew-symmetric decomposition reveals a conditioning that grows inversely proportional to the thickness. Then, a preconditioning based on this symmetric/skew-symmetric decomposition is proposed. It will be shown how the preconditioning corrects the bad scaling for a simple geometry consisting of two facing panels,

then for two pairs of panels on a symmetric airfoil, and two pairs of panels on a non-symmetric airfoil. Finally, the expression for the preconditioning in the general case is shown.

4. A PAIR OF PANELS

Consider first two identical flat square panels of side H, parallel to each other as in Figure 1, separated by a distance $d = H\epsilon$. This should be regarded as a very crude idealization of a slender airfoil, so that the normals are taken as shown in the figure. The matrix of interaction coefficients is:

$$\mathbf{A} = \begin{bmatrix} \frac{1}{2} & -g(\epsilon) \\ -g(\epsilon) & \frac{1}{2} \end{bmatrix}, \tag{7}$$

where the authors replaced $A_{11} = A_{22} = \frac{1}{2}$, as is usual for the self-interaction coefficients, and by symmetry $A_{21} = A_{12} = -g(\epsilon)$, where $g(z)$ is the potential produced by a unit distribution of double layer potential at a distance z on an axis perpendicular to the panel and passing through its centroid, z is taken as positive if the point is on the side with positive charge. The qualitative behavior for $g(z)$ is like that in Figure 2, it is antisymmetric with respect to $z = 0$, it has a jump of unit magnitude (i.e. equal to the double layer density) and continuous derivative at $z = 0$:

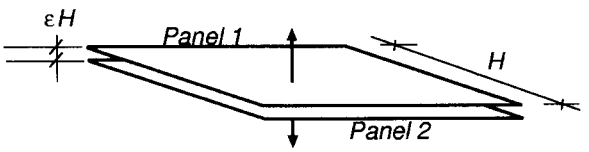

Figure 1. Crude realization of a thin airfoil with two panels.

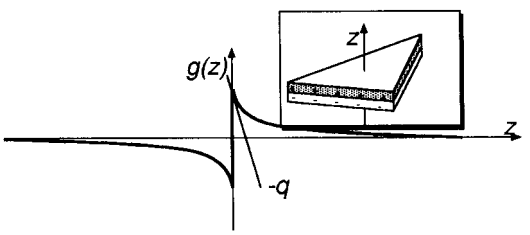

Figure 2. Potential across the panel produced by constant double layer density.

$$
\begin{aligned}
&g\big|_{z=0+} = +\tfrac{1}{2},\\
&g\big|_{z=0-} = -\tfrac{1}{2},\\
&\left.\frac{\mathrm{d}g}{\mathrm{d}z}\right|_{z=0+} = \left.\frac{\mathrm{d}g}{\mathrm{d}z}\right|_{z=0-} = -q.
\end{aligned}
\tag{8}
$$

Then, for small ϵ, you can put, to first-order:

$$
A = \begin{bmatrix} \frac{1}{2} & -\frac{1}{2}+q\epsilon \\ -\frac{1}{2}+q\epsilon & \frac{1}{2} \end{bmatrix}. \tag{9}
$$

As the problem is symmetric about the centerplane, it decouples in a symmetric/antisymmetric basis. Let

$$
\mathbf{S} = \frac{1}{\sqrt{2}}\begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix} \tag{10}
$$

be the change of basis matrix, then the transformed matrix is a diagonal one:

$$
\tilde{\mathbf{A}} = \mathbf{S}^{-1}\mathbf{A}\mathbf{S} = \begin{bmatrix} q\epsilon & 0 \\ 0 & 1-q\epsilon \end{bmatrix} \tag{11a,b}
$$

and the bad conditioning is clear, since for $\epsilon \to 0$, the first diagonal entry vanishes:

$$
\mathrm{cond}(\mathbf{A}) = \mathrm{cond}(\tilde{\mathbf{A}}) = \frac{1-q\epsilon}{q\epsilon} \sim \frac{1}{q\epsilon} \quad \text{for large } \epsilon. \tag{12}
$$

Let us take a closer look at what matrix $\mathbf{S}$ represents. If we take its columns as double layer density distributions, then the first column corresponds to $\mu_{1,2} = 1$, and the second one to $\mu_1 = -\mu_2 = 1$ (see Figure 3). Recall that a double layer density distribution can be thought of as two single layer distributions of equal strength but opposite sign in the limit when the distance between the single layer distributions tends to zero, keeping the product of the

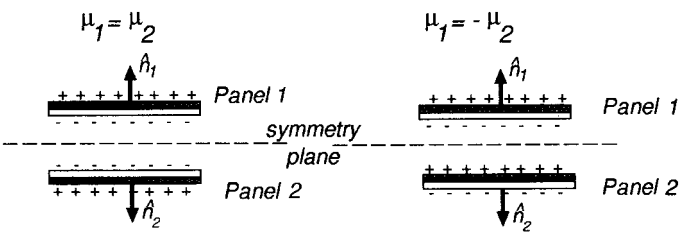

Figure 3. Symmetric and skew-symmetric distributions of charge.

distance and the strength of the single layer distributions constant. By convention, the normal vector points from the negative charge side to the positive one for a positive double layer distribution. Note that this distribution corresponds to a symmetric arrangement of charge about the symmetry plane, whereas the second column, corresponds to a skew-symmetric one. The coefficients in matrix $\tilde{\mathbf{A}}$ are the interaction coefficient between these two distributions. It is clear from symmetry arguments that the symmetric arrangement has a null interaction with the skew-symmetric one, and this explains why the off-diagonal entries in $\tilde{\mathbf{A}}$ are null. Now, regarding the diagonal entries, for $\epsilon \to 0$, the double layer distribution of panel 1 tends to cancel that one of panel 2, so that the field vanishes to zero. The first diagonal entry in $\tilde{\mathbf{A}}$, see Equation (11b), represents the self-interaction coefficient of this arrangement, and then vanishes for $\epsilon \to 0$. In contrast, in the skew-symmetric arrangement, the field of each panel tends to reinforce that of the other, and in the limit, a single panel with the original distribution that is twice as much is obtained. The second diagonal entry corresponds to the self-interaction coefficient for this charge distribution and then it approaches a non-null value for $\epsilon \to 0$. As for this very simple ($N_{\text{pan}} = 2$) case $\mathbf{S}^{-1}\mathbf{AS}$ is diagonal, then:

$$\mathbf{A} = \mathbf{S}\,\text{diag}(\mathbf{S}^{-1}\mathbf{AS})\mathbf{S}^{-1}, \tag{13}$$

where diag($\mathbf{X}$) stands for a diagonal matrix with the same diagonal entries of $\mathbf{X}$:

$$\text{diag}(\mathbf{X}) = \begin{bmatrix} X_{11} & 0 & \dots & & \\ 0 & X_{22} & 0 & \dots & \\ \vdots & 0 & X_{33} & 0 & \dots \\ & \vdots & \vdots & \vdots & \end{bmatrix}. \tag{14}$$

It is well-known that the best preconditioning $\mathbf{Q}$ is that one that most resembles $\mathbf{A}$ for a given computational effort in inverting a linear system for the preconditioning matrix. In this simple case, we can take $\mathbf{Q} = \mathbf{A}$, and from (13):

$$\mathbf{Q} = \mathbf{S}\,\text{diag}(\mathbf{S}^{-1}\mathbf{AS})\mathbf{S}^{-1}, \tag{15}$$

and this preconditioning will give a preconditioned matrix that is the identity, and then is optimal.

In Section 5, this preconditioning will be extended to the case with a large number of panels and it will be shown that it has good preconditioning properties. Note that the computational effort in solving a system for $\mathbf{Q}$ as in (13) is negligible, since the change of basis is performed by $N_{\text{pan}}/2$ sums and differences, and the inversion of the diagonal part of the transformed matrix involves $O(N_{\text{pan}})$ operations.

5. SYMMETRIC AIRFOIL

Consider now the case of a symmetric airfoil discretized with a large number of panels. It is easy to see that it suffices to consider the interaction coefficients between two pair of panels as shown in Figure 4. As the airfoil is symmetric, panels 2 and 4 are obtained from panels 1 and 3 by reflection about the centerplane of the airfoil. For simplicity, it will be assumed also that all the panels are parallel to the symmetry plane, and that panels 1 and 3 are identical (and then 2 and 4), but these assumptions are not essential. The structure of the matrix is then:

$$\mathbf{A} = \begin{bmatrix} \frac{1}{2} & -\frac{1}{2}+q\epsilon & 0 & -p\epsilon \\ -\frac{1}{2}+q\epsilon & \frac{1}{2} & -p\epsilon & 0 \\ 0 & -p\epsilon & \frac{1}{2} & -\frac{1}{2}+q\epsilon \\ -p\epsilon & 0 & -\frac{1}{2}+q\epsilon & \frac{1}{2} \end{bmatrix}. \tag{16}$$

Of course, the 2×2 diagonal matrix blocks are the same as in the 2-panel example. Regarding the off-diagonal terms, the in-plane coefficients like A_{13}, A_{31}, A_{24} and A_{42} are null, which can be easily shown from (6), whereas the others, like A_{14}, are $O(\epsilon)$ and negative, say $-p\epsilon$. Now, the change of basis matrix is written as:

$$\mathbf{S} = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 0 & 1 & 0 \\ 1 & 0 & -1 & 0 \\ 0 & 1 & 0 & 1 \\ 0 & 1 & 0 & -1 \end{bmatrix}. \tag{17}$$

Note that, again, the first two columns correspond to symmetric (with respect to the horizontal plane of symmetry) distribution of potentials, whereas the last two correspond to skew-symmetric ones. The transformed matrix is

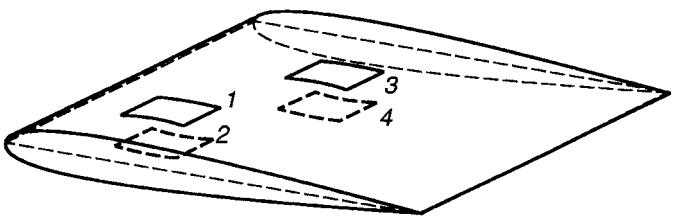

Figure 4. Two pairs of panels on a symmetric airfoil.

$$\tilde{\mathbf{A}} = \begin{bmatrix} q\epsilon & -p\epsilon & 0 & 0 \\ -p\epsilon & q\epsilon & 0 & 0 \\ 0 & 0 & 1-q\epsilon & p\epsilon \\ 0 & 0 & p\epsilon & 1-q\epsilon \end{bmatrix} \tag{18}$$

and it is verified that the off-diagonal 2×2 blocks are null. In addition, the first block diagonal entry corresponding to interaction coefficients between the symmetric distributions has terms $O(\epsilon)$, whereas the second block diagonal entry corresponding to the skew-symmetric mode is $O(1)$. The condition number is again $O(1/\epsilon)$ and can be corrected with the preconditioning defined by (15). Effectively,

$$\begin{aligned} \mathrm{cond}(\mathbf{A}\mathbf{Q}^{-1}) &= \mathrm{cond}(\mathbf{A}\mathbf{S}\,\mathrm{diag}(\mathbf{S}^{-1}\mathbf{A}\mathbf{S})^{-1}\mathbf{S}^{-1}) = \mathrm{cond}(\tilde{\mathbf{A}}\,\mathrm{diag}(\tilde{\mathbf{A}})^{-1}) \\ &= \begin{bmatrix} 1 & -p/q & 0 & 0 \\ -p/q & 1 & 0 & 0 \\ 0 & 0 & 1 & p\epsilon/(1-q\epsilon) \\ 0 & 0 & p\epsilon/(1-q\epsilon) & 1 \end{bmatrix} = O(1) \quad \text{for} \quad \epsilon \to 0. \end{aligned} \tag{19}$$

6. NON-SYMMETRIC AIRFOIL

In this case, the change of basis does not decouple the problem as cleanly as in the symmetric case, but it can still be shown how the proposed preconditioner gives an $O(1)$ condition number, which is verified afterwards with numerical examples (see Figure 5). As before, only two pairs of panels will be considered. For convenience, the following block decomposition of $\tilde{\mathbf{A}}$ is made:

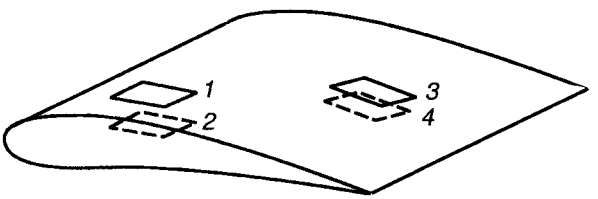

Figure 5. Two pairs of panels on a non-symmetric airfoil.

$$\tilde{\mathbf{A}} = \begin{bmatrix} \tilde{\mathbf{A}}_{++} & \tilde{\mathbf{A}}_{+-} \\ \tilde{\mathbf{A}}_{-+} & \tilde{\mathbf{A}}_{--} \end{bmatrix}, \tag{20}$$

where each of the submatrices are 2×2, and the order will be assessed with respect to ϵ of each of the elements in each of the blocks. This is done in Appendix A and the conclusions are:

$$\tilde{\mathbf{A}}_{++} = \epsilon \tilde{\mathbf{A}}^*_{++} + O(\epsilon^2), \tag{21}$$

$$\tilde{\mathbf{A}}_{-+} = \epsilon^2 \tilde{\mathbf{A}}^*_{-+} + O(\epsilon^3), \tag{22}$$

$$\tilde{\mathbf{A}}_{+-} = \tilde{\mathbf{A}}^*_{+-} + O(\epsilon), \tag{23}$$

$$\tilde{\mathbf{A}}_{--} = \tilde{\mathbf{A}}^*_{--} + O(\epsilon), \tag{24}$$

where $\tilde{\mathbf{A}}^*_{\cdots} = O(1)$.

7. EFFICIENCY OF THE PRECONDITIONING

From (21)–(24), it can be written:

$$\tilde{\mathbf{A}} \sim \begin{bmatrix} \epsilon \mathbf{A}^*_{++} & \tilde{\mathbf{A}}^*_{+-} \\ \epsilon^2 \tilde{\mathbf{A}}^*_{-+} & \tilde{\mathbf{A}}^*_{--} \end{bmatrix}. \tag{25}$$

It is clear from this expression that $\mathrm{cond}(\mathbf{A}) \sim 1/\epsilon$ (or worse). In contrast, if it is preconditioned with (15):

$$\begin{aligned} \mathbf{A}\mathbf{Q}^{-1} &= \mathbf{A}[\mathbf{S}\,\mathrm{diag}(\mathbf{S}^{-1}\mathbf{A}\mathbf{S})\mathbf{S}^{-1}]^{-1} = \mathbf{A}\mathbf{S}\,\mathrm{diag}(\mathbf{S}^{-1}\mathbf{A}\mathbf{S})^{-1}\mathbf{S}^{-1} = \mathbf{S}\tilde{\mathbf{A}}\,\mathrm{diag}(\tilde{\mathbf{A}})^{-1}\mathbf{S}^{-1} \\ &\sim \mathbf{S}\begin{bmatrix} \epsilon\tilde{\mathbf{A}}^*_{++} & \tilde{\mathbf{A}}^*_{+-} \\ \epsilon^2\tilde{\mathbf{A}}^*_{-+} & \tilde{\mathbf{A}}^*_{--} \end{bmatrix}\begin{bmatrix} \epsilon\,\mathrm{diag}(\tilde{\mathbf{A}}^*_{++}) & 0 \\ 0 & \mathrm{diag}(\tilde{\mathbf{A}}^*_{--}) \end{bmatrix}^{-1}\mathbf{S}^{-1} \\ &\sim \mathbf{S}\begin{bmatrix} \tilde{\mathbf{A}}^*_{++}\,\mathrm{diag}(\tilde{\mathbf{A}}^*_{++})^{-1} & \tilde{\mathbf{A}}^*_{+-}\,\mathrm{diag}(\tilde{\mathbf{A}}^*_{--})^{-1} \\ \epsilon\tilde{\mathbf{A}}^*_{-+}\,\mathrm{diag}(\tilde{\mathbf{A}}^*_{++})^{-1} & \tilde{\mathbf{A}}^*_{--}\,\mathrm{diag}(\tilde{\mathbf{A}}^*_{--})^{-1} \end{bmatrix}\mathbf{S}^{-1} \\ &\sim \mathbf{S}\begin{bmatrix} \tilde{\mathbf{A}}^*_{++}\,\mathrm{diag}(\tilde{\mathbf{A}}^*_{++}) & \tilde{\mathbf{A}}^*_{+-}\,\mathrm{diag}(\tilde{\mathbf{A}}^*_{--})^{-1} \\ 0 & \tilde{\mathbf{A}}^*_{--}\,\mathrm{diag}(\tilde{\mathbf{A}}^*_{--})^{-1} \end{bmatrix}\mathbf{S}^{-1}, \end{aligned} \tag{26a–d}$$

and then

$$\mathrm{cond}(\mathbf{A}\mathbf{Q}^{-1}) = O(1) \tag{27}$$

provided that the argument matrix in (26b) is non-singular. Again, the bad conditioning is caused by the symmetric modes. Due to the tendency to cancel the field of the facing panel when $\epsilon \to 0$, the field produced by these modes is $O(\epsilon)$ and so are the interaction coefficients (the first matrix column in (25)). The diagonal preconditioning (in the transformed basis) successfully corrects this behavior.

8. EXPLICIT EXPRESSION FOR REAL ($N_{\text{pan}} > 2$) MESHES

Now, the explicit expression for the change of basis matrix for the case where $N_{\text{pan}} > 4$ is given. It is assumed that the panels are numbered so that panel $2n-1$ and $2n$ are on opposite sides of the airfoil and they collapse to each other for $\epsilon \to 0$. The change of basis matrix is formed by putting first all the $N_{\text{pan}}/2$ symmetric modes and after the skew-symmetric ones:

$$\mathbf{S} = \frac{1}{\sqrt{2}} \left[\begin{bmatrix} 1 \\ 1 \end{bmatrix} \otimes \mathbf{I} \quad \begin{bmatrix} 1 \\ -1 \end{bmatrix} \otimes \mathbf{I} \right], \tag{28}$$

where $\mathbf{I}$ stands for the identity matrix of $(N_{\text{pan}}/2) \times (N_{\text{pan}}/2)$ and the Kronecker product $\otimes$ of two matrices is defined by

$$\mathbf{A} \otimes \mathbf{B} = \begin{bmatrix} B_{11}\mathbf{A} & B_{12}\mathbf{A} & B_{13}\mathbf{A} & \dots \\ B_{21}\mathbf{A} & B_{22}\mathbf{A} & B_{23}\mathbf{A} & \dots \\ \vdots & \vdots & \vdots & \ddots \end{bmatrix}. \tag{29}$$

With respect to the limits of applicability, the mesh on both surfaces have to be congruent, i.e. for $\epsilon \to 0$, the nodes and elements have to coincide with each other and the relative shift has to be normal to the surface. Small deviations or distortions are acceptable provided that they are small with respect to the average element size (see Figure 6).

9. MULTIPOLAR PRECONDITIONING

This preconditioning somewhat follows the lines of the modal one previously presented, but it is oriented towards a single surface instead. It also has some resemblance to multigrid methods. It will be seen in the next section how to combine it with the modal preconditioning. Consider, for instance, two pairs of panels as shown in Figure 7 and start with the same change of basis as before Equation (17). Assume that both panels are coplanar and have the same area. This is true or nearly true for highly structured meshes. It will be discussed later

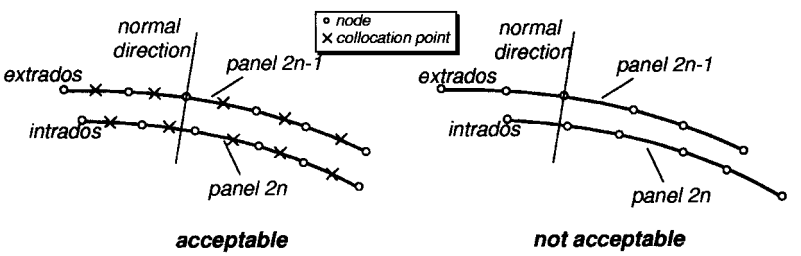

Figure 6. Effect of small distortions in the efficiency of the preconditioned.

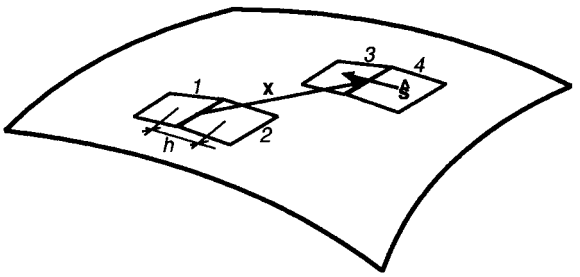

Figure 7. Two pairs of panels on a surface.

how to extend the multipolar preconditioning for unstructured ones. The same block split as in (20) is made and each block will be analyzed at a time. Consider an off-diagonal term in $\tilde{\mathbf{A}}_{\_\ \_}$ like

$$\tilde{A}_{34} = \tfrac{1}{2}(A_{31} - A_{32} - A_{41} + A_{42}). \tag{30}$$

As before, $A_{31} - A_{32}$ is the potential produced by a double layer density distribution of $\mu = +1$ on panel 1 and $\mu = -1$ on panel 2. Suppose now that panels 3 and 4 are far from panels 1 and 2, i.e.

$$|\mathbf{x}| = |\mathbf{x}_{34} - \mathbf{x}_{12}| \gg h, \tag{31}$$

where $\mathbf{x}_{34}$ is the centroid of the '*panel cluster*' composed of panels 3 and 4 and so on. Then, one can approximate $A_{31} - A_{32}$ by a quadrupole expansion:

$$A_{31} - A_{32} \approx -3ah\frac{xz}{r^5} = \phi_{12}(x, y, z), \tag{32}$$

where a is the area of the panels, the system xyz is chosen as in Figure 8, and $r^2 = x^2 + y^2 + z^2$. Now, let $\hat{\mathbf{s}}$ be a unit vector going from the centroid of panel 3 to that of panel 4. Then,

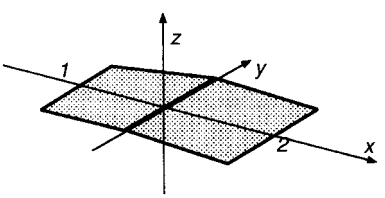

Figure 8. Local system at a panel cluster.

$$\tilde{A}_{34} = -\tfrac{1}{2}[\phi_{12}(\mathbf{x}_{34} + \tfrac{1}{2}h\hat{\mathbf{s}}) + \phi_{12}(\mathbf{x}_{34} - \tfrac{1}{2}h\hat{\mathbf{s}})] \approx \tfrac{1}{2}h(\hat{\mathbf{s}}\cdot\nabla\phi_{12}) = O(h^2/r^4). \tag{33}$$

Note that this is $O[(h/r)^2]$, smaller than the typical interaction coefficient between two panels, which has typically a decay rate of a dipole $O(1/r^2)$. In contrast, the diagonal elements are $O(1)$. In a similar way, it can be shown that the off-diagonal elements in $\tilde{\mathbf{A}}_{-+}$ and $\tilde{\mathbf{A}}_{+-}$ are $O(h/r^3)$. The diagonal elements vanish by symmetry. Finally, the elements in $\tilde{\mathbf{A}}_{++}$ are basically the same as if each pair of panels were added in a single larger panel, with two points of collocation, i.e. it has the same structure of the original matrix $\mathbf{A}$ but it has a dimension that is smaller by half. In brief,

$$\tilde{\mathbf{A}} = \begin{bmatrix} O(\mathbf{I}) + O(1/r^2) & O(h/r^3) \\ O(h/r^3) & O(\mathbf{I}) + O(h^2/r^4) \end{bmatrix}, \tag{34}$$

where $O(\mathbf{I}) + O(1/r^2)$ means that the diagonal elements are $O(1)$, whereas the other off-diagonal terms are $O(1/r^2)$. The preconditioning proposed is obtained retaining only the $\tilde{\mathbf{A}}_{++}$ block and the diagonal part of $\tilde{\mathbf{A}}_{--}$:

$$\mathbf{Q} = \mathbf{S}\begin{bmatrix} \tilde{\mathbf{A}}_{++} & 0 \\ 0 & \mathrm{diag}(\tilde{\mathbf{A}}_{--}) \end{bmatrix}\mathbf{S}^{-1}. \tag{35}$$

The computational effort in inverting this conditioning corresponds to inverting the $\tilde{\mathbf{A}}_{++}$ block. As this has half the dimension of the full matrix, the core memory requirement is one quarter smaller and the CPU time one eighth smaller than those for the full matrix, which is a significant saving. However, panel applications are limited strongly by the matrix size, and then as we are interested in preconditionings requiring smaller amounts of memory at the expense of higher CPU times, this is achieved using larger clusters. Consider a cluster of four panels (two such clusters are shown in Figure 9). The change of basis considered is now

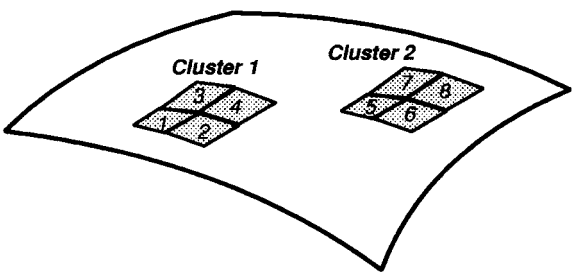

Figure 9. Two clusters of four panels each.

$$\mathbf{S}=\frac{1}{2}\begin{bmatrix} 1 & 0 & 1 & 1 & 1 & 0 & 0 & 0 \\ 1 & 0 & -1 & 1 & -1 & 0 & 0 & 0 \\ 1 & 0 & 1 & -1 & -1 & 0 & 0 & 0 \\ 1 & 0 & -1 & -1 & 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 1 & 1 & 1 \\ 0 & 1 & 0 & 0 & 0 & -1 & 1 & -1 \\ 0 & 1 & 0 & 0 & 0 & 1 & -1 & -1 \\ 0 & 1 & 0 & 0 & 0 & -1 & -1 & 1 \end{bmatrix}. \tag{36}$$

Columns 1, 3–5 correspond to double layer distributions on cluster 1, whereas the others correspond to cluster 2. Column 1 represents a constant distribution of double layer density, so its far-field expansion is a dipole, whereas columns 3–5 have a null total sum of double layer density and their far-field expansions are, therefore, at least of the order of a quadrupole, see Figure 10. The same occurs respectively for columns 2, 6–8 for cluster 2. If the first two columns are called the '$+$ *modes*' and the rest the '$-$ *modes*', a block split in the transformed matrix is induced. The matrices $\tilde{\mathbf{A}}_{++}$, $\tilde{\mathbf{A}}_{-+}$, $\tilde{\mathbf{A}}_{+-}$ and $\tilde{\mathbf{A}}_{--}$ have respectively, sizes of 2×2, 2×6, 6×2 and 6×6. The behavior with respect to h/r are as before or higher: for instance, $\tilde{A}_{58}$ corresponds to the interaction of an octupole on cluster 2 double differentiated at the center of cluster 1. This is then $O(h^4/r^6)$. The size of the $\tilde{\mathbf{A}}_{++}$ block in the preconditioning given by (35) is now 2×2, one quarter the size of $\mathbf{A}$.

In this way, by recursion, larger and larger clusters can be used and the dimension of the matrix to be inverted is $O(N_{\text{pan}}/m)$, where $m=2^n$ is the number of panels in a cluster. Consider a cluster of 2^n panels, then the '*in-cluster*' change of basis matrix $\mathbf{S}_{\text{cluster}}$ is defined recursively as

$$\mathbf{S}_{\text{cluster}}(2^n)=\mathbf{S}_{\text{cluster}}(2)\otimes\mathbf{S}_{\text{cluster}}(2^{n-1}),$$

$$\mathbf{S}_{\text{cluster}}(2)=\frac{1}{\sqrt{2}}\begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix}. \tag{37}$$

$\mathbf{S}_{\text{cluster}}$ is split in its first column, representing a constant distribution of double layer density and all the rest representing higher-order distributions,

$$\mathbf{S}_{\text{cluster}}=[\mathbf{S}^{+}_{\text{cluster}}\quad \mathbf{S}^{-}_{\text{cluster}}], \tag{38}$$

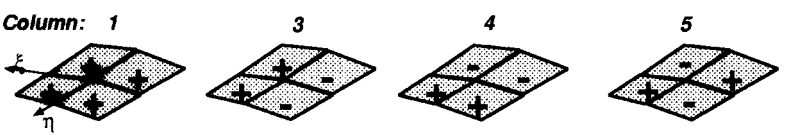

Figure 10. Distribution of double layer densities on the four panel cluster.

where

$$\mathbf{S}^{+}_{\text{cluster}} = 2^{-(n-1)}\begin{bmatrix}1 & 1 & \cdots & 1\end{bmatrix}^T \tag{39}$$

and $\mathbf{S}^{-}_{\text{cluster}}$ is a matrix of $2^n \times (2^n - 1)$ with the rest of the columns. The change of basis matrix is then

$$\mathbf{S}_{\text{mp}} = \begin{bmatrix}\mathbf{S}^{+}_{\text{cluster}} \otimes \mathbf{I} & \mathbf{S}^{-}_{\text{cluster}} \otimes \mathbf{I}\end{bmatrix}, \tag{40}$$

where $\mathbf{I}$ is an identity matrix of size $N_{\text{pan}}/2^n$. Finally, a few words on the implementation of the change of basis for large clusters. As long as not so large clusters are used, the cost of a change of basis is negligible. If very large clusters are used, and a '*naïve*' implementation (i.e. as a literal matrix–vector product) is used, the cost could affect performance. However, this kind of change of basis can be seen as the '*discrete Haar transform*', which is well-known in the theory of discrete signals. There are efficient implementations of this kind of transformation, in the same spirit of the well known *fast Fourier transform algorithm*.

10. THE MODAL AND MULTIPOLAR PRECONDITIONINGS COMBINED

To combine both preconditionings for a thin airfoil, you first transform the matrix to the modal symmetric/skew-symmetric basis, as explained in Sections 5 and 6. The change of basis matrix is

$$\mathbf{S}_{\text{mod}} = \frac{1}{\sqrt{2}}\begin{bmatrix}
1 & 0 & & \cdots & 1 & 0 & & \cdots \\
1 & 0 & \ddots & & -1 & 0 & \ddots & \\
0 & 1 & \ddots & \ddots & 0 & 1 & \ddots & \ddots \\
0 & 1 & 0 & & 0 & -1 & 0 & \\
\vdots & 0 & 1 & \ddots & \vdots & 0 & 1 & \ddots \\
 & 0 & 1 & \ddots & & 0 & -1 & \ddots \\
 & \vdots & 0 & \ddots & & \vdots & 0 & \ddots \\
 & & 0 & \ddots & & & 0 & \ddots \\
 & & \vdots & & & & \vdots &
\end{bmatrix}. \tag{41}$$

And the transformed matrix is

$$\tilde{\mathbf{A}}_{\text{mod}} = \mathbf{S}^{-1}_{\text{mod}}\mathbf{A}\mathbf{S}_{\text{mod}} = \begin{bmatrix}\tilde{\mathbf{A}}^{++}_{\text{mod}} & \tilde{\mathbf{A}}^{+-}_{\text{mod}} \\ \tilde{\mathbf{A}}^{-+}_{\text{mod}} & \tilde{\mathbf{A}}^{--}_{\text{mod}}\end{bmatrix}. \tag{42}$$

For thin airfoils ($\epsilon \to 0$) $\tilde{\mathbf{A}}^{--}_{\text{mod}}$ tends to a matrix with interaction coefficients for panels with double layer density distributions, as if the finite thin airfoil were replaced by a zero thickness airfoil at the centerplane. Then the multipolar preconditioning can be applied with, say,

4-panels clusters. The change of basis matrix is the identity for the symmetric modes and the multipolar matrix $\mathbf{S}_{\text{mp}}$ for $N_{\text{pan}}/2$ panels given by (40):

$$\mathbf{S} = \mathbf{S}_{\text{mod}} \begin{bmatrix} \mathbf{I} & \mathbf{0} \\ \mathbf{0} & \mathbf{S}_{\text{mp}} \end{bmatrix}. \tag{43}$$

The transformed matrix has a structure of the following form

$$\tilde{\mathbf{A}}_{\text{md/mp}} = \mathbf{S}^{-1}\mathbf{A}\mathbf{S} = \begin{bmatrix} \tilde{\mathbf{A}}^{++}_{\text{mod}} & \tilde{\mathbf{A}}^{+-}_{\text{mod}}\mathbf{S}_{\text{mp}} \\ \mathbf{S}_{\text{mp}}^{-1}\tilde{\mathbf{A}}^{-+}_{\text{mod}} & \mathbf{S}_{\text{mp}}^{-1}\tilde{\mathbf{A}}^{--}_{\text{mod}}\mathbf{S}_{\text{mp}} \end{bmatrix}. \tag{44}$$

The $\mathbf{S}_{\text{mp}}$ change of basis matrix induces a block decomposition (20) on $\tilde{\mathbf{A}}^{--}_{\text{mod}}$ and the same estimation of the order of coefficients as in Equation (34), Section 9, is valid.

$$\mathbf{S}_{\text{mp}}^{-1}\tilde{\mathbf{A}}^{--}_{\text{mod}}\mathbf{S}_{\text{mp}} = \begin{bmatrix} \tilde{\tilde{\mathbf{A}}}^{++}_{\text{md/mp}} & \tilde{\tilde{\mathbf{A}}}^{+-}_{\text{md/mp}} \\ \tilde{\tilde{\mathbf{A}}}^{-+}_{\text{md/mp}} & \tilde{\tilde{\mathbf{A}}}^{--}_{\text{md/mp}} \end{bmatrix}. \tag{45}$$

The combined modal and multipolar preconditioning is obtained in this basis by neglecting all the non-diagonal elements but those in $\tilde{\tilde{\mathbf{A}}}^{++}_{\text{mod}}$ so that the combined preconditioning is

$$\mathbf{Q}_{\text{md/mp}} = \begin{bmatrix} \text{diag}(\tilde{\mathbf{A}}^{++}_{\text{mod}}) & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \tilde{\tilde{\mathbf{A}}}^{++}_{\text{md/mp}} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \text{diag}(\tilde{\tilde{\mathbf{A}}}^{--}_{\text{md/mp}}) \end{bmatrix}. \tag{46}$$

The cost of this preconditioning is related to solving a system for $\tilde{\tilde{\mathbf{A}}}^{++}_{\text{mod}}$, which is a system $N_{\text{pan}}/2m$, i.e. $2m$ times smaller than that of the original one. In the limit of a very large cluster size ($2m = N_{\text{pan}}$) the modal/multipolar preconditioning approaches the pure modal one, since the preconditioning is diagonal (in a somewhat confuse notation, $2m$ is referred to as the size of the cluster so that a 'modal/multipolar preconditioning with 64 panels/cluster' stands for a modal preconditioning combined with a multipolar preconditioning with 32 panels per cluster). For simplicity, the authors have considered so far that the geometry is composed of a single thin airfoil. In the general case, some other elements not necessarily thin, as for instance a fuselage present in the geometry. In that case, the change of basis matrix is block diagonal, with the identity matrix for those elements not in the thin airfoil part, and the usual modal change of basis matrix for those which are on the thin airfoil. After this, the change of basis matrix for the multipolar preconditioning is constructed by clustering those elements in the thin airfoil, as well as those in the other (non-thin) parts. The extension to several thin airfoils is also simple.

11. NUMERICAL EXAMPLES

Firstly, several small scales (approximately 200 panels) will be shown. The interest of these examples is that the full eigenvalue distribution and condition number can be computed using

standard LAPACK routines (embedded in the high level scientific programming language Octave). Secondly, convergence curves for examples with large number of panels (>12000) will be presented. A thin biconvex spherical lens is considered as shown in Figure 11, of relative thickness t. The expression for the surface of the lens is

$$[|z| + R/\tan(\theta)]^2 + x^2 + y^2 = (R/\sin(\theta)^2, \tag{47}$$

where the angle θ is given by $t = 2\tan(\theta/2)$. Figure 12 presents the behavior of the condition number as a function of the relative thickness t for the modal/multipolar preconditioned ($\mathbf{Q}^{-1}\mathbf{A}$) and non-preconditioned ($\mathbf{A}$) matrices. Four values of thicknesses, ranging from 25% to 0.2%, have been considered. It is verified that the condition number for the non-preconditioned matrix is $O(1/\epsilon)$ for $\epsilon \to 0$, whereas it is almost independent of ϵ for the preconditioned matrix. Moreover, the condition number is nearly the same for all cluster sizes. Remember that the computational effort involved in the preconditioning is roughly inversely proportional to the cluster size. However, it will be seen that in practice the rate of convergence does depend on the cluster size. The mesh was composed of 192 triangular panels. Figure 13 presents the eigenvalue spectra for the lens for two thicknesses. The mesh was the same as in the previous paragraph, the modal/multipolar preconditioning has been used. Note that in all cases the eigenvalues are clearly separated in two branches for the non-preconditioned case, and each

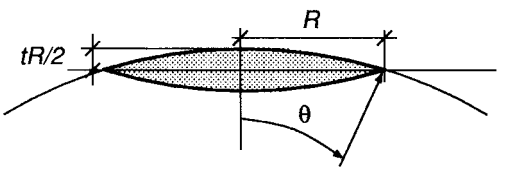

Figure 11. Geometrical description of the lens.

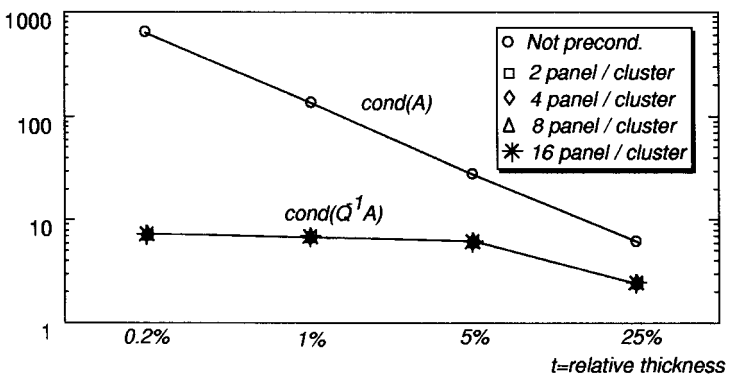

Figure 12. Condition number versus thickness for the lens with several cluster-sizes n_c.

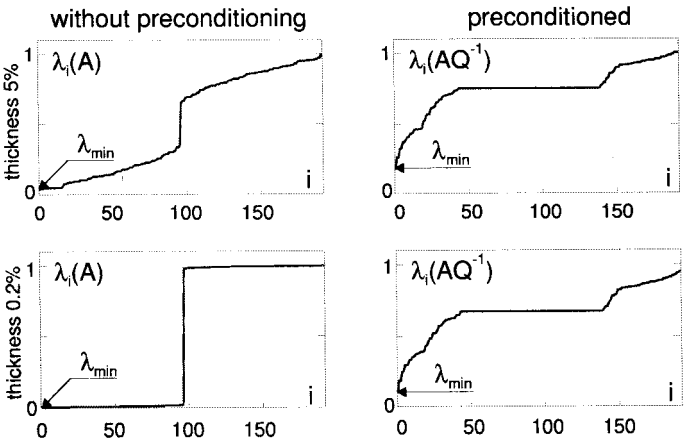

Figure 13. Eigenvalue distribution for the lens with thicknesses 5 and 0.2%.

branch has $N_{\text{pan}} = /2$ eigenvalues. This is typical for problems where a symmetry '*separates*' branches of eigenvalues. As the thickness gets smaller, the lower branch, i.e. that one with the lowest (in absolute value) eigenvalues, gets smaller too, and for very small thicknesses (as $t = 0.2\%$ in the figure), the lower branch is clustered near the origin, whereas the upper branch is clustered near $|\lambda| = 1$. In contrast, the spectra for the preconditioned matrices remains almost unaltered. Figure 14 presents the convergence history for the 5% thick lens at 0, 5 and 90° angles of attack, and cluster sizes n_c ranging from 2 to 64, where the Bi-conjugate gradient (B-CG) [18,19] algorithm was used. The curve labeled as 'J' stands for the non-preconditioned problem. For incidence at 0°, the flow is symmetric, and then the skew-symmetric part of the equations is not excited, then the rate of convergence is poorly improved by the preconditioning. In the other cases, the improvement in the rate of convergence monotonically decreases with cluster size. For a typical case of incidence at 5°, it is seen that even with a large cluster of 64 panels the improvement in the rate of convergence is significant. Similar plots for

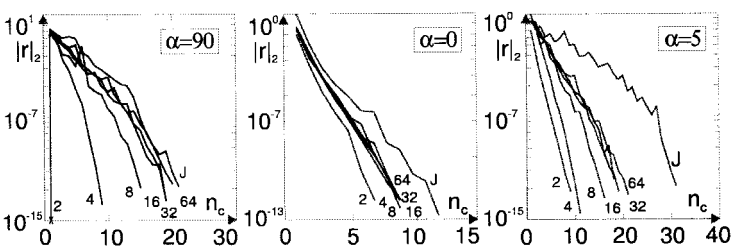

Figure 14. Convergence histories of residual for the 5% thick lens at several angles of attack and cluster size n_c, where 'J' stands for the non-preconditioned problem.

thickness of 0.2% are shown in Figure 15. The improvements are more significant for thinner lenses, as expected. Figures 1–18 show similar results but for a symmetric delta wing, with a wing section consisting of a symmetric Joukowski profile. Figure 19 represents curves of

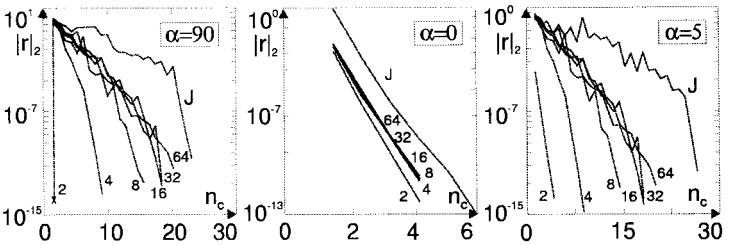

Figure 15. Same as Figure 14 for the 0.2% thick lens.

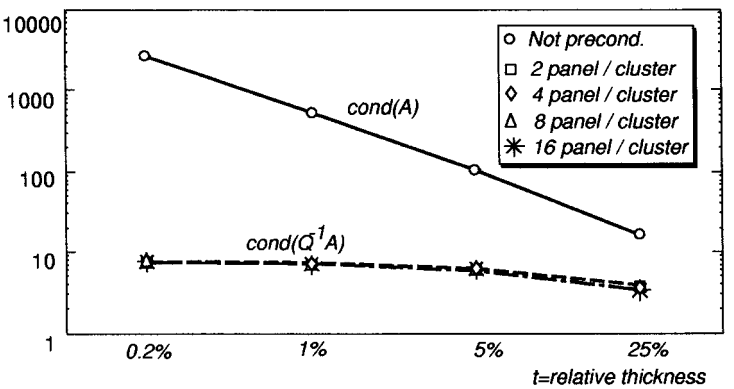

Figure 16. Condition number versus thickness t for several cluster sizes n_c.

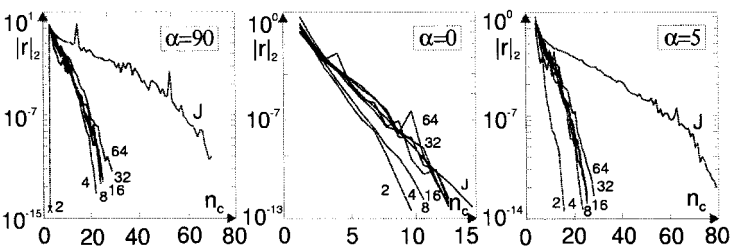

Figure 17. Convergence histories of residual for the 5% thick symmetric delta wing at several angles of attack and cluster size n_c.

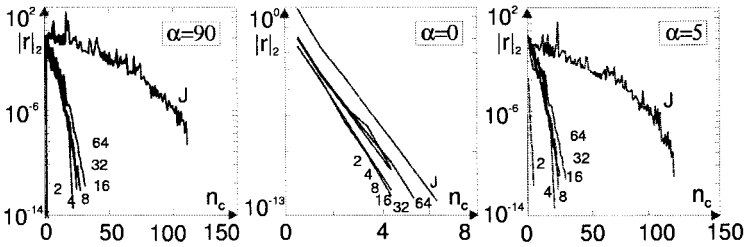

Figure 18. Same as Figure 17 for the 0.2% thick delta wing.

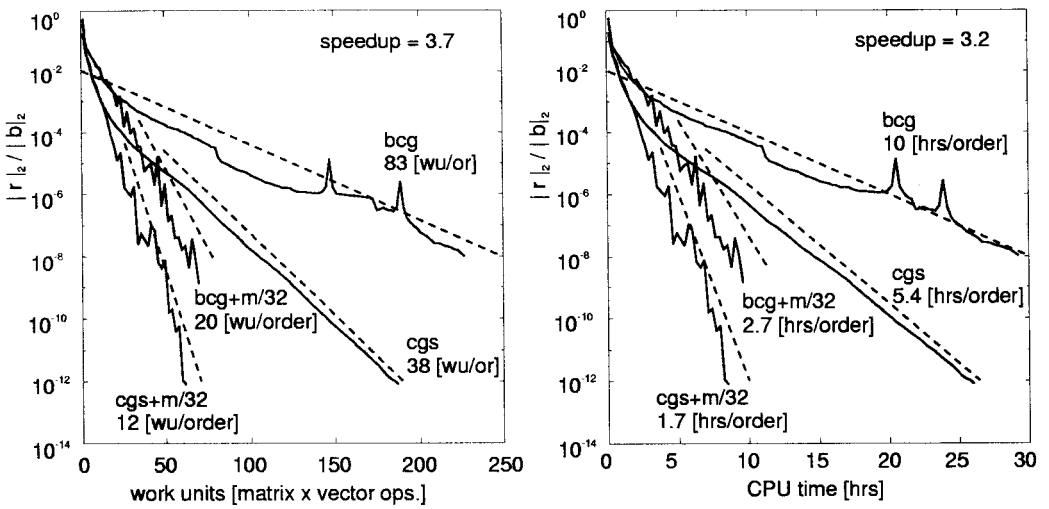

Figure 19. Convergence curves for symmetric delta wing 8% thick, 12288 panels, cluster size $n_c = 32$.

convergence for a delta wing with a typical section being a Joukowski profile, where 12288 panels were used and the airfoil thickness was 8%, see Figure 20. In this case both conjugate gradient squared (CGS) [20] and Bi-conjugate gradient (B-CG) methods were used without and with a preconditioning of 32 panels per cluster. It can be seen that the improvement in the rate of convergence by the preconditioning is significant with both methods and the performance with CGS is better than with B-CG. In all cases, convergence history versus number of matrix–vector products (which will be called from here on '*work units*') and CPU time are presented. A work unit involves computation of the interaction coefficients by columns and standard dot products or DAXPY (the LINPACK vector sum routine) operations. The expression for the interaction coefficients involve transcendental functions and is by far, the most time consuming part of the work unit. To estimate the speedup obtained with the preconditioning we estimated the rates of convergence as the reciprocal of the mean slope of

the curves with and without preconditioning (as the curves are often noisy, the choice of such a mean slope is somewhat arbitrary, the dashed lines in the figures represent the mean slope we adopted). In the case of the residual versus work units curves these rates are expressed in work units/order that is, the number of work units needed to reduce the residual by a factor of ten. In the case of the residual versus time, the corresponding unit is CPU hours per order. Work units rates have the advantage that are independent of processor speed but do not take into account the overhead needed by the preconditioning. It includes changes to and from the modal basis and a back-substitution for the preconditioning matrix. For the 12288 mesh with a 32 panels/cluster preconditioning, the overhead represents a 6.9% of the matrix–vector product operation but depends strongly on details of the implementation. For instance, on a vector processor, the overhead tends to be smaller, since the preconditioning operations (mainly the back-substitution) is much more prone to vectorization than the computation of coefficients. On the other hand, it is common to evaluate the interaction coefficients for distant panels using far-field expansions. In this case, the overhead tends to be higher. The speedup is defined as the quotient between the rates with and without preconditioning, and is a little higher for the work units rates than those based on CPU time

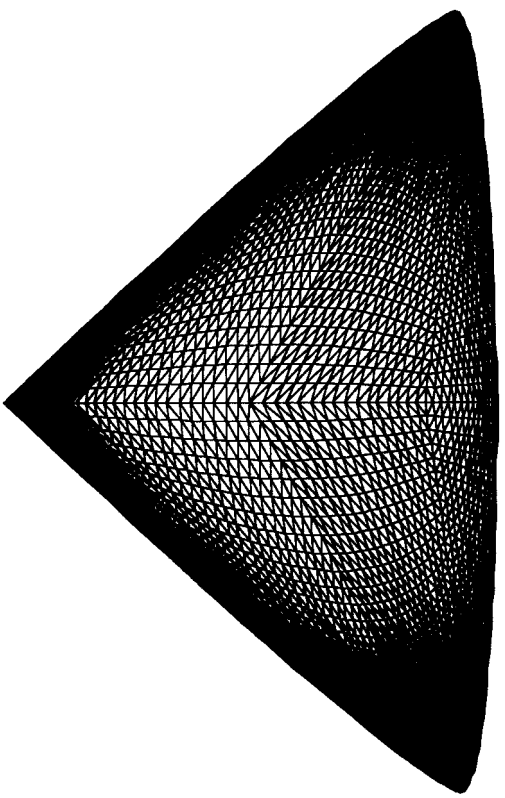

Figure 20. 12288 panels mesh for the 8% thick delta wing.

$$\frac{\text{Speedup (work units)}}{\text{Speedup (CPU time)}} = \frac{\text{CPU time for a work unit with prec.}}{\text{CPU time for a work without prec.}}$$
$$= 1 + \frac{\text{preconditioning CPU time overhead}}{\text{CPU time for a work unit without prec.}} > 1. \quad (48)$$

Note that solution of the full system matrix would require 1.2 Gbytes of RAM, whereas the preconditioning matrix only requires 1.2 Mbytes. Indeed, this problem was ran on a DEC ALPHA/200 workstation of 233 MHz where the core memory requirement was only of 5.2 Mbytes.

12. CONCLUSIONS

The modal/multipolar preconditioning successfully corrects the bad conditioning arising from thin airfoils. It has been shown that the condition number for the resulting linear system is $O(1/\epsilon)$, whereas for the preconditioned system it is almost constant. Numerical results for a biconvex lens and a symmetric airfoil are shown, and significants in condition number are obtained even for airfoils as thick as 25%.

ACKNOWLEDGMENTS

This work has received financial support from *Consejo Nacional de Investigaciones Científicas y Técnicas* (CONICET, Argentina) through grant BID 802/OC-AR PID No. 26, and from *Universidad Nacional del Litoral* (Argentina). The authors made extensive use of software distributed by the *Free Software Foundation/GNU-Project*: Linux ELF-OS, Octave, Fortran f2c compiler, Tgif and others.

APPENDIX A

A.1. $\tilde{\mathbf{A}}_{++}$ *block*

Consider for instance, the coefficient $\tilde{A}_{11}$ which from (11a) and (17) is

$$\tilde{A}_{11} = \tfrac{1}{2}(A_{11} + A_{12} + A_{21} + A_{22}). \quad (49)$$

Recall that A_{ij} is the potential produced by a constant distribution of double layer density on the i-panel on the centroid of the j-panel. If we call ϕ the potential produced by a distribution of double layer density $\mu = 1$ on panel 1 and $\mu = 1$ on panel 2, then

$$A_{i1} + A_{i2} = \phi(\mathbf{x}_i). \quad (50)$$

But, as we discussed previously, for $\epsilon \to 0$, the double layer densities tend to cancel each other and we have

$$\phi \to \epsilon\tilde{\phi}, \quad (51)$$

where $\tilde{\phi}$ is the potential produced by a constant '*quadruple layer*' density. Now,

$$\tilde{A}_{11} = \tfrac{1}{2}[\phi(\mathbf{x}_1) + \phi(\mathbf{x}_2)] \to \tfrac{1}{2}\epsilon[\tilde{\phi}(\mathbf{x}_1) + \tilde{\phi}(\mathbf{x}_2)] \to \tfrac{1}{2}\epsilon\tilde{\phi}(\mathbf{x}_{12}), \tag{52}$$

where $\mathbf{x}_{12}$ is the point lying in the middle of the segment joining $\mathbf{x}_1$ and $\mathbf{x}_2$. This is so for all the elements in the block, then (21) is obtained.

A.2. $\tilde{\mathbf{A}}_{-+}$ *block*

Consider first a non-diagonal coefficient:

$$\tilde{A}_{41} = \tfrac{1}{2}(A_{31} + A_{32} - A_{41} - A_{42}) = \tfrac{1}{2}[\phi(\mathbf{x}_3) - \phi(\mathbf{x}_4)] \to \tfrac{1}{2}\epsilon[\tilde{\phi}(\mathbf{x}_3) - \tilde{\phi}(\mathbf{x}_4)] \to \tfrac{1}{2}\epsilon^2 \frac{\partial\tilde{\phi}}{\partial z}(\mathbf{x}_{34}), \tag{53}$$

where z is a co-ordinate normal to the panels 3 and 4 and we have used the fact that $\tilde{\phi}$ is well-behaved at $\mathbf{x}_{3,4}$. This reasoning is valid for all the non-diagonal elements in $\tilde{\mathbf{A}}_{-+}$, as $\tilde{\mathbf{A}}_{32}$. For the diagonal elements

$$\tilde{\mathbf{A}}_{31} = \tfrac{1}{2}(A_{11} + A_{12} - A_{21} - A_{22}) = \tfrac{1}{2}[\phi(\mathbf{x}_1) - \phi(\mathbf{x}_2)] = \tfrac{1}{2}\epsilon[\tilde{\phi}(\mathbf{x}_1) - \tilde{\phi}(\mathbf{x}_2)] + O(\epsilon^2), \tag{54}$$

but now $\tilde{\phi}$ has some degree of discontinuity across the panel and it can not be expanded in power series. However, as mentioned above, the '*quadruple layer density*' is a symmetric distribution of density about the centerplane and then $\tilde{\phi}(\mathbf{x}_{12} + z\hat{\mathbf{n}})$ is even about $z = 0$, so that

$$\tilde{\mathbf{A}}_{31} = \tfrac{1}{2}\epsilon[\tilde{\phi}(\mathbf{x}_1) - \tilde{\phi}(\mathbf{x}_2)] + O(\epsilon^2) = \tfrac{1}{2}\epsilon[\tilde{\phi}(\mathbf{x}_{12} + \tfrac{1}{2}\epsilon\hat{\mathbf{n}}) - \tilde{\phi}(\mathbf{x}_{12} - \tfrac{1}{2}\epsilon\hat{\mathbf{n}})] + O(\epsilon^2) = O(\epsilon^2). \tag{55}$$

As all the elements (diagonal and non-diagonal) are $O(\epsilon^2)$, then (22) is obtained.

A.3. $\tilde{\mathbf{A}}_{+-}$ *block*

Consider the off-diagonal coefficients:

$$A_{23} = \tfrac{1}{2}(A_{31} - A_{32} + A_{41} - A_{42}) = \psi(\mathbf{x}_3) + \psi(\mathbf{x}_4), \tag{56}$$

where ψ is one-half the potential produced by a distribution of double layer density $\mu = 1$ on panel 1 and $\mu = -1$ on panel 2. In this case, the double layer densities tend to reinforce each other and

$$\psi \to \tilde{\psi} \quad \text{for } \epsilon \to 0. \tag{57}$$

In addition, ψ is well-behaved at $\mathbf{x}_{3,4}$ so that

$$A_{23} \to \tilde{\psi}(\mathbf{x}_3) + \tilde{\psi}(\mathbf{x}_4) \to 2\tilde{\psi}(\mathbf{x}_{34}) = O(1), \tag{58}$$

in contrast, the diagonal terms are of the form

$$A_{13} \to \tilde{\psi}(\mathbf{x}_1) + \tilde{\psi}(\mathbf{x}_2) = O(\epsilon), \tag{59}$$

since $\tilde{\psi}$ is odd across the panel. Then, (23) is obtained.

A.4. $\tilde{\mathbf{A}}_{\_\ \_}$ *block*

For the off-diagonal terms

$$\tilde{\mathbf{A}}_{34} = \tfrac{1}{2}(A_{31} - A_{32} - A_{41} + A_{42}) = \psi(\mathbf{x}_3) - \psi(\mathbf{x}_4) \to \tilde{\psi}(\mathbf{x}_3) - \tilde{\psi}(\mathbf{x}_4) \to \epsilon \frac{\partial \tilde{\psi}}{\partial z}(\mathbf{x}_{34}) = O(\epsilon), \tag{60}$$

since $\tilde{\psi}$ is well-behaved on $\mathbf{x}_{3,4}$. On the other hand, for the diagonal terms,

$$\tilde{\mathbf{A}}_{33} = \tfrac{1}{2}(A_{11} - A_{12} - A_{21} + A_{22}) = \psi(\mathbf{x}_1) - \psi(\mathbf{x}_2) \to \tilde{\psi}(\mathbf{x}_1) - \tilde{\psi}(\mathbf{x}_2) \to 2\tilde{\psi}(\mathbf{x}_1) = O(1), \tag{61}$$

since $\tilde{\psi}$ is odd across the panel. Then, (24) is obtained.

APPENDIX B. NOMENCLATURE

$\mathbf{A} \otimes \mathbf{B}$	Kronecker product of two matrices defined by (29)
$\mathbf{A}$	system matrix of interaction coefficients for double layer density distribution
$\tilde{\mathbf{A}}$	matrix $\mathbf{A}$ in the modal or multipolar basis
$\mathbf{C}$	system matrix of interaction coefficients for single layer density distribution
cond($\mathbf{X}$)	condition of matrix $\mathbf{X}$
d	distance between panels
diag( )	(i) for matrices diag($\mathbf{X}$) is the diagonal part of $\mathbf{X}$, see Equation (14), whereas (ii) diag(a, b, c, . . .) stands for a matrix whose diagonal entries are a, b, c, . . .
g	potential produced by a single panel on a normal axis passing through its centroid
H	length side of the square panels
m	number of panels in a cluster
$\hat{\mathbf{n}}$	normal unit vector
N_{pan}	number of panels
$\mathbf{Q}$	preconditioning matrix
$\mathbf{S}$	change of basis matrix
t	airfoil thickness
$\mathbf{u}_\infty$	undisturbed velocity

Greek letters

ϵ	vanishing parameter scaling the airfoil thickness
ϕ	velocity potential
μ	double layer density on the surface
σ	single layer density
Γ	airfoil surface
Ω	exterior domain to the airfoil

REFERENCES

1. J. Katz and A. Plotkin, *Low-Speed Aerodynamics, from Wing Theory to Panel Methods*, McGraw-Hill, New York, 1991.
2. J. D'Elía, M. Storti and S. Idelsohn, 'A CVBEM formulation for multiple profiles and cascades', *Appl. Mech. Rev.*, **48**, S203–S210 (1995).
3. J. D'Elía, M. Storti and S. Idelsohn, 'A 3D panel code for wave resistance calculations. Part I: General formulation and discretization', *Rev. Int. Métodos Numér. Cálc. Diseño Ing.*, **13**, 515–530 (1997).
4. M. Storti, J. D'Elía and S. Idelsohn, 'Algebraic discrete non-local (DNL) absorbing boundary condition for the ship wave resistance problem', *J. Comput. Phys.* (1998) to appear.
5. L. Morino and C.C. Kuo, 'Subsonic potential aerodynamics for complex configurations: a general theory', *AIAA J.*, **12**, 191–197 (1974).
6. L. Morino (ed.), *Computational Methods in Potential Aerodynamics*, Springer, Berlin, 1985.
7. M. Gennaretti and L. Morino, 'A boundary element method for the potential, compressible aerodynamics of bodies in arbitrary motion', *AIAA J.*, **96**, 15–19 (1992).
8. T. Maître, 'Modelisation de l'ecoulémént autour d'une hélice marine par la méthode du potential', *Ph.D. Tesis*, Institut National Polytechnique de Grenoble, 1988.
9. D.J. Evans (ed.), *Preconditioning Methods, Theory and Applications*, Gordon and Breach Science Publishers, 1983.
10. K.G. Prasad, J.H. Kane and D.E. Keyes, 'Iterative solution techniques in boundary element analysis', *Int. J. Numer. Methods Eng.*, **31**, 1511–1536 (1991).
11. M. Hriberšek, P. Škerget and H. Mang, 'Preconditioned conjugate gradient methods for boundary domain integral method', *Eng. Anal. Bound. Elem.*, **12**, 111–118 (1993).
12. W.J. Manzur, F.C. Araújo and E.B. Malaghini, 'Solution of BEM systems of equations via iterative techniques', *Int. J. Numer. Methods Eng.*, **33**, 1823–1841 (1992).
13. W. Hackbusch and Z.P. Nowak, 'On the fast matrix multiplication in the boundary element method by panel clustering', *Numer. Math.*, **54**, 463–491 (1989).
14. K. Nabors, F.T. Korsmayer, F.T. Leighton and J. White, 'Preconditioned, adaptative, multipole-accelerated iterative methods for three-dimensional first-kind integral equations of potential theory', *SIAM J. Sci. Comput.*, **15**, 713–735 (1994).
15. S.A. Vavasis, 'Preconditioning for boundary integral equations', *SIAM J. Matrix Anal. Appl.*, **13**, 905–925 (1992).
16. J. D'Elía, 'Numerical methods for the ship wave resistance problem', *Ph.D. Thesis*, 1997.
17. N.M. Nachtigal, S.C. Reddy and L.N. Trefethen, 'How fast are non-symmetric matrix iterations?', *SIAM J. Matrix Anal. Appl.*, **13**, 778–795 (1992).
18. C. Lanczos, 'Solution of systems of linear equations by minimized iteration', *J. Res. Nat. Bur. Stand.*, **49**, 33–53 (1952).
19. R. Fletcher, 'Conjugate gradient methods for indefinite systems', *Lecture Notes Math.*, **506**, 73–89 (1976).
20. P. Sonneveld, 'CGS, a fast Lanczos-type solver for non-symmetric linear systems', *SIAM J. Sci. Stat. Comput.*, **10**, 36–52 (1989).
21. Y. Yan, 'Sparse preconditioned iterative methods for dense linear systems', *SIAM J. Sci. Comput.*, **15**, 1190–1200 (1994).
22. M.R. Hestenes and E. Stiefel, 'Methods of conjugate gradients for solving linear systems', *J. Res. Nat. Bur. Stand.*, **49**, 409–436 (1952).

23. C.C. Paige and A. Saunders, 'LSQR: An algorithm for sparse linear equations and sparse least squares', *ACM Trans. Math. Software*, **8**, 43–71 (1982).
24. S.C. Eisentat, H.C. Helman and M.H. Schultz, 'Variational iterative methods for non-symmetric systems of linear equations', *SIAM J. Numer. Anal.*, **20**, 345–357 (1983).
25. Y. Saad and M.H. Schultz, 'Conjugate gradient-like algorithms for solving non-symmetric linear systems', *Math. Comput.*, **44**, 417–424 (1985).
26. M.H. Gutknecht, 'The unsymmetric Lanczos algorithms and their relations to Pade approximation, continued fractions and the qd algorithm', *Prelim. Proc. Copper Mountain Conference on Iterative Methods*, 1990.
27. H.A. Van der Vorst, 'Bi-CGSTAB: A fast and smoothly converging variant of Bi-CG for the solution of non-symmetric linear systems', *SIAM J. Sci. Stat. Comput.*, **13**, 631–644 (1992).
28. R.W. Freund, 'Conjugate gradient-type methods for linear systems with complex symmetric coefficient matrices', *SIAM J. Sci. Stat. Comput.*, **13**, 425–448 (1992).
29. M.A. Saunders, H.D. Simon and E.L. Yip, 'Two conjugate gradient-type methods for unsymmetric linear equations', *SIAM J. Numer. Anal.*, **25**, 927–940 (1988).
30. Y. Saad, 'Krylov subspace methods on supercomputers', *SIAM J. Sci. Stat. Comp.*, **10**, 1299–1232 (1989).
31. A. Greebaum, 'Estimating the attainable accuracy of recursively computed residual methods', *SIAM J. Matrix Anal. Appl.*, **18**, 535–551 (1997).

