

# Preconditioning harmonic unsteady potential flow calculations

Alistair L. Laird

Michael B. Giles  
giles@comlab.ox.ac.uk

*Oxford University Computing Laboratory, Oxford OX1 3QD*

This paper considers finite element discretisations of the Helmholtz equation and its generalisation arising from harmonic acoustics perturbations to a non-uniform steady potential flow. A novel elliptic, positive definite preconditioner, with a multigrid implementation, is used to accelerate the iterative convergence of Krylov subspace solvers.

Both theory and numerical results show that for a model 1D Helmholtz test problem the preconditioner clusters the discrete system's eigenvalues and lowers its condition number to a level independent of grid resolution. For the 2D Helmholtz equation, grid independent convergence is achieved using a QMR Krylov solver, significantly outperforming the popular SSOR preconditioner. Impressive results are also presented on more complex domains, including an axisymmetric aircraft engine inlet with non-stagnant mean flow and modal boundary conditions.

Oxford University Computing Laboratory  
Numerical Analysis Group  
Wolfson Building  
Parks Road  
Oxford, England OX1 3QD

January, 2005

# 1 Introduction

The Helmholtz equation

$$-\nabla^2 \phi - \omega^2 \phi = 0, \quad (1.1)$$

models the propagation of an acoustic wave with perturbation potential  $\phi e^{i\omega t}$  through a uniform stationary flow with unit speed of sound. A finite element approximation on a 3D problem with a high frequency will yield a very large discrete system of equations, for which direct solution will have excess memory and CPU requirements.

The alternative is to use an iterative solution method, with preconditioning to accelerate the iterative convergence. The problems associated with the iterative solution of the Helmholtz equation are however well documented, see for example [3] or [9]. As  $\omega$  increases, the discretisation matrix becomes highly indefinite and ill-conditioned. In addition the necessary inclusion of radiating boundary conditions produces a non-Hermitian<sup>1</sup> system matrix. The indefiniteness of the matrix has proved to be a stumbling block for multigrid approaches [6], leaving the more robust Krylov subspace methods as the usual choice for such problems. However, the non-Hermitian nature of the matrix rules out use of the CG method [16], whilst the ill-conditioning can prove highly problematic for other Krylov methods.

Consequently, much work has been done on preconditioning the discrete system to greatly improve the convergence rate of Krylov solvers. In 1983, Bayliss *et al.* set down what is often regarded as the benchmark, using a CGNR solver with a Laplacian preconditioner, inverted with one sweep of SSOR for the finite element solution of the Helmholtz equation [4]. This work was complemented the following year by Gozani *et al.* using multigrid for the inversion [15]. This issue of inexpensive preconditioner implementation also motivated the work initiated in 2D by Ernst and Golub [11] and generalised to 3D by Elman and O’Leary [9]. They selected a block preconditioner, based on the linear operator itself, and compatible with fast Poisson-type inversion. A very similar approach was also used by Otto and Larson [22]. Although all this work was carried out on simple domains, Ernst showed how it can be generalised onto more complex geometries using capacitance matrices [10].

Although basic multigrid is not a sensible option, several authors have looked into multigrid preconditioning for an outer Krylov iteration. This was first attempted by Shapira [30], who used red-black Gauss-Seidel smoothing down to the coarsest grid of two points per wavelength. However, Elman [8] found that using such classical iterative techniques failed to smooth all error modes on the coarser grids, and he therefore proposed to use GMRES as the smoother and a flexible GMRES algorithm for the outer loop on these troublesome grids, reverting back to the more traditional Jacobi when the mesh was sufficiently fine. His results were very encouraging, although it remains to be seen how such a method would compare to less memory intensive iterations.

Considering the more common preconditioners, Otto found that the Jacobi and ILU

---

<sup>1</sup>Hermitian refers to the conjugate transpose

preconditioners perform poorly [22], but Made has recently attempted to perturb the real part of  $A$  to make it positive definite, or at least less indefinite, before using an ILU preconditioner and a restarted GMRES solver [20]. His results show significant improvements against those of an ILU preconditioner on the unperturbed system. Malhorta *et al.* [21] used a block version of the QMR algorithm for problems with multiple right-hand sides and investigated the use of both SSOR and hierarchical basis functions. The SSOR preconditioner gave best performance even when the problem was reduced to only one right-hand side.

The Helmholtz equation can also be viewed as the stationary mean flow case of the more general full linearised harmonic potential equation

$$\nabla \cdot \left( \rho \nabla \phi - \frac{\rho}{c^2} (\nabla \bar{\phi} \cdot \nabla \phi + i\omega \phi) \nabla \bar{\phi} \right) - i\omega \frac{\rho}{c^2} (\nabla \bar{\phi} \cdot \nabla \phi + i\omega \phi) = 0, \quad (1.2)$$

for the harmonic perturbation potential superimposed upon a nonuniform steady mean flow defined by the potential  $\bar{\phi}$ , with density  $\rho$  and speed of sound  $c$  given by

$$c^2 = \rho^{\gamma-1} = 1 - \frac{\gamma-1}{2} (\nabla \bar{\phi} \cdot \nabla \bar{\phi} - M_\infty^2), \quad (1.3)$$

where  $M_\infty$  is the farfield Mach number. All quantities are non-dimensionalised using the density and speed of sound in the far field, and an appropriate reference length  $R$ . Consequently the acoustic potential is non-dimensionalised with respect to  $c_\infty R$  and frequency with respect to  $c_\infty/R$ . Such an equation has widespread engineering applications and presents us with the motivation for this work: the desire to perform 3D finite element aeroacoustic calculations for the radiation from turbofan engine inlets.

Although the difficulties faced in solving the full equation (1.2) are very similar to those of the Helmholtz equation (1.1), there has been little novel research into its iterative solution, with a preference for the traditional techniques, such as ILU preconditioning [26] and Jacobi and Schwartz decomposition preconditioners with the TFQMR algorithm [32]. This is in stark contrast to the work on the Helmholtz equation, but unfortunately most of these methods do not generalise onto the full equation on arbitrary geometries.

The new preconditioning matrix introduced in this paper is elliptic, positive definite, and closely related to the original problem. In the first part of the paper it is developed for the Helmholtz equation. When applied to a simple 1D test problem, a theoretical and computational analysis show it to both considerably cluster the system's eigenvalues and lower the condition number to a level independent of the grid resolution. To avoid the expense of applying the preconditioner exactly, the preconditioner is implemented approximately via a single multigrid V-cycle with optimised Jacobi smoothing. Numerical computations show such an approximation to have negligible adverse preconditioning effects over an exact implementation. When coupled with a variety of Krylov subspace schemes for a 2D model problem, this new preconditioning technique gives considerable convergence acceleration. The favoured scheme is Freund's complex symmetric QMR-SYM algorithm [14] which shows results comparable to the 'optimal' GMRES algorithm

but without the memory requirements. It gives iteration numbers only weakly dependent on grid resolution and scaling approximately linearly with the reduced frequency, which is a substantial improvement over the popular SSOR preconditioner.

The second part of the paper presents the generalisation of the preconditioner for the full equation (1.2). Results are presented on both an axisymmetric aircraft inlet geometry and rectangular and cylindrical ducts with base flow and modal boundary conditions. Such results continue to show little dependence on grid resolution and in many cases show considerable improvement in convergence. Final conclusions are drawn in Section 4.

## 2 Helmholtz equation

### 2.1 1D and 2D model problems

The two model problems are the Helmholtz equation defined on the unit interval and unit square, as shown in Figure 1. In each case, the inhomogeneous Neumann boundary condition at  $x = 1$  provide a sound source which generates an acoustic wave for which the Bayliss-Turkel radiating boundary condition [5] at  $x = 0$  is perfectly non-reflecting. The 2D model is completed with zero normal flow boundary conditions at  $y = 0, 1$ .

### 2.2 Finite element discretisation

For both problems we use a Galerkin finite element discretisation with isoparametric Lagrange elements on a uniform Cartesian grid. The starting point for such a formulation is the approximate solution

$$\phi = \sum_{j=1}^n \phi_j N_j \quad (2.1)$$

where  $\phi_j$  is the approximation to  $\phi$  at node  $j$ , and  $N_j$  is the corresponding first-order Lagrange shape function. This approximation is then substituted into the weak form of the Helmholtz equation,

$$\int_{\Omega} \nabla \psi \cdot \nabla \phi - \omega^2 \psi \phi \, d\Omega = \int_{\Gamma} \psi \frac{\partial \phi}{\partial \mathbf{n}} \, d\Gamma, \quad (2.2)$$

where  $\Omega$  is the problem's domain with boundary  $\Gamma$ , the normal derivative at the boundary  $\frac{\partial \phi}{\partial \mathbf{n}}$  is evaluated using the boundary conditions, and the test function  $\psi$  is taken to be  $N_j$  for  $j = 1, \dots, n$ . This yields the complex symmetric linear system

$$A \boldsymbol{\phi} \equiv \left( K - \omega^2 M + B \right) \boldsymbol{\phi} = \mathbf{f}, \quad (2.3)$$

where  $K$  is the usual symmetric positive semi-definite stiffness matrix,  $M$  is the symmetric positive definite mass matrix,  $B$  is a complex symmetric matrix arising from the

radiation boundary condition,  $\boldsymbol{\phi}$  is the vector of unknowns  $\phi_j$ , and  $\mathbf{f}$  comes from the inhomogeneous boundary condition at  $x=0$ .

The iterative convergence of Krylov methods will be dependent on the system's spectrum. Taking the  $l_2$ -condition number  $\kappa_2$ , defined in terms of a system's singular values ( $\sigma$ ) as

$$\kappa_2 = \frac{\sigma_{\max}}{\sigma_{\min}}, \quad (2.4)$$

then a standard result for the iterative convergence of the CG method (see for example [33]) gives the result that the convergence rate of the CGNR algorithm is faster for smaller values of  $\kappa_2$ . The GMRES algorithm on the other hand prefers the stronger condition that the system's eigenvalues ( $\lambda$ ) are clustered away from the origin [17]. In this work, we aim to construct a preconditioner which gives a preconditioned system with both a low value for  $\kappa_2$ , and eigenvalues clustered away from the origin.

### 2.3 1D results

The behaviour of both the eigenvalues and the condition number can be analysed in 1D. Supposing a uniform stepsize  $h$ , then in the limit  $h \rightarrow 0$ , with a fixed  $\omega \gg 1$ , Laird [18] derives the asymptotic behaviour

$$\lambda_{\min} = O(\omega), \quad \lambda_{\max} = O(h^{-2}) \quad \implies \quad \kappa_2(A) = O\left(\omega (n^\lambda)^2\right), \quad (2.5)$$

where  $n^\lambda = 2\pi/(\omega h)$  is the number of grid points per wavelength. This is typically in the range 8–16 for aeroacoustic applications, although in other applications it may be considerably larger, for example to resolve the details of a complex geometry. This grid dependence in both the eigenvalues and singular values is highly undesirable, potentially leading to a large spread of eigenvalues and a high condition number, considerably deteriorating convergence rates. This makes clear the need for effective preconditioning at higher frequencies.

To improve the rate of convergence when solving (2.3) iteratively the preconditioned equation

$$P^{-1}A\boldsymbol{\phi} = P^{-1}\mathbf{f} \quad (2.6)$$

is solved, where the preconditioner  $P$  is chosen so that the iterative convergence of matrix  $P^{-1}A$  is faster than that of  $A$ , whilst inverting  $P$  (or solving equations of the form  $P\mathbf{p} = \mathbf{d}$ ) is computationally much cheaper than inverting  $A$ .

Our new preconditioner is obtained from the Galerkin discretisation by simply omitting the boundary condition  $B$  and changing the sign of the mass matrix term, giving the symmetric positive definite matrix

$$P = K + \omega^2 M. \quad (2.7)$$

Laird's analysis [18] for the 1D problem as  $h \rightarrow 0$ , for fixed  $\omega \gg 1$ , gives

$$\lambda_{\min}(P^{-1}A) = O(1/\omega) \quad \lambda_{\max}(P^{-1}A) = O(1) \quad \implies \quad \kappa_2(P^{-1}A) = O(\omega). \quad (2.8)$$

with obvious benefits for iterative convergence. In contrast, a similar analysis of the Laplacian preconditioner used in [4, 15] gives  $\lambda_{\min} = O(1/\omega)$ ,  $\lambda_{\max} = O(\omega^2)$  and  $\kappa_2 = O(\omega^3)$ , resulting in a significant scattering of the spectrum and deterioration of the condition number at higher frequencies.

Figures 2 and 3 show numerical evidence to support the theory when the preconditioner  $P$  is implemented exactly. Figure 2 shows the clustering of eigenvalues into a disk in the complex plane of radius one, whilst Figure 3 shows the numerical values of  $\kappa_2$  for the 1D problem. The behaviour is as predicted with  $\kappa_2$  decreasing to a value nearly independent of grid resolution, scaling only linearly with  $\omega$ .

The excellent preconditioning properties of  $P$  have been demonstrated both theoretically and numerically in 1D, but to be efficiently utilised it must be possible to invert it with a computational cost which is much less than the direct solution of the original problem. At each step of a Krylov solver iteration, one has to solve systems of the form

$$P \mathbf{p} = \mathbf{d}. \quad (2.9)$$

If this is solved exactly by Gaussian elimination, the computational cost would be as great as the direct solution of the original unpreconditioned equation, and so nothing would have been gained. Instead, the key is to reduce the computational cost of the preconditioner through the use of an approximate solution. This is achieved by taking advantage of the symmetric positive-definite nature of  $P$  which makes it ideally suited for a number of highly efficient algorithms, in particular multigrid [12]. The advantage of multigrid over a direct method is that, for a given solution accuracy, it has complexity  $O(n)$  where  $n$  is the number of unknowns, which gives huge computational savings, especially on large 3D systems. Furthermore multigrid can be used to approximately implement  $P$  using only a limited number of cycles, whilst hopefully still retaining most of the favourable preconditioning properties. The idea of approximate multigrid preconditioning has been used extensively both on the Helmholtz equation ([15],[8]), and in a number of other applications (e.g. [25], [31]).

The work in this paper uses geometric multigrid with (bi\tri-)linear interpolation and relaxed Jacobi smoothing. The relaxation parameter is optimised analytically using a local Fourier analysis [18]. To minimise the cost, only one V-cycle multigrid iteration is used, with one pre- and post-smoothing step on the finest grid, and two on the coarser grids. This gives a cost which is approximately three times that of evaluating the residual error from the Galerkin equations, and an overall cost of a preconditioned matrix-vector multiplication of approximately the same cost as four unpreconditioned matrix-vector multiplications. Figure 4 compares the effects on the condition number of this approximate implementation of the preconditioner, compared to the exact implementation. It shows that almost all of the positive features of the preconditioner are retained.

## 2.4 2D results

We now investigate the performance of the preconditioner when applied to the 2D test problem using different Krylov subspace methods. For a non-Hermitian indefinite linear system, GMRES [29] represents the only ‘optimal’ Krylov subspace method. The optimality refers to the minimisation of some error norm at each iteration, thus (theoretically) guaranteeing convergence. In fact it minimises the  $l_2$ -norm of the residual  $\mathbf{r}^k = \mathbf{f} - A\phi^k$ , or  $P^{-1}\mathbf{r}^k$  with preconditioning, at each iteration  $k$ . The drawback of GMRES is that it is based on the memory intensive Arnoldi factorisation [1], leading to high memory requirements and a large computational workload at higher iteration numbers. The QMR algorithm [13] minimises an approximation to this residual norm at each iteration based on a short recurrence bi-Lanczos factorisation [19]. Furthermore due to the complex symmetric system arising in the Helmholtz equation the QMR-SYM algorithm [14] can be used, which approximately halves the amount of work of QMR to one matrix-vector multiplication (MVM) per iteration.

Figure 5 shows the iterative convergence of the relative residual for both GMRES and QMR-SYM on a typical problem with frequency  $\omega = 24$  and grid spacing  $h = 1/32$ . As the theory dictates, GMRES exhibits a monotonic smooth decrease in residuals, and QMR-SYM behaves similarly. Further evidence of this similar convergence is shown in Table 1 which presents results for a variety of frequencies and grid spacings. In each case, the stopping criterion for the iteration is  $\|\mathbf{r}^k\| / \|\mathbf{f}\| < 10^{-6}$ . Again there is little difference between the two methods, suggesting that for this application QMR-SYM has a convergence very close to optimal.

Looking now at the effectiveness of the preconditioner, Table 1 shows the iteration number to have only a mild dependence on grid resolution and a linear scaling with the reduced frequency. Both properties are highly desirable and are consistent with the 1D results. They are also competitive with the results of Elman [8], which are currently amongst the best in the literature for such a problem. This is demonstrated graphically by the continuous lines in Figure 6, showing the cost savings of using the new preconditioner. The cost measure is (fine grid) MVMs, and as explained earlier one can quantify a preconditioned iteration as being four times as expensive as an unpreconditioned iteration. We see substantial computational savings from our new preconditioner, especially for the higher resolution grid and at higher frequencies.

In Section 2.3 the theoretical advantages of the new preconditioner over the original Laplacian preconditioner were discussed. In Figure 7 we compare the numerical performance of our preconditioner to that of the SSOR preconditioner, a highly regarded preconditioner for this problem (see for example [21] and [22]). It is implemented via the efficient Eisenstat technique [7] giving the cost of a preconditioned MVM comparable to that of an unpreconditioned MVM. The problem is conducted with a number of different SSOR parameters (at intervals of 0.2) and the two most impressive cases are plotted. Again one can see tremendous advantages with the new preconditioner, most notably at the more troublesome high frequencies.

### 3 General harmonic potential equation

#### 3.1 3D and axisymmetric model problems

In this section we consider models based on the general harmonic equation, Eq. (1.2), for harmonic perturbations to a non-stationary mean flow. These demonstrate the extensibility of the new preconditioner.

The first two problems are 3D duct problems, one with a square cross-section and the other with a circular cross-section. Each has unit length in the  $x$ -direction, with boundaries at  $x = 0, 1$ , and hard wall boundary conditions on the boundaries  $y = 0, 1$ ,  $z = 0, 1$  for the rectangular duct and  $y^2 + z^2 = 1/4$  for the cylindrical duct. The mean flow is defined to have a uniform velocity  $u$  in the  $x$ -direction, and uniform speed of sound  $c$ .

The analytic solution of the general harmonic equation for such duct problems can be expressed as an infinite sum of the duct modes  $\phi_m$ :

$$\phi = \sum_m \left[ a_m^+ e^{ik_m^+(x-1)} + a_m^- e^{ik_m^-(x-1)} \right] \phi_m. \quad (3.1)$$

The coefficients  $a_m^+$  and  $a_m^-$  are the amplitudes of right- and left-travelling waves respectively and

$$k_m^+ = \begin{cases} \frac{u\omega - c\sqrt{\omega^2 - (c^2 - u^2)\mu_m}}{c^2 - u^2} & \text{if } \omega^2 \geq (c^2 - u^2)\mu_m \\ \frac{u\omega + ic\sqrt{(c^2 - u^2)\mu_m - \omega^2}}{c^2 - u^2} & \text{if } \omega^2 < (c^2 - u^2)\mu_m \end{cases}$$

$$k_m^- = \begin{cases} \frac{u\omega + c\sqrt{\omega^2 - (c^2 - u^2)\mu_m}}{c^2 - u^2} & \text{if } \omega^2 \geq (c^2 - u^2)\mu_m \\ \frac{u\omega - ic\sqrt{(c^2 - u^2)\mu_m - \omega^2}}{c^2 - u^2} & \text{if } \omega^2 < (c^2 - u^2)\mu_m. \end{cases}$$

for duct eigenvalues  $\mu_m$ , arising from the eigenvalue problem

$$-\nabla^2 \phi_m(y, z) = \mu_m \phi(y, z). \quad (3.2)$$

The real values of  $k^\pm$  signify propagating ‘cut-on’ modes whereas complex values signify cut-off modes. These latter modes decay exponentially along the duct, and modes with fast decay rates are thus of little importance for acoustic radiation. The duct eigenmodes are products of cosine waves in  $y$  and  $z$  for the rectangular problem, and a product of circumferential Fourier modes and radial Bessel functions in the cylindrical problem.

At  $x = 1$ , we specify for the rectangular duct problem that the amplitudes of the incoming modes are all zero except for the mode proportional to  $\cos \pi y \cos \pi z$ . Similarly, for the cylindrical duct problem a non-zero amplitude is specified for the eigenmode



$J_0(\lambda_0 r)$  where  $\lambda_0$  is the first non-zero root of  $J_0'(\lambda_0/2) = 0$ . The models are completed with the Bayliss-Turkel radiating boundary condition [5]

$$\frac{\partial \phi}{\partial x} = ik^- \phi, \quad (3.3)$$

at  $x = 0$ , where  $k^-$  corresponds to the input mode, making the boundary condition perfectly non-reflecting for this mode.

The final problem models the axisymmetric aircraft inlet depicted in Figure 8. The boundaries consist of a hard wall nacelle ( $\Gamma_n$ ) and central axis ( $\Gamma_a$ ), the fan ( $\Gamma_f$ ), a radial far field boundary ( $\Gamma_r$ ) and baffle ( $\Gamma_b$ ). The baffle is artificial and included only for numerical reasons to both lower the size of the discrete system and to produce a grid with low aspect ratio, important for effective geometric multigrid. The problem is defined in the axisymmetric coordinates  $(x, r)$ , periodicity then allows us to consider solutions of the form  $\phi(x, r) e^{i\kappa\theta}$  for a given circumferential wavenumber  $\kappa$ . This reduces Eq. (1.2) to an axisymmetric problem.

The mean flow is no longer uniform and is instead defined by solving the steady full potential flow equation subject to a homogeneous Dirichlet condition on  $\Gamma_r$  corresponding to stagnant flow in the far field, a homogeneous Neumann boundary condition on  $\Gamma_n$  to give zero mass flow through the nacelle, and the inhomogeneous Neumann boundary condition

$$\frac{\partial \bar{\phi}}{\partial \mathbf{n}} = \beta, \quad (3.4)$$

on  $\Gamma_f$  for a given mass-flux through the fan face.

For the acoustic problem, a uniform axial steady velocity is assumed on the fan face, and an incoming acoustic mode is specified as in the duct problems. On the nacelle, there is again no normal flow, whilst on  $\Gamma_r$  and  $\Gamma_b$  a Bayliss-Turkel radiating boundary condition is used

$$\frac{\partial \phi}{\partial \mathbf{n}} = -ik \phi, \quad (3.5)$$

with an appropriate non-reflecting value for  $k$  being constructed using high frequency ray theory [18].

## 3.2 Finite element discretisation

The duct calculations use a 3D finite element discretisation

$$\phi = \sum_{j=1}^n \phi_j N_j, \quad (3.6)$$

with the tri-quadratic shape functions  $N_j$  being used also as the finite element test functions. The incoming duct mode at  $x = 1$  is specified using a numerical technique first proposed by Astley [2] for the stagnant mean flow case, and then generalised by

Parrett and Eversman [23] for the case with uniform mean flow. The general solution at  $x=1$  is expressed as

$$\phi = \sum_{m=1}^{nev} [a_m^+ + a_m^-] \phi_m, \quad (3.7)$$

with normal derivative

$$\frac{\partial \phi}{\partial x} = i \sum_m^{nev} [k_m^+ a_m^+ + k_m^- a_m^-] \phi_m, \quad (3.8)$$

where  $a_m^- = 1$  for the input mode and  $a_m^- = 0$  for all other modes. The number  $nev$  of eigenmodes in the summation is chosen to include all of the cut-on modes and the leading four cut-off modes, and the eigenmodes  $\phi_m$  are computed through a finite element approximation (see for example [2] or [27]).

The insertion of (3.8) into the weak form of the general harmonic flow equations gives an under-determined system with an extra  $nev$  degrees of freedom. A corresponding number of additional equations are introduced by considering the residual form of (3.7) with respect to the weighting functions  $\phi_m$ . These extra equations are used to eliminate the coefficients  $a^+$ , giving the final system of equations

$$A \phi \equiv \left( K - \omega^2 M + i\omega C + B \right) \phi = \mathbf{f}, \quad (3.9)$$

where  $K$  is symmetric positive semi-definite,  $C$  is skew-symmetric,  $M$  is symmetric positive definite, and  $B$  is a complex matrix arising from the fan-face and far-field boundary conditions.

The iterative solution is again preconditioned using

$$P = K + \omega^2 M, \quad (3.10)$$

approximately implemented through one or more multigrid cycles. Because  $C$  is not symmetric, QMR-SYM is no longer an appropriate option and thus the more general QMR algorithm is used instead.

The axisymmetric aircraft inlet problem first requires the computation of the steady flow field. This uses the standard 2D finite element approximation yielding a non-linear system of equation. These are solved by Newton iteration, each step of which requires the solution of a linear system of equations involving a matrix which is symmetric positive definite. In 3D, these could be efficiently solved using multigrid, or multigrid-preconditioned CG, but for the axisymmetric model problem we use direct solution methods.

The acoustic solution for the inlet problem is represented as

$$\phi = \sum_{j=1}^n e^{i\kappa\theta} \phi_j N_j, \quad (3.11)$$

with bi-quadratic shape functions  $N_j$ . Using  $e^{i\kappa\theta} N_j$  as the test functions, and implementing the boundary conditions in a similar manner to the duct problems, yields a discrete linear system of equations of exactly the same form as for the duct problems.

### 3.3 3D duct results

Figure 9 shows the  $(y, z)$  cross-section through the grids for the two duct problems; uniform spacing is used in the  $x$  direction. The steady flow is prescribed to have velocity  $u = 0.25$  and speed of sound  $c = 1$ , giving a Mach number of 0.25.

Tables 2 and 3 show the number of iterations required for both duct problems. The isotropic grid is used for the rectangular duct, while the anisotropic grid is used for the cylindrical duct. The diagonals of the tables represent a fixed number of points per wavelength, starting at  $n^\lambda = 8$  on the main diagonal. In both cases the iteration number is only mildly dependent on grid resolution. The scalings with respect to  $\omega$  are impressive for the rectangular duct problem, but less so for the cylindrical problem. This is not surprising as the geometric multigrid is far more suited to the isotropic grid in the rectangular duct problem. Nevertheless, the savings are still considerable in both cases, especially on the finer grids.

In Table 4 it is shown that in the cylindrical duct case increasing the number of multigrid cycles in each  $P^{-1}$  approximate solution brings the number of QMR iterations back down to a level which scales very well with respect to  $\omega$ . Of course using such an approach would prove computationally expensive but it does illustrate that the preconditioning matrix is still effective, even though the approximate implementation is not. This is further emphasised in Figure 10 where the rectangular problem is solved using both the isotropic and anisotropic grids. It is clearly shown that the multigrid convergence on the isotropic grid is far more rapid than on the anisotropic grid. This suggests that further research is needed on more effective multigrid implementations for anisotropic grids, possibly based on algebraic multigrid methods [28].

### 3.4 Axisymmetric inlet results

The inlet grid shown in Figure 11 is used on both the steady and acoustic problems. It is designed to keep the aspect ratio of the cells as close to one as possible to improve the multigrid performance. The mass-flux parameter value  $\beta = 0.38$  ensures a fan Mach number of 0.35.

Table 5 gives the frequency and circumferential wavenumbers for three different problems, together with the approximate number of points per axial wavelength on the three grids which have been used; Grid A has 4961 nodes, Grid B 19521 nodes and Grid C 77441. Such conditions are representative of an ‘engine order’ investigation. The iterative results are given in Table 6 and show similarities to those of the cylindrical duct problem, a low dependence on grid resolution but a strong dependence on  $\omega$ . Of course it would again be hoped that a more efficient implementation of multigrid would improve the scaling with respect to frequency. This application highlights the difficulties associated with solving these problems at high frequencies. Without preconditioning, the linear equations are so ill-conditioned that the iterative solver fails to converge, whilst with the preconditioning the improvement in the system’s spectrum allows convergence

at a grid independent rate.

The issues connected to the grid type and multigrid merit further discussion. As seen the geometric multigrid is most efficient on an isotropic grid, but in practice it is hard to construct such grids. For a more general grid a possible solution to these difficulties is using algebraic multigrid [28], which requires no geometric information but instead exploits information within the coefficient matrix, with no need to tune any parameters. Its suitability to irregular finite element meshes has seen a recent surge in its popularity (see for example [24]), and it presents the most logical and promising way forward for this work.

## 4 Conclusions

This paper has introduced a novel preconditioner for the iterative solution of the Helmholtz and other similar equations.

For the 1D Helmholtz equation it has the excellent properties of giving a condition number which is independent of grid resolution, and considerable clustering of the eigenvalues. Furthermore, these features are largely retained in an efficient approximate implementation using a single multigrid cycle. When coupled with the QMR-SYM algorithm on a simple 2D Helmholtz problem the results are excellent, showing significant cost savings over both the unpreconditioned and SSOR preconditioned cases, with a number of iterations independent of grid resolution and scaling approximately linearly with the reduced frequency. QMR-SYM itself converges in a very similar manner to the optimal GMRES algorithm but without its substantial memory requirements.

These results gave the confidence to generalise the preconditioner for wave problems with a moving mean flow, using the QMR iterative solver. The grid independent convergence is retained, but for anisotropic grids the geometric multigrid is less effective, resulting in a poorer scaling with frequency than hoped for. Nevertheless, all of the results show good computational savings at the troublesome higher frequencies, and in the case of an axisymmetric inlet, convergence is achieved when not previously obtained. It is thought that future work using algebraic multigrid could solve the problems arising on more general grids, creating equally as impressive results on all grid types, and providing a solution to the problem of solving the large systems of equations arising from 3D high-frequency aeroacoustic applications.

## Acknowledgements

This research has been jointly sponsored by the Engineering and Physical Sciences Research Council and Rolls-Royce Plc. through a CASE studentship for the first author.

## References

- [1] W. E. Arnoldi, 1951. The principle of minimized iterations in the solution of the matrix eigenvalue problem. *Quarterley of Applied Mathematics*, **9**: 17–29.
- [2] R. J. Astley, 1983. Wave envelope and infinite elements for acoustical radiation. *International Journal for Numerical Methods in Fluids*, **3**: 507–526.
- [3] A. Bayliss, C. Goldstein and E. Turkell, 1985. On accuracy conditions for the numerical computation of waves. *Journal of Computational Physics*, **59**: 396–404.
- [4] A. Bayliss, C. I. Goldstein and E. Turkel, 1983. An iterative method for the Helmholtz equation. *Journal of Computational Physics*, **49**: 443–457.
- [5] A. Bayliss and E. Turkel, 1982. Far field boundary conditions for compressible flows. *Journal of Computational Physics*, **48**: 182–199.
- [6] J. H. Bramble, J. E. Pasciak and J. Xu, 1988. The analysis of multigrid algorithms for nonsymmetric and indefinite problems. *Mathematics of Computation*, **51**: 389.
- [7] S. C. Eisenstat, 1981. Efficient implementation of a class of preconditioned conjugate gradient methods. *SIAM Journal of Scientific and Statistical computing*, **2**: 1–4.
- [8] H. C. Elman, O. G. Ernst and D. P. O’Leary, 2001. A multigrid method enhanced by Krylov subspace iteration for discrete Helmholtz equations. *SIAM Journal on Scientific Computing*, **23**: 1291–1315.
- [9] H. C. Elman and D. P. O’Leary, 1998. Efficient iterative solution of the three-dimensional Helmholtz equation. *Journal of Computational Physics*, **142**(1): 163–181.
- [10] O. Ernst, 1994. *Fast Numerical Solution of Exterior Helmholtz Problems with Radiation Boundary Condition by Imbedding*. PhD thesis, Stanford University.
- [11] O. Ernst and G. H. Golub, 1994. A domain decomposition approach to solving the Helmholtz equation with a radiating boundary condition. In *Domain Decomposition in Science and Engineering*, pages 177–192. American Mathematical Society, Providence.
- [12] R. P. Fedorenko, 1964. The speed of convergence of one iterative process. *Computational Mathematics and Mathematical Physics*, **4**: 3.
- [13] R. W. Freund and N. Nachtigal, 1991. QMR: A quasi-minimal residual method for non-Hermitian linear systems. *Numerische Mathematik*, **60**: 315–339.
- [14] R.W. Freund, 1992. Conjugate gradient-type methods for linear systems with complex symmetric coefficient matrices. *SIAM Journal on Scientific and Statistical Computing*, **13**(1): 425–448.

- [15] J. Gozani, A. Nachson and E. Turkel, 1984. Conjugate gradient coupled with multigrid for an indefinite problem. *Advances in Computer Methods for Partial Differential Equations*, **5**: 425–427.
- [16] M. R. Hestenes and E. Steifel, 1952. Methods of conjugate gradients for solving linear systems. *Journal of Research of the National Bureau of Standards*, **49**: 409–436.
- [17] Campbell S. L., Ipsen I. C. F., Kelley C. T. and Meyer C. D., 1996. GMRES and the minimal polynomial. *BIT*, **36**: 664–675.
- [18] A.L. Laird, 2004. *A Hybrid Spectral Discretisation and Iterative Solution Methods for Acoustic Models in Potential Flow*. PhD thesis, University of Oxford.
- [19] C. Lanczos, 1950. An iteration method for the solution of the eigenvalue problem of linear differential and integral operators. *Journal of Research of the National Bureau of Standards*, **45**: 255–282.
- [20] M. M. M. Made, 2001. Incomplete factorized-based preconditioning for solving the Helmholtz equation. *International Journal for Numerical Methods in Engineering*, **50**: 1077–1101.
- [21] M. Malhotra, R. Freund and P. Pinsky, 1997. Iterative solution of multiple radiation and scattering problems in structural acoustics using a block quasi-minimal residual algorithm. *Computational Methods in Applied Engineering*, **146**: 173–196.
- [22] K. Otto and E. Larsson, 1999. Iterative solution of the Helmholtz equation by a second-order method. *SIAM Journal on Matrix Analysis and Applications*, **21**(1): 209–229.
- [23] A.V. Parrett and W. Eversman, 1986. Wave envelope and finite element approximations for turbofan noise radiation in flight. *AIAA Journal*, **24**(5): 753–760.
- [24] C. E. Powell and D. J. Silvester, 2003. Black-box preconditioning of self-adjoint elliptic PDEs. Number 35 in *Lecture Notes in Computational Science and Engineering*. Springer.
- [25] A. Ramage, 1999. A multigrid preconditioner for stabilised discretisations of advection-diffusion problems. *Journal of Computational and Applied Mathematics*, **101**: 187–203.
- [26] B.A. Regan and J.A. Eaton, 1999. Modelling the influence of acoustic liner non-uniformities on duct modes. *Journal of Sound and Vibration*, **219**(5): 859–879.
- [27] I. Danda Roy and W. Eversman, 1995. Improved finite element modelling of the turbofan engine inlet radiation problem. *Journal of Vibration and Acoustics*, **117**: 109–115.

- [28] J. W. Ruge and K. Stüben, 1985. Efficient solution of finite difference and finite element equations by algebraic multigrid (AMG). In D. J. Paddon and H. Holstein, editors, *Multigrid Methods for Integral and Differential Equations*. Clarendon Press, Oxford.
- [29] Y. Saad, 1982. The Lanczos biorthogonalisation algorithm and other oblique projection methods for solving large unsymmetric systems. *SIAM Journal on Numerical Analysis*, **19**: 485–506.
- [30] Y. Shapira, 1997. Multigrid techniques for highly indefinite equations. In *Seventh Copper Mountain Conference on Multigrid Methods, Volume 6, Electronic Transactions in Numerical Analysis*.
- [31] D. J. Silvester, H. E. Elman, D. Kay and A. J. Wathen, 2001. Efficient preconditioning of the linearized Navier-Stokes equations. *Journal of Computational and Applied Mathematics*, **128**: 261–279.
- [32] D. Stanescu, M. Y. Hussani and F. Farassat, 2002. Large scale frequency domain numerical simulation of aircraft tone noise radiation and scattering, aiaa paper 2002-2586. In *Proceedings of the 8th AIAA/CEAS Aeroacoustics Conference, Breckenridge, Colorado*.
- [33] L. N. Trefethen and D. Bau III, 1997. *Numerical Linear Algebra*. SIAM.

$$\begin{array}{c}
 \frac{d\phi}{dx} = i\omega\phi \quad \left| \begin{array}{c} -\frac{d^2\phi}{dx^2} - \omega^2\phi = 0 \\ \hline \frac{d\phi}{dx} = i\omega \end{array} \right. \quad \left| \frac{d\phi}{dx} = i\omega \right. \\
 0 \qquad \qquad \qquad 1 \\
 \\
 \begin{array}{c}
 y \\
 \uparrow \\
 1 \\
 \hline \\
 \hline \\
 \hline \\
 0 \\
 \hline \\
 0 \qquad \qquad \qquad 1 \\
 \left. \frac{\partial\phi}{\partial x} = i\sqrt{\pi^2 - \omega^2}\phi \right| \begin{array}{c} \frac{\partial\phi}{\partial y} = 0 \\ \hline -\nabla^2\phi - \omega^2\phi = 0 \\ \hline \frac{\partial\phi}{\partial x} = \cos(\pi y) \\ \hline \frac{\partial\phi}{\partial y} = 0 \end{array} \right| x
 \end{array}
 \end{array}$$

Figure 1: The test problems in 1D (above), and 2D (below).

$\omega \backslash h$	1/8	1/16	1/32	1/64	1/128
6	12(12)	13(13)	15(15)	15(15)	17(*)
12		19(19)	19(19)	21(20)	20(*)
24			27(27)	30(28)	33(*)
48				46(46)	51(*)

Table 1: The number of QMR-SYM (GMRES) iterations for the 2D preconditioned problem using bi-linear elements. A ‘\*’ indicates insufficient memory resources, whilst the blanks indicate under-resolved combinations.

$\omega \backslash h$	1/8	1/16	1/32	1/64
5	55(11)	82(13)	139(13)	256(15)
10		116(16)	191(18)	435(22)
20			477(30)	966(35)
40				2986(63)

Table 2: The number of QMR iterations for the unpreconditioned (preconditioned) rectangular duct problem with tri-quadratic elements.



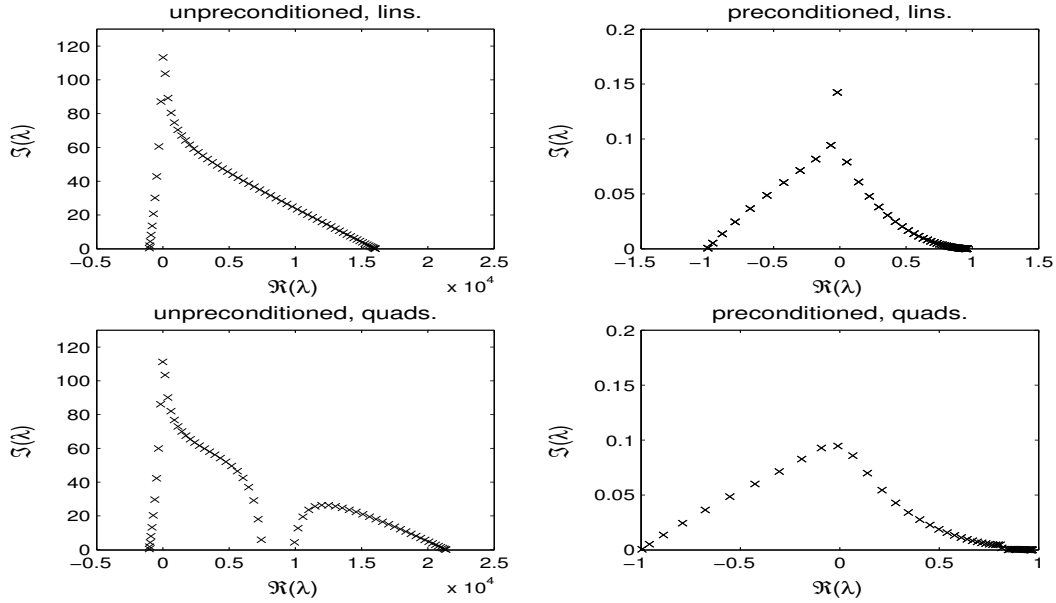


Figure 2: The effects of preconditioning on the eigenvalues for linear and quadratic finite element approximations of the 1D problem, with  $n^\lambda = 12.5$ ,  $\omega = 32$  and  $P$  implemented exactly.

$\omega \backslash h$	1/8	1/16	1/32	1/64
5	48(17)	114(19)	220(16)	406(17)
10		144(43)	317(52)	690(62)
20			941(271)	1930(220)
40				5288(732)

Table 3: The number of QMR iterations for the unpreconditioned (preconditioned) cylindrical duct problem with tri-quadratic elements.

$\omega$	Number of multigrid cycles							
	1	3	5	7	9	11	13	15
5	16	12						
10	52	42	28	18				
20	271	171	93	62	35	31	29	28

Table 4: The number of QMR iterations required for convergence of the cylindrical duct problem for different numbers of multigrid cycles, where  $h = 1/32$  and using tri-quadratic elements.

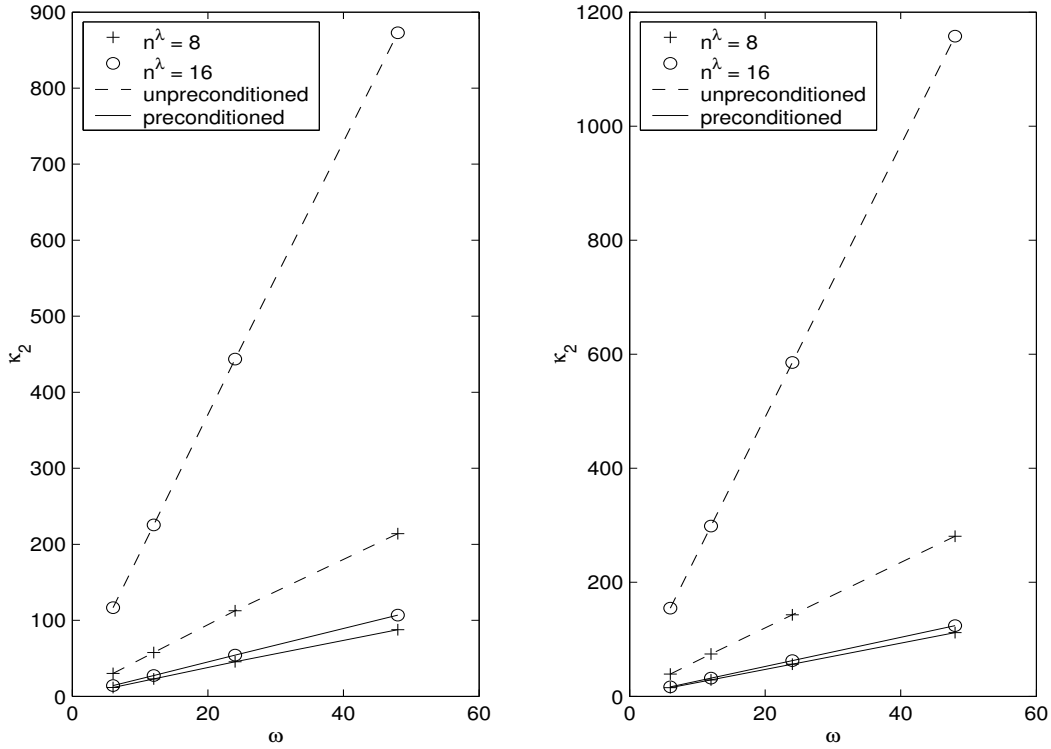


Figure 3: The effects on the 1D condition number of the preconditioning for linear elements (left), and quadratic elements (right), with  $P$  implemented exactly.

$\omega$	$\kappa$	A	B	C
15	13	10	20	40
30	26		10	20
60	52			10

Table 5: The number of points per wavelength for the inlet problem on grids A,B and C.

$\omega$	A	B	C
15	1356(457)	2722(647)	5711(613)
30		-(2584)	-(3328)
60			-(12591)

Table 6: The number of QMR iterations for the unpreconditioned (preconditioned) inlet problem with bi-quadratic elements. ‘-’ indicates a failure to converge in  $n$  iterations.




Figure 4: The effects on the 1D condition number of the exact (LU) and approximate multigrid (MG) preconditioning for linear elements (left), and quadratic elements (right).




Figure 5: Convergence of the relative residual using GMRES and QMR-SYM iteration applied to the 2D preconditioned problem with  $\omega = 24$  and  $h = 1/32$ .




Figure 6: Cost savings from the new preconditioner in 2D when using QMR-SYM.




Figure 7: Performance comparison of the new and SSOR preconditioners in 2D for  $n^\lambda = 8$  (left) and  $n^\lambda = 16$  (right).




Figure 8: The axisymmetric domain for the inlet problem.




Figure 9:  $(y, z)$  cross-sections through the  $h = 1/16$  grids used for the 3D duct problems, referred to as rectangular isotropic (left), rectangular anisotropic (middle) and cylindrical anisotropic (right).




Figure 10: Graph showing the QMR iterative convergence for the rectangular duct problem ( $\omega = 10$ ,  $u = 0.25$  and  $h = 1/16$ ) on the two types of grid.




Figure 11: Grid A for the inlet problem; with bi-quadratic elements, each cell has nodes at the corners, mid-edges, and cell centre.