BLOCK PRECONDITIONERS FOR A COUPLED PRESSURE-TEMPERATURE MODEL OF TRACE GAS SENSORS

BRIAN BRENNAN * AND ROBERT C. KIRBY †

Abstract. Trace gas sensors are currently used in many applications from leak detection to national security and may some day help with disease diagnosis. These sensors are modelled by a coupled system of complex elliptic partial differential equations for pressure and temperature. Solutions are approximated using the finite element method which yields a skew-Hermitian dominant discretization for which classical algebraic preconditioners quickly degrade. We develop a block preconditioner that requires scalar Helmholtz solutions to apply but gives a very low outer iteration count. We also present analysis showing eigenvalues of the preconditioned system are mesh-dependent but with a small coefficient. Numerical experiments confirm our theoretical discussion.

Key words. block preconditioners, finite element, multiphysics, trace gas sensors

AMS subject classifications. 65F08, 65N30, 97M50

1. Introduction. Currently, research on trace gas sensors is focused on the development of portable, efficient, and cost-effective sensor technologies that can be deployed in networks for large scale monitoring of carbon dioxide and atmospheric pollutants, as well as for non-invasive disease diagnosis using breath analysis [14, 16, 27]. One such trace gas sensor is the Quartz-Enhanced Photo-Acoustic Spectroscopy (QEPAS) sensor which employs a quartz tuning fork to detect the weak acoustic pressure waves generated by the interaction of laser radiation with a trace gas [13]. More specifically, QEPAS sensors are based on the following physical mechanisms. A laser generates optical radiation at a specific absorption wavelength of the gas to be detected. The laser beam is directed between the tines of the tuning fork. The optical energy that is absorbed by the trace gas is transformed into vibrational energy of the gas molecules, generating a temperature disturbance. If the interaction between the laser and the trace gas is sinusoidally modulated, this temperature disturbance is in the form of a thermal wave. In addition, vibrational to translational energy conversion processes in the gas molecules result in the generation of a weak acoustic pressure wave, which can be detected by the tuning fork. To amplify the signal detected by the tuning fork, the modulation frequency of the laser is chosen so as to excite a resonant vibration in the tuning fork. Finally, since quartz is a piezoelectric material, this mechanical vibration is converted to an electric current that can be measured. Because the entire process is linear, the measured current is proportional to the concentration of the trace gas. In some experimental regimes, the tuning fork can also be used to directly detect the thermal wave via the pyroelectric effect [20].

Most mathematical models of QEPAS sensors [8, 19, 25, 26] have included damping in the model using an *ad hoc* approach that involves making experimental measurements using the actual tuning fork being modeled. In [2], we began developing a more realistic model of the damping in a QEPAS sensor based on a coupled pressure-temperature system derived by Morse and Ingard [17]. This model of a QEPAS sensor incorporates the effects of viscous damping using a parameter that depends on the physical properties of the fluid and not the particular details the tuning fork. By elim-

*Department of Mathematics, Baylor University, One Bear Place #97328, Waco, TX 76798-7328 (Email: b_brennan@baylor.edu). This work was supported by NSF grant CCF-1325480,

†Department of Mathematics, Baylor University, One Bear Place #97328, Waco, TX 76798-7328 (Email: robert_kirby@baylor.edu). This work was supported by NSF grant CCF-1325480

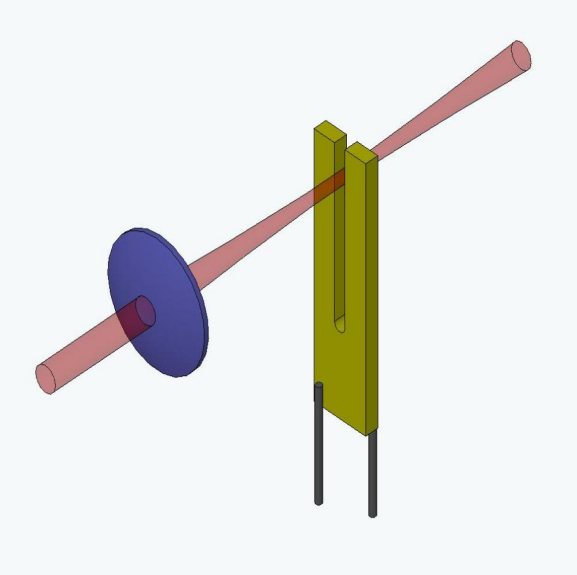

Fig. 1: Schematic diagram of the experimental setup for a QEPAS sensor showing the tuning fork (the U-shaped bar with two tines), two attached wires, and the laser source focused between the tines of the tuning fork.

inating these *ad hoc* parameters, we hope not only to obtain a higher-fidelity model, but also one more amenable to eventual design optimization of the tuning fork.

In this paper we present an effective block preconditioner for the pressure-temperature equations, which takes advantage of the problem structure and the particular regime of realistic physical parameters. We show that the cost applying the preconditioner is dominated by solving a pair of scalar indefinite Helmholtz problems with high wave number. Although this problem is quite challenging for classical iterative methods [6], thanks to recent progress on techniques such as complex shifts [5] this is at least tractable and far better than the original coupled system. In Section 2, we present our coupled mathematical model and its corresponding variational form and finite element discretization. Section 3 begins with a survey of some standard preconditioning methods and how they perform on our problem with various sets of physical data which motivates our new block preconditioner. Here we also give an eigenvalue bound which is mesh dependent but scaled by a very small constant due to the extreme data values. In Section 4 we describe how to apply our preconditioner in terms of solving scalar Helmholtz solves. In Section 5, we validate our implementation against a known plane wave solution to the coupled system, and also report on the performance of our preconditioner with respect to physical data and mesh refinement. In Section 6 we discuss concluding remarks and ideas for future work.

2. Mathematical Model and Discretization. The interaction of laser radiation with the trace gas generates an acoustic pressure wave, $\mathcal{P}$, and a thermal disturbance, $\mathcal{T}$. To model the effects of viscous damping and thermal conduction in the gas, Morse and Ingard derived a coupled system of pressure-temperature equations which generalizes the standard acoustic wave and heat equations. These equations

are given by

$$\frac{\partial}{\partial t}\left(\mathcal{T} - \frac{\gamma-1}{\gamma\alpha}\mathcal{P}\right) - \ell_h c\Delta\mathcal{T} = \mathcal{S}, \tag{2.1a}$$

$$\gamma\left(\frac{\partial^2}{\partial t^2} - \ell_v c\frac{\partial}{\partial t}\Delta\right)(\mathcal{P} - \alpha\mathcal{T}) - c^2\Delta\mathcal{P} = 0. \tag{2.1b}$$

Here ℓ_v and ℓ_h are characteristic lengths associated with the effects of fluid viscosity and thermal conduction, respectively, c is sound speed, γ is the ratio of the specific heat of the gas at constant pressure to that at constant volume, and $\alpha = \left(\frac{\partial P}{\partial T}\right)_v$ is the rate of change of ambient pressure with respect to ambient temperature at constant volume.

We model the interaction between the laser and the trace gas using the source term, $\mathcal{S}$, in equation (2.1a) given by

$$\mathcal{S}(x,t) = C\exp\left(-\frac{2[(x-x_s)^2 + (z-z_s)^2]}{\sigma^2}\right)\exp(-i\omega t) \tag{2.2}$$

where C is a constant that is proportional to the concentration of the trace gas to be detected, (x_s, z_s) are the coordinates of the axis of the cylindrically symmetric Gaussian power profile of the laser beam, σ is the beam width and ω is the frequency of the periodic interaction between the laser radiation and the trace gas. The modulation frequency, ω, is chosen so as to excite a resonant vibration in the tuning fork. Since $\mathcal{S}(\mathbf{x},t) = S(\mathbf{x})\exp(-i\omega t)$ is periodic in time, so are the pressure and temperature waves. Substituting $\mathcal{P}(\mathbf{x},t) = P(\mathbf{x})\exp(-i\omega t)$ and $\mathcal{T}(\mathbf{x},t) = T(\mathbf{x})\exp(-i\omega t)$ into equations (2.1a) and (2.1b) we obtain the coupled system of Helmholtz-type equations

$$-i\beta\omega\left(T - \frac{\gamma-1}{\gamma\alpha}P\right) - \beta\ell_h c\Delta T = S, \tag{2.3a}$$

$$-\gamma(\omega^2 - i\ell_v c\omega\Delta)(P - \alpha T) - c^2\Delta P = 0, \tag{2.3b}$$

where $\beta = \frac{\alpha^2\gamma^2\omega}{\gamma-1}$ has been chosen to maximize skewness on the antidiagonal in the operator as we see in (2.4). Since critical pieces of our numerical software stack do not currently handle complex arithmetic, we split the 2×2 complex-valued problem into a 4×4 system with real variables. Setting $T = T_1 + iT_2$, and $P = P_1 + iP_2$ while scaling the second of equations (2.3) by $-i$, we obtain a system of four partial differential equations of the form $Au = b$, where

$$\begin{aligned} A &= \begin{pmatrix} -\beta\ell_h c\Delta & \beta\omega & 0 & -\alpha\gamma\omega^2 \\ -\beta\omega & -\beta\ell_h c\Delta & \alpha\gamma\omega^2 & 0 \\ \alpha\gamma\ell_v c\omega\Delta & -\alpha\gamma\omega^2 & -\gamma\ell_v c\omega\Delta & \gamma\omega^2 + c^2\Delta \\ \alpha\gamma\omega^2 & \alpha\gamma\ell_v c\omega\Delta & -(\gamma\omega^2 + c^2\Delta) & -\gamma\ell_v c\omega\Delta \end{pmatrix} \\ &= \begin{pmatrix} -a_1\Delta & a_2 & 0 & -a_3 \\ -a_2 & -a_1\Delta & a_3 & 0 \\ a_4\Delta & -a_3 & -a_5\Delta & a_6 + a_7\Delta \\ a_3 & a_4\Delta & -(a_6 + a_7\Delta) & -a_5\Delta \end{pmatrix}. \end{aligned} \tag{2.4}$$

We should note that this model holds in all spatial dimensions.

The performance of any numerical method will be heavily dependent on the given physical data. Our goal is to solve the system using realistic parameters, such as

$$
\begin{aligned}
&\ell_h = 1.6 \cdot 10^{-6} \text{ m} && a_1 = \beta \ell_h c = 381.37 \\
&\ell_v = 10^{-6} \text{ m} && a_2 = \beta\omega = 4.19 \cdot 10^{11} \\
&c = 300 \text{ m/s} && a_3 = \alpha\gamma\omega^2 = 1.35 \cdot 10^{10} \\
&\omega = 33,000 \text{ Hz} && a_4 = \alpha\gamma\ell_v c\omega = 12.29 \\
&\gamma = 1.4 && a_5 = \gamma\ell_v c\omega = 1.38 \\
&\alpha = 8.8667 \text{ Pa/K} && a_6 = \gamma\omega^2 = 1.52 \cdot 10^9 \\
& && a_7 = c^2 = 9.0 \cdot 10^4.
\end{aligned} \tag{2.5}
$$

This PDE system admits a standard variational form where we seek $u \in V \subset H^1$, with appropriate boundary conditions enforced, such that

$$a(u, v) = L(v) \qquad \forall v \in \widehat{V}, \tag{2.6}$$

where V is the trial space and $\widehat{V}$ is the test space [1, 3, 12]. In our case,

$$
\begin{aligned}
a(u,v) = {} & \beta\ell_h c\langle \nabla T_1, \nabla v_1\rangle + \beta\omega\langle T_2, v_1\rangle - \alpha\gamma\omega^2\langle P_2, v_1\rangle - \beta\omega\langle T_1, v_2\rangle \\
& + \beta\ell_h c\langle \nabla T_2, \nabla v_2\rangle + \alpha\gamma\omega^2\langle P_1, v_2\rangle - \alpha\gamma\ell_v c\omega\langle \nabla T_1, \nabla v_3\rangle - \alpha\gamma\omega^2\langle T_2, v_3\rangle \\
& + \gamma\ell_v c\omega\langle \nabla P_1, \nabla v_3\rangle + \gamma\omega^2\langle P_2, v_3\rangle - c^2\langle \nabla P_2, \nabla v_3\rangle + \alpha\gamma\omega^2\langle T_1, v_4\rangle \\
& - \alpha\gamma\ell_v c\omega\langle \nabla T_2, \nabla v_4\rangle - \gamma\omega^2\langle P_1, v_4\rangle + c^2\langle \nabla P_1, \nabla v_4\rangle + \gamma\ell_v c\omega\langle \nabla P_2, \nabla v_4\rangle,
\end{aligned}
$$

and $L(v) = \langle S, v_1\rangle$.

On a finite-dimensional subspace $V_h \subseteq V$, the discrete approximation $u_h \in V_h$ is represented by

$$u_h(x) = \sum_{i=1}^{N} \eta_i \phi_i(x), \qquad x \in \Omega \tag{2.7}$$

where $\phi_i(x)$ is a continuous piecewise polynomial over the mesh. We will focus solely on P^1 elements. Our discretization will give linear systems with a natural block structure in terms of the mass and stiffness matrices, which are

$$K_{ij} = \int_\Omega \nabla\phi_i \cdot \nabla\phi_j dx, \tag{2.8}$$

and

$$M_{ij} = \int_\Omega \phi_i \phi_j dx. \tag{2.9}$$

By using mass lumping or equivalently, a vertex-based quadrature rule, M becomes diagonal.

In terms of M and K, the linear system arising from the finite element discretization has the form

$$A_h = \begin{pmatrix} a_1K & a_2M & 0 & -a_3M \\ -a_2M & a_1K & a_3M & 0 \\ -a_4K & -a_3M & a_5K & a_6M - a_7K \\ a_3M & -a_4K & a_7K - a_6M & a_5K \end{pmatrix}. \tag{2.10}$$

3. Preconditioning. Our goal is to develop an effective preconditioner that handles the block structure and severe physical parameters of our problem. Though a bit of a straw man for coupled systems, we will report on the standard incomplete LU factorization applied to the full, unblocked matrix. Besides this, we will consider two common block preconditioners based on the block diagonal (Jacobi) or block lower triangular (Gauss-Seidel) parts. We deal with a 2×2 partitioning of the system that respects the origin of the PDE as complex-valued. The block Jacobi preconditioner we consider is

$$P_{Jac} = \left(\begin{array}{cc:cc} a_1K & a_2M & 0 & 0 \\ -a_2M & a_1K & 0 & 0 \\ \hdashline 0 & 0 & a_5K & a_6M - a_7K \\ 0 & 0 & a_7K - a_6M & a_5K \end{array}\right), \tag{3.1}$$

and the block Gauss-Seidel is

$$P_{GS} = \left(\begin{array}{cc:cc} a_1K & a_2M & 0 & 0 \\ -a_2M & a_1K & 0 & 0 \\ \hdashline -a_4K & -a_3M & a_5K & a_6M - a_7K \\ a_3M & -a_4K & a_7K - a_6M & a_5K \end{array}\right). \tag{3.2}$$

Blocking with respect to the 4×4 system would be simpler, but less powerful.

In Tables 1 and 2 we compare the iteration counts of the *Generalized Minimal Residual Method* (GMRES) with no preconditioner, ILU(3) preconditioning, and the block Jacobi and block Gauss-Seidel preconditioners for a simple data set and the realistic data given by (2.5). Since we are simply testing the iteration count rather than trying to implement a practical method, we apply the preconditioners (3.1) and (3.2) with a direct method. In each case, we use unrestarted GMRES and iterate to a Euclidean norm tolerance of 10^{-10}.

For Table 1, we set the data to $\gamma = 1.5$, $\ell_h = 0.0016$, $\ell_v = 0.001$ and c, ω, α all to 1.0. Even with this simple data, unpreconditioned GMRES fails to converge as we refine the mesh. ILU(3) performs perhaps as expected, with the iteration count growing as the mesh is refined. We also see that the block preconditioners both give mesh-independent and very low iteration counts, so that deriving a practical algorithm for applying them would give a very powerful preconditioner indeed.

The same situation does not hold for the physical data of interest (2.5). The iteration counts for ILU(3) are a bit larger than on the simple data. Although the block preconditioners seem to have a mesh-independent bound on the iteration count, that these counts are so much larger motivates studying a preconditioner that respects the particular parameter regime of physical interest.

The effects of fluid viscosity and thermal conduction described by the characteristic lengths ℓ_v and ℓ_h, respectively, are particularly small in magnitude relative to all

N	GMRES	ILU(3)	Block Jacobi	Block Gauss-Seidel
32	1329	12	6	4
64	4514	21	6	4
128	10000+	37	6	4
256		69	6	4
512		122	6	4

Table 1: Iteration counts for the unpreconditioned GMRES method and GMRES coupled with the block Jacobi, block Gauss Seidel and Incomplete LU factorization preconditioners for simple data $\gamma = 1.5$, $\ell_h = 0.0016$, $\ell_v = 0.001$ and c, ω, α all set to 1.0.

N	GMRES	ILU(3)	Block Jacobi	Block Gauss-Seidel
32	1156	12	269	141
64	4320	14	291	154
128	10000+	44	257	143
256		235	245	134
512		429	231	127

Table 2: Iteration counts for the unpreconditioned GMRES method and GMRES coupled with the block Jacobi, block Gauss Seidel and Incomplete LU factorization preconditioners using realistic physical data.

other physical data. If we ignore these effects in the original model (2.3), we find the simplified system

$$-i\beta\omega\left(T - \frac{\gamma - 1}{\gamma\alpha}P\right) = S, \tag{3.3a}$$

$$-\gamma\omega^2(P - \alpha T) - c^2\Delta P = 0. \tag{3.3b}$$

If we solve $(3.3a)$ for either T or P and substitute this into $(3.3b)$, the system decouples to the standard inhomogeneous acoustic wave equation of a single variable. So we are essentially using the classical wave equation to precondition the coupled system of modified wave equations. Discretizing (3.3) gives the following block preconditioner,

$$P = \begin{pmatrix} 0 & a_2M & 0 & -a_3M \\ -a_2M & 0 & a_3M & 0 \\ 0 & -a_3M & 0 & a_6M - a_7K \\ a_3M & 0 & a_7K - a_6M & 0 \end{pmatrix}. \tag{3.4}$$

A tedious application of block Gaussian elimination shows that P is in fact in-

vertible, and we can calculate

(3.5)
$$P^{-1}A_h = \begin{pmatrix} I & -\hat{a}_1 M^{-1}K + \hat{a}_2 H^{-1}K & 0 & \hat{a}_4 H^{-1}K \\ \hat{a}_1 M^{-1}K - \hat{a}_2 H^{-1}K & I & -\hat{a}_4 H^{-1}K & 0 \\ 0 & \hat{a}_3 H^{-1}K & I & \hat{a}_5 H^{-1}K \\ -\hat{a}_3 H^{-1}K & 0 & -\hat{a}_5 H^{-1}K & I \end{pmatrix}$$

where $H = K - \kappa^2 M$ and the coefficients have been consolidated into

(3.6)
$$\hat{a}_1 = \frac{a_1}{a_2} = 9.091 \cdot 10^{-10} \qquad \hat{a}_4 = \frac{a_3 a_5}{a_2 a_7} = 7.443 \cdot 10^{-7}$$
$$\hat{a}_2 = \frac{a_3 a_4}{a_2 a_7} = 6.6 \cdot 10^{-6} \qquad \hat{a}_5 = \frac{a_5}{a_7} = 2.31 \cdot 10^{-5}$$
$$\hat{a}_3 = \frac{a_4}{a_7} - \frac{a_1 a_3}{a_2 a_7} = 6.827 \cdot 10^{-5} \qquad \kappa = \sqrt{\frac{a_6}{a_7} - \frac{a_3^3}{a_2 a_7}} = 110.0.$$

Here κ is the Helmholtz wave number.

We begin our analysis of this preconditioner with a simple comparison of the 2-norm condition numbers of A_h and $P^{-1}A_h$ for some very coarse meshes.

N	cond(A_h)	cond($P^{-1}A_h$)
8	1.1135e+05	1.0729
16	6.7555e+04	1.0003
32	1.8726e+05	1.0005

Table 3: 2−norm condition numbers of the unpreconditioned and preconditioned matrix for various $N \times N$ meshes.

A direct calculation of the condition number of $P^{-1}A_h$, although feasible only for coarse meshes, gives an indication in Table 3 of the efficiency of our preconditioner. We will soon develop an algorithm for applying P^{-1} to a given vector b.

A fuller picture comes from investigating the eigenvalues of the preconditioned operator. Krylov methods tend to perform well if the eigenvalues are clustered away from the origin in a small region of the complex plane. Our estimate reveals mesh-dependence but the mesh dependent term is scaled by a factor that, for our physical parameters, is very small. This suggests that our preconditioner will be practical on moderate mesh sizes, although a preconditioner combining mesh independence with small iteration counts will remain an open question for this problem.

We use the classical Gershgorin circle theorem, suitably generalized to block matrices [23].

THEOREM 3.1 (Block Gershgorin). *Let A be partitioned into A_{ij}, $i, j = 1, \ldots, N$ with $\|A_{ij}\|_2$ defined as the operator norm of A_{ij}. All eigenvalues of the block partitioned matrix A are contained in the set $G = \cup_{i=1}^{N} G_i$, where G_i is the set of all $\lambda \in \mathbb{C}$ satisfying,*

(3.7)
$$\left\|(A_{ii} - \lambda I)^{-1}\right\|_2^{-1} \leq \sum_{i \neq j} \|A_{ij}\|_2 .$$

COROLLARY 3.2. *All eigenvalues* λ *of* $P^{-1}A_h$*, as defined by* 3.5*, satisfy*

$$|\lambda - 1| \leq Ch^{-2}. \tag{3.8}$$

Proof. First, we assume the mesh is sufficiently refined so that the Helmholtz matrix, H, is positive definite. Since all diagonal blocks of $P^{-1}A_h$ are identity matrices, equation (3.7) immediately reduces to

$$|\lambda - 1| \leq r \tag{3.9}$$

where

$$r = \max\left\{\hat{a_1}\left\|M^{-1}K\right\|_2 + (\hat{a_2} + \hat{a_4})\left\|H^{-1}K\right\|_2, (\hat{a_3} + \hat{a_5})\left\|H^{-1}K\right\|_2\right\}. \tag{3.10}$$

We expect $\left\|H^{-1}K\right\|_2 = \mathcal{O}(1)$ by theory of equivalent operators [7]. However, $\left\|M^{-1}K\right\| \leq \frac{\lambda_{max}(K)}{\lambda_{min}(M)} = \mathcal{O}(h^{-2})$ in all spacial dimensions. ∎

Corollary 3.2 yields a mesh dependent bound on the eigenvalues clustered around 1.0. This is offset by two key observations. First, equation (3.10) can be reduced to

$$r \approx \hat{a_1}\mathcal{O}(h^{-2}) + (\hat{a_2} + \hat{a_4})\mathcal{O}(1) \tag{3.11}$$

where $\hat{a_1} = 9.091 \cdot 10^{-10}$ and $\hat{a_2} + \hat{a_4} = 7.34 \cdot 10^{-6}$. So unless we require an extremely fine mesh, the eigenvalues are well clustered about 1.0. Second, we can see from (3.5) that $P^{-1}A_h$ is a nearly skew perturbation of the identity. That is, $P^{-1}A_h = I + E$ where $E \approx -E^T$. This tells us that even as the eigenvalues of the preconditioned system move away from 1.0, they should do so almost entirely along the imaginary axis and more importantly, not towards 0.0. Both observations are illustrated in Figure 2.

4. Implementation. For the construction of the finite element matrices we employ the Python interface to FEniCS [15]. For the sparse linear algebra operations and Krylov solvers we use Numpy and SciPy [4, 21]. Finally, we rely on the algebraic multigrid preconditioners from PyAMG for solving our embedded Helmholtz problems with a complex shift [24].

Applying the preconditioner requires solving the system

$$\begin{pmatrix} 0 & a_2M & 0 & -a_3M \\ -a_2M & 0 & a_3M & 0 \\ 0 & -a_3M & 0 & a_6M - a_7K \\ a_3M & 0 & a_7K - a_6M & 0 \end{pmatrix} \begin{pmatrix} \hat{b_1} \\ \hat{b_2} \\ \hat{b_3} \\ \hat{b_4} \end{pmatrix} = \begin{pmatrix} b_1 \\ b_2 \\ b_3 \\ b_4 \end{pmatrix}. \tag{4.1}$$

We apply blockwise Gaussian elimination to obtain the block upper triangular system

$$\begin{pmatrix} -a_2M & 0 & a_3M & 0 \\ 0 & a_2M & 0 & -a_3M \\ 0 & 0 & K - \kappa^2 M & 0 \\ 0 & 0 & 0 & K - \kappa^2 M \end{pmatrix} \begin{pmatrix} \hat{b_1} \\ \hat{b_2} \\ \hat{b_3} \\ \hat{b_4} \end{pmatrix} = \begin{pmatrix} b_2 \\ b_1 \\ \frac{1}{a_7}(b_4 + \frac{a_3}{a_2}b_2) \\ -\frac{1}{a_7}(b_3 + \frac{a_3}{a_2}b_1) \end{pmatrix}, \tag{4.2}$$

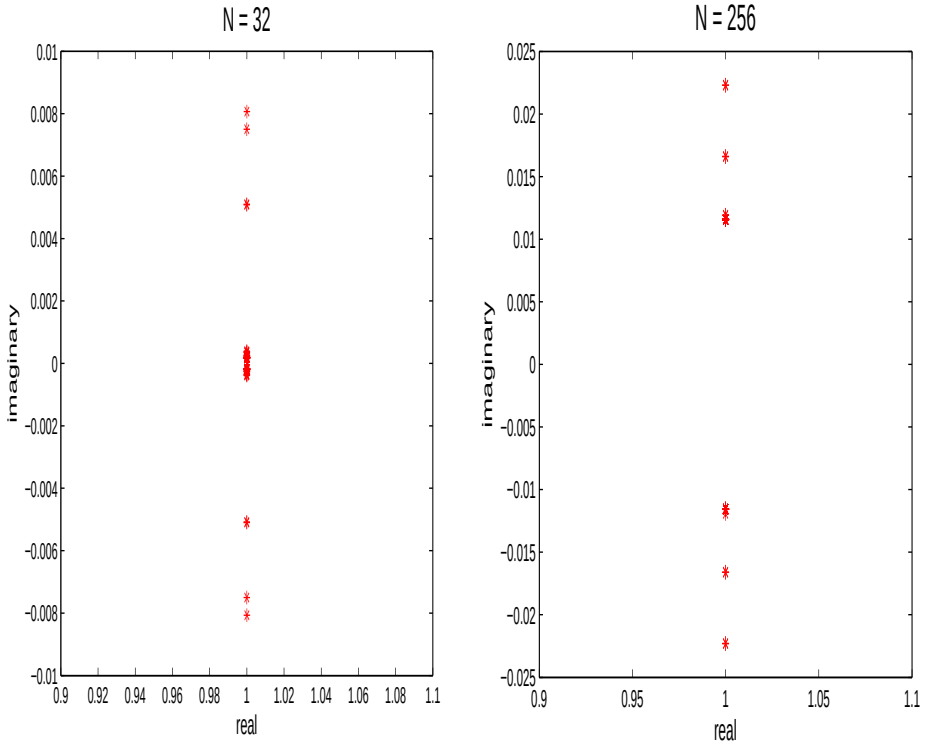

Fig. 2: A display of the 100 largest eigenvalues (in magnitude), of the preconditioned operator, $P^{-1}A_h$, for $N = 32$ on the left and $N = 256$ on the right. The smallest eigenvalues are clustered around $(1, 0)$.

where $\kappa = \sqrt{\frac{a_6}{a_7} - \frac{a_3^2}{a_2 a_7}} = 110.0$.

We will describe how we use algebraic multigrid, applied to a complex-shifted operator, for applying H^{-1}. Supposing this or some other effective Helmholtz solver is used, we have

$$\hat{b_3} = \frac{1}{a_7} H^{-1} \left(b_4 + \frac{a_3}{a_2} b_2 \right) \tag{4.3a}$$

$$\hat{b_4} = -\frac{1}{a_7} H^{-1} \left(b_3 + \frac{a_3}{a_2} b_1 \right). \tag{4.3b}$$

With a lumped mass matrix, we let $M = diag(\vec{M})$ and we can apply M or M^{-1} to a given vector by pure element-wise multiplication $(*)$ or division $(/)$, respectively.

$$\hat{b_1} = -(b_2 - a_3 \vec{M} * \hat{b_3})/(a_2 \vec{M}) \tag{4.4a}$$

$$\hat{b_2} = (b_1 + a_3 \vec{M} * \hat{b_4})/(a_2 \vec{M}). \tag{4.4b}$$

The key observation here is that the inversion of the full 4×4 block matrix P at each outer iteration has been reduced to inverting the discrete Helmholtz block, H, twice per iteration, plus some basic arithmetic.

Accuracy requirements in solving the Helmholtz problem are met by enforcing a sufficient number of mesh points per wavelength. A standard rule of thumb is $\kappa h \leq \frac{\pi}{5}$ where κ is the Helmholtz wave number [9]. For the rest of this paper we will assume this condition is met.

Even for sufficiently refined meshes, the indefinite Helmholtz equation is difficult to solve. Many preconditioners have been suggested for this problem. For this paper we apply the now common complex shift technique [5]. That is, we let $P_H = K - (\alpha + i\beta)\kappa^2 M$ where α and β are real numbers chosen such that multigrid methods are known to behave better on P_H than on H itself.

(α, β)	$\kappa = 25$	$\kappa = 50$	$\kappa = 100$
(1,0)	112	308	1000+
(0,0)	11	25	70
(1,0.5)	9	20	56

Table 4: Number of GMRES iterations for Helmholtz operator preconditioned with $P_H = K - (\alpha + i\beta)\kappa^2 M$. For each test, the number of mesh points per wavelength is kept constant by enforcing $\kappa h \leq \pi/5$.

In Table 4 we compare iteration counts for preconditioned GMRES using an algebraic multigrid preconditioner for P_H with various values of (α, β) applied to the Helmholtz equation over a range of wave numbers. By setting $(\alpha, \beta) = (1, 0)$ we are solving the unpreconditioned Helmholtz operator. As expected, GMRES fails to converge for large wave numbers. For $(\alpha, \beta) = (0, 0)$ we are simply preconditioning the Helmholtz operator with the Laplacian. While this shows great improvement over the unpreconditioned operator, optimal results are found by setting $(\alpha, \beta) = (1, 0.5)$. For all tests in this paper we will precondition the Helmholtz operator with

$$P_H = K - i(1 + 0.5i)\kappa^2 M, \tag{4.5}$$

with the widely used choice of $(\alpha, \beta) = (1, 0.5)$.

Algorithm 1 shows a brief outline of our method. Given a particular set of physical data, FEniCS is used to assemble the load vector b, the vector $\vec{M}$ containing the mass diagonal and the stiffness matrix K. Next we use Numpy, Scipy and PyAMG tools to build the global block operator A_h and the complex-shifted preconditioner P_H. For P_H, we use the smoothed aggregation solver of PyAMG with default parameters [18]. To apply the global preconditioner we require two calls to an unrestarted GMRES solver for the Helmholtz blocks. This is done at each iteration of the global unrestarted GMRES solver which returns the final solution $u = (T_1, T_2, P_1, P_2)$.

In Figure 3 we show the preconditioned residual norms of the GMRES iterations. Here we are solving $Hx = y$ with algebraic multigrid for the the complex-shifted Laplace operator as a preconditioner. Super linear convergence appears to be obtained independent of the mesh size but only after a significant number of iterations. This observation motivates future work with eigenvalue deflation [22].

Algorithm 1 Solve $A_h u = b$

Import physical data
Build FEM vectors b and $\vec{M}$ and the matrix K
Build global block operator A_h
$amg(P_H) \leftarrow$ algebraic multigrid PC for complex-shifted Laplacian

procedure $P_{inv}(b = (b_1, b_2, b_3, b_4))$ ▷ Apply P^{-1} to b
    $y_3 \leftarrow \left(b_4 - \frac{a_3}{a_2} b_2\right) / a_7$ ▷ Update RHS vectors
    $y_4 \leftarrow -\left(b_3 - \frac{a_3}{a_2} b_1\right) / a_7$

    $\hat{b_3} \leftarrow \texttt{gmres}(H, y_3,\ amg(P_H),\ tol = 10^{-10})$ ▷ Solve Helmholtz blocks
    $\hat{b_4} \leftarrow \texttt{gmres}(H, y_4,\ amg(P_H),\ tol = 10^{-10})$

    $\hat{b_1} \leftarrow -(b_2 - a_3 \vec{M} * \hat{b_3})/(a_2 \vec{M})$ ▷ Complete back-solve
    $\hat{b_2} \leftarrow (b_1 + a_3 \vec{M} * \hat{b_4})/(a_2 \vec{M})$

    $\hat{b} = (\hat{b_1}, \hat{b_2}, \hat{b_3}, \hat{b_4})$
    return $\hat{b}$ ▷ $\hat{b} \approx P^{-1} b$
end procedure

$u \leftarrow \texttt{gmres}(A_h, b,\ P_{inv},\ tol = 10^{-8})$ ▷ Solve global system

5. Results. Since our problem is highly ill-conditioned, it is important to make sure we are converging to the correct finite element solution and not succumbing to round-off error. For the purpose of testing our method, we can assume a plane wave pressure solution to equation (2.3) as presented by Morse and Ingard [17]. That is, we let

$$P(\mathbf{x}) = e^{i\mathbf{k}\cdot\mathbf{x}}, \tag{5.1}$$

where $\mathbf{k}$ is the complex wave vector. By setting $S = 0$ the temperature is also a plane wave, given by

$$T(\mathbf{x}) = \frac{i\omega(\gamma - 1)}{(i\omega - \ell_h c k^2)\gamma\alpha} e^{i\mathbf{k}\cdot\mathbf{x}}, \tag{5.2}$$

and we fix Dirichlet boundary conditions to coincide with the plane wave solution.

With these exact solutions, we can verify the accuracy of our method while also checking that it is converging to the exact solution at the expected rate. For polynomials of order p, we expect the error in the finite element solution to be $\mathcal{O}(h^{p+1})$ in the L^2 norm. In Figure 4 we show a log scale plot of the error with respect to mesh refinement while using piecewise linear basis functions. For this polynomial order of $p = 1$, we see the expected quadratic convergence rate. The algebraic system for this validation was solved using the preconditioner we have derived here.

In Figure 5 we show a log scale plot of the residual vector norms with respect to the outer GMRES iteration count for data set (2.5) on different levels of mesh refinement. The GMRES iteration count grows slightly as we refine the mesh.

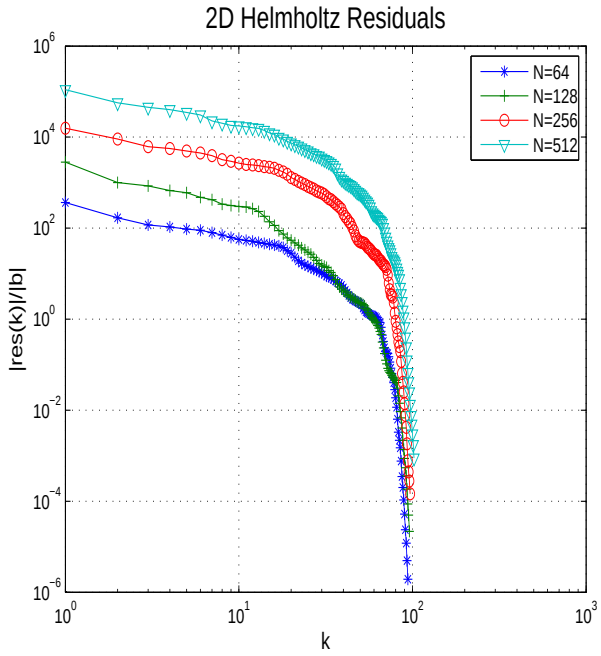

Fig. 3: Preconditioned residual norms at the k^{th} GMRES iteration for the 2D Helmholtz problem on various $N \times N$ meshes. Deflation or including plane wave modes in PyAMG could give future acceleration.

ω	$\hat{a_1}$	N = 32	N = 64	N = 128	N = 256
100000	1.6e-08	2	2	3	6
50000	3.2e-08	2	2	3	11
10000	1.6e-07	3	6	14	82
5000	3.2e-07	5	11	32	1000+
1000	1.6e-06	14	281	1000+	
500	3.2e-06	76	1000+		

Table 5: Number of GMRES iterations for the data sets $\ell_h = 0.0016$, $\ell_v = 0.001$, $c = 1$, $\alpha = 1$, $\gamma = 1.5$ with varying values of ω and N.

The proof of Corollary 3.2 tells us that a key physical parameter in determining the effectiveness of our preconditioner is $\hat{a_1} = \frac{a_1}{a_2} = \frac{\ell_h c}{\omega}$. If we fix the values of c and ℓ_h, we can examine the effectiveness of the preconditioner for different data sets by simply adjusting the value of ω. In Table 5 we see that the GMRES iteration count increases rapidly as we increase $\hat{a_1}$ and / or refine the mesh. So the extreme values our data, specifically the extremely small ratio of ℓ_h to ω, are what allow our preconditioner to be effective. When this ratio is not so extreme, Table 1 shows that more standard block preconditioners perform quite well.

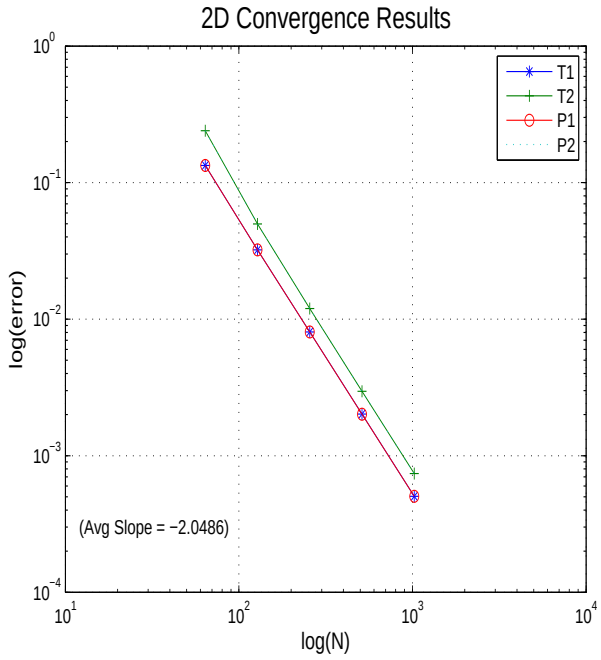

Fig. 4: Convergence results for the 2D problem on various $N \times N$ meshes.

6. Conclusion. In this paper we have presented an effective data-dependent block preconditioner for the coupled pressure-temperature equations used to model trace gas sensor technology. The need for such a preconditioner was demonstrated with a comparison of classical preconditioning techniques such as the incomplete LU, Jacobi and Gauss-Seidel methods. By taking advantage of the particular structure of the preconditioner and its sub-blocks, we have reduced the application of the global preconditioner to the inversion of two Helmholtz sub-blocks and two element-wise vector division operations.

All tests in this paper were run in serial. Future work on this topic will include high performance parallel libraries such as Trilinos [10]. The greatest room for improvement may still be with respect to the Helmholtz solver. Coupling the shifted Laplace preconditioner with a multigrid deflation algorithm has shown promising results [22].

We have been working with the 4×4 system of real variables. Another option for future work is to explore the complex 2×2 system using a software package capable of handling complex arithmetic, such as the Trilinos package Muelu [11].

Acknowledgments. The authors thank Dr. John Zweck, Dr. Susan E. Minkoff and Dr. Ronald Morgan for discussions regarding the mathematical model and linear solvers covered in this paper.

REFERENCES

[1] D. Braess, *Finite Elements: Theory, Fast Solvers, and Applications in Solid Mechanics*, Cambridge University Press, Cambridge, 2001.

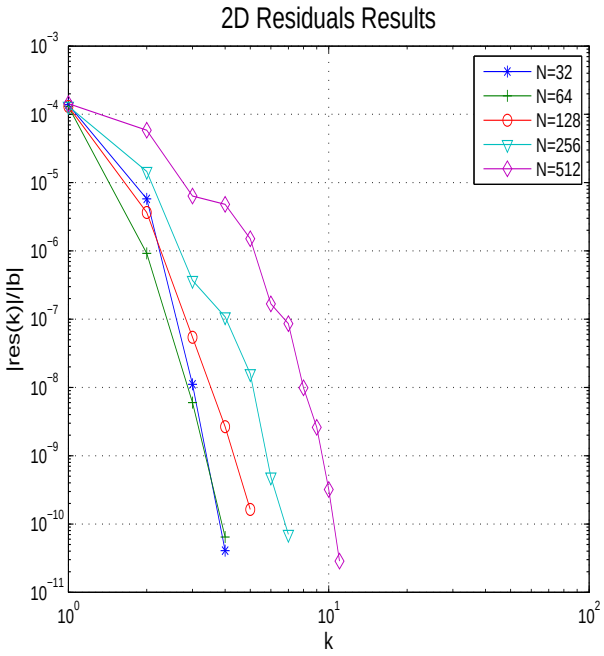

Fig. 5: Preconditioned residual norms at the k^{th} GMRES iteration for the 2D problem on various $N \times N$ meshes.

[2] B. Brennan, R.C. Kirby, J. Zweck, and S.E. Minkoff, *High-performance Python-based simulations of pressure and temperature waves in a trace gas sensor*, in PyHPC 2013: Python for High Performance and Scientific Computing, 2013.

[3] S.C. Brenner and L.R. Scott, *The Mathematical Theory of Finite Element Methods*, Springer, New York, NY, 2007.

[4] T. Oliphant E. Jones and P. Peterson, *SciPy: Open source scientific tools for Python.* `http://www.scipy.org/`, 2001.

[5] Y.A. Erlangga, C. Vuik, and C.W. Oosterlee, *On a class of preconditioners for solving the Helmholtz equation*, Applied Numerical Mathematics, 50 (2004), pp. 409–425.

[6] O.G. Ernst and M.J. Gander, *Why it is difficult to solve Helmholtz problems with classical iterative methods*, in Numerical Analysis of Multiscale Problems, Springer, 2012, pp. 325–363.

[7] V. Faber, T.A. Manteuffel, and S.V. Parter, *On the theory of equivalent operators and application to the numerical solution of uniformly elliptic partial differential equations*, Adv. in Appl. Math., 11 (1990), pp. 109–163.

[8] S.L. Firebaugh, E.A. Terray, and L. Dong, *Optimization of resonator radial dimensions for quartz enhanced photoacoustic spectroscopy systems*, Proc. SPIE 8600, Laser Resonators, Microresonators, and Beam Control XV, 86001S, (2013).

[9] O.G. Ernst H.C. Elman and D.P. O'Leary, *A multigrid method enhanced by Krylov subspace iteration for discrete Helmholtz equations*, SIAM Journal on scientific computing, 23 (2001), pp. 1291–1315.

[10] M. Heroux, R. Bartlett, V. Howle, R. Hoekstra, J. Hu, T. Kolda, R. Lehoucq, K. Long, R. Pawlowski, E. Phipps, A. Salinger, H. Thornquist, R. Tuminaro, J. Willenbring, and A. Williams, *An Overview of Trilinos*, Tech. Report 2003-2927, Sandia National Laboratories, 2003.

[11] J.J Hu, J. Gaidamour, C. Siefert, R.S. Tuminaro, and T. Wiesner, *Muelu-a flexible parallel multigrid framework.*, tech. report, Sandia National Laboratories Albuquerque, NM; Sandia National Laboratories (SNL-CA), Livermore, CA (United States), 2012.

[12] C. Johnson, *Numerical Solutions of Partial Differential Equations by the Finite Element*

Method, Dover Publications, Inc., Mineola, NY, 2009.
[13] A.A. Kosterev, Y.A. Bakhirkin, R.F. Curl, and F.K. Tittel, *Quartz-enhanced photoacoustic spectroscopy*, Optics Letters, 27 (2002), pp. 1902–1904.
[14] A.A. Kosterev, G. Wysocki, Y. Bakhirkin, S. So, R. Lewicki, M. Fraser, F.K. Tittel, and R. F. Curl, *Application of quantum cascade lasers to trace gas analysis*, Appl. Phys. B, 90 (2007), pp. 165–176.
[15] A. Logg, K. Mardal, and G.N. Wells, *Automated Solution of Differential Equations by the Finite Element Method*, Springer, 2012. 978-3-642-23098-1.
[16] M.R. McCurdy, Y. Bakhirkin, G. Wysocki, R. Lewicki, and F.K. Tittel, *Recent advances of laser-spectroscopy based techniques for applications in breath analysis*, Journal of Breath Research, 1.
[17] P.M. Morse and K.U. Ingard, *Theoretical acoustics*, Princeton University Press, Princeton, NJ, 1986.
[18] L. Olson N. Bell and J. Schroder, *Aggregation documentation.* `http://pyamg.googlecode.com/svn/branches/1.0.x/Docs/html/pyamg.aggregation.html`, 2009.
[19] A.A. Kosterev S.E. Minkoff N. Petra, J. Zweck and D. Thomazy, *Theoretical analysis of a quartz-enhanced photoacoustic spectroscopy sensor*, Applied Physics B: Lasers and Optics, 94 (2009), pp. 673–680.
[20] S.E. Minkoff A.A. Kosterev N. Petra, J. Zweck and J.H. Doty III, *Modeling and design optimization of a resonant optothermoacoutstic trace gas sensor*, SIAM J Appl Math, 71 (2001), pp. 309–332.
[21] T.E. Oliphant, *A Guide to NumPy*, vol. 1, Trelgol Publishing USA, 2006.
[22] A.H. Sheikh, D.J.P. Lahaye, and C. Vuik, *A scalable Helmholtz solver combining the shifted Laplace preconditioner with multigrid deflation*, Delft University of Technology, 2011.
[23] A. van der Sluis, *Gershgorin domains for partitioned matrices*, Linear Algebra and its Applications, 26 (1979), pp. 265–280.
[24] L.N. Olson W.N. Bell and J.B. Schroder, *PyAMG: Algebraic multigrid solvers in Python v2.0, 2011.* `http://www.pyamg.org`, 2011.
[25] M.D. Wojcik, M.C. Phillips, B.D. Cannon, and M.S. Taubman, *Gas-phase photoacoustic sensor at 8.41 m using quartz tuning forks and amplitude-modulated quantum cascade lasers*, Appl. Phys. B, 11 (2006), pp. 307–313.
[26] H. Yi, K. Liu, S. Sun, W. Zhang, and X. Gao, *Theoretical analysis of off beam quartz-enhanced photoacoustic spectroscopy sensor*, Optics Communications, (2012).
[27] Y. Zhang, J.A. Smith, A.P. Michel, M.L. Baeck, Z. Wang, J.D. Fast, and C. Gmachl, *Coupled monitoring and modeling of air quality and regional climate during the 2008 Beijing Olympic games*, San Francisco, 2009, American Geophysical Union, Fall Meeting 2009.