Electronic Transactions on Numerical Analysis.
Volume 60, pp. 197–220, 2024.

ISSN 1068–9613.
DOI: 10.1553/etna_vol60s197

ETNA
Kent State University and
Johann Radon Institute (RICAM)

SADDLE POINT PRECONDITIONERS FOR WEAK-CONSTRAINT 4D-VAR*

JEMIMA M. TABEART† AND JOHN W. PEARSON‡

Abstract. Data assimilation algorithms combine information from observations and prior model information to obtain the most likely state of a dynamical system. The linearised weak-constraint four-dimensional variational assimilation problem can be reformulated as a saddle point problem, which admits more scope for preconditioners than the primal form. In this paper we design new terms that can be used within existing preconditioners, such as block diagonal and constraint-type preconditioners. Our novel preconditioning approaches (i) incorporate model information and (ii) are designed to target correlated observation error covariance matrices. To our knowledge, (i) has not been considered previously for data assimilation problems. We develop a theory demonstrating the effectiveness of the new preconditioners within Krylov subspace methods. Linear and non-linear numerical experiments reveal that our new approach leads to faster convergence than existing state-of-the-art preconditioners for a broader range of problems than indicated by the theory alone. We present a range of numerical experiments performed in serial.

Key words. saddle point systems, variational data assimilation, preconditioning

AMS subject classifications. 65F08, 65F10, 65N21

1. Introduction.

Data assimilation has seen substantial interest in fields such as numerical weather prediction [6, 32], ecology [30, 31], and hydrology [7, 44] in recent decades. The variational data assimilation problem can be written mathematically as follows: For a given time window $[t_0, t_N]$, let $\mathbf{x}_i^t \in \mathbb{R}^s$ be the true state of a dynamical system of interest at time t_i, where s is the number of state variables. Data assimilation algorithms combine observations of a dynamical system $\mathbf{y}_i \in \mathbb{R}^{p_i}$ at times t_i with prior information from a model, $\mathbf{x}_b \in \mathbb{R}^s$, to find $\mathbf{x}_i \in \mathbb{R}^s$, the most likely state of the system at time t_i. The prior or background state is valid at initial time t_0 and can be written as an approximation of the true state via $\mathbf{x}_b = \mathbf{x}_0^t + \epsilon^b$. We assume that the background errors satisfy $\epsilon^b \sim \mathcal{N}(0, \mathbf{B})$, where $\mathbf{B} \in \mathbb{R}^{s\times s}$ is a background error covariance matrix. In order to compare observations made at different locations or of different variables than those in the state vector $\mathbf{x}_i$, we define a possibly non-linear observation operator $\mathcal{H}_i : \mathbb{R}^s \to \mathbb{R}^{p_i}$ that maps from the state variable space to the observation space at time t_i. Observations at time t_i are written as $\mathbf{y}_i = \mathcal{H}_i[\mathbf{x}_i^t] + \epsilon_i \in \mathbb{R}^{p_i}$, for $i = 0, 1, \ldots, N$, where the observation error satisfies $\epsilon_i \sim \mathcal{N}(0, \mathbf{R}_i)$ and $\mathbf{R}_i \in \mathbb{R}^{p_i \times p_i}$ are observation error covariance matrices.

In weak-constraint four-dimensional variational data assimilation (4D-Var) the state $\mathbf{x}_{i-1}$ at time t_{i-1} is propagated to the next observation time t_i using an imperfect forecast model $\mathcal{M}_i$ to obtain $\mathbf{x}_i = \mathcal{M}_i(\mathbf{x}_{i-1}) + \epsilon_i^m$. The model error at each time is given by $\epsilon_i^m \sim \mathcal{N}(0, \mathbf{Q}_i)$, where $\mathbf{Q}_i \in \mathbb{R}^{s\times s}$ is the model error covariance matrix at time t_i. It is typically assumed that the error covariance matrices are mutually uncorrelated across different types and different observation times. The analysis, or most likely state $\mathbf{x}_0$ at time t_0, minimises the weak-

*Received April 11, 2023. Accepted March 3, 2024. Published online on April 29, 2024. Recommended by Melina Freitag.
†Department of Mathematics and Computer Science, Eindhoven University of Technology, De Zaale, Eindhoven, 5612 AZ, The Netherlands (j.m.tabeart@tue.nl).
‡School of Mathematics, The University of Edinburgh, James Clerk Maxwell Building, The King's Buildings, Peter Guthrie Tait Road, Edinburgh, EH9 3FD, United Kingdom (j.pearson@ed.ac.uk).

ETNA
Kent State University and
Johann Radon Institute (RICAM)

constraint 4D-Var objective function given by

$$
\begin{aligned}
J(\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_N) = {} & (\mathbf{x}_0 - \mathbf{x}_b)^\top \mathbf{B}^{-1} (\mathbf{x}_0 - \mathbf{x}^b) \\
& + \sum_{i=0}^{N} (\mathbf{y}_i - \mathcal{H}_i[\mathbf{x}_i])^\top \mathbf{R}_i^{-1} (\mathbf{y}_i - \mathcal{H}_i[\mathbf{x}_i]) \\
& + \sum_{i=1}^{N} (\mathbf{x}_i - \mathcal{M}_i(\mathbf{x}_{i-1}))^\top \mathbf{Q}_i^{-1} (\mathbf{x}_i - \mathcal{M}_i(\mathbf{x}_{i-1})).
\end{aligned} \tag{1.1}
$$

Weak-constraint 4D-Var is used in numerical weather prediction (NWP) to estimate the initial condition for a weather forecast [42, 43]. Practical implementations of (1.1) pose mathematical and computational challenges. Firstly, the dimension of the problem can be large: in NWP [6] the dimension of the state can be of order 10^9, and the number of observations can be of order 10^6. This application is also time critical, so the time that can be allocated to the data assimilation procedure is very limited in operational situations. Computational efficiency is therefore vital, and designing techniques to ensure fast convergence of the minimisation of the objective function has been an ongoing area of research interest; see for instance [5, 9, 13, 21].

The weak-constraint objective function (1.1) is typically solved via an incremental approach, where a small number of non-linear outer loops and a larger number of linearised inner loops are solved [16, 17]. Standard Krylov subspace solvers can then be applied for the inner loops. An alternative approach involves solving the linearised inner loop using a saddle point formulation [8, 13, 15, 18], which admits a richer choice of preconditioning structures compared to the primal form. Prior work has developed specific preconditioners for the saddle point data assimilation problem [13, 15], typically focusing on approximations to the term containing information about the linearised model. Block diagonal preconditioners are appealing due to the potential to apply the MINRES algorithm and develop guaranteed theoretical insight about the convergence rates based on the eigenvalues of the preconditioned system, although inexact constraint preconditioners with GMRES have been found to yield better performance in data assimilation settings for a variety of problems [14, 15].

Recent research on the primal form of the variational data assimilation problem has revealed that the observation error covariance matrices $\mathbf{R}_i$ play an important role in determining the convergence of iterative methods [38, 40]. In the last decades, researchers have increasingly made use of observing systems that have correlated observation error covariance matrices [37, 45]. Previous saddle point preconditioners typically applied the exact inverses of $\mathbf{R}_i$ [15, 18], which is known to be computationally infeasible for many satellite observing systems [39, 45]. It is therefore expected that new terms within saddle point preconditioners which incorporate correlated information from $\mathbf{R}_i$ and are inexpensive to apply could be beneficial in terms of convergence for many observing systems. In what follows, we therefore consider preconditioners for the observation error covariance matrix that explicitly allow for full correlation matrices.

In this paper, we consider novel terms within existing preconditioning structures for the saddle point framework, with particular focus on the correlated observation error setting. To our knowledge, our preconditioners account explicitly for model information within the model term for the first time. We begin in Section 2 by introducing the saddle point formulation of the weak-constraint 4D-Var data assimilation problem and presenting existing state-of-the-art preconditioners for the saddle point setting. In Section 3 we prove bounds for the eigenvalues of a block diagonal preconditioned system in terms of constituent matrices of the saddle point problem. In Sections 4 and 5 we then analyse the effect of applying existing and new

ETNA
Kent State University and
Johann Radon Institute (RICAM)

preconditioners for the model term and observation error covariance term, respectively. In Section 6 we present our numerical framework for the experimental results in Section 7. For both the non-linear Lorenz 96 problem and the linear heat equation problem, we find that our new preconditioners result in a reduced iteration count compared to block diagonal and inexact constraint preconditioners in a variety of settings in the presence of correlated observation errors. Although the theoretical guarantees only apply to the block diagonal preconditioner, the qualitative behaviour is similar for the inexact constraint preconditioner with large reductions in iterations and matrix-vector products. Finally, in Section 8 we present our conclusions.

2. Background.

2.1. Data assimilation. In this section we introduce the saddle point formulation of the weak-constraint four-dimensional variational (4D-Var) data assimilation problem (1.1). Given a time window $[t_0, t_N]$ split into N subwindows, we wish to find the state $\mathbf{x} \in \mathbb{R}^s$ at time t_0 that is closest in a weighted norm sense both to the observations throughout the time window and to prior information at the initial time propagated by a model. The incremental primal formulation updates $\mathbf{x}_0^{(l+1)} = \mathbf{x}_0^{(l)} + \boldsymbol{\delta}\mathbf{x}^{(l)}$ by solving an objective function via a series of inner and outer loops to find a sequence of increments to the background state $\mathbf{x}_b = \mathbf{x}_0^{(0)}$. In the inner loop a linearised problem is solved, typically using iterative solvers such as the Conjugate Gradient method [23], and the outer loop is used to update linearisations of model and observation operators.

For each outer loop l, the inner loop minimises a quadratic objective function to find $\boldsymbol{\delta}\mathbf{x}^{(l)} \in \mathbb{R}^{(N+1)s}$, where $\boldsymbol{\delta}\mathbf{x}^{(l)} = \mathbf{x}^{(l+1)} - \mathbf{x}^{(l)}$. Writing $\boldsymbol{\delta}\mathbf{x} = (\boldsymbol{\delta}\mathbf{x}_0^\top, \boldsymbol{\delta}\mathbf{x}_1^\top, \ldots, \boldsymbol{\delta}\mathbf{x}_N^\top)^\top$, the full non-linear observation operator $\mathcal{H}_i$ (similarly, the model operator $\mathcal{M}_i$) is linearised around the current best guess $\mathbf{x}_i^{(l)}$ to obtain the linearised operator $\mathbf{H}_i^{(l)}$ (respectively $\mathbf{M}_i^{(l)}$). The updated initial guess $\boldsymbol{\delta}\mathbf{x}_0^{(l)}$ is propagated forward between the observation times by $\mathbf{M}_i^{(l)}$ to obtain $\boldsymbol{\delta}\mathbf{x}_{i+1}^{(l)} = \mathbf{M}_{i+1}^{(l)}\boldsymbol{\delta}\mathbf{x}_i^{(l)}$. We note that the time between the observations is likely to consist of multiple numerical model time-steps, hence $\mathbf{M}_i^{(l)}$ often corresponds to the composition of many discretised models for a single observation time-step.

Alternatively, the quadratic objective function in the inner loop may be replaced with a saddle point system. Following the notation of [18, Eq. (3.17)], we substitute the linearised objective function with the following saddle point system:

$$\begin{bmatrix} \mathbf{D} & \mathbf{0} & \mathbf{L} \\ \mathbf{0} & \mathbf{R} & \mathbf{H} \\ \mathbf{L}^\top & \mathbf{H}^\top & \mathbf{0} \end{bmatrix} \begin{bmatrix} \boldsymbol{\delta\eta} \\ \boldsymbol{\delta\nu} \\ \boldsymbol{\delta}\mathbf{x} \end{bmatrix} = \begin{bmatrix} \mathbf{b} \\ \mathbf{d} \\ \mathbf{0} \end{bmatrix}. \tag{2.1}$$

In this paper we focus on new preconditioners for the saddle point coefficient matrix:

$$\boldsymbol{\mathcal{A}} = \begin{bmatrix} \mathbf{D} & \mathbf{0} & \mathbf{L} \\ \mathbf{0} & \mathbf{R} & \mathbf{H} \\ \mathbf{L}^\top & \mathbf{H}^\top & \mathbf{0} \end{bmatrix} \in \mathbb{R}^{(N+1)(2s+p)\times(N+1)(2s+p)}, \tag{2.2}$$

where $\mathbf{D}$, $\mathbf{R}$, $\mathbf{H}$ are the following block diagonal matrices:

$$\begin{aligned} \mathbf{D} &= \texttt{blkdiag}\left(\mathbf{B}, \mathbf{Q}_1, \mathbf{Q}_2, \ldots, \mathbf{Q}_N\right) \in \mathbb{R}^{(N+1)s\times(N+1)s}, \\ \mathbf{R} &= \texttt{blkdiag}\left(\mathbf{R}_0, \mathbf{R}_1, \mathbf{R}_2, \ldots, \mathbf{R}_N\right) \in \mathbb{R}^{(N+1)p\times(N+1)p}, \\ \mathbf{H} &= \texttt{blkdiag}\left(\mathbf{H}_0^{(l)}, \mathbf{H}_1^{(l)}, \mathbf{H}_2^{(l)}, \ldots, \mathbf{H}_N^{(l)}\right) \in \mathbb{R}^{(N+1)p\times(N+1)s}, \end{aligned}$$

ETNA
Kent State University and
Johann Radon Institute (RICAM)

with the ith diagonal block in each case being the relevant matrix for time t_{i-1}. Here, the MATLAB-style notation '`blkdiag`' is used to describe a block diagonal matrix in terms of its block diagonal entries.

The matrix $\mathbf{L} \in \mathbb{R}^{(N+1)s \times (N+1)s}$ contains the linearised model information that evolves $\mathbf{x}_0$ between the N observation times, or subwindows, written as

$$\mathbf{L} = \begin{bmatrix} \mathbf{I} & & & & \\ -\mathbf{M}_1^{(l)} & \mathbf{I} & & & \\ & -\mathbf{M}_2^{(l)} & \mathbf{I} & & \\ & & \ddots & \ddots & \\ & & & -\mathbf{M}_N^{(l)} & \mathbf{I} \end{bmatrix}, \tag{2.3}$$

where $\mathbf{I}$ denotes the identity matrix. As we consider preconditioners for the inner loop only, for the remainder of the paper we drop the superscripts that denote the outer loop iteration and simply use $\mathbf{H}_i$ and $\mathbf{M}_i$.

As the saddle point system is indefinite, methods such as the Conjugate Gradient algorithm cannot be used. Depending on the choice of preconditioner, MINRES [28] or GMRES [35] are examples of viable algorithms for solving linear systems of the form (2.1). We note that one challenge of the saddle point formulation is that a monotonic decrease of the objective cost is no longer guaranteed [15]. This can be challenging in operational settings where a very limited number of iterations are performed. In this paper we design new preconditioners and study their numerical performance when iterative methods are permitted to reach convergence. The goal is the design of sufficiently effective and efficient preconditioners to allow convergence of MINRES- or GMRES-approaches in an operational setting.

Numerical methods for saddle point systems are well-studied in the optimisation literature (see [1] for a comprehensive survey). In order to devise suitable approximations of such systems, as in the forthcoming section, one powerful approach is to approximate the 'leading' $(1,1)$-block of the matrix, along with its Schur complement [25, 27]. For saddle point problems arising from data assimilation, not only does the $(1,1)$-block often have a complex structure, but the constraint block contains the evolution of the model terms forward/backward in time, leading to a Schur complement which is very difficult to approximate cheaply [13]. Therefore, approximating the constraint block cheaply is an important consideration for a good preconditioner of the matrix (2.2). The work presented here therefore attempts to combine suitable approximations for the $(1,1)$-block, the constraint block, and hence the Schur complement.

2.2. Preconditioners for saddle point systems from data assimilation. We now introduce some preconditioners that have been applied to the saddle point formulation of the data assimilation problem described above. We start by considering two classes of preconditioner, the block diagonal preconditioner and the inexact constraint preconditioner. Although the standard constraint approach [2, 3] would include $\widehat{\mathbf{H}}$, an approximation of $\mathbf{H}$, the use of $\widehat{\mathbf{H}} = \mathbf{0}$ is popular in the data assimilation setting where it is commonly called the 'inexact' constraint preconditioner [13, 14, 15]. In this paper we use the same forms that are considered in [15], which are given by

$$\mathcal{P}_D = \begin{bmatrix} \widehat{\mathbf{D}} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \widehat{\mathbf{R}} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \widehat{\mathbf{S}} \end{bmatrix}, \qquad \mathcal{P}_I = \begin{bmatrix} \mathbf{D} & \mathbf{0} & \widehat{\mathbf{L}} \\ \mathbf{0} & \widehat{\mathbf{R}} & \mathbf{0} \\ \widehat{\mathbf{L}}^\top & \mathbf{0} & \mathbf{0} \end{bmatrix}, \tag{2.4}$$

where $\widehat{\mathbf{D}}$ and $\widehat{\mathbf{R}}$ are approximations of $\mathbf{D}$ and $\mathbf{R}$ which are easier to apply than the original matrices, and $\widehat{\mathbf{L}}$ is an efficient approximation of $\mathbf{L}$. The exact (negative) Schur complement is

ETNA
Kent State University and
Johann Radon Institute (RICAM)

given by

$$\mathbf{S} = \mathbf{L}^\top \mathbf{D}^{-1} \mathbf{L} + \mathbf{H}^\top \mathbf{R}^{-1} \mathbf{H},$$

which we may approximate by $\widehat{\mathbf{S}}$. For instance, one may drop the second term and take an approximation $\widehat{\mathbf{L}}$ of $\mathbf{L}$, which is a Schur complement approximation of the form $\widehat{\mathbf{L}}^\top \mathbf{D}^{-1} \widehat{\mathbf{L}}$.

One attraction of the block diagonal preconditioner is that one may guarantee a fixed convergence rate based on the eigenvalues of the preconditioned system $\mathcal{P}_D^{-1}\mathcal{A}$. In Section 3 we present bounds for the eigenvalues of the preconditioned system for the block diagonal preconditioner. The inexact constraint preconditioner has been found to yield improved convergence for both toy- and operational-scale data assimilation problems, compared to the block diagonal preconditioner for the same choices of $\widehat{\mathbf{L}}$ and $\widehat{\mathbf{R}}$ [13, 14, 15]. However, it is difficult to develop any theoretical results apart from the unrealistic setting where $\widehat{\mathbf{L}} = \mathbf{L}$, $\widehat{\mathbf{H}} = \mathbf{H}$, $\widehat{\mathbf{R}} = \mathbf{R}$ [13].

In this paper we consider new preconditioners $\widehat{\mathbf{L}}$ and $\widehat{\mathbf{R}}$ for the model and observation error terms. We now present choices of $\widehat{\mathbf{L}}$ and $\widehat{\mathbf{R}}$ that have been studied before. Applying the exact choice of $\mathbf{L}^{-1}$ in a preconditioner can be prohibitively expensive as it requires serial products of model terms and their adjoints [13]. Evaluations of the model acting on a vector can dominate the computational cost of data assimilation algorithms, so reducing the total number of matrix-vector products of the form $\mathbf{M}_i\mathbf{v}$ or $\mathbf{M}_i^\top \mathbf{v}$ is an important consideration when designing new preconditioners. A common choice of preconditioner for $\mathbf{L}$ replaces all of the sub-diagonal terms $\mathbf{M}_i$ with $\mathbf{0}$ [15, 18]. We shall denote this choice of preconditioner as $\mathbf{L}_0$ and note that $\mathbf{L}_0$ is the $(N+1)s \times (N+1)s$ identity matrix, hence so is its inverse.

Another choice of preconditioner that was considered in [13, 15] is given by

$$\mathbf{L}_I = \begin{bmatrix} \mathbf{I} & & & & \\ -\mathbf{I} & \mathbf{I} & & & \\ & -\mathbf{I} & \mathbf{I} & & \\ & & \ddots & \ddots & \\ & & & -\mathbf{I} & \mathbf{I} \end{bmatrix},$$

where the sub-diagonal blocks are replaced with the negative $s \times s$ identity matrix. This was found to perform well experimentally when used in the inexact constraint preconditioner for a two-layer quasi-geostrophic model [13]. However, it does not include any model information, and may not be expected to perform well for all problems. In Section 4 we introduce a new preconditioner that incorporates model information.

Within the outer-inner loop structure of the weak-constraint solver, information from previous outer loops can be used to cheaply update the preconditioners in subsequent loops. This has been found to be computationally beneficial [12]. Such approaches could be combined with the method considered in this paper and are expected to provide value. However, we consider preconditioners for the first outer loop, where there is no additional information, and hence developing computationally feasible stand-alone preconditioners is crucial.

In most prior work the exact observation error covariance matrix has been used in preconditioners. This is due to the fact that the cost of applying $\mathbf{R}^{-1}$ is assumed not to be prohibitive due to the inherent block diagonal structure of $\mathbf{R}$. However, the rising use of more complex correlation structures (such as inter-channel observation errors, see [45]) means that applying $\mathbf{R}^{-1}$ exactly is not always affordable and may be a computational bottleneck. One option is to apply a computationally cheap approximation of $\mathbf{R}$ within the preconditioner, such as the diagonal of $\mathbf{R}$. However, this is typically an inaccurate approximation which can significantly delay convergence. In this paper we consider new classes of preconditioners for

ETNA
Kent State University and
Johann Radon Institute (RICAM)

the observation error covariance term which we expect to improve convergence in the presence of correlated observation error.

3. Bounds for the eigenvalues of the block diagonal preconditioner. For a symmetric positive definite block diagonal preconditioner of the form $\mathcal{P}_D$, it is possible to describe the convergence of the preconditioned MINRES algorithm by analysing the eigenvalues of the preconditioned system $\mathcal{P}_D^{-1}\mathcal{A}$ (see, e.g., [11, Chapter 4]). In the data assimilation setting, $\mathcal{P}_D$ is symmetric positive definite, as all proposed approximations $\widehat{\mathbf{D}}$, $\widehat{\mathbf{R}}$ as well as the Schur complement approximation $\widehat{\mathbf{S}}$ are themselves defined to be symmetric positive definite. We note that for the preconditioned GMRES algorithm, spectral information is insufficient to describe convergence (see [19], for instance). Therefore the theoretical bounds presented in Sections 3–5 are informative strictly for the block diagonal preconditioner, with experiments for the inexact constraint preconditioner $\mathcal{P}_I$ being presented in Section 7 for numerical comparison.

The matrices $\mathcal{A}$ and $\mathcal{P}_D$ defined in (2.2) and (2.4) can be written as

$$\mathcal{A} = \begin{bmatrix} \mathbf{\Phi} & \mathbf{\Psi}^\top \\ \mathbf{\Psi} & \mathbf{0} \end{bmatrix}, \qquad \mathcal{P}_D = \begin{bmatrix} \widehat{\mathbf{\Phi}} & \mathbf{0} \\ \mathbf{0} & \widehat{\mathbf{S}} \end{bmatrix}, \tag{3.1}$$

where

$$\mathbf{\Phi} = \begin{bmatrix} \mathbf{D} & \mathbf{0} \\ \mathbf{0} & \mathbf{R} \end{bmatrix}, \qquad \mathbf{\Psi} = \begin{bmatrix} \mathbf{L}^\top & \mathbf{H}^\top \end{bmatrix},$$

and with $\widehat{\mathbf{\Phi}}$ and $\widehat{\mathbf{S}}$ approximations of $\mathbf{\Phi}$ and the (negative) Schur complement $\mathbf{S} = \mathbf{\Psi}\mathbf{\Phi}^{-1}\mathbf{\Psi}^\top$. We now denote

$$\widetilde{\mathbf{S}} = \mathbf{L}^\top\mathbf{D}^{-1}\mathbf{L}, \qquad \widehat{\mathbf{S}} = \widehat{\mathbf{L}}^\top\mathbf{D}^{-1}\widehat{\mathbf{L}},$$

and $\widehat{\mathbf{D}}$, $\widehat{\mathbf{R}}$, $\widehat{\mathbf{L}}$ are approximations of $\mathbf{D}$, $\mathbf{R}$, $\mathbf{L}$. For the forthcoming theory, we suppose that $\mathbf{D}$, $\widehat{\mathbf{D}}$, $\mathbf{R}$, $\widehat{\mathbf{R}}$, $\mathbf{S}$, $\widetilde{\mathbf{S}}$, $\widehat{\mathbf{S}}$ are symmetric positive definite, with

$$\lambda(\widehat{\mathbf{D}}^{-1}\mathbf{D}) \in [\lambda_{\mathbf{D}}, \Lambda_{\mathbf{D}}], \qquad \lambda(\widehat{\mathbf{R}}^{-1}\mathbf{R}) \in [\lambda_{\mathbf{R}}, \Lambda_{\mathbf{R}}],$$
$$\lambda(\widetilde{\mathbf{S}}^{-1}\mathbf{S}) \in [\lambda_{\mathbf{S}}, \Lambda_{\mathbf{S}}], \qquad \lambda\left((\widehat{\mathbf{L}}^\top\widehat{\mathbf{L}})^{-1}(\mathbf{L}^\top\mathbf{L})\right) \in [\lambda_{\mathbf{L}}, \Lambda_{\mathbf{L}}],$$

where $\lambda(\cdot)$ denotes the eigenvalues of a matrix. We may then prove the following result:

THEOREM 3.1. *With the definitions as stated above, the eigenvalues of $\mathcal{P}_D^{-1}\mathcal{A}$ are real and satisfy*

$$\lambda(\mathcal{P}_D^{-1}\mathcal{A}) \in \left[\frac{\lambda_{\mathbf{\Phi}} - \sqrt{\lambda_{\mathbf{\Phi}}^2 + 4\Lambda_{\mathbf{\Phi}}\Lambda_{\mathbf{S}}\Lambda_{\mathbf{L}}\kappa(\mathbf{D})}}{2}, \frac{\Lambda_{\mathbf{\Phi}} - \sqrt{\Lambda_{\mathbf{\Phi}}^2 + \frac{4\lambda_{\mathbf{\Phi}}\lambda_{\mathbf{S}}\lambda_{\mathbf{L}}}{\kappa(\mathbf{D})}}}{2}\right]$$
$$\cup [\lambda_{\mathbf{\Phi}}, \Lambda_{\mathbf{\Phi}}] \cup \left[\frac{\lambda_{\mathbf{\Phi}} + \sqrt{\lambda_{\mathbf{\Phi}}^2 + \frac{4\lambda_{\mathbf{\Phi}}\lambda_{\mathbf{S}}\lambda_{\mathbf{L}}}{\kappa(\mathbf{D})}}}{2}, \frac{\Lambda_{\mathbf{\Phi}} + \sqrt{\Lambda_{\mathbf{\Phi}}^2 + 4\Lambda_{\mathbf{\Phi}}\Lambda_{\mathbf{S}}\Lambda_{\mathbf{L}}\kappa(\mathbf{D})}}{2}\right],$$

where $\lambda_{\mathbf{\Phi}} = \min\{\lambda_{\mathbf{D}}, \lambda_{\mathbf{R}}\}$, $\Lambda_{\mathbf{\Phi}} = \max\{\Lambda_{\mathbf{D}}, \Lambda_{\mathbf{R}}\}$, and $\kappa(\cdot)$ denotes the condition number of a matrix.

ETNA
Kent State University and
Johann Radon Institute (RICAM)

Proof. Applying well-known results (see [34, p. 2906], [33, Theorem 4.2.1]), we have

$$\lambda \in \left[\frac{\delta - \sqrt{\delta^2 + 4\Delta\Phi}}{2}, \frac{\Delta - \sqrt{\Delta^2 + 4\delta\phi}}{2}\right] \cup [\delta, \Delta] \cup \left[\frac{\delta + \sqrt{\delta^2 + 4\delta\phi}}{2}, \frac{\Delta + \sqrt{\Delta^2 + 4\Delta\Phi}}{2}\right], \tag{3.2}$$

where δ, ϕ denote the minimum eigenvalues of $\widehat{\mathbf{\Phi}}^{-1}\mathbf{\Phi}$, $\widehat{\mathbf{S}}^{-1}\mathbf{S}$ for a general block diagonal saddle point preconditioner (3.1), and Δ, Φ represent the corresponding maximum eigenvalues.

Clearly, for this problem $\delta = \lambda_{\mathbf{\Phi}}$ and $\Delta = \Lambda_{\mathbf{\Phi}}$. It remains to analyse the minimum and maximum eigenvalues of the preconditioned Schur complement, which we note is within the range of the Rayleigh quotient (for $\mathbf{v} \neq \mathbf{0}$):

$$\frac{\mathbf{v}^\top \mathbf{S}\mathbf{v}}{\mathbf{v}^\top \widehat{\mathbf{S}}\mathbf{v}} = \frac{\mathbf{v}^\top \mathbf{S}\mathbf{v}}{\mathbf{v}^\top \widetilde{\mathbf{S}}\mathbf{v}} \cdot \frac{\mathbf{v}^\top \widetilde{\mathbf{S}}\mathbf{v}}{\mathbf{v}^\top \widehat{\mathbf{S}}\mathbf{v}} = \frac{\mathbf{v}^\top \mathbf{S}\mathbf{v}}{\mathbf{v}^\top \widetilde{\mathbf{S}}\mathbf{v}} \cdot \frac{\mathbf{v}^\top \mathbf{L}^\top \mathbf{D}^{-1}\mathbf{L}\mathbf{v}}{\mathbf{v}^\top \mathbf{L}^\top \mathbf{L}\mathbf{v}} \cdot \frac{\mathbf{v}^\top \mathbf{L}^\top \mathbf{L}\mathbf{v}}{\mathbf{v}^\top \widehat{\mathbf{L}}^\top \widehat{\mathbf{L}}\mathbf{v}} \cdot \frac{\mathbf{v}^\top \widehat{\mathbf{L}}^\top \widehat{\mathbf{L}}\mathbf{v}}{\mathbf{v}^\top \widehat{\mathbf{L}}^\top \mathbf{D}^{-1}\widehat{\mathbf{L}}\mathbf{v}}. \tag{3.3}$$

Observing that

$$\frac{\mathbf{v}^\top \mathbf{S}\mathbf{v}}{\mathbf{v}^\top \widetilde{\mathbf{S}}\mathbf{v}} \in [\lambda_{\mathbf{S}}, \Lambda_{\mathbf{S}}], \qquad \frac{\mathbf{v}^\top \mathbf{L}^\top \mathbf{D}^{-1}\mathbf{L}\mathbf{v}}{\mathbf{v}^\top \mathbf{L}^\top \mathbf{L}\mathbf{v}} = \frac{\mathbf{y}^\top \mathbf{y}}{\mathbf{y}^\top \mathbf{D}\mathbf{y}} \in \left[\frac{1}{\Lambda_{\mathbf{D}}}, \frac{1}{\lambda_{\mathbf{D}}}\right],$$
$$\frac{\mathbf{v}^\top \mathbf{L}^\top \mathbf{L}\mathbf{v}}{\mathbf{v}^\top \widehat{\mathbf{L}}^\top \widehat{\mathbf{L}}\mathbf{v}} \in [\lambda_{\mathbf{L}}, \Lambda_{\mathbf{L}}], \qquad \frac{\mathbf{v}^\top \widehat{\mathbf{L}}^\top \widehat{\mathbf{L}}\mathbf{v}}{\mathbf{v}^\top \widehat{\mathbf{L}}^\top \mathbf{D}^{-1}\widehat{\mathbf{L}}\mathbf{v}} = \frac{\mathbf{z}^\top \mathbf{D}\mathbf{z}}{\mathbf{z}^\top \mathbf{z}} \in [\lambda_{\mathbf{D}}, \Lambda_{\mathbf{D}}],$$

where $\mathbf{y} = \mathbf{D}^{-1/2}\mathbf{L}\mathbf{v} \neq \mathbf{0}$, $\mathbf{z} = \mathbf{D}^{-1/2}\widehat{\mathbf{L}}\mathbf{v} \neq \mathbf{0}$, we may write

$$\frac{\mathbf{v}^\top \mathbf{S}\mathbf{v}}{\mathbf{v}^\top \widehat{\mathbf{S}}\mathbf{v}} \in \left[\frac{\lambda_{\mathbf{S}}\lambda_{\mathbf{L}}\lambda_{\mathbf{D}}}{\Lambda_{\mathbf{D}}}, \frac{\Lambda_{\mathbf{S}}\Lambda_{\mathbf{L}}\Lambda_{\mathbf{D}}}{\lambda_{\mathbf{D}}}\right] = \left[\frac{\lambda_{\mathbf{S}}\lambda_{\mathbf{L}}}{\kappa(\mathbf{D})}, \Lambda_{\mathbf{S}}\Lambda_{\mathbf{L}}\kappa(\mathbf{D})\right].$$

Therefore, it holds that $\phi \geq \frac{\lambda_{\mathbf{S}}\lambda_{\mathbf{L}}}{\kappa(\mathbf{D})}$ and $\Phi \leq \Lambda_{\mathbf{S}}\Lambda_{\mathbf{L}}\kappa(\mathbf{D})$. Substituting the bounds for δ, Δ, ϕ, Φ into (3.2) then gives the result. ∎

REMARK 3.2. Theorem 3.1 is an extension of known results, for example [34, p. 2906] and [33, Theorem 4.2.1], and in particular an application of this methodology to saddle point systems arising from weak-constraint 4D-Var. We highlight that eigenvalue results of this form are important because they lead to concrete convergence properties of MINRES. As in [11, Theorem 4.14], if $\lambda(\boldsymbol{\mathcal{P}}_D^{-1}\boldsymbol{\mathcal{A}}) \in [-\mu_1, -\mu_2] \cup [\mu_3, \mu_4]$, with $\mu_1, \mu_2, \mu_3, \mu_4 > 0$ such that $\mu_1 - \mu_2 = \mu_4 - \mu_3$, then after 2ℓ iterations of MINRES:

$$\|\mathbf{r}^{(2\ell)}\|_{\boldsymbol{\mathcal{P}}_D^{-1}} \leq 2\left(\frac{\sqrt{\mu_1\mu_4} - \sqrt{\mu_2\mu_3}}{\sqrt{\mu_1\mu_4} + \sqrt{\mu_2\mu_3}}\right)^{\ell} \|\mathbf{r}^{(0)}\|_{\boldsymbol{\mathcal{P}}_D^{-1}},$$

where $\mathbf{r}^{(\cdot)}$ denotes the residual vector at a given iteration. Note that both the positive or negative intervals for the eigenvalues of $\boldsymbol{\mathcal{P}}_D^{-1}\boldsymbol{\mathcal{A}}$ can always be stretched such that $\mu_1 - \mu_2 = \mu_4 - \mu_3$ holds, so this result holds without loss of generality. This allows one to ascertain the convergence behaviour of MINRES when we apply approximations of $\mathbf{L}^\top\mathbf{L}$, $\mathbf{R}$ (as well as $\mathbf{D}$), including those discussed in Sections 4 and 5.

REMARK 3.3. Due to the way the Rayleigh quotient is factored in (3.3), Theorem 3.1 gives a potentially very weak bound when $\mathbf{D}$ is ill-conditioned. We find the main features that affect the quality of the preconditioner, as predicted by the result, are the approximations of $\mathbf{D}$ and $\mathbf{R}$, the effect of dropping the second term of the Schur complement, and the quality of the approximation of $\mathbf{L}^\top\mathbf{L}$ (characterised by the eigenvalues of $\widehat{\mathbf{L}}^{-\top}\mathbf{L}^\top\mathbf{L}\widehat{\mathbf{L}}^{-1}$). The latter quantity is the subject of the forthcoming analysis.

ETNA
Kent State University and
Johann Radon Institute (RICAM)

4. Approximations $\widehat{\mathbf{L}}$. Theorem 3.1 suggests that the eigenvalues of the preconditioned system are influenced by the quality of the approximation of $\mathbf{L}^\top\mathbf{L}$ by $\widehat{\mathbf{L}}^\top\widehat{\mathbf{L}}$ in the block diagonal preconditioner. In this section we consider existing and new choices of $\widehat{\mathbf{L}}$ and analyse the eigenvalues and structure of $\widehat{\mathbf{L}}^{-\top}\mathbf{L}^\top\mathbf{L}\widehat{\mathbf{L}}^{-1}$, which is similar to $(\widehat{\mathbf{L}}^\top\widehat{\mathbf{L}})^{-1}(\mathbf{L}^\top\mathbf{L})$. The first choice of $\widehat{\mathbf{L}}$, $\mathbf{L}_0$, has previously been used for saddle point preconditioners for the data assimilation problem [15, 18]. We also propose $\mathbf{L}_M$, a new class of parallelisable preconditioners that depends on a user-defined parameter and incorporates model information.

4.1. A new preconditioner $\mathbf{L}_M$ and the eigenvalues of $\mathbf{L}_M^{-\top}\mathbf{L}^\top\mathbf{L}\mathbf{L}_M^{-1}$. We begin by defining our proposed preconditioner $\mathbf{L}_M$ that incorporates model information explicitly. For a user chosen parameter $1 \leq k \leq N+1$, every kth block sub-diagonal of $\mathbf{L}_M$ (i.e., $\mathbf{M}_k, \mathbf{M}_{2k}, \mathbf{M}_{3k}, \dots$) is set equal to $\mathbf{0}$. The other entries of $\mathbf{L}_M$ correspond to those of $\mathbf{L}$. Formally we write this as in the definition below.

DEFINITION 1. *Let $k \in \mathbb{N}$, and let $\mathbf{L}_M = \mathbf{L}_M(k) \in \mathbb{R}^{(N+1)s\times(N+1)s}$ be a block matrix made up of $s \times s$ blocks. For $i, j = 1, \dots, N+1$, we define*

$$\text{the } (i,j)\text{th block of } \mathbf{L}_M = \begin{cases} \mathbf{I} & \text{if } i = j, \\ -\mathbf{M}_j & \text{if } i = j+1 \text{ and } j \text{ is not divisible by } k, \\ \mathbf{0} & \text{otherwise.} \end{cases}$$

Similarly we can write the inverse of $\mathbf{L}_M$ as follows, using straightforward linear algebra:

LEMMA 4.1. *Let $k \in \mathbb{N}$, and let $\mathbf{L}_M = \mathbf{L}_M(k) \in \mathbb{R}^{(N+1)s\times(N+1)s}$ be a block matrix made up of $s \times s$ blocks. For $i, j = 1, \dots, N+1$, we may evaluate that*

$$\text{the } (i,j)\text{th block of } \mathbf{L}_M^{-1} = \begin{cases} \mathbf{I} & \text{if } i = j, \\ \prod_{m=1}^{i-j} \mathbf{M}_{i-m} & \text{if } 1 \leq i-j \leq (i-1) \bmod(k), \\ \mathbf{0} & \text{otherwise.} \end{cases}$$

We note that $\mathbf{L}_M^{-1}$ is lower triangular, and both the number of non-zero blocks of $\mathbf{L}_M^{-1}$ and the number of terms in each of the products of $\mathbf{L}_M^{-1}$ are controlled by the parameter k. Note that $\mathbf{L}$ defined in (2.3) satisfies $\mathbf{L} = \mathbf{L}_M(N+1)$ according to this notation. To further justify the effectiveness of this approximation, we now study the eigenvalues of $\mathbf{L}_M^{-\top}\mathbf{L}^\top\mathbf{L}\mathbf{L}_M^{-1}$ theoretically. We begin by stating the structure of $\mathbf{L}_M^{-\top}\mathbf{L}^\top\mathbf{L}\mathbf{L}_M^{-1}$ in terms of the linearised model matrices $\mathbf{M}_i$, using straightforward linear algebra.

LEMMA 4.2. *We can write $\mathbf{L}_M^{-\top}\mathbf{L}^\top\mathbf{L}\mathbf{L}_M^{-1} = \mathbf{I} + \mathbf{A}(\mathbf{M})$, where the block entries of $\mathbf{A}(\mathbf{M}) \in \mathbb{R}^{(N+1)s\times(N+1)s}$ are defined as follows: For $n = 1, \dots, \lfloor \frac{N}{k} \rfloor$,*

$$[\mathbf{A}(\mathbf{M})]_{i,j} = \begin{cases} (\prod_{t=i}^{nk} \mathbf{M}_t^\top)(\prod_{q=j}^{nk} \mathbf{M}_{nk-q+j}) & \text{for } (n-1)k+1 \leq i,j \leq nk, \\ -\prod_{t=j}^{nk} \mathbf{M}_{nk-t+j} & \text{for } i = nk+1, (n-1)k+1 \leq j \leq nk, \\ -\prod_{t=i}^{nk} \mathbf{M}_t^\top & \text{for } j = nk+1, (n-1)k+1 \leq i \leq nk, \\ \mathbf{0} & \text{otherwise,} \end{cases}$$

where $[\mathbf{A}(\mathbf{M})]_{i,j}$ denotes the (i,j)th block of $\mathbf{A}(\mathbf{M})$.

We now briefly describe the structure of $\mathbf{A}(\mathbf{M})$. The matrix is made up of $\lfloor \frac{N}{k} \rfloor$ overlapping diagonal blocks, where the size of each block is $(k+1)s \times (k+1)s$. Each block 'overlaps' at the $(nk+1, nk+1)$st block of $\mathbf{A}(\mathbf{M})$, meaning that the maximum number of non-zero blocks in any row or column is given by $2k+1$. We use this structure to demonstrate

ETNA
Kent State University and
Johann Radon Institute (RICAM)

that our new preconditioner $\mathbf{L}_M$ yields a number of unit eigenvalues for the preconditioned term $\mathbf{L}_M^{-\top}\mathbf{L}^\top\mathbf{L}\mathbf{L}_M^{-1}$.

PROPOSITION 1. *Let* $\mathbf{L}$ *be defined as in* (2.3) *and* $\mathbf{L}_M$ *be as in Lemma* 4.2. *For* $2 \le k \le N+1$, $\mathbf{L}_M^{-\top}\mathbf{L}^\top\mathbf{L}\mathbf{L}_M^{-1}$ *has at least* rs *unit eigenvalues, with* $r = N + 1 - 2\left\lfloor \frac{N}{k} \right\rfloor$.

Proof. From Lemma 4.2 we can construct eigenvectors corresponding to the zero eigenvalues of $\mathbf{A}(\mathbf{M})$, which will yield unit eigenvalues of $\mathbf{L}_M^{-\top}\mathbf{L}^\top\mathbf{L}\mathbf{L}_M^{-1}$. Let $\mathbf{e}_t$ define the canonical vector taking the unit value in position t and zero elsewhere. Observe that $\mathbf{A}(\mathbf{M})$ is a block diagonal matrix with a $(N - k\left\lfloor \frac{N}{k} \right\rfloor)s \times (N - k\left\lfloor \frac{N}{k} \right\rfloor)s$ zero block in the final position. We can construct $(N - k\left\lfloor \frac{N}{k} \right\rfloor)s$ linearly independent eigenvectors corresponding to the zero eigenvalue for this block using $\mathbf{e}_t$ for $t = (k\left\lfloor \frac{N}{k} \right\rfloor)s + 1, \ldots, Ns$.

For each value of n in Lemma 4.2 we obtain $(k-2)s$ eigenvectors corresponding to a zero eigenvalue of the form

$$(\mathbf{0}, \ldots, \mathbf{0}, \mathbf{v}_t^\top, -(\mathbf{M}_r\mathbf{v}_t)^\top, \mathbf{0}, \ldots, \mathbf{0})^\top,$$

for $rs + 1 \le t \le (r+1)s$ and $nk + 2 \le r \le (n+1)k - 1$. Therefore, these contribute $(k-2)s\left\lfloor \frac{N}{k} \right\rfloor$ linearly independent eigenvectors corresponding to the zero eigenvalue across the whole matrix. From the first block in the matrix we obtain additional s linearly independent eigenvectors corresponding to the zero eigenvalue via $(\mathbf{v}_t^\top, -(\mathbf{M}_1\mathbf{v}_t)^\top, \mathbf{0}, \ldots, \mathbf{0})^\top$, for $t = 1, \ldots, s$.

Combining the above reasoning, we obtain rs unit eigenvalues of $\mathbf{L}_M^{-\top}\mathbf{L}^\top\mathbf{L}\mathbf{L}_M^{-1}$ as required, where

$$r = 1 + (k-2)\left\lfloor \frac{N}{k} \right\rfloor + N - k\left\lfloor \frac{N}{k} \right\rfloor = N + 1 - 2\left\lfloor \frac{N}{k} \right\rfloor. \qquad \square$$

We see that r does not decrease as k increases, i.e., incorporating information from more timesteps will generally lead to a larger number of unit eigenvalues of the preconditioned model term. Increasing the number of observation times N will broadly lead to an increase in the number of unit eigenvalues of $\mathbf{L}_M^{-\top}\mathbf{L}^\top\mathbf{L}\mathbf{L}_M^{-1}$, but this behaviour is non-monotonic.

If we introduce assumptions on the spectral radii of the model operator terms, then we can obtain explicit bounds for the eigenvalues of $\mathbf{L}_M^{-\top}\mathbf{L}^\top\mathbf{L}\mathbf{L}_M^{-1}$.

PROPOSITION 2. *If* $\|\mathbf{M}_i\mathbf{M}_i^\top\|_2 \le 1$, *for all* i, *then the eigenvalues of* $\mathbf{L}_M^{-\top}\mathbf{L}^\top\mathbf{L}\mathbf{L}_M^{-1}$ *can be bounded above by* $k + 1 + 2\sqrt{k}$.

Proof. We bound the eigenvalues of $\mathbf{A}(\mathbf{M})$ by splitting the matrix into three sub-matrices $\mathbf{A}(\mathbf{M}) = \mathbf{A}_1 + \mathbf{A}_2 + \mathbf{A}_3$, where $\mathbf{A}_1, \mathbf{A}_2, \mathbf{A}_3$ are symmetric and will be defined explicitly in what follows. As all matrices being considered are symmetric, by using [4, Fact 5.12.2] we can bound the maximum eigenvalue of $\mathbf{A}(\mathbf{M})$ above by

$$\lambda_{\max}(\mathbf{A}(\mathbf{M})) \le \lambda_{\max}(\mathbf{A}_1) + \lambda_{\max}(\mathbf{A}_2) + \lambda_{\max}(\mathbf{A}_3).$$

Let $\mathbf{A}_1$ be a block diagonal matrix with blocks of size $nk \times nk$ and entries defined by

$$\text{the } (i,j)\text{th block of } \mathbf{A}_1 = \begin{cases} \text{the } (i,j)\text{th block of } \mathbf{A}(\mathbf{M}) & \text{for } (n-1)k+1 \le i,j \le nk, \\ \mathbf{0} & \text{otherwise.} \end{cases}$$

Each $ks \times ks$ block has rank s, as the first $(k-1)s$ rows are multiples of the final $s \times s$ rows. Substitution yields the eigenvalue problem

$$\left(\mathbf{M}_{nk}^\top \sum_{t=(n-1)k+1}^{nk} \left(\prod_{p=t}^{nk} \mathbf{M}_{nk-p+t} \right) \left(\prod_{q=t}^{nk-1} \mathbf{M}_q^\top \right) \right) \mathbf{v} = \mu\mathbf{v}.$$

ETNA
Kent State University and
Johann Radon Institute (RICAM)

We apply [24, Theorem 1.3.20], which states that exchanging the order of matrix multiplication for two compatible matrices has no effect on the non-zero eigenvalues of the product, so we instead consider

$$\left(\sum_{t=(n-1)k+1}^{nk}\left(\prod_{p=t}^{nk}\mathbf{M}_{nk-p+t}\right)\left(\prod_{q=t}^{nk}\mathbf{M}_q^\top\right)\right)\bar{\mathbf{v}}=\mu\bar{\mathbf{v}}.$$

We can separate the contribution of each individual term by applying [4, Fact 5.12.2] to obtain

$$\begin{aligned}\mu &\leq \lambda_{\max}\left(\sum_{t=(n-1)k+1}^{nk}\left(\prod_{p=t}^{nk}\mathbf{M}_{nk-p+t}\right)\left(\prod_{q=t}^{nk}\mathbf{M}_q^\top\right)\right)\\ &\leq \sum_{t=(n-1)k+1}^{nk}\prod_{p=t}^{nk}\|\mathbf{M}_p\mathbf{M}_p^\top\|_2 \leq \sum_{t=(n-1)k+1}^{nk} 1 = k.\end{aligned}$$

The maximum eigenvalue of each block is bounded above by k, and hence $\lambda_{\max}(\mathbf{A}_1)\leq k$.

Without loss of generality, assume that $\left\lfloor\frac{N}{k}\right\rfloor$ is odd. For $n=1,3,5,\ldots,2\left\lfloor\frac{N}{2k}\right\rfloor+1$, we define $\mathbf{A}_2$ by a block diagonal matrix with blocks of size $(k+1)s\times(k+1)s$ and entries given by

$$[\mathbf{A}_2]_{i,j}=\begin{cases}\text{the } (i,j)\text{th block of } \mathbf{A}(\mathbf{M}) & \text{for } i=nk+1,(n-1)k+1\leq j\leq nk,\\ \text{the } (i,j)\text{th block of } \mathbf{A}(\mathbf{M}) & \text{for } j=nk+1,(n-1)k+1\leq i\leq nk,\\ \mathbf{0} & \text{otherwise.}\end{cases}$$

The $(k+1)s\times(k+1)s$ blocks have rank $2s$ with non-zero eigenvalues that solve

$$\left(\sum_{t=(n-1)k+1}^{nk}\left(\prod_{p=t}^{nk}\mathbf{M}_{nk-p+t}\right)\left(\prod_{q=t}^{nk}\mathbf{M}_q^\top\right)\right)\mathbf{v}=\mu^2\mathbf{v}.$$

By the same argument as above we can bound $\mu^2\leq k$. Hence, $\lambda_{\max}(\mathbf{A}_2)\leq\sqrt{k}$.

For $n=2,4,6,\ldots,2\left\lfloor\frac{N}{2k}\right\rfloor$, the matrix $\mathbf{A}_3$ is a block diagonal matrix with blocks of size $(k+1)s\times(k+1)s$ and entries given by

$$[\mathbf{A}_3]_{i,j}=\begin{cases}\text{the } (i,j)\text{th block of } \mathbf{A}(\mathbf{M}) & \text{for } i=nk+1,(n-1)k+1\leq j\leq nk-1,\\ \text{the } (i,j)\text{th block of } \mathbf{A}(\mathbf{M}) & \text{for } j=nk+1,(n-1)k+1\leq i\leq nk-1,\\ \mathbf{0} & \text{otherwise.}\end{cases}$$

All blocks have the same structure as the blocks of $\mathbf{A}_2$ and hence have eigenvalues bounded above by $\sqrt{k}$.

The largest eigenvalue of $\mathbf{A}(\mathbf{M})$ is therefore bounded by $\lambda_{\max}(\mathbf{A}(\mathbf{M}))\leq k+2\sqrt{k}$. By adding the identity matrix, we obtain the upper bound for the eigenvalues in the proposition statement. ∎

REMARK 4.3. For smaller values of N, $\mathbf{A}_3$ is not involved. We can therefore apply a similar argument with $\mathbf{A}=\mathbf{A}_1+\mathbf{A}_2$ to obtain the tighter bounds

$$\lambda(\mathbf{L}_M^{-\top}\mathbf{L}^\top\mathbf{L}\mathbf{L}_M^{-1})\leq 1+k+\sqrt{k},\quad \text{if } k\leq N<2k,\ \text{ i.e., if } \left\lfloor\frac{N}{k}\right\rfloor=1.$$

REMARK 4.4. A similar approach is not illustrative when examining a lower bound for the eigenvalues as this would yield negative bounds, whereas all eigenvalues of $\mathbf{L}_M^{-\top}\mathbf{L}^\top\mathbf{L}\mathbf{L}_M^{-1}$ are clearly greater than zero by construction.

ETNA
Kent State University and
Johann Radon Institute (RICAM)

5. Approximations $\widehat{\mathbf{R}}_i$. Theorem 3.1 suggests that the eigenvalues of the preconditioned system are also influenced by the quality of the approximation of $\mathbf{R}$ by $\widehat{\mathbf{R}}$ in the block diagonal preconditioner. In this section we consider four choices of $\widehat{\mathbf{R}}_i$, each of which we apply blockwise to $\mathbf{R}_i$. Similarly to the preconditioners for $\mathbf{L}$, we consider an existing choice of preconditioner that is diagonal and three new choices of preconditioner that include correlation information. We expect the new preconditioners to be beneficial for highly correlated observation error covariance matrices. Correlated observation errors are currently implemented operationally at a number of numerical weather prediction centres for hyperspectral satellite instruments (see, e.g., [41, 45]) and Doppler Radar Winds (DRW) (e.g., [36]). Hyperspectral instruments have a block covariance structure with a number of highly correlated off-diagonal entries, and DRW error statistics are spatially correlated. We focus on the observation error covariance matrix in this section, as efficient approximations for this term have previously been overlooked. In the following we describe the approaches in terms of $\mathbf{R}$ without loss of generality, as the methods can also be applied to approximate the blocks of $\mathbf{D}$.

REMARK 5.1. We note that $\mathbf{R}$ has a block diagonal structure. Each of the preconditioners in this section is applied blockwise, yielding a block diagonal $\widehat{\mathbf{R}}$ with at least $N+1$ blocks. The two preconditioners presented in Section 5.1 *further* increase the sparsity of $\widehat{\mathbf{R}}$.

5.1. Choices of $\widehat{\mathbf{R}}_i$ that increase sparsity. Many previous studies of the saddle point data assimilation formulation assume that $\mathbf{R}$ is diagonal or easy to invert [12, 15]. For those instruments with uncorrelated or diagonally dominant observation error covariance matrices, it is likely that a diagonal approximation $\widehat{\mathbf{R}}$ will be sufficient. However, for more complicated structures, it is unlikely that the exact inverse of $\mathbf{R}$ can be applied efficiently in terms of storage or computation. The approaches we present here are designed to perform well in terms of effectiveness and efficiency for correlation structures that are currently used operationally.

The first two choices of $\widehat{\mathbf{R}}_i$ considered in this section allow for additional sparsification of the observation error component of the preconditioners. The first preconditioner, denoted $\mathbf{R}_{diag}$, takes the diagonal of the original observation error covariance matrix $\mathbf{R}_i$. This is often a first approximation of $\mathbf{R}_i$ or its inverse for simple covariance structures. This choice of $\widehat{\mathbf{R}}_i$ is cheap to apply and extremely sparse but is expected to perform poorly if there is significant off-diagonal correlation structure.

The second sparsifying choice of $\widehat{\mathbf{R}}_i$, denoted $\mathbf{R}_{block}$ in Algorithm 3 in Appendix A, is designed to exploit the existing block structure in $\mathbf{R}_i$. In applications, $\mathbf{R}_i$ itself often has a block structure with the strength of the off-diagonal correlations varying, e.g., for different instruments or measurement types (see, e.g., [45]). The idea of $\mathbf{R}_{block}$ is to retain the sub/super-diagonal blocks of $\mathbf{R}_i$ with the largest norm. Neglecting off-diagonal blocks of $\mathbf{R}_i$ with smaller norm ensures that $\mathbf{R}_{block}$ is decoupled into a block diagonal matrix. Let $\mathbf{R}_i \in \mathbb{R}^{p_i \times p_i}$ be a covariance matrix with an associated vector $\texttt{pvec} \in \mathbb{R}^{p_n}$ that specifies the size of the 'blocks' such that $\sum_{k=1}^{p_n} \texttt{pvec}_k = p_i$. Algorithm 3 returns a block covariance matrix $\mathbf{R}_{block}$ where only off-diagonal blocks with a scaled Frobenius norm larger than a user-defined tolerance `tol` are retained. In order to maximise computational efficiency, only the norms of the blocks on the first super-diagonal are computed. If two (or more) adjacent blocks are retained, then information from higher level super-diagonals is also included. This does not change the overall block structure of the new preconditioner but allows for the inclusion of more information from $\mathbf{R}_i$. This will not lead to a large increase in the cost of applying its inverse, and we deem that retaining the additional correlation information is likely to be beneficial for the preconditioner.

ETNA
Kent State University and
Johann Radon Institute (RICAM)

5.2. Preconditioning methods motivated by reconditioning. The next two choices of $\widehat{\mathbf{R}}_i$ are motivated by reconditioning methods [39]. These are commonly used in data assimilation implementations to mitigate the issues associated with ill-conditioned sample covariance matrices [45]. These methods do not increase the sparsity of $\widehat{\mathbf{R}}_i$ compared to $\mathbf{R}_i$ but can be applied to non-block matrices, such as the spatially varying error covariances used for Doppler Radar Winds. In this application we consider the use of such methods to develop new terms in the preconditioner only.

Algorithm 1: Ridge regression method.

Inputs: Matrix $\mathbf{R}_i$, target condition number $\kappa_{\max}$.
Define $\gamma = \frac{\lambda_{\max}(\mathbf{R}_i) - \lambda_{\min}(\mathbf{R}_i)\kappa_{\max}}{\kappa_{\max} - 1}$.
Set $\mathbf{R}_{RR} = \mathbf{R}_i + \gamma \mathbf{I}$.

Algorithm 2: Minimum eigenvalue method.

Inputs: Matrix $\mathbf{R}_i = \mathbf{V}\mathbf{\Lambda}\mathbf{V}^\top$, target condition number $\kappa_{\max}$.
Set $\lambda_{\max}(\mathbf{R}_{ME}) = \lambda_{\max}(\mathbf{R}_i)$.
Define $T = \frac{\lambda_{\max}(\mathbf{R}_i)}{\kappa_{\max}} > \lambda_{\min}(\mathbf{R}_i)$.
Set the remaining eigenvalues of $\mathbf{R}_{ME}$ via

$$\lambda_k(\mathbf{R}_{ME}) = \begin{cases} \lambda_k(\mathbf{R}_i) & \text{if } \lambda_k(\mathbf{R}_i) > T, \\ T & \text{if } \lambda_k(\mathbf{R}_i) \leq T. \end{cases}$$

Construct the reconditioned matrix via $\mathbf{R}_{ME} = \mathbf{V}\mathbf{\Lambda}_{ME}\mathbf{V}^\top$, where $\mathbf{\Lambda}_{ME}$ is a diagonal matrix with diagonal entries given by $\lambda_k(\mathbf{R}_{ME})$.

Algorithms 1 and 2 define parameter-dependent preconditioners. Typically in the reconditioning setting, γ and T are selected such that reconditioned matrices have condition number $\kappa_{\max}$. However, in the preconditioning approach we select γ and T directly, with larger parameter values yielding smaller condition numbers of $\widehat{\mathbf{R}}_i$.

Algorithms 1 and 2 can be used to construct preconditioners for $\mathbf{R}_i$ that retain much of the structure of the original matrix but are better conditioned. Additionally, we can prove how the eigenvalues of the preconditioned term that appears in Theorem 3.1 change as we vary γ and T, respectively. The following results determine the spectra of the preconditioned terms $\mathbf{R}_{RR}^{-1}\mathbf{R}_i$ and $\mathbf{R}_{ME}^{-1}\mathbf{R}_i$ for any choice of the parameters γ and T.

PROPOSITION 3. *Let $\lambda_k(\mathbf{R}_i)$ denote the eigenvalues of $\mathbf{R}_i$. The eigenvalues of $\mathbf{R}_{RR}^{-1}\mathbf{R}_i$ are given by*

$$\lambda_k(\mathbf{R}_{RR}^{-1}\mathbf{R}_i) = \frac{\lambda_k(\mathbf{R}_i)}{\lambda_k(\mathbf{R}_i) + \gamma}.$$

Proof. Let $\mathbf{R}_i = \mathbf{V}\mathbf{\Lambda}\mathbf{V}^\top$, with $\mathbf{\Lambda} = \texttt{diag}(\lambda_k)$, be the eigendecomposition of $\mathbf{R}_i$. Then

$$\mathbf{R}_{RR} = \mathbf{V}(\mathbf{\Lambda} + \gamma \mathbf{I})\mathbf{V}^\top,$$

and hence

$$\mathbf{R}_{RR}^{-1}\mathbf{R}_i = \mathbf{V}\left(\texttt{diag}\left(\frac{\lambda_k}{\lambda_k + \gamma}\right)\right)\mathbf{V}^\top. \qquad \square$$

ETNA
Kent State University and
Johann Radon Institute (RICAM)

We note that for any value of γ, $0 < \lambda(\mathbf{R}_{RR}^{-1}\mathbf{R}_i) < 1$. For small values of γ, all eigenvalues are closer to 1, and as γ increases, more eigenvalues move towards zero.

PROPOSITION 4. *The eigenvalues of* $\mathbf{R}_{ME}^{-1}\mathbf{R}_i$ *are given by*

$$\lambda_k(\mathbf{R}_{ME}^{-1}\mathbf{R}_i) = \begin{cases} 1 & \text{if } \lambda_k > T, \\ \frac{\lambda_k}{T} < 1 & \text{if } \lambda_k \leq T. \end{cases}$$

Proof. Let $\mathbf{R}_i = \mathbf{V}\mathbf{\Lambda}\mathbf{V}^\top$, with $\mathbf{\Lambda} = \texttt{diag}(\lambda_k)$, be the eigendecomposition of $\mathbf{R}_i$. Then

$$\mathbf{R}_{ME} = \mathbf{V}\mathbf{\Lambda}_{ME}\mathbf{V}^\top,$$

where $\mathbf{\Lambda}_{ME}$ is a diagonal matrix with diagonal entries given by $\max\{T, \lambda_k(\mathbf{R}_i)\}$. Hence,

$$\mathbf{R}_{ME}^{-1}\mathbf{R}_i = \mathbf{V}\left(\texttt{diag}\left(\frac{\lambda_k}{\max\{T, \lambda_k\}}\right)\right)\mathbf{V}^\top,$$

which has eigenvalues corresponding to the expression in the theorem statement. ∎

All of the eigenvalues of $\mathbf{R}_{ME}^{-1}\mathbf{R}_i$ lie in the range $(0, 1]$. For small values of T, most of the eigenvalues are units, with a few smaller than 1. As T increases, a larger number of eigenvalues become strictly smaller than 1.

Therefore, both choices of preconditioner maintain the ordering of the eigenvalues and yield eigenvalues of the preconditioned matrix lying between 0 and 1. For both approaches, increasing the parameter leads to larger differences between the eigenvalues of the preconditioned matrix and 1. Smaller parameter values may yield choices of $\widehat{\mathbf{R}}$ that are themselves ill-conditioned and hence expensive to evaluate as part of a preconditioner. Therefore, there is a balance to be struck when choosing a parameter value to avoid poor conditioning of either the preconditioner $\widehat{\mathbf{R}}$ or the preconditioned term $\widehat{\mathbf{R}}^{-1}\mathbf{R}$. A natural question is therefore how to select appropriate parameter values and how to implement this automatically. Heuristics for automated parameter selection are discussed in Section 6 for our numerical experiments, but it is likely that some initial investigation would be necessary to identify suitable methods for specific problems of interest. In Section 7 we compare the performance of the four approximations $\mathbf{R}_{diag}$, $\mathbf{R}_{block}$, $\mathbf{R}_{RR}$, and $\mathbf{R}_{ME}$, applied blockwise to each block $\mathbf{R}_i$ of $\mathbf{R}$.

6. Numerical framework.

In this section we introduce the numerical framework for the experiments presented in Section 7. We begin in Section 6.1 by defining the parameters for the data assimilation problem. The same data assimilation framework is used for all experiments. In Section 6.2 we discuss implementation aspects relating to the preconditioners. We note that all results are computed using MATLAB version 2019b on a machine with a 1.8GHz Intel quad-core i7 processor with 15GB RAM on an Ubuntu 20.04.2 LTS operating system.

6.1. Data assimilation parameters.

We now describe the data assimilation problem that is studied in Section 7. The size of the state space s is determined by the spatial discretisation $s = \frac{1}{\Delta x}$. For each choice of s, we fix the observation operator $\mathbf{H}_i \in \mathbb{R}^{p\times s}$ to be the same for all observation times t_i. We usually choose $p = \frac{s}{2}$, and observations of alternate state variables are smoothed equally over 5 adjacent state variables, with entries either being 0 or $\frac{1}{5}$. The full observation operator $\mathbf{H}$ is then assembled by taking $\mathbf{H} = \mathbf{I}_{N+1} \otimes \mathbf{H}_i$, where $\otimes$ denotes the Kronecker product.

ETNA
Kent State University and
Johann Radon Institute (RICAM)

We assume that the model error $\mathbf{Q}_i$ is the same at each observation time, i.e., $\mathbf{Q}_i \equiv \mathbf{Q}$, for $i = 1, \dots, N$. Although this is a simplified choice of model error, we note that treating $\mathbf{Q}_i$ is not the focus of this work. More complicated formulations could be taken into account for operational problems, and the preconditioning approaches for $\mathbf{R}$ discussed in Section 5 can also be applied to the blocks of $\mathbf{Q}$. Both $\mathbf{B}$ and $\mathbf{Q}$ are created using the same routine, based on a SOAR correlation matrix [46]. This routine constructs spatial local correlations whilst ensuring that the matrix has high sparsity. Both $\mathbf{B}$ and $\mathbf{Q}$ are $s \times s$ circulant matrices fully defined by a single row and symmetrised. The number of non-zero entries in each row is fixed irrespective of the value of s. The non-zero entries are computed using a modified SOAR function following the procedure in [22]:

$$c_i = \sigma \left(1 + \frac{2|\sin\left(\frac{i\theta}{2}\right)|}{L} \exp\left(-\frac{2|\sin(\frac{i\theta}{2})|}{L} \right) \right), \qquad \theta = \frac{\pi}{\texttt{maxval}}, \tag{6.1}$$

where L is the correlation lengthscale (0.6 for $\mathbf{B}$, 0.5 for $\mathbf{Q}$), `maxval` determines the number of non-zero entries (100 for $\mathbf{B}$, 120 for $\mathbf{Q}$), and σ is the amplitude of the correlation function (0.4 for $\mathbf{B}$, 0.2 for $\mathbf{Q}$). To ensure positive definiteness of the function, if the smallest eigenvalue of the full matrix is negative, a constant $\delta = |\lambda_{\min}(\mathbf{B})| + \psi$ is added to the diagonal, where ψ is a random number in $(0, 0.5]$. We then assemble $\mathbf{D}$ by taking the first block diagonal entry to be $\mathbf{B}$ and the remaining block diagonal entries to be N copies of $\mathbf{Q}$.

The spatial structure for $\mathbf{B}$ and $\mathbf{Q}$ introduced above is similar to numerical frameworks that have been considered previously for weak-constraint data assimilation experiments [8, 10]. However, we note that for realistic data assimilation systems, the true structure of $\mathbf{Q}_i$ is not well known. Improved understanding of model error covariance matrices and their estimation in preconditioners is of research interest but will not be considered here. For our experiments we apply the ridge regression preconditioner to $\mathbf{B}$ and $\mathbf{Q}$ given by Algorithm 1, for the parameter $\gamma = 0.01$ to obtain $\mathbf{B}_{RR}$ and $\mathbf{Q}_{RR}$. We apply $\widehat{\mathbf{D}}^{-1}$ using the incomplete Cholesky factors of $\mathbf{B}_{RR}$ and $\mathbf{Q}_{RR}$, respectively, computed with the MATLAB function `ichol` with zero-fill, i.e., using the same sparsity structure as $\mathbf{B}$ and $\mathbf{Q}$.

Many experiments that account for correlated observation error covariance matrices use spatial correlations and circulant matrix structures (similar to those we are using for $\mathbf{B}, \mathbf{Q}$). However, in NWP error correlations often arise from hyperspectral satellite-based instruments which have interchannel uncertainty structures (see for instance [37, 45]) that appear as block structures within a matrix. Therefore, for these experiments we construct a matrix with block structure, designed to replicate many of the properties of realistic interchannel error correlations.

To construct $\mathbf{R}_i \in \mathbb{R}^{p \times p}$ we define two vectors: $\texttt{pvec} \in \mathbb{R}^{\texttt{plen}}$ such that $p = \sum_{k=1}^{\texttt{plen}} \texttt{pvec}_k$, which gives the size of the blocks, and $\texttt{pcorr} \in \mathbb{R}^{\texttt{plen}(\texttt{plen}-1)/2}$, which gives the multiplication factor for each off-diagonal block (note that if $\texttt{pcorr}_k = 0$, then the corresponding off-diagonal block is uncorrelated). Diagonal blocks are correlated and constructed as the Hadamard product of a sparse random matrix and a sparse SOAR matrix (using the same approach as for $\mathbf{B}$ and $\mathbf{Q}$ above). Off-diagonal blocks are sparse random matrices with entries in $[0, \texttt{pcorr}_k)$. The matrix $\mathbf{R}_i$ is assembled by adding the diagonal blocks and the upper half of the matrix and then symmetrising. This also increases the weight of the diagonal blocks. A similar approach to that used for $\mathbf{B}$ and $\mathbf{Q}$ is applied to guarantee that $\mathbf{R}_i$ is positive definite, where if the minimum eigenvalue of $\mathbf{R}_i$ is less than zero, it is increased to a small positive value. This value is fixed at 0.41 in order to control the conditioning of $\mathbf{R}_i$ and ensure that the condition numbers of $\mathbf{D}$ and $\mathbf{R}$ are comparable. We assemble $\mathbf{R}$ by taking $\mathbf{R} = \mathbf{I}_{N+1} \otimes \mathbf{R}_i$, where $\otimes$ denotes the Kronecker product, and the choice of $\mathbf{R}_i \in \mathbb{R}^{p \times p}$ is fixed for a given data assimilation problem. This method of

ETNA
Kent State University and
Johann Radon Institute (RICAM)

TABLE 6.1
Summary of experimental design for Section 7. The modified SOAR function is described in (6.1), the block method to construct $\mathbf{R}_i$ is described in Section 6.1, as is the direct observation structure of $\mathbf{H}$. The final column shows the ratio of observations to state variables at each observation time.

Experiment	Model	$\mathbf{D}$	$\mathbf{R}_i$	$\mathbf{H}$	p/s
A	Lorenz 96	Modified SOAR	Block	Direct observations	0.5
B	Heat equation	Modified SOAR	Block	Direct observations	0.25/0.5

TABLE 6.2
Minimum and maximum eigenvalues of the preconditioned $\mathbf{R}_i$ for different choices of $\widehat{\mathbf{R}}_i$ computed using the `eigs` *function in* MATLAB.

	$\mathbf{R}_i$		$\mathbf{R}_{diag}^{-1}\mathbf{R}_i$		$\mathbf{R}_{block}^{-1}\mathbf{R}_i$		$\mathbf{R}_{RR}^{-1}\mathbf{R}_i$		$\mathbf{R}_{ME}^{-1}\mathbf{R}_i$
p	$\lambda_{\min}$	$\lambda_{\max}$	$\lambda_{\min}$	$\lambda_{\max}$	$\lambda_{\min}$	$\lambda_{\max}$	$\lambda_{\min}$	$\lambda_{\max}$	$\lambda_{\min}$
125	0.4100	108.63	0.0131	2.7466	0.6779	1.3221	0.2908	1.0000	0.2602
250	0.4100	200.55	0.0077	3.3466	0.7886	1.2114	0.2908	1.0000	0.6021
500	0.4100	312.05	0.0061	4.4236	0.8889	1.1110	0.2908	1.0000	0.1129
1250	0.4100	540.99	0.0044	5.5440	0.9335	1.0665	0.2908	1.0000	0.1077
2500	0.4100	650.67	0.0055	6.0254	0.7556	1.2444	0.2908	1.0000	0.1302
5000	0.4100	706.94	0.0057	6.8106	0.6058	1.3941	0.2908	1.0000	0.1725

construction ensures that we have adequate sparsity for high-dimensional experiments, that $\mathbf{R}_i$ is positive definite with significant correlations, and that it is well-conditioned. In practice, $\mathbf{R}$ may be very ill-conditioned compared to $\mathbf{D}$, in which case we expect selecting a good choice of $\widehat{\mathbf{R}}$ would be even more vital to ensure fast convergence. The combination of model and data assimilation parameters used in our experiments is summarised in Table 6.1.

Finally, we wish to apply the reconditioning-inspired preconditioner to $\mathbf{R}_i$ in an online way. One way to do this is to use the small eigenvalues of $\mathbf{R}_i$ to select γ and T. For the ridge regression approach, we run the method such that γ is fixed. For the minimum eigenvalue method, we compute the smallest two eigenvalue-eigenvector pairs and set the threshold equal to the second smallest, meaning only a single eigenvalue is changed. The small eigenvalues are computed using `eigs(R,2,'sr')` in MATLAB. For both approaches we use information from $\mathbf{R}_i$, but computing a small number of eigenvalues ensures that this is computationally efficient. Finally, for all correlated choices of $\widehat{\mathbf{R}}_i$, we apply $\widehat{\mathbf{R}}_i^{-1}$ using the incomplete Cholesky factors computed with the MATLAB function `ichol` with the same sparsity structure as $\mathbf{R}_i$. For $\mathbf{R}_{ME}$ this means that the reconditioning method is applied as a low-rank update to the Cholesky factors via the Woodbury identity, as this is more efficient in terms of storage.

Table 6.2 shows the extreme eigenvalues of $\widehat{\mathbf{R}}_i^{-1}\mathbf{R}_i$ for each of the preconditioners discussed in Section 5. We recall that the maximum eigenvalue of $\mathbf{R}_{ME}^{-1}\mathbf{R}_i$ is 1 by definition. We fix the smallest eigenvalue of $\mathbf{R}_i$, $\lambda_{\min}(\mathbf{R}_i) = 0.41$, to ensure that $\mathbf{R}_i$ is well-conditioned. We see that the maximum eigenvalue of $\mathbf{R}_i$ increases with p. Eigenvalues are more extreme for $\mathbf{R}_{diag}^{-1}\mathbf{R}_i$ than for any other preconditioned matrix. Including correlation information is beneficial in terms of the extreme eigenvalues. As the parameter choice for $\mathbf{R}_{RR}$ depends on the smallest eigenvalue of $\mathbf{R}_i$ (which is fixed), the bounds for the minimum and maximum eigenvalues of $\mathbf{R}_{RR}^{-1}\mathbf{R}_i$ do not change with increasing p. The block approach clusters both the minimum and the maximum eigenvalues of $\mathbf{R}_{block}^{-1}\mathbf{R}_i$ at either side of 1, whereas the reconditioning approaches lead to a maximum eigenvalue of the preconditioned matrix which is very close or equal to 1. For the approach used here, where T is set to the second smallest eigenvalue of $\mathbf{R}_i$, all eigenvalues bar the smallest of $\mathbf{R}_{ME}^{-1}\mathbf{R}_i$ are equal to 1.

ETNA
Kent State University and
Johann Radon Institute (RICAM)

6.2. Aspects of numerical linear algebra implementation. In this section we briefly discuss some of the numerical linear algebra aspects of our implementation to run the numerical experiments in Section 7.

We take advantage of the specific matrix structures in (2.2) and store only the non-zero blocks of the matrix, i.e., $\mathbf{R}_i \in \mathbb{R}^{p\times p}$, $\mathbf{B}, \mathbf{Q}_i \in \mathbb{R}^{s\times s}$, and $\mathbf{H}_i \in \mathbb{R}^{p\times s}$. Each of these (relatively) small matrices is stored as a sparse matrix. We also precompute and store $\widehat{\mathbf{R}}_i \in \mathbb{R}^{p\times p}$ prior to the iteration of the Krylov subspace method.

We compute the matrix-vector products $\boldsymbol{\mathcal{A}}\mathbf{v}$ and the preconditioner solves $\boldsymbol{\mathcal{P}}_D^{-1}\mathbf{v}$ and $\boldsymbol{\mathcal{P}}_I^{-1}\mathbf{v}$ via the action of a matrix on a vector rather than building the full matrices. We now describe this process briefly. Note first that

$$\boldsymbol{\mathcal{A}}\mathbf{v} = \boldsymbol{\mathcal{A}}\begin{bmatrix}\mathbf{v}_1\\ \mathbf{v}_2\\ \mathbf{v}_3\end{bmatrix} := \begin{bmatrix}\mathbf{c}_1\\ \mathbf{c}_2\\ \mathbf{c}_3\end{bmatrix} = \begin{bmatrix}\mathbf{D}\mathbf{v}_1 + \mathbf{L}\mathbf{v}_3\\ \mathbf{R}\mathbf{v}_2 + \mathbf{H}\mathbf{v}_3\\ \mathbf{L}^\top\mathbf{v}_1 + \mathbf{H}^\top\mathbf{v}_2\end{bmatrix}.$$

We compute $\mathbf{c}_1, \mathbf{c}_2, \mathbf{c}_3$ separately by looping over the $N+1$ blocks of $\mathbf{D}$, $\mathbf{R}$, and $\mathbf{H}$. The action of $\mathbf{L}$ and $\mathbf{L}^\top$ on a vector is applied via a function. We note that for each evaluation of the matrix-vector product we require one evaluation each of $\mathbf{L}^\top$ and $\mathbf{L}$.

To apply the inverse of the block diagonal preconditioner to a vector, we have that

$$\boldsymbol{\mathcal{P}}_D^{-1}\mathbf{v} = \begin{bmatrix}\widehat{\mathbf{D}}^{-1} & \mathbf{0} & \mathbf{0}\\ \mathbf{0} & \widehat{\mathbf{R}}^{-1} & \mathbf{0}\\ \mathbf{0} & \mathbf{0} & \widehat{\mathbf{S}}^{-1}\end{bmatrix}\begin{bmatrix}\mathbf{v}_1\\ \mathbf{v}_2\\ \mathbf{v}_3\end{bmatrix} = \begin{bmatrix}\widehat{\mathbf{D}}^{-1}\mathbf{v}_1\\ \widehat{\mathbf{R}}^{-1}\mathbf{v}_2\\ \widehat{\mathbf{L}}^{-1}\mathbf{D}\widehat{\mathbf{L}}^{-\top}\mathbf{v}_3\end{bmatrix}.$$

We see that each application of $\boldsymbol{\mathcal{P}}_D^{-1}$ to a vector requires one evaluation each with $\widehat{\mathbf{D}}^{-1}$, $\widehat{\mathbf{R}}^{-1}$, $\widehat{\mathbf{L}}^{-1}$, and $\widehat{\mathbf{L}}^{-\top}$. To apply $\widehat{\mathbf{D}}^{-1}$ and $\widehat{\mathbf{R}}^{-1}$, we loop over the $N+1$ blocks of the matrices, and for $\widehat{\mathbf{L}}^{-1}$ and $\widehat{\mathbf{L}}^{-\top}$, we loop over the k blocks of the relevant sub-matrices. Now,

$$\boldsymbol{\mathcal{P}}_I^{-1}\mathbf{v} = \begin{bmatrix}\mathbf{0} & \mathbf{0} & \widehat{\mathbf{L}}^{-\top}\\ \mathbf{0} & \widehat{\mathbf{R}}^{-1} & \mathbf{0}\\ \widehat{\mathbf{L}}^{-1} & \mathbf{0} & -\widehat{\mathbf{S}}^{-1}\end{bmatrix}\begin{bmatrix}\mathbf{v}_1\\ \mathbf{v}_2\\ \mathbf{v}_3\end{bmatrix} = \begin{bmatrix}\overbrace{\widehat{\mathbf{L}}^{-\top}\mathbf{v}_3}^{\mathbf{c}_1}\\ \widehat{\mathbf{R}}^{-1}\mathbf{v}_2\\ \widehat{\mathbf{L}}^{-1}(\mathbf{v}_1 - \mathbf{D}\mathbf{c}_1)\end{bmatrix}.$$

By computing the first component of $\boldsymbol{\mathcal{P}}_I^{-1}\mathbf{v}$, $\mathbf{c}_1$, prior to the final component, we can apply the inexact constraint preconditioner with the same dominant costs-per-iteration as an application of the block diagonal preconditioner. This saves the cost of a second application of $\widehat{\mathbf{L}}^{-\top}$. An additional advantage of the inexact constraint preconditioner is that it does not require the application of $\widehat{\mathbf{D}}^{-1}$ to a vector and hence may result in significant computational savings for challenging choices of $\mathbf{D}$.

We note that all of the experiments in Section 7 run the Krylov subspace methods to convergence to a relative tolerance of 10^{-6}. This is in contrast to most operational data assimilation implementations, where a small number of fixed iterations are applied. Due to the inclusion of the Lagrange multipliers in the residual of the saddle point formulation, we no longer have monotonic decrease of the objective cost. Running our experiments to convergence ensures a fair comparison between each preconditioner. We leave the investigation of the non-monotonicity for each preconditioner to future work.

ETNA
Kent State University and
Johann Radon Institute (RICAM)

TABLE 7.1
Experiment A: (Top:) Bounds for the negative and positive eigenvalues from Theorem 3.1 for $s = 500$, $N = 5$, $k = 3$ with parameters from Experiment A in Table 6.1; (Bottom:) Extreme negative and positive eigenvalues for this problem computed using the `eigs` *function in* MATLAB.

$\widehat{\mathbf{L}}$	$\widehat{\mathbf{R}}$	$\lambda_{\min}(\boldsymbol{\mathcal{P}}_D^{-1}\boldsymbol{\mathcal{A}})$	$\max(\lambda(\boldsymbol{\mathcal{P}}_D^{-1}\boldsymbol{\mathcal{A}}) < 0)$	$\min(\lambda(\boldsymbol{\mathcal{P}}_D^{-1}\boldsymbol{\mathcal{A}}) > 0)$	$\lambda_{\max}(\boldsymbol{\mathcal{P}}_D^{-1}\boldsymbol{\mathcal{A}})$
$\mathbf{L}_M$	$\mathbf{R}_{block}$	-15.3994	-0.1233	0.7886	16.4061
$\mathbf{L}_M$	$\mathbf{R}_{RR}$	-14.2002	-0.0574	0.2908	14.8536
$\mathbf{L}_M$	$\mathbf{R}$	-13.8536	-0.1773	1.0000	14.8536
$\mathbf{L}_0$	$\mathbf{R}$	-12.2360	-0.0551	1.0000	13.2360
$\mathbf{L}_M$	$\mathbf{R}_{block}$	-3.5516	-0.1773	0.8295	4.5535
$\mathbf{L}_M$	$\mathbf{R}_{RR}$	-2.5351	-0.1773	0.3878	3.2438
$\mathbf{L}_M$	$\mathbf{R}$	-3.6119	-0.1773	1.0000	4.6119
$\mathbf{L}_0$	$\mathbf{R}$	-2.0415	-0.0551	1.0000	3.0415

7. Numerical experiments. In this section we present numerical experiments for the two problems described in Section 6. We use the MINRES implementation of [29] for the block diagonal preconditioner, with a residual-based convergence criterion in the 2-norm. For the inexact constraint preconditioner we use the GMRES implementation of [20] with no restarts and a convergence criterion given by the relative residual in the 2-norm. We use a tolerance of 10^{-6} for both problems of interest. We note that all of our experiments converge within the maximum number of iterations allowed (1000).

7.1. Lorenz 96 model. The main problem of interest concerns the Lorenz 96 model [26], a non-linear problem that is often considered as a test problem for data assimilation applications (see for example [8, 18] for use within the saddle point formulation for data assimilation). The Lorenz 96 model consists of s coupled ordinary differential equations which are discrete in space and continuous in time. For notation purposes, we set $\Delta x = \frac{1}{s}$. For $i = 1, \ldots, s$,

$$\frac{\mathrm{d}x_i}{\mathrm{d}t} = (x_{i+1} - x_{i-2})x_{i-1} - x_i + 8,$$

where we have periodic boundary conditions (so $x_{-1} = x_{s-1}$, $x_0 = x_s$, $x_{s+1} = x_1$). The choice of the forcing constant $F = 8$ induces chaotic behaviour and is typical for data assimilation applications. We use the numerical implementation of [10], where the model is integrated in time with a fourth-order Runge–Kutta scheme. For all experiments we consider $N = 15$ subwindows. We consider $\Delta t = 10^{-4}$ for all the experiments, although similar results were obtained for other values of Δt and are not presented here.

As we are interested in assessing the performance of preconditioners within the linearised inner loops, we consider a single outer loop of the weak-constraint formulation. The Lorenz 96 example can be considered as a study of how well the proposed preconditioners perform in a realistic setting where the linearised model operators $\mathbf{M}_i$ differ in each subwindow. We note that the setting of Proposition 1 holds, and hence using $\mathbf{L}_M$ guarantees that $\mathbf{L}_M^{-\top}\mathbf{L}^{\top}\mathbf{L}\mathbf{L}_M^{-1}$ possesses a number of unit eigenvalues depending on the choice of k. However, we cannot bound the maximum eigenvalue of $\mathbf{L}_M^{-\top}\mathbf{L}^{\top}\mathbf{L}\mathbf{L}_M^{-1}$ using the theory of Section 4 since the assumptions of Proposition 2 are not satisfied, as $\lambda_{\max}(\mathbf{M}_{15}^{\top}\mathbf{M}_{15}) > 1$ for all choices of Δx that were studied.

Table 7.1 shows the values of the bounds from Theorem 3.1 (top four rows) and the computed eigenvalues (bottom four rows) when using $\boldsymbol{\mathcal{P}}_D$ with $\mathbf{L}_M(3)$ and a number of choices for $\widehat{\mathbf{R}}$ introduced in Section 5. We note that these experiments consider $\mathbf{D} = \mathbf{I}$, as the condition number of $\mathbf{D}$ appears in all terms in the bound. In the more realistic case where $\mathbf{B}, \mathbf{Q}_i \neq \mathbf{I}$, we expect the bounds given by Theorem 3.1 to be much weaker. We note

ETNA
Kent State University and
Johann Radon Institute (RICAM)

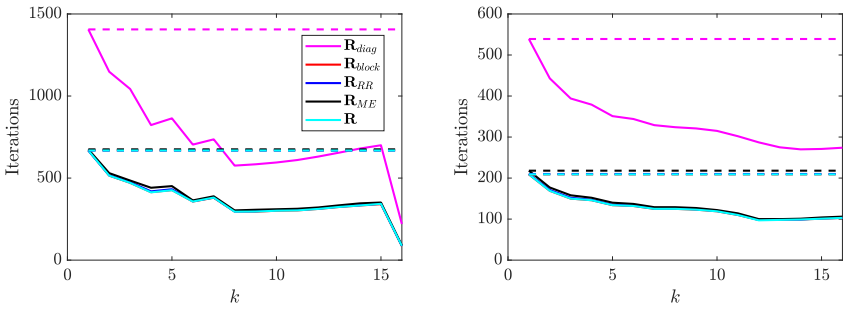

FIG. 7.1. *Experiment A: Number of iterations required for convergence of MINRES for different choices of* $\widehat{\mathbf{R}}$ *within the block diagonal preconditioner* $\boldsymbol{\mathcal{P}}_D$ *(left) and GMRES with the inexact constraint preconditioner* $\boldsymbol{\mathcal{P}}_I$ *(right) for the Lorenz 96 problem, with increasing* k*. The dimension of the problem is given by* $\boldsymbol{\mathcal{A}} \in \mathbb{R}^{100,000\times 100,000}$*. The dashed line shows the results for* $k = 1$*, i.e.,* $\mathbf{L}_M = \mathbf{I} = \mathbf{L}_0$*. We note that* $k = 16$ *corresponds to* $\mathbf{L}_M \equiv \mathbf{L}$*. Legend entries correspond to different choices of* $\widehat{\mathbf{R}}$ *as described in Section 5.*

TABLE 7.2
Experiment A: Total number of matrix-vector products with component matrices for $\boldsymbol{\mathcal{P}}_D$ *for increasing* k *for* $\mathbf{R}_{diag}$ *(left) and* $\mathbf{R}_{block}$ *(right).*

k	$\mathbf{R}_i$	$\mathbf{D}_i$	$\widehat{\mathbf{D}}_i^{-1}$	$\mathbf{M}_i/\mathbf{M}_i^\top$	$\mathbf{R}_i$	$\mathbf{R}_{block}^{-1}$	$\mathbf{D}_i$	$\widehat{\mathbf{D}}_i^{-1}$	$\mathbf{M}_i/\mathbf{M}_i^\top$
1	22496	44992	22496	42180	10704	10704	21408	10704	20070
2	18368	36736	18368	34440	8240	8240	16480	8240	15450
3	16688	33376	16688	47978	7536	7536	15072	7536	21666
4	13168	26336	13168	41150	6624	6624	13248	6624	20700
7	11776	23552	11776	41216	6080	6080	12160	6080	21280
10	9520	19040	9520	34510	4816	4816	9632	4816	17458
16	3520	7040	3520	12760	1376	1376	2752	1376	4988

that even with this choice of $\mathbf{D}$, the bounds provide pessimistic estimates of the eigenvalues. However, the qualitative behaviour of the bounds is similar to that of the computed eigenvalues for different choices of $\widehat{\mathbf{R}}$ and $\mathbf{L}_M$. The eigenvalues of smallest absolute value are mainly affected by the choice of $\mathbf{L}_M$ (for the negative eigenvalue with smallest magnitude) or $\widehat{\mathbf{R}}$ (for the smallest positive eigenvalue). This is not true for the bounds in the case of the negative eigenvalue with smallest magnitude. The computed eigenvalues of largest magnitude and the corresponding bounds are affected by changes to both $\mathbf{L}_M$ and $\widehat{\mathbf{R}}$, but these changes are rather small in all cases. As the largest magnitude computed eigenvalues are small for this problem (all less than five), the improvements to the small magnitude eigenvalues with $\mathbf{L}_M$ and correlated choices of $\widehat{\mathbf{R}}$ are likely to have the most significant influence on the convergence of the iterative methods.

We now study the performance of the proposed choices of preconditioners for the Lorenz 96 problem (Experiment A in Table 6.1). The left panel of Figure 7.1 shows how the number of iterations required for convergence changes with the choice of $\widehat{\mathbf{L}}$ and $\widehat{\mathbf{R}}$ within $\boldsymbol{\mathcal{P}}_D$ as k increases. Including some model information in $\widehat{\mathbf{L}}$ leads to a reduction in iterations compared to $k = 1$ for all choices of $\widehat{\mathbf{R}}$. For $k < 5$ including more model information leads to faster convergence. However, for $k \geq 5$ the change in iterations is non monotonic. Including correlation information in $\widehat{\mathbf{R}}$ results in large improvements to convergence. We note that the lines corresponding to the correlated choices of $\widehat{\mathbf{R}}$ ($\mathbf{R}_{block}$, $\mathbf{R}_{RR}$, and $\mathbf{R}_{ME}$) lie almost directly on top of the line for $\mathbf{R}$. Indeed, if using $\mathbf{R}_{diag}$, we require $k \geq 8$ to obtain fewer iterations than using $\mathbf{L}_0$ with an improved choice of $\widehat{\mathbf{R}}$. There is very little difference in

ETNA
Kent State University and
Johann Radon Institute (RICAM)

TABLE 7.3
Experiment A: Total number of matrix-vector products with component matrices for $\mathcal{P}_I$ for increasing k for $\mathbf{R}_{diag}$ (left) and $\mathbf{R}_{block}$ (right).

k	$\mathbf{R}_i$	$\mathbf{D}_i$	$\mathbf{M}_i/\mathbf{M}_i^\top$	$\mathbf{R}_i$	$\mathbf{R}_{block}^{-1}$	$\mathbf{D}_i$	$\mathbf{M}_i/\mathbf{M}_i^\top$
1	8624	17248	16170	3344	3344	6688	6270
2	7088	14176	13290	2704	2704	5408	5070
3	6304	12608	18124	2400	2400	4800	6900
4	6064	12128	18950	2336	2336	4672	7300
7	5264	10528	18424	2000	2000	4000	7000
10	5040	10080	18270	1904	1904	3808	6902
16	4384	8768	15892	1648	1648	3296	5974

performance between the correlated choices of $\widehat{\mathbf{R}}$. For all choices of $\widehat{\mathbf{R}}$, the smallest number of iterations occurs when using $k = N + 1$, i.e., the exact choice of $\mathbf{L}$. Table 7.2 shows the number of matrix-vector products required to reach convergence for the block diagonal preconditioner using $\mathbf{R}_{diag}$ and $\mathbf{R}_{block}$. Results using $\mathbf{R}_{RR}$ and $\mathbf{R}_{ME}$ are similar to those with $\mathbf{R}_{block}$. The number of matrix-vector products with $\mathbf{D}_i$, $\widehat{\mathbf{D}}_i^{-1}$, and $\mathbf{M}_i$ is reduced when using $\mathbf{R}_{block}$ compared to $\mathbf{R}_{diag}$. For some choices of k, the total number of evaluations with $\mathbf{R}_i$ and $\mathbf{R}_{block}^{-1}$ is slightly larger than in the $\mathbf{R}_{diag}$ case. Increasing k broadly decreases the number of matrix-vector products with the error covariance matrices and their inverses. For some choices of $k > 1$, more evaluations of $\mathbf{M}_i$ and $\mathbf{M}_i^\top$ are required than when using $\mathbf{L}_0$. However, this increase is small compared to the decrease in the other components.

The right panel of Figure 7.1 shows how the number of iterations required for convergence changes with the choice of $\widehat{\mathbf{L}}$ and $\widehat{\mathbf{R}}$ within $\mathcal{P}_I$ as k increases. We see a clear benefit of including model information in terms of a reduction in iterations. Increasing k leads to a reduction in the number of iterations required for convergence when using $\mathbf{L}_M$ for all choices of $\widehat{\mathbf{R}}$, unlike when using $\mathcal{P}_D$. The benefit of using an improved estimate of $\widehat{\mathbf{R}}$ is even more stark, with any choice of correlated $\widehat{\mathbf{R}}$ and $\mathbf{L}_0$ leading to fewer iterations than $\mathbf{R}_{diag}$ even when using $\widehat{\mathbf{L}} \equiv \mathbf{L}$. Again, the lines for the correlated choices of $\widehat{\mathbf{R}}$ ($\mathbf{R}_{block}$, $\mathbf{R}_{RR}$, and $\mathbf{R}_{ME}$) lie almost directly on top of the line for $\mathbf{R}$. Table 7.3 shows the total number of matrix-vector products required to reach convergence for the inexact constraint preconditioner for $\mathbf{R}_{diag}$ and $\mathbf{R}_{block}$. Results using $\mathbf{R}_{RR}$ and $\mathbf{R}_{ME}$ are similar to those with $\mathbf{R}_{block}$. In this case, using $\mathbf{R}_{block}$ leads to a large reduction in the number of matrix-vector products for all components. Increasing k can lead to increases in the number of model matrix-vector evaluations but generally leads to a decrease in the total number of matrix-vector products.

Table 7.4 shows the performance of the block diagonal preconditioner and the inexact constraint preconditioner, respectively, for a higher-dimensional problem when using $\mathbf{R}_{block}$, $\mathbf{R}_{RR}$, and $\mathbf{R}$ itself to approximate $\mathbf{R}$ within the preconditioner. Similarly to the smaller-dimensional problem considered in Figure 7.1, using $\mathbf{L}_M$ leads to improved convergence in terms of iterations compared to $\mathbf{L}_0$. The different choices of $\widehat{\mathbf{R}}$ lead to comparable iteration numbers, and we recall that $\mathbf{R}_{block}$ has additional sparsity structure. Increasing k leads to a slight reduction in the number of iterations, but increases the computational cost of each iteration. For this problem, choosing $k = 3$ or 4 allows decreased iteration counts compared to $\mathbf{L}_0$, without too many more matrix-vector products with $\mathbf{M}_i$. Iteration counts are much smaller for the inexact constraint preconditioner than the block diagonal preconditioner. Overall, using our new preconditioners $\mathbf{L}_M$ and correlated choices of $\widehat{\mathbf{R}}$ result in fewer iterations and matrix-vector products compared to those obtained when using $\mathbf{L}_0$ or $\mathbf{R}_{diag}$ for the Lorenz 96 problem.

ETNA
Kent State University and
Johann Radon Institute (RICAM)

TABLE 7.4
Experiment A: Number of iterations required for convergence of MINRES with the block diagonal preconditioner $\mathcal{P}_D$ (left) and GMRES with $\mathcal{P}_I$ (right) applied to the Lorenz 96 problem, using $\mathbf{R}_{block}$, $\mathbf{R}_{RR}$, $\mathbf{R}$ in combination with $\mathbf{L}_0$, $\mathbf{L}_M$ ($k = 3, 4, 5$). Here, $\boldsymbol{\mathcal{A}} \in \mathbb{R}^{1,600,000\times 1,600,000}$.

	$\mathbf{R}_{block}$	$\mathbf{R}_{RR}$	$\mathbf{R}$	$\mathbf{R}_{block}$	$\mathbf{R}_{RR}$	$\mathbf{R}$
$\mathbf{L}_0$	759	822	822	359	275	275
$\mathbf{L}_M, k = 3$	433	466	467	244	205	205
$\mathbf{L}_M, k = 4$	348	335	336	228	200	200
$\mathbf{L}_M, k = 5$	367	354	355	206	182	182

7.2. Heat equation with Dirichlet boundary conditions. The second problem of interest that we consider here is the one-dimensional heat equation

$$\frac{\partial u}{\partial t} = \alpha \frac{\partial^2 u}{\partial x^2} \tag{7.1}$$

with homogeneous Dirichlet boundary conditions. We discretise (7.1) using the forward Euler method in time and second-order centred differences in space. This means we can write the model evolution in matrix form for a single time model step as $\mathbf{u}^{t+\Delta t} = \mathbf{M}_{\Delta t}\mathbf{u}^t$, where $\mathbf{M}_{\Delta t}$ denotes the application of a single model time-step of length Δt to the heat equation with Dirichlet boundary conditions; this is given by

$$\mathbf{M}_{\Delta t} = \begin{bmatrix} 0 & 0 & 0 & 0 & \cdots & 0 \\ 0 & 1-2r & r & 0 & \cdots & 0 \\ 0 & r & 1-2r & \ddots & & \vdots \\ 0 & 0 & \ddots & \ddots & r & 0 \\ \vdots & \vdots & & r & 1-2r & 0 \\ 0 & 0 & \cdots & 0 & 0 & 0 \end{bmatrix},$$

where $r = \frac{\alpha \Delta t}{(\Delta x)^2}$. The case of non-homogeneous boundary conditions would follow similarly, by applying a source term to the model evolution equation. For the numerical experiments presented here we fix $\alpha = 1$ and vary spatial and temporal resolutions together, setting the ratio $r = \frac{\Delta t}{(\Delta x)^2} = 0.4$ for all experiments.

One advantage of the heat equation test problem is the ability to consider how our new preconditioners scale with the problem size. We now consider the best-performing preconditioners for a high-dimensional example, namely the inexact constraint preconditioner for $\mathbf{R}_{block}, \mathbf{R}_{RR}$, and $\mathbf{R}$. Experiment B studies the behaviour of the inexact constraint preconditioner for high-dimensional problems. In particular, for $\Delta x = 2\times 10^{-5}$, the dimension of the full saddle point problem is $750,000 \times 750,000$ for $s = 2p$. In Figure 7.2 we only consider the inexact constraint preconditioner and the three best choices of $\widehat{\mathbf{R}}$. We find that even for a high-dimensional problem the number of iterations is small, with only a modest difference between the results for $\mathbf{R}_{block}$ and $\mathbf{R}$ (as well as $\mathbf{R}_{RR}$). Similarly to the lower-dimensional case, using $\mathbf{L}_M$ requires fewer iterations than using $\mathbf{L}_0$. Figure 7.2 also considers two observation networks. We see that using a larger number of observations requires a larger number of iterations to reach convergence, which coincides with the findings of [8] for the unpreconditioned case. However, the qualitative behaviour across the different choices of $\widehat{\mathbf{L}}$ and $\widehat{\mathbf{R}}$ is the same for both observation networks.

ETNA
Kent State University and
Johann Radon Institute (RICAM)

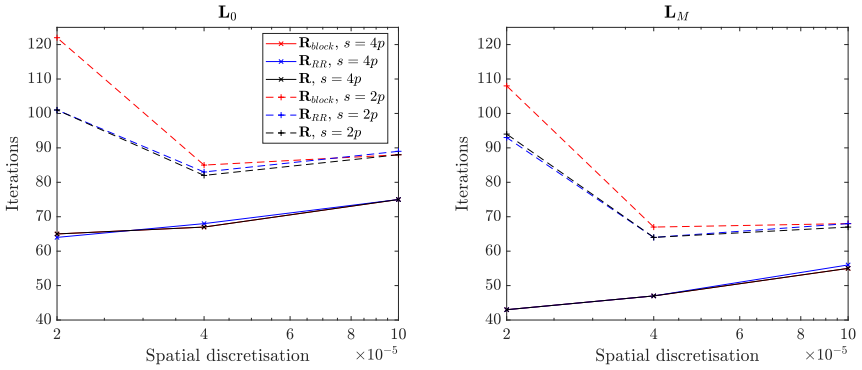

FIG. 7.2. *Experiment B: Number of iterations required for convergence using* $\mathbf{L}_0$ *(left) and* $\mathbf{L}_M$ *(right) for different choices of* $\widehat{\mathbf{R}}$ *within the inexact constraint preconditioner* $\boldsymbol{\mathcal{P}}_I$ *for the heat equation, for large choices of* s*. Two observation ratios are shown:* $p = \frac{s}{2}$ *(dashed line) and* $p = \frac{s}{4}$ *(solid line). The dimension of the problem ranges from* $\boldsymbol{\mathcal{A}} \in \mathbb{R}^{150,000\times 150,000}$ *(right) to* $\boldsymbol{\mathcal{A}} \in \mathbb{R}^{750,000\times 750,000}$ *(left) for* $p = \frac{s}{2}$*, and* $\boldsymbol{\mathcal{A}} \in \mathbb{R}^{135,000\times 135,000}$ *(right) to* $\boldsymbol{\mathcal{A}} \in \mathbb{R}^{675,000\times 675,000}$ *(left) for* $p = \frac{s}{4}$*. Legend entries correspond to different choices of* $\widehat{\mathbf{R}}$ *as described in Section 5.*

8. Conclusions. We proposed new preconditioners for the saddle point formulation of the weak-constraint 4D-Var data assimilation problem in the presence of correlated observation errors. Our approach for approximating the model term $\widehat{\mathbf{L}}$ incorporated model information for the first time. We also proposed a range of approaches that permit inclusion of computationally efficient correlation information within the observation error covariance term $\widehat{\mathbf{R}}$. In summary:

- We developed new bounds for the eigenvalues of the preconditioned saddle point system in the case of a block diagonal preconditioner.
- We investigated how the constituent terms within the bounds behave for existing and proposed choices of $\widehat{\mathbf{L}}$ and $\widehat{\mathbf{R}}$. Including model information via $\mathbf{L}_M$ yields many repeated unit eigenvalues of $\mathbf{L}_M^{-\top}\mathbf{L}^{\top}\mathbf{L}\mathbf{L}_M^{-1}$. Our new approaches yield eigenvalues of this matrix that are frequently bounded above by moderate numbers.
- We considered two numerical examples: the Lorenz 96 problem and the heat equation. Including model information via $\mathbf{L}_M$ reduced the iterations for both problems.

The use of preconditioners that account for correlated observation information led to a significant reduction in the iterations for all experiments. For many problems where $\mathbf{R}$ is very ill-conditioned, we would expect the improvements in performance to be even higher than in the experiments presented here. We find that including additional model information in $\mathbf{L}_M$ leads to reduced iterations but increases the computational expense of each iteration. We therefore suggest that selecting $k = 3$ or $k = 4$ represents a sensible trade-off. Future work for this problem includes developing efficient approximations of $\mathbf{D}$, multi-core implementations of our new preconditioners, and experiments within a full-scale operational NWP system.

Acknowledgements. We thank Adam El-Said for his code for the Lorenz 96 weak-constraint 4D-Var assimilation problem. We gratefully acknowledge funding from the Engineering and Physical Sciences Research Council (EPSRC) grant EP/S027785/1 and the London Mathematical Society Research in Pairs scheme.

Appendix A. Block preconditioner for $\widehat{\mathbf{R}}$.

In Algorithm 3, we state the algorithm used to apply the block preconditioner $\mathbf{R}_{block}$ for $\mathbf{R}$ as described in Section 5.1.

ETNA
Kent State University and
Johann Radon Institute (RICAM)

Algorithm 3: Block preconditioner for $\widehat{\mathbf{R}}$.

Inputs: $\mathbf{R}_i$, `pvec` = vector of block sizes, `tol` = tolerance for retaining blocks, `maxsize` = maximum size permitted on a single processor, `numproc` = number of available processors.
Compute `p = sum(pvec)`, `pn` = |`pvec`|.
Define `pst` = starting index for each new block.
Initialise $\mathbf{R}_{block} = \mathbf{R}_i$.
for $j =$ `1:pn-1`
  Compute scaled Frobenius norm of super-diagonal blocks via
  `normvec(j) = 1/sqrt(pvec(j)*pvec(j+1))*`
    `norm(R(pst(j):pst(j+1)-1,pst(j+1):pst(j+2)-1),'fro').`
end
Retain blocks where `normvec(j) >= tol`:
for $j =$ `1:pn-1`
  if `normvec(j) < tol`
    Set $\mathbf{R}_{block}$`(pst(j+1):p,1:pst(j+1)-1) = 0.`
    Set $\mathbf{R}_{block}$`(1:pst(j+1)-1,pst(j+1):p) = 0.`
  end
end
if size of largest block $>$ `maxsize`
  Split largest block into two components.
end
if number of distinct blocks $>$ `numproc`
  Combine two smallest adjacent blocks in $\mathbf{R}_{block}$.
elseif number of distinct blocks $<$ `numproc` $- 2$
  Split largest block of $\mathbf{R}_{block}$ into two components.
end

REFERENCES

[1] M. BENZI, G. H. GOLUB, AND J. LIESEN, *Numerical solution of saddle point problems*, Acta Numer., 14 (2005), pp. 1–137.
[2] L. BERGAMASCHI, J. GONDZIO, M. VENTURIN, AND G. ZILLI, *Inexact constraint preconditioners for linear systems arising in interior point methods*, Comput. Optim. Appl., 36 (2007), pp. 137–147.
[3] ———, *Erratum to: Inexact constraint preconditioners for linear systems arising in interior point methods*, Comput. Optim. Appl., 49 (2011), pp. 401–406.
[4] D. S. BERNSTEIN, *Matrix Mathematics*, 2nd ed., Princeton University Press, Princeton, 2009.
[5] M. BONAVITA, Y. TRÉMOLET, E. HOLM, S. T. K. LANG, M. CHRUST, M. JANISKOVÁ, P. LOPEZ, P. LALOYAUX, P. DE ROSNAY, M. FISHER, M. HAMRUD, AND S. ENGLISH, *A strategy for data assimilation*, Tech. Memo., European Centre for Medium-Range Weather Forecasts, Reading, 2017.
[6] A. CARRASSI, M. BOCQUET, L. BERTINO, AND G. EVENSEN, *Data assimilation in the geosciences: An overview of methods, issues, and perspectives*, WIREs Climate Change, 9 (2018), Art. e535, 50 pages.
[7] E. S. COOPER, S. L. DANCE, J. GARCIA-PINTADO, N. K. NICHOLS, AND P. J. SMITH, *Observation impact, domain length and parameter estimation in data assimilation for flood forecasting*, Env. Model. & Softw., 104 (2018), pp. 199–214.
[8] I. DAUŽICKAITĖ, A. S. LAWLESS, J. A. SCOTT, AND P. J. VAN LEEUWEN, *Spectral estimates for saddle point matrices arising in weak constraint four-dimensional variational data assimilation*, Numer. Linear Algebra Appl., 27 (2020), Art. e2313, 24 pages.
[9] A. EL AKKRAOUI, Y. TRÉMOLET, AND R. TODLING, *Preconditioning of variational data assimilation and the use of a bi-conjugate gradient method*, Quart. J. Royal Meteorol. Soc. A, 139 (2013), pp. 731–741.
[10] A. EL-SAID, *Conditioning of the Weak-Constraint Variational Data Assimilation Problem for Numerical Weather Prediction*, PhD. Thesis, University of Reading, Reading, 2015.

ETNA
Kent State University and
Johann Radon Institute (RICAM)

[11] H. C. ELMAN, D. J. SILVESTER, AND A. J. WATHEN, *Finite Elements and Fast Iterative Solvers: With Applications in Incompressible Fluid Dynamics*, 2nd ed., Oxford University Press, Oxford, 2014.

[12] M. FISHER, S. GRATTON, S. GÜROL, Y. TRÉMOLET, AND X. VASSEUR, *Low rank updates in preconditioning the saddle point systems arising from data assimilation problems*, Optim. Methods Softw., 33 (2018), pp. 45–69.

[13] M. FISHER AND S. GÜROL, *Parallelization in the time dimension of four-dimensional variational data assimilation*, Quart. J. Royal Meteorol. Soc. B, 143 (2017), pp. 1136–1147.

[14] M. A. FREITAG AND D. L. H. GREEN, *A low-rank approach to the solution of weak constraint variational data assimilation problems*, J. Comput. Phys., 357 (2018), pp. 263–281.

[15] S. GRATTON, S. GÜROL, E. SIMON, AND P. L. TOINT, *Guaranteeing the convergence of the saddle formulation for weakly constrained 4D-Var data assimilation*, Quart. J. Royal Meteorol. Soc. B, 144 (2018), pp. 2592–2602.

[16] S. GRATTON, A. S. LAWLESS, AND N. K. NICHOLS, *Approximate Gauss-Newton methods for nonlinear least squares problems*, SIAM J. Optim., 18 (2007), pp. 106–132.

[17] S. GRATTON, P. L. TOINT, AND J. TSHIMANGA ILUNGA, *Range-space variants and inexact matrix-vector products in Krylov solvers for linear systems arising from inverse problems*, SIAM J. Matrix Anal. Appl., 32 (2011), pp. 969–986.

[18] D. GREEN, *Model Order Reduction for Large-Scale Data Assimilation Problems*, PhD. Thesis, University of Bath, Bath, 2019.

[19] A. GREENBAUM, V. PTÁK, AND Z. STRAKOŠ, *Any nonincreasing convergence curve is possible for GMRES*, SIAM J. Matrix Anal. Appl., 17 (1996), pp. 465–469.

[20] C. GREIF, T. REES, AND D. B. SZYLD, *GMRES with multiple preconditioners*, SeMA J., 74 (2017), pp. 213–231.

[21] S. GÜROL, A. T. WEAVER, A. M. MOORE, A. PIACENTINI, H. G. ARANGO, AND S. GRATTON, *B-preconditioned minimization algorithms for variational data assimilation with the dual formulation*, Quart. J. Royal Meteorol. Soc. B, 140 (2014), pp. 539–556.

[22] S. A. HABEN, *Conditioning and Preconditioning of the Minimisation Problem in Variational Data Assimilation*, PhD. Thesis, Department of Mathematics and Statistics, University of Reading, Reading, 2011.

[23] M. R. HESTENES AND E. STIEFEL, *Methods of conjugate gradients for solving linear systems*, J. Research Nat. Bur. Standards, 49 (1952), pp. 409–436.

[24] R. A. HORN AND C. R. JOHNSON, *Matrix Analysis*, Cambridge University Press, Cambridge, 1985.

[25] Y. A. KUZNETSOV, *Efficient iterative solvers for elliptic finite element problems on nonmatching grids*, Russian J. Numer. Anal. Math. Modelling, 10 (1995), pp. 187–211.

[26] E. N. LORENZ, *Predictability: a problem partly solved*, in Seminar on Predictability, 4–8 September 1995, ECMWF, Reading, 1995.

[27] M. F. MURPHY, G. H. GOLUB, AND A. J. WATHEN, *A note on preconditioning for indefinite linear systems*, SIAM J. Sci. Comput., 21 (2000), pp. 1969–1972.

[28] C. C. PAIGE AND M. A. SAUNDERS, *Solutions of sparse indefinite systems of linear equations*, SIAM J. Numer. Anal., 12 (1975), pp. 617–629.

[29] C. C. PAIGE, M. A. SAUNDERS, S.-C. CHOI, D. ORBAN, U. E. VILLA, D. MADDIX, AND S. REGEV, *MINRES: Sparse symmetric equations*, Software, 2020. Available at https://web.stanford.edu/group/SOL/software/minres/.

[30] E. M. PINNINGTON, E. CASELLA, S. L. DANCE, A. S. LAWLESS, J. I. L. MORISON, N. K. NICHOLS, M. WILKINSON, AND T. L. QUAIFE, *Investigating the role of prior and observation error correlations in improving a model forecast of forest carbon balance using Four-dimensional Variational data assimilation*, Agricult. Forest Meteorol., 228–229 (2016), pp. 299–314.

[31] ———, *Understanding the effect of disturbance from selective felling on the carbon dynamics of a managed woodland by combining observations with model predictions*, J. Geophys. Res.: Biogeosciences, 122 (2017), pp. 886–902.

[32] F. RAWLINS, S. P. BALLARD, K. J. BOVIS, A. M. CLAYTON, D. LI, G. W. INVERARITY, A. C. LORENC, AND T. J. PAYNE, *The Met Office global four-dimensional variational data assimilation scheme*, Quart. J. Royal Meteorol. Soc. B, 133 (2007), pp. 347–362.

[33] T. REES, *Preconditioning Iterative Methods for PDE Constrained Optimization*, PhD. Thesis, New College, University of Oxford, Oxford, 2010.

[34] T. REES AND A. J. WATHEN, *Preconditioning iterative methods for the optimal control of the Stokes equations*, SIAM J. Sci. Comput., 33 (2011), pp. 2903–2926.

[35] Y. SAAD AND M. H. SCHULTZ, *GMRES: a generalized minimal residual algorithm for solving nonsymmetric linear systems*, SIAM J. Sci. Statist. Comput., 7 (1986), pp. 856–869.

[36] D. SIMONIN, J. A. WALLER, S. P. BALLARD, S. L. DANCE, AND N. K. NICHOLS, *A pragmatic strategy for implementing spatially correlated observation errors in an operational system: An application to Doppler radial winds*, Quart. J. Royal Meteorol. Soc. B, 145 (2019), pp. 2772–2790.

ETNA
Kent State University and
Johann Radon Institute (RICAM)

[37] L. M. STEWART, S. L. DANCE, AND N. K. NICHOLS, *Correlated observation errors in data assimilation*, Internat. J. Numer. Methods Fluids, 56 (2008), pp. 1521–1527.

[38] J. M. TABEART, S. L. DANCE, S. A. HABEN, A. S. LAWLESS, N. K. NICHOLS, AND J. A. WALLER, *The conditioning of least-squares problems in variational data assimilation*, Numer. Linear Algebra Appl., 25 (2018), Art. e2165, 22 pages.

[39] J. M. TABEART, S. L. DANCE, A. S. LAWLESS, N. K. NICHOLS, AND J. A. WALLER, *Improving the condition number of estimated covariance matrices*, Tellus A: Dyn. Meteorol. Oceanogr., 72 (2020), Art. 1696646, 19 pages.

[40] J. M. TABEART, S. L. DANCE, A. S. LAWLESS, N. K. NICHOLS, AND J. A. WALLER, *New bounds on the condition number of the Hessian of the preconditioned variational data assimilation problem*, Numer. Linear Algebra Appl., 29 (2022), Art. e2405, 19 pages.

[41] P. TANDEO, P. AILLIOT, M. BOCQUET, A. CARRASSI, T. MIYOSHI, M. PULIDO, AND Y. ZHEN, *A review of innovation-based methods to jointly estimate model and observation error covariance matrices in ensemble data assimilation*, Mon. Wea. Rev., 148 (2020), pp. 3973–3994.

[42] Y. TRÉMOLET, *Accounting for an imperfect model in 4D-Var*, Quart. J. Royal Meteorol. Soc. B, 132 (2006), pp. 2483–2504.

[43] ———, *Model-error estimation in 4D-Var*, Quart. J. Royal Meteorol. Soc. A, 133 (2007), pp. 1267–1280.

[44] J. A. WALLER, J. GARCÍA-PINTADO, D. C. MASON, S. L. DANCE, AND N. K. NICHOLS, *Assessment of observation quality for data assimilation in flood models*, Hydrol. Earth Sys. Sci., 22 (2018), pp. 3983–3992.

[45] P. P. WESTON, W. BELL, AND J. R. EYRE, *Accounting for correlated error in the assimilation of high-resolution sounder data*, Quart. J. Royal Meteorol. Soc. B, 140 (2014), pp. 2420–2429.

[46] A. M. YAGLOM, *Correlation Theory of Stationary and Related Random Functions, Volume I: Basic Results*, Springer, New York, 1987.