NUMERICAL LINEAR ALGEBRA WITH APPLICATIONS
Numer. Linear Algebra Appl. 0000; **00**:1–29
Published online in Wiley InterScience (www.interscience.wiley.com). DOI: 10.1002/nla

A generalized predictive analysis tool for multigrid methods

S. Friedhoff[1,2,*] and S. MacLachlan[1,3]

[1]*Department of Mathematics, Tufts University, 503 Boston Avenue, Medford, MA 02155, USA.*
[2]*Now at: Department of Computer Science, KU Leuven, Celestijnenlaan 200a - box 2402, 3001 Leuven, Belgium.*
[3]*Now at: Department of Mathematics and Statistics, Memorial University of Newfoundland, St. John's, NL, Canada.*

SUMMARY

Multigrid and related multilevel methods are the approaches of choice for solving the linear systems that result from discretization of a wide class of PDEs. A large gap, however, exists between the theoretical analysis of these algorithms and their actual performance. This paper focuses on the extension of the well-known local mode (often local Fourier) analysis approach to a wider class of problems. The semi-algebraic mode analysis (SAMA) proposed here couples standard local Fourier analysis approaches with algebraic computation to enable analysis of a wider class of problems, including those with strong advective character. The predictive nature of SAMA is demonstrated by applying it to the parabolic diffusion equation in one and two space dimensions, elliptic diffusion in layered media, as well as a two-dimensional convection-diffusion problem. These examples show that accounting for boundary conditions and heterogeneity enables accurate predictions of the short-term and asymptotic convergence behavior for multigrid and related multilevel methods. Copyright © 0000 John Wiley & Sons, Ltd.

Received ...

KEY WORDS: Local Fourier Analysis; Multigrid Methods; Space-Time discretizations

1. INTRODUCTION

The use of local mode and local Fourier analysis to estimate the performance of multigrid methods has a history dating back to the seminal work of Brandt in 1977 [1]. The idea behind this analysis is that global features, such as boundary conditions or interfaces, often play less of a role in determining the dominant convergence behavior and, as such, a better approach is to focus on the local character away from these features. Indeed, local mode analysis can be made rigorous if sufficient boundary relaxation, for example, is included in the multigrid algorithm for an elliptic problem [2,3], or if the matrices involved in the computation naturally satisfy appropriate boundary conditions [4]. As such, local mode analysis has been used in many contexts, including

- to examine relaxation strategies for strongly coupled systems of PDEs, including saddle-point problems [5–9];
- for methods on triangular and hexahedral grids, including semi-structured mesh approaches [10–14];
- analysing multigrid methods for stochastic PDEs [15];
- investigating the performance of preconditioned GMRES [16]; and

*Correspondence to: stephanie.friedhoff@alumni.tufts.edu

Contract/grant sponsor: This work was supported by the National Science Foundation under grant DMS-1015370 and partially supported by an NSERC discovery grant.

- designing solvers for discontinuous Galerkin discretizations [17].

The main downsides to local mode analysis are the obvious ones. First, there are some problems for which the details of the boundary conditions applied are crucial to the observed performance and, as such, neglecting them renders the analysis unhelpful. This is the case, for example, in applying local Fourier analysis to convection-dominated and parabolic problems [18–22]. Secondly, the analysis relies on (multilevel) Toeplitz or circulant matrix structure, greatly limiting its applicability. In particular, while a key advantage of the multigrid family of methodologies is their ability to robustly account for heterogeneous operators, such heterogeneity is poorly accounted for in local mode analysis. In this paper, we introduce a generalization of the local mode analysis technique that addresses both of these drawbacks, although not in full generality.

The classic view of applying multigrid to parabolic (space-time) problems is based on a time-marching approach: discretization of the PDE leads to a discrete elliptic problem at each time step when an implicit scheme is used for the time derivative. Multigrid is then used as an iterative solver for these elliptic equations. Parallelization in this approach is limited to parallelization in the elliptic (spatial) solver, since the time-stepping procedure is sequential. Simultaneously solving for multiple time steps increases the potential for parallelism in the solution process, but because time is sequential by nature, this idea is not intuitive. Yet, it is possible with work on this topic going back to as early as 1964 [23]. Research integrating multigrid concepts into time-evolution algorithms started about 30 years ago with "parabolic multigrid", a method proposed by Hackbusch in 1984 [24] that uses Jacobi-like parallel-in-time relaxation. Other "time-parallel" methods based on these ideas include [25–28]. In 1987, Lubich and Ostermann develop multigrid with "waveform relaxation" [29], a forward-in-time Gauss-Seidel approach forming the basis of many parallel multigrid approaches including [21, 30–33]. Current trends in computer architectures are leading towards systems with more, but not faster processors. As a consequence, faster compute speeds must come from greater parallelism resurging the interest in multigrid algorithms for parabolic problems that allow temporal parallelism [34, 35]. Regarding this development, a predictive analysis tool for these methods becomes highly relevant.

The primary focus of mode analysis for parabolic or convection-dominated problems has been on the "half-space" viewpoint [18–20,36–39], considering convergence on a discrete half-plane (in 2D) instead of the full infinite lattice used in local Fourier analysis. The half-space mode analysis point-of-view, first discussed in [18, 19] and used for analysis of convective flow problems in [20, 37, 38], considers the evolution of a Fourier mode posed on the inflow boundary over a half-space through each step of the multigrid process. By considering the maximum change in magnitude seen over a discrete set of Fourier modes, an estimate of the operator norm is derived. For waveform relaxation, considered in [36, 39], a related analysis technique is considered, based on the use of discrete Laplace transforms, carrying over from the continuous-time analysis in [40]. In combination with a Fourier expansion of the error in the non-advective direction, this again provides an estimate of the operator norm by maximizing over a discrete set of Fourier frequencies. A third applicable analysis approach is the "idealized relaxation/coarse-grid correction" approach proposed in [41]. This approach is readily applied to a wide range of problems; however, this broad applicability comes at the expense of being a non-rigorous approach. Here, we seek an approach that matches the near-rigorous performance of LFA, but which is applicable to a smaller selection of problems than the approach in [41].

Elliptic diffusion problems with discontinuous coefficients arise in many applications in computational science and engineering. When modeling diffusive groundwater flow, for example, domains are strongly heterogeneous, with large contrasts in permeability between materials such as sand and shale. While various multigrid and related multilevel methods such as the Black Box Multigrid (BoxMG) solver [42, 43] and deflation-based preconditioners [44, 45] are able to account for discontinuities in the permeability, LFA is not valid; leaving a large gap between the theoretical analysis of these algorithms and their actual performance.

The key approach proposed here is to mix analytical tools, such as LFA, with tractable computation to achieve a middle ground, avoiding over-generalization that ignores key problem features and yields poor predictions of performance but also realizing that direct spectral analysis of

DOI: 10.1002/nla

many interesting problems will always be out-of-reach. The proposed semi-algebraic mode analysis (SAMA) can be viewed as an essential generalization to the wide literature on local mode and local Fourier analysis, in that both of these exist as identifiable special cases of the more general technique. Equally important is the identification of a range of problems for which the analysis can be made rigorous, giving exact statements about the convergence of two-grid algorithms.

The remainder of this paper is laid out as follows. In §2, we describe two multigrid schemes for a parabolic model problem. One method based on waveform relaxation as in [21] and one based on temporal semicoarsening as in [35]. In §3, we illustrate the failure of LFA to generate accurate predictions of the convergence behavior for the waveform relaxation method applied to the parabolic model problem. Section 4 presents the key ideas of the semi-algebraic analysis proposed here, where each theoretical point is backed up by calculations for the model problem. Additionally, we include a discussion of the class of problems and multigrid methods for which the SAMA approach is rigorous. Section 5 focuses on the non-parabolic case, and is prefaced by a comment on the applicability of the analysis in that case. Conclusions are presented in §6.

2. MULTIGRID METHODS ON SPACE-TIME GRIDS FOR PARABOLIC PROBLEMS

2.1. The parabolic model problem in space-time

As a motivating example, we consider the diffusion equation in one space dimension,

$$u_t = u_{xx}, \quad x \in (a,b),\ t \in [0,T]; \tag{1}$$

further examples are discussed in Section 5. We prescribe u at $t = 0$, $u(x,0) = u_0(x)$, and impose either the periodicity condition

$$u(x \pm (b-a), t) \equiv u(x,t)$$

or homogeneous Dirichlet boundary conditions,

$$u(a,t) = u(b,t) = 0.$$

We discretize our parabolic model problem on a rectangular space-time grid consisting of N_x space intervals and N_t time intervals, using a spatial mesh size $h_x = (b-a)/N_x$, and a time step size $h_t = T/N_t$, where N_x and N_t are positive integers, and T denotes the final time. For the periodic case, the discretization results in a linear system of the form

$$A\mathbf{u} = (J \otimes I_{N_x} + I_{N_t} \otimes Q)\mathbf{u} = \mathbf{f}, \tag{2}$$

where $\mathbf{f} = (1/h_t)(u_0(x_{-N_x/2+1}), \ldots, u_0(x_{N_x/2}), 0, \ldots, 0)^T$ incorporates the discretized initial condition, in the unknowns $u_{i,j},\ i = -\frac{N_x}{2}+1, \ldots, \frac{N_x}{2}$, and $j = 1, \ldots, N_t$, I_{N_x} and I_{N_t} are identity matrices, J is the matrix describing the discretization in time, and Q represents the spatial discretization. We use first-order backward differences (BDF1, or implicit Euler) for the time discretization and periodic central differences for the spatial discretization, leading to

$$J = \frac{1}{h_t}\begin{bmatrix} 1 & & & & \\ -1 & 1 & & & \\ & -1 & 1 & & \\ & & -1 & 1 & \\ & & & \ddots & \ddots \end{bmatrix} \quad \text{and} \quad Q = \frac{1}{h_x^2}\begin{bmatrix} 2 & -1 & & & -1 \\ -1 & 2 & -1 & & \\ & \ddots & \ddots & \ddots & \\ & & -1 & 2 & -1 \\ -1 & & & -1 & 2 \end{bmatrix}. \tag{3}$$

In the case of Dirichlet boundary conditions, all Dirichlet boundary points are eliminated in the discretized operator, $A = J \otimes I_{N_x-1} + I_{N_t} \otimes Q$, where Q is the non-periodic analogue of Q in (3), i.e., removing the last row and last column of Q.

DOI: 10.1002/nla

2.2. Waveform relaxation

The multigrid waveform relaxation method was developed by Lubich and Ostermann in [29]. The method combines red-black zebra line-relaxation along lines parallel to the time axis with a semicoarsening strategy using coarsening only in the spatial dimension. We consider the two-level algorithm based on waveform relaxation, with two space-time grids, Ω_h and Ω_H, where the subscript h represents the pair (h_x, h_t). The grid Ω_H is derived from Ω_h by doubling the mesh size in the space dimension only, i.e., H represents the pair $(2h_x, h_t)$. We assume that grid points are in the order of increasing time and from left to right in space, and we permute A into a red-black block ordering by first considering all unknowns at red grid points, (ih_x, jh_t) with i odd, and then all unknowns at black grid points, (ih_x, jh_t) with i even. Then, the iteration (error-propagation) operator for the red-black relaxation on the fine grid, Ω_h, can be written in the form

$$S^{\text{RB}} = S^{\text{BLACK}} S^{\text{RED}} = (I - M_B A)(I - M_R A), \tag{4a}$$

with

$$A = \begin{bmatrix} A_{RR} & A_{RB} \\ A_{BR} & A_{BB} \end{bmatrix}, \quad M_R = \begin{bmatrix} M_{RR}^{-1} & 0 \\ 0 & 0 \end{bmatrix}, \quad M_B = \begin{bmatrix} 0 & 0 \\ 0 & M_{BB}^{-1} \end{bmatrix}. \tag{4b}$$

Our interest is in two red-black schemes, red-black Jacobi, where

$$M_{RR} = D_{RR} \quad \text{and} \quad M_{BB} = D_{BB},$$

with D_{RR} and D_{BB} denoting the diagonals of A_{RR} and A_{BB}, respectively, and red-black Gauss-Seidel, where

$$M_{RR} = D_{RR} - L_{RR} \quad \text{and} \quad M_{BB} = D_{BB} - L_{BB},$$

with $-L_{RR}$ and $-L_{BB}$ denoting the strictly lower triangular parts of A_{RR} and A_{BB}, respectively.

We use standard geometric coarse-grid correction in the spatial direction, using periodic linear interpolation, P, full-weighting restriction, $R = \frac{1}{2}P^T$, and rediscretization to get the coarse-grid operator. Taking the global space-time structure into account, the interpolation and restriction operators are defined by $P_t = I_{N_t} \otimes P$ and $R_t = I_{N_t} \otimes R$, respectively, and the coarse-grid operator, A_c, is defined analogously to (2) with h_x replaced by $2h_x$ in (3). The resulting two-level algorithm may be represented by the two-grid iteration matrix,

$$T_{\text{WR}}^{(1,2)} = S^{\text{RB}}(I - P_t A_c^{-1} R_t A) S^{\text{RB}}. \tag{5}$$

2.3. MGRIT

The multigrid-reduction-in-time (MGRIT) algorithm [35] is based on applying multigrid reduction (MGR) techniques [46, 47] to time integration. The method uses block smoothers for relaxation and employs a semicoarsening strategy, that, in contrast to waveform relaxation, coarsens only in the temporal dimension. To describe the MGRIT algorithm, we assume that grid points are in the order of increasing time and from left to right in space. We then write (2) using the time-integration operator, $\Phi_{h_t} = (I_{N_x} + h_t Q)^{-1}$, leading to the block-scaled system

$$\hat{A}\mathbf{u} \equiv \begin{bmatrix} I_{N_x} & & & \\ -\Phi_{h_t} & I_{N_x} & & \\ & \ddots & \ddots & \\ & & -\Phi_{h_t} & I_{N_x} \end{bmatrix} \mathbf{u} = h_t (I_{N_x} \otimes \Phi_{h_t}) \mathbf{f} \equiv \hat{\mathbf{f}}. \tag{6}$$

For simplicity, we only describe the two-level MGRIT algorithm; the multilevel scheme results from applying the two-level method recursively. We consider two space-time grids, Ω_h and Ω_H, where the subscript h represents the pair (h_x, h_t). The grid Ω_H is derived from Ω_h by increasing the mesh size in the time dimension by a positive integer factor m, i.e., H represents the pair (h_x, mh_t). We partition the grid Ω_h into C-points, given by the set of coarse time-scale points, (ih_x, jh_t) with $j \bmod m = 0$, and F-points, (ih_x, jh_t) with $j \bmod m \neq 0$. Reordering the fine-grid operator, $\hat{A}$, by

DOI: 10.1002/nla

first considering all unknowns at F-points, then all unknowns at C-points, and using the subscripts f and c to indicate the two sets of points, the error propagator for the block smoother, FCF-relaxation, on the fine grid, Ω_h, can be written in the form $S^{FCF} = S^F S^C S^F$ with

$$S^F = I - S_f(S_f^T \hat{A} S_f)^{-1} S_f^T \hat{A}, \quad S^C = I - S_c(S_c^T \hat{A} S_c)^{-1} S_c^T \hat{A}, \tag{7}$$

and

$$\hat{A} = \begin{bmatrix} A_{ff} & A_{fc} \\ A_{cf} & A_{cc} \end{bmatrix}, \quad S_f = \begin{bmatrix} I_f \\ 0 \end{bmatrix}, \quad S_c = \begin{bmatrix} 0 \\ I_c \end{bmatrix}.$$

Thus, FCF-relaxation consists of three sweeps: F-relaxation [48], then C-relaxation, and again F-relaxation. F-relaxation updates the unknowns at F-points, assuming that unknowns at C-points are frozen; C-relaxation updates the unknowns at C-points analogously. Note that the choice of FCF-relaxation naturally follows from applying MGR techniques [47]. The intergrid transfer operators of MGRIT are injection and "ideal" interpolation,

$$R_I = \begin{bmatrix} 0 & I_c \end{bmatrix}, \quad P_\Phi = \begin{bmatrix} -A_{ff}^{-1} A_{fc} \\ I_c \end{bmatrix},$$

and the coarse-grid operator, $\hat{A}_c$, is defined by the coarse-scale time-integration operator, $\Phi_{mh_t} = (I_{N_x} + mh_t Q)^{-1}$ as in (6). Thus, the two-level MGRIT algorithm may be represented by the two-level iteration matrix,

$$T_{\text{MGRIT}}^{(1,2)} = (I - P_\Phi \hat{A}_c^{-1} R_I \hat{A}) S^{FCF}. \tag{8}$$

When using F-relaxation, S^F, instead of FCF-relaxation, S^{FCF}, the resulting two-level algorithm is the parareal time integration method [49],

$$T_{\text{parareal}}^{(1,2)} = (I - P_\Phi \hat{A}_c^{-1} R_I \hat{A}) S^F. \tag{9}$$

Notice that both the MGRIT algorithm and the parareal method use only pre-relaxation.

3. FAILURE OF LOCAL FOURIER ANALYSIS FOR PARABOLIC PROBLEMS

Several authors have noted the failure of local Fourier analysis to produce its usual quality of predictive results for convection-dominated or parabolic problems [18, 50]. Two studies of the precise mechanisms of this failure are [21, 22]. In [21], it is observed that the maximum of the Fourier symbol for the multigrid waveform relaxation method has a sharp peak for values of the temporal Fourier frequency close to 0 and $\pm\pi$. Thus, the standard practice of choosing a grid in Fourier space that matches the dimensions of the physical space grid of interest fails to resolve the true maximum of the continuous Fourier symbol. This is confirmed and expanded upon in [22], where predictions made based upon a sufficiently resolved grid in Fourier space are still shown to be lacking.

Several additional factors in the poor performance of LFA for the waveform relaxation multigrid methods are identified in [22]. Primary among these is the length of the time-integration interval, which can be controlled somewhat independently by varying both the time-step size (on a uniform-in-time mesh) and the number of time steps considered. For relatively small time steps, many time steps are needed before the analysis offers good prediction; in contrast, for large time steps, accurate prediction is seen for relatively few time steps. Figure 1 shows the computed and LFA-predicted spectra for two-grid (0,1)-cycles of the waveform-relaxation scheme applied to the periodic parabolic model problem in $(-\pi, \pi) \times [0, T]$ discretized with fixed time step of $1/16$, as the number of time steps is increased. Table I confirms the complementary analysis of predictivity for fixed number of time steps and varying time-step size. In both cases, we see that only as the product, $T = h_t N_t$, of the time-step size and number of time steps gets large does the analysis become sharp. Unfortunately, when considered in terms of non-dimensionalized time units, this

Prepared using nlaauth.cls
DOI: 10.1002/nla

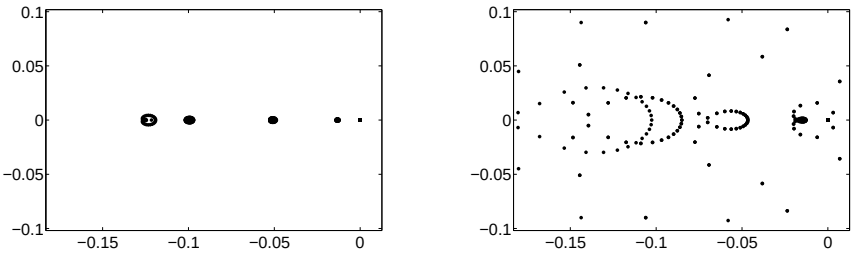

(a) Eigenvalues (in $\mathbb{C}$) on a 16×32 space-time grid.

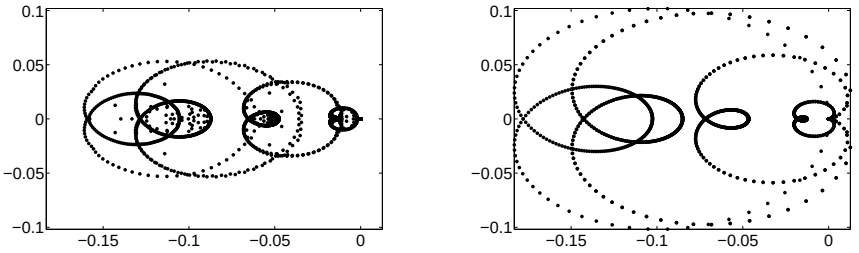

(b) Eigenvalues (in $\mathbb{C}$) on a 16×256 space-time grid.

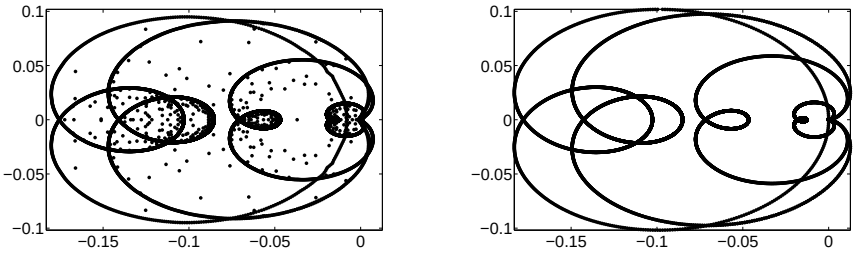

(c) Eigenvalues (in $\mathbb{C}$) on a 16×2048 space-time grid.

Figure 1. Eigenvalues (in $\mathbb{C}$) on space-time grids of size (a) 16×32 ($T = 2$), (b) 16×256 ($T = 16$), and (c) 16×2048 ($T = 128$) for two-grid (0,1)-cycles of the waveform-relaxation scheme applied to the periodic parabolic model problem on $(-\pi, \pi) \times [0, T]$ discretized with fixed time step $h_t = 1/16$. At left, eigenvalues of the analytically computed iteration matrix and at right, eigenvalues predicted by LFA. Results are similar to those of Figures 5.1 and 5.2 in [22].

h_t/h_x^2	$1/64$	$1/8$	1	8	65	512
$\overline{\rho}$	0.02	0.11	0.08	0.04	0.01	0.00
ρ_{LFA}	0.06	0.14	0.09	0.04	0.01	0.00
ρ_{analytic}	0.00	0.02	0.07	0.04	0.01	0.00

Table I. Average convergence factors per iteration (reported in [21]), $\overline{\rho}$, LFA predictions of spectral radii, ρ_{LFA}, and spectral radii of the analytically computed iteration matrix, ρ_{analytic}, for two-grid (1,1)-cycles of the waveform-relaxation scheme applied to the periodic parabolic model problem on $(-\pi, \pi) \times [0, T]$ discretized on a 128×128 space-time grid, reproduced from [22, Table 4]. Results are truncated after two digits past the decimal point because this corresponds to the precision of the results in [21].

renders the analysis useless except in the case of problems on time intervals that are too long to yield interesting physical behavior.

A secondary failure of LFA to be predictive that is exposed in [22] is the effect of multiplicities in the spectra of the multigrid iteration operator. While LFA appropriately resolves spectra of the iteration operators when ignoring the non-local character of the operators, i.e., under the long time-integration condition discussed above, it is based on the assumption that the iteration operator

DOI: 10.1002/nla

is diagonalizable and, as such, is unable to predict the effects of eigenvalues with degenerate geometric multiplicity. To some extent, this is a secondary problem for the waveform relaxation-based approach considered in [22]; however, for both the parareal [49] and multigrid-in-time [35] approaches, non-normality of the iteration operator is crucial: the iteration operator has only a single eigenvalue of zero, and all convergence information is contained in an analysis of the multiplicities of the invariant subspaces associated with this eigenvalue.

4. SEMI-ALGEBRAIC MODE ANALYSIS

In this section, we present the key ideas of SAMA. Analogously to classical mode analysis, exact predictions of the convergence behavior for a certain class of model problems and a certain class of multigrid algorithms can be made. The convergence behavior of problems and methods which do not belong to this class may be estimated by combining ideas of LFA with those of SAMA. In order to simplify the presentation of SAMA, we focus first on the multigrid methods for the parabolic model problem in space-time introduced in Section 2. Thus, in this section, we primarily consider the two-dimensional space-time scalar case. The generalization to higher dimensions is straighforward, and it is briefly outlined at the end of this section; applications of SAMA to other problems are discussed in Section 5.

4.1. The SAMA methodology

The semi-algebraic approach to mode analysis is motivated by the structure of the matrices of the iteration operator. As discussed above, the non-normal and non-local character of the operators seems to drive the failure of LFA. Considering the block lower-triangular structure, however, motivates a different approach. Using BDF1 for the time discretization given by time-differentiation matrix J in (3), the fine-grid operator, $A = J \otimes I_{N_x} + I_{N_t} \otimes Q$, is of the form

$$A = \begin{bmatrix} J_0 + Q & 0 & \cdots & 0 \\ J_{-1} & J_0 + Q & & 0 \\ & \ddots & \ddots & \vdots \\ & & J_{-1} & J_0 + Q \end{bmatrix}, \tag{10}$$

with $J_0 = (1/h_t) I_{N_x}$ and $J_{-1} = -(1/h_t) I_{N_x}$. In the case of periodic boundary conditions, it is natural to use periodic central differences for the spatial discretization. As a consequence, the fine-grid operator is block-Toeplitz with circulant blocks (BTCB). The idea of SAMA is to use a Fourier ansatz to resolve the circulant behavior and analytical or numerical computation for the non-circulant terms. More precisely, the circulant blocks can be diagonalized by the matrix, Ψ, of discretized spatial Fourier modes,

$$\boldsymbol{\psi}(x, \theta) = \mathrm{e}^{-\imath \theta x / h_x} \tag{11}$$

with N_x Fourier frequencies, θ, sampled on a uniform mesh in $(-\pi, \pi]$ with mesh size $2\pi/N_x$. Thus, the fine-grid operator, A, can be transformed to a block-Toeplitz matrix with diagonal blocks. Note that in the case of Dirichlet boundary conditions and central finite differences for the spatial discretization, the fine-grid operator is block-Toeplitz with Toeplitz blocks (BTTB) and no longer BTCB. However, instead of the exponential Fourier basis, the sine Fourier basis can be used to transform the resulting BTTB operator to a block-Toeplitz matrix with diagonal blocks as the sine Fourier modes are eigenfunctions of the spatial discretization matrix.

We now reorder this transformed block matrix, $\mathcal{F}^{-1} A \mathcal{F}$, where $\mathcal{F} = I_{N_t} \otimes \Psi$, from $N_t \times N_t$ blocks of size $N_x \times N_x$ to a block matrix with $N_x \times N_x$ blocks of size $N_t \times N_t$. This permutation

Prepared using nlaauth.cls
DOI: 10.1002/nla

results in a block-diagonal matrix with Toeplitz blocks,

$$\mathcal{P}^{-1}\mathcal{F}^{-1}A\mathcal{F}\mathcal{P} = \begin{bmatrix} B_1^{(A)} & 0 & \cdots & 0 \\ 0 & B_2^{(A)} & & 0 \\ & \ddots & \ddots & \vdots \\ 0 & \cdots & 0 & B_{N_x}^{(A)} \end{bmatrix} \tag{12}$$

and corresponds to gathering Fourier modes together, i.e., each diagonal block, $B_k^{(A)}$, $k = 1, \ldots, N_x$, corresponds to the evolution of one spatial Fourier mode over time,

$$B_k^{(A)} = \begin{bmatrix} j_0 + \lambda_k & 0 & \cdots & 0 \\ j_{-1} & j_0 + \lambda_k & & 0 \\ & \ddots & \ddots & \vdots \\ & & j_{-1} & j_0 + \lambda_k \end{bmatrix}, \tag{13}$$

where j_0 and j_{-1} are the diagonal entries of J_0 and J_{-1}, respectively, and λ_k denotes the (k,k)-entry of the diagonal matrix $\Psi^{-1}Q\Psi$ (the kth eigenvalue of Q; $\lambda_k = 4/h_x^2 \sin^2(k\pi/N_x)$ for our parabolic model problem with periodic boundary conditions). Thus, we can break up the computation of the spectrum of the fine-grid operator, A, to a mode-by-mode computation. Note that considering a backward time-discretization, the entries above the diagonal of each of the diagonal blocks are zero, while those below are constant on each subdiagonal. Therefore, SAMA provides exact eigenvalue expressions for each lower Jordan block, $B_k^{(A)}$.

Figure 2 demonstrates the predictive nature of SAMA for the fine-grid operator of the parabolic model problem with periodic boundary conditions. In contrast to the LFA prediction, the SAMA prediction of the spectrum is already sharp for a small number of time steps.

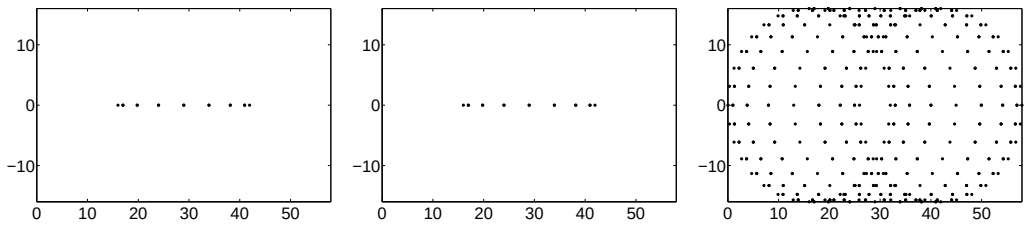

(a) Analytical spectrum (b) SAMA spectrum (c) LFA spectrum

Figure 2. Eigenvalues (in $\mathbb{C}$) for the periodic parabolic model problem on $(-\pi, \pi) \times [0, 2]$ discretized with fixed time step $h_t = 1/16$ on a space-time grid of size 16×32. The plots compare (a) the analytical spectrum to (b) the spectrum predicted by SAMA, and (c) the spectrum predicted by LFA.

Remark 1: If the fine-grid operator is block-circulant with circulant blocks (BCCB), each block of (12) can be diagonalized using Fourier modes. Thus, in the BCCB case, SAMA coincides with rigorous Fourier analysis.

Remark 2: While SAMA is most naturally considered for the BTCB case, we focus primarily on the special case of BTTB where the Toeplitz blocks are diagonalized exactly by the sine Fourier basis.

Remark 3: The use of circulant or block-circulant-structured preconditioners for non-circulant problems has a long history [51, 52], but is primarily focused on the case of problems that retain elliptic character; here, we consider the space-time discretization of parabolic problems, for which simple (multilevel) circulant preconditioners do not yield good performance.

4.2. Two-grid SAMA

A similar approach extends to all components of the multigrid process; however, the grid hierarchy affects the block size of the transformed operators. Considering temporal semicoarsening by a factor

DOI: 10.1002/nla

of m in the MGRIT approach and assuming exact spatial solves when applying the time-integration operator, the interpolation operator, P_Φ, for example, is a block matrix with $N_t \times N_t/m$ blocks of size $(N_x-1)\times(N_x-1)$, assuming Dirichlet boundary conditions. As a consequence, we reorder the Fourier-transformed block matrix, $\mathcal{F}^{-1}P_\Phi\mathcal{F}_c$, where $\mathcal{F}_c = I_{N_t/m}\otimes\Psi$, to a block matrix with $(N_x-1)\times(N_x-1)$ blocks of size $N_t \times N_t/m$,

$$\mathcal{P}^{-1}\mathcal{F}^{-1}P_\Phi\mathcal{F}_c\mathcal{P}_c = \begin{bmatrix} B_1^{(P_\Phi)} & 0 & \cdots & 0 \\ 0 & B_2^{(P_\Phi)} & & 0 \\ & \ddots & \ddots & \vdots \\ 0 & \cdots & 0 & B_{N_x-1}^{(P_\Phi)} \end{bmatrix}, \tag{14}$$

where

$$B_k^{(P_\Phi)} = I_{N_t/m}\otimes v \quad \text{with } v = \left[1, \lambda_k, \lambda_k^2, \ldots, \lambda_k^{m-1}\right]^T.$$

Similarly, the restriction operator, R_I, is transformed to a block matrix with $(N_x-1)\times(N_x-1)$ blocks of size $N_t/m \times N_t$, given by

$$B_k^{(R_I)} = \begin{bmatrix} 1 & & & & \\ \underbrace{\boxed{0 \quad \cdots \quad 0}}_{m} & 1 & & & \\ & \boxed{0 \quad \cdots \quad 0} & 1 & & \\ & & \ddots & & \\ & & & \boxed{0 \quad \cdots \quad 0} & \end{bmatrix}.$$

Denoting the diagonal blocks of the Fourier-transformed and permuted operators by $B_k^{(\cdot)}$ marking the respective operator in the superscript, the spectrum of the iteration matrix of the two-level MGRIT algorithm (8) can be computed mode-by-mode by considering the product

$$B_k^{(T_{\text{MGRIT}}^{(1,2)})} = \left(I - B_k^{(P_\Phi)}\left(B_k^{(\hat{A}_c)}\right)^{-1} B_k^{(R_I)} B_k^{(\hat{A})}\right) B_k^{(S^F)} B_k^{(S^C)} B_k^{(S^F)} \tag{15}$$

for each Fourier mode $k = 1, \ldots, N_x - 1$. For the two-level MGRIT algorithm with F-relaxation corresponding to the parareal time-integration method, we look at

$$B_k^{(T_{\text{parareal}}^{(1,2)})} = \left(I - B_k^{(P_\Phi)}\left(B_k^{(\hat{A}_c)}\right)^{-1} B_k^{(R_I)} B_k^{(\hat{A})}\right) B_k^{(S^F)}. \tag{16}$$

We note that the SAMA methodology could also be directly applied to methods that include post-relaxation as well as (or instead of) pre-relaxation.

The blocks $B_k^{(\hat{A})}$ of the transformed fine-grid operator $\hat{A}$ are given by

$$B_k^{(\hat{A})} = \begin{bmatrix} 1 & & & & \\ -\hat{\lambda}_k & 1 & & & \\ & -\hat{\lambda}_k & 1 & & \\ & & \ddots & \ddots & \\ & & & -\hat{\lambda}_k & 1 \end{bmatrix}, \tag{17}$$

with $\hat{\lambda}_k = \left(1 + 4h_t/h_x^2 \sin^2(k\pi/(2N_x))\right)^{-1}$. Since MGRIT uses temporal semicoarsening, the blocks $B_k^{(\hat{A}_c)}$ of the transformed coarse-grid operator, $\hat{A}_c$, are of the same form as those of the fine-grid operator given in (17) with the fine-grid eigenvalues, $\hat{\lambda}_k$, replaced by the coarse-grid eigenvalues, $\hat{\lambda}_{c;k} = \left(1 + 4mh_t/h_x^2 \sin^2(k\pi/2N_x)\right)^{-1}$. Considering F- and C-relaxation, the diagonal blocks $B_k^{(S^F)}$ and $B_k^{(S^C)}$ of the transformed iteration operators of these block smoothers

DOI: 10.1002/nla

are $N_t \times N_t$ block (bi-)diagonal matrices with blocks of size $m \times m$,

$$B_k^{(S^F)} = \begin{bmatrix} Z^{(S^F)} & & & \\ & Z^{(S^F)} & & \\ & & \ddots & \\ & & & Z^{(S^F)} \end{bmatrix} \quad \text{and} \quad B_k^{(S^C)} = \begin{bmatrix} Z^{(S^C)} & & \\ \hat{\lambda}_k \hat{e}_1 \hat{e}_m^T & Z^{(S^C)} & \\ & \ddots & \ddots \\ & & \hat{\lambda}_k \hat{e}_1 \hat{e}_m^T \quad Z^{(S^C)} \end{bmatrix}$$

respectively, with

$$Z^{(S^F)} = \begin{bmatrix} 1 & & & \\ \hat{\lambda}_k & 0 & & \\ \hat{\lambda}_k^2 & & 0 & \\ \vdots & & & \ddots \\ \hat{\lambda}_k^{m-1} & & & 0 \end{bmatrix}, \quad Z^{(S^C)} = \begin{bmatrix} 0 & & & \\ & 1 & & \\ & & 1 & \\ & & & \ddots \\ & & & & 1 \end{bmatrix},$$

$\hat{e}_1 = [1, 0, 0, \ldots, 0]^T$, and $\hat{e}_m = [0, \ldots, 0, 0, 1]^T$. This completes the definition of the diagonal blocks of the permuted operators of the parareal and MGRIT algorithms and, by using the expressions (15) and (16), this defines block matrices of the permuted iteration operators of the two-level methods as a whole. For all k, the diagonal block of these block matrices is zero and thus, for both the parareal and MGRIT approaches, the iteration operator has only a single eigenvalue of zero.

4.3. Operator norms and short-term convergence

The dynamics of short-term convergence behavior are more complex than predicted simply by the spectral radius. To gain insight into this aspect of performance, the operator norm of the iteration matrix, T, can be considered. More precisely, for any initial error, $\mathbf{e}_0$, $\|T^\nu \mathbf{e}_0\| \leq \|T^\nu\| \, \|\mathbf{e}_0\|$, where ν denotes the number of iterations of the multigrid scheme. Thus, $\|T^\nu\|$ is an upper bound for the error reduction after ν iterations which, itself, is naturally but not sharply bounded as $\|T^\nu\| \leq \|T\|^\nu$. The SAMA approach reduces the computation of these error reduction factors to the calculation of norms of products of block-diagonal matrices representing the evolution of a spatial Fourier mode over time for a discrete set of Fourier frequencies. Furthermore, since for any block-diagonal matrix B with blocks B_k, $k = 1, \ldots, n$, we have

$$\|T\|^2 = \sup_{\mathbf{x} \neq \mathbf{0}} \frac{\|B\mathbf{x}\|^2}{\|\mathbf{x}\|^2} = \sup_{\mathbf{x}_k \neq \mathbf{0}} \frac{\|B_1\mathbf{x}_1\|^2 + \cdots + \|B_n\mathbf{x}_n\|^2}{\|\mathbf{x}_1\|^2 + \cdots + \|\mathbf{x}_n\|^2} = \max_k \sup_{\mathbf{x}_k \neq \mathbf{0}} \frac{\|B_k\mathbf{x}_k\|^2}{\|\mathbf{x}_k\|^2} = \max_k \|B_k\|^2,$$

and likewise, $\|T^\nu\| = \max_k \|B_k^\nu\|$, we can estimate the operator norm of powers of the iteration operator by computing norms of the diagonal blocks and then maximize over the discrete set of Fourier frequencies.

Figure 3 compares SAMA-predicted error reduction factors (operator norms of powers of the iteration matrix) with experimentally measured error reduction factors for the two-level MGRIT algorithm with various coarsening factors, m. Results show that for all coarsening factors, SAMA tracks the short-term convergence behavior of MGRIT well. Additionally, SAMA predicts the following exactness property of MGRIT: Assuming exact arithmetic, one iteration of F-relaxation computes the exact solution at the first $m - 1$ time steps, i.e., at all F-points in the first coarse-scale time interval. Furthermore, one iteration of FCF-relaxation computes the exact solution at the first $2m - 1$ time steps. An interesting property of MGRIT resulting from this fact is that the algorithm solves for the exact solution in $N_t/(2m)$ iterations, corresponding to half the number of points on the coarse grid. Similarly, parareal solves for the exact solution in N_t/m iterations. Thus, the operator norm is exactly zero from iteration $N_t/(2m) + 1$ or $N_t/m + 1$, respectively, onwards. In the example considered in Figure 3, we look at $N_t = 256$. For $m = 32$ or $m = 16$, for example, the operator norm of the two-level MGRIT iteration matrix from the fifth or ninth power, respectively, onwards is zero according to the exactness property (and, thus, cannot be shown in log scale). This

DOI: 10.1002/nla

is predicted by SAMA as the SAMA-predicted error reduction factors of the fifth and ninth power of the iteration matrix for factor-32 and factor-16 coarsening, respectively, are zero.

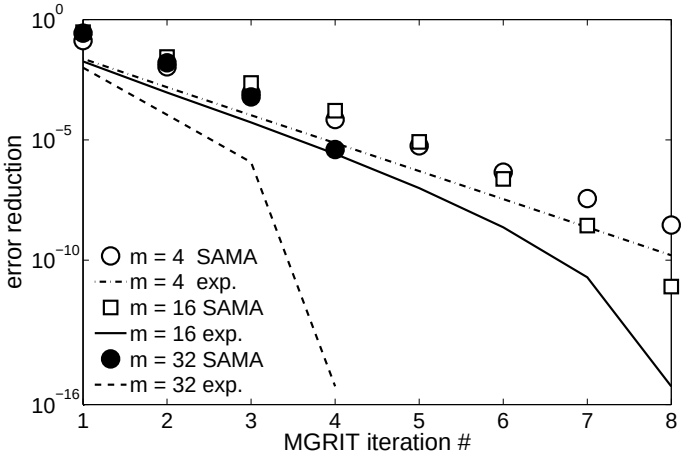

Figure 3. SAMA-predicted error reduction factors and measured error reduction factors for two-level cycles of the MGRIT algorithm with various coarsening factors applied to the parabolic model problem on $(0, \pi) \times [0, \pi^2/4]$ with Dirichlet BCs discretized on a 32×256 space-time grid.

4.4. Combining SAMA with LFA

The two-grid SAMA approach reduces the analysis of two-level methods to that of products of block-diagonal matrices, with each block representing the evolution of a spatial Fourier mode over time. The approach is rigorous precisely when the Fourier basis of spatial modes resolves the circulant or circulant-like structure of the operators of the iteration matrix. This is, however, an unreasonable assumption that fails to cover cases such as spatial semicoarsening or Gauss-Seidel-like relaxation schemes. In these cases, we use an LFA approach to account for *locally circulant* behavior. We extend the operators to an infinite grid in space and use the fact that any infinite-grid Toeplitz matrix can be diagonalized by a set of continuous Fourier modes, i.e., with Fourier frequencies, θ, that vary continuously in the interval $(-\pi, \pi]$. A discrete set of frequencies is then chosen, corresponding to a discrete mesh in θ, needed for the prediction of the performance of the two-grid method. In practice, we consider the discretized spatial Fourier modes (11).

Multigrid processing naturally mixes modes: Considering spatial semicoarsening with coarsening by a factor of two, as in the waveform relaxation-based method, for example, the restriction operator, R_t, maps two fine-grid functions, the Fourier harmonics, to one coarse-grid function. More precisely, these two functions are associated with the frequencies

$$\theta \in \left(-\frac{\pi}{2}, \frac{\pi}{2}\right] \quad \text{and} \quad \theta' = \theta - \text{sign}(\theta)\pi.$$

It can be verified that the fine-grid harmonic spaces are left invariant by the coarse-grid correction process [48]. As a result, using the matrix, Ψ_{harm}, of discretized Fourier modes, permuted in pairs of harmonics, $(\psi(x, \theta), \psi(x, \theta'))$, we can block-diagonalize each block of the restriction and interpolation operators, with block sizes 1×2 or 2×1, respectively, reflecting the coupling of the two Fourier harmonics. Thus, the restriction and interpolation operators, R_t and P_t, can be transformed to block-Toeplitz matrices with block-diagonal blocks. We then reorder these Fourier-transformed block matrices, $\mathcal{F}_c^{-1} R_t \mathcal{F}_{\text{harm}}$ and $\mathcal{F}_{\text{harm}}^{-1} P_t \mathcal{F}_c$, where $\mathcal{F}_{\text{harm}} = I_{N_t} \otimes \Psi_{\text{harm}}$ and $\mathcal{F}_c = I_{N_t} \otimes \Psi_c$, to block-diagonal matrices with $N_x/2 \times N_x/2$ blocks of size $N_t \times 2N_t$ or $2N_t \times N_t$,

DOI: 10.1002/nla

respectively. For the interpolation operator, P_t, for example, this permutation results in

$$\mathcal{P}^{-1}\mathcal{F}_{\text{harm}}^{-1}P_t\mathcal{F}_c\mathcal{P}_c = \begin{bmatrix} \mathcal{B}_1^{(P_t)} & 0 & \cdots & 0 \\ 0 & \mathcal{B}_2^{(P_t)} & & 0 \\ & \ddots & \ddots & \vdots \\ 0 & \cdots & 0 & \mathcal{B}_{N_x/2}^{(P_t)} \end{bmatrix},$$

with each block, $\mathcal{B}_k^{(P_t)}$, in contrast to (14), representing the interpolation of a spatial harmonic pair at all times. Similarly, since for red-black relaxation schemes, as in the waveform relaxation-based method, we are updating alternating points in space, we reorder the Fourier-transformed relaxation operators to block-diagonal matrices with $N_x/2 \times N_x/2$ blocks of size $2N_t \times 2N_t$ reflecting the coupling of pairs of "red" and "black" Fourier modes. Thus, the block calculation describing the evolution of a spatial Fourier mode over time needs to be extended.

4.4.1. Coarse-grid correction For the analysis of the coarse-grid correction process in the waveform relaxation approach, for example, instead of considering each Fourier mode separately, we look at each mode and its harmonic, denoted by the k-th and k'-th mode, respectively, with $k = 1, \ldots, N_x/2$, $k' = N_x - k$, and transform all operators to block-diagonal matrices with $N_x/2 \times N_x/2$ blocks. Taking into account that Fourier harmonics coincide on the coarse grid, we denote the diagonal blocks of the transformed fine-grid, coarse-grid, interpolation, and restriction operators by $\mathcal{B}_k^{(A)}$, $B_k^{(A_c)}$, $\mathcal{B}_k^{(P_t)}$, and $\mathcal{B}_k^{(R_t)}$, respectively. Then, the spectrum of the iteration matrix of the coarse-grid correction process in the waveform relaxation scheme, $I - P_t A_c^{-1} R_t A$, can be computed in terms of pairs of Fourier modes by considering the product

$$I - \mathcal{B}_k^{(P_t)} \left(B_k^{(A_c)}\right)^{-1} \mathcal{B}_k^{(R_t)} \mathcal{B}_k^{(A)}, \tag{18}$$

for each $k = 1, \ldots, N_x/2$ referring to the k-th pair of Fourier modes. The blocks $\mathcal{B}_k^{(A)}$ of the transformed fine-grid operator, A, can be written in terms of (13),

$$\mathcal{B}_k^{(A)} = \begin{bmatrix} B_k^{(A)} & 0 \\ 0 & B_{k'}^{(A)} \end{bmatrix},$$

with $j_0 = 1/h_t, j_{-1} = -1/h_t$, as well as

$$\lambda_k = \frac{4}{h_x^2}\sin^2\left(\frac{\pi k}{N_x} - \frac{\pi}{4}\right), \quad \text{and} \quad \lambda_{k'} = \frac{4}{h_x^2}\sin^2\left(\frac{\pi k}{N_x} + \frac{\pi}{4}\right)$$

for $B_k^{(A)}$ and $B_{k'}^{(A)}$, respectively, sampling the low Fourier frequency $\theta_k \in (-\pi/2, \pi/2]$ on a uniform mesh with mesh size $2\pi/N_x$ such that $\theta_k = -\pi/2 + 2\pi k/N_x$. The blocks $B_k^{(A_c)}$ of the transformed coarse-grid operator are also of the form (13) with $j_0 = 1/h_t, j_{-1} = -1/h_t$, and

$$\lambda_k = \frac{1}{h_x^2}\sin^2\left(\frac{2\pi k}{N_x} - \frac{\pi}{2}\right).$$

Finally, the blocks $\mathcal{B}_k^{(R_t)}$ and $\mathcal{B}_k^{(P_t)}$ of restriction and interpolation, respectively, are

$$\mathcal{B}_k^{(R_t)} = \begin{bmatrix} c_k I_{N_t} & s_k I_{N_t} \end{bmatrix} \quad \text{and} \quad \mathcal{B}_k^{(P_t)} = \left(\mathcal{B}_k^{(R_t)}\right)^T,$$

where c_k and s_k define the LFA symbols of linear interpolation and full-weighting restriction,

$$c_k = \frac{1}{2}\left(1 + \cos\left(\frac{\pi k}{N_x} - \frac{\pi}{4}\right)\right) \quad \text{and} \quad s_k = \frac{1}{2}\left(1 - \cos\left(\frac{\pi k}{N_x} - \frac{\pi}{4}\right)\right).$$

DOI: 10.1002/nla

4.4.2. Red-black relaxation For the analysis of the red-black relaxation schemes of the waveform relaxation-based two-grid method, it is natural to use a "red-black" Fourier basis instead of the "harmonic" Fourier basis considered for the analysis of coarse-grid correction. As a consequence, the $2N_t \times 2N_t$ diagonal blocks correspond to pairs of red and black Fourier modes instead of pairs of Fourier harmonics. However, a simple transformation allows combining the analyses of relaxation and coarse-grid correction to consider the waveform relaxation-based two-grid method as a whole. Details are given in Appendix A. Considering $N_x/2$ pairs of "red" and "black" Fourier modes, SAMA reduces the analysis of the iteration operator for the red-black relaxation defined in (4) to that of the product

$$\left(I - \mathcal{B}_k^{(M_B)}\mathcal{B}_k^{(A)}\right)\left(I - \mathcal{B}_k^{(M_R)}\mathcal{B}_k^{(A)}\right) \tag{19}$$

with $k = 1, \ldots, N_x/2$ referring to the k-th pair of Fourier modes. The blocks $\mathcal{B}_k^{(A)}$ of the transformed fine-grid operator are given by

$$\mathcal{B}_k^{(A)} = \left[\begin{array}{cccc:cccc} j_0+\lambda_k^{(RR)} & & & & \lambda_k^{(RB)} & & & \\ j_{-1} & j_0+\lambda_k^{(RR)} & & & & \lambda_k^{(RB)} & & \\ & \ddots & \ddots & & & & \ddots & \\ \hdashline \lambda_k^{(BR)} & & & & j_0+\lambda_k^{(BB)} & & & \\ & \lambda_k^{(BR)} & & & j_{-1} & j_0+\lambda_k^{(BB)} & & \\ & & \ddots & & & \ddots & \ddots & \end{array}\right],$$

with

$$j_0 = \frac{1}{h_t}, \qquad j_{-1} = -\frac{1}{h_t},$$
$$\lambda_k^{(RR)} = \lambda_k^{(BB)} = \frac{2}{h_x^2},\ \lambda_k^{(RB)} = -\frac{1}{h_x^2}\left(1 - \mathrm{e}^{-\imath(4\pi k/N_x)}\right),\ \lambda_k^{(BR)} = -\frac{1}{h_x^2}\left(1 - \mathrm{e}^{\imath(4\pi k/N_x)}\right).$$

Recall that

$$M_R = \begin{bmatrix} M_{RR}^{-1} & 0 \\ 0 & 0 \end{bmatrix}, \quad M_B = \begin{bmatrix} 0 & 0 \\ 0 & M_{BB}^{-1} \end{bmatrix}.$$

Since for red-black Jacobi in space-time $M_{RR} = D_{RR}$ and $M_{BB} = D_{BB}$ are already block-diagonal matrices with diagonal blocks, computing $\mathcal{B}_k^{(M_R)}$ and $\mathcal{B}_k^{(M_B)}$, $k = 1, \ldots, N_x/2$, is straightforward. For red-black Gauss-Seidel in space-time, we consider $M_{RR} = D_{RR} - L_{RR}$ and $M_{BB} = D_{BB} - L_{BB}$. In general, the blocks on the diagonal of these operators are Toeplitz and, thus, an LFA approach would need to be used. However, for the parabolic model problem there are no additional red-to-red or black-to-black connections besides the trivial ones. Hence, the blocks $\mathcal{B}_k^{(M_R)}$, $k = 1, \ldots, N_x/2$, of the reordered operator M_R are given by

$$\mathcal{B}_k^{(M_R)} = \left[\begin{array}{ccccc:c} \multicolumn{4}{c}{\left(\begin{array}{cccc} j_0+q_{11} & & & & \\ j_{-1} & j_0+q_{11} & & & \\ & \ddots & \ddots & & \\ & & j_{-1} & j_0+q_{11} & \end{array}\right)^{-1}} & & \\ & & & & & 0 \\ \hdashline \multicolumn{5}{c:}{0} & 0 \end{array}\right],$$

with $q_{11} = 2/h_x^2$ denoting the entry in the first row and first column of the matrix Q; a similar expression can be found for the blocks $\mathcal{B}_k^{(M_B)}$, $k = 1, \ldots, N_x/2$, of the reordered operator M_B.

Prepared using nlaauth.cls
DOI: 10.1002/nla

4.4.3. Application to multigrid waveform relaxation We consider the parabolic model problem (1) in the domain $(-\pi, \pi) \times [0, T]$, subject to an initial condition and periodic boundary conditions, and the two-level waveform relaxation-based multigrid method described in §2.2. For SAMA predictions, we use the expressions (18) and (19) for the analysis of coarse-grid correction and relaxation, respectively, along with the transformation given in Appendix A, which allows analyzing the two-grid method as a whole.

Figure 4 shows the computed and SAMA-predicted spectra on a 16×32 space-time grid for a two-grid (0,1)-cycle of the waveform-relaxation scheme applied to the parabolic model problem with fixed time step of $1/16$ considered in Figure 1(a). In contrast to the LFA prediction depicted in Figure 1(a), the SAMA prediction is already sharp when the time-integration interval is small. The differences in the computed and SAMA-predicted spectra result from the presence of lower Jordan blocks in the iteration matrix causing difficulties in the numerical computation of the spectrum. Jordan blocks can lead to clustering of computed eigenvalues in form of a circle in the complex plane. As opposed to the computed spectrum, SAMA generally predicts only one eigenvalue with multiplicity corresponding to the size of the lower Jordan block.

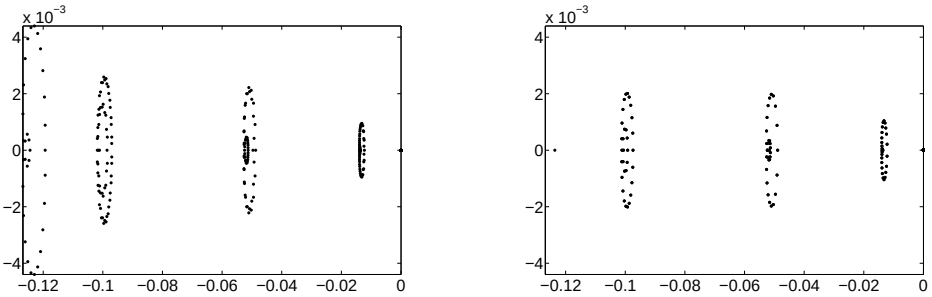

Figure 4. Eigenvalues (in $\mathbb{C}$) on a 16×32 space-time grid for a two-grid (0,1)-cycle of the waveform-relaxation scheme applied to the periodic parabolic model problem on $(-\pi, \pi) \times [0, 2]$ discretized with fixed time step $h_t = 1/16$. At left, eigenvalues of the analytically computed iteration matrix as in Figure 1(a) left side and at right, eigenvalues predicted by SAMA.

To separate effects of Jordan blocks in the numerical computation of the spectrum of the iteration matrix, we consider pseudospectra, where the ε-pseudospectrum of a matrix is defined as the set in the complex plane covered by the union of the spectrum of all perturbations of that matrix by terms less than ε in norm [53]. Figure 5 compares pseudospectra of the iteration matrix, computed with the MATLAB package EigTool [54] with SAMA and LFA predictions. In contrast to the eigenvalues of the LFA prediction, the eigenvalues of the SAMA prediction lie within the 10^{-15} pseudospectrum. Thus, we get a much better prediction of the spectrum of the iteration matrix using SAMA than when applying LFA. Furthermore, the structure of eigenvalue clustering is captured by the SAMA prediction.

Results of the performance of the two-grid waveform relaxation-based method are measured by the factor, ρ_ν, by which the l_2-norm of the residual is reduced in the νth iteration. The average convergence factor over k successive cycles, $\overline{\rho}_{\nu,\nu+k-1}$, is simply the geometric mean of $\rho_\nu, \rho_{\nu+1}, \ldots, \rho_{\nu+k-1}$. Since SAMA readily predicts the asymptotic convergence behavior, we average over the last 20 of 100 iterations. In order to be guaranteed to be able to run 100 iterations before convergence, we consider a zero right-hand side and a random initial guess. Table II shows these average convergence factors, $\overline{\rho}_{81,100}$, as well as SAMA-predicted spectral radii. We see that SAMA predictions correspond well to experimentally measured convergence factors for all grid sizes; however, SAMA generally predicts faster convergence than seen in practice, since large multiplicities lead to slower initial convergence.

To capture the effects of early convergence rates measured in experiments, we consider the operator norm of the iteration matrix. Figure 6 compares SAMA-predicted error reduction factors (operator norms of powers of the iteration matrix) with experimentally measured error reduction

DOI: 10.1002/nla

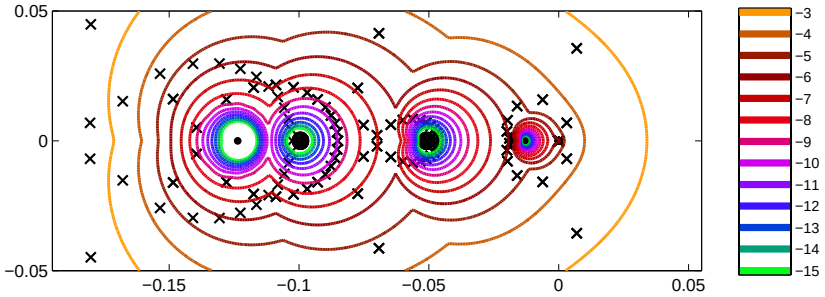

Figure 5. Eigenvalues predicted by SAMA (black dots) and LFA (black x's) as well as pseudospectra (in $\mathbb{C}$), represented by solid lines, of the iteration matrix of a two-grid (0,1)-cycle of the waveform-relaxation scheme applied to the periodic parabolic model problem on $(-\pi, \pi) \times [0, 2]$ discretized with fixed time step $h_t = 1/16$ on a 16×32 space-time grid. Note that the SAMA predictions are so close to one-another within the 10^{-15} pseudospectrum that one cannot distinguish the individual dots.

		$N_x = 32$	$N_x = 64$	$N_x = 128$	$N_x = 256$	$N_x = 512$
$\overline{\rho}_{81,100}$	$N_t = 32$	1.39e-01	8.54e-02	2.81e-02	7.16e-03	1.78e-03
	$N_t = 64$	1.56e-01	1.21e-01	5.72e-02	1.76e-02	4.67e-03
	$N_t = 128$	1.77e-01	1.54e-01	8.97e-02	3.29e-02	9.15e-03
ρ_{SAMA}	$N_t = 32$	1.18e-01	6.63e-02	2.20e-02	5.79e-03	1.50e-03
	$N_t = 64$	1.26e-01	1.07e-01	4.68e-02	1.39e-02	3.58e-03
	$N_t = 128$	1.31e-01	1.34e-01	8.31e-02	2.95e-02	8.08e-03

Table II. Average convergence factors per iteration in last 20 of 100 iterations, $\overline{\rho}_{81,100}$, and SAMA-predicted spectral radii, ρ_{SAMA}, of the iteration matrix of two-grid (0,1)-cycles of the waveform-relaxation scheme applied to the periodic parabolic model problem on $(-\pi, \pi) \times [0, 1]$ discretized with mesh sizes $h_x = 2\pi/N_x$ and $h_t = 1/N_t$ on $N_x \times N_t$ space-time grids.

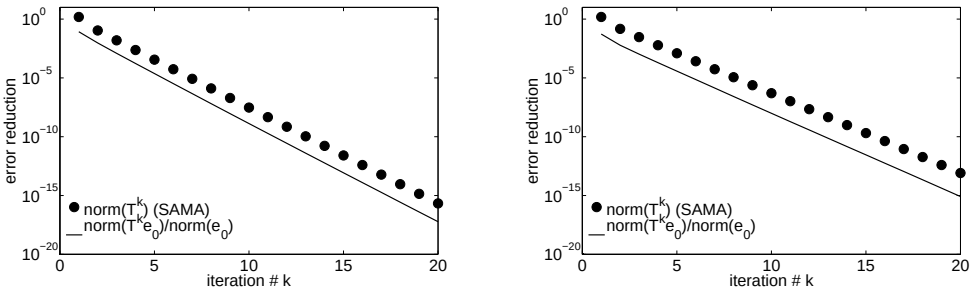

Figure 6. SAMA-predicted error reduction factors (dots) and experimentally measured error reduction factors (solid lines) of two-grid (0,1)-cycles of the waveform-relaxation scheme applied to the periodic parabolic model problem on $(-\pi, \pi) \times [0, 1]$. At left, error reduction factors on a 32×32 space-time grid and at right, error reduction factors on a 32×128 space-time grid.

factors. In addition to adequately predicting the asymptotic behavior as in Table II, Figure 6 shows that considering operator norms, SAMA accurately tracks the early performance of the method.

DOI: 10.1002/nla

The analysis for other discretization schemes is similar. Using periodic central differences for the spatial discretization and a general backward-difference-in-time scheme given by time-differentiation matrix

$$J = \begin{bmatrix} j_0 & 0 & \cdots & & 0 \\ j_{-1} & j_0 & & & 0 \\ j_{-2} & j_{-1} & j_0 & & \\ \vdots & \ddots & \ddots & \ddots & \vdots \\ j_{1-N_t} & \cdots & j_{-2} & j_{-1} & j_0 \end{bmatrix},$$

typically of fixed (small) bandwidth, and defining $J_i = j_i I_{N_x}$, $i = 0, -1, \ldots, 1 - N_t$, the fine-grid operator, $A = J \otimes I_{N_x} + I_{N_t} \otimes Q$, of the parabolic model problem is of the form

$$A = \begin{bmatrix} J_0 + Q & 0 & \cdots & 0 \\ J_{-1} & J_0 + Q & & 0 \\ & \ddots & \ddots & \vdots \\ J_{1-N_t} & \cdots & J_{-1} & J_0 + Q \end{bmatrix}.$$

The spectrum of this fine-grid operator can then be computed mode-by-mode by considering blocks of the form

$$B_k^{(A)} = \begin{bmatrix} j_0 + \lambda_k & 0 & \cdots & 0 \\ j_{-1} & j_0 + \lambda_k & & 0 \\ & \ddots & \ddots & \vdots \\ j_{1-N_t} & \cdots & j_{-1} & j_0 + \lambda_k \end{bmatrix}, \tag{20}$$

where λ_k again denotes the kth eigenvalue of Q and j_i, $i = 0, -1, \ldots, 1 - N_t$, are the diagonal entries of the matrices J_i. For second-order backward differences (BDF2) in time, for example, $j_0 = 3/2h_t, j_{-1} = -2/h_t, j_{-2} = 1/2h_t$, and $j_i = 0$ for all $i = -3, -4, \ldots, 1 - N_t$.

Figures 7 and 8 present similar results to Figures 4 and 5, respectively, for the two-grid waveform-relaxation method applied to the parabolic model problem using BDF2 instead of BDF1 for the discretization in time. Again, the SAMA prediction of the spectrum is already sharp for the short time-integration interval and it is much better than the LFA prediction.

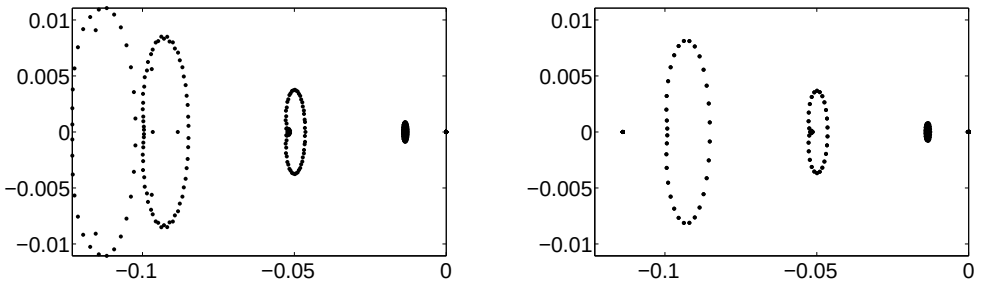

Figure 7. Eigenvalues (in $\mathbb{C}$) on a 16×32 space-time grid for a two-grid (0,1)-cycle of the waveform-relaxation scheme applied to the periodic parabolic model problem on $(-\pi, \pi) \times [0, 2]$ discretized with fixed time step $h_t = 1/16$ using BDF2 instead of BDF1 as in Figure 4. At left, eigenvalues of the analytically computed iteration matrix and at right, eigenvalues predicted by SAMA.

4.5. Three-grid SAMA

Similarly to the extension of LFA to a three-grid Fourier analysis presented in [55], incorporating a three-grid analysis into the SAMA approach is straightforward. In addition to broadening the applicability to a larger class of multigrid methods, three-grid SAMA can improve the accuracy

DOI: 10.1002/nla

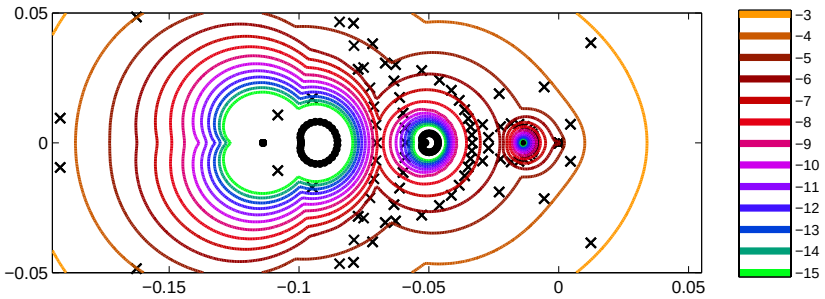

Figure 8. Eigenvalues predicted by SAMA (black dots) and LFA (black x's) as well as pseudospectra (in $\mathbb{C}$), represented by solid lines, of the iteration matrix of a two-grid (0,1)-cycle of the waveform-relaxation scheme applied to the periodic parabolic model problem on $(-\pi, \pi) \times [0, 2]$ discretized with fixed time step $h_t = 1/16$ using BDF2 instead of BDF1 as in Figure 5.

of asymptotic multigrid convergence predictions. Furthermore, a three-grid analysis can provide insight into extending optimal two-level algorithms to true multilevel algorithms while preserving optimality. As an example of this property, we consider the three-level MGRIT algorithm and the obvious generalization of the parareal method to three grids. In order to simplify the presentation, we focus on three-level V-cycles with only pre-relaxation as this variant is used in the MGRIT setting. More precisely, in this section three-level cycles refer to three-level $V(1,0)$-cycles. However, as in the LFA context using the recursion formula of L-grid cycles [4, 48, 56], iteration matrices of many variants, including three-grid W-cycles and different choices of the number of pre- and post-smoothing iterations, for example, can be defined and analyzed.

Instead of inverting $\hat{A}_c$ as in the two-level MGRIT algorithm (8), in the three-level method the system on the first coarse grid is approximated by a two-grid cycle with zero initial guess, i.e., the term $\hat{A}_c^{-1}$ in the two-grid iteration matrix (8) is replaced by $(I_c - T_{\text{MGRIT}}^{(2,3)})\hat{A}_c^{-1}$ to obtain the three-grid iteration matrix

$$T_{\text{MGRIT}}^{(1,3)} = \left(I - P_\Phi (I_c - T_{\text{MGRIT}}^{(2,3)})\hat{A}_c^{-1} R_I \hat{A}\right) S^{\text{FCF}}, \tag{21}$$

where $T_{\text{MGRIT}}^{(2,3)}$ is defined as in (8) with operators on the first and second coarse grids; the iteration matrix of the three-level parareal method can be derived analogously. The analysis of the three-level MGRIT algorithm by means of SAMA is then reduced to that of the product given in (15) with $\left(B_k^{(\hat{A}_c)}\right)^{-1}$ replaced by the product

$$\left(I - \left(I - B_k^{(P_{t;c})} \left(B_k^{(\hat{A}_{cc})}\right)^{-1} B_k^{(R_{t;c})} B_k^{(\hat{A}_c)}\right) B_k^{(S_c^{\text{FCF}})}\right) \left(B_k^{(\hat{A}_c)}\right)^{-1}, \tag{22}$$

where the subscripts c and cc indicate that operators are defined on the first and second coarse grid, respectively. Note that expressions for the diagonal blocks in (22), i.e., diagonal blocks of the transformed two-level iteration operator used for approximating the system on the first coarse grid, can be derived similarly to the expressions for the diagonal blocks of the transformed operators defining the two-level algorithms.

Figure 9 shows SAMA-predicted error reduction factors and experimentally measured error reduction factors for two- and three-level cycles of the parareal and MGRIT algorithms. The slopes of the best-fit lines of the plots can be used to compute the predicted and measured average error reduction per iteration, $\overline{\rho}_{\text{SAMA}}$ and $\overline{\rho}$, respectively. With superscripts denoting the number of grid

Prepared using nlaauth.cls
DOI: 10.1002/nla

levels, for parareal we obtain

$$\overline{\rho}^{(2)}_{\text{SAMA}} = 0.13, \quad \overline{\rho}^{(2)} = 0.11,$$
$$\overline{\rho}^{(3)}_{\text{SAMA}} = 0.22, \quad \overline{\rho}^{(3)} = 0.21,$$

and for MGRIT we get

$$\overline{\rho}^{(2)}_{\text{SAMA}} = 0.05, \quad \overline{\rho}^{(2)} = 0.05,$$
$$\overline{\rho}^{(3)}_{\text{SAMA}} = 0.07, \quad \overline{\rho}^{(3)} = 0.07.$$

The degradation from two-level to three-level in the analysis of parareal suggests that the obvious generalization of parareal to multiple levels may not produce an optimal method. If the algorithm were to be scalable, we would see not much (if any) degradation going from two-level to three-level, as is the case for MGRIT.

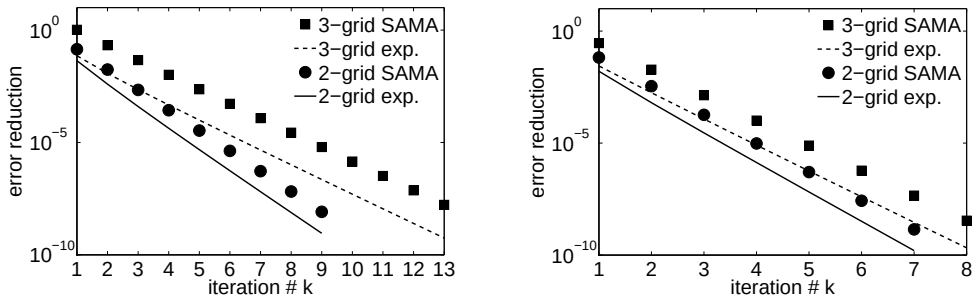

Figure 9. SAMA-predicted error reduction factors and measured error reduction factors for two- and three-level cycles of the parareal and MGRIT algorithms with factor-2 coarsening applied to the parabolic model problem on $(0, \pi) \times [0, \pi^2/4]$ with Dirichlet BCs discretized on a 32×256 space-time grid. At left, error reduction factors for parareal and at right, error reduction factors for MGRIT.

4.6. Nested two-grid SAMA

Another expansion of two-grid SAMA is to incorporate the analysis of nested two-grid methods. This is of natural interest for the analysis of the MGRIT algorithm due to the fact that the time-integration operator, Φ_{h_t}, in the block-scaled system (6) corresponds to a spatial solve. Considering exact spatial solves and Dirichlet boundary conditions, the sine Fourier basis can be used to diagonalize Φ_{h_t} as presented in Section 4.2. However in practice, the spatial problems are solved with an iterative method such as multigrid. We can incorporate using μ iterations of two-grid cycles for the spatial solves by replacing Φ_{h_t} by the expression

$$\tilde{\Phi}_{h_t} = \left(I - \left(I - T^{(1,2)}_{\text{space}} \Phi^{-1}_{h_t}\right)^{\mu}\right) \Phi_{h_t}, \tag{23}$$

where $T^{(1,2)}_{\text{space}}$ is the two-grid iteration matrix of the spatial multigrid method used. As a simple example, we consider using a spatial multigrid method consisting of two-grid (1,1)-cycles with coarsening by a factor of two, red-black Gauss-Seidel relaxation, full weighting restriction and linear interpolation. Since the sine Fourier modes no longer diagonalize the blocks of the operators of the iteration matrix, we consider an LFA approach. Using standard multigrid operators, a representation of their Fourier symbols can be found in the literature [4, 48, 56]. Since both red-black relaxation and coarse-grid correction mix modes, we consider pairs of Fourier harmonics. As a consequence, when analyzing two-level MGRIT with approximate spatial solves by means of SAMA, operators of the two-level MGRIT iteration matrix are transformed to block-diagonal matrices with $(N_x/2 - 1) \times (N_x/2 - 1)$ blocks as in Section 4.4, where each block corresponds to

DOI: 10.1002/nla

the evolution of a spatial harmonic pair over time; the remaining Fourier mode is in the nullspace of the red-black Gauss-Seidel relaxation scheme and, thus, does not contribute to the analysis. Hence, we consider the product

$$\left(I - \mathcal{B}_k^{(P_\Phi)} \left(B_k^{(\hat{A}_c)}\right)^{-1} \mathcal{B}_k^{(R_I)} \mathcal{B}_k^{(\hat{A})}\right) \mathcal{B}_k^{(S^F)} \mathcal{B}_k^{(S^C)} \mathcal{B}_k^{(S^F)} \tag{24}$$

with $k = 1, \ldots, N_x/2 - 1$ representing the k-th pair of Fourier harmonics. Expressions for the $\mathcal{B}_k^{(\cdot)}$'s in (24) correspond to those for the $B_k^{(\cdot)}$'s in the case of exact spatial solves, given in Section 4.2, with each scalar entry replaced by a matrix of size 2×2. More precisely, we replace each entry equal to 0 or 1 by the zero matrix or identity matrix of size 2×2, respectively. Furthermore, since we approximate the spatial solves on the fine grid by applying μ two-grid (1,1)-cycles of the spatial multigrid method described above, we replace each entry equal to $\hat{\lambda}_k$ by

$$\Lambda_k = \left(I - \left(\Lambda_k^{(S_x^{\text{RB}})} \Lambda_k^{(K_x^{(1,2)})} \Lambda_k^{(S_x^{\text{RB}})}\right)^\mu\right) \begin{bmatrix} \tilde{\lambda}_k & 0 \\ 0 & \tilde{\lambda}_{k'} \end{bmatrix}$$

with

$$\Lambda_k^{(S_x^{\text{RB}})} = \frac{1}{4} \begin{bmatrix} \beta+1 & \beta+1 \\ -\beta+1 & -\beta+1 \end{bmatrix} \begin{bmatrix} \beta+1 & -\beta-1 \\ \beta-1 & -\beta+1 \end{bmatrix}, \quad \beta = \frac{2\alpha}{1+2\alpha} \cos\left(\frac{k\pi}{N_x}\right), \; \alpha = \frac{h_t}{h_x^2},$$

and

$$\Lambda_k^{(K_x^{(1,2)})} = I - \tilde{\lambda}_{c;k} \begin{bmatrix} \cos^4(\theta_k) & -\cos^2(\theta_k)\sin^2(\theta_k) \\ -\cos^2(\theta_k)\sin^2(\theta_k) & \sin^4(\theta_k) \end{bmatrix} \begin{bmatrix} \tilde{\lambda}_k^{-1} & 0 \\ 0 & \tilde{\lambda}_{k'}^{-1} \end{bmatrix},$$

where $\theta_k = k\pi/(2N_x)$ and

$$\tilde{\lambda}_k = \left(1 + 4\frac{h_t}{h_x^2}\sin^2(\theta_k)\right)^{-1}, \quad \tilde{\lambda}_{c;k} = \left(1 + \frac{h_t}{h_x^2}\sin^2(\theta_k/2)\right)^{-1}$$

for $k = 1, \ldots, N_x/2 - 1$ and $k' = N_x - k$. Finally, considering exact spatial solves on the coarse grid, each entry of $B_k^{(\hat{A}_c)}$ equal to $\hat{\lambda}_{c;k}$ becomes

$$\begin{bmatrix} \hat{\lambda}_{c;k} & 0 \\ 0 & \hat{\lambda}_{c;k'} \end{bmatrix} \quad \text{for } k = 1, \ldots, N_x/2 - 1 \text{ and } k' = N_x - k.$$

Figure 10 shows SAMA-predicted error reduction factors and measured error reduction factors for two-level cycles of the MGRIT algorithm with FCF-relaxation using μ cycles of spatial multigrid to approximate the spatial solves on the fine grid for different values of μ. Increasing the number of two-level cycles for the spatial solves from one to two increases computation cost, but gives better average error reduction per iteration. However, increasing the number further to three cycles only increases computational cost without improving average error reduction per iteration. More precisely, the SAMA-predicted and actual average error reduction per iteration for $\mu = 1$ is about 0.11, whereas it is about 0.05 for $\mu = 2$ and $\mu = 3$.

4.7. SAMA for higher dimensional problems

The generalization of the SAMA approach to multi-dimensional problems is straightforward. Just as in the two-dimensional space-time scalar case, the SAMA approach reduces the analysis of multigrid methods applied to multi-dimensional problems to that of products of block-diagonal matrices, with each block representing the evolution of a spatial Fourier mode over time. However, for problems defined on a higher-dimensional spatial domain, the analysis requires considering more spatial Fourier modes, since Fourier frequencies, θ, are sampled on a higher-dimensional grid. As a consequence, in the rigorous case, i.e., in the case that the Fourier basis of spatial Fourier modes

DOI: 10.1002/nla

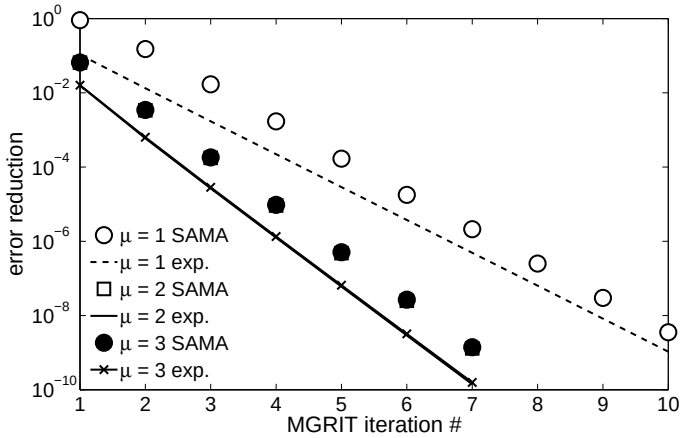

Figure 10. SAMA-predicted and measured error reduction factors for two-level cycles of the MGRIT algorithm with FCF-relaxation and factor-2 coarsening applied to the parabolic model problem on $(0,\pi)\times[0,\pi^2/4]$ with Dirichlet BCs discretized on a 32×256 space-time grid. The spatial solves on the fine grid are approximated by μ two-level (1,1)-cycles with coarsening by a factor of two, red-black Gauss-Seidel relaxation, full weighting restriction and linear interpolation. Note that error reduction factors almost coincide for $\mu=2$ and $\mu=3$.

resolves the circulant or circulant-like structure of the operators of the iteration matrix, only the number of blocks, $B_k^{(\cdot)}$, considered in the analysis increases.

Figure 11 shows SAMA-predicted error reduction factors and experimentally measured error reduction factors for two-level cycles of the MGRIT algorithm with FCF-relaxation and factor-2 coarsening applied to the parabolic model problem in two space dimensions with Dirichlet boundary conditions, and assuming exact spatial solves when applying the time-integration operator. The results show that SAMA provides good predictions of the short-term convergence behavior in the multi-dimensional case. Furthermore, the average error reduction per iteration is about 0.05 with rates only slowly degrading, as is typical for small grid sizes in a domain refinement study.

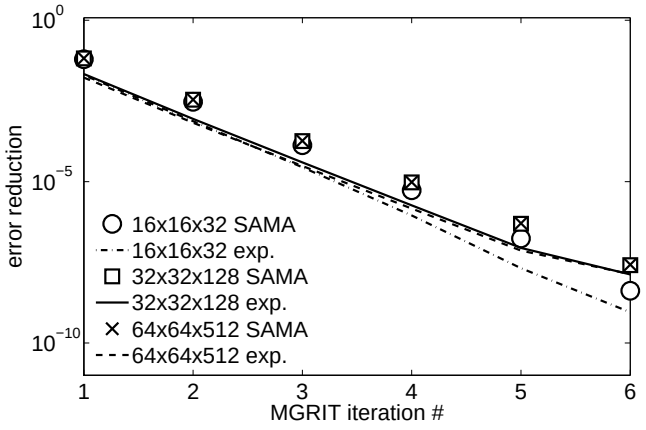

Figure 11. SAMA-predicted error reduction factors and measured error reduction factors for two-level cycles of the MGRIT algorithm applied to the parabolic model problem on $(0,\pi)^2\times[0,\pi^2/8]$ with Dirichlet BCs discretized on a space-time grid of various grid sizes.

In the non-rigorous case, we could again use an LFA approach to account for locally circulant behavior. For a parabolic problem in d space dimensions, every low frequency mode is coupled with 2^d-1 high frequency modes. Accordingly, in the analysis by means of SAMA, we would need to

DOI: 10.1002/nla

consider the evolution of 2^d spatial Fourier modes over time. We note that standard expressions for LFA of coarse-grid correction and relaxation processes are available for problems in two or three space dimensions [48] and, so, it is relatively straight-forward to account for these in the SAMA setting.

5. ADDITIONAL EXAMPLES

The SAMA approach reduces the analysis of multigrid methods to that of products of block-diagonal matrices representing, in the parabolic case, the evolution of one spatial Fourier mode or a pair of (or, in d spatial dimensions, 2^d) spatial Fourier modes over time for a discrete set of Fourier frequencies. In this section, we consider two non-parabolic model problems; one describing elliptic diffusion in layered media and one convection-diffusion problem. Both of these problems demonstrate that SAMA can be readily applied to non-parabolic problems, but is limited to cases where there is tensor-product structure with homogeneous elliptic-like behavior in one direction.

5.1. Elliptic diffusion in layered media

The first example deals with the two-scale layered diffusion problem

$$-\nabla \cdot (\kappa(y)\nabla u) = 0 \quad \text{on } [0,1]^2 \tag{25}$$

with permeability $\kappa(y)$ and subject to homogeneous Neumann boundary conditions at the boundaries $x = 0$, $x = 1$, and $y = 0$,

$$\frac{\partial u}{\partial x}(0,y) = \frac{\partial u}{\partial x}(1,y) = \frac{\partial u}{\partial y}(x,0) = 0,$$

and a homogeneous Dirichlet boundary condition at the boundary $y = 1$,

$$u(x,1) = 0.$$

The permeability is constructed based on dividing the unit square into layers of alternating permeability as depicted in Figure 12 for four layers and permeabilities 1 and 10^{-7}.

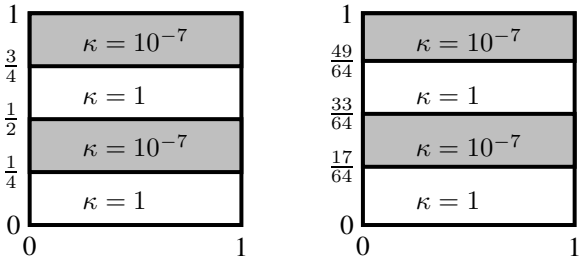

(a) Example 1 (b) Example 2

Figure 12. Permeability fields considered for two-scale layered diffusion problems.

We discretize our elliptic model problem (25) on a uniform Cartesian mesh using $N_x \times N_y$ bilinear quadrilateral finite elements of size $h_x \times h_y$, where $h_x = 1/N_x$ and $h_y = 1/N_y$. Then, the discretized operator can be written in the form $A = \mathcal{M}_y \otimes \mathcal{S}_x + \mathcal{S}_y \otimes \mathcal{M}_x$, where $\mathcal{M}_x$ and $\mathcal{M}_y$ are the mass matrices and $\mathcal{S}_x$ and $\mathcal{S}_y$ are the stiffness matrices in the x- and y-dimension, respectively, given in stencil form as

$$\mathcal{M}_x = h_x \begin{bmatrix} \frac{1}{6} & \frac{2}{3} & \frac{1}{6} \end{bmatrix}, \qquad \mathcal{S}_x = \frac{1}{h_x} \begin{bmatrix} -1 & 2 & -1 \end{bmatrix},$$

$$\mathcal{M}_y = h_y \begin{bmatrix} \frac{1}{6}\kappa_{i-1} & \frac{1}{3}(\kappa_{i-1}+\kappa_i) & \frac{1}{6}\kappa_i \end{bmatrix}, \qquad \mathcal{S}_y = \frac{1}{h_y} \begin{bmatrix} -\kappa_{i-1} & \kappa_{i-1}+\kappa_i & -\kappa_i \end{bmatrix},$$

DOI: 10.1002/nla

where $\kappa_i, i = 1, 2, \ldots, N_y$, is the permeability in the i-th element in the y-direction, and stencils at boundary points are adjusted appropriately. We eliminate Dirichlet boundary points in the discretized operator, whereas Neumann boundary points are retained. Thus, $\mathcal{M}_x$ and $\mathcal{S}_x$ are matrices of size $(N_x + 1) \times (N_x + 1)$ and $\mathcal{M}_y$ and $\mathcal{S}_y$ are matrices of size $N_y \times N_y$.

Efficient solvers for this, and more strongly heterogeneous, diffusion problems have attracted much attention in recent years, as they arise in many applications, including neutron transport and geophysical flow. One of the first effective algorithms was the Black Box Multigrid (BoxMG) method [42, 43], which introduced operator-induced interpolation into the multigrid literature, in combination with suitable relaxation and Galerkin coarsening techniques. Algebraic multigrid (AMG) [48, 57] is also an effective approach, although it is more suited to problems with much more heterogeneous structure. Specialized preconditioners have also been introduced, including geometric multigrid preconditioning [58], domain decomposition approaches [59,60], and deflation-based preconditioners [44, 45]. Here, we focus on the analysis of BoxMG and the deflation-based preconditioners; however, the analysis could also be applied to other approaches, as long as the circulant/Toeplitz structure is essentially maintained.

5.1.1. BoxMG The Black Box Multigrid method (BoxMG) was introduced in [42, 43]. Its key feature is Galerkin coarse-grid correction based on operator-induced interpolation that effectively accounts for the discontinuities in the permeability $\kappa(y)$ in (25). We consider the two-level algorithm based on BoxMG with weighted Jacobi relaxation and full coarsening by a factor of two in each dimension. BoxMG defines a two-dimensional interpolation scheme but, since the permeability only varies in y, BoxMG interpolation coincides with a tensor-product interpolation of standard linear interpolation in the x-direction and ideal one-dimensional interpolation in the y-direction. Thus, denoting interpolation in x by P_{lin} and interpolation in y by P_{1D}, given in stencil form as

$$P_{\text{lin}} = \frac{1}{2}\begin{bmatrix}1 & 2 & 1\end{bmatrix} \quad \text{and} \quad P_{\text{1D}} = \frac{1}{\kappa_{i-1}+\kappa_i}\begin{bmatrix}\kappa_{i-1} & \kappa_{i-1}+\kappa_i & \kappa_i\end{bmatrix},$$

where $\kappa_i, i = 1, 2, \ldots, N_y$, is the permeability in the i-th element in the y-direction, the interpolation operator can be written in the form $P = P_{\text{1D}} \otimes P_{\text{lin}}$. Choosing the transpose of interpolation for restriction and the Galerkin coarse-grid operator, $P^T AP$, the two-level method for solving (25) may be represented by the two-grid iteration matrix,

$$T = S_J(I - P(P^T AP)^{-1}P^T A)S_J, \quad \text{with } S_J = I - \frac{4}{5}D^{-1}A,\ D = \text{diag}(A). \tag{26}$$

We choose under-relaxation parameter $4/5$ as we would for the homogeneous problem, noting that the relaxation method remains convergent for this choice. Note that with a 64×64 element fine grid the permeability field in Figure 12(a) is perfectly resolved on both the fine and coarse grid, whereas the permeability field in Figure 12(b) cannot be represented on the coarse grid and thus, operator-induced interpolation and Galerkin coarsening are crucial ingredients to avoid poor multigrid performance.

In the finite-element setting, Neumann boundary conditions make it difficult to apply rigorous Fourier analysis. Therefore, we combine SAMA with LFA. More precisely, we use an LFA approach based on the exponential Fourier modes given in (11) to block-diagonalize the operators associated with the x-direction. Note that ∇u is continuous in the x-direction; the heterogeneity in y is taken into account by SAMA.

Table III compares BoxMG interpolation with bilinear interpolation on a 64×64 element grid with four layers of alternating permeability, not aligned with the coarse grid as depicted in Figure 12(b), of 1 and 10^p for different values of p. We see that SAMA predictions correspond well to both experimentally measured convergence factors and spectral radii of the analytically computed iteration matrix. SAMA accurately predicts that, in contrast to bilinear interpolation, using BoxMG interpolation produces an efficient two-level algorithm independent of the size of the jump in the permeability. Note that when considering coarse-grid aligned layers of varying permeability as depicted in Figure 12(a), the two interpolation schemes coincide.

DOI: 10.1002/nla

	$p=$	0	1	2	3	4
BoxMG interpolation	ρ_{SAMA}	1.60e-01	1.60e-01	1.60e-01	1.60e-01	1.60e-01
	$\overline{\rho}_{181,200}$	1.59e-01	1.60e-01	1.60e-01	1.60e-01	1.60e-01
	ρ_{analytic}	1.60e-01	1.60e-01	1.60e-01	1.60e-01	1.60e-01
bilinear interpolation	ρ_{SAMA}	1.60e-01	4.75e-01	6.15e-01	6.32e-01	6.34e-01
	$\overline{\rho}_{181,200}$	1.60e-01	4.75e-01	6.15e-01	6.32e-01	6.34e-01
	ρ_{analytic}	1.60e-01	4.75e-01	6.15e-01	6.32e-01	6.34e-01

Table III. Average convergence factors per iteration in last 20 of 200 iterations, $\overline{\rho}_{181,200}$, SAMA-predicted spectral radii, ρ_{SAMA}, and spectral radii of the analytically computed iteration matrix, ρ_{analytic}, for two-grid (1,1)-cycles applied to the elliptic model problem with four layers of alternating permeability, not aligned with the coarse grid, of 1 and 10^p for different values of p discretized on a 64×64 element grid.

In Table IV, we consider varying the number of layers of alternating permeability of 1 and 10^{-7} with layer boundaries at $y = 0$, $y = 1$, and $y = n/l + h_y,\ n = 1, \ldots, l-1$, where l denotes the number of layers. Table V looks at varying the grid size used for the discretization. Results show that SAMA accurately predicts that convergence of the two-level method using BoxMG interpolation is also independent of the number of layers and the grid size.

	number of layers	4	8	16
BoxMG interpolation	ρ_{SAMA}	1.60e-01	1.60e-01	1.60e-01
	$\overline{\rho}_{181,200}$	1.60e-01	1.60e-01	1.60e-01
	$\overline{\rho}_{\text{analytic}}$	1.60e-01	1.60e-01	1.60e-01
bilinear interpolation	ρ_{SAMA}	6.34e-01	6.34e-01	6.79e-01
	$\overline{\rho}_{181,200}$	6.34e-01	6.34e-01	6.79e-01
	$\overline{\rho}_{\text{analytic}}$	6.34e-01	6.34e-01	6.79e-01

Table IV. Average convergence factors per iteration in last 20 of 200 iterations, $\overline{\rho}_{181,200}$ and SAMA-predicted spectral radii, ρ_{SAMA}, and spectral radii of the analytically computed iteration matrix, ρ_{analytic}, for two-grid (1,1)-cycles applied to the elliptic model problem with layers of alternating permeability, not aligned with the coarse grid, of 1 and 10^{-7} discretized on a 64×64 element grid.

	number of elements	16×16	32×32	64×64	128×128
BoxMG interpolation	ρ_{SAMA}	1.60e-01	1.60e-01	1.60e-01	1.60e-01
	$\overline{\rho}_{181,200}$	1.60e-01	1.60e-01	1.60e-01	1.60e-01
bilinear interpolation	ρ_{SAMA}	6.79e-01	6.34e-01	6.34e-01	6.34e-01
	$\overline{\rho}_{181,200}$	6.79e-01	6.34e-01	6.34e-01	6.33e-01

Table V. Average convergence factors per iteration in last 20 of 200 iterations, $\overline{\rho}_{181,100}$ and SAMA-predicted spectral radii of the iteration matrix for two-grid (1,1)-cycles applied to the elliptic model problem with four layers of alternating permeability, not aligned with the coarse grid, of 1 and 10^{-7}.

In Figure 13, we consider the spectrum of the two-grid iteration matrix on a 64×64 element grid considering four layers of alternating permeability of 1 and 10^{-7}. Note that since we combine SAMA with LFA, for SAMA we analyze operators of size $64^2 \times 64^2$, whereas analytically computed operators are of size $65 \cdot 64 \times 65 \cdot 64$. Thus, the SAMA-predicted spectrum consists of 64^2 eigenvalues, whereas the spectrum of the analytically computed iteration matrix consists of $65 \cdot 64$ eigenvalues. In addition to the spectral radius considered in the above results, the plots show that the SAMA prediction of the spectrum also corresponds well to the spectrum of the analytically computed iteration matrix.

5.1.2. Deflation The deflation method (or deflated conjugate gradient method) [44] can be used to accelerate the convergence of the preconditioned conjugate gradient method (PCG). Here, motivated by [45], we consider a PCG method whose preconditioner is analogous to (but not equivalent

Prepared using nlaauth.cls
DOI: 10.1002/nla

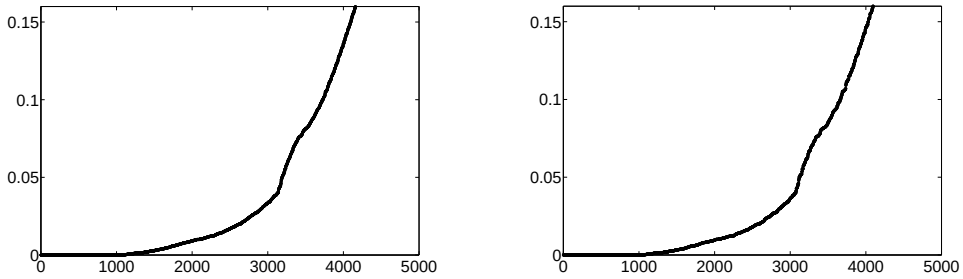

Figure 13. Eigenvalues on a 64×64 element grid for a two-grid (1,1)-cycle applied to the elliptic model problem with four layers of alternating permeability, not aligned with the coarse grid, of 1 and 10^{-7}. At left, eigenvalues of the analytically computed iteration matrix and at right, eigenvalues predicted by SAMA.

to) a two-grid (1,0)-cycle with Jacobi relaxation and a very aggressive coarsening strategy. More precisely, applied to (25), the restriction operator or so-called "deflation-subspace matrix" collapses all points within a layer to a single point with equal weighting. Thus, the number of coarse-grid points corresponds to the number of layers. Choosing interpolation, P, as the transpose of restriction, and denoting the diagonal of A by D, the two-level preconditioner is given by

$$\left(I - P\left(P^T A P\right)^{-1} P^T A\right) D^{-1}. \tag{27}$$

We note that the deflation preconditioner is, in fact, singular, and it is used in conjunction with a complementary solve over its nullspace to yield a solution process; see [45, 61] for more details.

Figure 14 shows the computed and SAMA-predicted spectra on a 64×64 element grid for the deflation preconditioner applied to the elliptic model problem with four layers of alternating permeability, as depicted in Figure 12(b), of 1 and 10^{-7}. Comparing the computed and SAMA-predicted spectra of the preconditioned system, SAMA accurately predicts that the two small eigenvalues resulting from the two layers with small permeability are shifted to zero. Thus, in this particular case of two layers with small permeability using the deflation preconditioner for the conjugate gradient (CG) method decreases the number of CG iterations by two. Note that this property makes deflation preconditioning interesting for problems with a large number of layers with small permeabilities.

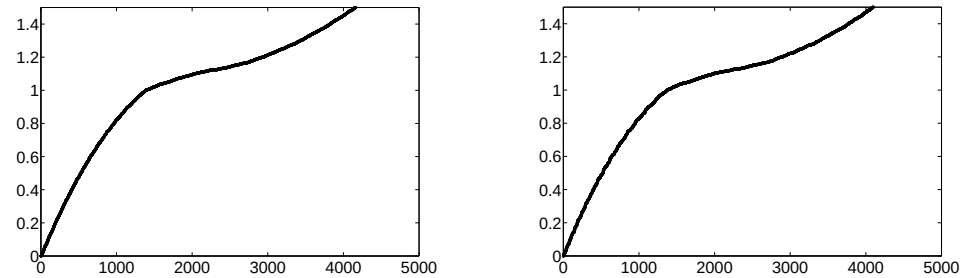

Figure 14. Eigenvalues of the preconditioned matrix for the deflation preconditioner applied to the elliptic model problem with four layers of alternating permeability of 1 and 10^{-7} discretized on a 64×64 element grid. At left, eigenvalues of the analytically computed matrix and at right, eigenvalues predicted by SAMA.

5.2. Convection-diffusion

As a second example, we consider the 2D grid-aligned convection-diffusion problem

$$-\varepsilon \Delta u + u_x = f \quad \text{on } \Omega = [0,1]^2 \tag{28}$$

DOI: 10.1002/nla

with $\varepsilon > 0$, determining the ratio between diffusion and convection, and subject to homogeneous Dirichlet boundary conditions, $u(x, y) = 0$ on $\partial\Omega$.

We discretize our convection-diffusion model problem on a uniform Cartesian mesh consisting of N_x intervals in the x-direction and N_y intervals in the y-direction, using mesh sizes $h_x = 1/N_x$ and $h_y = 1/N_y$, respectively, where N_x and N_y are positive integers. For simplicity, we only consider the case $N_x = N_y$ and mesh size $h = h_x = h_y$. To avoid instability issues arising for small ε when using standard central differencing for the discretization of the convective term [48], we use a first-order upwind discretization. Then, the discretized operator can be written in the form $A = (\varepsilon D_{xx} + D_x) \otimes I_{N_y-1} + I_{N_x-1} \otimes \varepsilon D_{yy}$, where we eliminate Dirichlet boundary points in the discretized operator, and where D_{xx}, D_{yy} and D_x denote the usual central finite difference discretization of $-u_{xx}$ and $-u_{yy}$, respectively, and the first-order upwind discretization of u_x, given in stencil form as

$$D_{xx} = \frac{1}{h^2}\begin{bmatrix} -1 & 2 & -1 \end{bmatrix}, \qquad D_{yy} = \frac{1}{h^2}\begin{bmatrix} -1 \\ 2 \\ -1 \end{bmatrix}, \qquad D_x = \frac{1}{h}\begin{bmatrix} -1 & 1 & 0 \end{bmatrix}.$$

While this problem is simple enough that there is no need to develop or analyze solvers for it specifically, it serves as a prototype model for solvers for general convection-diffusion problems, particularly as might arise in linearizations of high Reynold's number flow and other applications in fluid dynamics. As such, problems such as this one are common test cases in the design of improved iterative methods for non-symmetric matrices; see, for example, [62, 63]. Here, we illustrate the applicability of the SAMA approach by considering its ability to effectively predict both success and failure of two-grid methods, arising from varying the ordering of Gauss-Seidel relaxation sweeps. We consider two relaxation schemes, Gauss-Seidel with downwind ordering of the variables, defined by relaxation operator

$$S_{\text{GS-down}} = I - (D - L)^{-1}A,$$

with D and $-L$ denoting either the diagonal and strictly lower-triangular parts of A in the case of pointwise Gauss-Seidel, or the block-diagonal and strictly block-lower-triangular parts of A, when ordered by lines of nodes in the y-direction, in the case of blockwise relaxation, and Gauss-Seidel with upwind ordering of the variables, defined by relaxation operator

$$S_{\text{GS-up}} = I - (D - U)^{-1}A,$$

with $-U$ denoting the strictly upper-triangular part of A in the case of pointwise Gauss-Seidel, or the strictly block-upper-triangular part of A, when ordered by lines of nodes in the y-direction, in the case of blockwise relaxation. Note that the ordering of the grid in the case of downstream relaxation is according to the convection-direction, $(1, 0)^T$, and, thus, relaxation becomes an exact solver for $\epsilon \to 0$. We consider the two-level algorithm with upstream or downstream Gauss-Seidel relaxation and full standard coarsening by a factor of two in each dimension, using bilinear interpolation, P, and full weighting restriction, R, and rediscretization to get the coarse-grid operator, A_c. The resulting two-level algorithm may be represented by the two-grid iteration matrix,

$$T_{\text{TG-up}} = S_{\text{GS-up}}(I - PA_c^{-1}RA)S_{\text{GS-up}} \quad \text{or} \quad T_{\text{TG-down}} = S_{\text{GS-down}}(I - PA_c^{-1}RA)S_{\text{GS-down}}. \quad (29)$$

As the convection direction cannot be taken into account by rigorous Fourier analysis, we combine SAMA with mode analysis in the y-direction. More precisely, we consider both rigorous Fourier analysis, using the sine Fourier modes to block-diagonalize the operators in the y-direction, for the case of blockwise relaxation, and an LFA approach that uses the standard LFA analysis (in 1D) for the case of pointwise Gauss-Seidel. In both cases, the behavior in the convection direction is accounted for through the SAMA methodology.

Table VI compares the SAMA predictions and performance of two-level (1,1)-cycles based on downwind and upwind relaxation ordering, for different values of the convection coefficient, ε. In the case of block relaxation, SAMA predictions are exact, due to the use of rigorous mode analysis to treat the y-direction, while for point relaxation, an LFA approach is needed and, thus, SAMA is not rigorous. However, for both block and point relaxation, SAMA accurately predicts the efficiency of downstream relaxation as $\varepsilon \to 0$ as well as the failure of upwind ordering in this limit.

Prepared using nlaauth.cls
DOI: 10.1002/nla

	$\varepsilon =$	10^{-2}	10^{-4}	10^{-6}	10^{-8}
downstream	$\|T_{\text{TG-down}}\|_{\text{SAMA}}$	7.56e-01	2.37e-02	4.87e-06	4.91e-10
block Gauss-Seidel	$\|T_{\text{TG-down}}\|_2$	7.56e-01	2.37e-02	4.87e-06	4.91e-10
upstream	$\|T_{\text{TG-up}}\|_{\text{SAMA}}$	6.73e-01	1.13e+00	1.14e+00	1.14e+00
block Gauss-Seidel	$\|T_{\text{TG-up}}\|_2$	6.73e-01	1.13e+00	1.14e+00	1.14e+00
downstream	$\|T_{\text{TG-down}}\|_{\text{SAMA}}$	1.04e+00	4.09e-02	8.58e-06	8.65e-10
point Gauss-Seidel	$\|T_{\text{TG-down}}\|_2$	7.89e-01	4.64e-02	9.59e-06	9.66e-10
upstream	$\|T_{\text{TG-up}}\|_{\text{SAMA}}$	1.03e+00	1.11e+00	1.14e+00	1.14e+00
point Gauss-Seidel	$\|T_{\text{TG-up}}\|_2$	7.33e-01	1.11e+00	1.14e+00	1.14e+00

Table VI. Analytically computed error reduction factors and SAMA-predicted error reduction factors of two-grid (1,1)-cycles with various relaxation schemes applied to the convection-diffusion model problem with different values of ε discretized on a grid of size 64×64.

In Figure 15, we consider analytically computed and SAMA-predicted spectra on a 64×64 grid for a two-grid (1,1)-cycle with downstream block or point relaxation applied to the convection-diffusion model problem with $\varepsilon = 10^{-4}$. In addition to the error reduction rates measured by the operator norm considered in the above results, the plots show that, in the block-relaxation case, the SAMA prediction of the spectrum corresponds well to the spectrum of the analytically computed iteration matrix, whereas in the point-relaxation case it does not, due to the combination with LFA.

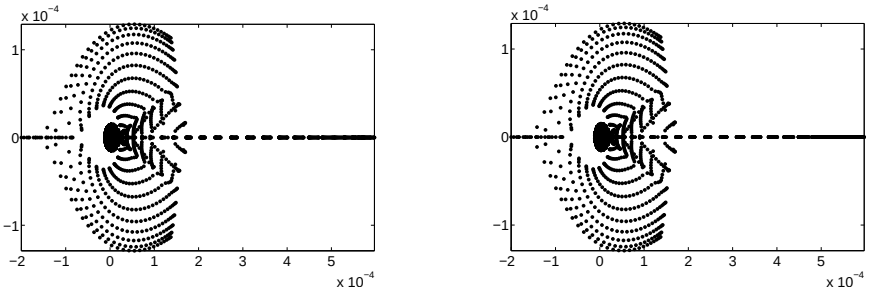

(a) Eigenvalues (in $\mathbb{C}$) for two-grid (1,1)-cycles with downstream block Gauss-Seidel relaxation.

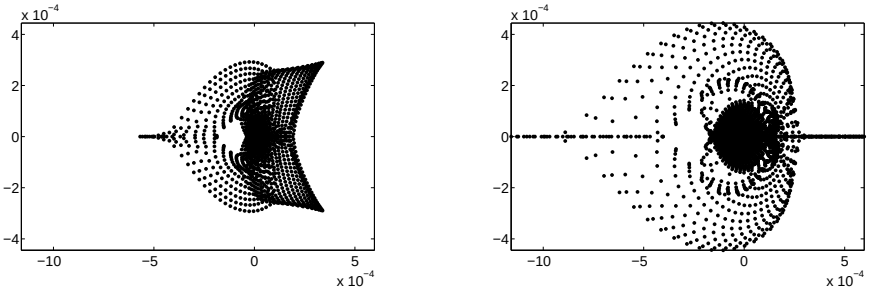

(b) Eigenvalues (in $\mathbb{C}$) for two-grid (1,1)-cycles with downstream point Gauss-Seidel relaxation.

Figure 15. Eigenvalues (in $\mathbb{C}$) for two-grid (1,1)-cycles with downstream (a) block or (b) point Gauss-Seidel relaxation applied to the convection-diffusion model problem with $\varepsilon = 10^{-4}$ discretized on a grid of size 64×64. At left, eigenvalues of the analytically computed iteration matrix and at right, eigenvalues predicted by SAMA.

6. CONCLUSIONS

A new predictive analysis tool for multigrid and related multilevel methods is proposed, which extends the local mode analysis approach to a class of problems for which standard LFA does not

DOI: 10.1002/nla

offer its usual predictivity. The method is applicable to problems with tensor-product structure, where only one "dimension" includes behaviour that is not amenable to either rigorous or local Fourier analysis. Thus, it is a natural analysis tool for parabolic problems in any number of spatial dimensions, where the time-direction has natural non-diagonalizable behaviour with structure akin to a Jordan block. It can also be applied to elliptic-type problems with one-dimensional structure that makes LFA inapplicable or unreliable, such as with layered media or grid-aligned convection. While this remains a restricted class of problems, it enlarges the applicability of mode analysis techniques to include many important problems that were previously intractable in this way.

The proposed SAMA approach, then, couples standard LFA with algebraic computation that accounts for the non-local or heterogeneous character of operators in this class. In addition to providing insight into asymptotic convergence behaviour, SAMA maintains multiplicities of eigenvalues of iteration operators and can be used effectively to examine worst-case short-term convergence through operator norms. We demonstrate that it clearly provides an advantage for parabolic problems, obtaining robust predictivity of performance independent of the length of the time domain, in sharp contrast to LFA, which only becomes predictive for extremely long time integration. We also demonstrate that it can be applied to non-traditional multigrid cycling schemes, such as in the case of simple deflation-based preconditioners.

A. COMBINING ANALYSES USING DIFFERENT FOURIER BASES

For the analysis of the red-black relaxation schemes, we use the red-black Fourier basis given by

$$\{\boldsymbol{\psi}_R(\theta_{c;k}), \boldsymbol{\psi}_B(\theta_{c;k})\}_{k=-\frac{N_x}{4}+1,\ldots,\frac{N_x}{4}}\ ,\ \ (\psi_R(\theta_{c;k}))_j = (\psi_B(\theta_{c;k}))_j = \mathrm{e}^{-\imath\theta_{c;k}j}, \tag{30}$$

with the discrete set of Fourier frequencies,

$$\theta_{c;k} = \frac{4\pi k}{N_x},\ k = -\frac{N_x}{4}+1,\ldots,\frac{N_x}{4}. \tag{31}$$

The analysis of coarse-grid correction is based on the harmonic Fourier basis given by

$$\{\boldsymbol{\psi}(\theta_k), \boldsymbol{\psi}(\theta_{k'})\}_{k=-\frac{N_x}{4}+1,\ldots,\frac{N_x}{4}}\ ,\ \psi_j(\theta_k) = \mathrm{e}^{-\imath\theta_k j}, \tag{32}$$

with the discrete set of harmonic Fourier frequencies,

$$\theta_k = \frac{2\pi k}{N_x},\ \theta_{k'} = \theta_k - \mathrm{sign}(\theta_k)\pi,\ k = -\frac{N_x}{4}+1,\ldots,\frac{N_x}{4}.$$

Denoting the Fourier matrix of red-black Fourier modes (30) by Ψ_{RB} and the Fourier matrix of harmonic Fourier modes (32) by Ψ_{harm}, the analyses of relaxation and coarse-grid correction can be combined by using the transformation

$$\Psi_{\text{harm}} = \mathcal{P}_{RB}\Psi_{RB}\mathcal{T},$$

where $\mathcal{P}_{RB}$ is the permutation matrix that puts Ψ_{RB} into red-black block ordering and $\mathcal{T}$ is given by

$$\mathcal{T} = \begin{bmatrix} 1 & 1 & & & & & & & \\ & & 1 & 1 & & & & & \\ & & & & \ddots & & & & \\ & & & & & & 1 & 1 & \\ \mathrm{e}^{\imath\theta_{-\frac{N_x}{4}+1}} & -e^{\imath\theta_{-\frac{N_x}{4}+1}} & & & & & & & \\ & & \mathrm{e}^{\imath\theta_{-\frac{N_x}{4}+2}} & -e^{\imath\theta_{-\frac{N_x}{4}+2}} & & & & & \\ & & & & \ddots & & & & \\ & & & & & & \mathrm{e}^{\imath\theta_{\frac{N_x}{4}}} & -e^{\imath\theta_{\frac{N_x}{4}}} \end{bmatrix}.$$

Prepared using nlaauth.cls
DOI: 10.1002/nla

REFERENCES

1. Brandt A. Multi–level adaptive solutions to boundary–value problems. *Mathematics of Computation* 1977; **31**:333–390.
2. Stevenson R. On the validity of local mode analysis of multi-grid methods. PhD Thesis, Utrecht University, Utrecht, the Netherlands 1990.
3. Brandt A. Rigorous quantitative analysis of multigrid. I. Constant coefficients two-level cycle with L_2-norm. *SIAM Journal on Numerical Analysis* 1994; **31**(6):1695–1730.
4. Stüben K, Trottenberg U. Multigrid methods: Fundamental algorithms, model problem analysis and applications. *Multigrid Methods*, *Lecture Notes in Mathematics*, vol. 960, Hackbusch W, Trottenberg U (eds.), Springer–Verlag: Berlin, 1982; 1–176.
5. Niestegge A, Witsch K. Analysis of a multigrid Stokes solver. *Applied Mathematics and Computation* 1990; **35**(3):291–303.
6. Sivaloganathan S. The use of local mode analysis in the design and comparison of multigrid methods. *Computer Physics Communications* 1991; **65**:246–252.
7. Wienands R, Gaspar FJ, Lisbona FJ, Oosterlee CW. An efficient multigrid solver based on distributive smoothing for poroelasticity equations. *Computing. Archives for Scientific Computing* 2004; **73**(2):99–119.
8. Boonen T, van Lent J, Vandewalle S. Local Fourier analysis of multigrid for the curl-curl equation. *SIAM Journal on Scientific Computing* 2008; **30**(4):1730–1755.
9. MacLachlan SP, Oosterlee CW. Local Fourier analysis for multigrid with overlapping smoothers applied to systems of PDEs. *Numerical Linear Algebra with Applications* 2011; **18**:751–774.
10. Zhou G, Fulton SR. Fourier analysis of multigrid methods on hexagonal grids. *SIAM Journal on Scientific Computing* 2009; **31**(2):1518–1538.
11. Gaspar F, Gracia J, Lisbona F. Fourier analysis for multigrid methods on triangular grids. *SIAM Journal on Scientific Computing* 2009; **31**(3):2081–2102.
12. Rodrigo C, Gaspar FJ, Lisbona FJ. Multigrid methods on semi-structured grids. *Archives of Computational Methods in Engineering. State of the Art Reviews* 2012; **19**(4):499–538.
13. Gmeiner B, Gradl T, Gaspar F, Rüde U. Optimization of the multigrid-convergence rate on semi-structured meshes by local Fourier analysis. *Computers & Mathematics with Applications. An International Journal* 2013; **65**(4):694–711.
14. Rodrigo C, Salinas P, Gaspar F, Lisbona F. Local Fourier analysis for cell-centered multigrid methods on triangular grids. *Journal of Computational and Applied Mathematics* 2014; **259, Part A**(0):35 – 47. Proceedings of the Sixteenth International Congress on Computational and Applied Mathematics (ICCAM-2012), Ghent, Belgium, 9-13 July, 2012.
15. Seynaeve B, Rosseel E, Nicolaï B, Vandewalle S. Fourier mode analysis of multigrid methods for partial differential equations with random coefficients. *Journal of Computational Physics* 2007; **224**(1):132–149.
16. Wienands R, Oosterlee CW, Washio T. Fourier analysis of $\mathrm{GMRES}(m)$ preconditioned by multigrid. *SIAM Journal on Scientific Computing* 2000; **22**(2):582–603.
17. Hemker PW, Hoffmann W, van Raalte MH. Fourier two-level analysis for discontinuous Galerkin discretization with linear elements. *Numerical Linear Algebra with Applications* 2004; **11**(5-6):473–491.
18. Brandt A. Multigrid solvers for non-elliptic and singular-perturbation steady-state problems. *Technical report*, Department of Applied Mathematics, The Weizmann Institute of Science 1981.
19. Brandt A. *Multigrid techniques: 1984 guide with applications to fluid dynamics*. GMD–Studien Nr. 85, Gesellschaft für Mathematik und Datenverarbeitung: St. Augustin, 1984.
20. Brandt A, Yavneh I. On multigrid solution of high-Reynolds incompressible entering flows. *Journal of Computational Physics* 1992; **101**(1):151–164.
21. Vandewalle S, Horton G. Fourier mode analysis of the multigrid waveform relaxation and time-parallel multigrid methods. *Computing* 1995; **54**(4):317–330.
22. Friedhoff S, MacLachlan S, Börgers C. Local Fourier analysis of space-time relaxation and multigrid schemes. *SIAM Journal on Scientific Computing* 2013; **35**(5):S250–S276.
23. Nievergelt J. Parallel methods for integrating ordinary differential equations. *Communications of the Association for Computing Machinery* 1964; **7**:731–733.
24. Hackbusch W. Parabolic multigrid methods. *Computing methods in applied sciences and engineering, VI (Versailles, 1983)*. North-Holland: Amsterdam, 1984; 189–197.
25. Bastian P, Burmeister J, Horton G. Implementation of a parallel multigrid method for parabolic partial differential equations. *Parallel Algorithms for PDEs, Proc. 6th GAMM Seminar Kiel, January 19-21, 1990*, Hackbusch W (ed.), Vieweg Verlag: Braunschweig, 1990; 18–27.
26. Burmeister J, Horton G. Time-parallel multigrid solution of the Navier-Stokes equations. *Multigrid methods, III (Bonn, 1990)*, *Internat. Ser. Numer. Math.*, vol. 98. Birkhäuser: Basel, 1991; 155–166.
27. Horton G. The time-parallel multigrid method. *Communications in Applied Numerical Methods* 1992; **8**(9):585–595.
28. Horton G, Knirsch R. A time-parallel multigrid-extrapolation method for parabolic partial differential equations. *Parallel Computing. Theory and Applications* 1992; **18**(1):21–29.
29. Lubich C, Ostermann A. Multigrid dynamic iteration for parabolic equations. *BIT. Numerical Mathematics* 1987; **27**(2):216–234.
30. Vandewalle S, Piessens R. Efficient parallel algorithms for solving initial-boundary value and time-periodic parabolic partial differential equations. *SIAM Journal on Scientific and Statistical Computing* 1992; **13**(6):1330–1346.
31. Vandewalle SG, Van de Velde EF. Space-time concurrent multigrid waveform relaxation. *Annals of Numerical Mathematics* 1994; **1**(1-4):347–360. Scientific computation and differential equations (Auckland, 1993).

DOI: 10.1002/nla

32. Janssen J, Vandewalle S. Multigrid waveform relaxation on spatial finite element meshes. *Contributions to multigrid (Amsterdam, 1993)*, *CWI Tract*, vol. 103. Math. Centrum Centrum Wisk. Inform.: Amsterdam, 1994; 75–86.
33. Horton G, Vandewalle S, Worley P. An algorithm with polylog parallel complexity for solving parabolic partial differential equations. *SIAM Journal on Scientific Computing* 1995; **16**(3):531–541.
34. Weinzierl T, Köppl T. A geometric space-time multigrid algorithm for the heat equation. *Numerical Mathematics. Theory, Methods and Applications* 2012; **5**(1):110–130.
35. Falgout R, Friedhoff S, Kolev T, MacLachlan S, Schroder J. Parallel time integration with multigrid 2013; Submitted.
36. Janssen J, Vandewalle S. Multigrid waveform relaxation on spatial finite element meshes: the discrete-time case. *SIAM Journal on Scientific Computing* 1996; **17**(1):133–155.
37. Brandt A, Diskin B. Multigrid solvers for the non-aligned sonic flow: the constant coefficient case. *Computers and Fluids* 1999; **28**(4–5):511 – 549.
38. Diskin B, Thomas JL. Half-space analysis of the defect-correction method for Fromm discretization of convection. *SIAM Journal on Scientific Computing* 2000; **22**(2):633–655.
39. J van Lent, Vandewalle S. Multigrid waveform relaxation for anisotropic partial differential equations. *Numerical Algorithms* 2002; **31**(1-4):361–380.
40. Janssen J, Vandewalle S. Multigrid waveform relaxation of spatial finite element meshes: the continuous-time case. *SIAM Journal on Numerical Analysis* 1996; **33**(2):456–474.
41. Diskin B, Thomas JL, Mineck RE. On quantitative analysis methods for multigrid solutions. *SIAM Journal on Scientific Computing* 2005; **27**(1):108–129.
42. Alcouffe RE, Brandt A, Dendy JE, Painter JW. The multigrid method for the diffusion equation with strongly discontinuous coefficients. *SIAM Journal on Scientific and Statistical Computing* 1981; **2**:430–454.
43. Dendy JE. Black box multigrid. *Journal of Computational Physics* 1982; **48**:366–386.
44. Nicolaides RA. Deflation of conjugate gradients with applications to boundary value problems. *SIAM Journal on Numerical Analysis* 1987; **24**(2):355–365.
45. Vuik C, Segal A, Meijerink JA. An efficient preconditioned CG method for the solution of a class of layered problems with extreme contrasts in the coefficients. *Journal of Computational Physics* 1999; **152**:385–403.
46. Ries M, Trottenberg U. MGR-ein blitzschneller elliptischer Löser. *Technical Report Preprint 277 SFB 72*, Universität Bonn 1979.
47. Ries M, Trottenberg U, Winter G. A note on MGR methods. *Linear Algebra and its Applications* 1983; **49**:1–26.
48. Trottenberg U, Oosterlee CW, Schüller A. *Multigrid*. Academic Press: London, 2001.
49. Lions JL, Maday Y, Turinici G. Résolution d'EDP par un schéma en temps "pararéel". *Comptes Rendus de l'Académie des Sciences. Série I. Mathématique* 2001; **332**(7):661–668.
50. Oosterlee CW, Gaspar FJ, Washio T, Wienands R. Multigrid line smoothers for higher order upwind discretizations of convection-dominated problems. *Journal of Computational Physics* 1998; **139**(2):274–307.
51. Chan RH, Ng MK. Conjugate gradient methods for Toeplitz systems. *SIAM Review. A Publication of the Society for Industrial and Applied Mathematics* 1996; **38**(3):427–482.
52. Lirkov ID, Margenov SD, Zikatanov LT. Circulant block-factorization preconditioning of anisotropic elliptic problems. *Computing* 1997; **58**(3):245–258.
53. Trefethen LN, Embree M. *Spectra and Pseudospectra*. Princeton University Press: Princeton, 2005.
54. Wright T. Eigtool. http://www.comlab.ox.ac.uk/pseudospectra/eigtool/ 2002.
55. Wienands R, Oosterlee CW. On three-grid Fourier analysis for multigrid. *SIAM Journal on Scientific Computing* 2001; **23**(2):651–671.
56. Hackbusch W. *Multigrid Methods and Applications*, *Computational Mathematics*, vol. 4. Springer–Verlag: Berlin, 1985.
57. Ruge JW, Stüben K. Algebraic multigrid (AMG). *Multigrid Methods*, *Frontiers in Applied Mathematics*, vol. 3, McCormick SF (ed.). SIAM: Philadelphia, PA, 1987; 73–130.
58. Xu J, Zhu Y. Uniform convergent multigrid methods for elliptic problems with strongly discontinuous coefficients. *Mathematical Models & Methods in Applied Sciences* 2008; **18**(1):77–105.
59. Graham IG, Scheichl R. Robust domain decomposition algorithms for multiscale PDEs. *Numerical Methods for Partial Differential Equations. An International Journal* 2007; **23**(4):859–878.
60. Graham IG, Lechner PO, Scheichl R. Domain decomposition for multiscale PDEs. *Numerische Mathematik* 2007; **106**(4):589–626.
61. Nabben R, Vuik C. A comparison of deflation and coarse grid correction applied to porous media flow. *SIAM Journal on Numerical Analysis* 2004; **42**(4):1631–1647.
62. Sala M, Tuminaro RS. A new Petrov-Galerkin smoothed aggregation preconditioner for nonsymmetric linear systems. *SIAM Journal on Scientific Computing* 2008; **31**(1):143–166.
63. Brezina M, Manteuffel T, McCormick S, Ruge J, Sanders G. Towards adaptive smoothed aggregation (αSA) for nonsymmetric problems. *SIAM Journal on Scientific Computing* 2010; **32**:14–39.

DOI: 10.1002/nla