

ROBUST SOLUTION OF SINGULARLY PERTURBED PROBLEMS USING MULTIGRID METHODS*

SCOTT MACLACHLAN[†] AND NIALL MADDEN[‡]

Abstract. We consider the problem of solving linear systems of equations that arise in the numerical solution of singularly perturbed ordinary and partial differential equations of reaction-diffusion type. Standard discretization techniques are not suitable for such problems and, so, specially tailored methods are required, usually involving adapted or fitted meshes that resolve important features such as boundary and/or interior layers. In this study, we consider classical finite difference schemes on the layer adapted meshes of Shishkin and Bakhvalov. We show that standard direct solvers exhibit poor scaling behavior, with respect to the perturbation parameter, when solving the resulting linear systems. We propose and prove optimality of a new block-structured preconditioning approach that is robust for small values of the perturbation parameter, and compares favorably with standard robust multigrid preconditioners for these linear systems. We also derive stopping criteria which ensure that the potential accuracy of the layer-resolving meshes is achieved.

Key words. boundary-fitted meshes, robust multigrid, preconditioning

AMS subject classifications. 65F10, 65N06, 65N22, 65N55

DOI. 10.1137/120889770

1. Introduction. This study addresses the problem of solving linear systems that arise when computing numerical solutions to certain linear singularly perturbed boundary value problems in one and two dimensions. These differential equations are characterized by a small positive parameter, usually denoted as ε , multiplying the highest derivative. The perturbation is “singular” in the sense that, as $\varepsilon \rightarrow 0$, the problem becomes ill-posed since the order of the differential equation is reduced, but the number of boundary conditions remains the same. (For a more formal definition, see [28, Chap. 1].)

The simplest example of a singularly perturbed problem is

$$(1.1) \quad -\varepsilon^2 u'' + a(x)u'(x) + b(x)u = f(x) \text{ on } (0, 1), \quad u(0) = 0, u(1) = 0.$$

When $b \equiv 0$ this is known as a *convection-diffusion* problem, whereas if $a \equiv 0$ and $b \neq 0$, it is of *reaction-diffusion* type. Such problems, and their higher-dimensional analogues, are common in mathematical models. Convection-diffusion problems are widespread in many formulations of fluid-flow problems (e.g., in linearized Navier–Stokes equations, and transport problems), and simulation of semiconductor devices; see [40] and [34] for a more exhaustive list. Coupled systems of reaction-diffusion equations are standard in many biological applications, simulation of chemical re-

*Submitted to the journal’s Methods and Algorithms for Scientific Computing section August 31, 2012; accepted for publication (in revised form) July 30, 2013; published electronically September 26, 2013. This research was supported in part by the Institute for Mathematics and its Applications with funds provided by the National Science Foundation.

<http://www.siam.org/journals/sisc/35-5/88977.html>

[†]Department of Mathematics, Tufts University, 503 Boston Avenue, Medford, MA 02155 (scott.maclachlan@tufts.edu). This author’s work was partially supported by the National Science Foundation under grant DMS-0811022.

[‡]School of Mathematics, Statistics and Applied Mathematics, National University of Ireland, Galway, Ireland (Niall.Madden@NUIGalway.ie). This author’s work was supported by the Science Foundation of Ireland under Grants 08/RFP/CMS1205 and Mathematics Initiative 07/MI/007.

actions, and in hydrodynamic stability. In each case, the solution to the singularly perturbed problem is characterized by the presence of boundary or interior *layers*: narrow regions of the domain where the solution changes rapidly.

The numerical solution of these problems is of significant mathematical interest. Classical numerical schemes that are suitable when ε is $\mathcal{O}(1)$ are often inappropriate as $\varepsilon \rightarrow 0$, unless the number of degrees of freedom in one dimension, N , satisfies a relation such as $N = \mathcal{O}(\varepsilon^{-1})$: without this, they may fail to resolve layers—usually the region of most interest; the order of convergence may be diminished for small ε , a phenomenon sometimes referred to as *locking* [3]; in the case of convection-diffusion problems, the method may become unstable and fail to yield any useful information.

Much of the difficulty in applying numerical methods to problems such as (1.1) stems from the fact that the derivatives of these solutions depend on negative powers of the perturbation parameter. Since estimates for the errors in numerical solution generated by classical schemes depend on bounds for these derivatives, they are not *parameter robust* meaning that they do not hold for arbitrarily small values of the perturbation parameter.

The development of algorithms that are robust with respect to the perturbation parameter, and resolve any layers present, is a very active field of endeavor. See, for example, [18, 28, 33, 40, 44], and the many references therein. Such so-called *parameter robust* (also known as *uniformly convergent* or “ ε -uniform”) methods guarantee that the computational effort required to obtain a certain accuracy is the same, for example, when $\varepsilon = 10^{-6}$ as when $\varepsilon = 10^{-2}$. The majority of these papers consider the application of standard schemes, such as finite difference and finite element methods, on specially constructed (a priori or a posteriori) fitted meshes, most often the piecewise uniform meshes introduced by Miller, O’Riordan, and Shishkin [33] or the graded meshes devised by Bakhvalov [4], described below in section 2.1.1.

These algorithms produce linear systems that must be solved, but it is notable that there are relatively few studies concerning their numerical solution. This is significant because, particularly for problems in more than one dimension, standard solvers are unlikely to be useful; the development of fast robust solvers is important and challenging. Moreover, most studies usually assume that the computational effort for solving the linear systems is independent of ε ; few have considered the issue of solving the linear systems with efficiency that is robust with respect to ε . We argue that this should not be taken for granted. First, the performance of direct solvers degrades badly with dimension, requiring $\mathcal{O}(N^3)$ memory and operations for two-dimensional problems on $N \times N$ meshes, with a higher penalty for problems in three (or more) dimensions. Moreover, it is usually assumed that the performance of direct solvers depend only on the matrix structure, and not its entries. We show in section 4.1 that this assumption does not hold for singularly perturbed PDEs and, for small ε , the performance degrades due to the presence of subnormal numbers. Therefore, direct methods are not parameter robust. Finally, the performance and analysis of most iterative schemes, particularly those involving robust preconditioners, is highly dependent on both the scheme and the underlying differential equation.

Therefore, it is surprising that there is so little in the literature on development of solvers for discretization schemes designed for singularly perturbed problems. For convection-diffusion problems, several studies exist, including [42], which considers the conditioning of the resulting discretizations on certain meshes and the effects of diagonal scaling; Farrell and Shishkin give a short analysis of a Gauss–Seidel method for a convection-diffusion problem in [19]; while, in [2], results of experiments with ILU-based preconditioners are reported. Multigrid methods for convection-diffusion

problems on Shishkin meshes are discussed in [20, 21], where a scalable multigrid scheme is introduced.

For reaction-diffusion problems, most of the multigrid literature focuses on the case of a singularly perturbed problem discretized on a uniform or quasi-uniform mesh. For example, in [36], it is shown that a standard multigrid method applied to the two- and three-dimensional analogues of (1.1) on a quasi-uniform mesh converges with bound independent of mesh size or ε ; only a small remark is made about lack of accuracy within the boundary layers. A hierarchical basis approach is discussed in [47]; however, the restriction in this work that the mesh aspect ratios be uniformly bounded is not satisfied by the tensor products of fitted meshes considered here. In contrast, there is extensive literature on the combination of multigrid methods with adaptive refinement algorithms, ranging from fundamental principles discussed in [7] to recent work on achieving efficiency in massively parallel environments [5]; clearly such approaches yield efficient solution algorithms for the problems considered here. Our interests, however, are in the cases where we have enough a priori knowledge to avoid the costs of grid adaptation, but still require efficient solvers on highly nonuniform fitted meshes. Overall, direct methods remain the methods of choice for singularly perturbed reaction-diffusion problems even though, as we show in section 4.1, they cannot be considered parameter robust. Instead, in order to achieve both scalability and robustness, we apply both existing and novel iterative approaches, based on multigrid principles, to these problems.

While we focus here on the case of singularly perturbed problems, we note that our approaches could also be applied to other problems where there is a substantial mismatch between the scaling of terms in the discrete equations over different parts of the domain. One such case, of key current interest, arises in the simulation of time-harmonic wave propagation. In the development of discretizations for either the Helmholtz equation or Maxwell's equations, special attention is always paid to the treatment of discrete approximations of the Sommerfeld radiation condition for outgoing waves. In many cases, in order to attain meaningful solutions over a physical domain of interest, a much larger volume needs to be modeled, implementing some form of absorbing boundary layer to prevent unphysical reflections of outgoing waves at the edges of the discretized domain. While this may be achieved simply by introducing attenuative terms within the layer [46, 50], a more common approach is to couple discretized attenuative terms with a significant increase in grid spacing [14, 16, 35, 54]. While some work has been done on the development of efficient multigrid approaches for these grid structures [39, 54], the methods proposed here may lead to simpler algorithms with improved efficiency for these grids.

1.1. Outline. The paper is organized as follows. Following a short description of notation in section 1.2, in section 1.3 we describe the hardware and libraries used to generate the numerous tables we provide. We introduce our model problems and examples in section 1.4. In section 2, we introduce the standard finite difference schemes for one- and two-dimensional problems, and introduce two standard parameter robust finite difference methods: one based on the piecewise uniform meshes of Shishkin, and the other on the graded meshes of Bakhvalov. In section 3, we review a geometric multigrid method for solving the one-dimensional problem, and propose a new preconditioner for this problem motivated by the graded meshes. Section 4 focuses on linear solvers for the two-dimensional problem, outlining the poor scaling seen in direct solvers and geometric multigrid, as well as optimal general-purpose solvers based on the algebraic and black-box multigrid methodologies. A new boundary-layer pre-

conditioning approach is proposed and analyzed in section 4.5, while stopping criteria are discussed in section 4.6, followed by the results of numerical experiments in section 4.7. A brief discussion of generalizations to higher-dimensional problems appears in section 5, with conclusions and some comments on other extensions of this work in section 6.

1.2. Notation. We denote by N the number of intervals in one dimension of a mesh, and by c and C generic positive constants that are independent of ε and N . Since we are mainly concerned with discretization by finite difference methods, we use $\|\cdot\|_\infty$ to denote the discrete maximum norm on a given mesh:

$$\|u\|_\infty = \begin{cases} \max_{x_i \in \omega^N} |u(x_i)| & \text{on the one-dimensional mesh } \omega^N, \\ \max_{(x_i, y_j) \in \omega^{N \times N}} |u(x_i, y_j)| & \text{on the two-dimensional mesh } \omega^{N \times N}. \end{cases}$$

The continuous analogue on the domain Ω is denoted $\|\cdot\|_{\Omega, \infty}$. Other norms used are the discrete ℓ_2 norm, $\|\cdot\|_2$, and the A -norm for symmetric and positive-definite matrices, A , $\|V\|_A = (V^T A V)^{1/2}$ for any vector, V . We reserve the use of $\|\cdot\|$ for denoting a generic norm, which may be any of the above.

We will use two parameters to measure the singularly perturbed nature of a discrete problem. Define $\delta_N = (\varepsilon N / \beta)^2$ to indicate if a problem is singularly perturbed relative to a mesh with N points, when $\delta_N \ll 1$. For the purposes of the theory developed in section 4.5, we use the term *boundary-fitted mesh* to mean a mesh that is uniform in the interior of the domain, but condenses near the boundary. This uniform mesh width away from boundaries is denoted h_I , and we define $\delta_h = (\varepsilon / (h_I \beta))^2$ to indicate the diagonal dominance of the matrix over the interior degrees of freedom. We note that, just as for δ_N , $\delta_h \ll 1$ when a problem is singularly perturbed relative to the mesh.

1.3. Computer implementation and timings. All numerical results in this paper are computed using code (both C and Fortran) compiled with full optimization, executed on a Beowulf cluster using a single core of a node with an AMD Opteron 2427, 2200 MHz processor with 32Gb of RAM. The main driver routines are written in C to compute the fitted meshes and assemble the matrices, while the CHOLMOD [11], BoxMG [15], and our own AMG libraries are linked to provide linear solvers. Both BoxMG and AMG provide their own PCG wrappers. The boundary-layer preconditioners in one and two dimensions are implemented (separately) in C within their own PCG wrappers.

1.4. Model problems and examples. In this paper, we consider *reaction-diffusion problems* in one and two dimensions. The general form of the one-dimensional problem is

$$(1.2a) \quad L_\varepsilon u := -\varepsilon^2 u'' + b(x)u = f(x) \quad \text{on } \Omega = (0, 1),$$

subject to the boundary conditions

$$(1.2b) \quad u(0) = 0, u(1) = 0.$$

We shall assume that $\varepsilon \in (0, 1]$ and that there exists β such that $0 < \beta^2 < b(x)$ for all $x \in \Omega$. It is easy to show that the operator L_ε satisfies a maximum principle (see, e.g., [38]), and that the problem possesses a unique solution. If $\varepsilon \ll 1/\beta$, then

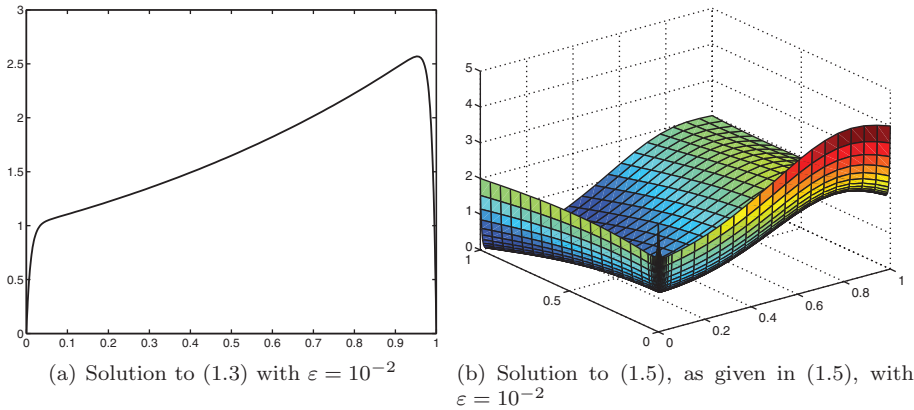


FIG. 1.1. Examples of solutions to one- and two-dimensional singularly perturbed reaction-diffusion problems.

the problem is singularly perturbed, and we expect its solution to exhibit boundary layers, with rapid changes in $u(x)$, near $x = 0$ and $x = 1$.

As an example, consider the problem

$$(1.3) \quad -\varepsilon^2 u'' + u = e^x \text{ on } (0, 1), \quad u(0) = u(1) = 0.$$

When $\varepsilon < 1$, the solution can be expressed as

$$u(x) = \frac{e^{-x/\varepsilon}(e^{1-1/\varepsilon} - 1) + e^{-(1-x)/\varepsilon}(e^{1/\varepsilon} - e)}{(1 - \varepsilon^2)(1 - e^{-2/\varepsilon})} + \frac{e^x}{1 - \varepsilon^2}.$$

One can consider the expressions $\exp(-x/\varepsilon)$ and $\exp(-(1-x)/\varepsilon)$ as representing the layer components. A typical solution exhibiting these layers is shown in Figure 1.1(a).

Our model two-dimensional problem is

$$(1.4) \quad -\varepsilon^2 \Delta u + b(x, y)u = f(x, y) \quad \text{on } \Omega := (0, 1)^2, \quad u(x, y) = g(x, y) \quad \text{on } \partial\Omega,$$

where, again, we assume that there is a positive β such that $0 < \beta^2 < b(x, y)$ for all $(x, y) \in \Omega$. Subject to sufficient regularity and compatibility of b , f , and g , this problem has a unique solution; we refer readers to, e.g., [25] for technical details. When ε is small, the solution may have boundary and corner layers.

As an example of a two-dimensional problem, we consider a variant on a standard test problem (see, e.g., [12]). Although, in general, one would expect solutions to (1.4) to have four boundary and four corner layers, for simplicity of exposition, we have constructed one that has only two boundary layers, near the edges $x = 0$ and $y = 0$, and a corner layer near $(0, 0)$. We take $b(x, y) = 1$ and choose f and g so that

$$(1.5) \quad u = x^3(1 + y^2) + \sin(\pi x^2) + \cos(\pi y/2) + (1 + x + y)(e^{-2x/\varepsilon} + e^{-2y/\varepsilon}).$$

This solution, in the case $\varepsilon = 10^{-2}$, is shown in Figure 1.1(b).

2. Parameter robust methods. The robust solution of singularly perturbed problems can be achieved using fitted operator schemes (i.e., specially designed methods, but used, say, on a uniform mesh) or fitted mesh methods—standard schemes employed on specially designed meshes. The latter approach has received the most

attention of late, not least because fitted meshes are easier to generalize to high-dimensional problems. These fitted mesh methods are categorized as either a priori or a posteriori (equivalently, “fitted” or “adaptive”). An a priori method is constructed based on a careful analysis of the asymptotic properties of the solution and its derivative; most published work considers such schemes. Alternatively, adaptive schemes may be generated based on a posteriori analysis; see, for example, [9, 26]. In this paper, we consider only fitted mesh methods. However, those meshes generated by adaptive schemes tend to be very similar to the Bakhvalov meshes we discuss below and, as such, we expect that similar techniques could be used for the meshes generated by adaptive schemes.

2.1. Finite difference scheme for one-dimensional problems. Given an arbitrary one-dimensional grid $\omega^N = \{0 = x_0 < x_1 < \dots < x_N = 1\}$, with $h_i = x_i - x_{i-1}$ for $i = 1, \dots, N$, the natural second-order finite difference discretization of problem (1.2) provides a potentially unsymmetric discretization of a Hermitian operator. Here, we consider the symmetrized finite difference method for the problem (1.2),

$$(2.1) \quad -\varepsilon^2 \left(\frac{U_{i+1} - U_i}{h_{i+1}} - \frac{U_i - U_{i-1}}{h_i} \right) + \bar{h}_i b(x_i) U_i = \bar{h}_i f(x_i) \quad \text{for } i = 1, \dots, N-1,$$

$$U_i = 0 \quad \text{for } i \in \{0, N\},$$

where $\bar{h}_i = (h_{i+1} + h_i)/2$. In matrix form, we write the symmetrically scaled equations, with boundaries eliminated, as $AU = F$.

If the mesh ω^N is uniform then, in both theory and practice, one must make the unreasonable assumption that N is $\mathcal{O}(\varepsilon^{-1})$ in order to obtain a convergent, layer-resolving method. Intuitively, it seems likely that a robust scheme could be generated if h_i is $\mathcal{O}(\varepsilon)$, but only in the region of the layers. This is indeed the case, but the construction of such meshes must be very careful if one is to rigorously prove robustness. We consider two examples of such meshes below: a simple piecewise uniform mesh, and a more accurate graded mesh.

2.1.1. Fitted meshes for one-dimensional problems. The most popular boundary-fitted mesh for singularly perturbed problems to be found in the mathematical literature is certainly the piecewise uniform mesh of Miller, O’Riordan, and Shishkin [33]. For a problem such as (1.2), it may be formed as follows: assuming the number of mesh intervals N is divisible by 4, define the *mesh transition point* to be

$$(2.2) \quad \tau_S = \min \left\{ \frac{1}{4}, 2 \frac{\varepsilon}{\beta} \ln N \right\},$$

and divide $[0, 1]$ into subintervals $[0, \tau_S]$, $[\tau_S, 1 - \tau_S]$ and $[1 - \tau_S, 1]$. A piecewise uniform mesh is constructed by subdividing $[\tau_S, 1 - \tau_S]$ into $N/2$ equidistant mesh intervals, and subdividing each of $[0, \tau_S]$ and $[1 - \tau_S, 1]$ into $N/4$ equidistant mesh intervals. See Figure 2.1(a) for the two-dimensional analogue of this mesh.

This mesh was first proposed in [45], and an accessible analysis of the uniform convergence of the finite difference method (2.1) applied to the linear reaction-diffusion problem (1.2) is given in [33, Chap. 6]. That analysis uses ideas involving decomposition of the solution into regular and layer parts and exploits the fact that the continuous and discrete operators satisfy maximum principles. A unified treatment based on discrete Green’s functions is given in [28], from which it quickly follows that there is a constant C which is independent of both N and ε such that that the solution

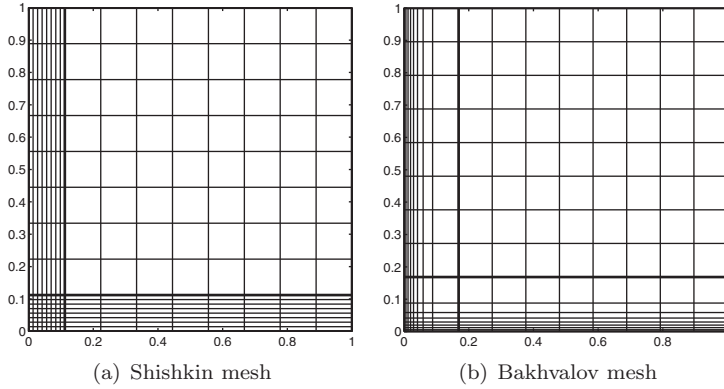


FIG. 2.1. *Fitted meshes for problem (1.5) when $N = 16$, $\varepsilon = 2 \times 10^{-2}$.*

to the finite difference scheme (2.1) satisfies

$$(2.3) \quad \|u - U\|_\infty = \max_{i=0,1,\dots,N} |u(x_i) - U_i| \leq CN^{-2} \ln^2 N.$$

Since C does not depend on ε , this qualifies as a parameter robust estimate. It is not immediately obvious from (2.3) that the numerical solution also resolves the boundary layers. However, one can also show (see, e.g., [28, Thm. 6.12]) that

$$\|u - \bar{U}\|_{\Omega,\infty} = \max_{0 \leq x \leq 1} |u(x) - \bar{U}(x)| \leq CN^{-2} \ln^2 N,$$

where \bar{U} is the piecewise linear interpolant to U .

In the case where ε is $\mathcal{O}(1)$, the scheme (2.1) applied on a uniform mesh should yield a numerical solution with error that is bounded by terms of $\mathcal{O}(N^{-2})$. The logarithmic factor that spoils (2.3) slightly is the price one pays for the simplicity of employing a piecewise uniform mesh. To regain full second-order convergence, one could use the more sophisticated nonuniform boundary-fitted mesh of Bakhvalov [4]. Like the Shishkin mesh, it is uniform over the interior of the domain (a fact which simplifies our analysis later in section 4.5), but is graded within the boundary layer; a sketch of such a mesh in two dimensions is shown in Figure 2.1(b). For full details on the generation of the mesh, see [28].

For the Bakhvalov mesh, it can be proved [28] that there is a constant C independent of ε and N such that

$$\|u - U\|_\infty \leq CN^{-2},$$

and, furthermore, as with the Shishkin mesh, the piecewise linear interpolant to the Bakhvalov solution is second-order convergent. While there is a substantial literature devoted to Shishkin meshes, we present numerical results here for Bakhvalov meshes, due to their improved accuracy.

2.2. Fitted methods for two-dimensional problems. For the two-dimensional problem (1.4), we employ the natural extension of the method (2.1): a standard finite difference scheme on a tensor-product grid. Let ω_x^N and ω_y^N be arbitrary meshes, each with N intervals on $[0, 1]$. Set $\omega^{N \times N} = \{(x_i, y_j)\}_{i,j=0}^N$ to be the Cartesian product of ω_x^N and ω_y^N . Taking $h_i = x_i - x_{i-1}$, $k_j = y_j - y_{j-1}$, $\bar{h}_i = (x_{i+1} - x_{i-1})/2$,

TABLE 2.1

Maximum pointwise errors for problem (1.5) solved by a finite difference method on a Bakhvalov mesh.

ε^2	$N = 2^7$	$N = 2^8$	$N = 2^9$	$N = 2^{10}$	$N = 2^{11}$	$N = 2^{12}$
1	9.762×10^{-5}	2.441×10^{-5}	6.103×10^{-6}	1.526×10^{-6}	3.814×10^{-7}	9.447×10^{-8}
10^{-2}	9.907×10^{-5}	2.479×10^{-5}	6.196×10^{-6}	1.549×10^{-6}	3.873×10^{-7}	9.685×10^{-8}
10^{-4}	2.858×10^{-5}	7.170×10^{-6}	1.794×10^{-6}	4.486×10^{-7}	1.122×10^{-7}	2.802×10^{-8}
10^{-6}	2.864×10^{-5}	7.188×10^{-6}	1.803×10^{-6}	4.511×10^{-7}	1.128×10^{-7}	2.818×10^{-8}
10^{-8}	2.880×10^{-5}	7.241×10^{-6}	1.814×10^{-6}	4.537×10^{-7}	1.135×10^{-7}	2.840×10^{-8}
10^{-10}	2.881×10^{-5}	7.245×10^{-6}	1.815×10^{-6}	4.543×10^{-7}	1.136×10^{-7}	2.841×10^{-8}
10^{-12}	2.881×10^{-5}	7.245×10^{-6}	1.815×10^{-6}	4.543×10^{-7}	1.137×10^{-7}	2.850×10^{-8}

and $\bar{k}_j = (y_{j+1} - y_{j-1})/2$, we define the symmetrized 5-point second-order central difference operator:

$$\Delta^N := \begin{pmatrix} & & & \frac{\bar{h}_i}{k_{j+1}} & & \\ & & & & & \\ \frac{\bar{k}_j}{h_i} & - \left(\bar{k}_j \left(\frac{1}{h_i} + \frac{1}{h_{i+1}} \right) + \bar{h}_i \left(\frac{1}{k_j} + \frac{1}{k_{j+1}} \right) \right) & & & \frac{\bar{k}_j}{h_{i+1}} & \\ & & & \frac{\bar{h}_i}{k_j} & & \end{pmatrix}.$$

The resulting numerical scheme is

$$(2.4) \quad \begin{aligned} (-\varepsilon^2 \Delta^N + \bar{h}_i \bar{k}_j b(x_i, y_j)) U_{i,j} &= \bar{h}_i \bar{k}_j f(x_i, y_j), & i = 1, \dots, N-1, j = 1, \dots, N-1, \\ U_{i,j} &= g(x_i, y_j), & i \in \{0, N\}, j \in \{0, N\}. \end{aligned}$$

Again, we write the linear system as $AU = F$.

As in one dimension, this scheme will not generate satisfactory numerical solutions if employed on a uniform mesh; cf. [18, Chap. 10]. Using a tensor-product Shishkin mesh, such as that shown in Figure 2.1(a), the error is bounded by $\mathcal{O}(N^{-2} \ln^2 N)$ [12]. This bound is sharp, and the logarithmic factor is obvious in numerical results [12, Table 1].

2.2.1. A Bakhvalov mesh for a two-dimensional problem. Take $\omega_x^N = \omega_y^N$ to be the one-dimensional graded Bakhvalov mesh described in section 2.1.1, adjusted to account for the fact that there is only one layer in each coordinate direction. Then, taking the Cartesian product mesh, $\omega^{N \times N}$, generates the two-dimensional Bakhvalov mesh, shown in Figure 2.1(b) (contrasted with the two-dimensional Shishkin mesh in Figure 2.1(a)).

Kellogg, Linss, and Stynes [24] prove the ε -uniform error estimate: there is a constant C independent of ε and N such that

$$(2.5) \quad \|u - U\|_\infty \leq CN^{-2}.$$

Table 2.1 shows the errors obtained when this method is applied to problem (1.5). Observe that, for fixed N , the error is uniformly bounded as $\varepsilon \rightarrow 0$. This demonstrates that (2.5) is indeed sharp. It also highlights the need to construct solvers that can match the optimal accuracy provided by Bakhvalov meshes.

3. Solving the one-dimensional system. From a practical viewpoint, there is no benefit to be gained in terms of computational time or memory requirements from

considering iterative approaches to solving discretizations of one-dimensional problems such as (1.1) by a standard 3-point scheme. The tridiagonal structure of the discretization matrices ensures that sparse direct solvers are optimal, and specialized approaches, such as cyclic reduction or the Thomas algorithm, are well known in the literature. It is, nonetheless, worth considering iterative approaches for these matrices to motivate the development of iterative approaches for two- (or multi-) dimensional problems. In particular, the one-dimensional boundary-layer preconditioner developed in section 3.2 develops the key ideas needed for its two-dimensional analogue in section 4.5.

3.1. Geometric multigrid. Multigrid methods are widely regarded as being among the most efficient iterative approaches for solving discretizations of elliptic PDEs, such as those considered here. In many cases, they are optimally efficient in the sense that the cost grows linearly with the number of degrees of freedom, and they are known to be highly parallelizable. The key to their efficiency lies in the combination of two processes, relaxation and coarse-grid correction, that effectively damp complementary components of the error in any approximation to the solution of the discrete system. For uniformly elliptic operators discretized on uniform meshes, simple analysis shows that standard relaxation approaches effectively damp errors that are oscillatory on the scale of the grid, while smooth error components can be resolved through relaxation on a hierarchy of coarser meshes. See [8, 49] for full details. In the context of singularly perturbed problems and nonuniform grids, more careful treatment must be given to the coarse-grid correction process to prove optimal efficiency, even for one-dimensional problems [6, 53]. In practice, however, only small changes are needed in the multigrid algorithm. In this paper, we consider iterations and preconditioners based on the multigrid V-cycle; thus, a single iteration of the multigrid cycle can be written in recursive form as follows.

ALGORITHM 1. $U^{(1)} = MG(U^{(0)}, F, N)$.

1. Apply relaxation to $AU = F$ with initial guess $U^{(0)}$, producing $U^{(r)}$.
2. Compute $F_c = R(F - AU^{(r)})$, for restriction matrix R .
3. Compute $U_c = MG(0, F_c, N/2)$.
4. Compute $U^{(c)} = U^{(r)} + PU_c$, for interpolation matrix P .
5. Apply relaxation to $AU = F$ with initial guess $U^{(c)}$, producing $U^{(1)}$.

Thus, the algorithm takes, as input, an initial guess, $U^{(0)}$, and right-hand side, F , of length $N - 1$ (in one dimension). On any level, relaxation is applied based on the matrix, A , which is taken to be the discretization of the given differential equation with $N - 1$ degrees of freedom after elimination of boundary conditions. A coarse-grid right-hand side is computed by restricting the fine-grid residual, and the algorithm is applied recursively to compute a coarse-grid representation of the error (which is the solution to $A_c U_c = F_c$), with a zero initial guess for U_c . The computed correction, U_c , is interpolated to the fine grid and added to the approximation, $U^{(r)}$, from after the initial relaxation sweep, and a second relaxation sweep is performed.

For one-dimensional problems, we use “red-black” Gauss–Seidel relaxation, where, in step 1 of Algorithm 1, the odd-indexed nodes are first processed (relative to 0-indexing in C), followed by the even-indexed nodes. The opposite ordering is used in step 5 to ensure symmetry. To create the hierarchy of meshes, we begin with a chosen fine mesh with N intervals (uniformly spaced or fitted), and create the first coarse mesh with $N/2$ intervals by aggregating intervals pairwise (or, equivalently, by discarding every other node in the fine mesh). This process is repeated recursively until a grid with fewer than 5 nodes is reached. We assume the initial N is chosen so

TABLE 3.1

CPU times for geometric multigrid preconditioned CG to achieve theoretically optimal error reduction on one-dimensional Bakhvalov meshes, aggregated over 500 runs. Corresponding iteration counts for a single run are given in parentheses.

ε^2	$N = 2^9$	$N = 2^{10}$	$N = 2^{11}$	$N = 2^{12}$	$N = 2^{13}$	$N = 2^{14}$
1	0.24 (5)	0.89 (6)	1.75 (7)	2.53 (7)	8.10 (7)	16.57 (7)
10^{-2}	0.46 (6)	0.88 (6)	2.00 (7)	4.06 (7)	8.22 (7)	18.91 (8)
10^{-4}	0.44 (6)	1.07 (7)	2.05 (7)	3.80 (7)	9.42 (8)	18.92 (8)
10^{-6}	0.54 (7)	1.08 (7)	2.03 (7)	4.58 (8)	9.25 (8)	21.21 (9)
10^{-8}	0.51 (7)	0.82 (7)	2.27 (8)	4.62 (8)	9.39 (8)	21.17 (9)
10^{-10}	0.53 (7)	1.08 (7)	2.27 (8)	4.64 (8)	10.46 (9)	21.06 (9)
10^{-12}	0.58 (7)	1.08 (7)	2.30 (8)	4.62 (8)	10.40 (9)	21.20 (9)

that this pairwise aggregation never fails. On each mesh, the matrix A is formed as in (2.1). The interpolation operator, P , between two such meshes is defined by linear interpolation to neighboring nodes, with the interpolation weights suitably adjusted for the unequal spacing of nodes. The restriction operator, R , is chosen to be the transpose of P , leading to a weighted averaging of residuals between neighboring nodes, emphasizing the closer neighbor on an unequally spaced mesh.

This approach offers excellent performance and scalability on both uniformly spaced and smoothly varying meshes, yielding errors at the level of discretization error in $\ln(N)$ iterations. For meshes with sharp contrasts in grid spacing, such as in the case of Shishkin meshes, acceleration by use of the preconditioned conjugate gradient iteration is an effective strategy, as in [52]. In this case, the V-cycle given by Algorithm 1 defines a preconditioning matrix, M , that yields a preconditioned system, $MAU = MF$, where the spectrum of MA has a tight cluster of eigenvalues around unity, with only a few outlying eigenvalues caused by the sharp transitions. Table 3.1 shows the aggregated solve times and (unaggregated) iteration counts for geometric multigrid preconditioned CG applied to Bakhvalov meshes, solving the discretized linear systems 500 times in succession, in order to achieve nontrivial timings. While the smooth variation in mesh sizes here does not preclude the use of the stationary iteration, we use PCG both for better comparison to the boundary-layer preconditioning approach proposed next and to enable use of stopping criteria similar to those discussed for the two-dimensional problems in section 4.6. We note that the solve times in Table 3.1 show the near-linear scaling in N expected of multigrid, as well as very slow growth with ε . Both of these observations are also demonstrated in the iteration counts, showing slight growth in the iteration counts as ε decreases and as N increases.

3.2. A one-dimensional boundary-layer preconditioning approach. In this section, we develop an algorithm that, in contrast to geometric multigrid, takes advantage of the singularly perturbed nature of these problems. In section 4.5, we develop the two-dimensional analogue and prove its optimality.

When $\delta_N \ll 1$, the fitted meshes for one-dimensional problems described in section 2.1.1 condense in the region of the two boundary layers and are uniform in the interior. From the linear algebraic point of view, the matrix corresponding to a discretization on such a mesh naturally partitions into three pieces: the left boundary layer, the interior (not including the transition points), and the right boundary layer. This gives a partitioned matrix of the form

$$(3.1) \quad A = \begin{bmatrix} A_{LL} & A_{LI} & 0 \\ A_{IL} & A_{II} & A_{IR} \\ 0 & A_{RI} & A_{RR} \end{bmatrix},$$

where we use the natural subscripts $L, I,$ and R to denote the three regions. Notably, in the symmetrically scaled system, we have the relations that $A_{LI} = A_{IL}^T$ and $A_{RI} = A_{IR}^T$, and we see that each of these matrices has only one nonzero entry, giving the coefficient that relates the last point in the layer (the transition point) to the first point in the interior. Defining the uniform interior mesh width to be h_I , these entries are all $-\varepsilon^2/h_I$.

Considering A_{II} , we notice that this matrix has diagonal entries given by $2\varepsilon^2/h_I + h_I b(x_j)$, for the row corresponding to node x_j , while the off-diagonal coefficients are $-\varepsilon^2/h_I$. In the singularly perturbed case, when $\delta_h = (\varepsilon/(h_I\beta))^2 \ll 1$, the diagonal of A_{II} strongly dominates all other entries in the rows and columns corresponding to interior (I) points. In this case, it is intuitive to approximate these rows and columns by just the dominant diagonal values given by the reaction term, $h_I b(x_j)$. Theorem 3.1 states that this approximation is accurate, in the spectral sense, when $\delta_h \ll 1$.

THEOREM 3.1. *Let h_I denote the (uniform) interior mesh spacing in the standard 3-point finite difference discretization of $-\varepsilon^2 u''(x) + b(x)u(x)$ on a boundary-fitted mesh, where $b(x) > \beta^2$ for all $x \in [0, 1]$, with symmetrized discretization matrix A . Order the rows and columns of A according to the boundary layer structure in (3.1). Define*

$$A_D = \begin{bmatrix} A_{LL} & 0 & 0 \\ 0 & D_{II} & 0 \\ 0 & 0 & A_{RR} \end{bmatrix},$$

where the partitioning matches that of A in (3.1), and the entries of diagonal matrix D_{II} are given by $h_I b(x_j)$ for the row/column corresponding to node x_j . Then

$$(1 - 2\delta_h)V^T A_D V \leq V^T A V \leq (1 + 6\delta_h)V^T A_D V$$

for all vectors V .

Proof. The proof uses similar techniques to those described in detail in the proof of Theorem 4.1. For full details, see [32]. \square

COROLLARY 3.2. *Under the assumptions of Theorem 3.1, if M_{LL} and M_{RR} are spectrally equivalent to A_{LL} and A_{RR} , respectively, meaning that there are constants c_0 and c_1 such that*

$$c_0 V_L^T M_{LL} V_L \leq V_L^T A_{LL} V_L \leq c_1 V_L^T M_{LL} V_L \text{ for all } V_L$$

and $c_0 V_R^T M_{RR} V_R \leq V_R^T A_{RR} V_R \leq c_1 V_R^T M_{RR} V_R$ for all V_R ,

then the matrix

$$A_M = \begin{bmatrix} M_{LL} & 0 & 0 \\ 0 & D_{II} & 0 \\ 0 & 0 & M_{RR} \end{bmatrix}$$

satisfies

$$\min(1 - 2\delta_h, c_0(1 - 2\delta_h))V^T A_M V \leq V^T A V \leq \max(1 + 6\delta_h, c_1(1 + 2\delta_h))V^T A_M V$$

for all V .

TABLE 3.2

Computed values of $\delta_h = \varepsilon^2/(h_1^2\beta^2)$ for Bakhvalov meshes on $[0, 1]$ with two boundary layers and $\beta = 0.99$.

ε^2	$N = 2^9$	$N = 2^{10}$	$N = 2^{11}$	$N = 2^{12}$	$N = 2^{13}$	$N = 2^{14}$
1	2.65×10^5	1.06×10^6	4.24×10^6	1.70×10^7	6.78×10^7	2.71×10^8
10^{-2}	2.65×10^3	1.06×10^4	4.24×10^4	1.70×10^5	6.78×10^5	2.71×10^6
10^{-4}	1.21×10^1	4.83×10^1	1.93×10^2	7.73×10^2	3.09×10^3	1.24×10^4
10^{-6}	7.28×10^{-2}	2.91×10^{-1}	1.17×10^0	4.66×10^0	1.86×10^1	7.46×10^1
10^{-8}	6.71×10^{-4}	2.68×10^{-3}	1.07×10^{-2}	4.29×10^{-2}	1.72×10^{-1}	6.87×10^{-1}
10^{-10}	6.63×10^{-6}	2.65×10^{-5}	1.06×10^{-4}	4.24×10^{-4}	1.70×10^{-3}	6.79×10^{-3}
10^{-12}	6.62×10^{-8}	2.65×10^{-7}	1.06×10^{-6}	4.24×10^{-6}	1.70×10^{-5}	6.78×10^{-5}

TABLE 3.3

CPU times for boundary layer preconditioned CG to achieve theoretically optimal error reduction on one-dimensional Bakhvalov meshes, aggregated over 500 runs. Corresponding iteration counts for a single run are given in parentheses.

ε^2	$N = 2^9$	$N = 2^{10}$	$N = 2^{11}$	$N = 2^{12}$	$N = 2^{13}$	$N = 2^{14}$
10^{-6}	0.35 (8)					
10^{-8}	0.24 (7)	0.53 (7)	1.34 (8)	3.04 (9)		
10^{-10}	0.26 (7)	0.66 (7)	1.28 (8)	2.61 (8)	6.10 (9)	12.0 (9)
10^{-12}	0.33 (7)	0.55 (7)	1.33 (8)	2.58 (8)	5.97 (9)	12.2 (9)

Corollary 3.2 is particularly relevant to the case where M_{LL} and M_{RR} are the inverses of effective multigrid preconditioners for A_{LL} and A_{RR} , respectively. If the stationary multigrid cycle represented by $I - M_{LL}^{-1}A_{LL}$ has spectral radius $\alpha < 1$, then

$$(1 - \alpha)V_L^T M_{LL} V_L \leq V_L^T A_{LL} V_L \leq (1 + \alpha)V_L^T M_{LL} V_L$$

for all V_L . If the same bounds hold for the right boundary layer, then the bound in Corollary 3.2 becomes

$$(1 - \alpha)(1 - 2\delta_h)V^T A_M V \leq V^T A V \leq \max(1 + 6\delta_h, (1 + \alpha)(1 + 2\delta_h))V^T A_M V$$

for all V . Since we typically expect $\alpha \approx 0.1$ and $\delta_h \ll 1$, a better rule of thumb would be

$$(1 - \alpha)(1 - 2\delta_h)V^T A_M V \leq V^T A V \leq (1 + \alpha)(1 + 2\delta_h)V^T A_M V$$

for all V . This suggests that, when $\delta_h \ll 1$, the convergence of a preconditioner consisting of one multigrid V-cycle applied to each of the boundary-layer regions plus diagonal scaling of the interior region should be very similar to that of multigrid applied directly to the boundary-layer regions alone. Table 3.2 shows the values of δ_h realized on the one-dimensional Bakhvalov meshes; following Theorem 3.1 and Corollary 3.2, we present results for the cases where $\delta_h < 0.1$, with CPU times aggregated over 500 runs, along with (unaggregated) iteration counts, in Table 3.3.

Comparing times between Tables 3.1 and 3.3, we see uniform improvements by factors of just under 2 for the boundary-layer preconditioning approach over standard geometric multigrid preconditioning. This is, in general, consistent with our cost expectation, since the boundary-layer approach is applying a regular geometric multigrid cycle (adapted only to account for the implied boundary condition at the transition points) to one-half of the grid, and a much cheaper diagonal scaling operation to the other half. We also see that the iteration counts in these tables are

essentially identical, with the boundary-layer preconditioner occasionally taking one more iteration to fulfill the stopping criterion in cases corresponding to larger δ_h .

4. Solving the two-dimensional system. In contrast to the tridiagonal matrices that arise with the one-dimensional problem, numerical solution of the two-dimensional discretizations poses more of a challenge. Despite the well-known complexity growth for matrices with these structures, the topic of efficient and robust solution algorithms for these problems is largely ignored in the literature of singular perturbation problems: we know only of the studies mentioned in the introduction. Most papers which report numerical results don't mention the solver used. Of the few that do, such as [29], a direct solver is employed. Thus, we begin with an explanation of why sparse direct solvers are unsuitable from two points of view, before investigating efficient iterative methods.

4.1. Direct solvers. As a "typical" serial, highly optimized sparse direct solver, we consider here CHOLMOD (supernodal sparse Cholesky factorization and update/down-date) Version 1.7.1 to solve the sparse symmetric linear systems; see [11, 13]. In Table 4.1, we show the time in seconds, averaged over three runs, required to solve the linear systems on Bakhvalov meshes that yield results given in Table 2.1. For a given ε , we see growth in these times that, for large N , scales as N^3 (increasing by factors of roughly eight when N is doubled), as expected given the known $\mathcal{O}(N^3)$ complexity of the nested dissection algorithm for these grids [22] and the $\mathcal{O}(N^3)$ lower bound for this complexity [23]. For fixed N , however, we observe that, rather remarkably, the amount of time required to solve the linear system depends quite badly on the perturbation parameter. This is in spite of the fact that, for a given N , the matrices for different ε have exactly the same size and structure; the only difference is the scaling of some entries due to the difference in both ε and the local mesh width. Similar results have been observed with other direct solvers, including MA57 [17], on different processors and architectures, and with different meshes, including uniform and Shishkin meshes. As we now explain, the degradation in the performance of the direct solvers is not related to their implementation but, rather, the specific nature of the discretized singularly perturbed problem.

Writing the linear system for the finite difference solution to (2.4) as $AU = F$, while neglecting boundary conditions, gives A as an $(N - 1)^2 \times (N - 1)^2$, banded, symmetric, and positive-definite matrix with bandwidth of $N - 1$. For problems with small off-diagonal entries, relative to the diagonal in each row, fill-in introduced during factorization results in many small values in the LU factors. Such values decay exponentially as factorization proceeds introducing successive fill-ins. If the matrix bandwidth is small, then this poses no problem for floating-point calculations. If, on the other hand, N is large, then floating point underflows may occur.

TABLE 4.1

Cholesky (CHOLMOD) solve times for linear systems generated by a finite difference method applied on a Bakhvalov mesh.

ε^2	$N = 2^7$	$N = 2^8$	$N = 2^9$	$N = 2^{10}$	$N = 2^{11}$	$N = 2^{12}$
1	0.07	0.39	2.65	18.29	195.87	1680.57
10^{-2}	0.06	0.38	2.66	18.27	196.18	1678.79
10^{-4}	0.06	0.38	2.66	18.39	196.23	1689.43
10^{-6}	0.07	0.97	11.83	89.03	860.62	7515.59
10^{-8}	0.15	1.25	10.62	71.40	478.32	2676.85
10^{-10}	0.19	1.16	8.34	46.22	343.66	1521.52
10^{-12}	0.18	1.10	6.72	36.11	257.12	1166.78

TABLE 4.2

Number of nonzero entries (top) and subnormal numbers (bottom) in Cholesky factors generated by CHOLMOD.

ε^2	$N = 2^7$	$N = 2^8$	$N = 2^9$	$N = 2^{10}$	$N = 2^{11}$	$N = 2^{12}$
1	350,112	1,833,813	9,425,559	45,671,436	183,759,251	831,532,333
	0	0	0	0	0	0
10^{-2}	350,112	1,833,813	9,425,559	45,671,436	183,759,251	831,532,333
	0	0	0	0	0	0
10^{-4}	350,112	1,833,813	9,425,559	45,671,436	183,759,251	831,532,333
	0	0	0	0	0	0
10^{-6}	350,112	1,828,215	9,293,727	44,499,256	179,511,201	808,690,367
	0	4,338	22,596	108,387	573,033	2,852,019
10^{-8}	347,351	1,717,341	8,266,871	37,946,547	147,162,291	625,420,613
	1,146	8,488	56,295	316,104	1,121,348	4,956,624
10^{-10}	335,322	1,614,213	7,535,505	33,695,760	130,437,185	544,870,886
	1,915	10,008	77,691	283,348	1,111,292	4,422,916
10^{-12}	322,935	1,534,747	7,019,889	31,076,314	120,736,814	504,478,967
	2,176	11,467	58,065	305,428	991,728	3,803,770

In IEEE standard double precision (cf. [37]), numbers are typically represented with 64 bits as $\pm X \times 2^{Y-1023}$, where 52 bits are used to store the significand, X , 11 bits are used to store the exponent, Y , and the remaining bit stores the sign of X . For “normal” numbers, with $0 < Y < 2047$, X is assumed to be a binary decimal with leading digit 1 (an implied 53rd bit). Thus, the smallest normal number occurs for $Y = 1$, $X = 0$, giving $2^{-1022} \approx 10^{-308}$. When $Y = 0$, the implied bit in X is taken to be a leading 0 and the exponent is fixed at -1022 , allowing representation of nonzero numbers as small as is $2^{-52} \times 2^{-1022} \approx 5 \times 10^{-324}$; anything smaller than this is rounded to zero (when $X = 0$ and $Y = 0$). The use of such “subnormal” numbers allows for gradual reduction in the precision of stored numbers; however, most processors do not provide hardware support for arithmetic with subnormal numbers and, instead, a compiler must rely on a software implementation, which is significantly slower [27]. The variation in timings seen with ε in Table 4.1 are due to the introduction of subnormal numbers in the fill-in generated by Cholesky.

To demonstrate this effect, in Table 4.2, we give the number of nonzero entries in the Cholesky factors produced by CHOLMOD for a range of values of N and ε , corresponding to those shown in Table 4.1, as well as the number of subnormal entries. For small ε and large N , we observe a significant increase in the number of subnormal numbers arising in the Cholesky factors, as well as a decrease in the number of nonzero numbers, due to underflow of subnormal entries. Although they are relatively few compared to the number of nonzero entries, they are sufficient to greatly increase the computation time. Such difficulties are likely to be exacerbated further for problems in three (or more) dimensions, where larger bandwidths occur.

It is worth noting that this behavior can be overcome with compiler directives. For the GCC family of compilers, the `-funsafe-math-optimizations` option turns on certain optimizations that violate the IEEE floating point standard, including the treatment of subnormal numbers. With this option enabled, we see general decreases in the factorization and solve times as ε decreases and the number of nonzero entries retained in the factors drops, with no observable degradation in the solution accuracy. However, we now see large (up to 60%) increases in these times for the smallest values of ε that are not consistent with the number of nonzeros in the factors. This suggests that simply ignoring subnormal entries in the factors may still not be enough to ensure a favorable scaling of direct solvers with ε , especially as doing so relies on compiler and architecture dependent implementations of variations to the IEEE standard.

TABLE 4.3

Iteration counts for geometric multigrid preconditioned CG to achieve theoretically optimal error reduction on Bakhvalov meshes.

ε^2	$N = 2^7$	$N = 2^8$	$N = 2^9$	$N = 2^{10}$
1	11	11	12	12
10^{-2}	14	16	18	19
10^{-4}	60	96	140	178
10^{-6}	76	152	300	557
10^{-8}	81	170	348	712
10^{-10}	88	179	364	748
10^{-12}	92	186	387	811

4.2. Geometric multigrid. The success of simple geometric multigrid with pointwise relaxation as a good preconditioner for the one-dimensional problem, with only mild dependence on N and ε , is not recreated in two dimensions, due to the nature of the tensor product of fitted meshes. As noted in [54, 55], the key parameter in determining the effectiveness of pointwise relaxation in geometric multigrid is the mesh aspect ratio; when the ratio of the largest to smallest edges of a mesh cell is much larger than unity, geometric multigrid convergence suffers, even when used as a preconditioner. Table 4.3 shows the number of iterations needed for geometric multigrid preconditioned CG to reduce the maximum norm of the error in the iterative solution to that of the direct solution of the linear system.

Two well-known antidotes to this degradation are to consider the use of mesh-adapted coarsening approaches or the use of linewise (coupled) relaxation techniques [8, 49]. In the former category, semicoarsening techniques are commonly used when the bad aspect ratios result from the use of tensor-product grids of vastly different mesh widths (as, for example, might be the case for a two-dimensional problem with a one-dimensional boundary layer). These techniques can be adapted to tensor products of graded or fitted meshes [54], but require much more involved development of the semicoarsening approach. Instead, in section 4.3, we consider an algebraic multigrid method that automatically adapts multigrid coarsening to the fitted mesh structure. In contrast, alternating-direction line relaxation techniques can be used to address the challenges posed by the tensor product of fitted meshes, with bad aspect ratios in two orientations. Within this class of approaches, we use the black box multigrid (BoxMG) method of Dendy [1, 15], coupling alternating-direction line relaxation with variational coarsening, discussed in section 4.4.

There are three natural advantages to the use of algebraic multigrid (AMG) and BoxMG. First, they both employ variational coarsening strategies based on Galerkin projection, so that variable reaction coefficients are naturally projected onto coarse meshes. Second, the same solvers can be easily applied to other discretizations, such as bilinear finite element discretizations on fitted meshes (see, e.g., [29] for an error analysis for a two-dimensional reaction diffusion problem on a Shishkin mesh). Finally, they are both readily available in software packages that can easily be incorporated into simulation codes beyond those used in the test cases presented in section 4.7. The downside of using these solvers is that they do not represent tuned software that takes full advantage of the structure of the differential operators, discretization, and fitted meshes considered here. Thus, small improvements in overall time to solution could be realized by more specialized software, but at the expense of applicability and ease of use. In contrast, in section 4.5, we present an algorithm that is tuned specifically for the fitted meshes considered here, but which is still applicable to both finite difference and finite element discretizations of singularly perturbed reaction-diffusion equations.

4.3. Algebraic multigrid. An alternate approach to using geometric multigrid is to adjust the interpolation operators to account for the sharp mesh transitions, thereby accurately approximating the associated error modes in the coarse-grid correction phase. A standard approach to developing operator-dependent interpolation schemes is the AMG method, which supplements the multigrid V-cycle given in Algorithm 1 with a preliminary setup stage in which the coarse meshes and interpolation, restriction, and coarse-grid operators are computed based on the given fine-grid operator.

Here, we use a standard AMG approach, as described in [8,43,48]. From the given fine-grid operator, the setup stage chooses a first coarse mesh, then interpolation, P , is defined, with restriction fixed as $R = P^T$, and the coarse-grid operator is given by the Galerkin triple product $A_c = P^T A P$. This process repeats recursively until a suitably small mesh is reached, where a direct solver can be cheaply applied. Beyond the setup stage, only a small change is made in the solve phase detailed in Algorithm 1, where the size of the coarse-grid problem is now determined by the number of coarse-grid points selected, changing the notation, but not the substance, of the recursive step. For relaxation, we use an ordered Gauss–Seidel sweep that, in step 1 of Algorithm 1, loops lexicographically over the points selected for the coarse mesh first, then over those that are not selected for the coarse mesh (known as C/F relaxation). In step 5 of Algorithm 1, the opposite ordering is used, to preserve the symmetry of the preconditioner.

The coarse-grid selection and interpolation algorithms of [43] are used here; for details, see [8,43]. Because the singularly perturbed nature of the problems considered here leads to matrices that have some rows that are strongly diagonally dominant, we adopt a slight variation to the standard approach, whereby rows that are strongly diagonally dominant (taken to be when the sum of the absolute values of the off-diagonals is less than 1/2 of the diagonal entry) are treated through fine-grid relaxation alone.

4.4. Black box multigrid. In contrast to AMG, the BoxMG algorithm [1,15] focuses on maintaining the regular, tensor-product grid structure of the fine mesh, and choosing more robust relaxation techniques, in combination with a similar operator-dependent definition of interpolation, to achieve robustness in the multigrid algorithm. The algorithmic choice to maintain structured grids on all levels of the multigrid hierarchy avoids the expensive graph processing and indirect addressing of AMG; as a result, BoxMG often achieves much faster total times to solution for problems where both AMG and BoxMG can be readily applied, about six times faster for some two-dimensional problems [10], and ten to fifteen times faster for some three-dimensional problems [31].

The key to maintaining robustness in BoxMG for two-dimensional problems is the use of alternating-direction line relaxation. The rectangular grid is decomposed twice, into alternating lines in both the x - and y -directions. In a single sweep of relaxation (step 1 of Algorithm 1), residuals are calculated first for alternating lines parallel to the x -axis, starting from the second row of the grid, and an update is calculated for each of these lines to simultaneously zero the residual at all points along these lines, then the same procedure is applied to alternating lines parallel to the x -axis, starting from the first row of the grid, then these two stages are repeated for alternating lines parallel to the y -axis. The opposite ordering is used in step 5 of Algorithm 1 to ensure symmetry. Interpolation is determined based on averaging of the operator coefficients along the structured mesh, achieving robustness without the indirect addressing of

AMG. The full specification of BoxMG interpolation is reviewed in [30], including a discussion of its relationship to AMG interpolation.

4.5. A two-dimensional boundary-layer preconditioning approach. For the two-dimensional problem with $\delta_N \ll 1$, we extend the preconditioning technique from section 3.2 to handle tensor products of fitted meshes, by partitioning the two-dimensional mesh into four pieces: the high-resolution corner, two edges with resolved meshes in one dimension (along the x - and y -axes), and the interior region. In the corner, connections in both coordinate directions are significant, requiring the full power of a multigrid method. In the edges, the problems are effectively one dimensional, and tridiagonal solvers can be used within an effective preconditioner. In the interior, as in one dimension, diagonal scaling is effective. Thus, the two-dimensional boundary-layer preconditioner can be expressed by partitioning the degrees of freedom in U and, consequently, the rows and columns of A as

$$(4.1) \quad A = \begin{bmatrix} A_{CC} & A_{CE} & 0 \\ A_{EC} & A_{EE} & A_{EI} \\ 0 & A_{IE} & A_{II} \end{bmatrix},$$

where the subscripts indicate the block structure of corners, C , edge layers, E , and interior points, I . The preconditioner, A_D , is defined in the same partitioning by

$$(4.2) \quad A_D = \begin{bmatrix} A_{CC} & 0 & 0 \\ 0 & T_{EE} & 0 \\ 0 & 0 & D_{II} \end{bmatrix}.$$

Here, matrix D_{II} has diagonal entries given by the scaled reaction coefficient. The matrix T_{EE} has three nonzero entries per row. In rows corresponding to nodes in the edge layer along the x -axis (where the mesh in the x -direction is uniformly spaced with width h_I), the diagonal entry is $h_I(1/k_j + 1/k_{j+1})\varepsilon^2 + h_I \bar{k}_j b(x_i, y_j)$. The two off-diagonal entries correspond to neighbors along the direction with the smallest mesh width, so their values are $-(h_I/k_j)\varepsilon^2$ and $-(h_I/k_{j+1})\varepsilon^2$. The entries of T_{EE} in rows corresponding to nodes in the edge layer along the y -axis are similar. Thus, T_{EE} could be itself partitioned into two parts, one corresponding to the edge layer along the x -axis, and one corresponding to the edge layer along the y -axis. Depending on the ordering of the degrees of freedom (in x then y , or in y then x), one of these blocks will be tridiagonal, and one will be block tridiagonal with diagonal blocks. The neglected entries in A_{II} and A_{IE} or A_{EI} are essentially the same as those treated in the one-dimensional case, but scaled by another factor of h_I from the symmetrization, giving $-\varepsilon$, relative to the diagonal values $h_I^2 b(x_i, y_j)$. We bound these using simple linear-algebraic arguments, similar to those used to prove Theorem 3.1 [32]. The neglected entries in A_{EE} , A_{CE} , and A_{EC} are slightly more complicated, due to the potentially uneven mesh spacing in the boundary layers, but similar linear-algebraic arguments bound their contribution.

THEOREM 4.1. *Let h_I denote the interior mesh spacing in the standard 5-point finite difference discretization of $-\varepsilon^2 \Delta u(x, y) + b(x, y)u(x, y)$ on a boundary-fitted mesh, where $b(x, y) > \beta^2$ for all $(x, y) \in [0, 1]^2$ and boundary conditions such that there are only two boundary layers, along the x - and y -axes (the south and west faces, respectively), with symmetrized discretization matrix A . Order the rows and columns of A according to the boundary-layer structure in (4.1) and define A_D in the same ordering as in (4.2), where the entries in the diagonal matrix D_{II} are given by the*

reaction coefficients, $h_I^2 b(x_i, y_j)$, for the row/column corresponding to node (x_i, y_j) and the off-diagonal coefficients of $A_{EE} - T_{EE}$ are of the form $-\varepsilon^2 \bar{k}_j / h_I$ (for edges along the x -axis) or $-\varepsilon^2 \bar{h}_j / h_I$ (for edges along the y -axis), with twice the magnitude of these values on the diagonal (so that $A_{EE} - T_{EE}$ is a zero row-sum operator except for the first and last rows). Then,

$$(1 - 3\delta_h)V^T A_D V \leq V^T A V \leq (1 + 9\delta_h)V^T A_D V$$

for all vectors V .

Proof. Writing generic vector $V = [V_C \quad V_E \quad V_I]$ (noting that here V_C is the component of V associated with the refined corner of the mesh, and should not be confused with the notation of sections 3.1 and 4.3 where V_c would denote the coarse-grid analogue of V) we see that

$$V^T A V = V_C^T A_{CC} V_C + 2V_C^T A_{CE} V_E + V_E^T A_{EE} V_E + 2V_E^T A_{EI} V_I + V_I^T A_{II} V_I,$$

using the symmetry of A . We seek to bound this above and below by a scalar (depending only on δ_h) times the quantity

$$V^T A_D V = V_C^T A_{CC} V_C + V_E^T T_{EE} V_E + V_I^T D_{II} V_I.$$

Bounding $V_I^T A_{II} V_I$ is straightforward, since $A_{II} = D_{II} + L_{II}$, where L_{II} is the symmetrized uniform-grid Laplacian matrix with diagonal entry $4\varepsilon^2$ and off-diagonal entries $-\varepsilon^2$. Thus, by Geršgorin's theorem,

$$V_I^T D_{II} V_I \leq V_I^T A_{II} V_I \leq V_I^T (D_{II} + (8\varepsilon^2)I) V_I \leq (1 + 8\delta_h)V_I^T D_{II} V_I$$

for any V_I , since $V_I^T V_I \leq 1/(\beta^2 h_I^2)V_I^T D_{II} V_I$.

Considering $V_E^T A_{EE} V_E$ is slightly more complicated because of the tridiagonal structure of $A_{EE} - T_{EE}$, we can, however, consider the spectral equivalence by analyzing a generic line in the edges of the mesh. To do this, first note that A_{EE} is reducible into two blocks, one containing connections within the layer along the x -axis, and one containing connections within the layer along the y -axis. Considering one of these pieces, along the x -axis, we decompose the stencil in (2.4) into two pieces, writing

$$\begin{aligned} & \begin{pmatrix} & & -\frac{h_I}{k_{j+1}}\varepsilon^2 & & \\ -\frac{\bar{k}_j}{h_I}\varepsilon^2 & \left(\frac{2\bar{k}_j}{h_I} + h_I\left(\frac{1}{k_j} + \frac{1}{k_{j+1}}\right)\right)\varepsilon^2 + \bar{k}_j h_I b(x_i, y_j) & & & -\frac{\bar{k}_j}{h_I}\varepsilon^2 \\ & & -\frac{h_I}{k_j}\varepsilon^2 & & \\ & & & & \end{pmatrix} \\ &= \begin{pmatrix} & & -\frac{h_I}{k_{j+1}}\varepsilon^2 & & \\ 0 & h_I\left(\frac{1}{k_j} + \frac{1}{k_{j+1}}\right)\varepsilon^2 + \bar{k}_j h_I b(x_i, y_j) & & & 0 \\ & & -\frac{h_I}{k_j}\varepsilon^2 & & \end{pmatrix} + \begin{pmatrix} & 0 & & & \\ -\frac{\bar{k}_j}{h_I}\varepsilon^2 & \frac{2\bar{k}_j}{h_I}\varepsilon^2 & & & -\frac{\bar{k}_j}{h_I}\varepsilon^2 \\ & & 0 & & \end{pmatrix}, \end{aligned}$$

where the first term on the right-hand side represents the terms kept in T_{EE} , and the second term represents $A_{EE} - T_{EE}$. From this decomposition, we immediately

see that $V_E^T T_{EE} V_E \leq V_E^T A_{EE} V_E$ for any V_E , since the remainder term is positive definite. To find the upper bound, note that the entries in $A_{EE} - T_{EE}$ depend only on the position in the y -direction, through \bar{k}_j , and not on the position in the x -direction. Thus, defining the diagonal matrix M , with dimension given by the number of points within the boundary layer and entries given by $(\bar{k}_j/h_I)\varepsilon^2$, for the row/column associated with the point j nodes in from the boundary, we can write

$$(A_{EE} - T_{EE})_x = \begin{bmatrix} 2M & -M & & & & & \\ -M & 2M & -M & & & & \\ & -M & 2M & -M & & & \\ & & & \ddots & \ddots & \ddots & \\ & & & & -M & 2M & -M \\ & & & & & -M & 2M \end{bmatrix},$$

assuming an ordering of points first in lines parallel to the y -axis, then in lines parallel to the x -axis, where the subscript x is used on $A_{EE} - T_{EE}$ to indicate that we consider only the edge layer along the x -axis. The matrix $(A_{EE} - T_{EE})_x$ has block dimension equal to the number of points on the grid that are *not* in the boundary layer, and the quantity $(V_E)_x^T (A_{EE} - T_{EE})_x (V_E)_x$ can easily be expressed by writing $(V_E)_x = [V_1 \ V_2 \ \dots]$, giving

$$(V_E)_x^T (A_{EE} - T_{EE})_x (V_E)_x = 2 \sum_{\ell} V_{\ell}^T M V_{\ell} - 2 \sum_{\ell} V_{\ell}^T M V_{\ell+1},$$

where the first sum extends over all ℓ , and the second over all but the last ℓ . Bounding the terms in the second sum, using the Cauchy–Schwarz inequality, gives

$$|V_{\ell}^T M V_{\ell+1}| \leq \|M^{1/2} V_{\ell}\|_2 \|M^{1/2} V_{\ell+1}\|_2 \leq \frac{1}{2} (V_{\ell}^T M V_{\ell} + V_{\ell+1}^T M V_{\ell+1});$$

this, in turn, bounds

$$|(V_E)_x^T (A_{EE} - T_{EE})_x (V_E)_x| \leq 2 \sum_{\ell} V_{\ell}^T M V_{\ell} + 2 \sum_{\ell} |V_{\ell}^T M V_{\ell+1}| \leq 4 \sum_{\ell} V_{\ell}^T M V_{\ell}.$$

For one of these terms, however, the diagonal entries in M , $\bar{k}_j/h_I\varepsilon^2$, are bounded above by $\delta_h \bar{k}_j h_I b(x_i, y_j)$ for any i . Thus,

$$V_{\ell}^T M V_{\ell} \leq \delta_h V_{\ell}^T D_{\ell} V_{\ell}$$

for all V_{ℓ} for any ℓ , where the diagonal matrices D_{ℓ} are taken to have entries $\bar{k}_j h_I b(x_{i_{\ell}}, y_j)$, where index i_{ℓ} corresponds to the line indexed by ℓ . Taken together, with the analogous bounds for the edge layer along the y -axis, this gives

$$V_E^T (A_{EE} - T_{EE}) V_E \leq 4\delta_h V_E^T D_{EE} V_E$$

for any V_E , where D_{EE} is the diagonal matrix with entries $\bar{k}_j h_I b(x_i, y_j)$ for nodes (i, j) in the edge layer along the x -axis and analogous values for nodes in the edge layer along the y -axis. Since $V_E^T D_{EE} V_E \leq V_E^T T_{EE} V_E$ for any V_E , which follows from the fact that the difference $T_{EE} - D_{EE}$ is positive definite, this gives

$$V_E^T T_{EE} V_E \leq V_E^T A_{EE} V_E \leq V_E^T T_{EE} V_E + 4\delta_h V_E^T D_{EE} V_E \leq (1 + 4\delta_h) V_E^T T_{EE} V_E.$$

We next look to bound the mixed terms, $2V_C^T A_{CE} V_E$ and $2V_E^T A_{EI} V_I$. Considering the first of these, we again employ the Cauchy–Schwarz inequality, writing

$$|V_C^T A_{CE} V_E| \leq \left\| D_{CC}^{-1/2} A_{CE} V_E \right\|_2 \cdot \left\| D_{CC}^{1/2} V_C \right\|_2$$

for any V_C and V_E , where D_{CC} is the diagonal matrix with entry $\bar{h}_i \bar{k}_j b(x_i, y_j)$ in the row/column corresponding to node (i, j) . Since $A_{CC} = D_{CC} + L_{CC}$ for positive-definite Laplacian matrix L_{CC} , $V_C^T D_{CC} V_C \leq V_C^T A_{CC} V_C$ and

$$|V_C^T A_{CE} V_E| \leq \left\| D_{CC}^{-1/2} A_{CE} V_E \right\|_2 \cdot (V_C^T A_{CC} V_C)^{1/2}.$$

To bound the remaining term, notice that there is at most one nonzero entry per row in A_{CE} and, consequently, in $D_{CC}^{1/2} A_{CE}$ and $A_{EC} D_{CC}^{-1} A_{CE}$. The nonzero entries in $A_{EC} D_{CC}^{-1} A_{CE}$ occur on the diagonals of the rows/columns corresponding to nodes that are adjacent to the highly resolved corner, and are given by $((\bar{k}_j/h_I)\varepsilon^2)/(\bar{h}_i \bar{k}_j b(x_i, y_j))$ for a node $(i+1, j)$ in the edge layer along the x -axis (so that node (i, j) is in the corner region), with a similar expression for a node in the edge layer along the y -axis. Since $x_{i+1} - x_i = h_I$, $\bar{h}_i \geq h_I/2$, giving

$$\begin{aligned} ((\bar{k}_j/h_I)\varepsilon^2)^2 / \bar{h}_i \bar{k}_j b(x_i, y_j) &\leq (\varepsilon^2/h_I)^2 \frac{2}{h_I \beta^2} \bar{k}_j \\ &\leq (\varepsilon^2/h_I)^2 \frac{2}{h_I^2 \beta^4} h_I \bar{k}_j b(x_{i+1}, y_j) = 2\delta_h^2 h_I \bar{k}_j b(x_{i+1}, y_j) \end{aligned}$$

with an analogous bound for nodes in the edge layer along the y -axis. Thus, $V_E^T A_{EC} D_{CC}^{-1} A_{CE} V_E \leq 2\delta_h^2 V_E^T D_{EE} V_E$ for all V_E and, thus, $V_E^T A_{EC} D_{CC}^{-1} A_{CE} V_E \leq 2\delta_h^2 V_E^T T_{EE} V_E$ for all V_E . This gives

$$|V_C^T A_{CE} V_E| \leq \sqrt{2}\delta_h (V_E^T T_{EE} V_E)^{1/2} (V_C^T A_{CC} V_C)^{1/2}$$

for all V_C and V_E . Since $2ab \leq a^2 + b^2$, this gives

$$(4.3) \quad 2|V_C^T A_{CE} V_E| \leq \delta_h V_E^T T_{EE} V_E + 2\delta_h V_C^T A_{CC} V_C$$

for all V_C and V_E .

Finally, we look to bound $2V_E^T A_{EI} V_I$ using similar techniques. The Cauchy–Schwarz bound becomes

$$|V_E^T A_{EI} V_I| \leq \left\| D_{EE}^{-1/2} A_{EI} V_I \right\|_2 \cdot \left\| D_{EE}^{1/2} V_E \right\|_2$$

for all V_E and V_I , giving

$$|V_E^T A_{EI} V_I| \leq \left\| D_{EE}^{-1/2} A_{EI} V_I \right\|_2 \cdot (V_E^T T_{EE} V_E)^{1/2}$$

for all V_E and V_I , since $V_E^T D_{EE} V_E \leq V_E^T T_{EE} V_E$ for all V_E . Again, there is only one entry per row of $A_{IE} D_{EE}^{-1} A_{EI}$, on the diagonal of rows/columns adjacent to the edge layers. For a node $(i, j+1)$ adjacent to the edge layer along the x -axis, the entry would be

$$\varepsilon^4 / (h_I \bar{k}_j b(x_i, y_j)) \leq \varepsilon^4 \frac{2}{h_I^2 \beta^2} \leq 2\delta_h^2 h_I^2 b(x_i, y_{j+1}),$$

where we use the bound $\bar{k}_j \geq h_I/2$, since $y_{j+1} - y_j = h_I$. This gives

$$|V_E^T A_{EI} V_I| \leq 2\delta_h^2 (V_I^T D_{II} V_I)^{1/2} (V_E^T T_{EE} V_E)^{1/2}$$

for all V_E and V_I . Consequently,

$$(4.4) \quad 2|V_E^T A_{EI} V_I| \leq \delta_h V_I^T D_{II} V_I + 2\delta_h V_E^T T_{EE} V_E$$

for all V_E and V_I .

Assembling these bounds, we have

$$\begin{aligned} V^T AV &= V_C^T A_{CC} V_C + 2V_C^T A_{CE} V_E + V_E^T A_{EE} V_E + 2V_E^T A_{EI} V_I + V_I^T A_{II} V_I \\ &\geq V_C^T A_{CC} V_C - 2|V_C^T A_{CE} V_E| + V_E^T T_{EE} V_E - 2|V_E^T A_{EI} V_I| + V_I^T D_{II} V_I \\ &\geq V_C^T A_{CC} V_C - (\delta_h V_E^T T_{EE} V_E + 2\delta_h V_C^T A_{CC} V_C) + V_E^T T_{EE} V_E \\ &\quad - (\delta_h V_I^T D_{II} V_I + 2\delta_h V_E^T T_{EE} V_E) + V_I^T D_{II} V_I \\ &= (1 - 2\delta_h) V_C^T A_{CC} V_C + (1 - 3\delta_h) V_E^T T_{EE} V_E + (1 - \delta_h) V_I^T D_{II} V_I \end{aligned}$$

for all V , establishing the lower bound. For the upper bound, we have

$$\begin{aligned} V^T AV &= V_C^T A_{CC} V_C + 2V_C^T A_{CE} V_E + V_E^T A_{EE} V_E + 2V_E^T A_{EI} V_I + V_I^T A_{II} V_I \\ &\leq V_C^T A_{CC} V_C + 2|V_C^T A_{CE} V_E| + (1 + 4\delta_h) V_E^T T_{EE} V_E \\ &\quad + 2|V_E^T A_{EI} V_I| + (1 + 8\delta_h) V_I^T D_{II} V_I \\ &\leq V_C^T A_{CC} V_C + (\delta_h V_E^T T_{EE} V_E + 2\delta_h V_C^T A_{CC} V_C) + (1 + 4\delta_h) V_E^T T_{EE} V_E \\ &\quad + (\delta_h V_I^T D_{II} V_I + 2\delta_h V_E^T T_{EE} V_E) + (1 + 8\delta_h) V_I^T D_{II} V_I \\ &= (1 + 2\delta_h) V_C^T A_{CC} V_C + (1 + 7\delta_h) V_E^T T_{EE} V_E + (1 + 9\delta_h) V_I^T D_{II} V_I \end{aligned}$$

for all V . \square

Note 1. The lower bound holds for all δ_h , but the bound is only useful in the case that $\delta_h < 1/3$, since a lower bound of $0 \leq V^T AV$ naturally holds. When $\delta_h < 1/3$, the theorem establishes spectral equivalence bounds that show that A_D is an excellent preconditioner for A as $\delta_h \rightarrow 0$.

Note 2. Slight variations in these bounds are easily derived by varying the division of the constants in (4.3) and (4.4). In particular, the upper bound can be slightly decreased at the expense of the constant in the lower bound and vice versa. However, since $\delta_h \ll 1$ is the case of interest, these variations are largely irrelevant.

COROLLARY 4.2. *Under the assumptions of Theorem 4.1, if M_{CC} is spectrally equivalent to A_{CC} , meaning that*

$$c_0 V_C^T M_{CC} V_C \leq V_C^T A_{CC} V_C \leq c_1 V_C^T M_{CC} V_C \text{ for all } V_C,$$

then the matrix

$$A_M = \begin{bmatrix} M_{CC} & 0 & 0 \\ 0 & T_{EE} & 0 \\ 0 & 0 & D_{II} \end{bmatrix}$$

satisfies

$$\min(1 - 3\delta_h, c_0(1 - 2\delta_h)) V^T A_M V \leq V^T AV \leq \max(1 + 9\delta_h, c_1(1 + 2\delta_h)) V^T A_M V$$

for all V .

Note 3. Corollary 4.2 is particularly relevant to the case where M_{CC} is the inverse of an effective multigrid preconditioner for A_{CC} . If the stationary multigrid cycle represented by $I - M_{CC}^{-1}A_{CC}$ has spectral radius $\alpha < 1$, then

$$(1 - \alpha)V_C^T M_{CC} V_C \leq V_C^T A_{CC} V_C \leq (1 + \alpha)V_C^T M_{CC} V_C$$

for all V_C . Under this assumption, the bound in Corollary 4.2 becomes

$$\begin{aligned} \min(1 - 3\delta_h, (1 - \alpha)(1 - 2\delta_h))V^T A_M V \\ \leq V^T A V \leq \max(1 + 9\delta_h, (1 + \alpha)(1 + 2\delta_h))V^T A_M V \end{aligned}$$

for all V . Since we typically expect $\alpha \approx 0.1$ (cf. [8, 49]) and $\delta_h \ll 1$, a better rule of thumb would be

$$(1 - \alpha)(1 - 2\delta_h)V^T A_M V \leq V^T A V \leq (1 + \alpha)(1 + 2\delta_h)V^T A_M V$$

for all V . This suggests that the convergence of a preconditioner consisting of one multigrid V-cycle applied to the corner region, appropriately ordered tridiagonal solves applied to the edge layers, plus diagonal scaling of the interior region should be very similar to that of multigrid applied directly to the boundary layer regions alone. In the case of Shishkin meshes, this means that we expect to recover the optimal scaling of multigrid on uniform meshes, while we expect similar behavior for Bakhvalov meshes when a robust multigrid approach is used.

In light of Corollary 4.2, in section 4.7, we present results for a boundary-layer preconditioner that makes use of a single V-cycle of BoxMG in the corner region, appropriately ordered tridiagonal solves to treat the edge layers, and diagonal scaling in the interior. For a tensor product of fitted meshes with $N/2$ points in each layer (including the transition points), this implies that one-quarter of the grid will be treated with a preconditioner that has the same cost as BoxMG, while the remaining three-quarters will be treated in a much cheaper manner. Estimating the work of a BoxMG V-cycle as four full line relaxations on each level times a factor of four-thirds (to account for the hierarchy of levels of sizes N^2 , $N^2/4$, $N^2/16$, etc.), this suggests that on three-quarters of the grid, the boundary-layer preconditioner does three-sixteenths of the work of the multigrid V-cycle. The total work estimate is, then, $1/4 + (3/4)(3/16) = 25/64$ times the work of a multigrid V-cycle on the full problem. This is, clearly, an overestimate, since there are extra savings from the lack of interpolation and restriction, and from the cost of diagonal scaling compared to line solves, but provides a rough estimate for the improvement possible from a preconditioner that is tuned to the full problem structure.

4.6. Stopping criteria. As is typical for iterative methods, we aim to stop the iteration as soon as we have an iterate, $U^{(k)}$, that approximates the “best” discrete approximation, U , of the continuum solution, u , to a degree of accuracy that is of the same order as the discretization error, $\|u - U\|$, giving

$$\|u - U^{(k)}\| \leq \|u - U\| + \|U - U^{(k)}\| \leq C\|u - U\|$$

for a moderate constant, C . In the case of finite difference approximations on Shishkin and Bakhvalov meshes, we know that the discretization error bound takes the form

$$\|u - U\|_\infty \leq g(N)$$

for $g(N) = CN^{-2} \ln^2(N)$ for Shishkin, and $g(N) = CN^{-2}$ for Bakhvalov, where the constants are independent of ε . We note that we use the discrete maximum norm here, as the natural norm for which error bounds are available.

Naturally, the discrete approximation error, $E^{(k)} = U - U^{(k)}$, is unknown during the iteration; thus, we must measure this error indirectly. For stationary iterations, the best indicator of this error is the residual, $R^{(k)} = F - AU^{(k)} = AE^{(k)}$, since $AU = F$. The simplest possible bound, then, would be to write $E^{(k)} = A^{-1}R^{(k)}$, giving

$$\|U - U^{(k)}\|_\infty \leq \|A^{-1}\|_\infty \|R^{(k)}\|_\infty.$$

Practical bounds on $\|A^{-1}\|_\infty$ arise from M-matrix theory; Theorem 2.7 in [40] states that

$$\|A^{-1}\|_\infty \leq \frac{\|V\|_\infty}{\min_k (AV)_k},$$

where V is a majorizing vector with $(AV)_k > 0$ for all k . For the discretization matrices considered here, the choice of V as the vector of all ones is natural, giving $\min_k (AV)_k$ as the Geršgorin bound on the smallest eigenvalue of A ; this is bounded from below by $\beta^2 h_{\min}^2$, where h_{\min} is the minimal mesh spacing on the grid. Numerical experiments show this bound to be sharp for small ε . With this bound, we can then guarantee a discrete approximation error comparable to the discretization error by asking for $\|R^{(k)}\|_\infty \leq \beta^2 h_{\min}^2 g(N)$, ensuring $\|U - U^{(k)}\|_\infty < g(N)$. For small ε , however, h_{\min} scales as $h_{\min} \sim \varepsilon \ln(N)/(N\beta)$ for Shishkin meshes, and as $h_{\min} \sim \varepsilon/(N\beta)$ for Bakhvalov meshes, requiring $\|R^{(k)}\|_\infty \leq C\varepsilon^2(\ln^3 N)/N^3$ to achieve discretization accuracy on Shishkin meshes, and $\|R^{(k)}\|_\infty \leq C\varepsilon^2/N^3$ to achieve discretization accuracy on Bakhvalov meshes. For small ε and moderate N , these values will be much less than machine precision, implying that it is difficult to compute $\|R^{(k)}\|_\infty$ accurately enough to guarantee convergence.

To overcome this limitation of stationary iterations, we turn to the preconditioned conjugate gradients algorithm and look to make use of the standard stopping criterion that bounds the inner product of $R^{(k)}$ with the preconditioned residual, $Z^{(k)} = MR^{(k)}$, where M represents the preconditioning matrix applied. When M is a good preconditioner in the spectral sense that MA (or, more precisely, $A^{1/2}MA^{1/2}$, where $A^{1/2}$ is the principal square root of A) is spectrally equivalent to the identity, this inner product accurately estimates $\|E^{(k)}\|_A^2$ as

$$\left(Z^{(k)}\right)^T R^{(k)} = \left(E^{(k)}\right)^T AMAE^{(k)} \approx \|E^{(k)}\|_A^2.$$

To bound $\|E^{(k)}\|_\infty$ by $\|E^{(k)}\|_A$, we make use of an intermediate bound, trading the max-norm for a discrete ℓ_2 -norm, noting that

$$\frac{1}{\sqrt{n}} \|E^{(k)}\|_2 \leq \|E^{(k)}\|_\infty \leq \|E^{(k)}\|_2$$

for vectors of length n . A reasonable, although not rigorous, estimate is to assume near equality with the lower bound, rather than the often pessimistic upper bound, under the assumption that the discrete approximation error is nearly equidistributed across the mesh, so that the error is not concentrated at relatively few mesh points.

With this assumption, a more tractable bound is available, writing

$$\|E^{(k)}\|_\infty \approx \frac{c}{\sqrt{n}} \|E^{(k)}\|_2 \leq \frac{c}{\sqrt{n}} \|A^{-1/2}\|_2 \|E^{(k)}\|_A,$$

where $A^{-1/2}$ is the principal square root of the symmetric and positive-definite matrix, A^{-1} . This bound presents two advantages. First, $\|E^{(k)}\|_A$ is both naturally estimated and minimized by the preconditioned conjugate gradient algorithm and, as such, makes a natural stopping criterion for the iterative approach. Second, as we show below, the natural bound on $\|A^{-1/2}\|_2$ needed to determine the stopping criterion is much smaller than that of $\|A^{-1}\|_\infty$ (and $\|A^{-1}\|_2$), yielding a stopping criterion that is also less susceptible to problems with round-off errors.

Since $A^{-1/2}$ is the principle square root of A^{-1} , we have $\|A^{-1/2}\|_2^2 = \|A^{-1}\|_2$, as the eigenvalues of A^{-1} are the squares of those of $A^{-1/2}$. Thus, we can estimate the needed norm by applying Geršgorin's theorem to A to estimate its smallest eigenvalue, deriving the same bound on $\|A^{-1}\|_2$ as was given for $\|A^{-1}\|_\infty$ above. This implies that $\|A^{-1/2}\|_2 \leq 1/(\beta h_{\min})$ for the two-dimensional problems. Thus, we have

$$\|E^{(k)}\|_\infty \approx \frac{c}{\sqrt{n}} \|E^{(k)}\|_2 \leq \frac{c}{N} \|A^{-1/2}\|_2 \|E^{(k)}\|_A \leq \frac{c}{N} \frac{1}{\beta h_{\min}} \|E^{(k)}\|_A.$$

Similar bounds can be derived for the one-dimensional problems for both $\|A^{-1}\|_\infty$ and $\|A^{-1/2}\|_A$, with changes that the matrices are of dimension $n \approx N$ instead of N^2 and a scaling of $h_{\min}\beta^2$ for the minimal eigenvalue estimate.

Given a bound of this form, the stopping criterion for $\|E^{(k)}\|_A$ to guarantee $\|E^{(k)}\|_\infty < g(N)$ (up to the assumption relating $\|E^{(k)}\|_\infty$ and $\|E^{(k)}\|_2$) is given by

$$\|E^{(k)}\|_A \leq C \frac{\beta N h_{\min}}{c} g(N);$$

with this, we have

$$\|u - U^{(k)}\|_\infty \lesssim Cg(N),$$

ensuring optimality of the approximate discrete solution, $U^{(k)}$.

What this analysis doesn't treat is the size of the constants, C and c , or their variation in N and ε . While these constants are uniformly bounded independently of N and ε , numerical experiments in two dimensions show that when ε is large, stricter convergence bounds are needed to obtain the best possible accuracy (that is, to be almost identical to that of a direct solver). Thus, for the results presented in this paper for two-dimensional problems, we use a graduated stopping tolerance of 10^{-3} times the above bounds when $\varepsilon = 1$, 10^{-2} when $\varepsilon = 10^{-1}$, and 10^{-1} when $\varepsilon = 10^{-2}$. For smaller values of ε , we take the constant, $c = 1$, with no further scaling. In all results, we take β as an input parameter, matching what is used for the grid generation, and use a computed value of h_{\min} . Finally, we note that although the required bounds on $\|E^{(k)}\|_A$ through this analysis are dependent on both N and ε , the multigrid-based methods used here lead to per-iteration reduction factors bounded below unity in each iteration, independently of N and ε . Thus, the number of iterations for convergence is bounded, at worst, by logarithmic factors of N and ε ; in practice, iteration counts grow quite slowly with these terms.

TABLE 4.4

CPU times for AMG preconditioned CG to achieve theoretically optimal error reduction on Bakhvalov meshes, averaged over 3 runs. Corresponding iteration counts for a single run are given in parentheses.

ε^2	$N = 2^7$	$N = 2^8$	$N = 2^9$	$N = 2^{10}$	$N = 2^{11}$	$N = 2^{12}$
10^0	0.07 (6)	0.38 (7)	1.81 (7)	8.06 (8)	36.09 (9)	166.10 (10)
10^{-2}	0.07 (6)	0.37 (6)	1.81 (7)	8.04 (8)	38.43 (10)	165.77 (10)
10^{-4}	0.05 (5)	0.32 (6)	1.62 (7)	7.01 (7)	33.77 (9)	147.10 (9)
10^{-6}	0.04 (5)	0.22 (5)	1.13 (6)	5.91 (7)	30.01 (8)	143.08 (9)
10^{-8}	0.04 (5)	0.22 (5)	1.13 (6)	5.28 (7)	24.15 (8)	113.79 (9)
10^{-10}	0.04 (5)	0.25 (7)	1.31 (8)	6.06 (9)	29.08 (11)	142.97 (13)
10^{-12}	0.05 (6)	0.25 (7)	1.31 (8)	6.05 (9)	29.07 (11)	142.93 (13)

4.7. Numerical results. We present here results for the test problem (1.5) on Bakhvalov meshes; results for Shishkin meshes are nearly identical and so omitted. We do not report errors for the solutions computed here. Instead, we fix convergence tolerances as discussed above (with a small constant factor for BoxMG, detailed below) such that the error in the iterative solutions generated by geometric multigrid, AMG, BoxMG, and, for suitably small δ_h , the boundary-layer preconditioners, matches that of the direct solver (as reported, for example, in Table 2.1 for the Bakhvalov mesh in two dimensions) to three significant digits.

Table 4.4 reports the times for the AMG solution of these linear systems, averaged over 3 runs, and the number of iterations required for a single run; since the iteration counts are the same for each run (as a zero initial guess and fixed right-hand side are used for each value of N and ε), no averaging is needed in these results. Comparing this with Table 4.1, we see that the AMG solution times are generally comparable to the best times for CHOLMOD for small N . For larger N , however, the near-optimal scaling of AMG coupled with the clearly suboptimal scaling of CHOLMOD results in much better performance when comparing AMG to the best-case solve times of CHOLMOD, and substantial improvements when CHOLMOD is slowed by subnormal arithmetic, as discussed in section 4.1.

Hidden in the AMG iteration times is some noteworthy variation in the number of PCG iterations needed for each problem. Considering the iteration counts in Table 4.4, we notice that, for fixed ε , there is a slight increase in the number of iterations needed to achieve the stopping tolerance with N , roughly proportional to $\log_2 N$. This is naturally predicted by the assumption of a fixed error reduction per iteration of AMG preconditioned CG, given that the stopping criteria decrease with N . Scaling with ε is less clear, where we see an initial decrease in iteration counts as ε decreases, followed by an increase as ε reaches 10^{-5} and 10^{-6} . For fitted meshes, decreases in ε match increases in the worst-case mesh aspect ratio, which are typically correlated with convergence of AMG. This variation in ε can, to some extent, be ameliorated by varying the value of the AMG parameter defining the relative magnitude of strong connections in each row. For worse mesh aspect ratios, larger values of the parameter (up to 1.0, where only the largest magnitude connections are taken to be strong) are needed to stabilize AMG convergence, while only a fixed value (0.25) has been used here.

As seen in other cases, BoxMG clearly outperforms AMG on these problems, with averaged solution times reported in Table 4.5. Comparing this with Table 4.4, we see that CPU times range from about 3 times faster for small grids to about 40% faster for large grids. The iteration counts, also shown in Table 4.5, again show $\log_2 N$ scaling

TABLE 4.5

CPU times for BoxMG preconditioned CG to achieve theoretically optimal error reduction on Bakhwalov meshes, averaged over 3 runs. Corresponding iteration counts for a single run are given in parentheses.

ε^2	$N = 2^7$	$N = 2^8$	$N = 2^9$	$N = 2^{10}$	$N = 2^{11}$	$N = 2^{12}$
10^0	0.02 (8)	0.13 (9)	0.77 (10)	4.44 (10)	22.38 (11)	110.39 (11)
10^{-2}	0.03 (8)	0.13 (9)	0.75 (10)	4.83 (11)	22.51 (11)	119.41 (12)
10^{-4}	0.02 (7)	0.12 (8)	0.68 (9)	4.45 (10)	24.35 (12)	119.42 (12)
10^{-6}	0.02 (6)	0.10 (7)	0.62 (8)	4.05 (9)	20.69 (10)	101.48 (10)
10^{-8}	0.02 (6)	0.11 (7)	0.62 (8)	4.05 (9)	18.86 (9)	101.48 (10)
10^{-10}	0.02 (6)	0.10 (7)	0.63 (8)	4.06 (9)	18.92 (9)	101.60 (10)
10^{-12}	0.02 (6)	0.10 (7)	0.62 (8)	4.06 (9)	18.96 (9)	101.82 (10)

TABLE 4.6

Computed values of $\delta_h = \varepsilon^2 / (h_I^2 \beta)$ for Bakhwalov meshes on $[0, 1]^2$ with a single boundary layer and $\beta = 0.99$.

ε^2	$N = 2^7$	$N = 2^8$	$N = 2^9$	$N = 2^{10}$	$N = 2^{11}$	$N = 2^{12}$
1	1.65×10^4	6.62×10^4	2.65×10^5	1.06×10^6	4.24×10^6	1.69×10^7
10^{-2}	1.24×10^2	4.99×10^2	2.00×10^3	7.98×10^3	3.19×10^4	1.28×10^5
10^{-4}	5.12×10^{-1}	2.05×10^0	8.19×10^0	3.27×10^1	1.31×10^2	5.24×10^2
10^{-6}	4.27×10^{-3}	1.71×10^{-2}	6.83×10^{-2}	2.73×10^{-1}	1.09×10^0	4.37×10^0
10^{-8}	4.15×10^{-5}	1.66×10^{-4}	6.65×10^{-4}	2.66×10^{-3}	1.06×10^{-2}	4.25×10^{-2}
10^{-10}	4.14×10^{-7}	1.66×10^{-6}	6.62×10^{-6}	2.65×10^{-5}	1.06×10^{-4}	4.24×10^{-4}
10^{-12}	4.14×10^{-9}	1.66×10^{-8}	6.62×10^{-8}	2.65×10^{-7}	1.06×10^{-6}	4.24×10^{-6}

as problem size increases, but notably less variation in ε . This is likely due to the effectiveness of the alternating-direction line relaxation used within BoxMG, although there is no existing theoretical verification of this scaling. Once again, the times are notably shorter (and better scaling) than those for the direct solver in Table 4.1, as well as showing no sensitivity to subnormal numbers (as one would expect, since a direct solver is only applied on the coarsest of grids). We note that, to achieve this performance, we use a slightly stricter stopping tolerance for BoxMG than we do for AMG or BLPCG, decreasing those described above by an experimentally determined factor of 10. Such small variations are expected, as the component bounds employed are true to within constant factors that have been neglected.

The boundary-layer preconditioner discussed in section 4.5 is expected to yield even faster times, but only when $\delta_h \ll 1$. Table 4.6 shows the values of δ_h for various choices of N and ε ; as expected, for fixed ε , δ_h increases with increasing N (since h_I decreases) while, for fixed N , δ_h decreases with decreasing ε . Note, however, that the decrease is not linear in ε^2 , as h_I is dependent on ε .

In Table 4.7, we only report results for cases where $\delta_h < 0.1$, consistent with the spectral equivalence bounds in Theorem 4.1 and Corollary 4.2. We note that, when $\delta_h < 0.01$, there is no notable difference between the errors in the iterative solutions generated by boundary-layer preconditioned CG (BLPCG) with the stopping criterion as given above from those generated by the direct solver. For $0.01 \leq \delta_h < 0.1$, the errors generated by BLPCG are somewhat larger, with observed relative increases of up to about 15%. These values could be decreased by a tighter stopping tolerance, but this would lead to oversolving in the cases where δ_h is smaller.

Comparing Table 4.7 with Tables 4.4 and 4.5, we see that, when the BLPCG approach is effective, it offers a substantial further reduction in the time to solution on large grids, by slightly more than a factor of 3 over BoxMG, and by a larger factor over AMG. This is slightly better than predicted by the rough cost estimate given in

TABLE 4.7

CPU times for preconditioned CG with the boundary-layer preconditioner to achieve (near) theoretically optimal error reduction on Bakhvalov meshes, averaged over 3 runs. Corresponding iteration counts for a single run are given in parentheses.

ε^2	$N = 2^7$	$N = 2^8$	$N = 2^9$	$N = 2^{10}$	$N = 2^{11}$	$N = 2^{12}$
10^{-6}	0.01 (6)	0.05 (6)	0.26 (8)			
10^{-8}	0.01 (6)	0.04 (6)	0.23 (7)	1.24 (8)	6.63 (8)	28.01 (8)
10^{-10}	0.00 (5)	0.04 (6)	0.23 (7)	1.24 (8)	6.63 (8)	28.20 (8)
10^{-12}	0.00 (5)	0.04 (6)	0.22 (7)	1.24 (8)	6.62 (8)	28.15 (8)

section 4.5, but not surprisingly so, especially as BLPCG typically converges in fewer iterations than BoxMG. The iteration counts in Table 4.7 are steady and show only small dependency on ε .

5. Extensions to three-dimensional problems. While we do not directly treat either the theory or practice of three-dimensional problems in this paper, we note that the techniques proposed here for two dimensions all naturally extend to three. Both AMG and BoxMG can be directly applied to the 7-point stencils of finite difference discretizations in three dimensions, and past comparisons show BoxMG to be about ten to fifteen times faster than AMG for some diffusion problems [31]. In three dimensions, BoxMG makes use of plane-based relaxation where, because of the nonoptimal scaling of direct solvers for two-dimensional problems, relaxation is composed of a single two-dimensional V-cycle along each plane of grid points in each direction of the mesh.

Similarly, a three-dimensional analogue of the boundary-layer preconditioner is also natural. Considering a problem with one boundary layer in each direction and a tensor product of meshes with $N/2$ points in the boundary layer and $N/2$ evenly spaced points in the interior, there is one highly resolved corner with $N^3/8$ points, to which we would apply a standard three-dimensional BoxMG V-cycle, with cost roughly $1/8$ of that of the V-cycle on all N^3 points. Along each of the x -, y -, and z -axes, away from the corner, there are layers that are highly resolved in two directions, but not the third. For example, along the x -axis away from the origin, there is high resolution in the y - and z -directions, but not in the x -direction. Each of these regions has $N^3/8$ points in it, and can be effectively treated by a single plane solve, instead of the six plane solves (in three alternating directions, both before and after coarse-grid correction) used in the standard BoxMG V-cycle. Thus, the cost of treating these regions at each iteration is $(3/8)(1/6)(7/8)$, where the factor of $7/8$ comes from the treatment of only the finest grid, and not grids of N^3 , $N^3/8$, $N^3/64$, etc., as in the multigrid V-cycle. Next, along each face in the xy -, xz -, and yz -planes, there is a layer that is highly resolved in one direction (orthogonal to the plane), but not along the plane. In these regions, line solves are sufficient. In the remaining $N^3/8$ points in the interior, diagonal scaling is sufficient, but we will estimate the cost of line solves here, too, for simplicity. Thus, for $N^3/2$ points, we make use of a single line solve (oriented in the direction of the mesh refinement) or point solve, with cost $1/4$ of the relaxation in a plane solve, which has relative complexity of $4/3$, and of which there are six in a three-dimensional BoxMG V-cycle, with relative complexity of $8/7$. This gives a cost of $(1/2)(1/4)(3/4)(1/6)(7/8)$ for the treatment of these regions, relative to the cost of a full multigrid V-cycle. Adding these costs gives

$$1/8 + (3/8)(1/6)(7/8) + (1/2)(1/4)(3/4)(1/6)(7/8) = 99/512$$

of the cost of a full three-dimensional BoxMG V-cycle. While the nonuniform coarsening used by AMG should naturally coarsen in ways consistent with this proposed boundary-layer preconditioner, the expectation that BoxMG by itself outperforms AMG implies that the boundary-layer preconditioner should be the fastest approach. We also note that the expected speedup in three dimensions is by a larger factor than that given in two dimensions (which, in turn, outperforms one dimension).

6. Conclusions. This paper focuses on the solution of the linear systems of equations that arise from finite difference discretizations of singularly perturbed reaction-diffusion equations on fitted meshes. We show that the commonly used direct solvers scale poorly with both the mesh size (as is well known) and the perturbation parameter (due to slow performance of IEEE floating-point arithmetic in this regime). In contrast, robust multigrid methods offer nearly parameter-uniform efficiency in solving these linear systems, and the carefully derived stopping criteria ensure accurate, and efficient, approximations to the discrete solution. Existing software implementing algebraic and black-box multigrid approaches are demonstrated to offer nearly scalable performance. A new boundary-layer preconditioner is also proposed and is shown, both theoretically and numerically, to offer near-optimal scaling at still lower cost. Overall, we see speedups of factors of 40 or greater for problems with small perturbation parameters.

The approaches proposed here can clearly be extended to a number of other situations. Bilinear finite element discretizations on fitted meshes generate similar stencils to those considered here and, as such, it is expected that simple extensions of these approaches can be applied in this case. Other layer-adapted meshes, such as hybrid Bakhvalov–Shishkin meshes (e.g., [41]), and variants due to Vulcanović (e.g., [51]), create similar structures; in these cases, very similar approaches should also yield near-optimal scaling solution algorithms for the resulting linear systems. The developed theory for the boundary-layer preconditioner, in particular, is directly applicable to finite difference discretizations on any fitted mesh that yields uniform mesh width away from the boundary and corner layers, while the stopping criterion derived depends only on the minimum mesh width and existing theoretical error estimates for the discretization.

Acknowledgments. The authors would like to thank Iain Duff for his helpful discussions that exposed the issue of subnormal numbers in causing the slowdown observed with sparse direct methods for these problems. They would also like to thank the anonymous referees for their helpful comments that improved the quality of this paper.

REFERENCES

- [1] R. E. ALCOUFFE, A. BRANDT, J. E. DENDY, JR., AND J. W. PAINTER, *The multi-grid method for the diffusion equation with strongly discontinuous coefficients*, SIAM J. Sci. Stat. Comput., 2 (1981), pp. 430–454.
- [2] A. R. ANSARI AND A. F. HEGARTY, *A note on iterative methods for solving singularly perturbed problems using non-monotone methods on Shishkin meshes*, Comput. Methods Appl. Mech. Engrg., 192 (2003), pp. 3673–3687.
- [3] I. BABUŠKA AND M. SURI, *On locking and robustness in the finite element method*, SIAM J. Numer. Anal., 29 (1992), pp. 1261–1293.
- [4] N. BAKHVALOV, *Towards optimization of methods for solving boundary value problems in the presence of boundary layers*, Zh. Vychisl. Mat. Mat. Fiz., 9 (1969), pp. 841–859 (in Russian).

- [5] B. BERGEN, G. WELLEIN, F. HÜLSEMANN, AND U. RÜDE, *Hierarchical hybrid grids: Achieving TERAFL0P performance on large scale finite element simulations*, Int. J. Parallel Emergent Distrib. Systems, 22 (2007), pp. 311–329.
- [6] J. H. BRAMBLE, J. E. PASCIAK, J. WANG, AND J. XU, *Convergence estimates for multigrid algorithms without regularity assumptions*, Math. Comp., 57 (1991), pp. 23–45.
- [7] A. BRANDT, *Multi-level adaptive techniques (MLAT) for singular-perturbation problems*, in Numerical Analysis of Singular Perturbation Problems, P. W. Hemker and J. J. H. Miller, eds., Academic, New York, 1979, pp. 53–142.
- [8] W. L. BRIGGS, V. E. HENSON, AND S. F. MCCORMICK, *A Multigrid Tutorial*, 2nd ed., SIAM, Philadelphia, 2000.
- [9] N. M. CHADHA AND N. KOPTEVA, *A robust grid equidistribution method for a one-dimensional singularly perturbed semilinear reaction-diffusion problem*, IMA J. Numer. Anal., 31 (2011), pp. 188–211.
- [10] D. CHEN, S. MACLACHLAN, AND M. KILMER, *Iterative parameter-choice and multigrid methods for anisotropic diffusion denoising*, SIAM J. Sci. Comput., 33 (2011), pp. 2972–2994.
- [11] Y. CHEN, T. A. DAVIS, W. W. HAGER, AND S. RAJAMANICKAM, *Algorithm 887: CHOLMOD, supernodal sparse Cholesky factorization and update/downdate*, ACM Trans. Math. Software, 35 (2008), pp. 22.
- [12] C. CLAVERO, J. GRACIA, AND E. O’RIORDAN, *A parameter robust numerical method for a two dimensional reaction-diffusion problem*, Math. Comp., 74 (2005), pp. 1743–1758.
- [13] T. A. DAVIS, *Direct Methods for Sparse Linear Systems*, Fundamentals of Algorithms 2, SIAM, Philadelphia, 2006.
- [14] S. DAVYDYCHEVA, V. DRUSKIN, AND T. HABASHY, *An efficient finite-difference scheme for electromagnetic logging in 3D anisotropic inhomogeneous media*, Geophysics, 68 (2003), pp. 1525–1536.
- [15] J. E. DENDY, *Black box multigrid*, J. Comput. Phys., 48 (1982), pp. 366–386.
- [16] V. DRUSKIN AND L. KNIZHNERMAN, *Gaussian spectral rules for the three-point second differences. I. A two-point positive definite problem in a semi-infinite domain*, SIAM J. Numer. Anal., 37 (1999), pp. 403–422.
- [17] I. S. DUFF, *MA57 – a code for the solution of sparse symmetric definite and indefinite systems*, ACM Trans. Math. Software, 30 (2004), pp. 118–144.
- [18] P. A. FARRELL, A. F. HEGARTY, J. J. H. MILLER, E. O’RIORDAN, AND G. I. SHISHKIN, *Robust Computational Techniques for Boundary Layers*, Appl. Math. 16, Chapman & Hall/CRC, Boca Raton, FL, 2000.
- [19] P. A. FARRELL AND G. I. SHISHKIN, *On the convergence of iterative methods for linear systems arising from singularly perturbed equations*, in Proceedings of the Copper Mountain Conference on Iterative Methods, 1998, pp. 1–7; available online from <http://www.cs.kent.edu/~farrell/papers/papers.html>.
- [20] F. GASPAR, C. CLAVERO, AND F. LISBONA, *Some numerical experiments with multigrid methods on Shishkin meshes*, J. Comput. Appl. Math., 138 (2002), pp. 21–35.
- [21] F. GASPAR, F. LISBONA, AND C. CLAVERO, *Multigrid methods and finite difference schemes for 2D singularly perturbed problems*, in Numerical Analysis and Its Applications, Lecture Notes in Comput. Sci. 1988, L. Vulkov, P. Yalamov, and J. Wasniewski, eds., Springer, Berlin, 2001, pp. 128–135.
- [22] A. GEORGE, *Nested dissection of a regular finite element mesh*, SIAM J. Numer. Anal., 10 (1973), pp. 345–363.
- [23] A. J. HOFFMAN, M. S. MARTIN, AND D. J. ROSE, *Complexity bounds for regular finite difference and finite element grids*, SIAM J. Numer. Anal., 10 (1973), pp. 364–369.
- [24] R. B. KELLOGG, T. LINSS, AND M. STYNES, *A finite difference method on layer-adapted meshes for an elliptic reaction-diffusion system in two dimensions*, Math. Comp., 77 (2008), pp. 2085–2096.
- [25] R. KELLOGG, N. MADDEN, AND M. STYNES, *A parameter-robust numerical method for a system of reaction-diffusion equations in two dimensions*, Numer. Methods Partial Differential Equations, 24 (2008), pp. 312–334.
- [26] N. KOPTEVA AND M. STYNES, *A robust adaptive method for a quasi-linear one-dimensional convection-diffusion problem*, SIAM J. Numer. Anal., 39 (2001), pp. 1446–1467.
- [27] O. LAWLOR, H. GOVIND, I. DOOLEY, M. BREITENFELD, AND L. KALE, *Performance degradation in the presence of subnormal floating-point values*, in Proceedings of the International Workshop on Operating System Interference in High Performance Applications, 2005; available online from <http://research.ihost.com/osihpa/>.
- [28] T. LINSS, *Layer-Adapted Meshes for Reaction-Convection-Diffusion Problems*, Lecture Notes in Math. 1985, Springer, Berlin, 2010.
- [29] F. LIU, N. MADDEN, M. STYNES, AND A. ZHOU, *A two-scale sparse grid method for a singularly*

- perturbed reaction-diffusion problem in two dimensions*, IMA J. Numer. Anal., 29 (2009), pp. 986–1007.
- [30] S. P. MACLACHLAN, J. D. MOULTON, AND T. P. CHARTIER, *Robust and adaptive multigrid methods: Comparing structured and algebraic approaches*, Numer. Linear Algebra Appl., 19 (2012), pp. 389–413.
- [31] S. P. MACLACHLAN, J. M. TANG, AND C. VUIK, *Fast and robust solvers for pressure-correction in bubbly flow problems*, J. Comput. Phys., 227 (2008), pp. 9742–9761.
- [32] S. MACLACHLAN AND N. MADDEN, *Robust Solution of Singularly Perturbed Problems using Multigrid Methods; Analysis and Numerical Results in One and Two Dimensions*, Technical report, Department of Mathematics, Tufts University, 2012.
- [33] J. J. H. MILLER, E. O’RIORDAN, AND G. I. SHISHKIN, *Fitted Numerical Methods for Singular Perturbation Problems: Error Estimates in the Maximum Norm for Linear Problems in One and Two Dimensions*, World Scientific, Singapore, 1996.
- [34] K. MORTON, *Numerical Solution of Convection-Diffusion Problems*, Applied Math. Math. Comput. 12, Chapman & Hall, London, 1996.
- [35] W. MULDER, *Geophysical modelling of 3D electromagnetic diffusion with multigrid*, Comput. Visual. Sci., 11 (2008), pp. 129–138.
- [36] M. A. OLSHANSKII AND A. REUSKEN, *On the convergence of a multigrid method for linear reaction-diffusion problem*, Computing, 65 (2000), pp. 193–202.
- [37] M. L. OVERTON, *Numerical Computing with IEEE Floating Point Arithmetic*, SIAM, Philadelphia, 2001.
- [38] M. H. PROTTER AND H. F. WEINBERGER, *Maximum Principles in Differential Equations*, Springer-Verlag, New York, 1984.
- [39] B. REPS, W. VANROOSE, AND H. BIN ZUBAIR, *On the indefinite Helmholtz equation: Complex stretched absorbing boundary layers, iterative analysis, and preconditioning*, J. Comput. Phys., 229 (2010), pp. 8384–8405.
- [40] H.-G. ROOS, M. STYNES, AND L. TOBISKA, *Robust Numerical Methods for Singularly Perturbed Differential Equations*, Springer Ser. Comput. Math. 24, 2nd ed., Springer-Verlag, Berlin, 2008.
- [41] H.-G. ROOS AND H. ZARIN, *A second-order scheme for singularly perturbed differential equations with discontinuous source term*, J. Numer. Math., 10 (2002), pp. 275–289.
- [42] H.-G. ROOS, *A note on the conditioning of upwind schemes on Shishkin meshes*, IMA J. Numer. Anal., 16 (1996), pp. 529–538.
- [43] J. W. RUGE AND K. STÜBEN, *Algebraic multigrid*, in Multigrid Methods, Front. Appl. Math. 3, S. F. McCormick, ed., SIAM, Philadelphia, 1987, pp. 73–130.
- [44] G. I. SHISHKIN AND L. SHISHKINA, *Difference Methods for Singular Perturbation Problems*, Chapman & Hall/CRC Monogr. Surv. Pure Appl. Math., Chapman & Hall/CRC, Boca Raton, FL, 2008.
- [45] G. I. SHISHKIN, *Grid Approximation of Singularly Perturbed Elliptic and Parabolic Equations*, Second doctoral thesis, Keldysh Institute, Moscow, 1991 (in Russian).
- [46] I. SINGER AND E. TURKEL, *A perfectly matched layer for the Helmholtz equation in a semi-infinite strip*, J. Comput. Phys., 201 (2004), pp. 439–465.
- [47] R. STEVENSON, *A robust hierarchical basis preconditioner on general meshes*, Numer. Math., 78 (1997), pp. 269–303.
- [48] K. STÜBEN, *An introduction to algebraic multigrid*, in Multigrid, U. Trottenberg, C. Oosterlee, and A. Schüller, eds., Academic, London, 2001, pp. 413–528.
- [49] U. TROTTENBERG, C. W. OOSTERLEE, AND A. SCHÜLLER, *Multigrid*, Academic, London, 2001.
- [50] N. UMETANI, S. P. MACLACHLAN, AND C. W. OOSTERLEE, *A multigrid-based shifted-Laplacian preconditioner for a fourth-order Helmholtz discretization*, Numer. Linear Algebra Appl., 16 (2009), pp. 603–626.
- [51] R. VULANOVIĆ, *A priori meshes for singularly perturbed quasilinear two-point boundary value problems*, IMA J. Numer. Anal., 21 (2001), pp. 349–366.
- [52] J. XU AND Y. ZHU, *Uniform convergent multigrid methods for elliptic problems with strongly discontinuous coefficients*, Math. Models Methods Appl. Sci., 18 (2008), pp. 77–105.
- [53] J. XU AND L. ZIKATANOV, *The method of alternating projections and the method of subspace corrections in Hilbert space*, J. Amer. Math. Soc., 15 (2002), pp. 573–597.
- [54] H. BIN ZUBAIR, S. P. MACLACHLAN, AND C. W. OOSTERLEE, *A geometric multigrid method based on L-shaped coarsening for PDEs on stretched grids*, Numer. Linear Algebra Appl., 17 (2010), pp. 871–894.
- [55] H. BIN ZUBAIR, C. W. OOSTERLEE, AND R. WIENANDS, *Multigrid for high-dimensional elliptic partial differential equations on non-equidistant grids*, SIAM J. Sci. Comput., 29 (2007), pp. 1613–1636.