

Optimal Polynomial Smoothers for Multigrid V-cycles

James Lottes¹

¹Google Research
jlottes@google.com

May 10, 2023

Abstract

The idea of using polynomial methods to improve simple smoother iterations within a multigrid method for a symmetric positive definite (SPD) system is revisited. A two-level bound going back to Hackbusch is optimized by a very simple iteration, a close cousin of the Chebyshev semi-iterative method, but based on the Chebyshev polynomials of the fourth instead of first kind. A full V-cycle bound for general polynomial smoothers is derived using the V-cycle theory of McCormick. The fourth-kind Chebyshev iteration is quasi-optimal for the V-cycle bound. The optimal polynomials for the V-cycle bound can be found numerically, achieving about an 18% lower error contraction factor bound than the fourth-kind Chebyshev iteration, asymptotically as the number of smoothing steps $k \rightarrow \infty$. Implementation of the optimized iteration is discussed, and the performance of the polynomial smoothers is illustrated with numerical examples.

1 Introduction

Given a multigrid method for solving a linear system that uses simple linear iterations for smoothers, a natural question is whether Krylov (i.e., polynomial) acceleration techniques can be applied to improve the

smoother iterations. General-purpose Krylov methods don't necessarily have good smoothing properties, since they are typically designed to optimize some norm of the total error and not just its "un-smooth" components. In particular, a method like conjugate gradients can dramatically worsen the performance of a smoother iteration on the "high-frequency" error components in return for a marginal improvement in the performance on the "smooth" error components, which should of course be left for the coarse grid correction to handle.

There are, however, polynomial methods specifically designed to be used as smoothers [1, 2, 7]. One popular strategy for symmetric positive definite (SPD) systems, suggested by Adams, Brezina, Hu, and Tuminaro [1], is to use the standard Chebyshev semi-iteration, but adjusted to target only an upper portion $[\lambda^*, \lambda_{\max}]$ of the range of eigenvalues $[\lambda_{\min}, \lambda_{\max}]$, with λ^* an ad-hoc parameter. Another possibility, suggested by Kraus, Vassilevski, and Zikatanov [7] is to use the best polynomial approximation to λ^{-1} in the uniform norm, again over the interval $[\lambda^*, \lambda_{\max}]$. We will see that we can design methods that avoid such ad-hoc parameters as λ^* , by picking polynomials that optimize multigrid convergence bounds. We will restrict our attention to SPD systems and single-step SPD smoothers (since polynomial methods tend to be most effective in the SPD setting), and examine both a two-level bound for general polynomial smoothers going back to Hackbusch [5], as well as develop a full V-cycle convergence bound for general polynomial smoothers starting from the V-cycle theory developed by McCormick [10]. We will see that a very simple Chebyshev-type iteration, which avoids the need for any ad-hoc parameter, optimizes the two-level bound and is quasi-optimal for the V-cycle bound. We will also see that a small modification to the iteration, though somewhat complicated to implement, optimizes the V-cycle bound. The effectiveness (or not) of these smoother acceleration techniques are demonstrated with numerical examples.

To fix notation, let A be a symmetric positive definite matrix, and consider the "coarse-to-fine" half V-cycle multigrid iteration $x_{n+1} = x_n + M(b - Ax_n)$ determined by the iteration matrix

$$E = I - MA = GK$$

where G is an iteration matrix for the smoother, e.g., $G = (I - BA)^k$ for k steps of the iteration $x_{i+1} = x_i + B(b - Ax_i)$, and K is the coarse grid correction

$$K := I - PM_c P^* A,$$

where the prolongation operator P is an injection from a space of smaller dimension, and where M_c is defined recursively in the same way as M , but for the system governed by the Galerkin coarse operator

$$A_c := P^*AP,$$

and with the recursion stopping by taking $M_c = A_c^{-1}$ at the coarsest level.

We will restrict our attention primarily to bounds of $\|E\|_A^2$, the square of the contraction factor in the energy norm of this half V-cycle. As is well-known (see, e.g., [10]), this suffices for the analysis of general V-cycles. If we denote the energy inner product adjoint for an operator T by T^\dagger , so that $AT^\dagger = T^*A$, i.e., $T^\dagger := A^{-1}T^*A$, then E^\dagger is the iteration matrix for a “fine-to-coarse” half V-cycle iteration, with equal norm $\|E^\dagger\|_A = \|E\|_A$. A general V-cycle E_V decomposes as some “fine-to-coarse” half V-cycle $(E')^\dagger$ followed by some “coarse-to-fine” half V-cycle E , $E_V = E(E')^\dagger$, so that

$$\|E_V\| \leq \|E\|_A \|E'\|_A.$$

For a symmetric V-cycle with $E' = E$ this inequality strengthens to the equality $\|E_V\|_A = \|E\|_A^2$.

Instead of using the simple smoothing iteration $G = (I - BA)^k$ we will consider polynomial iterations

$$G = p_k(BA)$$

for some polynomial $p_k(x)$ of degree k with $p_k(0) = 1$. Picking p_k to optimize $\|G\|_A$ leads to Krylov acceleration methods like the Chebyshev semi-iterative method and the method of conjugate gradients. Our central concern will be instead to investigate how to choose p_k to optimize the multigrid convergence rate, through achieving optimal or nearly optimal bounds on $\|E\|_A^2$. As polynomial methods tend to work best for symmetric iterations, we will confine our attention to the case of symmetric positive definite B . For example, Jacobi corresponds to $B = D^{-1}$ where D is the diagonal part of A . Furthermore, so as to remove some redundancy, we will assume B is scaled such that

$$\rho(BA) \leq 1$$

where ρ denotes the spectral radius. In particular, these assumptions imply that the single step ($k = 1$) smoother

$$G = I - \omega BA$$

has contraction factor $\|G\|_A = \rho(G) < 1$ whenever the damping parameter ω is in the range $0 < \omega < 2$.

Our results will all be expressed in terms of the approximation property

$$C := \|A^{-1} - PA_c^{-1}P^*\|_{A,B}^2 := \sup_{\|f\|_B \leq 1} \|(A^{-1} - PA_c^{-1}P^*)f\|_A^2. \quad (1)$$

Note that a different approximation property constant C is associated with each level in the V-cycle. Hackbusch [5] proved that the error contraction of the V-cycle $\|E\|_A^2$ is bounded by the largest value of

$$\frac{C}{C + 2k} \quad (2)$$

across all levels for the case $G = (1 - BA)^k$, i.e., with $\omega = 1$. Using as our main tool the general V-cycle theory developed by McCormick [10], we will generalize this to a bound depending on C and p_k . We will show that using a simple Chebyshev-type iteration for the smoother, based on a three-term recurrence, results in the V-cycle error contraction $\|E\|_A^2$ being bounded by the largest value of

$$\frac{C}{C + \frac{4}{3}k(k+1)} \quad (3)$$

across all levels.

The polynomials associated with this Chebyshev-type iteration are the lesser-known Chebyshev polynomials of the *third* and *fourth* kinds, also known as “airfoil” polynomials within the aerodynamics community, where they have been extensively studied [4, 8, 9]. Within the multigrid community, they have found extensive use in smoothed aggregation methods, originally as the optimal polynomials for constructing the smoothing operator to apply to some tentative prolongator, but also used for the smoothing iteration itself. The first appearance traces back to the thesis of Brezina [3], with continued development in papers by Vaněk, Brezina, Tezaur, and Krížková [11, 13, 14], particularly [12]. The use of these polynomials for smoothing iterations outside of the context of smoothed aggregation was suggested by Adams, Brezina, Hu, and Tuminaro [1], though at the time, the authors were apparently only aware of the existence of these polynomials for specific degrees (1, 4, and 13 being the first three). As mentioned above, Adams et al. also suggested the use of the usual Chebyshev semi-iteration based on the

polynomials of the first kind, adjusted to target the range of eigenvalues $[\lambda^*, \lambda_{\max}]$ instead of $[\lambda_{\min}, \lambda_{\max}]$, with λ^* an ad-hoc parameter.

Our main contribution is to point out that there is a very simple iteration based on the three-term recurrence satisfied by the fourth-kind Chebyshev polynomials, namely

$$\begin{aligned} x_k &= x_{k-1} + z_k, \\ z_0 &= 0, \\ z_k &= \frac{2k-3}{2k+1}z_{k-1} + \frac{8k-4}{2k+1} \frac{1}{\rho(BA)} B(b - Ax_{k-1}), \end{aligned}$$

and to give full multi-level V-cycle convergence bounds, particularly (3), when this iteration is used as the smoother. Note that the iteration is just as simple, if not simpler, than the standard (first-kind) Chebyshev semi-iterative method, and for use as a smoother avoids the ad-hoc choice of parameter λ^* . The fourth-kind Chebyshev polynomials and associated iteration deserves more wide-spread recognition in the broader multigrid community, as a simple method for accelerating an arbitrary multigrid V-cycle for SPD systems.

2 Fourth-kind Chebyshev Iteration

Let us first consider the bound due to Hackbusch [5] for the two-level iteration contraction factor (taking $M_c = A_c^{-1}$),

$$\begin{aligned} \|E\|_A &= \|E^\dagger\|_A = \|(I - PA_c^{-1}P^*A)G\|_A \\ &\leq \|A^{-1} - PA_c^{-1}P^*\|_{A,B} \|AG\|_{B,A} \\ &= C^{1/2} \sup_{\lambda \in \sigma(BA)} \lambda^{1/2} |p_k(\lambda)| \\ &\leq C^{1/2} \sup_{0 < \lambda \leq 1} \lambda^{1/2} |p_k(\lambda)|, \end{aligned} \tag{4}$$

where, recall, we have taken the smoothing iteration matrix to be $G = p_k(BA)$, with p_k a polynomial of degree k and $p_k(0) = 1$. This bound has its shortcomings: it is only for a two-level method, and it need not be smaller than one, whereas we always have $\|E\|_A < 1$ provided $\|G\|_A < 1$. We will be considering other bounds without these particular shortcomings later. The reason for considering the bound (4) is that the problem of picking optimal polynomials for it has an elegant solution.

Optimizing the bound (4) involves solving a *weighted* minimax polynomial approximation problem, with weight $w(x) = x^{1/2}$. Luckily, the solution is known. The unique monic polynomials minimizing the maximum of $(1+x)^{1/2}|p(x)|$ on $[-1, 1]$ are $2^{-n}V_n(x)$, where $V_n(x)$ are the Chebyshev polynomials of the *third* kind, see Mason [8, Property 1.1]. Up to scaling, these are Jacobi polynomials with $\alpha = -\frac{1}{2}$, $\beta = \frac{1}{2}$, but, like the Chebyshev polynomials of the first and second kinds, enjoy particularly simple forms of the various special properties of Jacobi polynomials (see [4, 8, 9]). For our purposes, it will be slightly more convenient to use the Chebyshev polynomials of the fourth kind $W_n(x)$, which are essentially the same: $W_n(x) = (-1)^n V_n(-x)$. They are given by

$$W_n(x) = \frac{\sin(n + \frac{1}{2})\theta}{\sin \frac{1}{2}\theta} \quad (5)$$

where $x = \cos \theta$. Let us map $x \in [-1, 1]$ to $\lambda \in [0, 1]$ via $x = 1 - 2\lambda$. The half angle formula for sine gives $\sin \frac{1}{2}\theta = \lambda^{1/2}$, and we see that the weighted polynomials

$$\lambda^{1/2}W_n(1 - 2\lambda) = \sin(n + \frac{1}{2})\theta \quad (6)$$

equioscillate between ± 1 . To enforce the constraint $p_k(0) = 1$ we need to scale by the inverse of $W_n(1) = 2n + 1$, giving

$$p_k(\lambda) = \frac{1}{2k + 1}W_k(1 - 2\lambda). \quad (7)$$

Using (6) we see that

$$\sup_{0 \leq \lambda \leq 1} \lambda^{1/2}|p_k(\lambda)| = \frac{1}{2k + 1}. \quad (8)$$

The first few $p_k(\lambda)$ are shown in Figure 1a and the corresponding weighted polynomials $\lambda^{1/2}p_k(\lambda)$, exhibiting the equioscillation property, in Figure 1b. Figure 1c anticipates the multilevel bound of the next section. Explicitly, the first four polynomials are

$$\begin{aligned} p_0(\lambda) &= 1, \\ p_1(\lambda) &= 1 - \frac{4}{3}\lambda, \\ p_2(\lambda) &= 1 - 4\lambda + \frac{16}{5}\lambda^2, \\ p_3(\lambda) &= 1 - 8\lambda + 16\lambda^2 - \frac{64}{7}\lambda^3. \end{aligned}$$

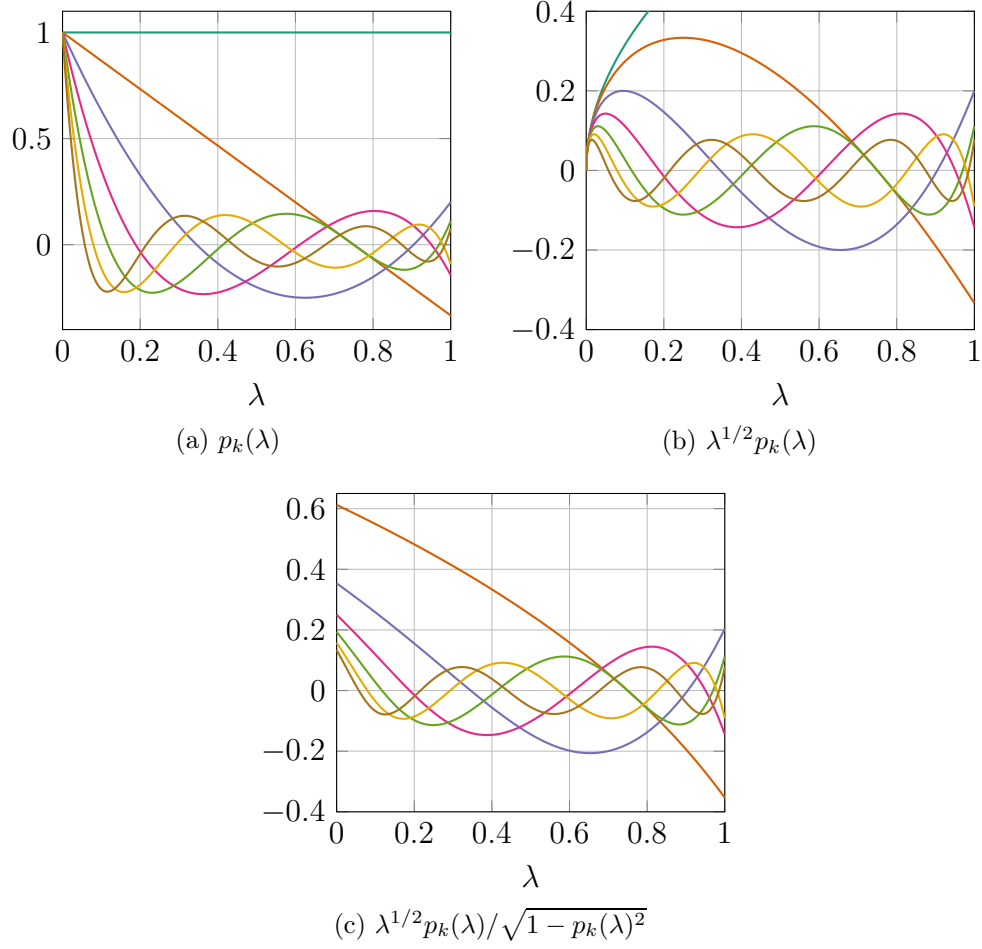


Figure 1: The first few scaled and shifted Chebyshev polynomials of the fourth kind, (a) unweighted, (b) weighted by $\lambda^{1/2}$, (c) weighted by $\sqrt{\lambda/(1 - p_k(\lambda)^2)}$.

The polynomials W_n satisfy the same recurrence as the Chebyshev polynomials T_n , U_n of the first and second kinds, just with a different initial condition [8].

$$W_n(x) = 2xW_{n-1}(x) - W_{n-2}(x), \quad W_0(x) = 1, W_1(x) = 2x + 1.$$

It follows that the scaled and shifted polynomials p_k satisfy the recurrence

$$p_k(\lambda) = \frac{4k-2}{2k+1}(1-2\lambda)p_{k-1}(\lambda) - \frac{2k-3}{2k+1}p_{k-2}(\lambda).$$

We can derive an iteration based on this recurrence, exactly as is done to derive the Chebyshev semi-iterative method. The result is the iteration

$$\begin{aligned} x_k &= x_{k-1} + z_k, \\ z_0 &= 0, \\ z_k &= \frac{2k-3}{2k+1}z_{k-1} + \frac{8k-4}{2k+1} \frac{1}{\rho(BA)} B(b - Ax_{k-1}), \end{aligned} \tag{9}$$

satisfying $A^{-1}b - x_k = p_k(BA)(A^{-1}b - x_0)$. For the sake of easy reference, we have included the scaling of B in (9) so that it works without our assumption $\rho(BA) \leq 1$. More generally, B can be scaled by any upper bound of $\rho(BA)$, with some degradation in performance. The bound

$$\|Ap_k(BA)\|_{B,A} \leq \frac{1}{2k+1} \tag{10}$$

follows from (8), leading to the *two-level* contraction factor estimate

$$\|E\|_A \leq \frac{C^{1/2}}{2k+1}. \tag{11}$$

3 Multilevel Analysis

We will base our multilevel analysis on the general V-cycle theory developed by McCormick [10]. While it is possible that an interesting analysis of polynomial smoothers could be based on the XZ identity [17], which considers all levels all at once, the main aim of our analysis will be to isolate the effect of the choice of polynomial from the choice of single-step smoother, so that isolating the different levels from each other is a natural first step. In retrospect, the McCormick theory can be seen as a (optimal in some sense) means of isolating the levels, making it a natural starting point for our analysis. The central idea

of the theory is to decompose the solution space $V = \mathbb{R}^n$ into the direct sum of the coarse space (the range of P) and its A -orthogonal complement,

$$V = R(P) \oplus R(P)^{\perp_A}.$$

Let us denote the pair of projections corresponding to this decomposition by

$$\pi_c := PA_c^{-1}P^*A, \quad \text{and} \quad \pi_f := I - \pi_c.$$

Note that $\pi_f = I - PA_c^{-1}P^*A$ is the coarse grid correction with an exact coarse grid solve, and that the approximation property constant C defined by (1) can be written as $C = \|\pi_f A^{-1}\|_{A,B}^2$. Alternately, we can write C as

$$C = \sup_{u \in R(\pi_f) \setminus \{0\}} \frac{\|u\|_{B^{-1}}^2}{\|u\|_A^2} \geq 1. \quad (12)$$

To see this, take any A -orthogonal basis of $R(P)^{\perp_A} = R(\pi_f)$ assembled as the columns of a matrix Q so that $Q^*AQ = I$ and $R(Q) = R(\pi_f)$. Then $\pi_f = QQ^*A$ and

$$\begin{aligned} C &= \|\pi_f A^{-1}\|_{A,B}^2 = \rho(B^{-1}A^{-1}\pi_f^*A\pi_f A^{-1}) \\ &= \rho(B^{-1}QQ^*) = \rho(Q^*B^{-1}Q) \\ &= \sup_{v \neq 0} \frac{\|Qv\|_{B^{-1}}^2}{\|v\|_2^2} = \sup_{v \neq 0} \frac{\|Qv\|_{B^{-1}}^2}{\|Qv\|_A^2} = \sup_{u \in R(Q) \setminus \{0\}} \frac{\|u\|_{B^{-1}}^2}{\|u\|_A^2}. \end{aligned}$$

The lower bound of 1 follows from our assumption $\rho(BA) \leq 1$.

The McCormick bound involves the approximation constant associated with the *symmetrized smoother* N , defined so that the corresponding iteration matrix is

$$I - NA = G^\dagger G = A^{-1}G^*AG, \quad (13)$$

i.e., the smoother followed by its adjoint. Explicitly, $N := A^{-1} - A^{-1}G^*AGA^{-1}$ and, provided $\|G\|_A < 1$, we see that

$$0 < N \leq A^{-1}. \quad (14)$$

That is, N is SPD and $A^{-1} - N$ is symmetric positive semidefinite.

Lemma 3.1 (McCormick V-cycle bound). *If the smoother on each level is a contraction in the energy norm, $\|G\|_A < 1$, then the symmetric V-cycle contraction factor is bounded by*

$$\|E\|_A^2 \leq 1 - C_N^{-1}$$

where C_N is the largest value across all levels of the approximation property constant

$$\|\pi_f A^{-1}\|_{A,N}^2 = \sup_{u \in \mathbf{R}(\pi_f) \setminus \{0\}} \frac{\|u\|_{N^{-1}}^2}{\|u\|_A^2} \geq 1$$

defined in terms of the symmetrized smoother $N := A^{-1} - A^{-1}G^*AGA^{-1}$.

This is essentially McCormick's Lemma 2.3 [10], except that we have included the alternate expression analogous to (12) for C_N (corresponding to the reciprocal of McCormick's β), which anticipates some of the ideas from McCormick's Theorem 3.4. In this case, the lower bound of 1 for C_N follows from (14). McCormick's V-cycle bound in the same form can also be found in Vassilevski [15, p. 146] as Theorem 5.16 together with Lemma 5.17.

It is now fairly straight-forward to prove the following V-cycle bound for our main case of interest, polynomial smoothers $G = p(BA)$. Like the two-level bound (4), this bound only depends on $C = \|\pi_f A^{-1}\|_{A,B}^2$ and p . The proof is quite similar to that of Theorem 3.4 of McCormick [10].

Lemma 3.2. *Let the smoother iteration (on each level) be given by*

$$G = p(BA)$$

where B is symmetric positive definite and scaled such that $\rho(BA) \leq 1$, and $p(x)$ is a polynomial satisfying $p(0) = 1$ and $|p(x)| < 1$ for $0 < x \leq 1$, possibly different on each level. Then the V-cycle contraction factor $\|E\|_A^2$ is bounded by the largest value across all levels of

$$\frac{C}{C + \gamma^{-1}}$$

where C is the approximation property constant $\|\pi_f A^{-1}\|_{A,B}^2$ and

$$\gamma = \sup_{0 < \lambda \leq 1} \frac{\lambda p(\lambda)^2}{1 - p(\lambda)^2}. \quad (15)$$

Proof. We have

$$\begin{aligned} C_N &= \sup_{u \in \mathbf{R}(\pi_f) \setminus \{0\}} \frac{\|u\|_{N^{-1}}^2}{\|u\|_A^2} = 1 + \sup_{u \in \mathbf{R}(\pi_f) \setminus \{0\}} \frac{\|u\|_{N^{-1}}^2 - \|u\|_A^2}{\|u\|_A^2} \\ &\leq 1 + \left(\sup_{u \neq 0} \frac{\|u\|_{N^{-1}}^2 - \|u\|_A^2}{\|u\|_{B^{-1}}^2} \right) \left(\sup_{u \in \mathbf{R}(\pi_f) \setminus \{0\}} \frac{\|u\|_{B^{-1}}^2}{\|u\|_A^2} \right) \\ &\leq 1 + \gamma C \end{aligned}$$

since

$$\begin{aligned}
\sup_{u \neq 0} \frac{\|u\|_{N^{-1}}^2 - \|u\|_A^2}{\|u\|_{B^{-1}}^2} &= \rho(B[N^{-1} - A]) \\
&= \rho(BA[(I - p(BA)^2)^{-1} - I]) \\
&= \sup_{\lambda \in \sigma(BA)} \lambda [(1 - p(\lambda)^2)^{-1} - 1] \leq \gamma
\end{aligned}$$

where we have used that $N^{-1} = A(I - G^2)^{-1}$, which follows from (13) as $G^\dagger = G$. The result now follows by applying the previous lemma. \square

As a quick sanity check, we note that, as a direct consequence, this lemma recovers Hackbusch's bound (2) on the contraction factor for the V-cycle with a simple iteration.

Corollary 3.3. *Let B be SPD and scaled such that $\rho(BA) \leq 1$ (on each level). If the smoother is k steps of a simple damped iteration,*

$$G = (I - \omega BA)^k,$$

for possibly different $0 < \omega < 2$ and k on each level, then $\|E\|_A^2$ is bounded by the largest value of

$$\frac{C}{C + 2\omega k}$$

provided

$$(1 - \omega)^{2k} \leq \frac{1}{1 + 2\omega k}. \quad (16)$$

In particular, $\omega \leq \frac{3}{2}$ is necessary for $k = 1$ and sufficient for all k . Asymptotically, the largest permissible ω for the bound to hold is

$$\omega = 2 - \frac{\log 4k}{2k} + o(k^{-1}) \quad \text{as } k \rightarrow \infty.$$

Proof. Here $p(\lambda) = (1 - \omega\lambda)^k$ and $f(\lambda) := \lambda p(\lambda)^2 / (1 - p(\lambda)^2)$ is convex on $[0, 1]$, so that the supremum of f is attained at either $\lambda = 1$ where $f(1)^{-1} = (1 - \omega)^{-2k} - 1$ or (in the limit) at $\lambda = 0$, where $\lim_{\lambda \rightarrow 0} f(\lambda)^{-1} = 2\omega k$. The latter is smaller when (16) holds. \square

3.1 Fourth-kind Chebyshev Iteration

The bound for the V-cycle using the fourth-kind Chebyshev iteration as smoother also follows directly from Lemma 3.2. As illustrated in Figure 1c, the supremum determining γ for the scaled and shifted Chebyshev polynomials of the fourth kind is attained in the limit $\lambda \rightarrow 0$. We prove this rigorously in Appendix A. The following bound then follows by an application of L'Hôpital's rule.

Corollary 3.4. *Let B be SPD and scaled such that $\rho(BA) \leq 1$ (on each level). If (9) is used for the smoother iteration, so that*

$$G = p_k(BA),$$

where p_k are the scaled and shifted Chebyshev polynomials of the fourth kind given by (7), then $\|E\|_A^2$ is bounded by the largest value of

$$\frac{C}{C + \frac{4}{3}k(k+1)}. \quad (17)$$

Proof. See Appendix A. □

3.2 Optimized Polynomials

The scaled and shifted Chebyshev polynomials of the fourth kind (7) are not optimal for the multilevel bound of Lemma 3.2, as they minimize the supremum of $\lambda p(\lambda)^2$ and not $\lambda p(\lambda)^2 / (1 - p(\lambda)^2)$. The optimal polynomials for the multilevel bound can be found numerically (see Appendix B for details). Explicitly, the first two are

$$\begin{aligned} p_1(\lambda) &= 1 - \frac{3}{2}\lambda, \\ p_2(\lambda) &= (1 - \frac{\sqrt{5}}{2}\lambda)(1 - \frac{5+\sqrt{5}}{2}\lambda) = 1 - (\frac{5}{2} + \sqrt{5})\lambda + \frac{5}{4}(1 + \sqrt{5})\lambda^2. \end{aligned}$$

The first few optimal polynomials are shown in Figure 2a, and they can be seen to exhibit the equioscillation property in Figure 2b. Numerical evidence (see Table 1) supports the conjecture that, asymptotically, these polynomials achieve

$$\gamma_{\text{opt}}^{-1} \sim 4 \left(\frac{2k+1}{\pi} \right)^2 - \frac{2}{3} + \frac{1}{60} \left(\frac{2k+1}{\pi} \right)^{-2} + O(k^{-4}) \quad \text{as } k \rightarrow \infty,$$

with, moreover, the first two terms of the series giving a fairly good lower bound of γ_{opt}^{-1} . Assuming this conjecture holds, we have the following V-cycle convergence bound.




Figure 2: The optimal polynomials p_k for the bound of Lemma 3.2, (a) un-weighted, (b) weighted by $\sqrt{\lambda/(1-p_k(\lambda)^2)}$.

Conjecture 3.5. *Let B be SPD and scaled such that $\rho(BA) \leq 1$ (on each level). Let the smoother on each level be given by $G = p_k(BA)$ (for possibly different k) where p_k are the optimal polynomials for Lemma 3.2,*

$$p_k = \arg \min_{p \in \mathbb{P}_k, p(0)=1} \sup_{0 < \lambda \leq 1} \frac{\lambda p(\lambda)^2}{1 - p(\lambda)^2}.$$

Here \mathbb{P}_k denotes the space of real polynomials of degree at most k . Then $\|E\|_A^2$ is bounded by the largest value of

$$\frac{C}{C + \frac{4}{\pi^2}(2k+1)^2 - \frac{2}{3}} \quad (18)$$

across all levels.

As hinted at in Table 1, we have verified the conjecture numerically for all practical values of k (into the hundreds). In the limit $k \rightarrow \infty$, the conjectured bound (18) for the optimized polynomials is an improvement over the bound (17) for the fourth-kind Chebyshev iteration by the factor $\frac{\pi^2}{12} \approx 0.822$, i.e., a reduction in the contraction factor bound by about 18%.

k	γ_{opt}^{-1}	$\frac{4}{\pi^2}(2k+1)^2 - \frac{2}{3}$	Difference	$\frac{\pi^2}{60}(2k+1)^{-2}$
1	3.	2.981	1.91×10^{-2}	1.83×10^{-2}
2	9.4721	9.465	6.68×10^{-3}	6.58×10^{-3}
3	19.1957	19.1923	3.38×10^{-3}	3.36×10^{-3}
4	32.1634	32.1614	2.04×10^{-3}	2.03×10^{-3}
5	48.3742	48.3728	1.36×10^{-3}	1.36×10^{-3}
10	178.0643	178.0639	3.73×10^{-4}	3.73×10^{-4}
100	16373.241899	16373.241895	4.07×10^{-6}	4.07×10^{-6}
1000	1.623×10^6	1.623×10^6	4.11×10^{-8}	4.11×10^{-8}

Table 1: Optimal γ^{-1} for Lemma 3.2.

4 Implementation Issues

The iteration (9) is the natural implementation for a smoother using the Chebyshev polynomials of the fourth kind, derived from the three-term recurrence they satisfy. The numerically determined optimal polynomials for the V-cycle bound of Lemma 3.2 do not satisfy a three-term recurrence, and so require a different implementation.

One obvious approach is to translate the factorization into monomials involving the roots r_i ,

$$p_k(\lambda) = \prod_{i=1}^k \left(1 - \frac{\lambda}{r_i}\right),$$

directly into the iteration

$$x_{i+1} = x_i + \frac{1}{r_i} B(b - Ax_i).$$

While probably fine for small k , a concern of this approach is that the polynomials corresponding to the intermediate steps are poorly behaved. Depending on the chosen ordering of roots, the iteration may first dampen some error modes by several orders of magnitudes before then amplifying them by several orders of magnitude, or vice versa.

Instead, we propose using an iteration based on the expansion in the basis of fourth-kind Chebyshev polynomials

$$p_k(\lambda) = \sum_{i=0}^k \alpha_i W_i(1 - 2\lambda),$$

exploiting the fact that we expect the orthogonal (w.r.t. a weight function) Chebyshev polynomials of the fourth kind W_i to be a good basis for representing the optimal p_k and the fact that we have a nice iteration, namely (9), for computing the basis W_i . The Gaussian quadrature rule associated with the polynomials W_i can be used to determine the coefficients α_i , requiring only that p_k be evaluated at the quadrature nodes (see Appendix B for details). We propose then the iteration

$$\begin{aligned} z_0 &= 0, \quad r_0 = b - Ax_0, \\ z_i &= \frac{2i-3}{2i+1}z_{i-1} + \frac{8i-4}{2i+1} \frac{1}{\rho(BA)} Br_{i-1}, \\ x_i &= x_{i-1} + \beta_i z_i, \quad r_i = r_{i-1} - Az_i. \end{aligned} \tag{19}$$

This is equivalent to (9) when we take $\beta_i = 1$. Note that β_i does not appear in the “residual” update, so that it is *not* the case that $r_i = b - Ax_i$. Instead, r_i tracks the residual of the unmodified fourth-kind Chebyshev iteration, and the vectors z_i are identical in both iterations. From the properties of (9), we know that $z_i = [\frac{1}{2i-1}W_{i-1}(BA) - \frac{1}{2i+1}W_i(BA)]e_0$ where $e_0 = A^{-1}b - x_0$. It follows that the polynomial $p_k(\lambda)$ corresponding to the iteration (19), i.e., so that $A^{-1}b - x_k = p_k(BA)e_0$, is given by

$$p_k(\lambda) = \sum_{i=0}^k \frac{\beta_i - \beta_{i+1}}{2i+1} W_i(1-2\lambda)$$

provided we let $\beta_0 = 1$ and $\beta_{k+1} = 0$. Hence, the coefficients β_i can be found from the coefficients α_i via the recursion

$$\beta_{i+1} = \beta_i - (2i+1)\alpha_i.$$

For the polynomials that optimize the multilevel bound (3.2), the resulting β_i (which depend on k) are in the interval $[1, 1.6)$, for k up to several hundred. We can thus think of the optimized iteration (19) as an “over-relaxed” version of the fourth-kind Chebyshev iteration (9). Explicit Matlab code for computing the coefficients β_i is given in Appendix B.

5 Numerical Results

We examine the effectiveness of the considered polynomial accelerations of smoothers by looking at three numerical examples, all variants

of geometric multigrid applied to a 2D Poisson problem discretized by the finite element method (FEM). In each case, the problem to be solved consists of a 1024×1024 grid of rectangular bilinear finite elements, and we consider diagonal single-step smoothers (e.g., damped Jacobi). We use simple bilinear interpolation in every case and use the Galerkin construction for the coarse problems (which is identical to the direct discretization of the coarse problem). For the first example, we take the grid spacing to be uniform, the diffusion coefficient to be constant, while varying the aspect ratio $\Delta y/\Delta x$, which has the effect of varying the approximation constant C . For our second and third examples, we group the elements into a grid of square “macroelements” each consisting of N by N finite elements, resulting in $1024/N$ macroelements in each direction. In the second example, we consider coefficient jumps, alternating the diffusion coefficients of the macroelements in a checkerboard pattern. For our third example, we consider instead nonuniform mesh spacing, taking the nodes within each macroelement to be given by an affine mapping of the Chebyshev nodes. This last example is closely related to low-order preconditioners for high-order elements (considered, e.g., in [6]). For all examples, we will primarily be looking at the asymptotic symmetric V-cycle error contraction factor $\|E\|_A^2$, measured by applying the symmetric V-cycle in a generalized Lanczos iteration [16].

5.1 Uniform grid with varying aspect ratio

We first consider the case of a uniform mesh with constant diffusion coefficient and varying aspect ratio, with Dirichlet boundary conditions on each boundary. In practice, the degradation of performance associated with high aspect ratios would probably be better addressed with a different multigrid strategy such as semi-coarsening. Here we use the aspect ratio to control the approximation property constant C , for a problem amenable to relatively simple analysis.

Figure 3 illustrates the performance of the fourth-kind Chebyshev iteration “cheb4” (9), the optimized iteration “opt”, and the simple iteration with damping parameters $\omega = 4/3$ and $\omega = 3/2$ as smoothers in a symmetric multigrid V-cycle. The damping parameters for the simple iterations were chosen so that they coincide with the two polynomial smoothers for $k = 1$ smoothing step. We use (rescaled) Jacobi as the single-step smoother, which in this case is the same as Richardson as the diagonal of A is constant. The approximation constant C

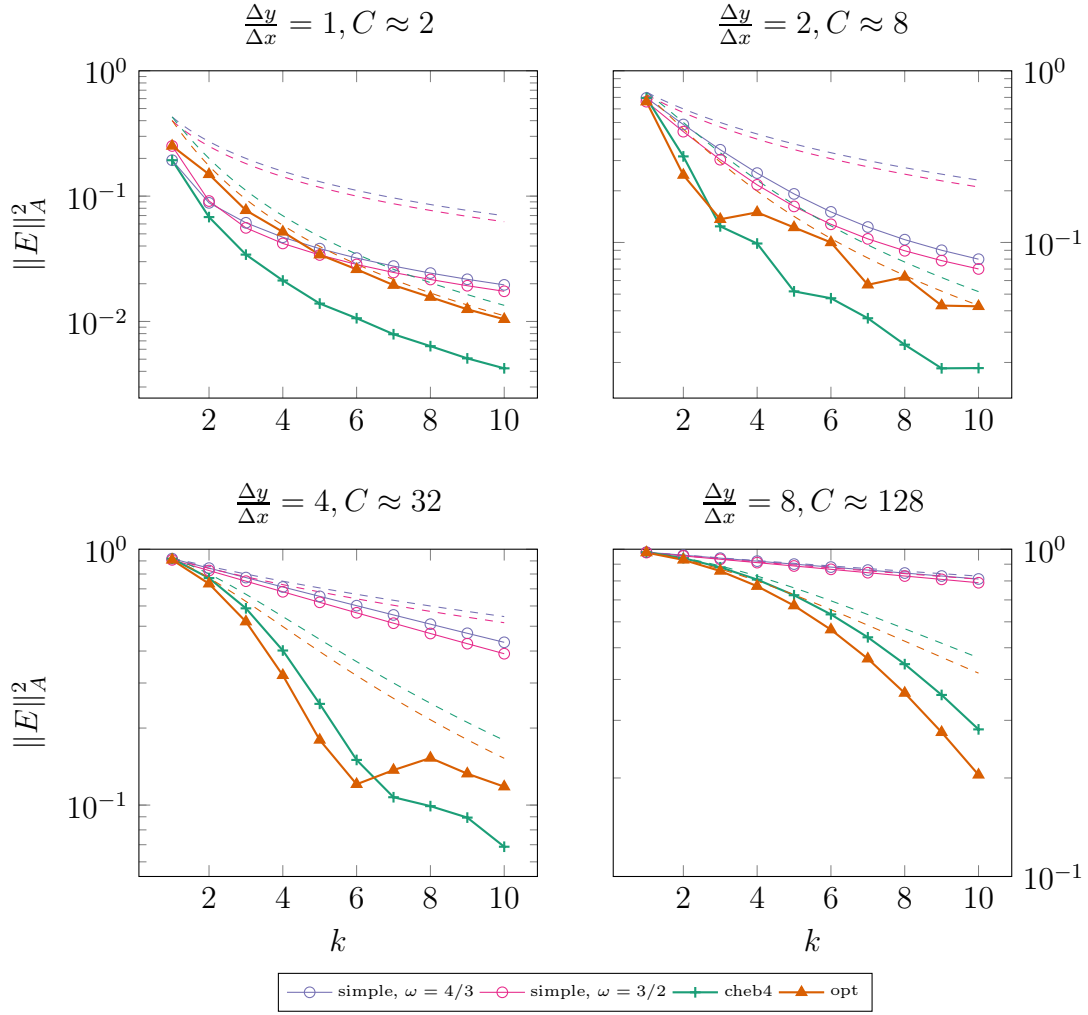


Figure 3: V-Cycle error contraction factors for Poisson discretized on a uniform grid.

for this problem asymptotes to twice the element aspect ratio squared as the grid size increases (as can be seen from a Fourier analysis), and the four subfigures correspond to aspect ratios of 1, 2, 4, and 8, and thus to $C \approx 2, 8, 32, 128$ on the first level. (The values of C taper off on the coarser levels.) The solid lines with markers are the measured (via Lanczos) V-cycle contraction factors for different numbers of smoothing steps k , taken to be the same on each level. For each solid line, the dashed line of the same color is the bound derived from Lemma 3.2, i.e., Corollary 3.3 for the simple iteration, Corollary 3.4 for the Chebyshev iteration, and (18) for the optimized iteration.

In all cases, the fourth-kind Chebyshev iteration significantly outperforms the simple iterations. While there is a significant gap between the measured contraction factors and the corresponding bounds, apart from the optimized iteration the measured contraction factors and bounds qualitatively follow the same trends, for the most part. The optimized iteration performs worse than even the simple iteration for the ideal case of aspect ratio 1 for $k < 5$. However, it does deliver the expected improvement over the fourth-kind Chebyshev iteration as suggested by the improved bound, for C large enough and k small enough. While the performance of the optimized iteration can be relatively poor for small C and/or large k , we observe that it approaches the bound derived from Lemma 3.2 quite closely at some points, showing that the bound can be rather sharp in some cases.

The non-monotonic behavior of some of the curves as k increases is perhaps surprising. As we shall see, this is ultimately due to the fact that the polynomials $p_k(\lambda)$ are not themselves monotonic in k for fixed λ (apart from the simple iterations), together with the fact that the convergence rate is more sensitive to the values of the polynomial at some eigenvalues than at others.

In Figure 4 we compare the fourth-kind Chebyshev and optimized iterations with two other polynomial methods. The first, marked “cheb1,” is the standard Chebyshev iteration based on polynomials of the first kind, targeting the range of eigenvalues $[\kappa^{-1}, 1]$ for $\kappa = 3, 10, 30, 100, 300$ (recall that we have scaled B so that $\rho(BA) = 1$). The second, marked “ufm,” is the method suggested by Kraus, Vassilevski, and Zikatanov [7], which takes $p_k(x) = 1 - q_{k-1}(x)x$ where $q_{k-1}(x)$ is the best degree $k - 1$ polynomial approximation to x^{-1} in the uniform norm on the interval $[\kappa^{-1}, 1]$. Equivalently, p_k for this

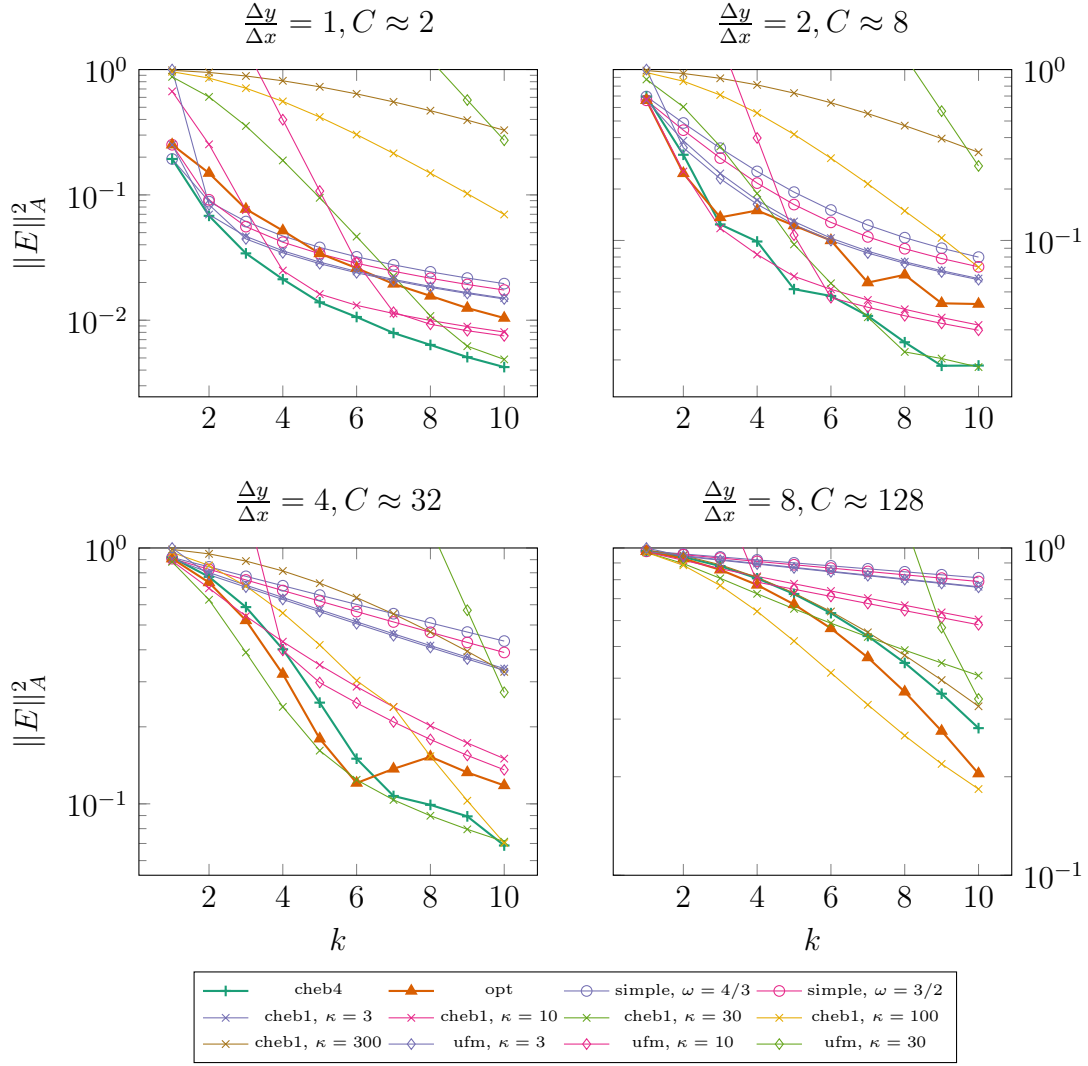


Figure 4: V-Cycle error contraction factors for Poisson discretized on a uniform grid, comparison of different polynomials.

method is characterized as minimizing

$$\sup_{\kappa^{-1} \leq \lambda \leq 1} \lambda^{-1} |p_k(\lambda)|.$$

Recall that the (scaled and shifted) Chebyshev polynomials of the fourth kind are characterized as minimizing a similar supremum, but over the whole interval $[0, 1]$ and with weight $\lambda^{1/2}$ instead of λ^{-1} . The Chebyshev polynomials of the first kind minimize the same supremum but with weight 1. The weight λ^{-1} is somewhat counter-intuitive, in that it places even more emphasis on damping the eigenmodes associated with small eigenvalues than the Chebyshev polynomials of the first kind. It also causes the polynomial iteration by itself to be divergent unless the degree is high enough,

$$k \geq 2 \log \frac{\sqrt{2}}{\sqrt{\kappa} + 1} \bigg/ \log \frac{\sqrt{\kappa} - 1}{\sqrt{\kappa} + 1}.$$

For the considered parameters $\kappa = 3, 10, 30$, we must have $k \geq 2, 4, 9$ respectively for these polynomials to give convergent iterations.

Figure 4 illustrates that the optimal parameter κ for both the “cheb1” and “ufm” cases depends both on the specific problem as well as the number of smoothing steps k . The same parameter κ can be optimal in one circumstance and very far from optimal in another. Compare, for example, the performance of $\kappa = 100$ in the case of aspect ratio 8 to its performance when the aspect ratio is 1. On the other hand, the parameter-free fourth kind Chebyshev and optimized polynomials, while not always optimal, are competitive or at least not very suboptimal in every case. For the “ufm” case, we clearly see the effect of a minimum degree required for convergence. For fixed κ , as k is increased the “ufm” polynomials eventually perform slightly better than the “cheb1” polynomials.

In Figure 5 we consider the McCormick bound $\|E\|_A^2 \leq 1 - C_N^{-1}$ of Lemma 3.1 for a subset of the polynomial smoothers considered in Figure 4, in order to gain more insight into their behavior. The markers correspond to measured values of C_N , determined using a generalized Lanczos iteration with N^{-1} implemented by transforming to the eigenbasis. Many of the general trends in Figure 4 are reflected in the upper bounds plotted in Figure 5, though there are some discrepancies resulting from the variance in the sharpness of the bound.

At first glance the curves in Figure 5 might appear to connect the markers. They are actually plots of $1 - \lambda^{-1}$ where λ is the maximum

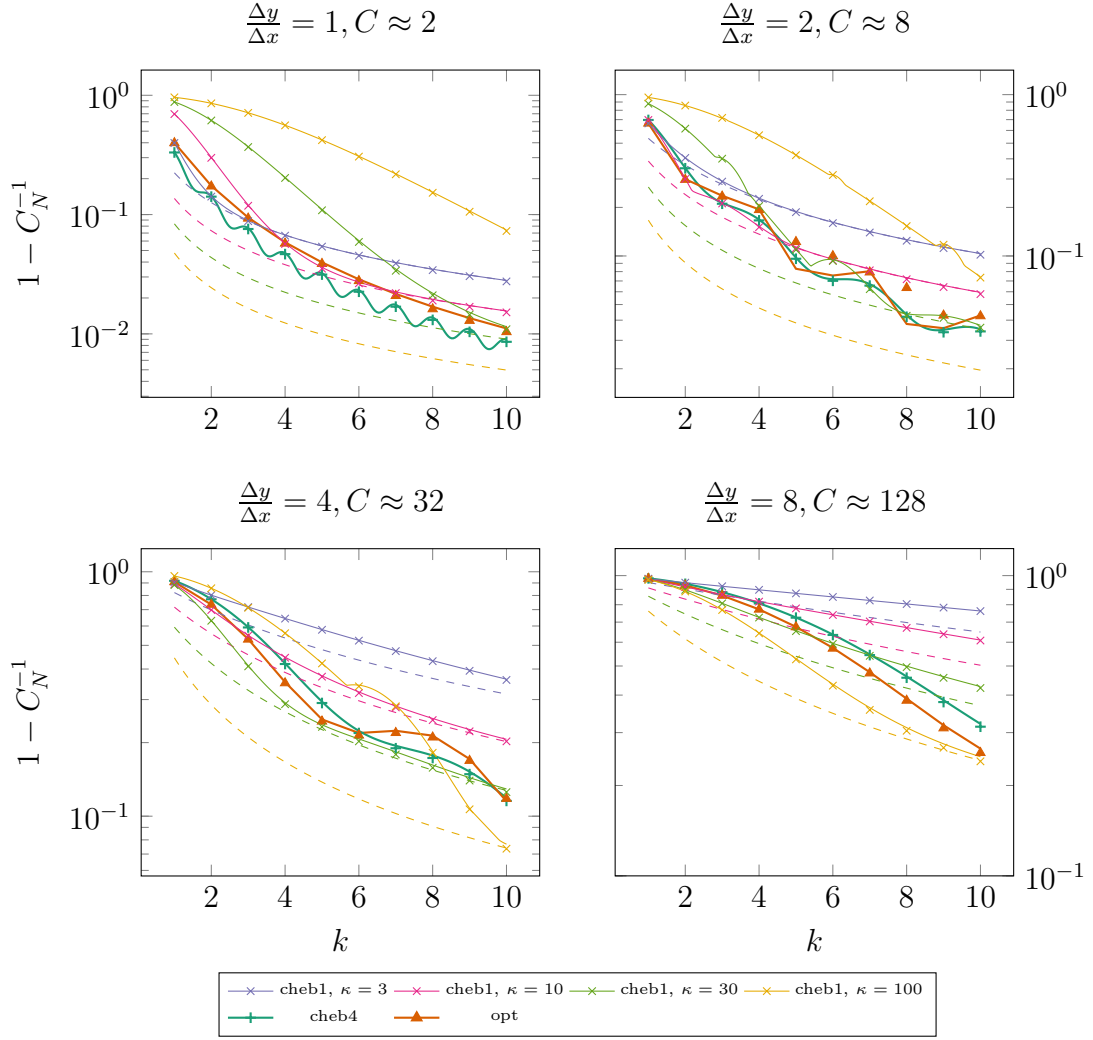


Figure 5: McCormick V-cycle bounds for Poisson discretized on a uniform grid.

of

$$\begin{aligned}\lambda_1 &= 1 - \frac{\alpha^2}{2p'_k(0)} + \frac{p_k(\alpha^{-2})^2}{1 - p_k(\alpha^{-2})^2}, \text{ and} \\ \lambda_2 &= 1 - \frac{1}{2p'_k(0)} + \frac{p_k(1)^2}{1 - p_k(1)^2},\end{aligned}\tag{20}$$

where $\alpha = \Delta y / \Delta x \geq 1$ is the aspect ratio. These are two asymptotic eigenvalues of $N^{-1}\pi_f A^{-1}$ associated with the smallest horizontal and vertical wavenumbers in the limit as the number of grid points goes to infinity, determined by a Fourier analysis. Note that C_N is the spectral radius of $N^{-1}\pi_f A^{-1}$. There's no assurance that the global maximum eigenvalue of $N^{-1}\pi_f A^{-1}$ is associated with the smallest wavenumbers, but the close correspondence of the curves in Figure 5 to the markers indicates this to likely be the case in most of the instances considered. The notable exceptions are the “opt” polynomials at aspect ratio $\alpha = 2$ and degrees $k = 5, 6, 8, 9$. Indeed a close examination of these cases reveals the global maximum eigenvalues to occur at wavenumbers far from the smallest.

To the extent that the McCormick bound reflects the actual multi-grid convergence, we can explain the convergence behavior by examining λ_1 and λ_2 , at least in those cases where their maximum well approximates C_N . First note that the terms involving $p'_k(0)$ and $p_k(1)$ behave monotonically with respect to k , for all the polynomial classes considered. E.g., with any parameter being fixed, $p_k(1)$ always decreases as k increases. On the other hand, the term involving $p_k(\alpha^{-2})$ can exhibit oscillations with k . Much of the interesting and perhaps unexpected behavior can be attributed to this term. Note that we have chosen to interpolate the plots of $1 - \lambda^{-1}$ to non-integer k using analytic expressions where available, in order to highlight the oscillatory behavior. For example, for the Chebyshev polynomials of the fourth kind, $p_k(\alpha^{-2})$ is proportional to $\sin[(k + \frac{1}{2}) \arccos(1 - 2\alpha^{-2})]$ by (5). As a result, the wavelength of the oscillations increases with the aspect ratio α . As another example, for the “cheb1” case with $\kappa = 100$ and aspect ratio $\alpha = 2$, λ is mostly equal to λ_2 except within two narrow intervals near $k = 6$ and $k = 9$ where the peaks of the oscillations of λ_1 stick out, explaining the subtle kinks in the plot of $1 - C_N^{-1}$ at these locations.

For the polynomial smoothers based on the Chebyshev polynomials of the first kind, Figure 4 illustrates the convergence often improves rapidly with k at first before settling into a slower asymptotic regime. This is reflected in the behavior of the McCormick upper bound in

Figure 5, which can be explained as follows. For the Chebyshev polynomials of the first kind,

$$p'_k(0) = k\sqrt{\kappa} \tanh \left[k \operatorname{arccosh} \frac{\kappa+1}{\kappa-1} \right].$$

For fixed κ , as k increases, in particular when it becomes large compared to $1/\operatorname{arccosh} \frac{\kappa+1}{\kappa-1} \sim \sqrt{\kappa}/2$, \tanh approaches 1 exponentially fast so that $p'_k(0) \sim k\sqrt{\kappa}$ plus exponentially small terms (e.s.t.). Meanwhile $p_k(\alpha^{-2})$ and $p_k(1)$ both eventually become exponentially small as k increases, so that $\lambda \sim 1 + \frac{\alpha^2}{2k\sqrt{\kappa}}$ plus e.s.t. Hence, since $C \approx 2\alpha^2$, for fixed κ we expect the McCormick bound $1 - C_N^{-1}$ for the Chebyshev polynomials of the first kind to eventually behave approximately like

$$\frac{C}{C + 4k\sqrt{\kappa}}$$

as k increases. The dashed curves in Figure 5 are plots of this expression, with colors matching the corresponding plots of λ^{-1} . Many of the “cheb1” curves can be seen to fall into this asymptotic behavior; the others can be expected to do so at higher values of k . The asymptotic behavior explains two general features of the McCormick bound reflecting features of the multigrid convergence rates. For a given aspect ratio, the first-kind Chebyshev method with a higher κ always eventually has a better McCormick bound than a lower κ , at high enough k , though the cut-off may be at impractically large k . Secondly, the fourth-kind Chebyshev and optimized polynomials always eventually have lower McCormick bounds than any given first-kind Chebyshev method with fixed κ , again for large enough k , since Corollary 3.4 implies an upper bound for $1 - C_N^{-1}$ for the fourth-kind and optimized polynomials that is $O(k^{-2})$.

In Figure 6, we consider the efficiency of the polynomial smoothers, taking into account the higher costs of using more smoother steps. Since the asymptotic convergence rate of the symmetric V-cycle is $\|E\|_A^2$, asymptotically $-\log_{10}\|E\|_A^2$ symmetric V-cycles are needed per decimal digit (reduction of the error by a factor of 10). Since each V-cycle requires $2k$ applications of the single-step smoother, $-k/\log_{10}\|E\|_A$ thus gives the asymptotic number of smoother steps required per digit. We see in Figure 6 that the optimal number of smoothing steps k increases as the aspect ratio and thereby the approximation constant C increases. At aspect ratio 1, there’s no benefit to using $k > 1$, at least in terms of total smoothing steps, whereas at aspect ratio 8, the

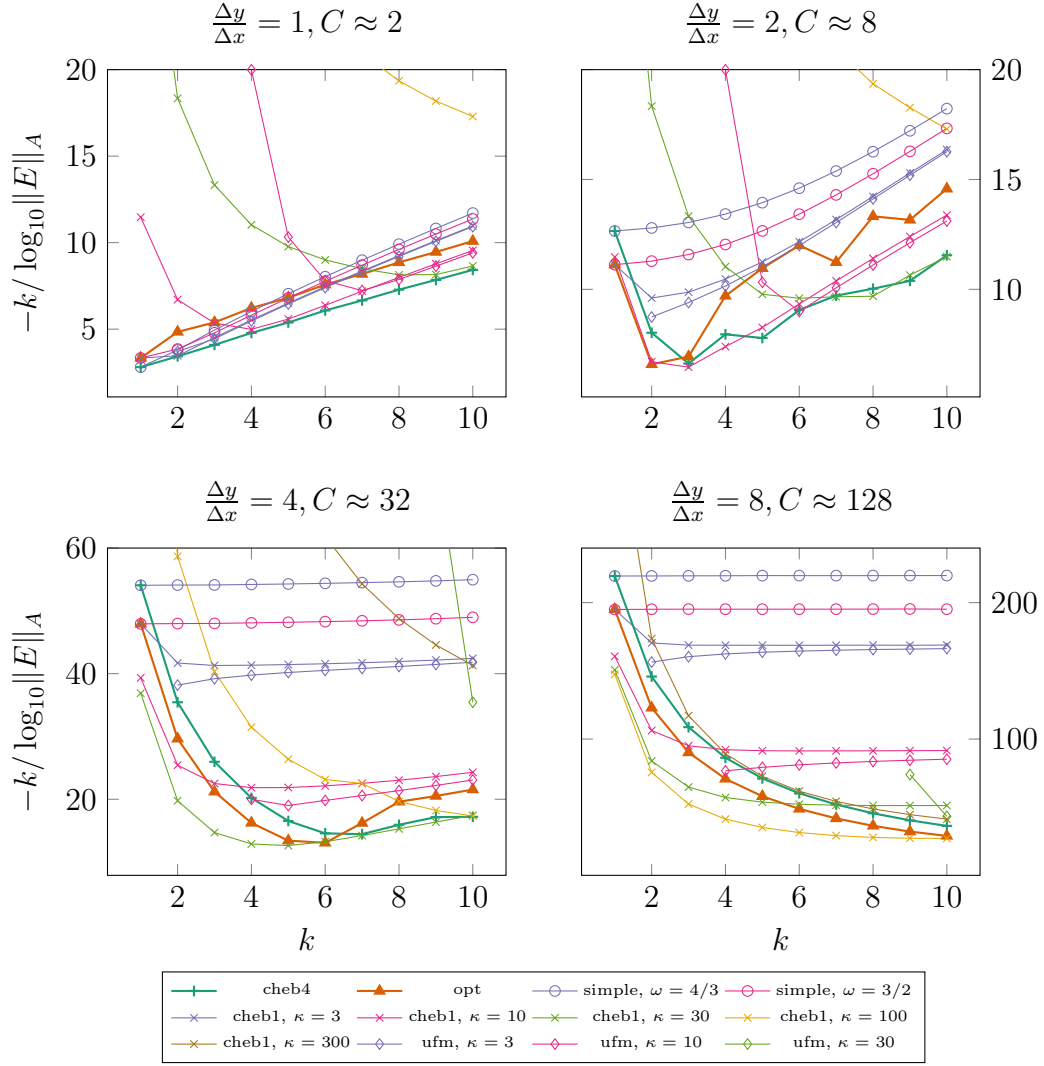


Figure 6: Smoother steps per decimal digit for Poisson discretized on a uniform grid.

optimum appears to be beyond even $k = 10$. The simple smoothers show no benefit to $k > 1$ with this metric in any case. We should note that considering only the total number of smoother steps ignores the costs of the prolongation and restriction operations, which scales with the number of V-cycles, hence further penalizing small k , and making the estimate of optimal k here conservative. Interestingly, while the fourth-kind and optimized polynomials can be somewhat suboptimal at specific k , we note that their minimum values of $-k/\log_{10}\|E\|_A$ are quite competitive for every aspect ratio.

5.2 Coefficient jumps

We next consider a problem involving coefficient jumps. We use the same grid as before, 1024 by 1024 bilinear square finite elements in a uniform grid. We now fix the aspect ratio to be 1, and group the elements into square macroelements each consisting of $N = 32$ by N elements, resulting an array of $M = 32$ by M macroelements. We set the diffusion coefficient to 1 in half of the macroelements and to a parameter a in the other half, arranged in a checkerboard pattern. Note that by taking N to be a power of 2, the usual coarse grids resolve all of the macroelement boundaries (the interfaces of coefficient jumps) until at some level 2 by 2 squares of macroelements are finally merged into a single element (with the Galerkin construction effectively making the diffusion coefficient in the merged element a harmonic average). As a result, as can be seen in Figure 7, the simple multigrid V-cycle with standard bilinear interpolation and damped Jacobi smoothing has a convergence rate that appears to be bounded independently of the problem size (number of macroelements) and the size of the coefficient jump characterized by a . Results for intermediate values of a are similar, with convergence rates appearing to asymptote for large a (making $a = 10^8$ representative of the very large coefficient jump regime).

Strikingly, for this problem, we see very little improvement of the convergence rates when we increase the polynomial degree k , and $k = 1$ is the most efficient in terms of minimizing the total number of smoothing steps. How does this conform, e.g., to the bound of Corollary 3.4 for the fourth-kind Chebyshev smoother, which asymptotically behaves like $\frac{3}{4}Ck^{-2}$? As we see in Table 2, the approximation constant C for this problem, measured using a generalized Lanczos iteration with an embedded multigrid solver, is very large and appears to be unbounded with problem size (the number of macroelements M in each direction),

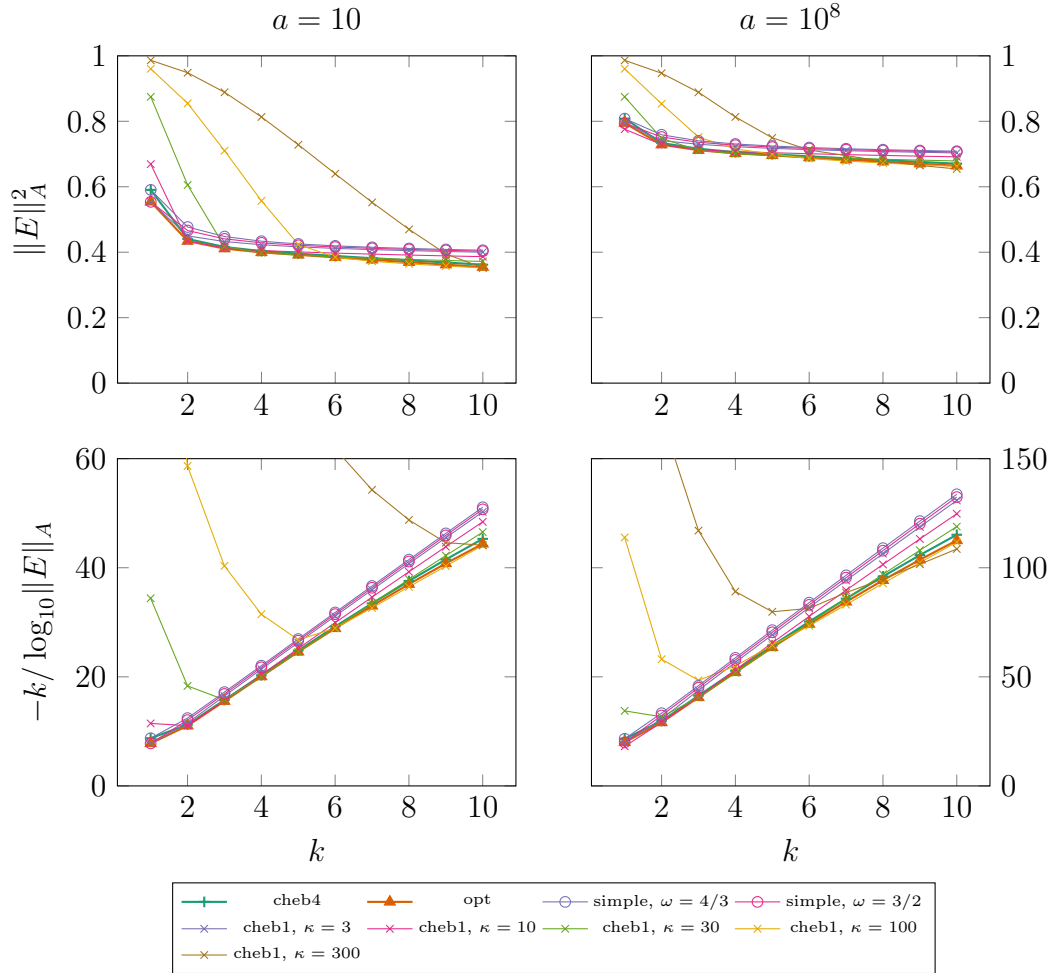


Figure 7: V-Cycle error contraction factors and smoother steps per digit for the discretized Poisson problem with coefficient jumps using (rescaled) Jacobi for the single-step smoother.

N	M	C			
		$a = 10$	$a = 10^2$	$a = 10^4$	$a = 10^8$
2	256	1.77e4	2.97e4	3.16e4	3.16e4
2	512	7.10e4	1.19e5	1.27e5	1.27e5
4	128	8.65e3	2.21e4	2.51e4	2.51e4
4	256	3.48e4	8.91e4	1.01e5	1.01e5
8	64	4.99e3	1.89e4	2.30e4	2.31e4
8	128	2.02e4	7.69e4	9.36e4	9.38e4
16	32	2.93e3	1.68e4	2.21e4	2.22e4
16	64	1.20e4	6.89e4	9.10e4	9.13e4
32	16	1.67e3	1.44e4	2.09e4	2.09e4
32	32	6.98e3	6.09e4	8.80e4	8.84e4

Table 2: Approximation constant C for multigrid on the discretized Poisson problem with coefficient jumps.

though, interestingly, it does appear to be bounded with respect to the parameter a . As such, the bounds based on Lemma 3.2 such as Corollary 3.4 are uselessly pessimistic. On the other hand, we see that having a bounded approximation constant C , while not a necessary condition for bounded multigrid V-cycle convergence, is a useful sufficient condition for the effectiveness of higher degree polynomial smoothers in a V-cycle, in accordance with Lemma 3.2. It is possible that a different multigrid strategy for this problem (e.g., AMG) could lead to bounded constants C , in which case higher degree polynomial smoothers might be effective.

We note that similar results are obtained if ℓ_1 -Jacobi, explained in the next section, is used instead of Jacobi as the single-step smoother.

5.3 Chebyshev spaced grid

For our final example, we again consider Poisson discretized by the FEM on a 1024 by 1024 grid of bilinear rectangular elements. As with our previous example, we group the elements into N by N square macroelements. In this case, instead of altering the diffusion coefficient, we use Chebyshev nodal spacing within each macroelement.




Figure 8: N by N macroelement with Chebyshev grid spacing, $N = 8$.

That is, element boundaries are located at the Chebyshev nodes $x_j = \cos \frac{j}{N}\pi$, $j = 0, \dots, N$, affinely mapped to each macroelement. A single macroelement for the case $N = 8$ is illustrated in Figure 8. We take homogeneous Dirichlet boundary conditions for the horizontal boundaries and Neumann for the vertical. We also use ℓ_1 -Jacobi [2] as the single-step smoother. This diagonal smoother has components $B_{ii}^{-1} = \sum_{j=1}^n |A_{ij}|$, and has the advantage of satisfying $\rho(BA) \leq 1$ by construction (as can be seen from the Gershgorin circle theorem), so that it need not be rescaled.

We note that this problem is of practical interest as it serves as a proxy for a low-order based preconditioner for a high-order element problem, considered, e.g., in [6]. Also note that, while the high aspect ratio elements of our first model problem could be handled by a semi-coarsening strategy, such an approach is not viable here. Approaches based on algebraic multigrid can result in hierarchies with small approximation constants C at each level, but typically at the expense of increased operator complexity. As a result, geometric multigrid can be competitive, depending on specifics of the problem. The use of polynomial acceleration of the smoother iteration, as we will see, can further increase the competitiveness of a geometric multigrid approach.

Asymptotic convergence rates $\|E\|_A^2$ for the symmetric V-cycle as well as asymptotic smoother steps per decimal digit $-k/\log_{10}\|E\|_A$ are shown in Figure 9 for the cases $N = 8$ and $N = 16$. Overall, results are quite similar to those for our first model problem. We have also included plots of time to solution—the time to reduce the 2-norm of the same random initial residual by the factor 10^{10} using a simple Matlab implementation of the symmetric multigrid V-cycle running on a desktop. We include the plot mainly to prove the point that the metric $-k/\log_{10}\|E\|_A$ is indeed indicative of real performance. There are two principal differences. For one, the total number of smoothing

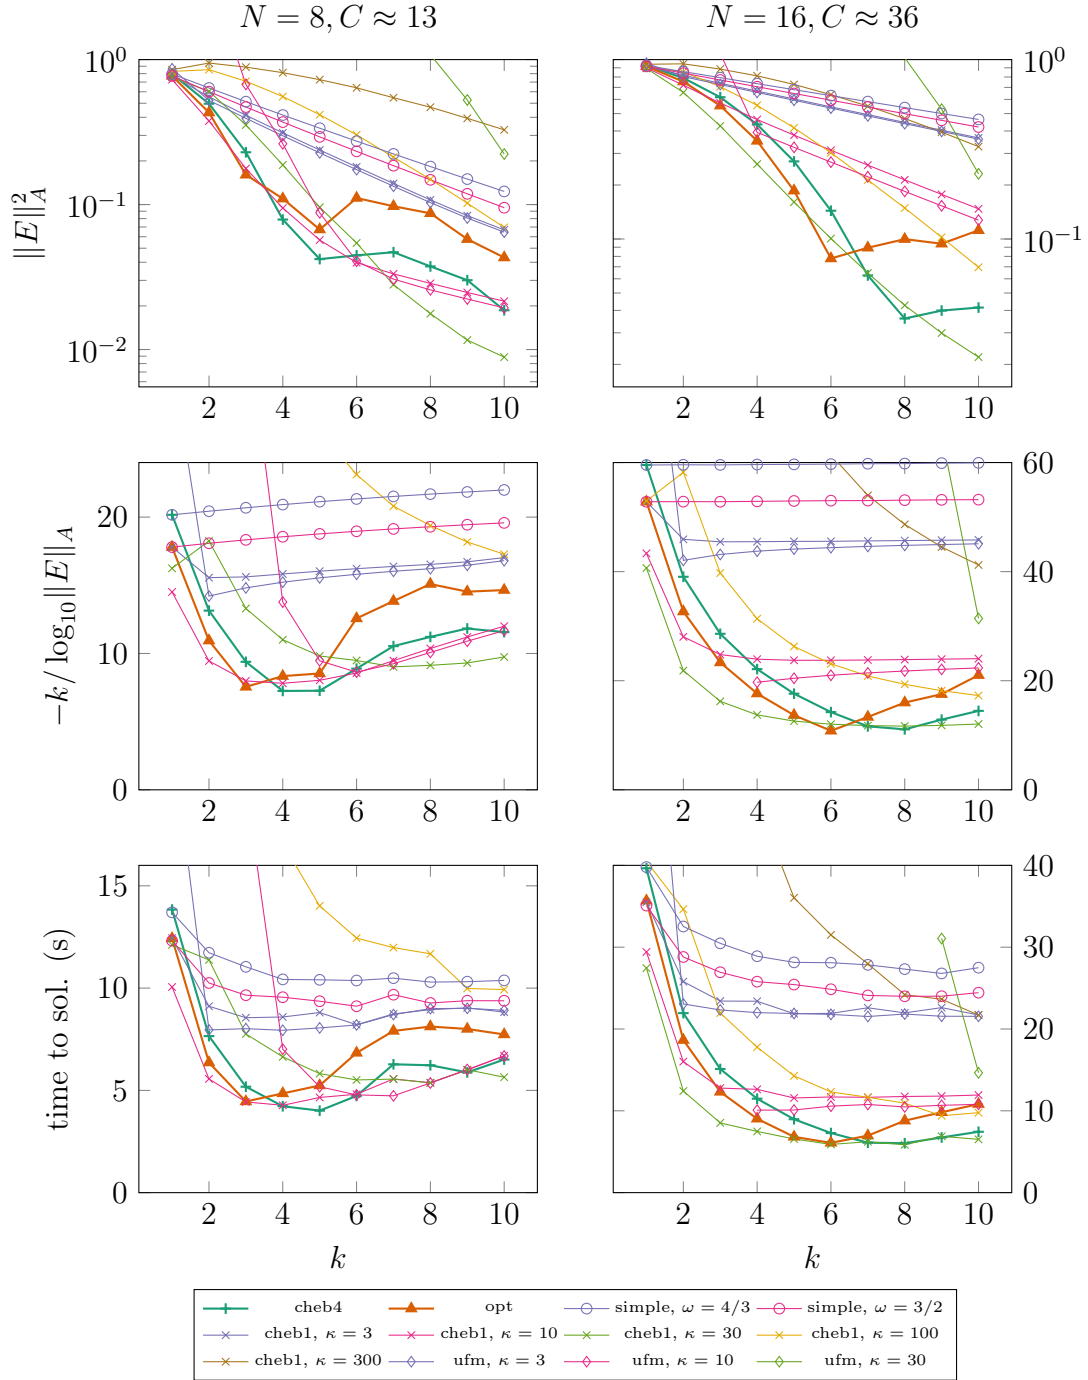


Figure 9: V-Cycle error contraction factors and smoother steps per digit for Poisson discretized on a Chebyshev grid using ℓ_1 -Jacobi smoothing.

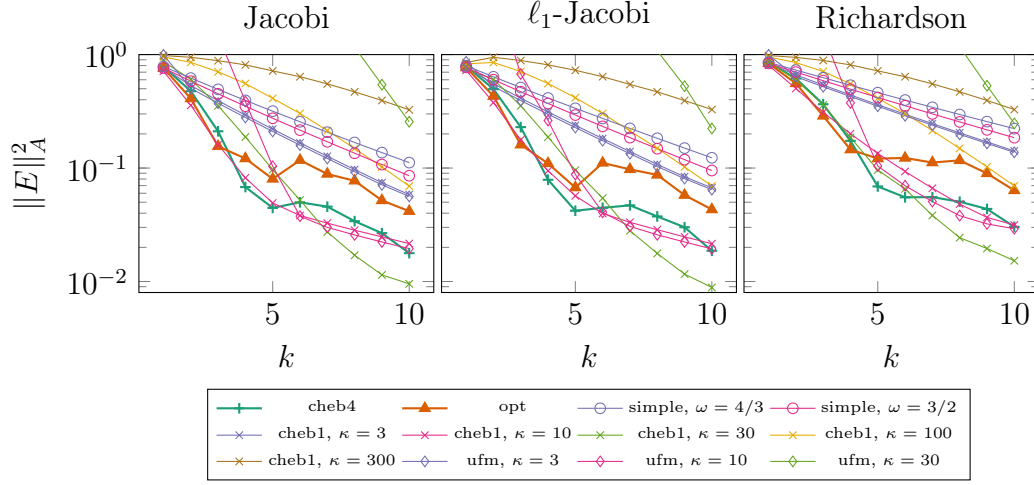


Figure 10: V-Cycle error contraction factors for different single-step smoothers for Poisson discretized on a Chebyhev grid with $N = 8$.

steps in an actual run is necessarily a multiple of $2k$, leading to a “quantization” effect. More significantly, the time to solution includes the time taken by the prolongation and restriction steps, which scale with the number of V-cycles, thus typically decreasing with k . In particular, this causes even the simple smoothers to benefit from increasing k . More generally, the effect is that the $-k/\log_{10}\|E\|_A$ metric may underpredict the optimal k . The time to solution plots also clearly demonstrate the benefit of using polynomial smoothers for this problem, and in going to higher polynomial degrees for the more difficult case, a trend that would continue if we increase N . Moreover, we see that the Chebyshev polynomials of the first kind with a well chosen parameter κ , the Chebyshev polynomials of the fourth kind, and the optimized polynomials are all competitive in terms of achieving the minimal time to solution.

In Figure 10, we compare the convergence rates when using three different diagonal single-step smoothers. Richardson performs somewhat worse than the other two, while Jacobi and ℓ_1 -Jacobi perform comparably, neither one being definitively better in all cases. As mentioned, ℓ_1 -Jacobi has the advantage of not needing to estimate the largest eigenvalue for rescaling. In particular, when combined with the fourth-kind Chebyshev or optimized iterations, the result is a truly parameter-free method.

6 Conclusion

The fourth-kind Chebyshev iteration (9) is a very simple way to accelerate the simple smoother iteration $(I - \omega BA)^k$ in a multigrid V-cycle for an SPD system, provided the smoother B is also SPD. All that is required is an estimate of the largest eigenvalue of BA , $\rho(BA)$, which can be obtained from a few iterations of conjugate gradient or a generalized Lanczos iteration. No further parameters are involved. In particular, compared to the usual (first-kind) Chebyshev semi-iteration, the fourth-kind iteration avoids the need to pick a minimum eigenvalue determining the range of eigenvalues to target. The use of a single-step smoother like ℓ_1 -Jacobi, which satisfies $\rho(BA) \leq 1$ by construction, can further eliminate even the need to estimate the largest eigenvalue.

We have given to our knowledge the first multilevel V-cycle bound for general (SPD) polynomial smoothers, isolating the polynomial from the single-step smoother. The asymptotic convergence rate of the symmetric V-cycle is bounded by the largest value across levels of

$$\frac{C}{C + \gamma^{-1}}, \quad \gamma := \sup_{0 < \lambda \leq 1} \frac{\lambda p(\lambda)^2}{1 - p(\lambda)^2},$$

where p is the polynomial, assumed to correspond to a convergent iteration with $p(0) = 1$, and where C is the approximation property constant

$$C = \|A^{-1} - PA_c^{-1}P^*\|_{A,B}^2 = \sup_{u \in \mathbf{R}(P)^{\perp A} \setminus \{0\}} \frac{\|u\|_{B^{-1}}^2}{\|u\|_A^2} \geq 1,$$

involving the single-step smoother B , assumed to be scaled such that $\rho(BA) \leq 1$.

For the fourth-kind Chebyshev iteration, we have $\gamma^{-1} = \frac{4}{3}k(k+1)$ where k is the number of smoothing steps (the degree of the polynomial), which is a significant improvement over $\gamma^{-1} = 2\omega k$ for the simple damped iteration, with damping parameter ω not too large. Moreover, it is possible to numerically determine the polynomials that minimize γ , which appear to give $\gamma^{-1} \geq \frac{4}{\pi^2}(2k+1)^2 - \frac{2}{3}$, giving roughly an 18% improvement in the convergence bound asymptotically as $k \rightarrow \infty$. The optimized polynomials can be implemented as a slight variation of the fourth-kind Chebyshev iteration, using a small set of coefficients computed and tabulated off-line.

In all our numerical experiments, whenever multiple smoothing steps were used, it was always the case that some polynomial per-

formed better than the simple smoother iteration. Increasing smoothing steps was ineffective for our example with coefficient jumps, which featured an unbounded approximation constant C while still attaining a V-cycle convergence rate bounded below 1. Our theory indicates that a bounded constant C is a sufficient condition for polynomial methods to effectively improve the convergence rate as the number of smoothing steps increases. Moreover, we saw that with polynomial methods, the improved convergence with more smoothing steps can more than make up for the increased cost, with the optimal number of smoothing steps typically increasing with the approximation constant C .

In our experiments, the standard first-kind Chebyshev iteration was often able to perform somewhat better than the fourth-kind Chebyshev iteration and its optimized variant. However, this depended on a good choice of the parameter κ determining the range of eigenvalues targeted, which we saw depended both on the problem and on the number of smoothing steps k , typically increasing both with C and k . An optimal parameter in one context could perform quite poorly in another. In contrast, the two parameter-free methods tended to stay within “striking distance” of the best method, never performing very poorly. Moreover, they remained especially competitive, at least for the problems we considered, when considering only the most efficient number of smoothing steps k . In practice, if it is possible to spend time tuning (by amortizing the cost across solving many right-hand-sides), we would recommend trying all options, tuning the parameters involved, and picking the fastest method. Otherwise, we would argue that the fourth-kind Chebyshev iteration or its optimized variant make sensible defaults.

References

- [1] Mark Adams, Marian Brezina, Jonathan Hu, and Ray Tuminaro. “Parallel multigrid smoothing: polynomial versus Gauss-Seidel”. In: *Journal of Computational Physics* 188.2 (July 2003), pp. 593–610.
- [2] Allison H. Baker, Robert D. Falgout, Tzanio V. Kolev, and Ulrike Meier Yang. “Multigrid Smoothers for Ultraparallel Computing”. In: *SIAM J. Sci. Comput.* 33.5 (Jan. 2011), pp. 2864–2887.

- [3] Marian Brezina. “Robust Iterative Methods on Unstructured Meshes”. PhD thesis. University of Colorado at Denver, 1997.
- [4] Robert N. Desmarais and Samuel R. Bland. *Tables of properties of airfoil polynomials*. NASA Reference Publication 1343. Sept. 1, 1995.
- [5] Wolfgang Hackbusch. “Multi-grid convergence theory”. In: *Multigrid Methods*. Jan. 1, 1982, pp. 177–219.
- [6] J. J. Heys, T. A. Manteuffel, Steve F. McCormick, and L. N. Olson. “Algebraic multigrid for higher-order finite elements”. In: *Journal of computational Physics* 204.2 (2005), pp. 520–532.
- [7] Johannes Kraus, Panayot Vassilevski, and Ludmil Zikatanov. “Polynomial of Best Uniform Approximation to and Smoothing in Two-level Methods”. In: *Computational Methods in Applied Mathematics* 12.4 (Jan. 1, 2012), pp. 448–468.
- [8] J.C. Mason. “Chebyshev polynomials of the second, third and fourth kinds in approximation, indefinite integration, and integral transforms”. In: *Journal of Computational and Applied Mathematics* 49.1 (Dec. 1993), pp. 169–178.
- [9] John C. Mason and David C. Handscomb. *Chebyshev polynomials*. CRC press, 2002.
- [10] S. F. McCormick. “Multigrid methods for variational problems: general theory for the V-cycle”. In: *SIAM J. Numer. Anal.* 22.4 (Aug. 1, 1985), pp. 634–643.
- [11] Radek Tezaur and Petr Vaněk. “Improved convergence bounds for two-level methods with an aggressive coarsening and massive polynomial smoothing”. In: *ETNA - Electronic Transactions on Numerical Analysis* 48 (Jan. 1, 2018), pp. 264–285.
- [12] Petr Vaněk and Marian Brezina. “Nearly optimal convergence result for multigrid with aggressive coarsening and polynomial smoothing”. In: *Appl Math* 58.4 (Aug. 1, 2013), pp. 369–388.

- [13] Petr Vaněk, Marian Brezina, and Radek Tezaur. “Two-grid method for linear elasticity on unstructured meshes”. In: *SIAM Journal on Scientific Computing* 21.3 (1999), pp. 900–923.
- [14] Petr Vaněk, Radek Tezaur, Marian Brezina, and J. Krížková. “Two-level method with coarse space size independent convergence”. In: *Domain Decomposition Methods in Sciences and Engineering, 8th International Conference, Beijing, PR China*. 1996, pp. 233–240.
- [15] Panayot S. Vassilevski. *Multilevel Block Factorization Preconditioners: Matrix-based Analysis and Algorithms for Solving Finite Element Equations*. Springer, Oct. 22, 2008. 527 pp.
- [16] H. A. van der Vorst. “A generalized Lánczos scheme”. In: *Math. Comp.* 39.160 (1982), pp. 559–561.
- [17] Jinchao Xu and Ludmil Zikatanov. “The method of alternating projections and the method of subspace corrections in Hilbert space”. In: *J. Amer. Math. Soc.* 15.3 (Apr. 8, 2002), pp. 573–597.

Appendix A Fourth-kind Chebyshev V-cycle Bound

The V-cycle convergence bound of Corollary 3.4 for the 4th kind Chebyshev polynomial smoothers follows from the fact that the scaled and shifted Chebyshev polynomials of the 4th kind

$$p_k(\lambda) = \frac{1}{2k+1} W_k(1-2\lambda)$$

satisfy

$$\sup_{0 < \lambda \leq 1} \frac{\lambda p_k(\lambda)^2}{1 - p_k(\lambda)^2} = \frac{1}{\frac{4}{3}k(k+1)},$$

which we prove below.

Proof of Corollary 3.4. If we let $n = 2k + 1 \geq 3$ and $\phi = \frac{1}{2}\theta$ where $\cos \theta = 1 - 2\lambda$, then we can write

$$p_k(\lambda) = \frac{1}{n} \frac{\sin n\phi}{\sin \phi}.$$

For ϕ an integer multiple of $\frac{\pi}{n}$, $p_k(\lambda) = 0$. Otherwise,

$$\frac{\lambda p_k(\lambda)^2}{1 - p_k(\lambda)^2} = \frac{1}{f_n(\phi)}, \quad f_n(\phi) := n^2 \csc^2 n\phi - \csc^2 \phi$$

and it suffices to show

$$\inf_{0 < \phi \leq \frac{\pi}{2}, \phi \notin \frac{\pi}{n}\mathbb{Z}} f_n(\phi) = \frac{4}{3}k(k+1) = \frac{1}{3}(n^2 - 1).$$

For $\frac{\pi}{n} < \phi \leq \frac{\pi}{2}$, since $\sin \phi \geq \frac{2}{\pi}\phi \geq \frac{2}{n}$,

$$f_n(\phi) \geq n^2 - \left(\frac{n}{2}\right)^2 = \frac{3}{4}n^2 \geq \frac{1}{3}(n^2 - 1).$$

So we can assume $0 < \phi < \frac{\pi}{n}$. By differentiating the classical partial fraction expansion of cotangent,

$$\cot z = \frac{1}{z} + 2z \sum_{j=1}^{\infty} \frac{1}{z^2 - j^2\pi^2},$$

we see that

$$\csc^2 z = \frac{1}{z^2} + 2 \sum_{j=1}^{\infty} \frac{z^2 + j^2\pi^2}{(z^2 - j^2\pi^2)^2}$$

so

$$\frac{1}{2}f_n(\phi) = \sum_{j=1}^{\infty} \frac{\phi^2 + j^2(\frac{\pi}{n})^2}{(\phi^2 - j^2(\frac{\pi}{n})^2)^2} - \sum_{j=1}^{\infty} \frac{\phi^2 + j^2\pi^2}{(\phi^2 - j^2\pi^2)^2}.$$

Since

$$\lim_{\phi \rightarrow 0} \frac{1}{2}f_n(\phi) = \frac{n^2 - 1}{\pi^2} \sum_{j=1}^{\infty} \frac{1}{j^2} = \frac{1}{6}(n^2 - 1),$$

it suffices to show that $f_n(\phi)$ is increasing on the interval $0 < \phi < \frac{\pi}{n}$. Differentiating term by term, we find

$$\frac{1}{2}f'_n(\phi) = 2\phi \sum_{j=1}^{\infty} a_{n,j}(\phi)$$

where

$$a_{n,j}(\phi) := \frac{3j^2(\frac{\pi}{n})^2 + \phi^2}{(j^2(\frac{\pi}{n})^2 - \phi^2)^3} - \frac{3j^2\pi^2 + \phi^2}{(j^2\pi^2 - \phi^2)^3}.$$

If each term $a_{n,j}(\phi)$ in the convergent sum is nonnegative, it will follow that $f'_n(\phi) > 0$. We can write

$$a_{n,j}(\phi) = \frac{j^2(n^2 - 1)}{n^6(j^2(\frac{\pi}{n})^2 - \phi^2)^3(j^2\pi^2 - \phi^2)^3} b_{n,j}(\phi)$$

where

$$b_{n,j}(\phi) := 6n^4\pi^2\phi^6 - 3n^2(n^2+1)j^2\pi^4\phi^4 + j^4(n^4 - 8n^2 + 1)\pi^6\phi^2 + 3(n^2+1)j^6\pi^8.$$

We can further write $b_{n,j}(\phi)$ as

$$b_{n,j}(\phi) = \left[\sqrt{6n^2\pi}\phi^3 - \sqrt{\frac{3}{8}}(n^2+1)j^2\pi^3\phi \right]^2 + j^4\pi^6c_{n,j}(\phi)$$

where

$$c_{n,j}(\phi) := \frac{5}{8}(n^4 - 14n^2 + 1)\phi^2 + 3(n^2+1)j^2\pi^2.$$

Since $n^4 - 14n^2 + 1 > 0$ for $n \geq 4$, it follows that $c_{n,j}(\phi) > 0$ for $n \geq 4$. On the other hand, for $n = 3$,

$$c_{3,j}(\phi) = 30j^2\pi^2 - \frac{55}{2}\phi^2 > 30\pi^2 - \frac{55}{2}\left(\frac{\pi}{3}\right)^2 = \frac{485\pi^2}{18} > 0$$

since $j \geq 1$ and $0 < \phi < \frac{\pi}{3}$. So, in either case, $c_{n,j}(\phi) > 0$, so that $b_{n,j}(\phi) > 0$, whereby $a_{n,j}(\phi) > 0$, and thus $f_n(\phi)$ is increasing for $0 < \phi < \frac{\pi}{n}$. \square

Appendix B Computing Optimal Polynomials

The optimal polynomials for Lemma 3.2 can be found numerically by applying Newton's method to a system of nonlinear equations expressing the equioscillation property. Let us represent the polynomial p by its k roots r_1, \dots, r_k ,

$$p = \prod_{i=1}^k \left(1 - \frac{x}{r_i}\right).$$

Note that p satisfies the required condition $p(0) = 1$. We seek to minimize $\gamma = \sup_{0 < x \leq 1} f^2$ where

$$f := w^{1/2}p, \quad w := \frac{x}{1-p^2}, \quad x > 0,$$

and where we take $w(0) := [-2p'(0)]^{-1}$ so that w and f are continuous at $x = 0$. For the optimal p , f will equioscillate,

$$f(0) = |f(x_1)| = \dots = |f(x_{k-1})| = |f(1)| \quad (21)$$

where x_i are the $k - 1$ extrema of f , $f'(x_i) = 0$, located between the k roots r_i of p . For convenience, take

$$x_k := 1$$

and let $F : \mathbb{R}^k \rightarrow \mathbb{R}^k$ be the function of the vector of roots $r = [r_1 \ \dots \ r_k]^T$ with components

$$F_i = f(0) - |f(x_i)|.$$

Then (21) is equivalent to $F(r) = 0$ and the optimal r can be found with Newton's method. A natural initial guess for r is the vector of roots of the shifted Chebyshev polynomial of the fourth kind,

$$r_i^{(0)} = \frac{1}{2} - \frac{1}{2} \cos \frac{i}{k + \frac{1}{2}} \pi, \quad i = 1, \dots, k.$$

A second, nested, Newton iteration can be used to compute the extrema x_1, \dots, x_{k-1} of f . For the first outer iteration, a good initial guess for the extrema are those of the initial p ,

$$x_i^{(0)} = \frac{1}{2} - \frac{1}{2} \cos \frac{i + \frac{1}{2}}{k + \frac{1}{2}} \pi, \quad i = 1, \dots, k - 1.$$

Subsequent iterations can use the extrema from the previous iteration for the initial guess.

The total logarithmic derivative of p is

$$\frac{dp}{p} = \sum_{i=1}^k \frac{dx - \frac{x}{r_i} dr_i}{x - r_i} = q dx - \sum_{i=1}^k \frac{x}{x - r_i} \frac{dr_i}{r_i},$$

where

$$q(x) := \frac{p'(x)}{p(x)} = \sum_{i=1}^k \frac{1}{x - r_i},$$

and the total logarithmic derivative of f is

$$\frac{df}{f} = \frac{dx}{2x} + \frac{w}{x} \frac{dp}{p} = \left(\frac{1}{2} + wq\right) \frac{dx}{x} - \sum_{i=1}^k \frac{w}{x - r_i} \frac{dr_i}{r_i}. \quad (22)$$

It follows that the extrema x_i of f are determined by the condition $\frac{1}{2} + wq = 0$, or equivalently by $g(x_i) = 0$ where

$$g(x) := \frac{x}{2w(x)} + xq(x) = \frac{1 - p(x)^2}{2} + xq(x).$$

The negative derivative of $g(x)$ is

$$-g'(x) = \sum_{i=1}^k \frac{1}{x - r_i} \left[p(x)^2 + \frac{r_i}{x - r_i} \right],$$

giving the Newton step to update the extrema, $x_i \leftarrow x_i - g(x_i)/g'(x_i)$.

Suppose x is constrained to equal x_i , i.e., to be an extremum of f when $i < k$ or else fixed at 1 when $i = k$. In the first case $\frac{1}{2} + wq = 0$ and in the second case $dx = 0$, so that in both cases (22) reduces to

$$\frac{df}{f} = - \sum_{j=1}^k \frac{w}{x - r_j} \frac{dr_j}{r_j},$$

from which we can read off $\partial f(x_i)/\partial r_j$. Combining with $\partial f(0)/\partial r_j = f(0)^3/r_j^2$, which can be derived from $f(0) = [-2q(0)]^{-1/2}$, we find that the components of the Jacobian J of F are

$$J_{ij} = \frac{f(0)^3}{r_j^2} + \frac{w(x_i)|f(x_i)|}{r_j(x_i - r_j)}.$$

The Newton step for updating the vector of roots is then just $r \leftarrow r + \Delta r$ where $J\Delta r = -F$. A full Matlab implementation of the algorithm is shown in Figure 11. Empirically, the initial guesses suggested above are good enough for convergence of the pair of Newton iterations, at least up to $k = 200$.

Once the roots of the optimal $p_k(\lambda)$ have been determined, we can compute the coefficients α_i of the expansion

$$p_k(\lambda) = \sum_{i=0}^k \alpha_i W_i(1 - 2\lambda) \quad (23)$$

of p_k in the fourth-kind Chebyshev polynomial basis, so that

$$\beta_0 = 1, \quad \beta_{i+1} = \beta_i - (2i + 1)\alpha_i$$

then gives the coefficients β_i for implementing p_k in the iteration (19). By combining the orthogonality relation for the fourth-kind Chebyshev polynomials [4]

$$\frac{1}{\pi} \int_{-1}^1 \sqrt{\frac{1-x}{1+x}} W_i(x) W_j(x) dx = \delta_{ij}$$

```

function r = opt_roots(k)
    function [w,f,g,ngp] = vars(r, x)
        p = prod(1 - x' ./ r, 1)';
        w = x ./ (1 - p.^2);
        f = sqrt(w) .* p;
        q = sum(1 ./ (x' - r), 1)';
        g = x .* (1./(2*w) + q);
        ngp = sum((p'.^2 + r./(x'-r))./(x'-r), 1)';
    end

    r = .5 - .5*cos([1:k]'/(k+.5) * pi);
    x = .5 - .5*cos((.5+[1:k-1]')/(k+.5) * pi);

    dr = r; drsize = 1;
    while drsize > 128*eps
        dx = x; dxsize = 1;
        while dxsize > 128*eps
            dxsize = norm(dx, inf);
            [~,~,g,ngp] = vars(r,x);
            dx = g ./ ngp; x = x + dx;
        end
        x1 = [x;1]; [w,f,~,~] = vars(r,x1);
        f0 = sqrt(.5/sum(1./r));
        J = f0^3./r'.^2 + w.*abs(f)./(r'.*(x1-r'));
        drsize = norm(dr, inf);
        dr = -J \ (f0 - abs(f)); r = r + dr;
    end
end

```

Figure 11: A Matlab implementation of Newton's method for finding the roots of the optimal polynomials for Lemma 3.2.

```

function beta = opt_coef(k)
    x = cos([1:k]'*pi/(k+.5)); % Quadrature nodes
    w = (1-x)/(k+.5); % Quadrature weights

    % 4th kind Chebyshev polynomials W evaluated at x
    W = zeros(k, k);
    W(:, 1) = 1;
    if k >= 2; W(:, 2) = 2*x+1; end
    for i=3:k
        W(:,i) = (2*x) .* W(:,i-1) - W(:,i-2);
    end

    r = opt_roots(k); % Roots of optimal polynomial p
    lambda = (1-x)/2; % Nodes transformed to [0, 1]
    p = prod(1-lambda'./r, 1)'; % p(lambda)
    alpha = W' * (w .* p);
    beta = 1 - cumsum((2*[0:k-1]'+1) .* alpha, 1);
end

```

Figure 12: Matlab code for computing the coefficients β_i for the optimized iteration (19).

with the associated Gaussian quadrature rule [4],

$$\frac{1}{\pi} \int_{-1}^1 \sqrt{\frac{1-x}{1+x}} p(x) dx = \sum_{i=1}^k \frac{1-x_i}{k+\frac{1}{2}} p(x_i), \quad x_i = \cos \frac{i}{k+\frac{1}{2}} \pi,$$

which holds for any polynomial $p(x)$ of degree at most $2k-1$, it follows that if (23) holds then

$$\alpha_j = \sum_{i=1}^k \frac{1-x_i}{k+\frac{1}{2}} W_j(x_i) p\left(\frac{1-x_i}{2}\right) \quad \text{for } j \leq k-1.$$

Note that α_k is not needed to compute β_k . A Matlab implementation of the computation of the coefficients β_i is shown in Figure 12.