

ALGEBRAIC ERROR ANALYSIS FOR MIXED-PRECISION MULTIGRID SOLVERS*

STEPHEN F. MCCORMICK[†], JOSEPH BENZAKEN[‡], AND RASMUS TAMSTORF[‡]

Abstract. This paper establishes the first theoretical framework for analyzing the rounding-error effects on multigrid methods using mixed-precision iterative-refinement solvers. While motivated by the sparse symmetric positive definite (SPD) matrix equations that arise from discretizing linear elliptic PDEs, the framework is purely algebraic such that it applies to matrices that do not necessarily come from the continuum. Based on the so-called *energy* or A norm, which is the natural norm for many problems involving SPD matrices, we provide a normwise forward error analysis, and introduce the notion of progressive precision for multigrid solvers. Each level of the multigrid hierarchy uses three different precisions that each increase with the fineness of the level, but at different rates, thereby ensuring that the bulk of the computation uses the lowest possible precision. The theoretical results developed here in the energy norm differ notably from previous theory based on the Euclidean norm in important ways. In particular, we show that simply rounding an exact result to finite precision causes an error in the energy norm that is proportional to the square root of κ , the associated matrix condition number. (By contrast, this error is of order 1 when measured in the Euclidean norm.) Given this observation, we show that the limiting accuracy for both V-cycles and full multigrid is optimal in the sense that it is also proportional to $\kappa^{1/2}$ in energy. Additionally, we show that the loss of convergence rate due to rounding grows in proportion to $\kappa^{1/2}$, but argue that this loss is insignificant in practice. The theory presented here is the first forward error analysis in the energy norm of iterative refinement and the first rounding error analysis of multigrid in general.

Key words. mixed precision, multigrid, rounding error analysis, progressive precision

AMS subject classifications. 65F10, 65G50, 65M55

DOI. 10.1137/20M1348571

1. Introduction. Most computing systems today are power limited. This is true all the way from battery operated edge devices to the top supercomputers of the world. Current trends in computer architectures therefore favor computation with low precision arithmetic because it allows higher throughput and reduces the amount of data that must be moved through the memory hierarchy. As an example, using single precision allows for roughly four times higher throughput than double precision,[13], while the savings due to memory traffic can be even larger depending on how close to the processor the data is stored [24]. Unfortunately, ill-conditioned problems are often intractable in very low precision using standard methods. As a result, much research over the years has focused on algorithms using mixed-precision computation, with most operations performed in low precision while select operations use higher precision to achieve higher accuracy. This approach was first introduced in [29], and later analyzed in [30] and [23]. Since then, it has become a well-known technique and, with the introduction of GPUs, much research has focused on combining single and double precision [20, 7, 6, 3]. More recently, the addition of hardware support for half-precision (FP16) computations has led to a push toward its use [18, 17]. Mixed-precision approaches are also beneficial for minimizing the need for extended precision when dealing with problems that are inherently so ill-conditioned that they are difficult to solve even using today’s high-precision arithmetic. Examples of such

*Received by the editors July 1, 2020; accepted for publication (in revised form) February 18, 2021; published electronically June 24, 2021.

[†]University of Colorado at Boulder, Boulder, CO (stephen.mccormick@colorado.edu).

[‡]Walt Disney Animation Studios, Burbank, CA (Joseph.Benzaken@disneyanimation.com, Rasmus.Tamstorf@disneyanimation.com).

ill-conditioned problems include high-resolution discretizations of high-order partial differential equations (PDEs).

As described in [18], many mixed-precision methods fall within the broad class of nested *inner-outer* iterative methods. In this context, *outer* refers to an upper-level solver that calls the *inner* method to perform essential solver tasks. In this paper, we consider a particular algorithm belonging to this class of methods. Specifically, we use iterative refinement (IR) with three precisions as the outer solver (similar to [9]) and a *restarted* multigrid V-cycle as the inner solver. As an important contribution, we extend the traditional notion of mixed-precision IR to progressive precision by allowing all precision levels to increase with each new fine level in the multigrid hierarchy. Our goal is to establish a purely *algebraic* theory for this method that guarantees convergence when the matrix condition number times unit roundoff is about one at each level of the hierarchy. The estimates confirm that convergence is achieved at nearly the same rate as in infinite precision. In the case of discretized elliptic PDEs, the resulting accuracy is comparable to discretization accuracy. In the more general case, we obtain an algebraic notion of discretization error that the algorithm attains even if the matrices do not originate from PDEs. The theory also establishes optimal convergence to this abstract discretization accuracy for progressive-precision full multigrid (FMG). While mixed-precision methods have previously been considered for a variety of multigrid algorithms (e.g., [26, 14, 10]), convergence theory has to the best of our knowledge not been established before.

Our forward rounding error analysis for iterative refinement parallels that in [8], with the crucial difference that we focus on (sparse) symmetric positive definite (SPD) matrices. This allows us to pose the theory in the energy norm, which is the natural and desired norm for a large class of elliptic PDEs. Physical principles are often posed in terms of minimizing some physical quantity such as energy, leading to PDEs of this class equipped with a naturally induced energy norm.

Extending the results of [8] from the infinity to the energy or A norm is critical because it enables access to standard variational theory for V-cycles and FMG, and the resulting forward analysis enables direct estimation of the rounding-error effects involving $\kappa^{1/2}$ instead of κ , where κ is the matrix condition number. (Note that all norms on a *fixed* Euclidean space are theoretically equivalent, but the constants really matter: those relating the energy and Euclidean norms depend on κ . As such, the disparity in these constants grows rapidly up through the hierarchy of levels that we consider.) We also exploit matrix sparsity to eliminate the direct dependence of the error bounds on the size of the matrix, thus avoiding an even higher power of the condition number in the estimates. These advantages for the sparse SPD case allow us to sharpen existing theory for general matrices, most notably those in [8].

The results developed here in the energy norm differ notably from previous theory based on the Euclidean norm in other important ways. In particular, we show that simply rounding an exact result to finite precision causes an error in the energy norm that is proportional to $\kappa^{1/2}$. (By contrast, this error is of order 1 when measured in the Euclidean norm.) Given this observation, the limiting accuracy for both V-cycles and full multigrid proves to be optimal in the sense that it is also proportional to $\kappa^{1/2}$ in energy. The loss of convergence rate from rounding also grows in proportion to $\kappa^{1/2}$, but we argue that this loss is insignificant in practice.

By definition, multigrid coarsening in correction form applies to the residual equation, so it is naturally an iterative-refinement process. Each step of a residual-based relaxation scheme can also be interpreted as iterative refinement. A mixed-precision version of multigrid therefore requires very little change in the algorithm itself, and

that change is primarily in the choice of where to invoke higher precision. Our choice here is to apply a simple multigrid algorithm in low precision to the residuals computed in higher precision between cycles, where the cycles we study involve one relaxation sweep on each level. (We establish an extension to multiple sweeps in the supplemental material.)

We begin in the next section by describing tools used in our theory in the form of basic error-analysis estimates for floating point arithmetic. Section 3 reviews the classical iterative-refinement process while detailing the notation, terminology, and certain principles that are used throughout the remaining sections. We develop a general theory in Section 4 that provides error bounds for iterative refinement assuming an abstract energy-convergent inner solver. Section 5 then introduces a reformulation of the two-grid method that leverages mixed-precision iterative refinement. The V-cycle follows in Section 6. Theoretical results are presented in Section 7 and 8 that, together with our iterative refinement estimates, provide a framework for determining convergence of mixed-precision two-grid and V-cycle solvers for specific applications. The notion of progressive precision is introduced in Section 8 as part of the analysis of the V-cycle. Progressive-precision FMG is then introduced in Section 9 and analyzed in Section 10. The paper concludes in Section 11. For easy reference, we have gathered all the symbols used throughout the paper in Appendix A.

2. Floating Point Estimates. We consider three basic floating point environments: “standard” precision with unit roundoff ε , referred to here as ε -precision; “high” precision with unit roundoff $\bar{\varepsilon}$, referred to here as $\bar{\varepsilon}$ -precision; and “low” precision with unit roundoff $\acute{\varepsilon}$, referred to here as $\acute{\varepsilon}$ -precision. Initially, one can think of $\bar{\varepsilon}$ -precision as being double, and $\acute{\varepsilon}$ -precision as being half, that of ε -precision, but we only formally require $\bar{\varepsilon} \leq \varepsilon \leq \acute{\varepsilon}$.

Our analysis is based on existing models of error due to rounding. See [19] for a thorough discussion of these models and other aspects of rounding error analysis. This section describes only those estimates used in our theory. We do not consider computation using posits [15] or other unums [16] because this would require different bounds than those presented in this section. We leave these issues as an interesting avenue for future work.

Let \circ stand for one of the four basic binary operations $+$, $-$, $*$, $/$ between two scalars x and y , and let $\text{fl}(x \circ y)$ stand for the ε -precision computed value of the exact operation $x \circ y$. The standard model [19, Eq. (2.4)] that accounts for rounding error (in ε -precision for illustration) is then given by

$$(2.1) \quad \text{fl}(x \circ y) = (x \circ y)(1 + \delta), \quad |\delta| \leq \varepsilon,$$

where $\text{fl}()$ denotes the result of the floating point operation. (The definition of δ with and without subscripts may change meaning at each occurrence in this paper, especially here and from one proof to the next.) An alternative to (2.1) is the variant [19, Eq. (2.5)], which we write as

$$(2.2) \quad \text{fl}(x \circ y) = x \circ y + \delta \text{fl}(x \circ y), \quad |\delta| \leq \varepsilon.$$

Assuming now that $A \in \mathbb{R}^{n \times n}$ and $x \in \mathbb{R}^n$, then the dot product model [19, Eq. (3.4)] implies the following matrix-vector product estimate in ε -precision:

$$(2.3) \quad \text{fl}(Ax) = Ax + \delta, \quad |\delta| \leq \frac{n\varepsilon}{1 - n\varepsilon} |A| \cdot |x|,$$

where $|\cdot|$ denotes the vector or matrix with entries replaced by their absolute values and relations between vectors and matrices are defined componentwise. With $\|\cdot\|$

denoting the Euclidean norm for a vector and its induced matrix norm (together with the Euclidean inner product $\langle \cdot, \cdot \rangle$), we also frequently use the fact that $\| |z| \| = \|z\|$ for any vector z (although this is not generally true for matrices). If $b \in \mathbb{R}^n$, then an estimate for computing the residual $Ax - b$ in ε -precision is given by

$$(2.4) \quad \text{fl}(Ax - b) = Ax - b + \delta, \quad |\delta| \leq \frac{(n+1)\varepsilon}{1 - (n+1)\varepsilon} (|b| + |A| \cdot |x|).$$

Finally, computing the residual in $\bar{\varepsilon}$ -precision and rounding it to ε -precision results in the following bound [8, Eq. (2.3)]:

$$(2.5) \quad \text{fl}(Ax - b) = Ax - b + \delta, \quad |\delta| \leq \varepsilon |Ax - b| + (1 + \varepsilon) \frac{(n+1)\bar{\varepsilon}}{1 - (n+1)\bar{\varepsilon}} (|b| + |A| \cdot |x|).$$

We have used fl here to indicate computed quantities, but it would be too cumbersome to continue with this notation in the analysis that follows. Instead of fl or any other special way to denote computed quantities, we add δ 's to exact expressions, with various subscripts and bounds, to denote quantities computed in finite precision. For example, the *computed* residual r in (2.5) would be written as

$$r = Ax - b + \delta_r, \quad |\delta_r| \leq \varepsilon |Ax - b| + (1 + \varepsilon) \frac{n\bar{\varepsilon}}{1 - n\bar{\varepsilon}} (|b| + |A| \cdot |x|)$$

and the *computed* solution of $Ax = b$ would be written as $A^{-1}b + \delta$.

Following the usual convention in rounding-error analyses, we assume throughout that the system matrix A and right-hand side (RHS) b for our target problem (see (3.1) below) are exact. To rein in the complexity of our convergence bounds, we take this assumption further by assuming exactness of all of the multigrid components: the intergrid transfer, the system matrix, and the RHS on all levels. While this may be a reasonable assumption in certain ideal cases (e.g., simple discretizations of Poisson's equation), these quantities are only approximate in most applications, such as when they are constructed via finite elements using quadrature for a weak form applied to individual basis elements. Even if these approximations are very accurate, A and b must ultimately be stored in finite precision. While it may be argued that errors in the target problem are not at issue here since our focus is on the discrete problem we are given, accuracy of the coarse components within the multigrid solvers is important. This is particularly true for FMG since its goal is to deliver discretization-order accuracy, which makes it the most error-sensitive multigrid scheme. See [27] for an analysis of the effects of finite precision on the coarse FMG components.

3. Iterative Refinement. Consider the $n \times n$ matrix equation

$$(3.1) \quad Ax = b,$$

where n is a positive integer, $A \in \mathbb{R}^{n \times n}$ is SPD, $x \in \mathbb{R}^n$ is unknown, and $b \in \mathbb{R}^n$ is given. Assume that A is sparse, with at most m_A nonzeros per row. The first mixed-precision approach we analyze in this paper uses IR as the outer loop and an approximate linear solver (to be specified later) as the inner loop. The pseudocode \mathcal{IR} is given in Algorithm 3.1 below and illustrated in Fig. 1. As noted by [12, Sec. 2.5], \mathcal{IR} is effectively Newton's method applied to the function $f(x) = Ax - b$, where $f(x)$ is considered as a nonlinear function due to rounding errors. It can also be




FIG. 1. An illustration of the steps associated with one iteration of \mathcal{IR} . Similar to the text in Algorithm 3.1, the image is color-coded with respect to precision. Specifically, green corresponds to reduced precision, blue to standard precision, and red to high precision. Note that the arrows extending from one precision to another are stroked with a gradient to convey that the precision associated with the entity under consideration is either extended or reduced. In the subsequent figures, if no input is provided for the correction, y , then it is assumed to be zero.

interpreted in many other ways, including preconditioned Richardson iteration [6], defect correction [1], and double discretization [4].

The floating-point operations in \mathcal{IR} here use all three precisions. The full residual $r = Ax - b$ between successive calls to the inner solver is evaluated in $\bar{\epsilon}$ -precision (red font), while the inner solver itself uses $\hat{\epsilon}$ -precision (green font). All other operations, most notably the update step, are evaluated using ϵ -precision (blue font). Note that step 2 of \mathcal{IR} is in red and blue font because the residual is computed in $\bar{\epsilon}$ -precision but rounded to ϵ -precision. A critical feature of \mathcal{IR} is that the inner loop is applied not to the full equation $Ax = b$ for the full approximation x , but rather to the \mathcal{IR} residual equation $Ay = r$ for the correction y .

Algorithm 3.1 Iterative Refinement (\mathcal{IR})

Input: A, b, x initial guess, $\text{tol} > 0$ convergence tolerance.
 1: $r \leftarrow Ax - b$ ▷ Compute \mathcal{IR} Residual and Round
 2: **if** $\|r\| < \text{tol}$ **then**
 3: **return** x ▷ Return Solution of $Ax = b$
 4: **end if**
 5: $y \leftarrow \text{InnerSolve}(A, r)$ ▷ Compute Approximate Solution of $Ay = r$
 6: $x \leftarrow x - y$ ▷ Update Approximate Solution of $Ax = b$
 7: **goto** 1

For simplicity, \mathcal{IR} here uses the residual error in its stopping criterion. When the matrix is highly ill-conditioned, the residual can be a poor indicator of convergence because most solvers tend to produce errors whose components are predominantly in the lower spectrum of the matrix and are thus hidden in the residual. Fortunately, in practice, the residual is usually a good indicator for multigrid: while a few relaxation sweeps generally lead to the lower end of the spectrum dominating the error, coarsening tends to reduce those components so that the upper end is exposed; taken together, relaxation and coarsening tend to produce errors that are balanced across the spectrum. Nevertheless, the stopping criterion for mixed-precision algorithms should be chosen carefully as [3] recommends. This issue is discussed further in [2] and [11], and we develop a stopping criterion carefully for full multigrid in the related paper [27] based on estimates of the differences between approximations on successive levels.

4. Convergence Theory: Iterative Refinement. The statement and proof of our first theorem follow a course similar to that in [8]. However, instead of the Euclidean norm $\|\cdot\|$, our rounding-error estimates are in terms of the discrete *energy* norm $\|\cdot\|_A$ defined by $\|x\|_A = \|A^{\frac{1}{2}}x\|$, $x \in \mathbb{R}^n$. Our bounds are also tighter in that they exploit the sparsity of A . We state this theorem in terms of a general iterative solver satisfying a convergence bound with generic factor ρ . Later, we consider specific multigrid solvers for the inner loop.

Several parameters appear below in our convergence estimates, such as the condition number of A denoted by $\kappa(A) = \|A\| \cdot \|A^{-1}\|$. For convenience, denote the norm of $|A|$ by $\psi(A) = \||A|\|$, the resulting modified condition number of A by $\underline{\kappa}(A) = \psi\|A^{-1}\|$, and a term that arises in several of our estimates by

$$\bar{m}_A^+ = \frac{m_A + 1}{1 - (m_A + 1)\bar{\varepsilon}}.$$

We henceforth suppress the explicit dependence of $\kappa, \underline{\kappa}$, and ψ on A with possible subscripts that remain (e.g., $\kappa_j = \kappa(A_j)$), but otherwise include its dependence explicitly.

All bounds obtained herein are expressed in terms of the parameters $\dot{\tau} = \kappa^{\frac{1}{2}}\dot{\varepsilon}$, $\tau = \kappa^{\frac{1}{2}}\varepsilon$, and $\bar{\tau} = \kappa\bar{\varepsilon}$ that represent fundamental scales. Quantities representing rounding errors in convergence factors (e.g., $\delta_{\rho_{ir}}$ and $\delta_{\rho_{tg}}$) and \mathcal{IR} 's limiting accuracy (i.e., χ) are functions of these parameters, but this is suppressed in the notation when that dependence is clear. To simplify the bounds below, let

$$(4.1) \quad \gamma = \frac{\kappa^{\frac{1}{2}} + \underline{\kappa}}{\kappa},$$

which is of order 1 for discrete Laplacians and many other discrete elliptic PDEs.

Throughout the paper, we use the notation of the pseudo algorithms we present, with superscripts used occasionally within the iterations to keep track of various quantities. Thus, $x^{(i)}$ denotes the i^{th} \mathcal{IR} iterate and $x^{(0)}$ the \mathcal{IR} initial guess.

THEOREM 4.1. \mathcal{IR} . *Let $x^{(i)}$ be the iterate at the start of the i^{th} cycle of \mathcal{IR} and $r = Ax^{(i)} - b$ its residual computed in $\bar{\varepsilon}$ -precision and rounded to ε -precision. Suppose that a $\rho < 1$ exists such that, for any $r \in \mathbb{R}^n$, the solver used in the inner loop of Algorithm 3.1 (line 5) is guaranteed to compute a correction y that satisfies*

$$(4.2) \quad \|y - A^{-1}r\|_A \leq \rho\|A^{-1}r\|_A.$$

Then $x^{(i+1)}$ approximates the solution $A^{-1}b$ of (3.1) with the relative error bound

$$(4.3) \quad \frac{\|x^{(i+1)} - A^{-1}b\|_A}{\|A^{-1}b\|_A} \leq \rho_{ir} \frac{\|x^{(i)} - A^{-1}b\|_A}{\|A^{-1}b\|_A} + \chi,$$

where the \mathcal{IR} convergence factor is given by

$$(4.4) \quad \rho_{ir} = \rho + \delta_{\rho_{ir}}, \quad \delta_{\rho_{ir}} = \frac{(1 + 2\rho)\tau + \gamma(1 + \rho)(1 + \varepsilon)\bar{m}_A^+\bar{\tau}}{1 - \tau},$$

and the limiting accuracy is given by

$$(4.5) \quad \chi = \frac{\tau + \gamma(1 + \rho)(1 + \varepsilon)\bar{m}_A^+\bar{\tau}}{1 - \tau}.$$

If $\rho + \delta_{\rho_{ir}} < 1$, then the error after $N \geq 1$ \mathcal{IR} cycles starting with initial guess $x^{(0)}$ satisfies

$$(4.6) \quad \frac{\|x^{(N)} - A^{-1}b\|_A}{\|A^{-1}b\|_A} \leq (\rho + \delta_{\rho_{ir}})^N \frac{\|x^{(0)} - A^{-1}b\|_A}{\|A^{-1}b\|_A} + \frac{\chi}{1 - (\rho + \delta_{\rho_{ir}})}.$$

Proof. Using (2.5) allows us to write the \mathcal{IR} residual computed in $\bar{\varepsilon}$ -precision and rounded to ε -precision as

$$(4.7) \quad r = \underbrace{Ax^{(i)} - b}_{\text{exact residual}} + \underbrace{\delta_1}_{\varepsilon\text{-}\bar{\varepsilon} \text{ error}}, \quad |\delta_1| \leq \varepsilon \|Ax^{(i)} - b\| + (1 + \varepsilon) \bar{m}_A^+ \bar{\varepsilon} \left(|b| + |A| \cdot |x^{(i)}| \right).$$

Note that

$$\begin{aligned} \left\| A^{-\frac{1}{2}} \left(|b| + |A| \cdot |x^{(i)}| \right) \right\| &\leq \|A^{-\frac{1}{2}}\| \left(\|AA^{-1}b\| + \psi \|x^{(i)}\| \right) \\ &\leq \|A^{-\frac{1}{2}}\| \cdot \|A^{\frac{1}{2}}\| \cdot \|A^{-1}b\|_A + \psi \|A^{-1}\| \cdot \|x^{(i)}\|_A \\ &= \kappa^{\frac{1}{2}} \|A^{-1}b\|_A + \underline{\kappa} \|x^{(i)}\|_A \\ &\leq (\kappa^{\frac{1}{2}} + \underline{\kappa}) \|A^{-1}b\|_A + \underline{\kappa} \|x^{(i)} - A^{-1}b\|_A. \end{aligned}$$

Hence,

$$(4.8) \quad \begin{aligned} \|A^{-1}\delta_1\|_A &= \|A^{-\frac{1}{2}}\delta_1\| \\ &\leq \varepsilon \|A^{-\frac{1}{2}}\| \cdot \|Ax^{(i)} - b\| + (1 + \varepsilon) \bar{m}_A^+ \bar{\varepsilon} \left\| A^{-\frac{1}{2}} \left(|b| + |A| \cdot |x^{(i)}| \right) \right\| \\ &\leq \varepsilon \kappa^{\frac{1}{2}} \|x^{(i)} - A^{-1}b\|_A \\ &\quad + (1 + \varepsilon) \bar{m}_A^+ \bar{\varepsilon} \left((\kappa^{\frac{1}{2}} + \underline{\kappa}) \|A^{-1}b\|_A + \underline{\kappa} \|x^{(i)} - A^{-1}b\|_A \right). \end{aligned}$$

By (4.2) and (4.7), we thus have that

$$(4.9) \quad \begin{aligned} \|y - A^{-1}r\|_A &\leq \rho \|A^{-1}r\|_A \\ &= \rho \|A^{-1} \left(Ax^{(i)} - b + \delta_1 \right)\|_A \\ &\leq \rho \left[\|x^{(i)} - A^{-1}b\|_A + \|A^{-1}\delta_1\|_A \right] \\ &\leq \rho \left[(1 + \varepsilon \kappa^{\frac{1}{2}}) \|x^{(i)} - A^{-1}b\|_A \right. \\ &\quad \left. + (1 + \varepsilon) \bar{m}_A^+ \bar{\varepsilon} \left((\kappa^{\frac{1}{2}} + \underline{\kappa}) \|A^{-1}b\|_A + \underline{\kappa} \|x^{(i)} - A^{-1}b\|_A \right) \right]. \end{aligned}$$

Using (2.2) allows us to write the \mathcal{IR} update computed in ε -precision as

$$x^{(i+1)} = \underbrace{x^{(i)} - y}_{\text{exact update}} + \underbrace{\delta_2}_{\varepsilon \text{ error}}, \quad |\delta_2| \leq \varepsilon |x^{(i+1)}|.$$

Using (4.8) and (4.9) thus leads to

$$\begin{aligned}
\|x^{(i+1)} - A^{-1}b\|_A &= \|x^{(i)} - A^{-1}b - A^{-1}r - (y - A^{-1}r) + \delta_2\|_A \\
&\leq \|A^{-1}\delta_1 + (y - A^{-1}r)\|_A + \|\delta_2\|_A \\
&\leq \|A^{-1}\delta_1\|_A + \|y - A^{-1}r\|_A + \varepsilon\|A^{\frac{1}{2}}\| \cdot \|x^{(i+1)}\| \\
&\leq \left(\rho + (1 + \rho)\varepsilon\kappa^{\frac{1}{2}}\right) \|x^{(i)} - A^{-1}b\|_A \\
&\quad + (1 + \rho)(1 + \varepsilon)\bar{m}_A^+ \bar{\varepsilon} \left((\kappa^{\frac{1}{2}} + \underline{\kappa}) \|A^{-1}b\|_A + \underline{\kappa} \|x^{(i)} - A^{-1}b\|_A \right) \\
&\quad + \varepsilon\kappa^{\frac{1}{2}} (\|x^{(i+1)} - A^{-1}b\|_A + \|A^{-1}b\|_A) \\
&= \varepsilon\kappa^{\frac{1}{2}} \|x^{(i+1)} - A^{-1}b\|_A + (\rho + \delta_3) \|x^{(i)} - A^{-1}b\|_A \\
(4.10) \quad &+ \left(\varepsilon\kappa^{\frac{1}{2}} + (1 + \rho)(1 + \varepsilon)\bar{m}_A^+ \bar{\varepsilon} (\kappa^{\frac{1}{2}} + \underline{\kappa}) \right) \|A^{-1}b\|_A,
\end{aligned}$$

where

$$\delta_3 = (1 + \rho) \left(\varepsilon\kappa^{\frac{1}{2}} + (1 + \varepsilon)\bar{m}_A^+ \bar{\varepsilon} \underline{\kappa} \right).$$

Bound (4.3) now follows by subtracting $\varepsilon\kappa^{\frac{1}{2}} \|x^{(i+1)} - A^{-1}b\|_A$, dividing by $(1 - \varepsilon\kappa^{\frac{1}{2}}) \|A^{-1}b\|_A$ on both sides of (4.10), appealing to (4.1), and noting that

$$\frac{\rho + \delta_3}{1 - \varepsilon\kappa^{\frac{1}{2}}} = \rho + \frac{\delta_3 + \rho\varepsilon\kappa^{\frac{1}{2}}}{1 - \varepsilon\kappa^{\frac{1}{2}}} = \rho + \delta_{\rho_{ir}}.$$

Bound (4.6) follows by tracing (4.3) back to the initial error and noting that

$$\sum_{i=0}^{N-1} (\rho + \delta_{\rho_{ir}})^i \leq \frac{1}{1 - (\rho + \delta_{\rho_{ir}})}.$$

□

Remark 4.2. Limiting Accuracy. Theorem 4.1 establishes convergence in energy of \mathcal{IR} if $\delta_{\rho_{ir}} < 1 - \rho$, but only until the relative error reaches χ . If $\rho \ll 0.9$ for example, then we just want $\delta_{\rho_{ir}} \leq 0.1$ to ensure that $\rho_{ir} \ll 1$. For $\bar{\varepsilon}$ small enough that the terms involving $\bar{\tau} = \kappa\bar{\varepsilon}$ can be ignored, both χ and $\delta_{\rho_{ir}}$ are of order $\tau = \kappa^{\frac{1}{2}}\varepsilon$, so χ tends to be the limiting factor long before the loss in $\delta_{\rho_{ir}}$ becomes an issue. It should also be noted that the inner solver *only* impacts the analysis by way of its convergence factor. Its precision thus has little effect as long as $\rho \ll 1$. It is also important to note that this performance of \mathcal{IR} is optimal in that it reaches an accuracy comparable to that of the finite element solution that has simply been rounded to ε -precision. That is, simply rounding the exact finite element solution of the PDE induces an error of the same order as the limiting accuracy. The mechanism at play here is the discontinuity of fl that allows an arbitrarily small-energy perturbation of a function to induce a large-energy perturbation in the last significant bit. As a simple illustration, suppose that (3.1) is derived from applying Rayleigh-Ritz on a uniform grid h to the *no-flow* 1D Poisson-reaction equation $-u'' + u = f$, $0 \leq x \leq 1$, $u'(0) = u'(1) = 0$. Suppose that the exact solution is $x = 1 + y$, where y is a vector with very small energy: $\|y\|_A < \varepsilon$. Suppose also that y oscillates from positive to negative values from one grid point to the next. Letting $1 + z$ denote x rounded towards zero in ε -precision, then z must oscillate between $-\varepsilon$ and 0. For sufficiently small $\bar{\varepsilon}$, we thus have that

$$\frac{\|1 + z - 1 - y\|_A}{\|1 + z\|_A} \approx \|z\|_A \approx \sqrt{2}h^{-1}\varepsilon = \mathcal{O}(\kappa^{\frac{1}{2}}\varepsilon).$$

This is the same level of error that \mathcal{IR} achieves at its limiting accuracy.

5. Two-grid solver. Analysis of the full multigrid method proceeds in several stages. We begin here by writing the two-grid solver in correction form applied to the \mathcal{IR} residual equation $Ay = r$. (See [5] for an introduction to multigrid methods and principles.)

To understand multigrid solvers and the development that follows, it is important to be clear about terminology. The term *residual* could be confusing if the equation it references is not fully understood. Even for standard multigrid solvers based on the correction scheme, care is needed in using the term *residual* in reference to an equation: a level in the middle of the hierarchy gets its equation as an approximation to the finer-level residual equation, but it too must pass the residual of its approximate residual equation to the next-coarser level. While such ambiguity is usually avoided in the literature, it becomes more crucial here because we need to compute residuals related to residual equations even on the fine level. To avoid this potential confusion, we use the terminology \mathcal{TG} *residual* for those computed within a two-grid cycle, \mathcal{V} *residual* for those computed within a V-cycle, and \mathcal{IR} *residual* for those formed from the full approximation x computed between cycles, and similarly for the equations to which they refer.

The two basic components of any multigrid algorithm are *relaxation* and *prolongation*. We first assume that relaxation is given by the convergent stationary linear iteration $x \leftarrow x - M(Ax - b)$, where $M \in \mathbb{R}^{n \times n}$ is a nonsingular matrix. Matrix M is meant to be an easily computed approximation to A^{-1} , such as the inverse of the diagonal or lower-triangular part of A . (While it is standard in the matrix splitting literature to use M^{-1} , we use M here for simplicity.) Next, assume that prolongation (i.e., interpolation) is given by a matrix $P \in \mathbb{R}^{n \times n_c}$ defined in terms of a coarse level of $n_c < n$ variables. Note that we use subscript c here to signify a coarse-level quantity with no subscript for fine-level quantities. We use this convention when there is no risk of ambiguity, but for more than two levels, subscripts involving level j become necessary.

All computations in the pseudocode \mathcal{TG} in Algorithm 5.1 below are performed in low $\hat{\epsilon}$ -precision (green font), except for the exact, infinite-precision coarse-level solve (black font). Accordingly, since the input RHS may be in higher precision, the cycle is initialized with a rounding step. The solver then proceeds by relaxing on the initial guess (assumed to be zero here) to the solution of the \mathcal{IR} residual equation and then improving the result by a coarse-level correction based on prolongation.

Algorithm 5.1 Two-Grid (\mathcal{TG}) Correction Scheme

Input: $A, r, P, M.$	
1: $r \leftarrow r$	▷ Round RHS and Initialize \mathcal{TG}
2: $y \leftarrow Mr$	▷ Relax on Current Approximation ($y = 0$)
3: $r_{\text{tg}} \leftarrow Ay - r$	▷ Evaluate \mathcal{TG} Residual
4: $b_c \leftarrow P^t r_{\text{tg}}$	▷ Restrict \mathcal{TG} Residual to Coarse-Level
5: $d_c \leftarrow B_c(P^t AP)^{-1} b_c$	▷ Solve Coarse-Level Equation
6: $d \leftarrow Pd_c$	▷ Interpolate Correction to Fine Level
7: $y \leftarrow y - d$	▷ Update Approximate Solution of $Ay = r$
8: return y	▷ Return Approximate Solution of $Ay = r$

\mathcal{TG} incorporates $B_c \in \mathbb{R}^{n_c \times n_c}$ in the coarse-level solution process to allow our theory to apply to the case of approximate solution of the coarse-level equation (as opposed to exact inversion of $P^t AP$ assumed for the standard two-grid case). Note that if $B_c = I_c$, the coarse-level identity matrix, then \mathcal{TG} reverts to a standard two-grid scheme that uses an exact solve on the coarse level. The essential assumption

that we make about \mathcal{TG} when $B_c = I_c$ is that it converges in infinite precision in energy, that is, its error propagation matrix TG is bounded according to

$$(5.1) \quad \|TG\|_A < 1,$$

where $G = I - MA$ is the error propagation matrix for relaxation and $T = I - P(P^tAP)^{-1}P^tA$ is the error propagation matrix for coarse-level correction. Vectors in the ranges of T and P are called *algebraically oscillatory* and *algebraically smooth*, respectively, because they tend to correspond to the respective geometrically oscillatory vectors targeted by relaxation and geometrically smooth vectors targeted by coarse-level correction. See [5]. For the case of general B_c , we assume that $\|B_c - I_c\|_{A_c} < 1$, where $A_c = P^tAP$ and I_c is the coarse-level identity, and we let ρ_{tg}^* denote a bound on the norm of the resulting error propagation matrix:

$$(5.2) \quad \|(I - PB_c(P^tAP)^{-1}P^tA)G\|_A \leq \rho_{tg}^*,$$

where the subscript tg signifies two-grid and the superscript asterisk indicates that it is the infinite-precision factor. To see that bound (5.1) allows us to choose $\rho_{tg}^* < 1$, note that since $I - T$ and T are energy-orthogonal projections onto the range of P and its energy-orthogonal complement, respectively (c.f. [5]), then for any fine-level vector y with $\|y\|_A = 1$ we have that

$$\begin{aligned} & \|(I - PB_c(P^tAP)^{-1}P^tA)Gy\|_A^2 \\ &= \|TGy\|_A^2 + \|P(B_c - I_c)(P^tAP)^{-1}P^tA)Gy\|_A^2 \\ &\leq \|TGy\|_A^2 + \|B_c - I_c\|_{A_c}^2 \|(P^tAP)^{-1}P^tA)Gy\|_{A_c}^2 \\ &= \|TGy\|_A^2 + \|B_c - I_c\|_{A_c}^2 \|(I - T)Gy\|_A^2 \\ &= (1 - \|B_c - I_c\|_{A_c}^2) \|TGy\|_A^2 + \|B_c - I_c\|_{A_c}^2 (\|TGy\|_A^2 + \|(I - T)Gy\|_A^2) \\ &\leq (1 - \|B_c - I_c\|_{A_c}^2) \|TG\|_A^2 + \|B_c - I_c\|_{A_c}^2 < 1. \end{aligned}$$

Until we treat the multigrid solver, it is probably best for the reader to keep in mind that choosing $B_c = I_c$ reverts to the standard two-grid case.

Assume further that relaxation converges monotonically in energy in the sense that $\|G\|_A < 1$ and that P has full rank with at most m_P nonzero entries per row or column. We have in mind that $\|G\|_A$ may be very close to 1 so that relaxation is a poor solver by itself, but that ρ_{tg}^* and, therefore, $\|B_c - I_c\|_{A_c}$ and $\|TG\|_A$ are enough less than 1 that coarse-level correction effectively eliminates algebraic errors that relaxation cannot properly attenuate. However, these conditions are neither specified nor required in what follows.

The objective of \mathcal{TG} is to approximate the exact solution, $A^{-1}r$, of the \mathcal{IR} residual equation

$$(5.3) \quad Ay = r,$$

starting from the initial approximation $y = 0$. Note our use of $A^{-1}r$ here to denote the exact solution of (5.3). As mentioned above, we refer to exact quantities by using expressions like this that characterize them, thereby avoiding the need to introduce additional notation to distinguish between exact and computed quantities. So, while $y = 0$ is of course exact, the \mathcal{TG} residuals and iterates are assumed to be computed quantities subject to rounding errors when and as specified in what follows.

6. Multigrid V-cycle. To correspond to \mathcal{TG} with one relaxation sweep per cycle, we also analyze the so-called $V(1,0)$ -cycle shown in Algorithm 6.1 and illustrated in Fig. 2. Here, processing begins on the finest level and proceeds down through the hierarchy to the coarsest level, with one relaxation sweep performed on each level along the way. Our focus is on one V-cycle defined recursively by the pseudocode \mathcal{V} below on a nested hierarchy of ℓ levels from the coarsest $j = 1$ to the finest $j = \ell$. Each level $j \in \{1, 2, \dots, \ell\}$ is equipped with a system matrix A_j , with $A_\ell = A$. For each $j \in \{2, 3, \dots, \ell\}$, let P_j be the interpolation matrix that maps from level $j - 1$ to level j and let $T_j = I_j - P_j A_{j-1}^{-1} P_j^t A_j$, where I_j is the level j identity matrix. For the coarsest level that involves one relaxation sweep and no further coarsening, we set $P_1 = 0$ so that T_1 is just the identity matrix I_1 . Assume that the following *Galerkin condition* [5] is exactly satisfied on all coarse levels:

$$(6.1) \quad A_{j-1} = P_j^t A_j P_j, \quad 2 \leq j \leq \ell.$$

See [27] for an analysis of the rounding-error effects on computing (6.1).

Algorithm 6.1 $V(1,0)$ -Cycle (\mathcal{V}) Correction Scheme

Input: $A, r, P, \ell \geq 1$ \mathcal{V} levels.

1: $r \leftarrow r$	▷ Round RHS and Initialize \mathcal{V}
2: $y \leftarrow Mr$	▷ Relax on Current Approximation ($y = 0$)
3: if $\ell > 1$ then	▷ Check for Coarser Level
4: $r_v \leftarrow Ay - r$	▷ Evaluate \mathcal{V} Residual
5: $r_{\ell-1} \leftarrow P^t r_v$	▷ Restrict \mathcal{V} Residual to Coarse-Level
6: $d_{\ell-1} \leftarrow \mathcal{V}(A_{\ell-1}, r_{\ell-1}, P_{\ell-1}, \ell - 1)$	▷ Compute Correction from Coarser Levels
7: $d \leftarrow P d_{\ell-1}$	▷ Interpolate Correction to Fine Level
8: $y \leftarrow y - d$	▷ Update Approximate Solution of $Ay = r$
9: end if	
10: return y	▷ Return Approximate Solution of $Ay = r$

The theory in [21] and the references cited therein establish optimal energy convergence in infinite precision of Algorithm 6.1 under fairly general conditions for fully regular elliptic PDEs discretized by standard finite elements. We simply assume this to be the case by supposing that the error propagation matrix V_j for level j is bounded by some $\rho_v^* \in [0, 1)$ for all $j \in \{1, 2, \dots, \ell\}$. Specifically, noting that the error propagation matrices are defined recursively [22] by

$$(6.2) \quad V_1 = G_1 \text{ and } V_j = (P_j V_{j-1} A_{j-1}^{-1} P_j^t A_j + T_j) G_j, \quad 2 \leq j \leq \ell,$$

then we assume that

$$(6.3) \quad \|V_j\|_{A_j} \leq \rho_v^*, \quad 1 \leq j \leq \ell.$$

Because $\|V_1\|_{A_1} = \|G_1\|_{A_1}$, a requirement implied by (6.3) is that $\|G_1\|_{A_1} \leq \rho_v^*$. Such a bound holds in cases that use standard multigrid methods with a sufficiently small coarse level applied to a well-posed coarse-level matrix equation. While this case does not include all potential multigrid applications, it is beyond the scope of the present work to consider situations where full coarsening is difficult or the coarse-level matrices are ill-conditioned.

The principal aim of this paper is an abstract algebraic theory that applies to both algebraic and geometric multigrid when applied to a large class of PDEs. Accordingly, we have in mind matrices whose condition numbers depend on the mesh size, h . (While we do not explicitly exclude coarsening in terms of the degree, p , of the discretization,




FIG. 2. An illustration of the V-cycle correction scheme outlined in Algorithm 6.1 combined with \mathcal{IR} . Once again, the figure is color-coded to denote the precisions associated with the algorithm. The number of boxes in the V-cycle depict the fact that as the multigrid hierarchy is traversed to coarser levels, the required precision can continually be reduced.

our focus is on coarsening in terms of h .) To abstract this h -dependence, define the *pseudo mesh size* $h_j = \kappa_j^{-\frac{1}{2m}}$, $1 \leq j \leq \ell$, where m is a positive integer, and the *pseudo mesh coarsening factor*

$$(6.4) \quad \theta_j = \frac{h_{j-1}}{h_j}, \quad 2 \leq j \leq \ell.$$

In the geometric setting, θ_j and $2m$ correspond to the mesh-refinement factor and order of the PDE, respectively. Under standard assumptions for finite element discretizations, classical theory shows that the condition number on a given grid is bounded by a constant (depending on the element order) times h_{min}^{-2m} , where h_{min} is the smallest element size on that grid (see [25, Sec. 5.2]). Our abstract parameter h_j is therefore bounded by that constant times the grid j mesh size.

To allow a *progressive-precision* V-cycle, where precision is tailored to each level in the hierarchy, assume now that $\dot{\epsilon}$ varies by letting $\dot{\epsilon}_j$ denote the unit roundoff used on level j , $1 \leq j \leq \ell$. (We use $\dot{\epsilon}$ without subscripts when the level is understood.) Specifically, $\dot{\epsilon}_j$ -precision is used on level j to store the data, perform relaxation, transfer residuals to level $j - 1$ and corrections to level $j + 1$, and round residuals transferred from level $j + 1$. Define the *precision coarsening factor* by

$$(6.5) \quad \dot{\zeta}_j = \frac{\dot{\epsilon}_{j-1}}{\dot{\epsilon}_j}, \quad 2 \leq j \leq \ell.$$

To accommodate our use of a geometric series involving the rounding-error effects on each level of the V-cycle, let $\vartheta = \min_{1 \leq j \leq \ell} \{\theta_j \dot{\zeta}_j^{-\frac{1}{m}}\}$. The estimate for the overall effect of rounding error on V-cycle convergence obtained in Theorem 8.1 suggests that ϑ should be substantially larger than 1 (i.e. $\dot{\zeta}_j \ll \theta_j^m$); this scale of ϑ would ensure that the V-cycle convergence factors are bounded nicely in terms of the two-grid factors. (We are not suggesting here that $\bar{\epsilon}_j$ should be large, which could worsen rounding-error effects, but rather that ϵ_j should be sufficiently larger than $\bar{\epsilon}_j$.) However, our only formal assumption for our abstract theory is simply that $\vartheta > 1$.

Only the low precision varies by level in the V-cycle because its finest level is fixed. But FMG's outer loop uses progressively finer levels for the inner loop's finest levels, thus enabling variable ϵ_j and $\bar{\epsilon}_j$, $1 \leq j \leq \ell$, where ℓ is assumed to be the very finest level used in the FMG scheme.

7. Convergence Theory: Two-grid. In what follows, we develop certain bounds that are written in terms of various parameters such as the convergence estimate in (5.2) involving ρ_{tg}^* . Nothing other than what is initially stated about these

bounds is assumed until we make conclusions about when and how well \mathcal{IR} actually converges. However, useful prototypes for this work are the matrix equations associated with a hierarchy of grids that arise from discretizing a PDE such as the model biharmonic equation treated in [27]. In this case, we have in mind parameters that are fixed constants so that the bounds hold uniformly in the mesh size. Our references to “optimal” and “optimally” here are only meant to suggest the loose concept that the corresponding bounds and parameters should be nice in some way. For example, ρ_{tg}^* should not be very close to 1. However, in the PDE context, we mean to suggest that these qualifiers also connote a sense that the bounds and parameters hold uniformly in the mesh size. In any event, our aim is to provide a framework that can be used to confirm optimal performance for specific applications.

Several additional parameters appear below in our convergence estimates, such as the condition number of P^tP denoted by $\kappa(P^tP) = \|P^tP\| \cdot \|(P^tP)^{-1}\|$, and the following quantities that are defined just for convenience:

$$\dot{m}_P^+ = \frac{m_P}{1 - m_P\dot{\epsilon}}, \quad \dot{m}_A^+ = \frac{m_A + 1}{1 - (m_A + 1)\dot{\epsilon}}, \quad \text{and} \quad \mu = 3\dot{\zeta}\kappa^{\frac{1}{2}}(P^tP)\dot{m}_P^+.$$

To account for rounding errors in relaxation, suppose that a constant α_M exists such that computing Mz for a vector $z \in \mathbb{R}^n$ in $\dot{\epsilon}$ -precision yields the result

$$(7.1) \quad Mz + \delta_M, \quad \|\delta_M\| \leq \alpha_M\dot{\epsilon}\|z\|.$$

Our initial two-grid result, in the form of a lemma, acts as a stepping stone for the analysis of the full \mathcal{TG} algorithm that follows. For the simpler “Limited \mathcal{TG} ” scheme considered here, we assume that finite precision is used only in the transfers between levels. We therefore ignore the initial rounding step and assume that the coarse-level solve and all fine-level computations are done exactly. For this lemma and our second theorem, we only analyze one cycle of \mathcal{TG} , with initial guess $y = 0$. The exact initial algebraic error is therefore just $y - A^{-1}r = -A^{-1}r$, so relaxation yields the new iterate Mr with error $-GA^{-1}r$, and the \mathcal{TG} residual to be transferred to the coarse level is just $-AGA^{-1}r$. To accommodate progressive precision, assume that the coarse level uses $\dot{\epsilon}_c = \dot{\zeta}\dot{\epsilon}$ precision, where $\dot{\zeta} \geq 1$ (see (6.5)).

LEMMA 7.1. *Limited \mathcal{TG} . Consider a limited version of the two-grid correction scheme, where all computations are in infinite precision except for $\dot{\epsilon}_c$ -precision computation of b_c in step 4 and d in step 6 of Algorithm 5.1. Then the result of one such limited \mathcal{TG} cycle yields a result y with error $y - A^{-1}r$ that satisfies*

$$\|y - A^{-1}r\|_A \leq \rho_{itg}\|A^{-1}r\|_A,$$

where $\rho_{itg} = \rho_{tg}^* + \delta_{\rho_{itg}}$ and $\delta_{\rho_{itg}} = \rho_{tg}^*\mu\dot{\tau} + 3\mu\dot{\tau} + 2\mu^2\dot{\tau}^2$.

Proof. Superscripts (0), $(\frac{1}{2})$, and (1) are used here to keep track of the errors e and iterates y , with $e^{(0)} = -A^{-1}r$ and $y^{(0)} = 0$. The proof proceeds by treating in turn three cases based on where rounding error is assumed to occur:

- *Case 1 (Restriction):* All computations are in infinite precision except for finite-precision computation of b_c in step 4 of \mathcal{TG} .
- *Case 2 (Interpolation):* All computations are in infinite precision except for finite-precision computation of d in step 6 of \mathcal{TG} .
- *Case 3 (Restriction and Interpolation):* All computations are in infinite precision except for finite-precision computations in steps 4 and 6 of \mathcal{TG} .

Beginning with *Case 1*, let each entry of the vector δ be the error in computing the corresponding entry of $P^t r_{\text{tg}}$. This task is just an inner product between P 's i^{th} column and r_{tg} , which, appealing to just the nonzero entries of P , are vectors of length at most m_P . In analogy to (2.5) and remembering that $b = 0$ and $r_{\text{tg}} = AGe^{(0)}$ is exact, we have the loose bound

$$\begin{aligned} |\delta| &\leq (\dot{\epsilon}_c + (1 + \dot{\epsilon}_c)\dot{m}_P^+ \dot{\epsilon}) |P^t AGe^{(0)}| \\ &\leq (\dot{\zeta} \dot{\epsilon} + 2\dot{m}_P^+ \dot{\epsilon}) |P^t AGe^{(0)}| \\ &\leq 3\dot{\zeta} \dot{m}_P^+ \dot{\epsilon} P^t |AGe^{(0)}|. \end{aligned}$$

Since $\|P^t\| = \|P\| = \|P^t P\|^{\frac{1}{2}}$, we therefore have that

$$(7.2) \quad \|\delta\| \leq 3\dot{\zeta} \dot{m}_P^+ \dot{\epsilon} \|P^t P\|^{\frac{1}{2}} \|AGe^{(0)}\| \leq 3\dot{\zeta} \dot{m}_P^+ \dot{\epsilon} \|P^t P\|^{\frac{1}{2}} \|A\|^{\frac{1}{2}} \|Ge^{(0)}\|_A.$$

Since $\|Ge^{(0)}\|_A \leq \|e^{(0)}\|_A$ by assumption, then

$$(7.3) \quad \|\delta\| \leq 3\dot{\zeta} \dot{m}_P^+ \dot{\epsilon} \|P^t P\|^{\frac{1}{2}} \|A\|^{\frac{1}{2}} \|e^{(0)}\|_A.$$

Remembering that $y^{(\frac{1}{2})} = Mr^{(0)}$ is the intermediate iterate formed by relaxation on $y^{(0)} = 0$, note that the final computed update in the \mathcal{TG} cycle is given by

$$\begin{aligned} y^{(1)} &= y^{(\frac{1}{2})} - d \\ &= y^{(\frac{1}{2})} - PB_c(P^t AP)^{-1}(P^t AGe^{(0)} + \delta) \\ &= \underbrace{y^{(\frac{1}{2})} - PB_c(P^t AP)^{-1}P^t AGe^{(0)}}_{\text{exact update}} - \underbrace{PB_c(P^t AP)^{-1}\delta}_{\text{propagated error}}. \end{aligned}$$

Subtracting the exact solution, $A^{-1}r^{(0)}$, of the \mathcal{IR} residual equation from both sides and noting that $e^{(\frac{1}{2})} = Ge^{(0)}$ yields

$$\begin{aligned} e^{(1)} &= e^{(\frac{1}{2})} - PB_c(P^t AP)^{-1}P^t AGe^{(0)} - PB_c(P^t AP)^{-1}\delta \\ &= (I - PB_c(P^t AP)^{-1}P^t A)Ge^{(0)} - PB_c(P^t AP)^{-1}\delta. \end{aligned}$$

Taking energy norms of both sides yields

$$\begin{aligned} \|e^{(1)}\|_A &\leq \|(I - PB_c(P^t AP)^{-1}P^t A)Ge^{(0)}\|_A + \|B_c(P^t AP)^{-1}\delta\|_{A_c} \\ &\leq \|(I - PB_c(P^t AP)^{-1}P^t A)Ge^{(0)}\|_A + \|B_c\|_{A_c} \|(P^t AP)^{-1}\delta\|_{A_c} \\ (7.4) \quad &\leq \|(I - PB_c(P^t AP)^{-1}P^t A)Ge^{(0)}\|_A + 2\|(P^t AP)^{-\frac{1}{2}}\| \cdot \|\delta\|. \end{aligned}$$

But $\|(P^t AP)^{-\frac{1}{2}}\| \leq \|A^{-\frac{1}{2}}\| \cdot \|(P^t P)^{-\frac{1}{2}}\|$ follows from noting that if $z \neq 0$ is a coarse-level eigenvector of $P^t AP$ belonging to the smallest eigenvalue, λ_c , of $P^t AP$, then the smallest eigenvalue, λ , of A satisfies

$$\lambda \leq \frac{\langle APz, Pz \rangle}{\langle Pz, Pz \rangle} = \frac{\langle APz, Pz \rangle}{\langle z, z \rangle} \cdot \frac{\langle z, z \rangle}{\langle Pz, Pz \rangle} = \lambda_c \frac{\langle (P^t P)^{-1}Pz, Pz \rangle}{\langle Pz, Pz \rangle} \leq \lambda_c \|(P^t P)^{-1}\|.$$

Remembering that $\mu = 3\dot{\zeta}\kappa^{\frac{1}{2}}(P^t P)\dot{m}_P^+$ and $\dot{\tau} = \kappa\dot{\epsilon}$, then (5.2), (7.3), and (7.4) combine-level to yield

$$(7.5) \quad \|e^{(1)}\|_A \leq \left(\rho_{tq}^* + 2\kappa^{\frac{1}{2}}(P^t P)\kappa^{\frac{1}{2}}3\dot{\zeta}\dot{m}_P^+ \dot{\epsilon} \right) \|e^{(0)}\|_A = (\rho_{tq}^* + 2\mu\dot{\tau})\|e^{(0)}\|_A.$$

The energy convergence factor for *Case 1* is therefore bounded by $\rho_1 = \rho_{tg}^* + 2\mu\dot{\tau}$.

Consider now *Case 2*. The symbols we use here are the same as before, but defined differently now to suit this case. Accordingly, let δ be the error in the correction term d incurred due to $\dot{\epsilon}_c$ -precision interpolation of d_c in step 6 of \mathcal{TG} . From (2.3), we then have that $|\delta| \leq \dot{m}_P^+ \dot{\epsilon}_c P |d_c|$, which yields a bound similar to but simpler than (7.3) because there is no need for rounding: $\|\delta\| \leq \dot{m}_P^+ \dot{\epsilon}_c \|P^t P\|^{\frac{1}{2}} \|d_c\|$. The squared relative energy norm of the error caused by this computed correction is then bounded loosely as follows:

$$(7.6) \quad \frac{\langle A\delta, \delta \rangle}{\langle APd_c, Pd_c \rangle} \leq \|A\| \cdot \|(P^t AP)^{-1}\| \frac{\|\delta\|^2}{\|d_c\|^2} \leq \kappa(P^t P) \kappa \zeta^2 (\dot{m}_P^+)^2 \dot{\epsilon}_c^2 \leq \mu^2 \dot{\tau}^2.$$

With $e^{(0)}$ and $e^{(1)}$ again denoting the respective initial and final errors and $e^{(\frac{1}{2})}$ the error after relaxation, then the computed coarse-level update becomes

$$e^{(1)} = \underbrace{e^{(\frac{1}{2})} - Pd_c}_{\text{exact update}} + \underbrace{\delta}_{\dot{\epsilon}_c \text{ error}}.$$

Here we invoke the only property of d_c that we need, that is, that its interpolation and correction to the fine-level iterate reduces the error by a factor of at least ρ_{tg}^* :

$$\|e^{(\frac{1}{2})} - Pd_c\|_A \leq \rho_{tg}^* \|e^{(0)}\|_A.$$

We thus obtain the following bound on the energy convergence factor:

$$\begin{aligned} \|e^{(1)}\|_A &= \|e^{(\frac{1}{2})} - Pd_c + \delta\|_A \\ &\leq \|e^{(\frac{1}{2})} - Pd_c\|_A + \|\delta\|_A \\ &\leq \rho_{tg}^* \|e^{(0)}\|_A + \mu\dot{\tau} \|Pd_c\|_A \\ &= \rho_{tg}^* \|e^{(0)}\|_A + \mu\dot{\tau} \|e^{(\frac{1}{2})} - Pd_c - e^{(\frac{1}{2})}\|_A \\ &\leq \rho_{tg}^* \|e^{(0)}\|_A + \mu\dot{\tau} \left(\rho_{tg}^* \|e^{(0)}\|_A + \|e^{(\frac{1}{2})}\|_A \right) \\ &\leq ((1 + \mu\dot{\tau}) \rho_{tg}^* + \mu\dot{\tau}) \|e^{(0)}\|_A. \quad \square \end{aligned}$$

To summarize, the result for *Case 1* is that the two-grid solver converges in energy with factor bounded by $\rho_1 = \rho_{tg}^* + 2\mu\dot{\tau}$, while for *Case 2* the factor is bounded by $\rho_2 = (1 + \mu\dot{\tau}) \rho_{tg}^* + \mu\dot{\tau}$. But *Case 2* made no use of the specific form of d_c : the only property of z that was actually used was that it led to a reduction in the energy norm by a factor bounded by ρ_{tg}^* . So we can easily analyze *Case 3* by just replacing ρ_{tg}^* by ρ_1 in the expression for ρ_2 , thus proving the lemma.

To rein in complexity in what follows, we are occasionally loose with upper bounds. Accordingly, for the two-grid estimate, let σ be any constant such that

$$(7.7) \quad \sigma \geq (1 + \dot{\epsilon}) \max\{\alpha_M \|A\|, \psi \alpha_M, \psi \|M\|\}.$$

THEOREM 7.2. \mathcal{TG} . *One cycle of Algorithm 5.1 applied to (5.3) converges according to*

$$(7.8) \quad \|y - A^{-1}r\|_A \leq \rho_{tg} \|A^{-1}r\|_A, \quad \rho_{tg} = \rho_{tg}^* + \delta_{\rho_{tg}},$$

provided $\dot{\tau}$ is small enough that $\delta_{\rho_{tg}} < 1 - \rho_{tg}^*$, where $\delta_{\rho_{tg}} = \delta_{\rho_{tg}}(\dot{\tau}) = a_1 \dot{\tau} + a_2 \dot{\tau}^2 + a_3 \dot{\tau}^3$ is a cubic polynomial in $\dot{\tau}$ with coefficients given by $a_1 = 4 + \sigma + 4\mu$, $a_2 = 2(3 + \sigma + \dot{m}_A^+(1 + \sigma))\mu + 2 + \sigma$, and $a_3 = 2(1 + \sigma + \dot{m}_A^+(1 + \sigma))\mu$.

Proof. The \mathcal{TG} cycle in Algorithm 5.1 uses $\hat{\epsilon}$ -precision for all finite-precision computations. Its objective is to approximate the solution $A^{-1}r$ of (5.3). Using superscripts as in the previous proof, we first estimate how the \mathcal{TG} cycle approximates the solution $A^{-1}r^{(0)}$ of $Ay = r^{(0)}$, where $r^{(0)} = r + \delta_0$, $\|\delta_0\| \leq \hat{\epsilon}\|r\|$, is r rounded to $\hat{\epsilon}$ -precision. Specifically, this proof is mostly devoted to showing that the final computed approximation $y^{(1)}$ satisfies

$$(7.9) \quad \|y^{(1)} - A^{-1}r^{(0)}\|_A \leq \hat{\rho}_{tg} \|A^{-1}r^{(0)}\|_A, \quad \hat{\rho}_{tg} = \rho_{tg}^* + \hat{\delta}_{\rho_{tg}},$$

where

$$(7.10) \quad \hat{\delta}_{\rho_{tg}} = (\dot{\tau} + \mu\dot{\tau} + \mu\dot{\tau}^2)(\rho_{tg}^* + 2\beta\mu\dot{\tau}) + 2\beta\mu\dot{\tau} + (1 + \dot{\tau})\mu\dot{\tau} + (1 + \hat{\epsilon})\alpha_M \|A\| \hat{\epsilon} + \dot{\tau},$$

with $\beta = 1 + (\alpha_M \|A\| + \dot{m}_A^+ (1 + \psi \|M\| + \psi \alpha_M \hat{\epsilon})) \hat{\epsilon}$. We would be done if we could establish (7.9)-(7.10) because we could then use (5.2) and the observations that $\|r\| \leq \|A^{\frac{1}{2}}\| \cdot \|A^{-1}r\|_A$, $\hat{\epsilon} \leq \dot{\tau}$, and $\rho_{tg}^* + 2\beta\mu\dot{\tau} \leq \rho_{tg} < 1$ to prove (7.8) as follows:

$$\begin{aligned} \|y^{(1)} - A^{-1}r\|_A &\leq \|y^{(1)} - A^{-1}r^{(0)}\|_A + \|A^{-1}(r - r^{(0)})\|_A \\ &\leq \hat{\rho}_{tg} \|A^{-1}r^{(0)}\|_A + \|A^{-1}(r - r^{(0)})\|_A \\ &\leq \hat{\rho}_{tg} \|A^{-1}r\|_A + (1 + \hat{\rho}_{tg}) \|A^{-1}(r - r^{(0)})\|_A \\ &\leq \hat{\rho}_{tg} \|A^{-1}r\|_A + (1 + \hat{\rho}_{tg}) \|A^{-\frac{1}{2}}\| \hat{\epsilon} \|r\| \\ &\leq \hat{\rho}_{tg} \|A^{-1}r\|_A + (1 + \hat{\rho}_{tg}) \dot{\tau} \|A^{-1}r\|_A \\ &= (\rho_{tg}^* + (1 + \rho_{tg}^*)\dot{\tau} + (1 + \dot{\tau}) [(\dot{\tau} + \mu\dot{\tau} + \mu\dot{\tau}^2)(\rho_{tg}^* + 2\beta\mu\dot{\tau}) \\ &\quad + 2\beta\mu\dot{\tau} + (1 + \dot{\tau})\mu\dot{\tau} + (1 + \hat{\epsilon})\alpha_M \|A\| \hat{\epsilon} + \dot{\tau}]) \|A^{-1}r\|_A \\ &\leq (\rho_{tg}^* + 2\dot{\tau} + (1 + \dot{\tau}) [\dot{\tau} + \mu\dot{\tau} + \mu\dot{\tau}^2 + 2\mu\dot{\tau} \\ &\quad + 2(\sigma + \dot{m}_A^+(1 + \sigma))\mu\dot{\tau}^2 + (1 + \dot{\tau})\mu\dot{\tau} + (1 + \sigma)\dot{\tau}]) \|A^{-1}r\|_A \\ &= \rho_{tg} \|A^{-1}r\|_A. \end{aligned}$$

To establish (7.9)-(7.10), we begin by accounting for the rounding-error effects in relaxation in step 2 and residual computation in step 3. Since $y^{(0)} = 0$, (5.3) becomes $Ay = -Ae^{(0)}$, where the error in $y^{(0)}$ is $e^{(0)} = y^{(0)} - A^{-1}r^{(0)} = -A^{-1}r^{(0)}$. Using $y^{(0)} = 0$ again shows that relaxation computed in $\hat{\epsilon}$ -precision yields

$$(7.11) \quad y^{(\frac{1}{2})} = \underbrace{Mr^{(0)}}_{\text{exact iterate}} + \underbrace{\delta_M}_{\hat{\epsilon} \text{ error}}, \quad \|\delta_M\| \leq \alpha_M \hat{\epsilon} \|r^{(0)}\|.$$

The next step is to compute the \mathcal{TG} residual in $\hat{\epsilon}$ -precision for transfer to the coarse level. Using (2.4), we have that

$$(7.12) \quad r^{(\frac{1}{2})} = \underbrace{Ay^{(\frac{1}{2})} - r^{(0)}}_{\text{exact residual}} + \underbrace{\delta_1}_{\hat{\epsilon} \text{ error}}, \quad |\delta_1| \leq \dot{m}_A^+ \hat{\epsilon} (|r^{(0)}| + |A| \cdot |y^{(\frac{1}{2})}|).$$

But the algebraic errors $y^{(\frac{1}{2})} - A^{-1}r^{(0)}$ and $e^{(0)} = y^{(0)} - A^{-1}r^{(0)}$ are related according to $y^{(\frac{1}{2})} - A^{-1}r^{(0)} = Ge^{(0)} + \delta_M$, so $Ay^{(\frac{1}{2})} - r^{(0)} = A(Ge^{(0)} + \delta_M)$. Letting $\delta_2 = A\delta_M + \delta_1$, we can therefore rewrite the computed \mathcal{TG} residual as

$$r^{(\frac{1}{2})} = \underbrace{AGe^{(0)}}_{\text{exact residual}} + \underbrace{\delta_2}_{\text{propagated } \hat{\epsilon} \text{ error}},$$

where

$$|\delta_2| \leq |A\delta_M| + |\delta_1| \leq |A\delta_M| + \dot{m}_A^+ \dot{\epsilon} \left(|r^{(0)}| + |A| \cdot |Mr^{(0)} + \delta_M| \right).$$

By (7.11), the observation that $\|r^{(0)}\| = \|Ae^{(0)}\| \leq \|A\|^{\frac{1}{2}} \|e^{(0)}\|_A$, other similar matrix norm bounds, and the triangle inequality, we then have that

$$\begin{aligned} \|\delta_2\| &\leq \alpha_M \dot{\epsilon} \|A\| \cdot \|r^{(0)}\| + \dot{m}_A^+ \dot{\epsilon} (1 + \psi \|M\| + \psi \alpha_M \dot{\epsilon}) \|r^{(0)}\| \\ &\leq (\alpha_M \dot{\epsilon} \|A\| + \dot{m}_A^+ \dot{\epsilon} (1 + \psi \|M\| + \psi \alpha_M \dot{\epsilon})) \|A\|^{\frac{1}{2}} \|e^{(0)}\|_A. \end{aligned}$$

For the rounding error δ_3 incurred in transferring the computed residual to the coarse level, we can thus mimic (7.2)-(7.3) with δ replaced by δ_3 and $AGe^{(0)}$ replaced by $AGe^{(0)} + \delta_2$ to obtain

$$\|\delta_3\| \leq 3\dot{\zeta} \dot{m}_P^+ \dot{\epsilon} \|P^t P\|^{\frac{1}{2}} \|AGe^{(0)} + \delta_2\|,$$

which leads directly to the bound

$$(7.13) \quad \|\delta_3\| \leq 3\beta \dot{\zeta} \dot{m}_P^+ \dot{\epsilon} \|P^t P\|^{\frac{1}{2}} \|A\|^{\frac{1}{2}} \|e^{(0)}\|_A.$$

Note that (7.13) is just (7.3) with the extra leading factor β . We can therefore proceed as we did after (7.3), but with μ replaced by $\beta\mu$ and the convergence factor corresponding to Case 1 of Lemma 7.1 now reading $\rho_1 = \rho_{tg}^* + 2\beta\mu\dot{\tau}$. The rest of the lemma applies directly with μ unchanged, so letting d denote the computed correction from the coarse level, we could then conclude that the exact update $Mr^{(0)} - d$ of the exact intermediate iterate $Mr^{(0)}$ converges to $A^{-1}r^{(0)}$ in energy with factor bounded by $\rho_3 = (1 + \mu\dot{\tau})(\rho_{tg}^* + 2\beta\mu\dot{\tau}) + \mu\dot{\tau}$:

$$(7.14) \quad \|Mr^{(0)} - d - A^{-1}r^{(0)}\|_A \leq \rho_3 \|e^{(0)}\|_A.$$

However, to bound the error in the computed quantities, we must now account for rounding errors in $y^{(\frac{1}{2})}$ due to the use of M and in the subtraction $y^{(\frac{1}{2})} - d$.

To this end, note that the error in the computed iterate $y^{(1)}$ can be written as

$$e^{(1)} = y^{(1)} - A^{-1}r^{(0)} = (Mr^{(0)} + \delta_M - d)(1 + \delta_4) - A^{-1}r^{(0)},$$

where the rounding error δ_4 due to the subtraction satisfies $|\delta_4| \leq \dot{\epsilon}$ (see (2.1)). Taking energy norms of both sides and rearranging terms yields

$$(7.15) \quad \begin{aligned} \|e^{(1)}\|_A &= \|(Mr^{(0)} - d - A^{-1}r^{(0)})(1 + \delta_4) + \delta_M(1 + \delta_4) + \delta_4 A^{-1}r^{(0)}\|_A \\ &\leq \|(Mr^{(0)} - d - A^{-1}r^{(0)})(1 + \delta_4)\|_A + \|\delta_M(1 + \delta_4)\|_A + \|\delta_4 A^{-1}r^{(0)}\|_A. \end{aligned}$$

Using (7.14) and the fact that $\|w\| \leq \|A^{-\frac{1}{2}}\| \cdot \|w\|_A$ for any w , then the first term on the right is bounded according to

$$(7.16) \quad \begin{aligned} &\|(Mr^{(0)} - d - A^{-1}r^{(0)})(1 + \delta_4)\|_A \\ &\leq \|(Mr^{(0)} - d - A^{-1}r^{(0)})\|_A + \|(Mr^{(0)} - d - A^{-1}r^{(0)})\delta_4\|_A \\ &\leq \rho_3 \|e^{(0)}\|_A + \|A^{\frac{1}{2}}\| \cdot \|(Mr^{(0)} - d - A^{-1}r^{(0)})\delta_4\| \\ &\leq \rho_3 \|e^{(0)}\|_A + \dot{\epsilon} \|A^{\frac{1}{2}}\| \cdot \|Mr^{(0)} - d - A^{-1}r^{(0)}\| \\ &\leq \rho_3 \|e^{(0)}\|_A + \dot{\tau} \|Mr^{(0)} - d - A^{-1}r^{(0)}\|_A \\ &\leq (1 + \dot{\tau}) \rho_3 \|e^{(0)}\|_A. \end{aligned}$$

Since $r^{(0)} = -Ae^{(0)}$, then the remaining two terms are bounded according to

$$\begin{aligned}
\|\delta_M(1 + \delta_4)\|_A + \|\delta_4 A^{-1} r^{(0)}\|_A &\leq \|A^{\frac{1}{2}}\| \left(\|\delta_M(1 + \delta_4)\| + \|\delta_4 A^{-1} r^{(0)}\| \right) \\
&\leq (1 + \dot{\epsilon}) \dot{\epsilon} \alpha_M \|A^{\frac{1}{2}}\| \cdot \|r^{(0)}\| + \dot{\tau} \|A^{-\frac{1}{2}} r^{(0)}\| \\
(7.17) \qquad \qquad \qquad &\leq ((1 + \dot{\epsilon}) \dot{\epsilon} \alpha_M \|A\| + \dot{\tau}) \|e^{(0)}\|_A. \quad \square
\end{aligned}$$

Bounds (7.15)-(7.17) combine to establish (7.9)-(7.10) and thus prove the theorem.

Remark 7.3. Mixed Precision Rationale. Theorems 4.1 and 7.2 suggest why mixed precision is needed and why it works. Both the inner and outer iterations compute residuals: \mathcal{TG} in (7.12) and \mathcal{IR} in (4.7). However, while the estimates for \mathcal{TG} only involve $\kappa^{\frac{1}{2}}$ so that low precision suffices, the appearance of κ in the estimates for \mathcal{IR} indicates its need for high precision. The key point that makes multigrid cycles much less sensitive to rounding error than \mathcal{IR} and possibly any other solver is that it consists of a fixed number of simple arithmetic steps. In fact, because \mathcal{TG} starts with a zero initial guess, the \mathcal{TG} iterate is just M times the RHS of its target equation (5.3). The sequence of relatively simple steps of this relaxation scheme and the subsequent intergrid transfers is the only source of rounding error in this simple multigrid cycle. In contrast, while successive \mathcal{IR} iterates can be written in terms of b in (3.1), that expression becomes increasingly complex as the iterations proceed due to accumulation of the many calculations over successive \mathcal{IR} iterations. This complexity of \mathcal{IR} accounts for the need for high precision, which coupled with the simplicity of the V-cycle explains the success of the mixed-precision solver.

8. Convergence Theory: Multigrid V-cycle. Our next theorem extends Theorem 7.2 to the multilevel case with progressive precision by confirming that one V-cycle reduces the error optimally toward the solution of the target equation $Ay = r$ provided the derived expression for δ_{ρ_v} is less than $1 - \rho_v^*$. This proviso means that coarsening in the level hierarchy should be fast enough and progression of the precision should be slow enough that the factors in (6.4) and (6.5), respectively, should be small enough to ensure that $\vartheta \gg 1$. It further requires that $\kappa_j \ll \dot{\epsilon}_j^{-2}$, $1 \leq j \leq \ell$. Theorem 4.1 then confirms that the mixed-precision version of \mathcal{IR} with a V-cycle as the inner loop converges optimally to the solution of (3.1) until the limiting accuracy χ is reached.

To proceed, we need to make assumptions about the parameters in Theorem 7.2 and how they behave on all levels. Writing α_{M_j} for the parameter α_M on level $j \in \{1, 2, \dots, \ell\}$, it is easy to see that α_{M_j} is of order $\frac{1}{\|A_j\|}$ for Richardson iteration, where $M_j = s_j I_j$: we assume that the energy norm of its error propagation matrix $G_j = I_j - s_j A_j$ is less than 1, so s_j must be positive and bounded according to

$$s_j = \frac{\|s_j A_j\|_{A_j}}{\|A_j\|} \leq \frac{\|I_j\|_{A_j} + \|I_j - s_j A_j\|_{A_j}}{\|A_j\|} < \frac{2}{\|A_j\|},$$

from which by (2.1) and (7.1) it follows that $\alpha_{M_j} < \frac{2}{\|A_j\|}$. The error propagation operator for Jacobi is $G_j = I_j - s_j D_j^{-1} A_j$, where D_j is the diagonal of A_j and $s_j = \frac{\omega_j}{\|D_j^{-1/2} A_j D_j^{-1/2}\|}$ for some fixed positive relaxation parameter $\omega_j < 2$. Applying (2.1) twice, it is easy to see that $\text{fl}(x \circ \text{fl}(y \circ z)) = x \circ y \circ z(1 + \delta)$, where $|\delta| \leq \epsilon(2 + \epsilon)$, from which it follows that $\alpha_{M_j} = \mathcal{O}\left(\frac{\kappa(D_j)}{\|A_j\|}\right)$.

These estimates are typical of relaxation methods applied to elliptic PDEs because they can only effectively target the upper end of the spectrum of A with the simple way

they approximate A^{-1} . Considering any stronger relaxation scheme that captures the lower end of the spectrum (think of A^{-1} in the extreme case) would beg the question as to how one would solve the resulting linear system that must accompany this stronger method. To exclude such a consideration, assume now that (7.7) holds for fixed σ on all levels (e.g., $\sigma \geq (1 + \hat{\epsilon}_j) \max\{\alpha_{M_j} \|A_j\|, \alpha_{M_j} \| \|A_j\| \|, \|M_j\| \cdot \| \|A_j\| \| \}, 1 \leq j \leq \ell$). We also now redefine the following parameters to mean their maxima over all levels:

$$\begin{aligned} \kappa(P^t P) &= \max_{1 \leq j \leq \ell} \kappa(P_j^t P_j), \dot{m}_A^\dagger = \max_{1 \leq j \leq \ell} \frac{m_{A_j}}{1 - m_{A_j} \hat{\epsilon}_j}, \dot{m}_P^\dagger = \max_{1 \leq j \leq \ell} \frac{m_{P_j}}{1 - m_{P_j} \hat{\epsilon}_j}, \\ \bar{m}_A^\dagger &= \max_{1 \leq j \leq \ell} \frac{m_{A_j}}{1 - m_{A_j} \bar{\epsilon}_j}, m_P^\dagger = \max_{1 \leq j \leq \ell} \frac{m_{P_j}}{1 - m_{P_j} \epsilon_j}, \mu = 3(\max_{1 \leq j \leq \ell} \zeta_j) \kappa^{\frac{1}{2}}(P^t P) \dot{m}_P^\dagger. \end{aligned}$$

These assumptions are valid for a very general class of elliptic PDEs. They allow us to interpret the coefficients of $\delta_{\rho_{tg}} = \delta_{\rho_{tg}}(\hat{\tau})$ in Theorem 7.2 to be the maxima over j . Note then that all of the coefficients in $\delta_{\rho_{tg}}$ are positive. The same expressions hold for each grid level j , but with $\hat{\tau}$ replaced by $\hat{\tau}_j = \kappa_j^{\frac{1}{2}} \hat{\epsilon}_j$. Let $\delta_{\rho_{tg}}(\hat{\tau}_j)$ denote $\delta_{\rho_{tg}}$ with the implied argument $\hat{\tau}$ replaced by $\hat{\tau}_j$, $1 \leq j \leq \ell$. Since $\delta_{\rho_{tg}}(\hat{\tau}_j)$ consists of a strictly linear term in $\hat{\tau}_j$ plus higher-order terms, it is easy to see by (6.4) and (6.5) that

$$(8.1) \quad \delta_{\rho_{tg}}(\hat{\tau}_{j-1}) = \delta_{\rho_{tg}}(h_{j-1}^{-m} \hat{\epsilon}_{j-1}) = \delta_{\rho_{tg}}(\theta_j^{-m} \zeta_j h_j^{-m} \hat{\epsilon}_j) \leq \vartheta^{-m} \delta_{\rho_{tg}}(\hat{\tau}_j), \quad 2 \leq j \leq \ell.$$

THEOREM 8.1. *\mathcal{V} . If $\hat{\tau}_j$ is small enough to ensure that $\delta_{\rho_v}(\hat{\tau}_j) = \frac{\vartheta^m}{\vartheta^m - 1} \delta_{\rho_{tg}}(\hat{\tau}_j) < 1 - \rho_v^*$, $1 \leq j \leq \ell$, then one cycle of the progressive-precision version of Algorithm 6.1 for solving the \mathcal{IR} residual equation converges according to*

$$(8.2) \quad \|y - A^{-1}r\|_A \leq \rho_v \|A^{-1}r\|_A, \rho_v = \rho_v^* + \delta_{\rho_v}, \delta_{\rho_v} = \delta_{\rho_v}(\hat{\tau}_\ell) = \frac{\vartheta^m}{\vartheta^m - 1} \delta_{\rho_{tg}}(\hat{\tau}_\ell).$$

Proof. Note from (6.2) and the definition of T_j that

$$V_j = (I_j - P_j(I_{j-1} - V_{j-1})A_{j-1}^{-1}P_j^t)G_j,$$

which is the error propagation matrix for \mathcal{TG} on the “finest” level j with $B_{j-1} = I_{j-1} - V_{j-1}$. Thus, step 6 of \mathcal{V} is step 5 of \mathcal{TG} with this choice of $B_c = I_{j-1} - V_{j-1}$ and \mathcal{V} is \mathcal{TG} with the coarse level solved recursively. Note that we satisfy our requirement on B_c here because $\|B_{j-1} - I_{j-1}\|_{A_{j-1}} = \|V_{j-1}\|_{A_{j-1}} \leq \rho_v^* < 1$. Note also that ρ_v^* serves as the \mathcal{TG} convergence bound for $\|V_j\|$, so we can take $\rho_{tg}^* = \rho_v^*$.

We could now use (6.3) and Theorem 7.2 to account for the effects of $\hat{\epsilon}_\ell$ -precision on the finest level, but we also need to assess the effects of the rounding error accumulated from all of the coarse levels. This we do by recursive use of bound (7.8), exploiting its form as a perturbation from the infinite-precision bound. A key point that enables this recursion is that $\delta_{\rho_{tg}}(\hat{\tau}_j)$ does not depend on ρ_{tg}^* if $\rho_{tg}(\hat{\tau}_j) < 1$, which holds because of our assumption on $\hat{\tau}_j$ and because $\rho_{tg}^* = \rho_v^*$. (Using (7.8) recursively means that we are assuming that initial rounding occurs on all coarse levels. While this assumption is valid on the finest \mathcal{V} level, the coarse levels have already been accounted for by the use of δ_0 in Theorem 7.2, meaning that the convenience of using recurrence here leads to a slightly looser but nevertheless acceptable bound on the convergence estimate.)

Another key point that enables recursion is that we keep these perturbations separate on all levels by writing the error propagation matrices in perturbation form.

More precisely, for each $j \in \{1, 2, \dots, \ell\}$, let the error propagation matrix for the result of the computed V-cycle be denoted by $V_j + E_j$, where V_j remains as the exact error propagation operator and E_j incorporates the rounding errors it accumulates. To assess the impact that E_j has on ρ_v^* in (6.3), from (6.2) and the fact that $P_j A_{j-1}^{-1} P_j^t A_j$ is an energy-orthogonal projection onto the range of P_j , we have that the exact level j V-cycle error propagation matrix based on the computed coarse-level V-cycle has the new bound

$$\begin{aligned}
\|V_j\|_{A_j} &= \|(P_j(V_{j-1} + E_{j-1})A_{j-1}^{-1}P_j^t A_j + T_j)G_j\|_{A_j} \\
&\leq \|(P_j V_{j-1} A_{j-1}^{-1} P_j^t A_j + T_j)G_j\|_{A_j} + \|P_j E_{j-1} A_{j-1}^{-1} P_j^t A_j\|_{A_j} \|G_j\|_{A_j} \\
&\leq \rho_v^* + \|E_{j-1} A_{j-1}^{-1} P_j^t A_j\|_{A_{j-1}} \\
&\leq \rho_v^* + \|E_{j-1}\|_{A_{j-1}} \|P_j A_{j-1}^{-1} P_j^t A_j\|_{A_j} \\
(8.3) \quad &\leq \rho_v^* + \|E_{j-1}\|_{A_{j-1}}.
\end{aligned}$$

We first need to estimate $\|E_j\|_{A_j}$ for the coarsest level ($j = 1$), which only uses relaxation to approximate the solution of $A_1 y_1 = r_1$. To this end, using iteration superscripts as in the previous two proofs, note by the definition of α_M that the final error on the coarsest level computed in $\hat{\epsilon}_1$ -precision satisfies

$$e_1^{(1)} = y_1^{(\frac{1}{2})} - A_1^{-1} r_1^{(0)} = M_1 r_1^{(0)} - A_1^{-1} r_1^{(0)} + \delta, \quad |\delta| \leq \alpha_{M_1} \hat{\epsilon}_1 |r_1^{(0)}|.$$

Remembering that $\dot{\tau}_j = \kappa_j^{\frac{1}{2}} \hat{\epsilon}_j$, we thus have that

$$\|e_1^{(1)}\|_{A_1} \leq \|M_1 r_1^{(0)} - A_1^{-1} r_1^{(0)}\|_{A_1} + \|\delta\|_{A_1} \leq (\rho_v^* + \alpha_{M_1} \dot{\tau}_1) \|e_1^{(0)}\|_{A_1},$$

from which follows the bound

$$(8.4) \quad \|E_1\|_{A_1} \leq \alpha_{M_1} \dot{\tau}_1.$$

Bound (8.3) implies that the basic convergence factor that we should be using in (7.8) is not just ρ_v^* , but rather $\rho_v^* + \|E_{j-1}\|_{A_{j-1}}$, so we can now use Theorem 7.2 to assert that $\rho_v = \rho_v^* + \|E_{j-1}\|_{A_{j-1}} + \delta_{\rho_{tg}}$ bounds the level j V-cycle convergence factor. Restating this conclusion in terms of the rounding-error matrices yields

$$(8.5) \quad \|E_j\|_{A_j} \leq \|E_{j-1}\|_{A_{j-1}} + \delta_{\rho_{tg}}(\dot{\tau}_j).$$

This inequality provides the recursion we need. Noting that (8.4) and Theorem 7.2 lead to the loose bound $\|E_1\|_{A_1} \leq \delta_{\rho_{tg}}(\dot{\tau}_1)$, then we can start with $j = \ell$ and recurse (8.5) back down through the coarse levels, using (8.1) to conclude that

$$\|E_\ell\|_{A_\ell} \leq \|E_{\ell-1}\|_{A_{\ell-1}} + \delta_{\rho_{tg}} \leq \sum_{j=1}^{\ell} \delta_{\rho_{tg}}(\dot{\tau}_j) \leq \delta_{\rho_{tg}} \sum_{j=1}^{\ell} \vartheta^{-m(\ell-j)} \leq \frac{\vartheta^m}{\vartheta^m - 1} \delta_{\rho_{tg}},$$

thus proving the theorem. \square

9. Full multigrid algorithm. While stopping criteria for iterative solvers applied to general linear equations is not entirely settled and is often application dependent, matters clarify for solvers applied specifically to discretized PDEs. If $A_j x_j = b_j$ represent increasingly accurate approximations to a given PDE, then the error in the

numerical solution of (3.1) incorporates errors from algebraic and discretization processes (not to mention other sources). Investing effort to compute an approximation with an algebraic error that is much smaller than the discretization error is usually not productive. Full multigrid (see [5]) is based on this premise. Properly designed and applied, full multigrid is a direct method in the sense that it targets the PDE: a single iteration through its outer loop guarantees a solution with algebraic accuracy comparable to that of the discretization. It can be the most efficient solver for this purpose, with optimal complexity proportional to the size of the finest grid. The key to its success is to capitalize on a consistent form of discretization accuracy in the multigrid hierarchy and to aim at delivering an approximation that is comparable to that accuracy on every level. The basic full multigrid approach is to start on the coarsest grid and proceed to the finest, making sure that each grid along the way is solved by V-cycles with accuracy commensurate to the discretization. In essence, if grid $j - 1$ is solved to within an error of Ch_{j-1}^q for some positive constants C and q , then using that result as an initial guess on grid j means that the initial error on grid j is bounded by some small multiple (depending on θ_j) of Ch_j^q . This in turn means that only a small number of V-cycles are needed to obtain discretization accuracy on grid j (i.e., error below Ch_j^q).

The full multigrid algorithm based on $N \geq 1$ inner \mathcal{IR} cycles using one \mathcal{V} each is given by the pseudocode described \mathcal{FMG} below and illustrated in Fig. 3. Note that \mathcal{FMG} amounts to three nested loops: outer \mathcal{FMG} , middle \mathcal{IR} , and inner \mathcal{V} . The choice of N here is critical to guarantee convergence to within discretization accuracy on each level. The goal of the next section is to determine N in the presence of rounding errors.

Algorithm 9.1 FMG(1, 0)-Cycle (\mathcal{FMG})

Input: $A, b, P, N \geq 1$ \mathcal{IR} cycles (using one $\mathcal{V}(1, 0)$ each), $\ell \geq 1$ \mathcal{FMG} levels.

1: $x \leftarrow 0$	▷ Initialize \mathcal{FMG}
2: if $\ell > 1$ then	▷ Check for Coarser Level
3: $x_{\ell-1} \leftarrow \mathcal{FMG}(A_{\ell-1}, b_{\ell-1}, P_{\ell-1}, \ell - 1, N)$	▷ Compute Coarse-Level Approximation
4: $x \leftarrow Px_{\ell-1}$	▷ Interpolate Approximation to Fine Level
5: end if	
6: $i \leftarrow 0$	▷ Initialize \mathcal{IR}
7: while $i < N$ do	
8: $r \leftarrow Ax - b$	▷ Update \mathcal{IR} Residual and Round
9: $y \leftarrow \mathcal{V}(A, r, P, \ell)$	▷ Compute Correction by \mathcal{V}
10: $i \leftarrow i + 1$	▷ Increment \mathcal{IR} Cycle Counter
11: $x \leftarrow x - y$	▷ Update Approximate Solution of $Ax = b$
12: end while	
13: return x	▷ Return Approximate Solution of $Ax = b$

10. Convergence Theory: Full Multigrid. While the theory developed here has elliptic PDEs in mind, the intent is to provide a more general algebraic theory that is removed from specific applications. To obtain such an abstract theory for \mathcal{FMG} applied to (3.1), we need to extract a sense of discretization accuracy directly from the matrix hierarchy itself. To do so, we need to establish a relationship between the exact solutions $A_j^{-1}b_j$ on all levels $j \in \{1, 2, \dots, \ell\}$, as we do in (10.1) below. Establishing how well level $j - 1$ solutions approximate level j solutions then leads us to a sense of discretization accuracy on the finest level, as expressed in (10.6) below.

To this end and commensurate with our assumption on the exactness of the Galerkin condition (6.1), assume now that $b_{j-1} = P_j^t b_j$, $2 \leq j \leq \ell$, are also computed




FIG. 3. An illustration of the FMG algorithm, shaded to convey the mixed precisions used throughout the processes. Only one V-cycle per \mathcal{IR} level is shown for clarity. The number of boxes in the intermediate V-cycles convey the precision reduction associated with the low-precision operations on the coarser levels in the multigrid hierarchy. The size of the triangles provide a qualitative depiction of the amount of precision required for \mathcal{IR} on each level.

exactly. Since $P_j A_{j-1}^{-1} P_j^t A_j$ is an energy-orthogonal projection, we have that

$$\begin{aligned}
 \|A_{j-1}^{-1} b_{j-1}\|_{A_{j-1}} &= \langle A_{j-1}^{-1} P_j^t b_j, P_j^t b_j \rangle^{\frac{1}{2}} \\
 &= \langle (P_j A_{j-1}^{-1} P_j^t A_j) A_j^{-1} b_j, A_j A_j^{-1} b_j \rangle^{\frac{1}{2}} \\
 &= \|(P_j A_{j-1}^{-1} P_j^t A_j) A_j^{-1} b_j\|_{A_j} \\
 (10.1) \qquad \qquad \qquad &\leq \|A_j^{-1} b_j\|_{A_j}.
 \end{aligned}$$

We can thus characterize the relative accuracy of adjacent levels in the hierarchy by assuming that C and q are positive constants such that

$$(10.2) \qquad \|P_j A_{j-1}^{-1} b_{j-1} - A_j^{-1} b_j\|_{A_j} \leq C h_{j-1}^q \|A_j^{-1} b_j\|_{A_j}, \quad 2 \leq j \leq \ell.$$

For a fully regular PDE of order $2m$, sufficient smoothness of the solution u , and finite element functions of a fixed order k (e.g., polynomials of fixed degree $p = k - 1$), it can be shown (see [25, Sec. 2.2]) that

$$(10.3) \qquad \|u - u_h\|_s = \mathcal{O}\left(h^{k-s} + h^{2(k-m)}\right),$$

where u_h denotes the discrete solution and s is the highest-order derivative in the Sobolev or Sobolev-equivalent norm $\|\cdot\|_s$. For the energy norm, $s = m$, so we assume that $k > m$ to ensure convergence. The error is therefore dominated by the first term such that the relative accuracy of the discrete approximation on grid j is h_j^q , where $q = k - m$. It is easy to show by the triangle inequality that this confirms that (10.2) holds in this case. This bound can also be established from existing matrix theory without reference to a PDE. Specifically, remembering the definitions of h_{j-1} and θ_j , then (10.2) holds with $C = \theta_j^{-m} C_1 C_2$ and $q = m + 1$ directly from the so-called *strong approximation property* (cf. [28]) given by

$$(10.4) \qquad \|P_j A_{j-1}^{-1} b_{j-1} - A_j^{-1} b_j\|_{A_j} \leq \frac{C_1}{\|A_j^{\frac{1}{2}}\|} \|A_j(A_j^{-1} b_j)\|,$$

which is assumed to hold for any RHS, and the *algebraic smoothness property* of our specific RHS given by

$$(10.5) \qquad \|b_j\| = \|A_j(A_j^{-1} b_j)\| \leq \frac{C_2}{\|A_j^{-\frac{1}{2}}\|} \|A_j^{-1} b_j\|_{A_j}.$$

The strong approximation property says loosely that $A_j^{-1}b_j$ can be accurately approximated when its expansion in terms of the eigenvectors of A_j are predominantly in the lower spectrum of A_j , which is what algebraic smoothness asserts.

While (10.2) characterizes the relative error in a coarse-level solution with respect to the next finer level, it also suggests the following definition. We say that x_j solves $A_j x_j = b_j$ to the order of discretization error or simply to discretization accuracy if

$$(10.6) \quad \|x_j - A_j^{-1}b_j\|_{A_j} \leq Ch_j^q \|A_j^{-1}b_j\|_{A_j}, \quad 1 \leq j \leq \ell.$$

Assume that (10.6) holds on the coarsest level as a result of one \mathcal{FMG} cycle there, that is, when x_1 is the result of N cycles of $\mathcal{IR}\text{-}\mathcal{V}$ on level $j = 1$ starting with a zero initial guess. (Note that \mathcal{V} consists of just one relaxation sweep.) This assumption holds (similar to (6.3) with $j = 1$) in the common case that A_1 is well-conditioned enough that relaxation converges quickly and rounding effects are negligible.

THEOREM 10.1. \mathcal{FMG} . *Assume that $\rho_v + \delta_{\rho_{ir}} < 1$ and that χ is small enough and N is large enough that the following holds on all levels $j \in \{1, 2, \dots, \ell\}$:*

$$(10.7) \quad (\rho_v + \delta_{\rho_{ir}})^N \left((\sqrt{2} + \mu\tau)\theta^q Ch^q + \mu\tau \right) + \frac{\chi}{1 - (\rho_v + \delta_{\rho_{ir}})} \leq Ch^q,$$

where $h = h_j$, $\theta = \theta_j$, $\tau = \tau_j$, and $\mu = \mu_j = \kappa^{\frac{1}{2}}(P_j^t P_j)m_P^+$, and (with subscript j understood) the parameters ρ_v , $\delta_{\rho_{ir}}$, and χ are given by (8.2), (4.4), and (4.5), respectively. Then Algorithm 9.1 solves (3.1) to the order of discretization error on each level.

Proof. We proceed by induction. Knowing that (10.6) holds for the coarsest level $j = 1$ by assumption, suppose that $j > 1$ and that (10.6) holds on level $j - 1$, where x_{j-1} is now the result of \mathcal{FMG} applied on that level. We need to account for the errors in transferring x_{j-1} to the fine level, but we first follow the effects of exact computation. To this end, note that

$$\begin{aligned} \langle P_j x_{j-1} - P_j A_{j-1}^{-1} b_{j-1}, A_j (P_j A_{j-1}^{-1} b_{j-1} - A_j^{-1} b_j) \rangle \\ &= \langle x_{j-1} - A_{j-1}^{-1} b_{j-1}, P_j^t A_j P_j A_{j-1}^{-1} b_{j-1} - P_j^t A_j A_j^{-1} b_j \rangle_{A_j} \\ &= \langle x_{j-1} - A_{j-1}^{-1} b_{j-1}, b_{j-1} - P_j^t b_j \rangle = 0. \end{aligned}$$

This orthogonality, the induction hypothesis, and (10.1) allow us to conclude that

$$\begin{aligned} \|P_j x_{j-1} - A_j^{-1} b_j\|_{A_j} &= \|P_j x_{j-1} - P_j A_{j-1}^{-1} b_{j-1} + P_j A_{j-1}^{-1} b_{j-1} - A_j^{-1} b_j\|_{A_j} \\ &= (\|P_j x_{j-1} - P_j A_{j-1}^{-1} b_{j-1}\|_{A_j}^2 + \|P_j A_{j-1}^{-1} b_{j-1} - A_j^{-1} b_j\|_{A_j}^2)^{\frac{1}{2}} \\ &\leq Ch_{j-1}^q (\|A_{j-1}^{-1} b_{j-1}\|_{A_{j-1}}^2 + \|A_j^{-1} b_j\|_{A_j}^2)^{\frac{1}{2}} \\ (10.8) \quad &\leq \sqrt{2} Ch_{j-1}^q \|A_j^{-1} b_j\|_{A_j}. \end{aligned}$$

We can account for rounding error in $P_j x_{j-1}$ with logic similar to that leading up to (7.6), but now in ε -precision: we can write its computed value as $P_j x_{j-1} + \delta_j$, $\|\delta_j\|_{A_j} \leq \mu_j \tau_j \|x_{j-1}\|_{A_{j-1}}$. The induction hypothesis and (10.1) again yields

$$\begin{aligned} \|x_{j-1}\|_{A_{j-1}} &\leq \|A_{j-1}^{-1} b_{j-1}\|_{A_{j-1}} + \|x_{j-1} - A_{j-1}^{-1} b_{j-1}\|_{A_{j-1}} \\ &\leq (1 + Ch_{j-1}^q) \|A_{j-1}^{-1} b_{j-1}\|_{A_{j-1}} \\ &\leq (1 + Ch_{j-1}^q) \|A_j^{-1} b_j\|_{A_j}, \end{aligned}$$

which, with (10.8) and (6.4), leads to

$$\begin{aligned} \|P_j x_{j-1} + \delta_j - A_j^{-1} b_j\|_{A_j} &\leq \sqrt{2} C h_{j-1}^q \|A_j^{-1} b_j\|_{A_j} + \|\delta_j\|_{A_j} \\ &\leq (\sqrt{2} C h_{j-1}^q + (1 + C h_{j-1}^q) \mu_j \tau_j) \|A_j^{-1} b_j\|_{A_j} \\ &\leq \left((\sqrt{2} + \mu_j \tau_j) \theta_j^q C h_j^q + \mu_j \tau_j \right) \|A_j^{-1} b_j\|_{A_j}. \end{aligned}$$

Applying N cycles of \mathcal{IR} to the initial approximation $P_j x_{j-1} + \delta_j$ thus yields

$$\|x_j - A_j^{-1} b_j\|_{A_j} \leq c_j \|A_j^{-1} b_j\|_{A_j},$$

where appealing to Theorems 4.1 and 8.1 allows us to take

$$c_j = (\rho_v + \delta_{\rho_{ir}})^N \left((\sqrt{2} + \mu_j \tau_j) \theta_j^q C h_j^q + \mu_j \tau_j \right) + \frac{\chi}{1 - (\rho_v + \delta_{\rho_{ir}})},$$

thus proving the theorem. \square

Remark 10.2. Understanding the Error Bounds. While the proof of Theorem 10.1 is fairly simple, its statement nevertheless involves the many parameters that are needed to account for errors in all three of the nested solver loops. We can, however, parse condition (10.7) somewhat by writing it loosely as follows, assuming that $\hat{\epsilon} \geq \epsilon \geq \bar{\epsilon}$ and remembering that $\hat{\tau} = \kappa^{\frac{1}{2}} \hat{\epsilon}$ and $\tau = \kappa^{\frac{1}{2}} \epsilon$:

$$(10.9) \quad (\rho_v^* + \mathcal{O}(\hat{\tau}))^N \left(\sqrt{2} \theta^q C h^q + \mathcal{O}(\tau) \right) + \mathcal{O}(\tau) \leq C h^q,$$

where we assumed as in Remark 4.2 that $\bar{\tau}$ is small enough to allow us to replace χ by $\mathcal{O}(\tau)$. We must have $\mathcal{O}(\hat{\tau}) < 1 - \rho_v^*$ to satisfy Theorem 10.1. We also require τ to be $\mathcal{O}(h^q)$. (Otherwise, (10.9) could not be obtained no matter how many \mathcal{IR} iterations are taken.) This assumption means that the term $\sqrt{2} \theta^q C h^q + \mathcal{O}(\tau)$ in (10.9) is bounded by a multiple of $C h^q$, which in turn means that $(\rho_v^* + \mathcal{O}(\hat{\tau}))^N$ only needs to be small enough to reduce that multiple to below $C h^q$ itself. On the other hand, if $\mathcal{O}(\hat{\tau})$ and $\mathcal{O}(\tau)$ are small enough to be negligible in (10.9), then this condition is satisfied by N being only so large that $(\rho_v^*)^N \sqrt{2} \theta^q < 1$, that is, $N > \frac{0.5+q \log_2(\theta)}{|\log_2(\rho_v^*)|}$. This requirement is independent of h , which ensures that \mathcal{FMG} converges optimally in the sense that it obtains accuracy to the level of the discretization in a uniformly bounded number of floating point operations per fine-level unknown.

11. Conclusions. This paper has established an abstract theory for analyzing the effects of rounding error on two-grid, V-cycle, and FMG solvers applied to sparse SPD matrices. We have not accounted for the effects of inexact system matrices, RHS's, and interpolation operators, but these issues are treated in the related paper [27] along with numerical results for a model biharmonic equation. To avoid further complexity, we have also not assessed the effects of overflow or underflow, post smoothing, or non-polynomial smoothers such as Gauss-Seidel. Nevertheless, the abstract algebraic framework we have developed should be applicable to any self-adjoint discretization of an elliptic PDE as well as virtually any other SPD matrix equation. The essential requirement is that the multigrid schemes treated here converge optimally in energy. The estimates we established then hold under the assumptions imposed on the various parameters.

We have shown that, in the energy norm, the relative algebraic error grows faster for fixed than for mixed precision, which illustrates the benefit of mixed-precision

$\mathcal{IR}\text{-}\mathcal{V}$. By comparing algebraic and discretization errors, we also conclude that in the absence of progressive precision, multigrid is ultimately dominated by rounding errors. Additionally, we have shown that, in the energy norm, the computed approximation is affected more by convergence stalling at limiting accuracy rather than by convergence factors degrading. In fact, the precision used in iterative refinement’s inner solve has little effect on the resulting accuracy compared to that of the limiting accuracy determined by the higher precision. The results also show that V-cycles and FMG are capable of leveraging progressive precision by using increasingly lower precision on levels that are increasingly coarser, and thus decreasingly accurate. In this way, progressive precision allows FMG to obtain discretization accuracy using a minimal amount of resources.

Appendix A. Summary of notation. The basic quantities and abbreviations used throughout the paper are as follows:

PDE	partial differential equation
RHS	right-hand side
SAP	strong approximation property
\mathcal{IR}	iterative refinement algorithm
\mathcal{TG}	two-grid multigrid algorithm
\mathcal{V}	V-cycle algorithm
\mathcal{FMG}	full multigrid algorithm
ε	“standard” precision unit roundoff
$\bar{\varepsilon}$	“high” precision unit roundoff
$\hat{\varepsilon}$	“low” precision unit roundoff
$\text{fl}(x \circ y)$	finite-precision result of the operation $x \circ y$
δ	perturbation of a scalar or vector due to finite precision
$n = n_j$	dimension of the fine level
$n_c = n_{j-1}$	dimension of the coarse level
ℓ	number of levels in the hierarchy
$A \in \mathbb{R}^{n \times n}$	system matrix
$D \in \mathbb{R}^{n \times n}$	diagonal of A
$M \in \mathbb{R}^{n \times n}$	approximation of A^{-1}
$P \in \mathbb{R}^{n \times n_c}$	prolongation or interpolation matrix
$B_c \in \mathbb{R}^{n_c \times n_c}$	expresses inaccuracy in Galerkin coarse-level matrix
$G \in \mathbb{R}^{n \times n}$	error propagation matrix for relaxation
$T \in \mathbb{R}^{n \times n}$	error propagation matrix for coarse-level correction
$V \in \mathbb{R}^{n \times n}$	error propagation matrix for a V-cycle
$Ax = b$	target problem
P^t	transpose of P
$\ \cdot\ $	Euclidean norm
$\ \cdot\ _A = \ A^{\frac{1}{2}} \cdot\ $	energy norm
$\kappa(\cdot)$	condition number
$\kappa = \kappa(A)$	condition number of A
m_A	bound on the number of nonzero row entries of A
m_P	bound on the number of nonzero row entries of P
N	number of \mathcal{IR} cycles
ρ	generic energy convergence factor
ρ_{ir}	energy convergence factor for \mathcal{IR}
ρ_{tg}^*	exact energy convergence factor for \mathcal{TG}

ρ_v^*	exact energy convergence factor for \mathcal{V}
$\rho_{ltg} = \rho_{tg}^* + \delta_{\rho_{ltg}}$	computed energy convergence factor for limited \mathcal{TG}
$\rho_{tg} = \rho_{tg}^* + \delta_{\rho_{tg}}$	computed energy convergence factor for \mathcal{TG}
$\rho_v = \rho_v^* + \delta_{\rho_v}$	computed energy convergence factor for \mathcal{V}
χ	limit accuracy in energy
h	pseudo mesh size
θ	pseudo mesh-refinement factor
ζ	precision coarsening factor
δ_M	rounding error in computing Mz
α_M	constant in $\mathcal{O}(\dot{\epsilon}\ z\)$ bound on $\ \delta_M\ $
$2m$	order of the PDE
k	finite element polynomial order
p	finite element polynomial degree
q	mesh-size exponent of coarse-level approximation order

Additionally, the following convenience parameters are used:

$$\psi = \|A\|, \quad \underline{\kappa} = \psi\|A^{-1}\|, \quad \bar{m}_A^+ = \frac{m_A + 1}{1 - (m_A + 1)\bar{\epsilon}}, \quad \dot{m}_P^+ = \frac{m_P}{1 - m_P\bar{\epsilon}},$$

$$\dot{m}_A^+ = \frac{m_A + 1}{1 - (m_A + 1)\dot{\epsilon}}, \quad \dot{\tau} = \kappa^{\frac{1}{2}}\dot{\epsilon}, \quad \tau = \kappa^{\frac{1}{2}}\epsilon, \quad \bar{\tau} = \kappa\bar{\epsilon}, \quad \gamma = \frac{\kappa^{\frac{1}{2}} + \underline{\kappa}}{\kappa},$$

$$\sigma \geq (1 + \dot{\epsilon}) \max\{\alpha_M\|A\|, \psi\alpha_M, \psi\|M\|\}, \quad \vartheta = \min_{1 \leq j \leq \ell} \{\theta_j \zeta_j^{-\frac{1}{m}}\},$$

$$\mu = 3(\max_{1 \leq j \leq \ell} \zeta_j) \kappa^{\frac{1}{2}} (P^t P) \dot{m}_P^+,$$

$$\delta_{\rho_{tg}} = \delta_{\rho_{tg}}(\dot{\tau}) = a_1 \dot{\tau} + a_2 \dot{\tau}^2 + a_3 \dot{\tau}^3, \quad \delta_{\rho_v} = \delta_{\rho_v}(\dot{\tau}_\ell) = \frac{\vartheta^m}{\vartheta^m - 1} \delta_{\rho_{tg}}(\dot{\tau}_\ell),$$

$$a_1 = 4 + \sigma + 4\mu, \quad a_2 = 2(3 + \sigma + \dot{m}_A^+(1 + \sigma))\mu + 2 + \sigma, \quad a_3 = 2(1 + \sigma + \dot{m}_A^+(1 + \sigma))\mu.$$

Subscript j on various quantities indicates the current (fine) level and $j - 1$ or c its adjacent coarse level in the hierarchy, with $j = 1$ the coarsest and $j = \ell$ the finest. These subscripts may be suppressed when there is no risk of ambiguity. Superscripts in parentheses like (1) and (i) are used to denote iteration count.

REFERENCES

- [1] M. AHUÉS AND F. CHATELIN, *The Use of Defect Correction to Refine the Eigenlements of Compact Integral Operators*, SIAM Journal on Numerical Analysis, 20 (1983), pp. 1087–1093, <https://doi.org/10.1137/0720077>.
- [2] M. ARIOLI, J. DEMMEL, AND I. DUFF, *Solving Sparse Linear Systems with Sparse Backward Error*, SIAM Journal on Matrix Analysis and Applications, 10 (1989), pp. 165–190, <https://doi.org/10.1137/0610013>.
- [3] M. BABOULIN, A. BUTTARI, J. DONGARRA, J. KURZAK, J. LANGOU, J. LANGOU, P. LUSZCZEK, AND S. TOMOV, *Accelerating scientific computations with mixed precision algorithms*, Computer Physics Communications, 180 (2009), pp. 2526–2533, <https://doi.org/10.1016/j.cpc.2008.11.005>. 40 YEARS OF CPC: A celebratory issue focused on quality software for high performance, grid and novel computing architectures.
- [4] A. BRANDT AND O. LIVNE, *Multigrid Techniques. 1984 Guide with Applications to Fluid Dynamics, Revised Ed.*, SIAM, Philadelphia, 2001, <https://doi.org/10.1137/1.9781611970753>.
- [5] W. L. BRIGGS, V. E. HENSON, AND S. F. MCCORMICK, *A Multigrid Tutorial*, SIAM Books, Philadelphia, 2000, <https://doi.org/10.1137/1.9780898719505>. Second edition.

- [6] A. BUTTARI, J. DONGARRA, J. KURZAK, P. LUSZCZEK, AND S. TOMOV, *Using Mixed Precision for Sparse Matrix Computations to Enhance the Performance While Achieving 64-bit Accuracy*, ACM Trans. Math. Softw., 34 (2008), pp. 17:1–17:22, <https://doi.org/10.1145/1377596.1377597>.
- [7] A. BUTTARI, J. DONGARRA, J. LANGOU, J. LANGOU, P. LUSZCZEK, AND J. KURZAK, *Mixed Precision Iterative Refinement Techniques for the Solution of Dense Linear Systems*, The International Journal of High Performance Computing Applications, 21 (2007), pp. 457–466, <https://doi.org/10.1177/1094342007084026>.
- [8] E. CARSON AND N. J. HIGHAM, *A New Analysis of Iterative Refinement and Its Application to Accurate Solution of Ill-Conditioned Sparse Linear Systems*, SIAM Journal on Scientific Computing, 39 (2017), pp. A2834–A2856, <https://doi.org/10.1137/17M1122918>.
- [9] E. CARSON AND N. J. HIGHAM, *Accelerating the Solution of Linear Systems by Iterative Refinement in Three Precisions*, SIAM Journal on Scientific Computing, 40 (2018), pp. A817–A847, <https://doi.org/10.1137/17M1140819>.
- [10] M. CLARK, R. BABICH, K. BARROS, R. BROWER, AND C. REBBI, *Solving lattice QCD systems of equations using mixed precision solvers on GPUs*, Computer Physics Communications, 181 (2010), pp. 1517–1528, <https://doi.org/10.1016/j.cpc.2010.05.002>.
- [11] J. DEMMEL, Y. HIDA, W. KAHAN, X. S. LI, S. MUKHERJEE, AND E. J. RIEDY, *Error Bounds from Extra-precise Iterative Refinement*, ACM Trans. Math. Softw., 32 (2006), pp. 325–351, <https://doi.org/10.1145/1141885.1141894>.
- [12] J. W. DEMMEL, *Applied Numerical Linear Algebra*, Society for Industrial and Applied Mathematics, 1997, <https://doi.org/10.1137/1.9781611971446>.
- [13] S. GALAL, *Energy efficient floating-point unit design*, PhD thesis, Stanford University, Nov 2012, <http://purl.stanford.edu/tf297yq9849>.
- [14] D. GÖDDEKE, *Fast and Accurate Finite-Element Multigrid Solvers for PDE Simulations on GPU Clusters*, PhD thesis, Technischen Universität Dortmund, May 2010, <http://hdl.handle.net/2003/27243>.
- [15] GUSTAFSON AND YONEMOTO, *Beating Floating Point at Its Own Game: Posit Arithmetic*, Supercomputing Frontiers and Innovations: an International Journal, 4 (2017), p. 71–86, <https://doi.org/10.14529/jsfi170206>.
- [16] J. L. GUSTAFSON, *The End of Error - Unum Computing*, Chapman and Hall/CRC, 2015, <https://www.routledge.com/The-End-of-Error-Unum-Computing/Gustafson/p/book/9781482239867>.
- [17] A. HAIDAR, S. TOMOV, J. DONGARRA, AND N. J. HIGHAM, *Harnessing GPU Tensor Cores for Fast FP16 Arithmetic to Speed Up Mixed-precision Iterative Refinement Solvers*, in Proceedings of the International Conference for High Performance Computing, Networking, Storage, and Analysis, SC '18, Piscataway, NJ, USA, 2018, IEEE Press, pp. 47:1–47:11, <https://doi.org/10.1109/SC.2018.00050>.
- [18] A. HAIDAR, P. WU, S. TOMOV, AND J. DONGARRA, *Investigating Half Precision Arithmetic to Accelerate Dense Linear System Solvers*, in Proceedings of the 8th Workshop on Latest Advances in Scalable Algorithms for Large-Scale Systems, ScalA '17, New York, NY, USA, 2017, ACM, pp. 10:1–10:8, <https://doi.org/10.1145/3148226.3148237>.
- [19] N. J. HIGHAM, *Accuracy and Stability of Numerical Algorithms, 2nd. Ed.*, SIAM, Philadelphia, 2002, <https://doi.org/10.1137/17M1122918>.
- [20] J. LANGOU, J. LANGOU, P. LUSZCZEK, J. KURZAK, A. BUTTARI, AND J. DONGARRA, *Exploiting the Performance of 32 bit Floating Point Arithmetic in Obtaining 64 bit Accuracy (Revisiting Iterative Refinement for Linear Systems)*, in SC '06: Proceedings of the 2006 ACM/IEEE Conference on Supercomputing, Nov 2006, pp. 50–50, <https://doi.org/10.1109/SC.2006.30>.
- [21] J. MANDEL, S. MCCORMICK, AND R. BANK, *Variational Multigrid Theory*, SIAM, Philadelphia, 1987, ch. 5, pp. 131–177, <https://doi.org/10.1137/1.9781611971057.ch5>.
- [22] S. F. MCCORMICK, *Multigrid Methods for Variational Problems: General Theory for the V-Cycle*, SIAM Journal on Numerical Analysis, 22 (1985), pp. 634–643, <https://doi.org/10.1137/0722039>.
- [23] C. B. MOLER, *Iterative Refinement in Floating Point*, J. ACM, 14 (1967), pp. 316–321, <https://doi.org/10.1145/321386.321394>.
- [24] A. PEDRAM, S. RICHARDSON, M. HOROWITZ, S. GALAL, AND S. KVATINSKY, *Dark Memory and Accelerator-Rich System Optimization in the Dark Silicon Era*, IEEE Design Test, 34 (2017), pp. 39–50, <https://doi.org/10.1109/MDAT.2016.2573586>.
- [25] G. STRANG AND G. FIX, *An Analysis of the Finite Element Method*, Wellesley-Cambridge Press, second ed., 2008.
- [26] R. STRZODKA AND D. GÖDDEKE, *Mixed Precision Methods for Convergent Iterative Schemes*,

- in Proceedings of the 2006 Workshop on Edge Computing Using New Commodity Architectures, May 2006, pp. D-59-60, <http://asc.ziti.uni-heidelberg.de/sites/default/files/research/papers/public/StGo06convIter.pdf>.
- [27] R. TAMSTORF, J. BENZAKEN, AND S. F. MCCORMICK, *Discretization-error-accurate mixed-precision multigrid solvers*, SIAM Journal on Scientific Computing, 43 (2021), pp. S420–S447, <https://doi.org/10.1137/20M1349230>.
- [28] P. S. VASSILEVSKI, *Multilevel Block Factorization Preconditioners*, Prentice-Hall series in automatic computation, Springer-Verlag, New York, 2008, <https://doi.org/10.1007/978-0-387-71564-3>.
- [29] J. H. WILKINSON, *Progress report on the Automatic Computing Engine*, Tech. Report MA/17/1024, Mathematics Division, Department of Scientific and Industrial Research, National Physical Laboratory, Teddington, UK, April 1948.
- [30] J. H. WILKINSON, *Rounding Errors in Algebraic Processes*, Prentice-Hall series in automatic computation, Prentice-Hall, Englewood Cliffs, N. J., 1963.