

## THE NUMERICAL STABILITY ANALYSIS OF PIPELINED CONJUGATE GRADIENT METHODS: HISTORICAL CONTEXT AND METHODOLOGY\*

ERIN C. CARSON<sup>†</sup>, MIROSLAV ROZLOŽNÍK<sup>‡</sup>, ZDENĚK STRAKOŠ<sup>§</sup>, PETR TICHÝ<sup>§</sup>,  
AND MIROSLAV TŮMA<sup>§</sup>

**Abstract.** Algebraic solvers based on preconditioned Krylov subspace methods are among the most powerful tools for large-scale numerical computations in applied mathematics, sciences, technology, as well as in emerging applications in social sciences. As the name suggests, Krylov subspace methods can be viewed as a sequence of projections onto nested subspaces of increasing dimension. They are therefore by their nature implemented as synchronized recurrences. This is the fundamental obstacle to efficient parallel implementation. Standard approaches to overcoming this obstacle described in the literature involve reducing the number of global synchronization points and increasing parallelism in performing arithmetic operations within individual iterations. One such approach, employed by the so-called pipelined Krylov subspace methods, involves overlapping the global communication needed for computing inner products with local arithmetic computations. Inexact computations in Krylov subspace methods, due to either floating point roundoff error or intentional action motivated by savings in computing time or energy consumption, have two basic effects, namely, slowing down convergence and limiting attainable accuracy. Although the methodologies for their investigation are different, these phenomena are closely related and cannot be separated from one another. The study of mathematical properties of Krylov subspace methods, in both the cases of exact and inexact computations, is a very active area of research and many issues in the analytic theory of Krylov subspace methods remain open. Numerical stability issues have been studied since the formulation of the conjugate gradient method in the middle of the last century, with many remarkable results achieved since then. Recently, the issues of attainable accuracy and delayed convergence caused by inexact computations became of interest in relation to pipelined conjugate gradient methods and their generalizations. In this contribution we recall the related early results and developments in synchronization-reducing conjugate gradient methods, identify the main factors determining possible numerical instabilities, and present a methodology for the analysis and understanding of pipelined conjugate gradient methods. We derive an expression for the residual gap that applies to any conjugate gradient method variant that uses a particular auxiliary vector in updating the residual, including pipelined conjugate gradient methods, and show how this result can be used to perform a full-scale analysis for a particular implementation. The paper concludes with a brief perspective on Krylov subspace methods in the forthcoming exascale era.

**Key words.** pipelined Krylov subspace methods, pipelined conjugate gradient methods, numerical stability, inexact computations, delay of convergence, maximal attainable accuracy, exascale computations

**AMS subject classifications.** 65F10, 65N12, 65G50, 65Y05, 65N22

**DOI.** 10.1137/16M1103361

---

\*Submitted to the journal's Methods and Algorithms for Scientific Computing section November 14, 2016; accepted for publication (in revised form) August 23, 2018; published electronically October 18, 2018.

<http://www.siam.org/journals/sisc/40-5/M110336.html>

**Funding:** The second author's work was partially supported by Grant Agency of the Czech Republic project GA17-1295S and by institutional support RVO 67985840. The third and fourth authors' work was partially supported by Grant Agency of the Czech Republic project GA17-04150J. The fifth author's work was partially supported by Grant Agency of the Czech Republic projects GA13-06684S and GA18-12719S.

<sup>†</sup>Courant Institute of Mathematical Sciences, New York University, New York, NY 10012 (carson@karlin.mff.cuni.cz).

<sup>‡</sup>Institute of Mathematics, Czech Academy of Sciences, Prague, 11567, Czech Republic (miro@math.cas.cz).

<sup>§</sup>Department of Numerical Mathematics, Faculty of Mathematics and Physics, Charles University, Prague, 18675, Czech Republic (strakos@karlin.mff.cuni.cz, ptichy@karlin.mff.cuni.cz, mirektuma@karlin.mff.cuni.cz).

**1. Introduction and background.** This paper considers the problem of using Krylov subspace methods in solving linear algebraic systems  $Ax = b$ , where  $A$  is a real  $N \times N$  nonsingular matrix and  $b$  is a real vector of length  $N$ . The restriction to real problems is for simplicity of notation and has no other role with respect to the studied phenomena. Given an initial approximation  $x_0$  to  $x$  and initial residual  $r_0 = b - Ax_0$ , Krylov subspace methods construct *mathematically*, i.e., assuming exact computation, a sequence of approximate solutions using the nested Krylov subspaces

$$\mathcal{K}_1(A, r_0) \subset \mathcal{K}_2(A, r_0) \subset \cdots \subset \mathcal{K}_n(A, r_0) \subset \cdots$$

defined as

$$\mathcal{K}_i(A, r_0) = \text{span}\{r_0, Ar_0, \dots, A^{i-1}r_0\}, \quad i = 1, 2, \dots$$

The  $n$ th approximation  $x_n \in x_0 + \mathcal{K}_n(A, r_0)$  to the solution  $x$  is constructed using the orthogonality of the  $n$ th residual to the  $n$ -dimensional constraint space  $\mathcal{C}_n$ , i.e.,  $r_n = b - Ax_n \perp \mathcal{C}_n$ , where the choice of  $\mathcal{C}_n$  along with properties of  $A$  distinguish the various Krylov subspace methods. As an example studied in this contribution, for symmetric positive definite (SPD) matrices and  $\mathcal{C}_n = \mathcal{K}_n(A, r_0)$  we get the conjugate gradient method (CG), where the orthogonality constraint  $r_n \perp \mathcal{K}_n(A, r_0)$  is equivalent to the minimization of the energy norm of the error over the Krylov subspaces, i.e.,

$$(1.1) \quad \|x - x_n\|_A = \min_{z \in x_0 + \mathcal{K}_n(A, r_0)} \|x - z\|_A,$$

where the energy norm  $\|u\|_A = (u, u)_A^{1/2} = (Au, u)^{1/2}$  is induced by the energy inner product defined by the SPD matrix  $A$ .

From the given orthogonality conditions it is clear that Krylov subspace methods are based on *linear projections onto highly nonlinear nested subspaces*. They are therefore highly nonlinear in the input data  $A$  and  $b$  defining the problem. This nonlinearity is what allows Krylov subspace methods to adapt to the problem as the iteration proceeds but makes analysis of the methods a challenge; see, e.g., [32] and the references therein. In particular, such analyses must respect the relationships between the studied phenomena, making them highly complex and unamenable to easy simplification. We will now explain the difficulty of the analysis of Krylov subspace methods, specifically of CG, in more detail.

Consider the Lanczos process for computing an orthonormal basis of the sequence of Krylov subspaces for a matrix  $A$  and starting vector  $v_1 = r_0/\|r_0\|$ , which is described by the matrix equations

$$(1.2) \quad AV_n = V_n T_n + \delta_{n+1} v_{n+1} e_n^T, \quad T_n = V_n^* A V_n,$$

where  $V_n = [v_1, \dots, v_n]$  is the  $N \times n$  matrix storing the orthonormal basis vectors as its columns,  $V_n^*$  denotes the transpose of the matrix  $V_n$ , and  $T_n$  denotes the Jacobi matrix storing in its columns the orthogonalization and normalization coefficients. The Jacobi matrix  $T_n$  depends on the  $2n - 1$  moments  $(A^i v_1, v_1)$ ,  $i = 1, 2, \dots, 2n - 1$ . The CG approximation is then given by solving the projected problem with the  $n \times n$  SPD tridiagonal matrix  $T_n$ , which determines the coefficients of the linear combination of the Lanczos basis functions, i.e.,

$$(1.3) \quad T_n y_n = \|r_0\| e_1, \quad x_n = x_0 + V_n y_n.$$

Summarizing, given an initial approximation  $x_0$ , the sequence of subsequent CG approximations  $x_1, x_2, \dots, x_n$  to the solution of the linear algebraic system  $Ax = b$

that represents the original model is obtained by solving the sequence of reduced models (1.3) of increasing dimensionality. Again, we emphasize that the described model reduction is highly nonlinear in the data  $A$  and  $b$ . The Jacobi matrix  $T_n$  is related to the short recurrences characteristic of the Lanczos method and also has connections with orthogonal polynomials, the Stieltjes problem of moments, Gauss–Christoffel quadrature, continued fractions, etc.; see [14], [32, Chapter 3].

As for practical computations, Krylov subspace methods such as MINRES and GMRES are indeed commonly implemented in a way that closely resembles the idea of a sequence of projection processes; see, e.g., [32, section 2.5.5] and the references therein. This is not the case for CG. The standard CG implementation described by Hestenes and Stiefel [26] (see also [32]) uses three two-term recurrences for updating the approximate solution  $x_i$ , the direction vector  $p_i$ , and the residual vector  $r_i$  (which is mathematically related to the Lanczos basis vectors by  $r_i = (-1)^i \|r_i\| v_{i+1}$ ), as follows.

ALGORITHM 1.1. HS conjugate gradient method.

**Input:** SPD matrix  $A \in R^{N \times N}$ , right-hand-side vector  $b \in R^N$ , initial approximation  $x_0 \in R^N$ , maximum number of iterations  $nmax$ .

**Output:** Approximate solution  $x_n$  after the algorithm has been stopped.

0. **Initialization:**  $r_0 = b - Ax_0$ ,  $p_0 = r_0$
1. **for**  $i = 1 : nmax$  **do**
2.      $\alpha_{i-1} = \frac{(r_{i-1}, r_{i-1})}{(p_{i-1}, Ap_{i-1})}$
3.      $x_i = x_{i-1} + \alpha_{i-1} p_{i-1}$
4.      $r_i = r_{i-1} - \alpha_{i-1} Ap_{i-1}$
5.     *evaluate the stopping criterion*
6.      $\beta_i = \frac{(r_i, r_i)}{(r_{i-1}, r_{i-1})}$
7.      $p_i = r_i + \beta_i p_{i-1}$
8. **end do**

Here the choice of  $\alpha_{i-1}$  ensures the minimization of  $\|x - x_i\|_A$  along the line

$$(1.4) \quad x(\alpha) = x_{i-1} + \alpha p_{i-1}.$$

The mathematical elegance and power of CG is given by the following crucial condition. Provided that the direction vectors are mutually orthogonal with respect to the energy inner product defined by the SPD matrix  $A$ , i.e.,

$$(1.5) \quad p_i \perp_A p_j \quad \text{for } i \neq j,$$

the one-dimensional line minimizations at the individual steps 1 to  $i$  result in the  $i$ -dimensional minimization over the *whole* shifted Krylov subspace

$$x_0 + \mathcal{K}_i(A, r_0) = x_0 + \text{span}\{p_0, p_1, \dots, p_{i-1}\}.$$

Analogous to the link between the Jacobi matrix and the short recurrences mentioned above, the line search (1.4) and the orthogonality condition (1.5) lead to short recurrences due to the symmetry of the matrix  $A$  or, equivalently, due to the relationship with the orthogonal polynomials that define the algebraic residuals and search vectors.

In each iteration of HS CG (Algorithm 1.1), updating the recurrences for  $x_i$ ,  $r_i$ , and  $p_i$  requires computing the inner products  $(p_{i-1}, Ap_{i-1})$  and  $(r_i, r_i)$  in order to

form the scalars  $\alpha_{i-1}$  and  $\beta_i$ . When Algorithm 1.1 is implemented on a parallel machine, computing each inner product requires a global synchronization point; i.e., the computation cannot proceed until all processors have finished their local computation and communicated the result to other processors. It is well-known that for large-scale sparse problems on large-scale machines, the cost of synchronization between parallel processors can dominate the run-time (see, e.g., the exascale computing report [12, p. 28]). This limits the potential speed in performing individual iterations attainable by an implementation of CG. Note that in this work, we will use the term *synchronization* to refer to a global synchronization required to compute an inner product; i.e., all processors must exchange data and wait for all communication to finish before proceeding with the computation. In MPI terminology, our use of the term synchronization is synonymous with blocking collective communication.

The goal of removing performance bottlenecks caused by synchronization has motivated work on pipelined Krylov subspace methods (see, e.g., [17], [18]). The key feature of pipelined CG [18] (and other pipelined Krylov subspace methods) is that computing the necessary inner products only requires a single nonblocking communication per iteration. Furthermore, the computation of the inner products can be overlapped with the matrix-vector product and other local computations, effectively hiding the latency of global communication. This overlapping is enabled by the introduction of one or more auxiliary vectors which are updated via additional recurrences in each iteration.

Other Krylov subspace method variants have been developed for the case where the matrix-vector products required in each iteration dominate the run-time. One example is the so-called inexact Krylov subspace methods (see, e.g., [2], [50], [51], [62], [63]) in which the accuracy of the matrix-vector product is intentionally relaxed throughout the iterations according to a function of the residual norm. We note that, in contrast to the terminology introduced in the works mentioned above, in this work we use the term “inexact” in a more general sense, referring to methods in which computed quantities can deviate from their exact mathematical counterparts due to intentional approximation or relaxation of accuracy, or simply due to floating point roundoff error.

In exact arithmetic, the pipelined CG method produces the same sequence of approximate solutions and residual vectors as the standard CG method. In finite precision, however, the introduction of additional recurrences causes the numerical behavior of the standard and pipelined CG and of their implementation variants to differ. It has indeed been observed (see, e.g., [18]) that in finite precision computations, the convergence rate and attainable accuracy of pipelined CG can be worse than in standard CG, with the attainable accuracy reduced by orders of magnitude in some cases. When pipelined Krylov subspace methods are used in practical settings, these differences in numerical behavior must then be taken into account, ideally balancing the application-specific accuracy requirements with potential performance improvements. Analyses which provide insight into the numerical behavior of pipelined Krylov subspace methods are thus vital to their use in practice.

In general, analyses of inexact CG computations (again, due to either roundoff error or an intentional relaxation of accuracy in performing the most costly operations) must address the question of *how the previously discussed relationships that hold for exact standard CG are changed in the presence of inaccuracies*. After deriving recurrences for the *computed quantities*, i.e., the mathematical equalities analogous to (1.2) and to the lines in Algorithm 1.1, the inaccuracies of individual local computations may be accounted for by adding local error terms to lines 2–4 and 6–7 of

Algorithm 1.1 (in the Lanczos process, to each column of (1.2)); see [39, section 4, Theorem 4.1, and section 5, relation (5.1)]. These local error terms and other subsequent local inaccuracies at individual iteration steps (such as the loss of orthogonality among the consecutive residuals and direction vectors) can be small. Their effect, however, can be drastically amplified within few iterations and can cause, e.g., a significant loss of global orthogonality or even numerical linear independence among the computed vectors generating the associated subspaces. The Jacobi matrices resulting from the computed data can quickly become very different from their exact arithmetic counterparts.

Much work has been done in the area of analyzing standard CG in finite precision arithmetic. In finite precision CG (as well as in general inexact CG methods), the CG optimality property (1.1) does not have a clear meaning with respect to the subspaces generated by the computed residual (or direction) vectors. Greenbaum showed in 1989 [19] that for the finite precision Lanczos process it does have, however, a rigorously defined meaning with respect to some particular distribution functions defined by the original data and the rounding errors in iterations 1 to  $n$ . Greenbaum proved, using the results of Paige on the numerical behavior of the Lanczos method published within 1971–1980 [42], [43], [44], [45], that the Jacobi matrix *computed in iteration  $n$  in finite precision arithmetic* can be considered a left principal submatrix of a certain larger Jacobi matrix having all its eigenvalues close to the eigenvalues of the original matrix  $A$ ; for a summary and exposition on the consequences of this fundamental result we refer to [39, section 5]. Let us assume, for a moment, that the effects of rounding errors in computing  $y_n$  (see (1.3)) and subsequently forming  $x_n$  are negligible in comparison to the effects of rounding errors in computing the Lanczos vectors  $v_1, \dots, v_n$ . Then the first  $n$  iterations of the given finite precision CG computation using (1.3) are mathematically equivalent to evaluating *exact* Gauss quadrature for a certain distribution function (that depends on  $n$ ) that has tight clusters of points of increase around the original eigenvalues of  $A$ ; see, e.g., [19], [39, section 5.2], [41], [32, section 5.9.1]. It is clear that determining how the convergence of the practically computed approximations to the solution  $x$  is slowed down in comparison to the mathematical formulation of CG indeed represents a challenge.

Another question is related to the maximal accuracy that can be attained in practical computations, which becomes important in solving problems for which high accuracy of the computed approximation is required. While in exact arithmetic the exact solution  $x$  must be reached by CG at or before step  $N$ , with inexact computations this is not the case. The basic methodology for analyzing the attainable accuracy involves comparing the residuals  $\bar{r}_i$ , iteratively computed using, e.g., Algorithm 1.1, with the residuals mathematically defined as  $b - A\bar{x}_i$ , where  $\bar{r}_i$  and  $\bar{x}_i$  denote the computed quantities. The *mathematical* difference

$$(1.6) \quad f_i \equiv \bar{r}_i - (b - A\bar{x}_i)$$

is then evaluated as a function of the increasing index  $i$  and is subsequently bounded in an appropriate norm. Using the argument that the iteratively computed residual  $\bar{r}_i$  eventually becomes negligible (even in finite precision computations), this gives for  $i$  large enough an estimate of the maximal attainable accuracy. This methodology is rigorously supported by formal proofs for some special CG implementations (see, e.g., [20] and the surveys in [27, section 17.1], [39, section 5.4]). Otherwise it is based on justified heuristics and observations. In order to avoid confusion that has appeared in the literature, it should be stressed that  $f_i$  in (1.6) represents a *mathematical quantity*, whereas  $\bar{r}_i$  and  $\bar{x}_i$  are the outputs of inexact computations. For details of the rounding

error analysis of the classical variant HS CG we refer to the survey [39] and to the update in [32, Chapter 5].

From the previous argument it follows that any analysis of inexact Krylov subspace methods, including pipelined variants, as well as any potential technique for counteracting the numerical effects of inexact computations, must very carefully consider the links between the studied phenomena. The assumptions and simplifications used must be justified. A clear, thorough understanding of how inexact computations affect numerical behavior is imperative in balancing the trade-offs between accuracy and speed in high-performance implementations; this becomes especially important for computations at the exascale level.

The following section will recall some earlier efforts and recent developments toward increasing parallel efficiency in CG computations. Section 3 then describes and numerically illustrates how modifications of the CG implementation and the corresponding changes in the individual sources of instabilities contribute to the overall loss of accuracy and delay of convergence. In section 4, we present a methodology for the analysis of pipelined CG and related variants in finite precision. We begin by deriving an expression for the residual gap that applies to CG variants like pipelined CG and then demonstrate how the size of this quantity can be appropriately bounded in terms of quantities produced during the iterations for a particular implementation. We also outline essentials of methodology for understanding and analyzing the delay of convergence. We conclude with brief remarks on the use of Krylov subspace methods for exascale-level computations.

Preconditioning represents a crucial part of any practical iterative computation, and thus the analysis of applications of Krylov subspace methods must always include consideration of preconditioning. Here, however, we take the luxury of formulating most of the text using unpreconditioned CG; as it turns out, the particular preconditioning approach does not affect the description of the basic stability analysis building blocks. As will be mentioned in section 2, preconditioning can easily be included in the described framework.

**2. Toward parallel efficiency: Historical context and recent work.** The term “standard CG” is generally associated with the seminal paper by Hestenes and Stiefel [26] and with the HS CG implementation described in Algorithm 1.1. Nevertheless, it is instructive to mention some other early contributions and ideas that lead to this culmination. This approach will enable us to demonstrate that the early development of CG contained many ideas that have been repeatedly reinvented with the emergence of modern computing architectures.

One of the basic driving forces behind the development of CG was a careful investigation of the connection between solving systems of linear algebraic equations and minimizing quadratic functionals. As stated by Hestenes [25], ideas of this type had been pursued previously by many authors. In the middle of the previous century, these ideas were considered in relation to the state-of-the-art relaxation methods of the time. A general framework of the related class of the so-called  $n$ -step methods can be found in the paper by Stiefel [54] (they should not be confused with the  $s$ -step Krylov subspace methods mentioned below). Investigation of theoretical properties of the  $n$ -step methods as well as connections to the steepest gradient method in Stiefel [55] reveals another important concept.

**2.1. Three-term recurrence variants.** The ST method (or ST CG for short), shown in Algorithm 2.1, is based on *three-term recurrences* for updating the approximate solutions and corresponding residuals. Its derivation was motivated by the

relation to three-term recurrences for orthogonal polynomials with a specific choice of density function [61, Chapter 3]. The algorithm is often attributed to Rutishauser [13]; see, e.g., [46]. As mentioned above, it was actually developed much earlier, around the same time as HS CG; see also Rosser [47], who attributes the described results to the joint work with Forsythe, Hestenes, Lanczos, Motzkin, and Paige.

ST CG is mathematically equivalent to HS CG given in Algorithm 1.1, but ST CG and HS CG exhibit rather different behavior in finite precision arithmetic. We will return to this point in section 3.

ALGORITHM 2.1. ST (three-term) conjugate gradient method.

**Input:** SPD matrix  $A \in R^{N \times N}$ , right-hand-side vector  $b \in R^N$ , initial approximation  $x_0 \in R^N$ , maximum number of iterations  $nmax$ .

**Output:** Approximate solution  $x_n$  after the algorithm has been stopped.

0. **Initialization:**  $r_0 = b - Ax_0$ ,  $x_{-1} = x_0$ ,  $r_{-1} = r_0$ ,  $e_{-1} = 0$
1. **for**  $i = 1 : nmax$  **do**
2.  $q_{i-1} = \frac{(r_{i-1}, Ar_{i-1})}{(r_{i-1}, r_{i-1})} - e_{i-2}$
3.  $x_i = x_{i-1} + \frac{1}{q_{i-1}}[r_{i-1} + e_{i-2}(x_{i-1} - x_{i-2})] \equiv x_{i-1} + \frac{1}{q_{i-1}}[r_{i-1} + e_{i-2}\Delta x_{i-1}]$
4.  $r_i = r_{i-1} + \frac{1}{q_{i-1}}[-Ar_{i-1} + e_{i-2}(r_{i-1} - r_{i-2})] = r_{i-1} + \frac{1}{q_{i-1}}[-Ar_{i-1} + e_{i-2}\Delta r_{i-1}]$
5. *evaluate the stopping criterion*
6.  $e_{i-1} = q_{i-1} \frac{(r_i, r_i)}{(r_{i-1}, r_{i-1})}$
7. **end do**

While the mathematical equivalence of ST CG and HS CG is well-known, we will state this relationship formally to point out the mutually equivalent quantities used in these algorithms and show relations among  $e_i, q_i, \alpha_i$ , and  $\beta_i$ . We do this in Observation 2.1.

*Observation 2.1.* Algorithms 1.1 and 2.1 produce the same sequences of iterates  $x_i$  and residuals  $r_i$ , respectively, for  $i = 0, 1, \dots$ . Moreover, we have  $q_i = 1/\alpha_i$  and  $e_i = q_i\beta_{i+1}$  for  $i = 0, 1, \dots$ .

*Proof.* Let us first relate the scalar coefficients of the two considered algorithms. Consider  $q_i$  and  $e_i$  from Algorithm 2.1. Clearly,  $q_0 = 1/\alpha_0$  and  $e_0 = \beta_1q_0$ . Consider  $i > 1$ . Using notation from Algorithm 1.1, basic orthogonality relations between the exact quantities and the induction assumption, we get

$$\begin{aligned} q_{i-1} &= \frac{(r_{i-1}, Ar_{i-1})}{(r_{i-1}, r_{i-1})} - e_{i-2} = \frac{(p_{i-1}, Ap_{i-1})}{(r_{i-1}, r_{i-1})} + \frac{(\beta_{i-1}p_{i-2}, \beta_{i-1}Ap_{i-2})}{(r_{i-1}, r_{i-1})} - q_{i-2}\beta_{i-1} \\ &= \frac{1}{\alpha_{i-1}} + \frac{\beta_{i-1}^2}{\beta_{i-1}\alpha_{i-2}} - q_{i-2}\beta_{i-1} = \frac{1}{\alpha_{i-1}}. \end{aligned}$$

The three-term recurrences from Algorithm 2.1 then easily follow. Consider, for example, the recursion for approximate residuals. From Algorithm 1.1 we have  $r_i = r_{i-1} - \alpha_{i-1}Ap_{i-1}$ . Observing that the direction vector update provides  $Ap_i = Ar_i + \beta_iAp_{i-1}$ , clearly,

$$\begin{aligned} r_i &= r_{i-1} - \alpha_{i-1}Ap_{i-1} \\ &= r_{i-1} - \alpha_{i-1}Ar_{i-1} - \alpha_{i-1}\beta_{i-1}Ap_{i-2} \\ &= r_{i-1} - \alpha_{i-1}Ar_{i-1} + \alpha_{i-1}\beta_{i-1}/\alpha_{i-2}(r_{i-1} - r_{i-2}) \\ &= r_{i-1} - 1/q_{i-1}Ar_{i-1} + 1/q_{i-1}e_{i-2}(r_{i-1} - r_{i-2}). \end{aligned}$$

The three-term update for the approximate solution  $x_i$  can be obtained analogously.  $\square$

The fact that a method using three-term recurrences can behave in finite precision arithmetic very differently from Algorithm 1.1 has been demonstrated by Gutknecht and Strakoš in [23]; see also [37, Chapters 6 and 7]. Nevertheless, it should be noted that the authors analyzed a slightly different algorithm and thus the detailed floating-point analysis of Algorithm 2.1 is still to be done. The same is true for other existing mathematically equivalent approaches.

**2.2. Early approaches to reducing synchronization in Krylov subspace methods.** A renewed interest in the numerical properties of different variants of CG and other Krylov subspace methods came with the advent of modern massively parallel computer architectures. Several modifications of the basic algorithms have been proposed in order to deal with the problem of communication cost. Chronopoulos and Gear proposed in [6], [7] the so-called s-step extensions of Krylov subspace methods. A couple of substantial further developments along this line followed; see, e.g., the Ph.D. theses by Hoemmen [28], Ballard [1], and Carson [4]. The main idea of the s-step methods is to enable more parallelism by relaxing the sequential construction of the basis of the Krylov subspace in Algorithm 1.1 and computing more direction vectors simultaneously. For further information on the state of the art and open questions we refer to [4].

As mentioned in the introduction, the effort toward improving parallel efficiency in individual CG iterations by reducing the synchronization cost has recently led to the creation of pipelined Krylov subspace methods [17], [18]. Analogous ideas were presented in numerous early works. The recent approaches to reducing the number of synchronization points in fact represent a variation of the three-term recurrences given in Algorithm 2.1 that dates back to the 1950s. Before returning to more recent work, we recall some other early approaches.

An idea independently given by Johnsson [29], [30], van Rosendale [67] (see also [15]), and Saad [48] is based on the identity that in the notation of Algorithm 1.1 reads

$$(2.1) \quad \|r_i\|^2 + \|r_{i-1}\|^2 = \alpha_{i-1}^2 \|Ap_{i-1}\|^2.$$

This enables the computation of  $\beta_i$  from  $\alpha_{i-1}$  and  $Ap_{i-1}$  as well as the computation of the vectors  $x_i$ ,  $r_i$ , and  $p_i$  in one block. As demonstrated by Saad in [49], this approach is unstable and cannot be used in a straightforward way. Despite this, based on the proposal of Meurant (see [36], [35], [5]), it is possible to use the value of  $\|r_i\|^2$  computed via (2.1) in updating the direction vector. The quantity  $\beta_i$  may be recomputed once the corresponding residual vector  $r_i$  is available at the expense of an additional inner product. Meurant claims in [36] that this works well in all his test examples. This idea might deserve further investigation. A similar approach has been proposed in [34], where the computation is based on the formula

$$(2.2) \quad \|r_i\|^2 = \|r_{i-1}\|^2 + \alpha_{i-1}^2 \|Ap_{i-1}\|^2 + 2\alpha_{i-1} r_{i-1} Ap_{i-1},$$

although this work lacks a supporting numerical stability analysis.

Another strategy for computing the inner products using only one synchronization point is to derive  $\alpha_i$  from  $\beta_{i-1}$  using the relation

$$(2.3) \quad \alpha_i = \frac{(r_i, r_i)}{(Ar_i, r_i) - (\beta_{i-1}/\alpha_{i-1})(r_i, r_i)}.$$



Observation 2.1 implies that Algorithm 2.1 does exactly this. Probably the first attempt to use Algorithm 2.1 to reduce the number of synchronization points in CG implemented on parallel computers has been published in [59]; see also [58]. Since then, the idea of computing  $\alpha_i$  based on  $\beta_{i-1}$  has been reinvented a couple of times in parallel computing, sometimes cited as an alternative approach, and often complemented by additional modifications motivated by particular computer architectures. Many later variants slightly differ in the details of the computation of both the scalar quantities and the vector updates, which may influence their actual numerical behavior; see section 3. In order to make precise our terminology, we introduce the following definition.

**DEFINITION 2.1.** *We use three-term CG to refer to any CG implementation mathematically equivalent to Algorithm 2.1 that computes  $\alpha_i$  based on  $\beta_{i-1}$  using (2.3), where this computation can be accomplished by a single synchronization point per iteration.*

Chronopoulos and Gear in [7] in fact reintroduced in their Algorithm 2.2 a three-term CG very closely related to ST CG (Algorithm 2.1). They in addition use recursive computation of the products  $Ap_i$  according to the formula

$$(2.4) \quad Ap_i = Ar_i + \beta_{i-1}Ap_{i-1};$$

see the following algorithm.

**ALGORITHM 2.2.** ChG conjugate gradient method.

**Input:** SPD matrix  $A \in R^{N \times N}$ , right-hand-side vector  $b \in R^N$ , initial approximation  $x_0 \in R^N$ , maximum number of iterations  $nmax$ .

**Output:** Approximate solution  $x_n$  after the algorithm has been stopped.

0. **Initialization:**  $r_0 = b - Ax_0$ ,  $p_0 = r_0$ ,  $s_0 = Ap_0$ ,  $\alpha_0 = \frac{(r_0, r_0)}{(p_0, s_0)}$
1. **for**  $i = 1 : nmax$  **do**
2.    $x_i = x_{i-1} + \alpha_{i-1}p_{i-1}$ ,
3.    $r_i = r_{i-1} - \alpha_{i-1}s_{i-1}$
4.   *evaluate the stopping criterion*
5.    $w_i = Ar_i$
6.    $\beta_i = \frac{(r_i, r_i)}{(r_{i-1}, r_{i-1})}$
7.    $\alpha_i = \frac{(r_i, r_i)}{(w_i, r_i) - (\beta_i/\alpha_{i-1})(r_i, r_i)}$
8.    $p_i = r_i + \beta_i p_{i-1}$
9.    $s_i = w_i + \beta_i s_{i-1}$
10. **end do**

Although the recurrences in Algorithm 2.2 look quite similar to the two-term recurrences used in Algorithm 1.1, this is only a formal resemblance. Substitution of rows 8 and 9 into the solution and residual updates on rows 2 and 3 immediately reveals that Algorithm 2.2 is based on three-term recurrences as Algorithm 2.1. Moreover, using notation in Algorithms 2.1 and 2.2, we have

$$(2.5) \quad q_i = \frac{1}{\alpha_i} = \frac{(w_i, r_i)}{(r_i, r_i)} - \frac{\beta_i(r_i, r_i)}{\alpha_{i-1}(r_i, r_i)} = \frac{(w_i, r_i)}{(r_i, r_i)} - q_{i-1}\beta_i = \frac{(w_i, r_i)}{(r_i, r_i)} - e_{i-1}.$$

Algorithm 2.2 uses the additional recurrence (2.4) for the  $A$ -multiple of the direction vector,  $s_i = Ap_i$ . This may bring in an additional source of instability in ChG CG (Algorithm 2.2) in comparison to ST CG (Algorithm 2.1).

There were some other works related to the original contribution by Stiefel [54], [55]. The idea of computing  $\alpha_i$  based on  $\beta_{i-1}$  has been used by D’Azevedo, Eijkhout, and Romine [10]; see also [9]. Here the connection to the work of Stiefel has not been noticed. The proposed scheme from [10] and Method 1 from [9] are similar to the ChG CG method in that they compute  $A$  multiples of the direction vectors  $p_i$  using an additional recurrence.

**2.3. Recent work in pipelined Krylov subspace methods.** An algorithm that combines reducing the number of synchronization points and overlapping inner products with other computations has been developed by Ghysels and Vanroose [18] under the name pipelined conjugate gradient. We present here its unpreconditioned version.

ALGORITHM 2.3. GV (pipelined) conjugate gradient method.

**Input:** SPD matrix  $A \in R^{N \times N}$ , right-hand-side vector  $b \in R^N$ , initial approximation  $x_0 \in R^N$ , maximum number of iterations  $nmax$ .

**Output:** Approximate solution  $x_n$  after the algorithm has been stopped.

```

0. Initialization:  $r_0 = b - Ax_0$ ,  $p_0 = r_0$ ,  $s_0 = Ap_0$ ,  $w_0 = Ar_0$ ,  $z_0 = Aw_0$ ,  $\alpha_0 = \frac{(r_0, r_0)}{(p_0, s_0)}$ 
1. for  $i = 1 : nmax$  do
2.    $x_i = x_{i-1} + \alpha_{i-1}p_{i-1}$ ,
3.    $r_i = r_{i-1} - \alpha_{i-1}s_{i-1}$ 
4.    $w_i = w_{i-1} - \alpha_{i-1}z_{i-1}$ 
5.   evaluate the stopping criterion
6.    $q_i = Aw_i$ 
7.    $\beta_i = \frac{(r_i, r_i)}{(r_{i-1}, r_{i-1})}$ 
8.    $\alpha_i = \frac{(r_i, r_i)}{(w_i, r_i) - (\beta_i/\alpha_{i-1})(r_i, r_i)}$ 
9.    $p_i = r_i + \beta_i p_{i-1}$ 
10.   $s_i = w_i + \beta_i s_{i-1}$ 
11.   $z_i = q_i + \beta_i z_{i-1}$ 
12. end do

```

Comparing pipelined CG, denoted here as GV CG (Algorithm 2.3), with ChG CG (Algorithm 2.2), we can see that both use the same computation of scalars  $\alpha_i$  and the same computation of  $Ap_i$ . In addition, GV CG introduces a recurrence for computing  $Ar_i$  iteratively. As described in section 1, the potential benefit to performance comes from the fact that the matrix vector product in line 6 can be executed in parallel with the computation of inner products in lines 7 and 8, since the computations of the scalars  $\alpha_i$  and  $\beta_i$  do not depend on the vector  $q_i$ . However, the auxiliary recurrences that enable overlapping the inner products with the matrix-vector multiplication do not explicitly enforce close local orthogonality among the computed vectors by recomputing the associated recurrence coefficients. Consequently, a substantial source of potential numerical instability is introduced.

Another recent approach called asynchronous CG has been proposed by Gropp [22]; see also [18]. In its preconditioned form, asynchronous CG keeps the two

synchronization points as in HS CG but uses additional recurrences to compute the preconditioned residual  $M^{-1}r_i$  and the vector  $Ap_i$ . Here substantial numerical instabilities may again arise from the additional recursively computed quantities.

Let us emphasize here the role of preconditioning mentioned at the end of the previous section. Some approaches mentioned above that use overlapping of operations conveniently involve preconditioning in balancing the associated computation load. In pipelined CG, preconditioning is implemented at the expense of an additional recursion. This inevitably influences finite precision computations. From a methodological point of view, preconditioning presents no obstacle to performing numerical stability analysis.

The idea of overlapping the computation of inner products with other local computations at the price of auxiliary vector updates (ideas that were used in earlier Krylov subspace method implementations) has not been, to our knowledge, widely publicized prior to [17], [18]. In [59], overlapping was not used because of computer architecture limitations. An example of such overlapping can be found, e.g., in the preconditioned HS algorithm proposed by van der Vorst [64]. Here the preconditioning operation is split into forward and backward solves using the incomplete Cholesky factorization combined with delaying the update of the approximate solution.

The pipelined CG method and other pipelined Krylov subspace methods have attracted recent attention for their potential use in high-performance computing. A large number of subsequent papers and reports emphasize various features of these methods with respect to parallel performance. It is not our task to list them all here. Our goal is to point out the need for a thorough rigorous numerical stability analysis and to outline the way toward doing this.

**3. Individual sources of instabilities and their effect on delay and attainable accuracy.** In this section we identify individual sources of instability and use simple examples to demonstrate their effect on numerical behavior. Using the notation standard in numerical stability analysis, the quantities computed in finite precision arithmetic will be decorated with bars. The residual vectors computed using the recurrence  $r_i = r_{i-1} - \alpha_{i-1}Ap_{i-1}$  satisfy

$$\bar{r}_i = \bar{r}_{i-1} - \bar{\alpha}_{i-1}A\bar{p}_{i-1} + \delta r_i,$$

where  $\bar{r}_{i-1}$ ,  $\bar{r}_i$ , and  $\bar{p}_{i-1}$  are the computed vectors,  $\bar{\alpha}_{i-1}$  is the computed coefficient, and  $\delta r_i$  is a vector that accounts for local roundoff errors. We will now recall the basic methodology for numerical stability analysis of CG through the example of investigating the maximal attainable accuracy.

It is well known that there exists a limit on the accuracy of approximate solutions  $\bar{x}_i$  computed in finite precision arithmetic. There is always a gap (1.6) between the true residual  $b - A\bar{x}_i$  and the recursively updated vectors  $\bar{r}_i$  (often called updated residuals). We denote this gap by  $f_i$ , i.e.,  $f_i \equiv \bar{r}_i - (b - A\bar{x}_i)$ . While the norm of the recursively updated residual  $\bar{r}_i$  eventually decreases beyond the level of unit roundoff (see the introduction for more details and references), the norm of the true residual  $b - A\bar{x}_i$  for  $i$  sufficiently large stagnates, and the size of  $f_i$  determines the *maximal attainable accuracy* (measured by the norm of the true residual).

It is important to investigate the maximal attainable accuracy of iterative methods and to use some measure for improving it, if needed. The effects of rounding errors, however, go beyond affecting only the maximal attainable accuracy. In particular, rounding errors can cause a significant *delay of convergence* that is of *primary importance* in challenging real-world problems, where computations are in most cases

stopped much before the maximal attainable accuracy is reached. Moreover, measures taken for improving maximal attainable accuracy can negatively affect delay of convergence; this is discussed further in section 3.5.

**3.1. Coupled two-term recurrences.** Let us demonstrate how  $\|f_i\|$  can be bounded in the HS CG (Algorithm 1.1). Given an initial approximation  $\bar{x}_0$ , the computed vectors satisfy

$$\bar{p}_0 = \bar{r}_0 = b - A\bar{x}_0 + f_0$$

and

$$(3.1) \quad \bar{x}_i = \bar{x}_{i-1} + \bar{\alpha}_{i-1}\bar{p}_{i-1} + \delta x_i,$$

$$(3.2) \quad \bar{r}_i = \bar{r}_{i-1} - \bar{\alpha}_{i-1}A\bar{p}_{i-1} + \delta r_i,$$

$$(3.3) \quad \bar{p}_i = \bar{r}_i + \bar{\beta}_i\bar{p}_{i-1} + \delta p_i,$$

where the individual  $\delta$ -terms account for the local roundoff errors. Assume the standard model of floating point arithmetic with unit roundoff  $\epsilon$ ; see, e.g., [27, relations (2.4)]. Then the vectors of local roundoff errors can be bounded in terms of the dimension  $N$  of the problem, machine precision  $\epsilon$ , the norm of  $A$ , and the norms of the associated vectors. To simplify the notation we give these bounds in the form

$$\begin{aligned} \|\delta x_i\| &\leq \mathcal{O}(\epsilon) \max\{\|\bar{x}_{i-1}\|, \|\bar{x}_i\|\}, \\ \|\delta r_i\| &\leq \mathcal{O}(\epsilon) \max\{\|\bar{r}_{i-1}\|, \|\bar{\alpha}_{i-1}A\bar{p}_{i-1}\|\}, \\ \|\delta p_i\| &\leq \mathcal{O}(\epsilon) \max\{\|\bar{r}_i\|, \|\bar{\beta}_i\bar{p}_{i-1}\|\}, \end{aligned}$$

where  $\mathcal{O}(\epsilon)$  denotes a term that can be bounded by  $\epsilon$  and a low degree polynomial<sup>1</sup> in  $N$ . The details can be found, e.g., in [56].

Combining (3.1) and (3.2) we can express the gap  $f_i$  in terms of the gap  $f_{i-1}$  and the local roundoff errors:

$$\begin{aligned} f_i &= \bar{r}_i - (b - A\bar{x}_i) \\ &= \bar{r}_{i-1} - \bar{\alpha}_{i-1}A\bar{p}_{i-1} + \delta r_i - (b - A(\bar{x}_{i-1} + \bar{\alpha}_{i-1}\bar{p}_{i-1} + \delta x_i)) \\ (3.4) \quad &= f_{i-1} + \delta r_i + A\delta x_i. \end{aligned}$$

Hence by induction it can be shown that

$$(3.5) \quad f_n = f_0 + \sum_{i=1}^n \delta r_i + A \sum_{i=1}^n \delta x_i.$$

The maximal attainable accuracy of the coupled two-term recurrences was analyzed, e.g., by Sleijpen, van der Vorst, and Fokkema [53] and Greenbaum [20]. In these

<sup>1</sup>It may seem that this standard simplification devalues the bounds obtained for large  $N$ . For an excessively large  $N$ , say,  $N = 10^{10}$  or more, the presence of terms like  $N^2$  can indeed devalue some worst-case assumptions concerning the conditioning of the problem, and this issue must be considered for emerging exascale computations. On the other hand, one should not forget the main principle stated very clearly by Wilkinson in the visionary paper [68] published in 1971: “*There is still a tendency to attach too much importance to the precise error bounds obtained by a priori error analysis. In my opinion, the bound itself is usually the least important part of it. The main object of such an analysis is to expose the potential instabilities, if any, of an algorithm so that, hopefully, from the insight thus obtained one might be led to improved algorithms. . .*” This principle has been emphasized many times since then, including in Parlett’s beautiful essay [40], and we aim to follow it in our paper.

works the gap  $f_i$  is analyzed in terms of the local errors associated with the two-term recurrences and it is shown that in this case the local errors simply accumulate. In particular, the authors of [53] prove that

$$(3.6) \quad \|f_n\| \leq \mathcal{O}(\epsilon)\|A\| \max_{i=0,\dots,n} \|x - \bar{x}_i\|,$$

and in [20] it is shown that

$$(3.7) \quad \|f_n\| \leq \mathcal{O}(\epsilon)\|A\| \left( \|x\| + \max_{i=0,\dots,n} \|\bar{x}_i\| \right).$$

The delay of convergence of CG in finite precision arithmetic has been analyzed, e.g., in papers by Greenbaum [19], Strakoš [60], Greenbaum and Strakoš [21], and Gergelits and Strakoš [16]; for more detailed comments and references see the introduction.

We now turn our attention to other variants of CG and investigate the individual sources of instabilities. In other words, we are interested in the question of how the attainable accuracy and delay of convergence are affected by various modifications to Algorithm 1.1. To illustrate the numerical effects, throughout this section we use the example of the linear system  $Ax = b$  with the matrix `bcsstk03` from the set `BCSSTRUC1` (BCS structural engineering matrices) in the SuiteSparse matrix collection [11], of order  $N = 112$ . We set  $x_0$  to be the zero vector. The right-hand-side  $b$  has been chosen such that  $b$  has equal components in the eigenvector basis and  $\|b\| = 1$ . The minimal and maximal eigenvalues of  $A$  are  $\lambda_{\min} \approx 2.9\text{e}+04$  and  $\lambda_{\max} \approx 2.0\text{e}+11$ , respectively, so that  $\kappa(A) \approx 6.9\text{e}+06$ . Since large outlying eigenvalues are present in the spectrum, one can expect a significant delay of convergence of CG in finite precision arithmetic; for more details, see, e.g., [60], [56], [16].

We note that the example problem we have chosen is a particularly difficult one for the CG method. In Figure 3.1 it can be seen that even exact CG requires  $N$  iterations for convergence and HS CG takes nearly  $7N$ . Our running example is not meant to be representative of a real-world problem but was rather chosen for its ability to expose the changes to attainable accuracy and convergence rate caused by the use of different CG recurrences in finite precision.

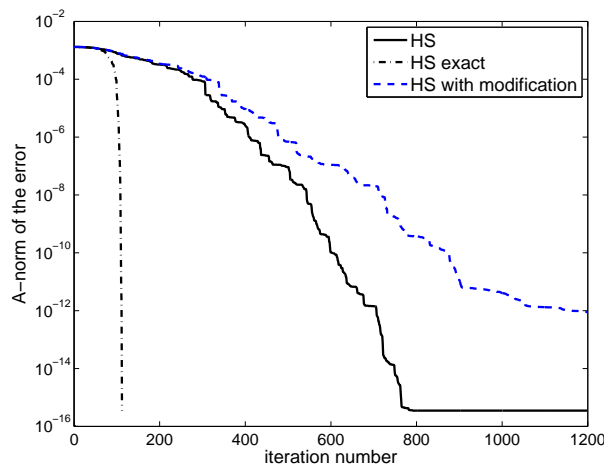


FIG. 3.1. Replacing (3.8) by (3.9) in Algorithm 1.1 can significantly change behavior in finite precision arithmetic.

This is not to say that such effects will not be observed in practical situations. In section 4.3 we present experiments demonstrating the numerical behavior of preconditioned HS CG and variants thereof on larger real-world problems. The decrease in attainable accuracy due to modified recurrences in GV CG observed in practical scenarios was the motivation behind the analysis and residual replacement strategy recently developed in [8], in which the reader can find additional examples of large-scale problems in which the attainable accuracy of GV CG (both preconditioned and unpreconditioned) deviates significantly from the HS CG method.

Although the investigation of maximal attainable accuracy is formulated in terms of residuals and their norms, in our experiments we purposely plot the  $A$ -norm of the error rather than the Euclidean norm of the residuals. The reason is twofold. First, the convergence behavior viewed in terms of the Euclidean norm of the residual can be oscillatory. Second, and even more importantly, as stated clearly in the seminal paper by Hestenes and Stiefel [26], the Euclidean norm of the residual is not an appropriate indicator of convergence for problems where  $\kappa(A)$  is substantially larger than one.

**3.2. Three-term recurrences.** Before considering the implementation using three-term recurrences, we present the following numerical illustration. Consider Algorithm 1.1, and replace the recurrence for updating the direction vector

$$(3.8) \quad p_i = r_i + \beta_i p_{i-1}$$

by the recurrence

$$(3.9) \quad p_i = r_i + \frac{\beta_i}{\alpha_{i-1}} (x_i - x_{i-1}).$$

At first glance, this looks like a benign change that should not cause any trouble since  $x_i$  is computed using the recurrence  $x_i = x_{i-1} + \alpha_{i-1} p_{i-1}$ . However, as we can observe in Figure 3.1, replacing (3.8) by (3.9) can have a substantial influence on the behavior of the algorithm in finite precision arithmetic, with respect to both the rate of convergence and the attained accuracy. This example shows that a seemingly trivial change can lead to substantially different numerical behavior. We also include in Figure 3.1 the result of exact HS CG computation (simulated via double reorthogonalization of the residual vectors). It shows that in the problem used for our illustration the effect of roundoff error is indeed substantial even for HS CG.

The coupled two-term recurrences used in Algorithm 1.1 can, in mathematical abstraction, be considered to be equivalent to the three-term recurrences (see section 2)

$$(3.10) \quad x_i = \alpha_{i-1} r_{i-1} + x_{i-1} + \frac{\alpha_{i-1} \beta_{i-1}}{\alpha_{i-2}} (x_{i-1} - x_{i-2}),$$

$$(3.11) \quad r_i = -\alpha_{i-1} A r_{i-1} + r_{i-1} + \frac{\alpha_{i-1} \beta_{i-1}}{\alpha_{i-2}} (r_{i-1} - r_{i-2}).$$

Various three-term CG implementations differ in the way the coefficients are computed in (3.10)–(3.11). The ST CG (see Algorithm 2.1) updates  $\alpha_{i-1}$  using the relation

$$(3.12) \quad \frac{1}{\alpha_{i-1}} = \frac{r_{i-1}^T A r_{i-1}}{\|r_{i-1}\|^2} - \frac{\beta_{i-1}}{\alpha_{i-2}}.$$

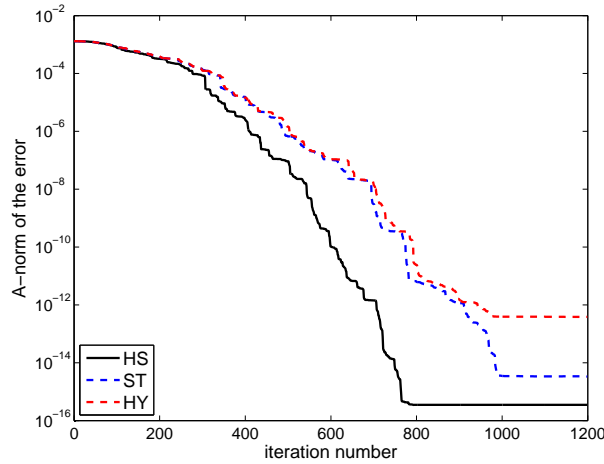


FIG. 3.2. HS CG (Algorithm 1.1), ST CG (Algorithm 2.1), and Hageman–Young (HY) [24, p. 143] implementations.

Once  $\alpha_{i-1}$  is known, one can compute  $r_i$  and evaluate  $\beta_i$  as in Algorithm 1.1. This is probably the most straightforward way of determining the coefficients in (3.10)–(3.11). Note that the implementation by Hageman and Young [24, p. 143] also uses (3.12) (in a scaled form) to derive the updating formula for the coefficient

$$\rho_{i-1} = 1 + \frac{\alpha_{i-1}\beta_{i-1}}{\alpha_{i-2}}.$$

This results in more complicated updating formulas. It appears that the Hageman–Young implementation can be more affected by rounding errors than ST CG; see Figure 3.2.

The gap  $f_i$  between the recursively updated residual  $\bar{r}_i$  and the true residual  $b - A\bar{x}_i$  in three-term recurrence implementations of CG was analyzed by Gutknecht and Strakoš [23]. They showed that the local roundoff errors incurred using (3.10)–(3.11) can be significantly amplified as they propagate through the recurrences. The amplification factors can be essentially expressed in terms of the coefficients in the recurrences (3.10)–(3.11) that can be related to the coefficients in (3.1)–(3.3), which can exhibit, in general, rather irregular behavior. It is therefore shown that using two three-term recurrences can lead to worse maximal attainable accuracy than using three two-term recurrences (Algorithm 1.1). In addition, numerical experiments indicate that using three-term recurrences can also significantly delay convergence in comparison to coupled two-term recurrences; see Figure 3.2.

**3.3. Adding auxiliary recurrences.** One can also use the updating formula (3.12) to compute the coefficients  $\alpha_i$  in HS CG. This allows the computations to be rearranged in a way that reduces the number of synchronization points to one. However, in HS CG, the vectors  $Ar_{i-1}$  are not available. So, either we have to perform two matrix-vector products per iteration or we need to use the additional recurrence for updating the vectors  $s_i \equiv Ap_i$

$$(3.13) \quad Ap_i = Ar_i + \beta_i Ap_{i-1}, \quad \text{i.e.,} \quad s_i = Ar_i + \beta_i s_{i-1};$$

cf. also Algorithm 2.2. A CG variant that uses the latter option is shown in Algorithm 3.1.

ALGORITHM 3.1. Conjugate gradient method with the recursively updated direction vectors.

**Input:** SPD matrix  $A \in R^{N \times N}$ , right-hand-side vector  $b \in R^N$ , initial approximation  $x_0 \in R^N$ , maximum number of iterations  $nmax$ .

**Output:** Approximate solution  $x_n$  after the algorithm has been stopped.

0. **Initialization:**  $r_0 = b - Ax_0$ ,  $p_0 = r_0$ ,  $s_0 = Ap_0$

1. **for**  $i = 1 : nmax$  **do**

$$2. \quad \alpha_{i-1} = \frac{(r_{i-1}, r_{i-1})}{(p_{i-1}, s_{i-1})}$$

$$3. \quad x_i = x_{i-1} + \alpha_{i-1} p_{i-1}$$

$$4. \quad r_i = r_{i-1} - \alpha_{i-1} s_{i-1}$$

5. evaluate the stopping criterion

$$6. \quad \beta_i = \frac{(r_i, r_i)}{(r_{i-1}, r_{i-1})}$$

$$7. \quad p_i = r_i + \beta_i p_{i-1}$$

$$8. \quad s_i = Ar_i + \beta_i s_{i-1}$$

9. **end do**

Adding a new recurrence without recomputing the recurrence coefficient  $\beta_i$  so that the computed vectors satisfy local orthogonality relations can negatively affect the numerical behavior of the resulting algorithm. Algorithm 3.1 is closely related to ChG CG (Algorithm 2.2), the only difference being in the way the coefficient  $\alpha_{i-1}$  is computed, which will be discussed in the next subsection. Note that other CG implementations that allow more parallelism than Algorithm 3.1 like GV CG (Algorithm 2.3) use even more auxiliary recurrences. In particular, Algorithm 2.3 updates the vectors  $s_i = Ap_i$ ,  $z_i = A^2 p_i$ ,  $w_i = Ar_i$ , and  $q_i = A^2 r_i$ . In problems sensitive to changes in the underlying distribution function that enlarge its support (see [41]), this can lead to a numerical disaster with respect to both delay of convergence and maximal attainable accuracy; see Figure 3.3. It should be noted that the sensitive problems mentioned above are in particular those with large outlying eigenvalues; for a detailed discussion see [16].

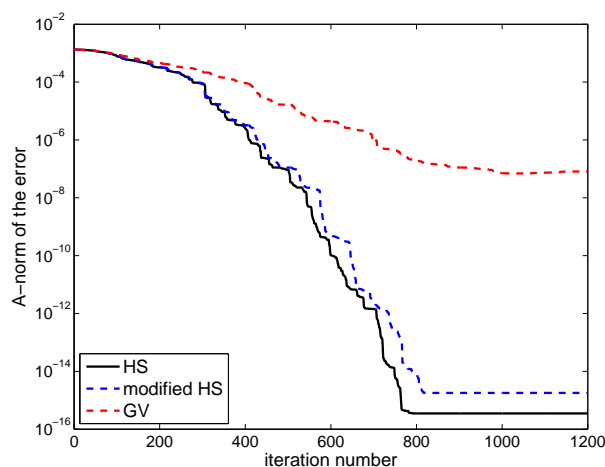


FIG. 3.3. HS CG (Algorithm 1.1), modified HS CG with the recursive update (3.13) (Algorithm 3.1), and GV CG (Algorithm 2.3).



**3.4. Recurrence coefficients.** As mentioned above, the recurrence coefficients in two-term as well as in three-term variants of CG can be computed in many different ways. We will illustrate the effects of changing the computation of the recurrence coefficients using a modified version of Algorithm 1.1 where the formula for updating  $\alpha_{i-1}$  is replaced by (3.12) without adding an auxiliary recurrence for  $Ar_{i-1}$  (the vectors  $Ar_{i-1}$ ,  $Ap_{i-1}$  are computed by explicit matrix-vector multiplication). This variant is shown in Algorithm 3.2.

ALGORITHM 3.2. Conjugate gradient method with modified coefficient computation.

**Input:** SPD matrix  $A \in R^{N \times N}$ , right-hand-side vector  $b \in R^N$ , initial approximation  $x_0 \in R^N$ , maximum number of iterations  $nmax$ .

**Output:** Approximate solution  $x_n$  after the algorithm has been stopped.

0. **Initialization:**  $r_0 = b - Ax_0$ ,  $p_0 = r_0$
1. **for**  $i = 1 : nmax$  **do**
2.      $\alpha_{i-1} = \left( \frac{r_{i-1}^T Ar_{i-1}}{\|r_{i-1}\|^2} - \frac{\beta_{i-1}}{\alpha_{i-2}} \right)^{-1}$
3.      $x_i = x_{i-1} + \alpha_{i-1} p_{i-1}$
4.      $r_i = r_{i-1} - \alpha_{i-1} Ap_{i-1}$
5.     *evaluate the stopping criterion*
6.      $\beta_i = \frac{(r_i, r_i)}{(r_{i-1}, r_{i-1})}$
7.      $p_i = r_i + \beta_i p_{i-1}$
9. **end do**

Notice that in this case, (3.4) still holds, so the gap  $f_i$  can be bounded analogously to (3.7). In fact, (3.4) will hold regardless of how  $\alpha_{i-1}$  is computed as long as the same  $\alpha_{i-1}$  is used in updating the recurrence both for  $x_i$  and for  $r_i$ ; this has been pointed out by, e.g., Greenbaum [20]. In Figure 3.4 we observe a significant delay of convergence of Algorithm 3.2 in comparison to HS CG. The maximal attainable accuracy seems to be similar for both algorithms.

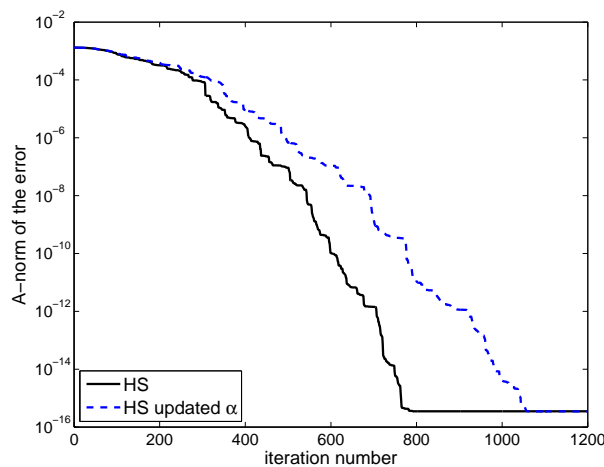


FIG. 3.4. HS CG (Algorithm 1.1), and the modified HS CG with explicit matrix-vectors multiplications  $Ar_{i-1}$ ,  $Ap_{i-1}$  and the formula (3.12) for updating  $\alpha_{i-1}$ ; see Algorithm 3.2.

**3.5. Residual replacement strategy.** To attain a higher accuracy in terms of the norm of the residual vector, several authors suggest reconciling the recursively computed residual  $\bar{r}_i$  and the true residual  $b - A\bar{x}_i$  by replacing  $\bar{r}_i$  by a direct computation of  $b - A\bar{x}_i$  in some iterations; see, e.g., [52], [66], [65], [3].

Since the direct computation of  $b - A\bar{x}_i$  is rather expensive, the residual replacement should occur only at a few iteration steps. As also emphasized in [66], these steps should be selected so that the residual replacement does not have an effect on the rate of convergence. Based on that, the authors in [66] develop an updating strategy where they set a threshold and carry out the residual replacement when a certain quantity containing the terms  $\mathcal{O}(\varepsilon)(\|A\|\|\bar{x}_i\| + \|\bar{r}_i\|)$  reaches the given threshold. The justification of the proposed heuristics assumes that the matrix composed of the computed residual vectors is of full column rank, which may not hold in practical computations.

In Figure 3.5 we plot the  $A$ -norm of the error in the exact computation using HS CG simulated as above (see Figure 3.1). The orthogonality and numerical linear independence among the computed residual vectors is lost in HS CG (solid line) very quickly, which results in a significant delay of convergence. Here, just to illustrate the possible effects of the residual replacement on the rate of convergence, we first carry out three residual replacements at steps  $n = 10$ ,  $n = 30$ , and  $n = 50$  before the linear independence of the computed residual vectors is lost. In this case, we indeed do not observe a significant difference in the convergence rate (see the red dashed line in Figure 3.5). However, three replacements after the loss of linear independence (at steps  $n = 400$ ,  $n = 450$ , and  $n = 500$ ) cause a significant delay of convergence in comparison to HS CG due to the introduction of large perturbations (relative to the residual norm at these steps) in the recurrence relations; see the blue dashed line. In Figure 3.5 we also plot the  $A$ -norm of the error in the algorithm with full residual replacement, in which  $\bar{r}_i$  is replaced by the direct computation of  $b - A\bar{x}_i$  in each iteration step. This costly strategy keeps the gap between the true and updated residuals on the level  $\mathcal{O}(\varepsilon)\|A\|(\|\bar{x}_i\| + \|x\|)$ . As expected based on the results in [66], this causes a significant deterioration of convergence rate. Similar behavior in terms

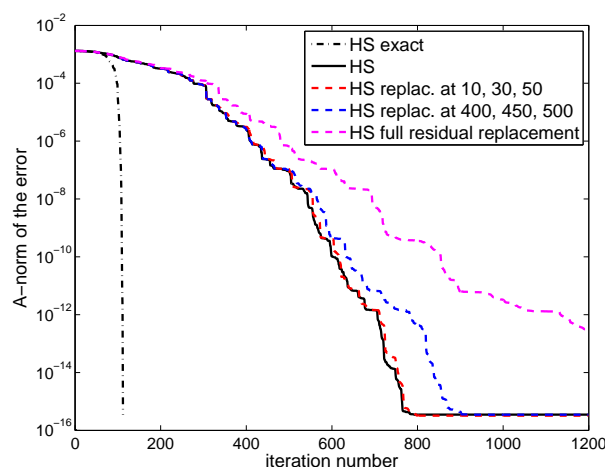


FIG. 3.5. *The HS CG implementation and the modified CG implementations with residual replacement.*

of the delay of convergence when using various residual replacement strategies has also been observed by Meurant [38, section 6.4], who suggests an optional residual replacement heuristic. Residual replacement strategies deserve further attention and thorough analysis.

We point out that for this chosen test problem, HS CG already attains an accurate approximate solution, and therefore we do not see in Figure 3.5 further improvement from the use of residual replacement.

**4. Methodology for numerical stability analysis.** In the following we describe a methodology for the numerical stability analysis of CG variants that use an auxiliary recurrence for the matrix-vector product  $Ap_i$ , which includes pipelined CG (GV CG) [18]. Since the methodologies for analysis of the maximal attainable accuracy and delay of convergence are very different, we will divide this section into two parts.

**4.1. Maximal attainable accuracy.** A full analysis of the attainable accuracy for all CG variants described in this paper would require many pages of technical text, which would make the paper unacceptably long. We will therefore start by deriving a general expression for the residual gap that applies to any coupled two-term recurrence version of CG that performs updates to  $x_i$  and  $r_i$  using the recurrences

$$(4.1) \quad x_i = x_{i-1} + \alpha_{i-1}p_{i-1},$$

$$(4.2) \quad r_i = r_{i-1} - \alpha_{i-1}s_{i-1},$$

where  $s_i \equiv Ap_i$  is also recursively updated. We then demonstrate how one can use this result to bound the maximum attainable accuracy by doing so for Algorithm 3.1, which is the CG variant “closest” to HS CG (Algorithm 1.1) that uses recurrences (4.1) and (4.2). The purpose is to provide a methodology that can be used, with more or less straightforward extensions, to perform full-scale analysis of algorithms like GV CG. An example of such an extension that uses the methodology suggested in this paper (see also [23]) can be found, e.g., in [8].

We will show that the maximal attainable accuracy of Algorithm 3.1 can be significantly worse than in HS CG (Algorithm 1.1) and, like ST CG (Algorithm 2.1), can suffer from amplification of local errors if there are large oscillations of the residual norm. It is worth pointing out that in Figure 3.3, this difference is not substantial, which just underlines the need for extensive experiments in order to gather reliable numerical evidence about the behavior of any implementation for a particular class of problems.

Consider a variant of HS CG that performs updates to  $x_i$  and  $r_i$  using the recurrences (4.1) and (4.2) for  $i = 1, \dots, n - 1$ . Note that Algorithms 3.1, 2.2 (ChG CG), and 2.3 (GV CG) all fall into this category. In finite precision arithmetic (denoted here with bars), the updates satisfy

$$(4.3) \quad \bar{x}_i = \bar{x}_{i-1} + \bar{\alpha}_{i-1}\bar{p}_{i-1} + \delta x_i, \quad \|\delta x_i\| \leq \mathcal{O}(\epsilon) \max\{\|\bar{x}_{i-1}\|, \|\bar{x}_i\|\},$$

$$(4.4) \quad \bar{r}_i = \bar{r}_{i-1} - \bar{\alpha}_{i-1}\bar{s}_{i-1} + \delta r_i, \quad \|\delta r_i\| \leq \mathcal{O}(\epsilon) \max\{\|\bar{r}_{i-1}\|, \|\bar{\alpha}_{i-1}\bar{s}_{i-1}\|\}.$$

In iteration  $i$ , the gap  $f_i$  between the recursively updated residual  $\bar{r}_i$  and the true residual  $b - A\bar{x}_i$  can be written in the form

$$(4.5) \quad \begin{aligned} f_i &= \bar{r}_i - (b - A\bar{x}_i) \\ &= \bar{r}_{i-1} - \bar{\alpha}_{i-1}\bar{s}_{i-1} + \delta r_i - (b - A(\bar{x}_{i-1} + \bar{\alpha}_{i-1}\bar{p}_{i-1} + \delta x_i)) \\ &= f_{i-1} - \bar{\alpha}_{i-1}g_{i-1} + \delta r_i + A\delta x_i, \end{aligned}$$

where

$$g_{i-1} = \bar{s}_{i-1} - A\bar{p}_{i-1}$$

denotes the gap between the recursively computed vector  $\bar{s}_{i-1}$  and the computed direction vector  $\bar{p}_{i-1}$  multiplied by  $A$ . Hence, by induction, after  $n$  iterations,

$$(4.6) \quad f_n = f_0 - \sum_{i=1}^n \bar{\alpha}_{i-1} g_{i-1} + \sum_{i=1}^n \delta r_i + A \sum_{i=1}^n \delta x_i.$$

We see from (4.6) that  $f_n$  is a superposition of the effects of local errors  $\delta x_i$  and  $\delta r_i$  and the differences  $g_{i-1}$  multiplied by the coefficients  $\bar{\alpha}_{i-1}$  throughout all iteration steps  $i = 1, \dots, n$ . Introducing the notation  $G_n = (g_0, \dots, g_{n-1})$  and  $d_n = (\bar{\alpha}_0, \dots, \bar{\alpha}_{n-1})^T$ , the formula (4.6) can be written as

$$(4.7) \quad f_n = f_0 + \sum_{i=1}^n (\delta r_i + A\delta x_i) - G_n d_n.$$

Comparing the above formula to (3.5), it is clear that the quantity  $G_n$  plays a key role in determining how the accuracy of a particular variant compares with HS CG. We stress that the above formula for the residual gap applies to *any* CG variant that uses the recurrences (4.1) and (4.2) to update  $x_i$  and  $r_i$ . From here, the maximum attainable accuracy of a particular variant, e.g., GV CG, can be analyzed by obtaining a bound on the size of  $G_n$ .

We will now demonstrate this process by bounding  $\|G_n\|$  for the CG variant closest to HS CG that uses recurrences (4.1) and (4.2), Algorithm 3.1. We will show that the sizes of the  $g_i$  terms may play a decisive role in the maximum attainable accuracy of Algorithm 3.1. In particular, we will show that the size of  $g_i$  can be significantly larger than the error in the floating point computation  $\text{fl}(A\bar{p}_i)$ , i.e.,

$$\|\text{fl}(A\bar{p}_i) - A\bar{p}_i\| \leq \mathcal{O}(\epsilon)\|A\|\|\bar{p}_i\|.$$

Introducing the notation

$$P_n = (p_0, \dots, p_{n-1}), \quad R_n = (r_0, \dots, r_{n-1}), \quad \text{and} \quad S_n = (s_0, \dots, s_{n-1}),$$

the recurrences (3.8) and (3.13) (lines 7 and 8 in Algorithm 3.1) for  $i = 1, \dots, n-1$  can be written

$$(4.8) \quad R_n = P_n U_n, \quad A R_n = S_n U_n,$$

where  $U_n$  is the  $n$ -by- $n$  unit upper bidiagonal matrix of the form

$$(4.9) \quad U_n = \begin{pmatrix} 1 & -\beta_1 & 0 & 0 \\ 0 & 1 & \ddots & 0 \\ \vdots & \ddots & 1 & -\beta_{n-1} \\ 0 & \dots & 0 & 1 \end{pmatrix}.$$

The quantities computed in finite precision arithmetic (denoted here with bars<sup>2</sup>) satisfy

<sup>2</sup>Mathematically (assuming exact arithmetic),  $s_i = Ap_i$ . In finite precision computations, however,  $\bar{s}_i$  computed via (3.13) can no longer be identified with the multiplication  $A\bar{p}_i$ . This point is very important. Examples of confusion mixing mathematically equivalent but computationally different quantities can be seen in recently published literature.

$$(4.10) \quad \bar{p}_i = \bar{r}_i + \bar{\beta}_i \bar{p}_{i-1} + \delta p_i, \quad \|\delta p_i\| \leq \mathcal{O}(\epsilon) \max\{\|\bar{r}_i\|, \|\bar{\beta}_i \bar{p}_{i-1}\|\},$$

$$(4.11) \quad \bar{s}_i = A\bar{r}_i + \bar{\beta}_i \bar{s}_{i-1} + \delta s_i, \quad \|\delta s_i\| \leq \mathcal{O}(\epsilon) \max\{\|A\|\|\bar{r}_i\|, \|\bar{\beta}_i \bar{s}_{i-1}\|\},$$

along with (4.3) and (4.4), where  $i = 1, \dots, n - 1$ . Similarly to (4.8), the recurrences (4.10) and (4.11) for  $i = 1, \dots, n - 1$  can be rewritten into

$$(4.12) \quad \bar{R}_n = \bar{P}_n \bar{U}_n - \Delta P_n, \quad A\bar{R}_n = \bar{S}_n \bar{U}_n - \Delta S_n,$$

where  $\Delta P_n = (\delta p_0, \dots, \delta p_{n-1})$  and  $\Delta S_n = (\delta s_0, \dots, \delta s_{n-1})$ . It is clear from (4.10) and (4.11) that

$$\begin{aligned} \|\delta p_i\| &\leq \mathcal{O}(\epsilon)(\|\bar{r}_i\| + \|\bar{\beta}_i \bar{p}_{i-1}\|) \leq \mathcal{O}(\epsilon)(\|\bar{R}_n\| + \|\bar{P}_n\| \|\bar{U}_n\|), \\ \|\delta s_i\| &\leq \mathcal{O}(\epsilon)(\|A\|\|\bar{r}_i\| + \|\bar{\beta}_i \bar{s}_{i-1}\|) \leq \mathcal{O}(\epsilon)(\|A\|\|\bar{R}_n\| + \|\bar{S}_n\| \|\bar{U}_n\|) \end{aligned}$$

for each  $i = 0, \dots, n - 1$ . Since  $\|\Delta P_n\| \leq \|\Delta P_n\|_F \leq \sqrt{n} \max_{i=0, \dots, n-1} \|\delta p_i\|$  and  $\|\Delta S_n\| \leq \|\Delta S_n\|_F \leq \sqrt{n} \max_{i=0, \dots, n-1} \|\delta s_i\|$ , we obtain the bounds

$$\begin{aligned} \|\Delta P_n\| &\leq \mathcal{O}(\epsilon)(\|\bar{R}_n\| + \|\bar{P}_n\| \|\bar{U}_n\|), \\ \|\Delta S_n\| &\leq \mathcal{O}(\epsilon)(\|A\|\|\bar{R}_n\| + \|\bar{S}_n\| \|\bar{U}_n\|). \end{aligned}$$

From (4.12) we have

$$G_n = \bar{S}_n - A\bar{P}_n = (\Delta S_n - A\Delta P_n)\bar{U}_n^{-1}.$$

Taking the norms and assuming that the matrix  $\bar{U}_n$  is numerically nonsingular with  $\mathcal{O}(\epsilon)\kappa(\bar{U}_n) < 1$ , we get

$$\begin{aligned} \|G_n\| &\leq (\|\Delta S_n\| + \|A\|\|\Delta P_n\|) \|\bar{U}_n^{-1}\| \\ &\leq \mathcal{O}(\epsilon) [(\|\bar{S}_n\| + \|A\|\|\bar{P}_n\|)\|\bar{U}_n\| + 2\|A\|\|\bar{R}_n\|] \|\bar{U}_n^{-1}\| \\ &\leq \mathcal{O}(\epsilon) [(\|G_n\| + 2\|A\|\|\bar{P}_n\|)\|\bar{U}_n\| + 2\|A\|\|\bar{R}_n\|] \|\bar{U}_n^{-1}\| \\ (4.13) \quad &\leq \frac{\mathcal{O}(\epsilon)}{1 - \mathcal{O}(\epsilon)\kappa(\bar{U}_n)} [\kappa(\bar{U}_n)\|A\|\|\bar{P}_n\| + \|A\|\|\bar{R}_n\|] \|\bar{U}_n^{-1}\|. \end{aligned}$$

The first term in the bound (4.13) is larger than  $\mathcal{O}(\epsilon)\|A\|\|\bar{P}_n\|$  by a factor of  $\kappa(\bar{U}_n)$  and it suggests that the local errors at the individual steps contributing to the size of the difference  $f_n = \bar{r}_n - (b - A\bar{x}_n)$  in (4.7) can potentially be amplified by the entries in  $\bar{U}_n^{-1}$ . Since the matrix  $\bar{U}_n$  is unit upper bidiagonal, the entries of its inverse can be expressed in terms of the products of the coefficients  $\bar{\beta}_i$ , for  $i = 1, \dots, n - 1$ , i.e.,

$$(4.14) \quad \bar{U}_n^{-1} = \begin{pmatrix} 1 & \bar{\beta}_1 & \dots & \dots & \bar{\beta}_1 \bar{\beta}_2 \dots \bar{\beta}_{n-1} \\ 0 & 1 & \bar{\beta}_2 & \dots & \bar{\beta}_2 \dots \bar{\beta}_{n-1} \\ \vdots & \ddots & \ddots & \ddots & \vdots \\ \vdots & & \ddots & 1 & \bar{\beta}_{n-1} \\ 0 & \dots & \dots & 0 & 1 \end{pmatrix}.$$

In exact arithmetic, the coefficient  $\beta_i$  in Algorithm 3.1 is equal to  $\beta_i = \|r_i\|^2 / \|r_{i-1}\|^2$  and the multiplicative factors  $\beta_i \beta_{i+1} \dots \beta_j$  are equal to

$$\beta_i \beta_{i+1} \dots \beta_j = \frac{\|r_j\|^2}{\|r_{i-1}\|^2}, \quad i < j.$$

Consequently, due to possible oscillations of the CG residual norms, the factors  $\|r_j\|^2/\|r_{i-1}\|^2$  can for some  $i < j$  be rather large. Similar reasoning can also be applied to the norm of the updated residuals  $\|\bar{r}_\ell\|$ , leading potentially to large products of the computed coefficients  $\bar{\beta}_i\bar{\beta}_{i+1}\dots\bar{\beta}_j$  in (4.14) and to the ill-conditioned matrix  $\bar{U}_n$ . Consequently, a dramatic amplification of the local errors can be expected and the maximum attainable accuracy in computations using Algorithm 3.1 can be significantly worse than when the standard Algorithm 1.1 is used.

This analysis resembles the results for the three-term CG implementations presented in [23]. It demonstrates that a seemingly innocuous change in the CG implementation can turn the two-term recurrence CG into a three-term recurrence CG with substantially different (and often worse) numerical behavior. Generalization of the presented analysis to pipelined CG implementations will use the same methodology. It will, however, require a detailed and cautious evaluation of the various terms and a thorough discussion of the obtained bounds as well as comparison with properly designed numerical experiments. For the first step in this direction we refer to [8].

**4.2. Delay of convergence.** As explained in the introduction, CG is closely related to the Lanczos process (1.2) and it is mathematically equivalent to the so-called Lanczos-CG formulated by (1.2)–(1.3). Using Algorithm 1.1, the entries of the Jacobi matrix

$$T_n = \begin{pmatrix} \gamma_1 & \delta_2 & & & \\ \delta_2 & \ddots & \ddots & & \\ & \ddots & \ddots & \delta_n & \\ & & & \delta_n & \gamma_n \end{pmatrix}$$

are determined via the formulas (here  $\alpha_{-1} = 1$ ,  $\beta_0 = 0$ )

$$(4.15) \quad \gamma_i = \frac{1}{\alpha_{i-1}} + \frac{\beta_{i-1}}{\alpha_{i-2}}, \quad i = 1, 2, \dots, n,$$

$$(4.16) \quad \delta_i = \frac{\sqrt{\beta_i}}{\alpha_{i-1}}, \quad i = 2, 3, \dots, n;$$

see, e.g., [21, section 2]. The sequence of Jacobi matrices  $T_n$ ,  $n = 1, 2, \dots$ , determines the sequence of the distribution functions  $\omega^{(n)}(\lambda)$  with the points of increase  $\theta_\ell^{(n)}$ ,  $\ell = 1, \dots, n$ , given by the eigenvalues of  $T_n$  and the weights  $\omega_\ell^{(n)}$ ,  $\ell = 1, \dots, n$ , given by the squared first elements of the associated (normalized) eigenvectors; see, e.g., [32, sections 3.4 and 3.5]. The sequence  $\omega^{(n)}(\lambda)$  approximates in the sense of the Gauss–Christoffel quadrature the distribution function  $\omega(\lambda)$ , which is determined by the eigenvalues of the matrix  $A$  and the squared size of the projections of the normalized initial residual  $r_0/\|r_0\|$  onto the associated invariant subspaces; see, e.g., [32, section 3.5].

In finite precision CG computations, the computed quantities  $\bar{\alpha}_i$  and  $\bar{\beta}_i$  differ from their exact precision counterparts, and this difference can easily grow to many orders of magnitude. Therefore, the counterparts to (4.15) and (4.16),

$$(4.17) \quad \bar{\gamma}_i \equiv \frac{1}{\bar{\alpha}_{i-1}} + \frac{\bar{\beta}_{i-1}}{\bar{\alpha}_{i-2}} \quad \text{and} \quad \bar{\delta}_i \equiv \frac{\sqrt{\bar{\beta}_i}}{\bar{\alpha}_{i-1}},$$

respectively, which define the Jacobi matrices associated with the finite precision computations, can also differ many orders of magnitude from the exact arithmetic

quantities  $\gamma_i$  and  $\delta_i$ . In order to avoid misunderstanding we point out that the issue is not in finite precision computations of  $\bar{\gamma}_i$  and  $\bar{\delta}_i$  from  $\bar{\alpha}_i$  and  $\bar{\beta}_i$ ; here  $\bar{\gamma}_i$  and  $\bar{\delta}_i$  are determined from  $\bar{\alpha}_i$  and  $\bar{\beta}_i$  exactly using the mathematical equalities (4.17).

Now we come to the main point that builds upon the results of Greenbaum and others recalled in the introduction. In exact arithmetic, the CG computations in steps 1 to  $n$  can be associated with the sequence of nested Jacobi matrices

$$T_1, T_2, \dots, T_n$$

determined by (4.15) and (4.16). For each  $n = 1, 2, \dots$ , the matrix  $T_n$  defines the distribution function  $\omega^{(n)}(\lambda)$  that approximates the distribution function  $\omega(\lambda)$  determined by the original data  $A$ ,  $b$ , and  $x_0$ , in the sense of the  $n$ th Gauss–Christoffel quadrature. Moreover, the CG method represents in fact the matrix formulation of the Gauss–Christoffel quadrature (see [32, section 3.5]) and, quantitatively, the squared  $A$ -norm of the CG error for the function  $f(\lambda) = \lambda^{-1}$  is given as the scaled quadrature error (see [32, relation (3.5.9)]),

$$(4.18) \quad \int \lambda^{-1} d\omega(\lambda) = \sum_{\ell=1}^n \omega_\ell^{(n)} \left\{ \theta_\ell^{(n)} \right\}^{-1} + \frac{\|x - x_n\|_A^2}{\|r_0\|^2}.$$

Summarizing, it is clear that in exact arithmetic the sequence of the distribution functions  $\omega^{(n)}$ ,  $n = 1, 2, \dots$ , determines convergence behavior of CG in the steps  $n = 1, 2, \dots$ .

In finite precision arithmetic, the CG computations in steps 1 to  $n$  can be associated with the sequence of Jacobi matrices

$$\bar{T}_1, \bar{T}_2, \dots, \bar{T}_n$$

determined via (4.17). Despite possible loss of orthogonality and even loss of (numerical) linear independence among the computed CG residual and direction vectors, we thus have a *mathematical structure of the sequence of Jacobi matrices that is preserved regardless of the numerical inaccuracies*. Here  $\bar{T}_n$  again defines the distribution function  $\bar{\omega}^{(n)}(\lambda)$ . This distribution function is, however, in general very different from  $\omega^{(n)}(\lambda)$  determined by exact computations, and it depends very significantly on the particular CG implementation. In particular, in finite precision CG computations using HS CG (Algorithm 1.1) the distribution function  $\bar{\omega}_{\text{HS}}^{(n)}(\lambda)$  can be very different from the distribution function  $\bar{\omega}_{\text{GV}}^{(n)}(\lambda)$  determined by the finite precision computations using GV CG (Algorithm 2.3). This is illustrated in Figure 4.1, which depicts the difference in the entries of the Jacobi matrices  $\bar{T}_n$  for different CG implementations using the same problem `bcsstk03` (see section 3.2).

The crucial question is whether  $\bar{\omega}^{(n)}(\lambda)$  determined by the finite precision CG computations in steps 1 through  $n$  can be associated with some distribution function  $\bar{\omega}_{1:n}(\lambda)$  *meaningfully related to the original distribution function*  $\omega(\lambda)$  such that we can write, analogously to (4.18),

$$(4.19) \quad \int \lambda^{-1} d\omega(\lambda) \approx \int \lambda^{-1} d\bar{\omega}_{1:n}(\lambda) = \sum_{\ell=1}^n \bar{\omega}_\ell^{(n)} \left\{ \bar{\theta}_\ell^{(n)} \right\}^{-1} + \frac{\|x - \bar{x}_n\|_A^2}{\|r_0\|^2} + F_n,$$

where  $F_n$  is small in comparison to the error term  $\|x - \bar{x}_n\|_A^2 / \|r_0\|^2$ . The positive answer was given by Greenbaum for Algorithm 1.1 in [19]; see also [21] and the expositions in [41], [39], [56, section 5], and [32, section 5.9.1]. Such a distribution

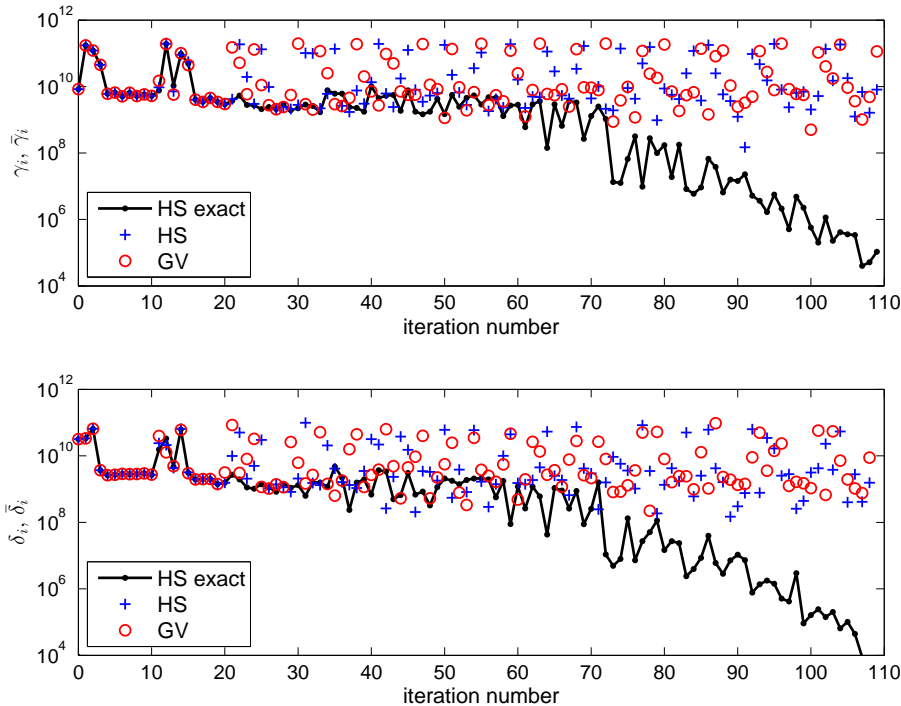


FIG. 4.1. Differences in the entries of the Jacobi matrices for HS and GV CG implementations. Solid line gives the values of  $\gamma_i$ ,  $\delta_i$  for the exact HS CG (Algorithm 1.1) determined via (4.15)–(4.16). Pluses (+) give the values of  $\bar{\gamma}_i$ ,  $\bar{\delta}_i$  for the finite precision computations using HS CG (Algorithm 1.1) determined via (4.17). Circles (o) give the analogous quantities for GV CG (Algorithm 2.3).

function  $\bar{\omega}_{1:n}(\lambda)$  indeed exists for every  $n = 1, 2, \dots$ , and its points of increase are close to eigenvalues of  $A$ . We point out that this fundamental result is highly nontrivial and is based on the work of Paige [42], [43], [44], [45].

An analogous analysis for the implementations of CG like Algorithm 2.3, which must include derivation of the term  $F_n$  in (4.19), are yet to be done. While this is beyond the scope of the present work, we claim such an analysis will reveal that in some cases, one can expect Algorithm 2.3 to exhibit a substantially larger delay of convergence due to roundoff than Algorithm 1.1. As explained in [32, section 5.9.1], such cases are related to the distribution of eigenvalues of  $A$  having large gaps in the spectrum and/or outlying large eigenvalues. For illustration, the eigenvalues of the matrix `bcsstk03` are plotted in Figure 4.2.

It is well known that the delay of convergence is related to the approximation of some eigenvalues by multiple Ritz values, and this is determined by the underlying (case-dependent) distribution function. A detailed explanation with references to extensive literature on the topic is given, e.g., in [32, section 5.9.1]. Due to the high nonlinearity of the involved phenomena, there is no simple rule on how to decide whether the gaps in the spectrum or a separation of large eigenvalues is significant or not. *Any practical occurrence* of these is potentially important and therefore it should be of concern. For the example with the matrix `bcsstk03` the situation is obvious, as very many repeated Ritz value approximations of the largest eigenvalues are formed as the computation proceeds; see Figure 4.3.



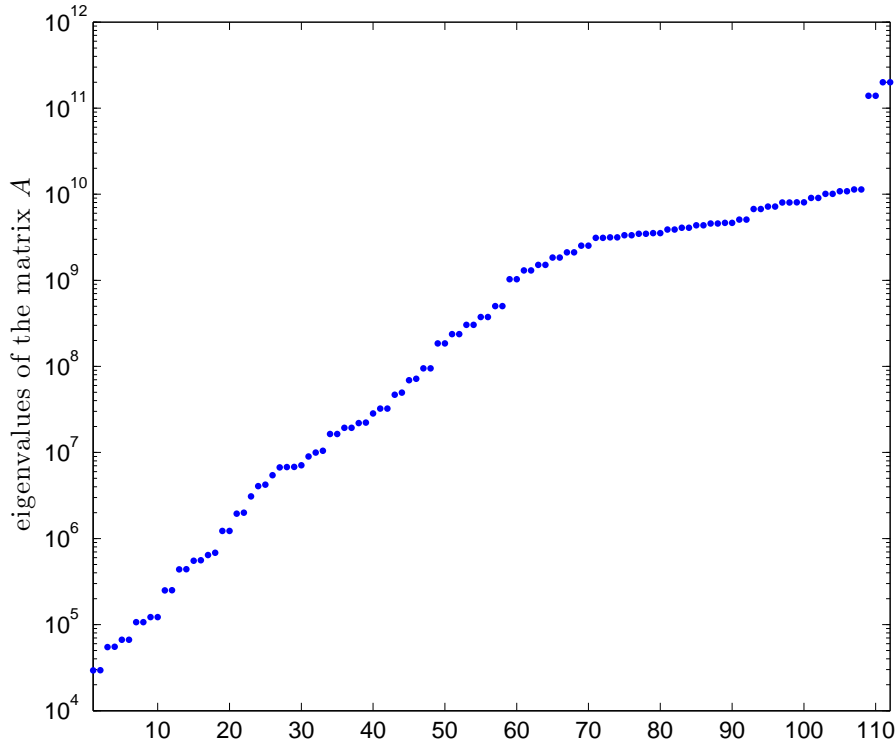


FIG. 4.2. Distribution of the eigenvalues of the matrix `bcsstk03`.

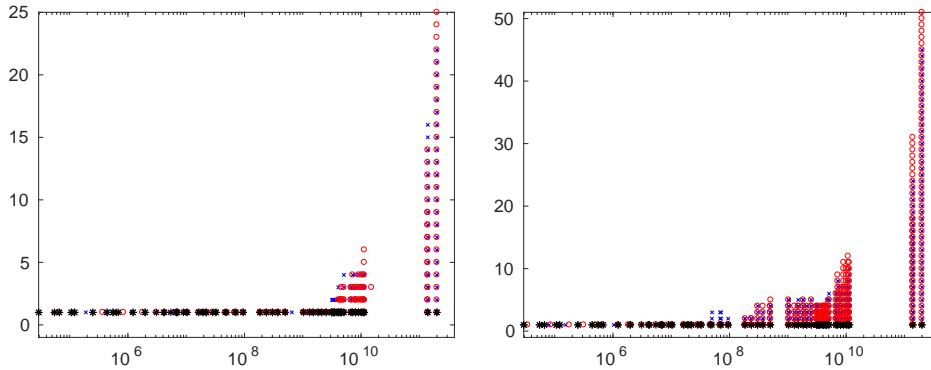


FIG. 4.3. Approximation of the eigenvalues of the matrix `bcsstk03` (denoted by stars) by the Ritz values, i.e., the eigenvalues of the Jacobi matrices for HS CG (denoted by crosses) and GV CG (denoted by circles) implementations in double precision; see also Figure 4.1. The left plot corresponds to iteration 200 and the right plot to iteration 400. The Ritz values plotted on the x-axis form vertical columns whenever they approximate the same eigenvalue, with the number of the symbols in the column giving the multiplicity of the approximation. As expected, we observe the devastating effect of the gap in the spectrum that is only about one order of magnitude (this may wrongly be considered insignificant relative to the large condition number). More multiple Ritz values are produced by GV CG than HS CG, which corresponds to its worse convergence behavior.

**4.3. Numerical examples.** To demonstrate the effects of finite precision on the behavior of CG variants, we consider two linear systems arising from real-world

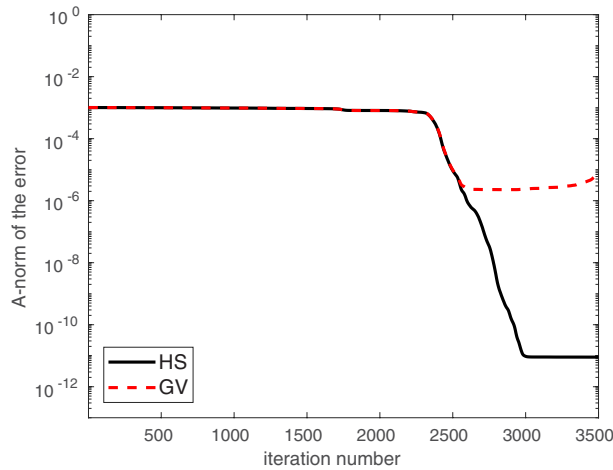


FIG. 4.4. A comparison of the numerical behavior of HS CG and GV CG for the system with matrix `s3dkt3m2` with incomplete Cholesky preconditioning.

applications: the first system with the matrix `s3dkt3m2`, and the second one with the matrix `tube`. For both systems we use preconditioning as described below; for more details and experiments focused on error estimation in preconditioned CG, see [57]. The experiments were run using double precision in MATLAB 9.2.0.538062 (R2017a).

The matrix `s3dkt3m2` of order  $N = 90,449$ , with  $\kappa(A) \approx 3.6e + 11$ , can be downloaded from the CYLSHELL collection in the Matrix Market library (<https://math.nist.gov/MatrixMarket/>), which contains matrices that represent low-order finite-element discretizations of a shell element test, the pinched cylinder. Only the last element of the right-hand-side vector  $b$  is nonzero, which corresponds to the given physical problem; for more details see [31]. The factor  $L$  in the preconditioner  $M = LL^T$  is determined by the incomplete Cholesky (`ichol`) factorization with threshold dropping, `type = 'ict'`, `droptol = 1e-5`, and with the global diagonal shift `diagcomp = 1e-2`. Note that here  $\text{nnz}(A) = 3,686,223$  and  $\text{nnz}(L) = 6,541,916$ .

The second system `tube` was provided to us by R. Kouhia, the author of the CYLSHELL collection. The tube is a cylindrical shell with constant wall thickness, loaded with an axial stress distribution at both ends. The mesh is refined at the center, and it is almost uniform toward the ends. The order of the matrix  $A$  is  $N = 21,498$ ,  $\kappa(A) \approx 5.6e + 08$ , and  $\text{nnz}(A) = 894,490$ . The factor  $L$  in the preconditioner  $M = LL^T$  is determined again by `ichol` with threshold dropping, `type = 'ict'`, `droptol = 1e-5`, and with the global diagonal shift `diagcomp = 1e-3`. The resulting matrix  $L$  has  $\text{nnz}(L) = 3,415,487$  nonzero entries.

In Figures 4.4 and 4.5 we compare HS CG and GV CG for the problems with `s3dkt3m2` and `tube`, respectively. We use the standard preconditioned HS CG and the preconditioned version of GV CG given in the original paper [18, Algorithm 4, p. 229]. Note that preconditioned GV CG uses two additional auxiliary recurrences in comparison to GV CG.

We can see in Figure 4.4 that while the GV CG method initially follows the same convergence trajectory as HS CG, after the  $A$ -norm of the error reaches around  $10^{-5}$  the convergence of GV CG stagnates. This represents a loss of accuracy of about 5 orders of magnitude. In Figure 4.5, we can see both a loss of attainable accuracy and a slight delay of convergence. The curve for GV CG begins to deviate from HS CG

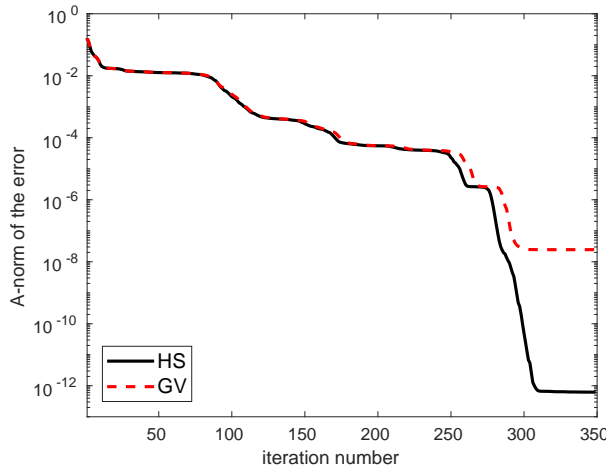


FIG. 4.5. A comparison of the numerical behavior of HS CG and GV CG for the system with matrix `tube` with incomplete Cholesky preconditioning.

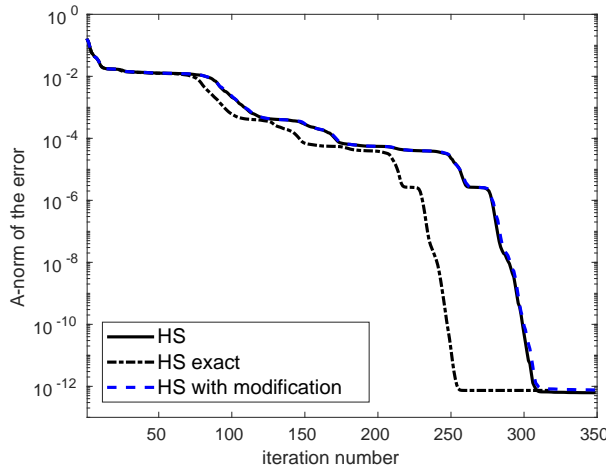


FIG. 4.6. A comparison of the numerical behavior of HS CG and HS CG with slight modification (cf. Figure 3.1) for the system with matrix `tube` with incomplete Cholesky preconditioning.

after about 250 iterations and stagnates soon after the  $A$ -norm of the error reaches  $10^{-7}$  (which is again in this case a loss of accuracy of about 5 orders of magnitude). It is clear from these examples that preconditioning does not serve to eliminate the described effects of finite precision error.

As an additional illustration, in Figure 4.6 we compare exact CG, HS CG, and HS CG with the slight modification described in section 3.2 where (3.8) is replaced by (3.9) for the system with matrix `tube`. We use the same incomplete Cholesky preconditioner for the `tube` matrix previously described. Comparing this plot with Figure 3.1 for the unpreconditioned linear system with matrix `bcsstk03`, we see that the effect of using the modified recurrence is much less drastic in this case (although this can change for larger sizes of the problem). This emphasizes the point that determining whether the observable convergence delay and loss of accuracy due to recurrence modifications are significant enough to affect practical use is *entirely problem dependent*.

**5. Conclusions.** As mentioned in the introduction, Krylov subspace methods are highly nonlinear in the input data  $A$  and  $b$ . This makes them incredibly powerful methods in terms of being efficient solvers but also makes them inherently difficult to analyze and understand. In this work, we present an overview of the historical context of the analysis of Krylov subspace method variants in finite precision. We study the effects of various modifications to the standard CG method on convergence rate and attainable accuracy for an example problem. The modifications studied include using three-term recurrences, adding auxiliary recurrences, changing the way in which the recurrence coefficients are computed, and using residual replacement. Finally, we outline a methodology for analyzing the maximum attainable accuracy and convergence delay in variants of CG that use an auxiliary recurrence for the matrix-vector product  $Ap_i$ . We give an expression for the residual gap which is applicable to any such variant, including the pipelined CG method of Ghysels and Vanroose [18]. We show that, similar to three-term recurrence CG, any variant of CG which uses an auxiliary recurrence for  $Ap_i$  in updating the residual vector can suffer from amplification of local rounding errors due to large residual oscillations.

As we as a community push toward the goal of exascale-level computational science, it is important that we not forget the bigger picture of what we want to accomplish: enabling scientific insight, analysis, and discovery through the use of computation. New variants of Krylov subspace methods are often borne out of efforts to minimize time per iteration. While an important piece of the puzzle, one must be cautious about using time per iteration as a metric of efficiency. Alone, this fails to capture the whole picture of the method's effectiveness within the context of a scientific application. If the modifications made to the method cause a convergence delay with a greater effect than the per-iteration performance improvement, the modified method may actually be slower than the standard approach in a practical setting. If the modifications to the method cause a significant decrease in attainable accuracy, then the method may have no use in some scientific applications, making any gains in iteration speed for naught.

It is clear then that in the landscape of high-performance computing, the design and implementation of iterative solvers requires a *holistic* approach. In selecting the right method and parameters to use for a given problem, we must consider the expected time per iteration on various problem sizes and architectures as well as the numerical stability and convergence properties, which we have seen depend heavily on the particular recurrences used in the method as well as numerical properties of the input data. We must also consider the performance and use of the Krylov subspace method within the context of the overall scientific application. We stress that the derivation of a new method or technique, the evaluation of its performance on various problems and architectures, the theoretical analysis of its behavior in finite precision, and the application to practical problems need not be accomplished simultaneously; this process is in practice usually iterative and highly collaborative in nature.

In light of this, we give a few recommendations on what must be done to sufficiently analyze the finite precision behavior of a new Krylov subspace method variant (or a new technique to be used together with an existing variant, e.g., residual replacement). Both the maximum attainable accuracy and delay of convergence should be studied by using, e.g., the methodology outlined in section 4. Ideally, such studies should be done both by deriving theoretical results and by performing a thorough experimental evaluation of the method's numerical behavior. We stress that theoretical results on convergence rate *assuming exact computation* cannot be applied to inexact computations unless such application is rigorously justified by a thorough

analysis (this point is of particular importance for methods using short recurrences). It is also worth investigating whether the new variant is fundamentally different from other variants previously studied in the literature, or whether such differences are only superficial. For example, as we have shown in sections 3.3 and 4, certain modified variants of the coupled two-term recurrence version of CG very closely resemble the three-term recurrence method ST CG.

Related to this, it is very rare that Krylov subspace methods can be used in practice without some transformation of the problem, generally called preconditioning. Technically, this involves introducing additional matrix-vector operations within each iteration step (typically solving a linear algebraic system with the matrix representing the preconditioner). In numerical stability analysis, preconditioning can be included using the standard methodology developed for the unpreconditioned methods. Since preconditioning is closely linked with discretization (see, e.g., [33]), high-performance computing efficiency may stimulate new developments that will consider this link instead of elaborating only on the level of the algebraic system arising from the separated discretization step.

We are optimistic that pipelined Krylov subspace methods and other synchronization-reducing variants can certainly be useful in achieving practical speedups for certain problems in a number of application domains; there are likely many problems for which the convergence delay in methods like GV CG does not negate the savings from reduced synchronization and for which the resulting attainable accuracy is acceptable. However, as numerical analysts, it is our task to identify how significant these effects can be in finite precision, to identify classes of problems for which such methods can be used, and to improve our overall understanding of the trade-offs between reducing synchronization and maintaining the numerical properties of the method. As we work toward enabling computational science at the exascale level, it is our hope that this manuscript helps guide the way toward the future design and analysis of efficient Krylov subspace methods.

## REFERENCES

- [1] G. M. BALLARD, *Avoiding Communication in Dense Linear Algebra*, Ph.D. thesis, University of California, Berkeley, 2013.
- [2] A. BOURAS AND V. FRAYSSÉ, *Inexact matrix-vector products in Krylov methods for solving linear systems: A relaxation strategy*, SIAM J. Matrix Anal. Appl., 26 (2005), pp. 660–678.
- [3] E. CARSON AND J. DEMMEL, *A residual replacement strategy for improving the maximum attainable accuracy of  $s$ -step Krylov subspace methods*, SIAM J. Matrix Anal. Appl., 35 (2014), pp. 22–43.
- [4] E. CARSON, *Communication-Avoiding Krylov Subspace Methods in Theory and Practice*, Ph.D. thesis, University of California, Berkeley, 2015.
- [5] Y. CHAUVET AND G. MEURANT, *Multitasking on the CRAY X-MP*, J. Systems Software, 2 (1986), pp. 17–20.
- [6] A. T. CHRONOPOULOS AND C. W. GEAR, *On the efficient implementation of preconditioned  $s$ -step conjugate gradient methods on multiprocessors with memory hierarchy*, Parallel Comput., 11 (1989), pp. 37–53.
- [7] A. T. CHRONOPOULOS AND C. W. GEAR,  *$s$ -step iterative methods for symmetric linear systems*, J. Comput. Appl. Math., 25 (1989), pp. 153–168.
- [8] S. COOLS, E. F. YETKIN, E. AGULLO, L. GIRAUD AND W. VANROOSE, *Analyzing the effect of local rounding error propagation on the maximal attainable accuracy of the pipelined conjugate gradients method*, SIAM J. Matrix Anal. Appl., 39 (2018), pp. 426–450.
- [9] E. D’AZEVEDO, V. ELJKHOUT, AND C. ROMINE, *LAPACK Working Note 56: Reducing Communication Costs in the Conjugate Gradient Algorithm on Distributed Memory Multiprocessors*, Technical report UT-CS-93-185, University of Tennessee, Knoxville, 1993.

- [10] E. D'AZEVEDO AND C. ROMINE, *Reducing Communication Costs in the Conjugate Gradient Algorithm on Distributed Memory Multiprocessors*, Technical report ORNL/TM-12192, Oak Ridge National Laboratory, 1992.
- [11] T. A. DAVIS AND Y. HU, *The university of Florida sparse matrix collection*, ACM Trans. Math. Software, 38 (2011), pp. 1:1–1:25.
- [12] J. DONGARRA ET AL., *The international exascale software project roadmap*, Int. J. High Perform. Comput. Appl., 25 (2011), pp. 3–60.
- [13] M. ENGELI, TH. GINSBURG, H. RUTISHAUSER, AND E. STIEFEL, *Refined iterative methods for computation of the solution and the eigenvalues of self-adjoint boundary value problems*, Mitt. Inst. Angew. Math. Zürich. 8 (1959), p. 107.
- [14] B. FISCHER, *Polynomial Based Iteration Methods for Symmetric Linear Systems*, Wiley-Teubner Ser. Adv. Numer. Math., Wiley, Chichester, UK, 1996.
- [15] D. GANNON AND J. VAN ROSENDALE, *Parallel architectures for iterative methods on adaptive, block structured grids*, in Elliptic Problem Solvers, II, Academic Press, Orlando, FL, 1984, pp. 93–104.
- [16] T. GERGELITS AND Z. STRAKOŠ, *Composite convergence bounds based on Chebyshev polynomials and finite precision conjugate gradient computations*, Numer. Algorithms, 65 (2014), pp. 759–782.
- [17] P. GHYSELS, T. ASHBY, K. MEERBERGEN, AND W. VANROOSE, *Hiding global communication latency in the GMRES algorithm on massively parallel machines*, SIAM J. Sci. Comput., 35 (2013), pp. C48–C71.
- [18] P. GHYSELS AND W. VANROOSE, *Hiding global synchronization latency in the preconditioned conjugate gradient algorithm*, Parallel Comput., 40 (2014), pp. 224–238.
- [19] A. GREENBAUM, *Behavior of slightly perturbed Lanczos and conjugate-gradient recurrences*, Linear Algebra Appl., 113 (1989), pp. 7–63.
- [20] A. GREENBAUM, *Estimating the attainable accuracy of recursively computed residual methods*, SIAM J. Matrix Anal. Appl., 18 (1997), pp. 535–551.
- [21] A. GREENBAUM AND Z. STRAKOŠ, *Predicting the behavior of finite precision Lanczos and conjugate gradient computations*, SIAM J. Matrix Anal. Appl., 13 (1992), pp. 121–137.
- [22] W. GROPP, *Update on Libraries for Blue Waters, A Presentation on the Third Workshop of the INRIA-Illinois Joint-Laboratory on Petascale Computing*, 2010, <http://jointlab-pc.ncsa.illinois.edu/events/workshop3/agenda.html>.
- [23] M. H. GUTKNECHT AND Z. STRAKOŠ, *Accuracy of two three-term and three two-term recurrences for Krylov space solvers*, SIAM J. Matrix Anal. Appl., 22 (2000), pp. 213–229.
- [24] L. A. HAGEMAN AND D. M. YOUNG, *Applied Iterative Methods*, Comput. Sci. Appl. Math., Academic Press, New York, 1981.
- [25] M. R. HESTENES, *Iterative methods for solving linear equations*, J. Optim. Theory Appl., 11 (1951), pp. 323–334.
- [26] M. R. HESTENES AND E. STIEFEL, *Methods of conjugate gradients for solving linear systems*, J. Res. National Bureau of Standards, 49 (1952), pp. 409–435.
- [27] N. J. HIGHAM, *Accuracy and Stability of Numerical Algorithms*, SIAM, Philadelphia, 2002.
- [28] M. HOEMMEN, *Communication-Avoiding Krylov Subspace Methods*, Ph.D. thesis, University of California, Berkeley, 2010.
- [29] L. JOHNSON, *Highly Concurrent Algorithms for Solving Linear Systems of Equations*, Technical report 5079-TR-83, California Institute of Technology, 1983.
- [30] L. JOHNSON, *Highly concurrent algorithms for solving linear systems of equations*, in Elliptic Problem Solvers, II, Academic Press, Orlando, FL, 1984, pp. 105–126.
- [31] R. KOUHIA, *Description of the CYLSHELL set*, in The University of Florida Sparse Matrix Collection, ACM Trans. Math. Software 38, ACM, New York, NY, 2011, pp. 1:1–1:25.
- [32] J. LIESEN AND Z. STRAKOŠ, *Krylov Subspace Methods: Principles and Analysis*, Oxford University Press, Oxford, NY, 2013.
- [33] J. MÁLEK AND Z. STRAKOŠ, *Preconditioning and the Conjugate Gradient Method in the Context of Solving PDEs*, SIAM Spotlight Series, SIAM, Philadelphia, 2015.
- [34] K. MCMANUS, S. JOHNSON, AND M. CROSS, *Communication latency hiding in a parallel conjugate gradient method*, in Proceedings of the 11th International Conference on Domain Decomposition Methods, London, 1998, DDM.org, Augsburg, 1999, pp. 306–313.
- [35] G. MEURANT, *Numerical Experiments for the Preconditioned Conjugate Gradient Method on the Cray-X-MP/2*, Technical report LBL-18023, University of California, Berkeley, 1984.
- [36] G. MEURANT, *Multitasking the conjugate gradient method on the CRAY X-MP/48*, Parallel Comput., 5 (1987), pp. 267–280.
- [37] G. MEURANT, *Computer Solution of Large Linear Systems*, Elsevier, Amsterdam, 1999.

- [38] G. MEURANT, *The Lanczos and Conjugate Gradient Algorithms: From Theory to Finite Precision Computations*, Software Environ. Tools 19, SIAM, Philadelphia, 2006.
- [39] G. MEURANT AND Z. STRAKOŠ, *The Lanczos and conjugate gradient algorithms in finite precision arithmetic*, Acta Numer., 15 (2006), pp. 471–542.
- [40] S. G. NASH, ED., *A History of Scientific Computing*, ACM Press, New York, 1990.
- [41] D. P. O’LEARY, Z. STRAKOŠ, AND P. TICHÝ, *On sensitivity of Gauss-Christoffel quadrature*, Numer. Math., 107 (2007), pp. 147–174.
- [42] C. PAIGE, *The Computation of Eigenvalues and Eigenvectors of Very Large Sparse Matrices*, Ph.D. thesis, London University, 1971.
- [43] C. PAIGE, *Computational variants of the Lanczos method for the eigenproblem*, IMA J. Appl. Math., 10 (1972), pp. 373–381.
- [44] C. PAIGE, *Error analysis of the Lanczos algorithm for tridiagonalizing a symmetric matrix*, IMA J. Appl. Math., 18 (1976), pp. 341–349.
- [45] C. PAIGE, *Accuracy and effectiveness of the Lanczos algorithm for the symmetric eigenproblem*, Linear Algebra Appl., 34 (1980), pp. 235–258.
- [46] J. K. REID, *On the method of conjugate gradients for the solution of large sparse systems of linear equations*, in Large Sparse Sets of Linear Equations (Proc. Conf., St. Catherine’s Coll., Oxford, 1970), Academic Press, London, 1971, pp. 231–254.
- [47] J. B. ROSSER, *Rapidly converging iterative methods for solving linear equations*, in Simultaneous Linear Equations and the Determination of Eigenvalues, National Bureau of Standards Appl. Math. Ser., 29, U.S. Government Printing Office, Washington, DC, 1953, pp. 59–64.
- [48] Y. SAAD, *Practical use of polynomial preconditionings for the conjugate gradient method*, SIAM J. Sci. Statist. Comput., 6 (1985), pp. 865–881.
- [49] Y. SAAD, *Krylov subspace methods on supercomputers*, SIAM J. Sci. Statist. Comput., 10 (1989), pp. 1200–1232.
- [50] V. SIMONCINI AND D. SZYLD, *Theory of inexact Krylov subspace methods and applications to scientific computing*, SIAM J. Sci. Comput., 25 (2003), pp. 454–477.
- [51] V. SIMONCINI AND D. SZYLD, *On the occurrence of superlinear convergence of exact and inexact Krylov subspace methods*, SIAM Rev., 47 (2005), pp. 247–272.
- [52] G. L. G. SLEIJPEN AND H. A. VAN DER VORST, *Reliable updated residuals in hybrid Bi-CG methods*, Computing, 56 (1996), pp. 141–163.
- [53] G. L. G. SLEIJPEN, H. A. VAN DER VORST, AND D. R. FOKKEMA, *BiCGstab(l) and other hybrid Bi-CG methods*, Numer. Algorithms, 7 (1994), pp. 75–109.
- [54] E. STIEFEL, *Ausgleichung ohne Aufstellung der Gaußschen Normalgleichungen*, Wiss. Z. Technische Hochschule Dresden, 2 (1952/53), pp. 441–442.
- [55] E. STIEFEL, *Relaxationsmethoden bester Strategie zur Lösung linearer Gleichungssysteme*, Comment. Math. Helv., 29 (1955), pp. 157–179.
- [56] Z. STRAKOŠ AND P. TICHÝ, *On error estimation in the conjugate gradient method and why it works in finite precision computations*, Electron. Trans. Numer. Anal., 13 (2002), pp. 56–80.
- [57] Z. STRAKOŠ AND P. TICHÝ, *Error estimation in preconditioned conjugate gradients*, BIT, 45 (2005), pp. 789–817.
- [58] Z. STRAKOŠ, *Performance of the EC 2345 array processor*, Computers Artificial Intelligence, 4 (1985), pp. 273–284.
- [59] Z. STRAKOŠ, *Effectivity and optimizing of algorithms and programs on the host-computer/array-processor system*, Parallel Comput., 4 (1987), pp. 189–207.
- [60] Z. STRAKOŠ, *On the real convergence rate of the conjugate gradient method*, Linear Algebra Appl., 154/156 (1991), pp. 535–549.
- [61] G. SZEGÖ, *Orthogonal Polynomials*, 4th ed., Amer. Math. Soc. Colloq. Publ., AMS, Providence, RI, 1939.
- [62] J. VAN DEN ESHOF, *Nested Iteration Methods for Nonlinear Matrix Problems*, Ph.D. thesis, Utrecht University, 2003.
- [63] J. VAN DEN ESHOF AND G. SLEIJPEN, *Inexact Krylov subspace methods for linear systems*, SIAM J. Matrix Anal. Appl., 26 (2004), pp. 125–153.
- [64] H. VAN DER VORST, *Parallel Iterative Solution Methods for Linear Systems Arising from Discretized PDE’s, Special Course on Parallel Computing in CFD*, Technical report URN:NBN:NL:UI:10-1874-1640, Utrecht University, 2001, <http://dspace.library.uu.nl/handle/1874/1640>.
- [65] H. A. VAN DER VORST, *Iterative Krylov Methods for Large Linear Systems*, Cambridge Monogr. Appl. Comput. Math. 13, Cambridge University Press, Cambridge, UK, 2003.

- [66] H. A. VAN DER VORST AND Q. YE, *Residual replacement strategies for Krylov subspace iterative methods for the convergence of true residuals*, SIAM J. Sci. Comput., 22 (2000), pp. 835–852.
- [67] J. VAN ROSENDALE, *Minimizing Inner Product Data Dependencies in Conjugate Gradient Iteration*, Technical report 172178, ICASE-NASA, 1983.
- [68] J. WILKINSON, *Modern error analysis*, SIAM Rev., 13 (1971), pp. 548–568.