

ACCELERATED INEXACT NEWTON SCHEMES FOR LARGE SYSTEMS OF NONLINEAR EQUATIONS*

DIEDERIK R. FOKKEMA[†], GERARD L. G. SLEIJPEN[‡], AND
HENK A. VAN DER VORST[‡]

Abstract. Classical iteration methods for linear systems, such as Jacobi iteration, can be accelerated considerably by Krylov subspace methods like GMRES. In this paper, we describe how inexact Newton methods for nonlinear problems can be accelerated in a similar way and how this leads to a general framework that includes many well-known techniques for solving linear and nonlinear systems, as well as new ones. Inexact Newton methods are frequently used in practice to avoid the expensive exact solution of the large linear system arising in the (possibly also inexact) linearization step of Newton’s process. Our framework includes acceleration techniques for the “linear steps” as well as for the “nonlinear steps” in Newton’s process. The described class of methods, the accelerated inexact Newton (AIN) methods, contains methods like GMRES and GMRESR for linear systems, Arnoldi and Jacobi–Davidson for linear eigenproblems, and many variants of Newton’s method, like damped Newton, for general nonlinear problems. As numerical experiments suggest, the AIN approach may be useful for the construction of efficient schemes for solving nonlinear problems.

Key words. nonlinear problems, Newton’s method, inexact Newton, iterative methods

AMS subject classification. 65H10

PII. S1064827595296148

1. Introduction. Our goal in this paper is twofold. A number of iterative solvers for linear systems of equations, such as the full orthogonalization method (FOM) [23], the generalized minimal residual method (GMRES) [26], the generalized congruent residual method (GCR) [31], the flexible GMRES method [25], the GMRES recursive method (GMRESR) [29], and the GCR orthogonal method (GCRO) [7], are in structure very similar to iterative methods for linear eigenproblems, like shift and invert Arnoldi [1, 24], Davidson [6, 24], and Jacobi–Davidson [28]. We will show that all these algorithms can be viewed as instances of an accelerated inexact Newton (AIN) scheme (cf. Algorithm 3) when applied to either linear equations or linear eigenproblems. This observation may help us in the design and analysis of algorithms by “transporting” algorithmic approaches from one application area to another. Moreover, our aim is to identify efficient AIN schemes for nonlinear problems as well, and we will show how we can learn from the algorithms for linear problems.

To be more specific, we will be interested in the numerical approximation of the solution u of the nonlinear equation

$$(1.1) \quad F(u) = 0,$$

where F is some smooth (nonlinear) map from a domain in \mathbb{R}^n (or \mathbb{C}^n) that contains the solution u , into \mathbb{R}^n (or \mathbb{C}^n), where n is typically large.

Some special types of systems of equations will play an important motivating role in this paper.

*Received by the editors December 8, 1995; accepted for publication (in revised form) April 15, 1996.

<http://www.siam.org/journals/sisc/19-2/29614.html>

[†]ISE Integrated Systems Engineering AG, Technopark Zurich, Technoparkstr. 1, CH-8005 Zurich, Switzerland (fokkema@ise.ch).

[‡]Mathematical Institute, Utrecht University, P. O. Box 80.010, NL-3508 TA Utrecht, the Netherlands (sleijpen@math.ruu.nl, vorst@math.ruu.nl).

The first type is the linear system of equations

$$(1.2) \quad Ax = b,$$

where A is a nonsingular matrix and b , x are vectors of appropriate size; A and b are given, x is unknown. The dimension n of the problem is typically large, and A is often sparse. With $u = x$ and $F(u) := b - Au$, (1.2) is equivalent to (1.1). This type will serve as the main source of inspiration for our ideas.

The second type concerns the generalized linear eigenproblem

$$(1.3) \quad Av = \lambda Bv.$$

With $u = (v, \lambda)$ we have that, for normalized v , with $F(u) := Av - \lambda Bv$, (1.3) is equivalent to (1.1). This type is an example of a mildly nonlinear system and will serve as an illustration for the similarity between various algorithms, seen as instances of AIN (see section 6).

However, the AIN schemes, that we will discuss, will be applicable to more general nonlinear problems like, for instance, equations that arise from discretizing nonlinear partial differential equations of the form

$$(1.4) \quad -\nabla(a\nabla u) + bg(u)\nabla u + ch(u) = f \text{ on } \Omega,$$

where Ω is a domain in \mathbb{R}^2 or \mathbb{R}^3 , $a, b \in C^1(\bar{\Omega})$, $c \in C(\bar{\Omega})$, $h, g \in C^1(\mathbb{R})$, and $f \in L^2(\Omega)$, and u satisfies suitable boundary conditions. An example of (1.4) is, for instance,

$$(1.5) \quad -\Delta u - \lambda e^u = 0 \text{ on } \Omega,$$

where Ω is some domain in \mathbb{R}^2 and $u = 0$ on $\partial\Omega$ (see also section 8).

Guided by the known approaches for the linear system (cf. [25, 29, 7]) and the eigenproblem (cf. [28, 27]), we will define AIN schemes for more general nonlinear systems. This leads to a combination of Krylov subspace methods for inexact Newton (cf. [16, 4] and also [8]) with acceleration techniques (as in [2]) and offers us an overwhelming choice of techniques for further improving the efficiency of Newton-type methods. As a side-effect this leads to a surprisingly simple framework for the identification of many well-known methods for linear and nonlinear problems and eigenproblems.

Our numerical experiments for nonlinear problems, like problem (1.5), serve as an illustration for the usefulness of our approach.

The rest of this paper is organized as follows. In section 2 we briefly review the ideas behind the inexact Newton method. In section 3 we introduce the AIN methods. We will examine how iterative methods for linear problems are accelerated, and we will distinguish between a Galerkin approach and a minimal residual (MR) approach. These concepts are then extended to the nonlinear case. In section 4 we make some comments on the implementation of AIN schemes. In section 5 we show how many well-known iterative methods for linear problems fit into the AIN framework. In sections 6 and 7 we consider instances of AIN for the mildly nonlinear generalized eigenproblem and for more general nonlinear problems. In section 8 we present our numerical results, and some concluding remarks are made in section 9.

2. Inexact Newton methods. Newton-type methods are very popular for solving systems of nonlinear equations as represented, for instance, by (1.4). If u_k is the approximate solution at iteration number k , Newton's method requires, for the next

1. Set $k = -1$ and choose an initial approximation u_0 .
2. Repeat until u_k is accurate enough:
 - (a) $k = k + 1$.
 - (b) Compute the residual $r_k = F(u_k)$.
 - (c) Compute an approximation J_k for the Jacobian $F'(u_k)$.
 - (d) *Solve the correction equation (approximately).* Compute an (approximate) solution p_k of the correction equation

$$J_k p = -r_k.$$

- (e) *Update.* Compute the new approximation:
 $u_{k+1} = u_k + p_k$.

ALGORITHM 1. *Inexact Newton.*

approximate solution of (1.1), the evaluation of the Jacobian $J_k := F'(u_k)$ and the solution of the *correction equation*

$$(2.1) \quad J_k p = -r_k, \text{ where } r_k := F(u_k).$$

Unfortunately, it may be very expensive, or even practically impossible, to determine the Jacobian and/or to solve the correction equation exactly, especially for larger systems.

In such situations one aims for an approximate solution of the correction equation, possibly with an approximation for the Jacobian (see, e.g., [8]). Algorithm 1 is an algorithmical representation of the resulting inexact Newton scheme.

If, for instance, Krylov subspace methods are used for the approximate solution of the correction equation (2.1), then only directional derivatives are required (cf. [8, 4]); there is no need to evaluate the Jacobian explicitly. If v is a given vector, then the vector $F'(u)v$ can be approximated using the fact that

$$\frac{F(u + \varepsilon v) - F(u)}{\varepsilon} \rightarrow F'(u)v \text{ for } \varepsilon \rightarrow 0.$$

The combination of a Krylov subspace method with directional derivatives combines steps 2(c) and 2(d) in Algorithm 1.

For an initial guess u_0 sufficiently close to a solution, Newton's method has asymptotically at least quadratic convergence behavior. However, this quadratic convergence is usually lost if one uses inexact variants, and often the convergence is not much faster than linear. In the next section we make suggestions on how this (linear) speed of convergence may be improved.

Note. It is our aim to restore, as much as possible, the asymptotic convergence behavior of exact Newton; we do not address the question of global convergence.

3. Accelerating inexact Newton methods. Newton's method is a one-step method; that is, in each step, Newton's method updates the approximate solution with information from the previous step only. However, in the computational process, subspaces have been built gradually that contain useful information concerning the problem. This information may be exploited to improve the current approximate

1. Set $k = -1$ and choose an initial approximation x_0 .
2. Repeat until x_k is accurate enough:
 - (a) $k = k + 1$.
 - (b) $r_k = b - Ax_k$.
 - (c) $p_k = \text{diag}(A)^{-1}r_k$.
 - (d) $x_{k+1} = x_k + p_k$.

ALGORITHM 2. *Jacobi iteration.*

solution, and this is what we propose to do. More precisely, we will consider alternative update strategies for step 2(e) of Algorithm 1.

3.1. Acceleration in the linear case. The linear system $Ax = b$ can be written as $F(x) := b - Ax = 0$, and $-A$ is the Jacobian of F . When the approximate solution p_k is computed as $p_k = M^{-1}r_k$, where M is some preconditioning matrix (approximating A), then the inexact Newton algorithm, Algorithm 1, reduces to a standard Richardson-type iteration process for the splitting $A = M - R$, for instance, the choice $M = D$, where $D = \text{diag}(A)$ leads to Jacobi iteration (see Algorithm 2).

One may improve the convergence behavior of standard iteration schemes by

- using more sophisticated preconditioners M , and/or
- applying acceleration techniques in the update step.

Different preconditioners and different acceleration techniques lead to different algorithms, some of which are well known.

Examples of iteration schemes that use more sophisticated preconditioners are, for instance, Gauss–Seidel iteration, where $M = L + D$, with L being the strict lower triangular part of A and $D = \text{diag}(A)$, and successive overrelaxation (SOR), where $M = \omega L + D$ and ω is a relaxation parameter.

Examples of iteration schemes that use acceleration techniques are algorithms that take their updates to the approximate solution as a linear combination of previous directions p_j . Preferable updates $\tilde{p}_k := \sum_{j \leq k} \gamma_j p_j$ are those for which $b - Ax_{k+1}$, where $x_{k+1} = x_k + \tilde{p}_k$, is minimal in some sense: e.g., $\|b - Ax_{k+1}\|_2$ is minimal, as in GMRES [26] and GCR [31], or $b - Ax_{k+1}$ is orthogonal to the p_j for $j \leq k$, as in FOM or generalized conjugate gradient (GENCG) [23], or $b - Ax_{k+1}$ is “quasi-minimal,” as in biconjugate gradient (Bi-CG) [17] and quasi-minimal residual (QMR) [11].

Of course, the distinction between preconditioning and acceleration is not a clear one. Acceleration techniques with a limited number of steps can be seen as a kind of dynamic preconditioning, as opposed to the static preconditioning with fixed M . In this view one is again free to choose an acceleration technique. Examples of such iteration schemes are flexible GMRES [25], GMRESR [29], and GCRO [7].

All these accelerated iteration schemes for linear problems construct approximations $x_{k+1} = x_0 + V_k y_k$, $V_k := [p_1, p_2, \dots, p_k]$, with y_k the solution of a *smaller* or an *easier* projected problem. For example, GMRES computes y_k such that $\|b - A(x_k + V_k y_k)\|_2$ is minimal, or equivalently $(AV_k)^*(b - A(x_k + V_k y_k)) = 0$; FOM computes y_k such that $V_k^*(b - A(x_k + V_k y_k)) = 0$, whereas Bi-CG and QMR compute y_k as the solution of a larger tridiagonal problem obtained with oblique projections.

For stability (and efficiency) reasons one usually constructs another basis for the span of V_k with certain orthogonality properties, depending on the selected approach.

3.2. Acceleration in the nonlinear case. We are interested in methods for finding a zero of a general nonlinear mapping F . For the linear case, the methods

mentioned above are, apart from the computation of the residual, essentially a mix of two components: (1) the computation of a new search direction (which involves the residual), and (2) the update of the approximation (which involves the current search direction and possibly previous search directions, and the solution y_k of a projected problem). The first component may be interpreted as preconditioning, while the second component may be seen as the acceleration.

Looking more carefully on how y_k in the linear case is computed, we can distinguish between two approaches based on two different conditions. With $G_k(y) := F(x_k + V_k y)$, the FOM and the other “oblique” approaches lead to methods that compute y such that (for appropriate W_k) $W_k^* G_k(y) = 0$ (a *Galerkin* condition), whereas the GMRES approach leads to methods that compute y such that $\|G_k(y)\|_2$ is minimal (a MR condition).

From these observations for the linear case we now can formulate iteration schemes for the nonlinear case.

The inexact Newton iteration can be accelerated in a similar way as the standard linear iteration. This acceleration can be accomplished by updating the solution by a correction \tilde{p}_k in the subspace spanned by all correction directions p_j ($j \leq k$).

To be more precise, the update \tilde{p}_k for the approximate solution is given by $\tilde{p}_k = V_k y$, where $V_k = [v_1, v_2, \dots, v_k]$ for some basis (v_j) of the *search space* \mathcal{V}_k spanned by p_1, p_2, \dots, p_k . Furthermore, with $G_k(y) = F(u_k + V_k y)$, we propose to determine y by

- a *Galerkin* condition on $G_k(y)$: y is a solution of

$$(3.1) \quad W_k^* G_k(y) = 0,$$

where W_k is some matrix of the same dimensions as V_k ;

- a MR condition on $G_k(y)$: y is a solution of

$$(3.2) \quad \min_y \|G_k(y)\|_2;$$

- or a mix of both, a *restricted minimal residual* (RMR) condition on $G_k(y)$: y is a solution of

$$(3.3) \quad \min_y \|W_k^* G_k(y)\|_2.$$

Equation (3.1) generalizes the FOM approach, while (3.2) generalizes the GMRES approach.

Solving (3.1) means that the component of the residual r_{k+1} in the subspace \mathcal{W}_k (spanned by W_k) vanishes. For W_k one may choose, for instance, $W_k = V_k$ (as in FOM), or $W_k = [W_{k-1}, w_k]$ where w_k is the component of $J_k p_k$ orthogonal to W_{k-1} (as in GMRES: $w_k = (I - W_{k-1} W_{k-1}^*) J_k p_k$). For linear equations, the MR and the Galerkin approaches coincide for the last choice.

As is known from the linear case, a complication of the Galerkin approach is that (3.1) may have no solution, which means that this approach may lead to breakdown of the method. In order to circumvent this shortcoming to some extent, we have formulated the RMR approach (3.3). Compared to (3.1) this formulation is also attractive for another reason: one can apply standard Gauss–Newton [9] schemes for solving general nonlinear least squares problems to it. One might argue that a drawback of a Gauss–Newton scheme is that it may converge slowly (or not at all). However, for least squares problems with zero residual solutions, the asymptotic speed of convergence of a Gauss–Newton method is that of Newton’s method. This means that if the Galerkin problem (3.1) has a solution, a Gauss–Newton scheme applied to (3.3) will find it fast and efficient (see also section 8).

1. Set $k = -1$ and choose an initial approximation u_0 .
2. $V = []$, $W = []$.
3. Repeat until u_k is accurate enough:
 - (a) $k = k + 1$.
 - (b) Compute the residual $r_k = F(u_k)$.
 - (c) Compute an approximation J_k for the Jacobian $F'(u_k)$.
 - (d) *Solve the correction equation (approximately)*. Compute an (approximate) solution p_k for the correction equation

$$J_k p = -r_k.$$

- (e) *Expand the search space*. Select a v_k in the $\text{span}(V_{k-1}, p_k)$ that is linearly independent of V_{k-1} and update $V_k = [V_{k-1}, v_k]$.
- (f) *Expand the shadow space*. Select a w_k that is linearly independent of W_{k-1} and update $W_k = [W_{k-1}, w_k]$.
- (g) *Solve the projected problem*. Compute nontrivial solutions y of the projected system

$$W_k^* G_k(y) = 0.$$

- (h) *Update*. Select a y_k (from the set of solutions y) and update the approximation: $u_{k+1} = u_k + V_k y_k$.
- (i) *Restart*. $\ell = \dim(\text{span}(V_k))$. If ℓ is too large, select an $\ell' < \ell$, select $\ell \times \ell'$ matrices R_V and R_W and compute $V_k = V_k R_V$, $W_k = W_k R_W$ (i.e., take suitable combinations of the columns of V_k and W_k).

ALGORITHM 3. Accelerated inexact Newton.

Note that (3.1)–(3.3) represent nonlinear problems in only k variables, which may be much easier to solve than the original problem. If these smaller nonlinear problems can be formulated cheaply, then the costs for an update step may be considered relatively small.

Note also that since (3.1)–(3.3) are nonlinear, they may have more than one solution. This fact may be exploited to steer the computational process to a specific preferable solution of the original problem.

Accelerated inexact Newton. For the Galerkin approach, step 2(e) in Algorithm 1 is replaced by four steps in which

- the search subspace \mathcal{V}_{k-1} is expanded by an approximate “Newton correction” and a suitable basis is constructed for this subspace,
- a shadow space \mathcal{W}_k is selected on which we project the original problem,
- the projected problem (3.1) is solved,
- and the solution is updated.

This is represented by steps 3(e)–3(h) in Algorithm 3. The MR approach and the RMR approach can be represented in a similar way.

4. Computational considerations. In this section we make some comments on implementation details that mainly focus on limiting computational work and memory space.

4.1. Restart. For small k , problems (3.1)–(3.3) are of small dimension and may often be solved at relatively low computational costs (e.g., by some variant of Newton’s method).

For larger k they may become a serious problem in themselves. In such a situation, one may wish to restrict the subspaces \mathcal{V} and \mathcal{W} to subspaces of smaller dimension (see Algorithm 3, step 3(i)). Such an approach limits the computational costs per iteration, but it may also have a negative effect on the speed of convergence.

For example, the simplest choice, restricting the search subspace to a one-dimensional subspace, leads to damped inexact Newton methods, where, for instance, the damping parameter α is the solution of $\min_{\alpha} \|G_k(\alpha)\|_2$, where $G_k(\alpha) = F(u_k + p_k\alpha)$.

Of course, a complete restart is also feasible, say after each m th step (cf. step 3(i) of Algorithm 3):

$$3(i): \quad \text{If } \dim(\text{span}(V_k)) > m, \text{ then } V_k = [] \text{ and } W_k = [].$$

The disadvantage of a complete restart is that we have to rebuild subspace information again. Usually it leads to a slower speed of convergence.

It seems like an open door to suggest that parts of subspaces better be retained at restart, but in practical situations it is very difficult to predict what those parts should be. A meaningful choice would depend on spectral properties of the Jacobian as well as on the current approximation. When solving linear equations with GMRESR [29], good results have been reported in [30] when selecting a number of the first and the last columns (cf. step 3(i) of Algorithm 3); e.g.,

$$3(i): \quad \text{If } \dim(\text{span}(V_k)) > 10, \text{ then } V_k = V_k R_V \text{ and } W_k = W_k R_W, \text{ with} \\ R_V = R_W = [e_1, \dots, e_5, e_7, \dots, e_{11}].$$

In [7], a variant of GMRESR, called GCRO, is proposed, which implements another choice. For subspaces of dimension $l + m$, the first l columns are retained, together with a combination of the last m columns. This combination is taken such that the approximate solution, induced by a MR condition, is the same for the subspaces of both dimensions $l + m$ and $l + 1$. To be more specific, if $u_{k+1} = u_k + V_k y_k$, where y_k solves $\min_y \|F(u_k + V_k y)\|_2$, then V_k is replaced by $[V_{k_l}, V_{k_m} y_{k_m}]$ (denoting $V_k = [V_{k_l}, V_{k_m}]$ and $y_k = [y_{k_l}^T, y_{k_m}^T]^T$).

4.2. Update. In the update step (step 3(h) of Algorithm 3), a solution y_k of the projected problem has to be selected from the set of solutions y . Selection may be necessary since many nonlinear problems have more than one solution. Sometimes this may be the reason for poor convergence of (inexact) Newton: the sequence of approximate solutions “wavers” between different exact solutions. For larger search subspaces, the search subspace may contain good approximations for more than one solution. This may be exploited to steer the sequence of approximate solutions to the wanted solution, and it may help to avoid wavering convergence behavior.

The selection of y_k should be based on additional properties of the solution. For instance, we may look for the solution largest in norm, or, as in the case of eigenvalue problems, for a solution of which one component is close to some specific value. (For instance, if one is interested in eigenvalues close to, say, zero, the Ritz vector with Ritz value closest to zero will be chosen.)

4.3. The projected problem. Even though problems of small dimension can be solved with relatively low computational costs, step 3(g) in Algorithm 3 is not necessarily inexpensive. The projected problem is embedded in the large subspace and it may require quite some computational effort to represent the problem in a small subspace (to which y belongs) of dimension $\ell := \dim(\text{span}(V_k))$ ($y \in \mathbb{C}^\ell$). In the case of linear equations (or linear eigenvalue problems) the computation of an $\ell \times \ell$ -matrix as $W_k^* A V_k$ requires ℓ^2 inner products. For this type of problem, and for many others as well, one may save on the computational costs by re-using information from previous iterations.

4.4. Expanding the search subspace. The AIN algorithm breaks down if the search subspace is not expanded. This happens when p_k belongs to the span of V_{k-1} (or, in finite precision arithmetic, if the angle between p_k and this subspace is very small). As for GMRES, one may then replace p_k by $J_k v_\ell$, where v_ℓ is the last column vector of the matrix V_{k-1} .

With approximate solution of the correction equation, a breakdown will also occur if the new residual r_k is equal to the previous residual r_{k-1} . We will have such a situation if $y_{k-1} = 0$. Then, instead of modifying the expansion process in iteration number k , one may also take measures in iteration number $k-1$ in order to avoid $y_{k-1} = 0$. In [29] a few steps by least square QR (LSQR) are suggested when the linear solver is a Krylov subspace method: $p_k = J_{k-1}^T r_{k-1}$ may already cure the stagnation.

5. How linear solvers fit in the AIN framework. In this section we will show how some well-known iterative methods for the solution of linear systems fit in the AIN framework. The methods that follow from specific choices in AIN are equivalent to well-known methods only in the sense that, at least in exact arithmetic, they produce the same basis vectors for the search spaces, the same approximate solutions, and the same Newton corrections (in the same sense as that in which GMRES and ORTHODIR are equivalent).

With $u = x$, $F(u) := b - Au$, the linear equation (1.2) is equivalent to the one in (1.1) and $J_k = -A$. In this section, M denotes a preconditioning matrix for A (i.e., for a vector v , $M^{-1}v$ is easy to compute and approximates $A^{-1}v$).

5.1. GCR. With the choice, $p_k = M^{-1}r_k$, $w_k = Av_k$, and v_k such that $w_k \perp \text{span}(W_{k-1})$, the AIN algorithm, Algorithm 3 (without restart), is equivalent to preconditioned GCR [31].

5.2. FOM and GMRES. The choice $p_k = M^{-1}r_k$ and v_k such that $v_k \perp \text{span}(V_{k-1})$ gives algorithms that are related to FOM and GMRES [26]. With the additional choice $w_k = v_k$, Algorithm 3 is just FOM, while the choice $w_k = Av_k$ gives an algorithm that is equivalent to GMRES.

5.3. GMRESR. Taking $w_k = Av_k$ and v_k such that w_k is perpendicular to $\text{span}(W_{k-1})$ as in GCR, and taking p_k as an approximate solution of the equation $Ap = r_k$, Algorithm 3 is equivalent to the GMRESR algorithms [29]. One might compute p_k by a few steps of GMRES, for instance.

6. AIN schemes for mildly nonlinear problems. In this section we will discuss numerical methods for the iterative solution of the generalized eigenproblem (1.3). We will show that they also fit in the general framework of the AIN Algorithm 3 methods.

As already mentioned, these AIN methods consist of two parts. In one part an approximate solution of the correction equation (cf. step 3(d) in Algorithm 3) is used

to extend the search space. In the other part a solution of the projected problem (cf. step 3(g) in Algorithm 3) is used to construct an update for the approximate solution.

We will start with the derivation of a more suitable form for the (Newton) correction equation for the generalized eigenproblem. After that, we will make some comments on how to solve the projected problem.

The correction equation. In order to avoid some of the complications that go with complex differentiation, we will mainly focus on the numerical computation of eigenvectors with a fixed component in some given direction (rather than on the computation of eigenvectors with a fixed norm).

First, let \tilde{u} be a fixed vector with a nontrivial component in the direction of the desired eigenvector x . We want to compute approximations u_k for x with a normalized component in the \tilde{u} -direction: $(u, \tilde{u}) = (x, \tilde{u}) = 1$. We will select ϑ_k such that the residual $r_k := (A - \vartheta_k B)u_k$ is orthogonal to w , where w is another fixed nontrivial vector; i.e., the approximate eigenvalue ϑ_k is given by $\vartheta_k := w^* A u_k / w^* B u_k$.

Consider the map F given by

$$F(u) := (A - \vartheta B)u, \text{ where } \vartheta := \frac{w^* A u}{w^* B u},$$

and u belongs to the hyperplane $\{y \in \mathbb{C}^n \mid (y, \tilde{u}) = 1\}$. The Jacobian $J_k = F'(u_k)$ is then given by

$$J_k = \left(I - \frac{B u_k w^*}{w^* B u_k} \right) (A - \vartheta_k B)|_{\tilde{u}^\perp},$$

and the correction equation reads as

$$(6.1) \quad p \perp \tilde{u}, \text{ such that } \left(I - \frac{B u_k w^*}{w^* B u_k} \right) (A - \vartheta_k B)p = -r_k.$$

Since $r_k \perp w$, this equation is equivalent to

$$(6.2) \quad \begin{bmatrix} A - \vartheta_k B & B u_k \\ \tilde{u}^* & 0 \end{bmatrix} \begin{bmatrix} p \\ \varepsilon \end{bmatrix} = \begin{bmatrix} -r_k \\ 0 \end{bmatrix};$$

that is, p is the solution of (6.1) if and only if p is the solution of (6.2).

The projected problem. For the generalized eigenvalue problem we are in the fortunate position that all the solutions of problems of moderate size can be computed by standard methods such as, for instance, the QZ [19] method. However, before we can apply these methods we have to reformulate the projected problem because of the exceptional position of u_k in $W^* F(u_k + V_k y)$.

The key to this reformulation is the observation that in the methods we consider the affine subspace $u_k + \text{span}(V_k)$ is equal to V_k because V_k contains u_k by itself. Now, as an alternative to step 3(g) in Algorithm 3, we may also compute all the solutions y of

$$W_k^* F(V_k y) = 0.$$

This problem can now be solved by, for instance, the QZ method, and after selecting y_k , a new approximation u_{k+1} is given by $u_{k+1} := V_k y_k$.

6.1. Arnoldi's method. We consider the simplified case where $B = I$, i.e., the standard eigenproblem. If we do only one step of a Krylov subspace method (Krylov dimension 1) for the solution of the correction equation (6.1), then we obtain for the correction p_k

$$p_k = -\alpha r_k.$$

Hence, $p_k = -\alpha(A - \vartheta_k I)u_k$. Note that this may be a (very) poor approximation because, in general, $r_k \not\perp \tilde{u}$. The approximate eigenvector u_k belongs to the search subspace $\text{span}(V_{k-1})$, and expanding the search subspace by the component of p_k orthogonal to $\text{span}(V_{k-1})$ is equivalent to expanding this space with the orthogonal component of Au_k , which would be the “expansion” vector in Arnoldi's method. Hence, the search subspace is precisely the Krylov subspace generated by A and u_0 . Apparently, Arnoldi's method is an AIN method (with a “very inexact Newton step”) without restart.

The choice $W_k = V_k$ corresponds to the standard one in Arnoldi and produces ϑ 's that are called Ritz values, while the choice $W_k = AV_k$ leads to harmonic Ritz values [22].

6.2. Davidson's method. As in Arnoldi's method, Davidson's method [6] also carries out only one step of a Krylov subspace method for the solution of the correction equation. However, in contrast to Arnoldi's method, Davidson also incorporates a preconditioner.

He suggests solving (6.1) approximately by p_k with

$$p_k = -M^{-1}r_k,$$

where M is the inverse of the diagonal of $A - \vartheta_k B$. Other choices have been suggested as well (cf., e.g., [5, 21]). Because of the preconditioner, even if $B = I$, the search space is not simply the Krylov subspace generated by A and u_0 . This may lead to an advantage of Davidson's method over Arnoldi's method.

For none of the choices of the preconditioner has proper care been taken of the projections (see (6.1)): the preconditioner should approximate the inverse of the projected matrix (see (6.1)) as a map from \tilde{u}^\perp onto w^\perp rather than of $A - \vartheta_k B$.

However, if M is the diagonal of $A - \vartheta B$, and we choose \tilde{u} and w equal to the same arbitrary standard basis vector (as Davidson does [6]), then

$$\left(I - \frac{Bu_k w^*}{w^* Bu_k}\right) M \left(I - \frac{u_k w^*}{w^* u_k}\right) p = -r_k$$

whenever $w \perp r_k$ and p solves $Mp = -r_k$. Note that $p \perp w$, because M is diagonal and $r_k \perp w$. Therefore, for this particular choice of w (and \tilde{u}), the diagonal M may be expected to be a good preconditioner for the correction equation (including the projections) in the cases where M is a good preconditioner for $A - \vartheta_k B$. Observe that this argument does not hold for nondiagonal preconditioners M .

6.3. Jacobi–Davidson. Davidson methods with a nondiagonal preconditioner do not take proper care of the projections in the correction equation (6.1). This observation was made in [28], and a new algorithm was proposed for eigenproblems by including the projections in the Davidson scheme. In addition, these modified schemes allow for more general approximate solutions p_k than $p_k = -M^{-1}r_k$. For instance, the use of ℓ steps of a preconditioned Krylov subspace method for the correction

equation is suggested, leading to Arnoldi-type methods in which the variable polynomial preconditioning is determined efficiently and the projections have been included correctly. The new methods have been called *Jacobi–Davidson* methods (Jacobi took proper care of the projections, but did not build a search subspace as Davidson did (see [28] for details and further references)).

The analysis and results in [3, 27] show that these Jacobi–Davidson methods can also be effective for solving generalized eigenproblems, even without any matrix inversion.

The Jacobi–Davidson methods allow for a variety of choices that may improve efficiency of the steps and speed of convergence and are good examples of AIN methods in which the projected problem (3.1) is used to steer the computation.

For an extensive discussion, we refer to [27].

7. AIN schemes for general nonlinear problems. In this section we summarize some iterative methods for the solution of nonlinear problems that have been proposed by different authors, and we show how these methods fit in the AIN framework.

Brown and Saad [4] describe a family of methods for solving nonlinear problems. They refer to these methods as *nonlinear Krylov subspace projection methods*. Their modifications to Newton’s method are intended to enhance robustness and are heavily influenced by ideas presented in [9]. One of their methods is a variant of damped inexact Newton, in which they approximate the solution of the correction equation by a few steps of Arnoldi or GMRES and determine the damping parameter α by a “linesearch backtracking technique.” So this is just another AIN scheme, with a special one-dimensional subspace acceleration. They also propose a model trust region approach, where they take their update to the approximation from the Krylov subspace \tilde{V}_m generated by m steps of (preconditioned) Arnoldi or GMRES as $p_k = \tilde{V}_m \tilde{y}_k$, where \tilde{y}_k is the point on the dogleg curve for which $\|\tilde{y}_k\|_2 = \tau$, the trust region size: \tilde{y}_k is an approximation for $\min_y \|F(u_k + \tilde{V}_m y)\|_2$. This could be considered as a block version of the previous method.

In [2] Axelsson and Chronopoulos propose two nonlinear versions of a (truncated) GCR type of method. Both methods fit in the AIN framework. The first method, nonlinear generalized conjugate gradient (NGCG), is a MR AIN method with $p_k = -r_k$ and V_k orthonormal; in other words the correction equation is not solved. The second method, Newton direction NGCG (NNGCG), differs from NGCG in that p_k is now computed as an approximate solution (by some method) of the correction equation (2.1), where the accuracy is such that $\|F(u_k) - F'(u_k)p_k\| \leq \rho_k \|F(u_k)\|$ for some nonincreasing sequence (ρ_j) , $0 \leq \rho_j < 1$ (see also [8]). So the method NNGCG is a MR AIN method. It can be viewed as a generalization of GMRESR [29]. Under certain conditions on the map F they prove global convergence.

In [15], Kaporin and Axelsson propose a class of nonlinear equation solvers, Gauss–Newton (direction/augmented) Krylov subspace (GNKS), in which the ideas presented in [4] and [2] are combined. There, the direction vectors p_k are obtained as linear combinations of the columns of \tilde{V}_m and V_k . To be more precise, $p_k = [\tilde{V}_m, V_k]y_k$, where y_k solves $\min_y \|F(u_k + [\tilde{V}_m, V_k]y)\|_2$. This problem is then solved by a special Gauss–Newton iteration scheme, which avoids excessive computational work by taking into account the acute angle between r_k and $J_k p_k$, and the rate of convergence. The method generalizes GCRO [7].

8. Numerical experiments. In this section we test several AIN schemes and present results of numerical experiments on three different nonlinear problems. For tests and test results with methods for linear problems and eigenproblems we refer to their references. The purpose of this presentation is to show that acceleration may also be useful in the nonlinear case. By *useful* we mean that additional computational cost is compensated for by faster convergence.

Different AIN schemes distinguish themselves by the way they (approximately) solve the correction equation and the projected problem (cf. sections 3.2 and 7). Out of the overwhelming variety of choices we have selected a few possible combinations, some of which lead to AIN schemes that are equivalent to already proposed methods and some of which lead to new methods. We compare the following (existing) MR AIN schemes:

- LINESEARCH, the backtracking linesearch technique [4, p. 458];
- DOGLEG, the model trust region approach as proposed in [4, p. 462];
- NNGCG, a variant of the method proposed in [2], solving (3.2) by the Levenberg–Marquardt algorithm [20];
- GNKS, the method proposed in [15];
- and the (new) RMR AIN schemes:
- RMR A, choosing $W_k = V_k$; and
- RMR B, choosing $W_k = [W_{k-1}, w_k]$, where w_k is the component of $J_k p_k$ orthogonal to W_{k-1} .

For these last two schemes, the minimization problem (3.3) was solved by the Gauss–Newton variant described in [15]. The necessary subspaces for the direction p_k or the projected problem were obtained by 10 steps of GMRES, or (in the third example) also by at most 50 iterations of the generalized conjugate gradient squared (CGS) variant, CGS2 [10].

In all cases the exact Jacobian was used. Furthermore, we used orthonormal matrices V_k and W_k , obtained from a modified Gram–Schmidt process and restricted to the last 10 columns, in an attempt to save computational work. The computations were done on a Sun Sparc 20 in double precision and the iterations were stopped when $\|r_k\|_2 \leq 10^{-6}$. A method *failed* either when the convergence was too slow, i.e., when $|\|r_{k+1}\|_2 - \|r_k\|_2| < 10^{-6} \|r_{k+1}\|_2$, or when the number of nonlinear iterations (per step) exceeded 200.

Since the computational cost of the methods is approximately proportional to the costs of the number of function evaluations and matrix multiplications, the following counters are given in the tables:

- NI, the number of nonlinear iterations;
- FE, the number of function evaluations;
- MV, the number of multiplications by the Jacobian;
- PRE, the number of applications of the preconditioner;
- TOTAL, the sum of FE, MV, and PRE.

8.1. A one-dimensional Burgers’s equation. As a first test problem we consider the following one-dimensional Burgers’s equation [14]:

$$(8.1) \quad \begin{aligned} \frac{\partial u}{\partial t} + \sin(2u) \frac{\partial u}{\partial x} &= \mu \frac{\partial^2 u}{\partial x^2}, \\ u(x, 0) &= g(x), \quad x \in \Omega, \\ u(x, t) &= \psi(x, t), \quad x \in \partial\Omega, \end{aligned}$$

where $\Omega = [0, 1]$.

TABLE 8.1
Results for Burgers's equation.

Method	NI	FE	MV	TOTAL
LINESEARCH	594	1818	4853	6671
DOGLEG	644	3559	6140	9699
NNGCG	229	4426	11769	16195
GNKS	187	852	7067	7919
RMR A	230	926	4471	5397
RMR B	236	856	4606	5462




FIG. 8.1. Solution of Burgers's equation.

We discretized the spatial variable x with finite differences with 64 grid points and for the time derivative we used

$$\theta \dot{u}_{n+1} + (1 - \theta) \dot{u}_n = \frac{u_{n+1} - u_n}{\Delta t},$$

with $\theta = \frac{2}{3}$ and $\Delta t = 10^{-2}$. u_n denotes the solution at time $t_n = t_0 + n\Delta t$. For this test the solution u_n was computed for $n = 1, 2, \dots, 30$, and as an initial guess to u_{n+1} , we took u_n . No preconditioning was used.

In Table 8.1 we show the results for problem (8.1) with

$$g(x) = \pi - 2\pi x, \quad \psi(x, t) = g(x), \quad \text{and } \mu = 10^{-2}.$$

A plot of the solutions u_1, u_2, \dots, u_{30} is given in Fig. 8.1. The table shows the cumulative value of the counters for each method after completing the computation of u_{30} .

If we look at the number of nonlinear iterations (NI), we see that acceleration indeed reduces this number. However, in the case of GNKS, this does not result in less work because the number of matrix multiplications (MV) increases too much. Here both the Galerkin approaches RMR A and RMR B are less expensive than all the other methods, with RMR A being the winner.

8.2. The Bratu problem. As a second test problem we consider the Bratu problem [12, 4]. We seek a solution (u, λ) of the nonlinear boundary value problem:

$$(8.2) \quad -\Delta u - \lambda e^u = 0 \text{ in } \Omega, \quad u = 0 \text{ on } \partial\Omega.$$

TABLE 8.2

Results for the Bratu problem, solved by the arc length continuation method.

Method	NI	FE	MV	PRE	TOTAL
LINESEARCH	391	1013	3732	3421	8166
DOGLEG	381	2664	3010	3010	8684
NNGCG	361	1297	4243	3091	8631
GNKS	358	1056	6896	2780	10732
RMR A	389	539	4005	3399	7943
RMR B	361	414	3806	3091	7311

TABLE 8.3

Single solve of the Bratu problem, $u_0 = 0, \lambda = 6.8$.

Method	NI	FE	MV	PRE	TOTAL
LINESEARCH	29	85	336	308	729
DOGLEG	27	151	494	260	905
NNGCG	9	49	196	88	333
GNKS	38	119	1806	370	2295
RMR A	6	13	77	55	145
RMR B	6	12	79	55	146

For Ω we took the unit square and we discretized with finite differences on a 31×31 regular grid. It is known (cf. [12]) that there exists a critical value λ^* such that for $0 < \lambda < \lambda^*$, problem (8.2) has two solutions and for $\lambda > \lambda^*$ problem (8.2) has no solutions. In order to locate this critical value we use the arc length continuation method as described in [12, sections 2.3 and 2.4]. Problem (8.2) is replaced by a problem of the form

$$\begin{cases} F(u_s, \lambda(s)) = 0, \\ \ell(u_s, \lambda(s), s) = 0, \end{cases}$$

where ℓ , a scalar-valued function, is chosen such that s is some arc length on the solution branch and u_s is the solution of (8.2) for $\lambda = \lambda(s)$. We preconditioned GMRES by ILU(0) [18] of the discretized Laplace operator Δ .

Table 8.2 shows the results after a full continuation run: starting from the smallest solution (u, λ) with $\lambda = 1$ the solution branch is followed along the (discretized) arc with $s_n = s_0 + n\Delta s$ for step size $\Delta s = 1$ and $n = 1, 2, \dots, 80$. Again we see that acceleration may be useful, in spite of the fact that there is little room for it because, on average, approximately only 4.5 Newton iterations were necessary to compute the solution per continuation step. In this example RMR B performs better than RMR A.

Table 8.3 shows the results for the case where we solve (8.2) for fixed $\lambda = 6.8$ (near the critical value). In this case Galerkin acceleration is even more useful and the differences are more pronounced.

The sup norm of the solution for the different values of λ are plotted in Fig. 8.2. The two solutions at $\lambda \approx 4$ along the diagonal of the unit square are shown in Fig. 8.3.

8.3. The driven cavity problem. In this section we present test results for the classical *driven cavity* problem from incompressible fluid flow. We follow closely the presentations in [12, 4]. In stream function-vorticity formulation the equations




FIG. 8.2. *Sup norms of the solution u along the arc.*




FIG. 8.3. *Solutions u at $\lambda \approx 4$ along the diagonal of the domain.*

are

$$\begin{aligned} \nu \Delta \omega + (\psi_{x_2} \omega_{x_1} - \psi_{x_1} \omega_{x_2}) &= 0 \text{ in } \Omega, \\ -\Delta \psi &= \omega \text{ in } \Omega, \\ \psi &= 0 \text{ on } \partial\Omega, \\ \frac{\partial \psi}{\partial n}(x_1, x_2)|_{\partial\Omega} &= \begin{cases} 1 & \text{if } x_2 = 1, \\ 0 & \text{if } 0 \leq x_2 < 1, \end{cases} \end{aligned}$$

where Ω is the unit square and the viscosity ν is the reciprocal of the Reynolds number Re . In terms of ψ alone this can be written as

$$\nu \Delta^2 \psi + (\psi_{x_2} (\Delta \psi)_{x_1} - \psi_{x_1} (\Delta \psi)_{x_2}) = 0 \text{ in } \Omega,$$

subject to the same boundary conditions. This equation was discretized with finite differences on a 25×25 grid; see Fig. 8.4. The grid lines are distributed as the roots of the Chebyshev polynomial of degree 25. As preconditioner we used the modified ILU(2) [13] decomposition of the biharmonic operator Δ^2 . Starting from the solution for $\text{Re} = 0$, we computed several solutions, using the arc length continuation method (cf. the previous example and [12]) with step sizes $\Delta s = 100$ for $0 \leq \text{Re} \leq 1400$, $\Delta s = 200$ for $1400 < \text{Re} \leq 2600$, and $\Delta s = 400$ for $\text{Re} = 3000$.

Table 8.4 shows the results of this test when using 10 steps of GMRES and CGS2 [10] for the correction equation. In the case of CGS2 we approximately solved the correction equation to a relative residual norm precision of 2^{-k} , where k is the current Newton step [8], with a maximum of 50 steps. Clearly, the methods using (the basis produced by) 10 steps of GMRES perform very poorly for this example. Only GNKS is able to complete the full continuation run, but requires a large number of Newton steps. If we look at the results for the AIN schemes for which CGS2 is used, we see that, except for the LINESEARCH method, these methods perform much better. The RMR methods are again the most efficient ones.

This test also reveals a possible practical drawback of methods like DOGLEG and GNKS. These methods exploit an affine subspace to find a suitable update for the approximation. This may fail when the problem is hard or when the preconditioner is not good enough. In that case the dimension of the affine subspace must be large, which may be, because of storage requirements and computational overhead, not

TABLE 8.4

Results for the driven cavity problem, solved by the arc length continuation method for $Re = 100, \dots, 3000$.

GMRES					
Method	NI	E	MV	PRE	TOTAL
LINESEARCH	fails at $Re = 400$ after a TOTAL of 545				
DOGLEG	fails at $Re = 100$ after a TOTAL of 113				
NNGCG	fails at $Re = 2200$ after a TOTAL of 19875				
GNKS	641	2315	30206	6210	38731
RMR A	fails at $Re = 2000$ after a TOTAL of 13078				
RMR B	fails at $Re = 800$ after a TOTAL of 7728				
CGS2					
Method	NI	FE	MV	PRE	TOTAL
LINESEARCH	fails at $Re = 1300$ after a TOTAL of 4342				
NNGCG	141	555	7003	6119	13677
RMR A	137	297	6266	5969	12532
RMR B	137	167	6277	5937	12381




FIG. 8.4. Grid for the driven cavity problem, 25×25 .




FIG. 8.5. Stream lines of the driven cavity problem, $Re = 100$.

feasible. For the schemes that use approximate solutions of the correction equation, delivered by some arbitrary, iterative method, e.g., CGS2, one can easily adapt the precision, which leaves more freedom.

Plots of the stream lines for the values

$$\psi = -0.12, -0.1, -0.08, -0.06, -0.04, -0.02, \\ 0.0, 0.0025, 0.001, 0.0005, 0.0001, 0.00005$$

(cf. [12]) are given in Figs. 8.5–8.9. The plots show virtually the same solutions as in [12].

9. Conclusions. We have shown how the classical Newton iteration scheme for nonlinear problems can be accelerated in a similar way as standard Richardson-type iteration schemes for linear equations. This leads to the AIN framework in which many well-known iterative methods for linear problems, eigenproblems, and general nonlinear problems fit. From this framework an overwhelming number of possible




FIG. 8.6. Stream lines of the driven cavity problem, $Re = 400$.




FIG. 8.7. Stream lines of the driven cavity problem, $Re = 1600$.




FIG. 8.8. Stream lines of the driven cavity problem, $Re = 2000$.




FIG. 8.9. Stream lines of the driven cavity problem, $Re = 3000$.

iteration schemes can be formulated. We have selected a few and shown by numerical experiments that the RMR methods, especially, can be very useful for further reducing computational costs.

REFERENCES

- [1] W. E. ARNOLDI, *The principle of minimized iterations in the solution of the matrix eigenvalue problem*, Quart. Appl. Math., 9 (1951), pp. 17–29.
- [2] O. AXELSSON AND A. T. CHRONOPOULOS, *On nonlinear generalized conjugate gradient methods*, Numer. Math., 69 (1994), pp. 1–16.
- [3] J. G. L. BOOTEN, H. A. VAN DER VORST, P. M. MEIJER, AND H. J. J. TE RIELE, *A Preconditioned Jacobi-Davidson Method for Solving Large Generalized Eigenvalue Problems*, Report NM-R9414, Dept. Num. Math., CWI, Amsterdam, 1994.
- [4] P. N. BROWN AND Y. SAAD, *Hybrid Krylov methods for nonlinear systems of equations*, SIAM J. Sci. Statist. Comput., 11 (1990), pp. 450–481.

- [5] M. CROUZEIX, B. PHILIPPE, AND M. SADKANE, *The Davidson method*, SIAM J. Sci. Comput., 15 (1994), pp. 62–76.
- [6] E. R. DAVIDSON, *The iterative calculation of a few of the lowest eigenvalues and corresponding eigenvectors of large real symmetric matrices*, J. Comput. Phys., 17 (1975), pp. 87–94.
- [7] E. DE STURLER AND D. R. FOKKEMA, *Nested Krylov methods and preserving the orthogonality*, in Sixth Copper Mountain Conference on Multigrid Methods, N. D. Melson, T. A. Manteuffel, and S. F. McCormick, eds., NASA Conference Publication 3324, Part 1, NASA Langley Research Center, Hampton, VA, 1993, pp. 111–126.
- [8] R. S. DEMBO, S. C. EISENSTAT, AND T. STEIHAUG, *Inexact Newton methods*, SIAM J. Numer. Anal., 19 (1982), pp. 400–408.
- [9] J. E. DENNIS, JR. AND R. B. SCHNABEL, *Numerical Methods for Unconstrained Optimization and Nonlinear Equations*, Prentice–Hall, Englewood Cliffs, NJ, 1983.
- [10] D. R. FOKKEMA, G. L. G. SLEIJPEN, AND H. A. VAN DER VORST, *Generalized conjugate gradient squared*, J. Comput. Appl. Math., 71 (1996), pp. 125–146.
- [11] R. W. FREUND AND N. M. NACHTIGAL, *QMR: A quasi minimal residual method for non-Hermitian linear systems*, Numer. Math., 60 (1991), pp. 315–339.
- [12] R. GLOWINSKI, H. B. KELLER, AND L. REINHART, *Continuation-conjugate gradient methods for the least squares solution of nonlinear boundary value problems*, SIAM J. Sci. Statist. Comput., 6 (1985), pp. 793–832.
- [13] I. GUSTAFSSON, *A class of first order factorizations methods*, BIT, 18 (1978), pp. 142–156.
- [14] E. HOPF, *The partial differential equation $u_t + uu_x = \mu u_{xx}$* , Comm. Pure Appl. Math., 3 (1950), pp. 201–230.
- [15] I. E. KAPORIN AND O. AXELSSON, *On a class of nonlinear equation solvers based on the residual norm reduction over a sequence of affine subspaces*, SIAM J. Sci. Comput., 16 (1995), pp. 228–249.
- [16] T. KERKHOVEN AND Y. SAAD, *Acceleration Techniques for Decoupling Algorithms in Semiconductor Simulation*, Report UIUCDCS-R-87-1363, Dept. of Comp. Sci., University of Illinois, Urbana, IL, 1987.
- [17] C. LANCZOS, *Solution of systems of linear equations by minimized iteration*, J. Res. Nat. Bur. Standards, 49 (1952), pp. 33–53.
- [18] J. A. MEIJERINK AND H. A. VAN DER VORST, *An iterative solution method for linear systems of which the coefficient matrix is a symmetric M-matrix*, Math. Comp., 31 (1977), pp. 148–162.
- [19] C. B. MOLER AND G. W. STEWART, *An algorithm for generalized matrix eigenvalue problems*, SIAM J. Numer. Anal., 10 (1973), pp. 241–256.
- [20] J. J. MORÉ, *The Levenberg-Marquardt algorithm: Implementation and theory*, in Numerical Analysis, G. A. Watson, ed., Lecture Notes in Mathematics 630, Springer-Verlag, Berlin, Heidelberg, New York, 1977, pp. 105–116.
- [21] R. B. MORGAN, *Generalizations of Davidson’s method for computing eigenvalues of large non-symmetric matrices*, J. Comput. Phys., 101 (1992), pp. 287–291.
- [22] C. C. PAIGE, B. N. PARLETT, AND H. A. VAN DER VORST, *Approximate solutions and eigenvalue bounds from Krylov subspaces*, Linear Algebra Appl., 2 (1995), pp. 115–133.
- [23] Y. SAAD, *Krylov subspace method for solving large unsymmetric linear systems*, Math. Comp., 37 (1981), pp. 105–126.
- [24] Y. SAAD, *Numerical Methods for Large Eigenvalue Problems*, Manchester University Press, Manchester, UK, 1992.
- [25] Y. SAAD, *A flexible inner-outer preconditioned GMRES algorithm*, SIAM J. Sci. Comput., 14 (1993), pp. 461–469.
- [26] Y. SAAD AND M. H. SCHULTZ, *GMRES: A generalized minimum residual algorithm for solving nonsymmetric linear systems*, SIAM J. Sci. Statist. Comput., 7 (1986), pp. 856–869.
- [27] G. L. G. SLEIJPEN, J. G. L. BOOTEN, D. R. FOKKEMA, AND H. A. VAN DER VORST, *Jacobi-Davidson Type Methods for Generalized Eigenproblems and Polynomial Eigenproblems*, BIT, 36 (1996), pp. 595–633.
- [28] G. L. G. SLEIJPEN AND H. A. VAN DER VORST, *A Jacobi-Davidson iteration method for linear eigenvalue problems*, SIAM J. Matrix Anal. Appl., 17 (1996), pp. 401–425.
- [29] H. A. VAN DER VORST AND C. VUIK, *GMRESR: A family of nested GMRES methods*, Linear Algebra Appl., 1 (1994), pp. 369–386.
- [30] C. VUIK, *Further Experiences with GMRESR*, Tech. report 92-12, Faculty of Technical Mathematics and Informatics, Delft University of Technology, Delft, the Netherlands, 1992.
- [31] D. M. YOUNG AND K. C. JEA, *Generalized conjugate-gradient acceleration of nonsymmetrizable iterative methods*, Linear Algebra Appl., 34 (1980), pp. 159–194.