

Efficient Solution of $Ax^{(k)} = b^{(k)}$ Using A^{-1}

Adi Ditkowski,^{1,2} Gadi Fibich,¹ and Nir Gavish¹

Received August 20, 2006; accepted (in revised form) September 17, 2006

In this work, we consider the problem of solving $Ax^{(k)} = b^{(k)}$, $k = 1, \dots, K$, where $b^{(k+1)} = f(x^{(k)})$. We show that when A is a full $n \times n$ matrix and $K \geq cn$, where $c \ll 1$ depends on the specific software and hardware setup, it is faster to solve $Ax^{(k)} = b^{(k)}$ for $k = 1, \dots, K$ by explicitly evaluating the inverse matrix A^{-1} rather than through the LU decomposition of A . We also show that the forward error is comparable in both methods, regardless of the condition number of A .

KEY WORDS: Matrix inversion; linear systems.

1. INTRODUCTION

One of the first things we learn in a basic numerical linear algebra course is that in order to solve the linear system $Ax = b$, we should not calculate the inverse matrix A^{-1} . For example, we quote from Matlab's user guide:

In practice, it is seldom necessary to form the explicit inverse of a matrix. A frequent misuse of inverse arises when solving the system of linear equations. One way to solve this is with $x = \text{inv}(A) \cdot b$. A better way, from both an execution time and numerical accuracy standpoint, is to use the matrix division operator $x = A \backslash b$. This produces the solution using Gaussian elimination, without forming the inverse.

Similar statements appear in the classical numerical analysis textbooks. For example, Golub and van Loan [1, p. 121] say:

... As a final example we show how to avoid the pitfall of explicit inverse computation... when a matrix inverse is encountered in a formula, we must think in terms of solving equations rather than in terms of explicit inverse formation.

Similarly, Conte and de Boor [2, p. 166] say:

¹Department of Applied Mathematics, School of Mathematical Sciences, Tel-Aviv University, Tel-Aviv, Israel.

²To whom correspondence should be addressed. E-mail: adid@post.tau.ac.il

...given this simple prescription for calculating the inverse of a matrix, we hasten to point out that there is usually no good reason for ever calculating the inverse. ... whenever A^{-1} is needed merely to calculate a vector $A^{-1}b$ (as in solving $Ax = b$) or a matrix product $A^{-1}B$, A^{-1} should never be calculated explicitly.

Higham [3, p. 262] also states:

Not only is the inversion approach three times more expensive, but it is much less stable. ...we see that matrix inversion is likely to give a much larger residual than Gaussian elimination with partial pivoting...

Thus, the reasons for preferring Gaussian elimination over explicit calculation of A^{-1} to solve $Ax = b$ are performance, and accuracy and stability.

In this study we consider the case where we want to solve the equations

$$Ax^{(k)} = b^{(k)}, \quad k = 1, \dots, K, \quad (1.1)$$

when

1. A is a full $n \times n$ matrix.
2. Equations (1.1) have to be solved sequentially, e.g., when $b^{(k+1)} = f(x^{(k)})$.

The standard, LU-based approach involves a preprocessing stage in which the LU decomposition of the matrix A is calculated:

$$[L, U] = \text{lu}(A). \quad (1.2)$$

Then, for each right-hand-side the linear system is solved using a forward substitution and a backward substitution:

$$y^{(k)} = L \setminus b^{(k)}, \quad x^{(k)} = U \setminus y^{(k)}, \quad k = 1, \dots, K. \quad (1.3)$$

A second, A^{-1} -based approach for solving problem (1.1) is to calculate the inverse of A in a preprocessing stage:

$$\text{inv}A = \text{inv}(A). \quad (1.4)$$

Then, for each right-hand-side the linear system is solved by a single matrix-vector multiplication (Fig. 1):

$$x^{(k)} = \text{inv}A \cdot b^{(k)}, \quad k = 1, \dots, K. \quad (1.5)$$

In this study we show that it is better to solve Eq. (1.1) using A^{-1} rather than through LU decomposition. In Sect. 2 we consider the performance issue. We first show that the two methods require the same number of arithmetic operations. However, the actual performance has a lot

Efficient Solution of $Ax^{(k)} = b^{(k)}$ Using A^{-1}




Fig. 1. Average CPU time for solving $Ax = b^{(k)}$ using the backward and forward substitution (Eq. (1.3), \times) and through multiplication by A^{-1} (Eq. (1.5), \square).

to do with the way that optimized computer codes, such as Intel's MKL Basic Linear Algebra Support (BLAS) package, handle computer architecture issues, such as caching and memory allocation. Indeed, because the data structure of a square matrix (A^{-1}) is inherently simpler than that of two triangular matrices (L and U), in our numerical tests we observe that it is faster to solve these equations using A^{-1} than through the LU decomposition of A . In Sect. 2.4 we show that there are more efficient implementations of the LU algorithm than the one given in Eqs. (1.2) and (1.3). Even with these improvements, however, the A^{-1} algorithm is still more efficient.

As noted, the second argument against using A^{-1} is accuracy and stability. This argument is based on analysis which showed that the backward error in the A^{-1} -based method is larger than in the LU-based method, and that the ratio between the two increases with the condition number of A . While there have been no such analytical results for the forward error, it was implicitly assumed that it would behave similarly. However, in Sec. 3, our simulations with randomly-chosen matrices, as well as with three families of ill-conditioned matrices, suggest that this implicit assumption is wrong. Indeed, we observe that the *forward error is of the same magnitude for both the A^{-1} -based method and the LU-based method, regardless of the condition number of A .*

In Sec. 4 we list various applications that are expected to run faster when implemented with the A^{-1} -based method (quasi-Newton iterations, inverse-power method, and spectral differentiation methods). We conclude by showing that the inverse power method for finding the smallest eigenvalue of a 100×100 matrix can be 11 times faster with the A^{-1} -based method than with the LU-based method, and just as accurate.

Our conclusion that using A^{-1} is superior to using LU decomposition applies to problems in which:

1. A is full, since when A is sparse then L and U are sparse but A^{-1} is full.
2. Equations (1.1) have to be solved sequentially, hence only level 2 BLAS optimization (matrix-vector operations) can be used. Indeed, when all the $b^{(k)}$ are known in advance, all the right-hand-sides can be simultaneously solved with a block LU algorithm, i.e., replacing (1.3) with

$$Y = L \setminus B, \quad X = U \setminus Y,$$

where

$$B = [b^{(1)} \dots b^{(K)}], \quad X = [x^{(1)} \dots x^{(K)}], \quad Y = [y^{(1)} \dots y^{(K)}].$$

In this case, the block LU algorithm is more efficient than using $X = A^{-1} \cdot B$, since one can exploit level 3 BLAS optimization (matrix-matrix operations).

2. PERFORMANCE

2.1. Arithmetic Operations Count

We begin with the standard count of arithmetic operations:

Lemma 2.1. The overall number of arithmetic operations needed to solve the system (1.1) in the LU-based method is

$$N_{lu} = \underbrace{\frac{2}{3}n^3 + \mathcal{O}(n^2)}_{\text{eq. (1.2)}} + \underbrace{2Kn^2 + \mathcal{O}(Kn)}_{\text{eq. (1.3)}}$$

and in the A^{-1} -based method is

$$N_{A^{-1}} = \underbrace{\frac{8}{3}n^3 + \mathcal{O}(n^2)}_{\text{eq. (1.4)}} + \underbrace{2Kn^2 + \mathcal{O}(Kn)}_{\text{eq. (1.5)}}.$$

Proof. For the LU-based method the preprocessing stage involves finding the LU decomposition of A , which requires $\frac{2}{3}n^3 + \mathcal{O}(n^2)$ arithmetic operations (see [4, p. 152]). Then, solving each linear system involves

Efficient Solution of $Ax^{(k)} = b^{(k)}$ Using A^{-1}

a backward and forward substitution. Each substitution requires $n^2 + \mathcal{O}(n)$ arithmetic operations. Therefore, an overall of $2Kn^2 + \mathcal{O}(Kn)$ arithmetic operations are needed to solve (1.3).

Similarly, for the A^{-1} -based method the preprocessing stage involves finding A^{-1} , which require $\frac{8}{3}n^3 + \mathcal{O}(n^2)$ arithmetic operations (see [4, p. 161]). Then, solving each linear system requires multiplying $b^{(k)}$ by A^{-1} , which requires $2n^2 + \mathcal{O}(n)$ arithmetic operations. Therefore, $2Kn^2 + \mathcal{O}(Kn)$ arithmetic operations are needed to solve (1.5). \square

Observation 2.1 shows that the number of arithmetic operations for the LU-based method is smaller than for the A^{-1} -based method for $K \leq n$, and is the same for $K \gg n$. More precisely,

$$\begin{aligned} N_{lu} &\approx \frac{1}{4}N_{A^{-1}}, & K \ll n \\ N_{lu} &= \frac{2+6\lambda}{8+6\lambda}N_{A^{-1}} & K = \lambda n, \quad \lambda = \mathcal{O}(1) \\ N_{lu} &\approx N_{A^{-1}} & K \gg n. \end{aligned} \tag{2.1}$$

2.2. Numerical Tests

The standard approach to estimate run-time of algorithms is to count the number of arithmetic operations. Hence, Lemma 2.1 and Equ. (2.1) seem to suggest that the LU-based method is faster than the A^{-1} -based method. We now show that this is not always the case.

To demonstrate this, we solved linear systems of the form (1.1) using the LU-based and A^{-1} -based methods. For each method and for each $n = 2^j$, $j = 5, \dots, 12$, we randomly choose three different $n \times n$ matrices, for each matrix we solve the linear system $Ax = b^{(k)}$ 1,000 times, sequentially, and calculate the average CPU time per RHS. This measurement of CPU time does not include the preprocessing time of LU decomposition or of finding the inverse.

The tests were conducted on a Pentium 4 with 1 Mb cache and 2 GB RAM using Matlab 7.0 R14 for Linux. In the first test we used Intel's MKL BLAS package. The average CPU time per RHS with the A^{-1} -based method is smaller by a factor of 15–35 compared with the LU-based method (see Fig. 2), for matrix sizes varying from 32×32 to 4096×4096 .

Next, we repeated the same test, but changed the software package that Matlab uses for the linear algebra subroutines from Intel's MKL BLAS to ATLAS BLAS. In this case, the average CPU time for a forward and a backward substitution did not change, but the average CPU time for multiplication by A^{-1} increased by a factor of ≈ 1.2 – 1.9 , hence the




Fig. 2. Also plotted, is the average CPU time using ATLAS BLAS for solving $Ax = b^{(k)}$ using the forward and backward substitution (Eq. (1.3), \circ) and through multiplication by A^{-1} (Eq. (1.5), \diamond).

A^{-1} -based method was “only” 8–16 times faster than the LU-based method, (see Fig. 2).

Finally, we ran a benchmark test in which we implemented the multiplication by A^{-1} and the backward and forward substitution using a naive straightforward non-vectorized code of a double loop. In that case, the average CPU time is about the same for forward and backward substitution and for multiplication by A^{-1} (see Fig. 2). As expected, the naive implementation was slower than the MKL BLAS implementation. The matrix-vector multiplication by A^{-1} was 35–50 times faster (using MKL BLAS) than with the naive implementation. On the other hand, forward and backward substitution using MKL BLAS was only 1.4–2.6 times faster than forward and backward substitution with the naive implementation.

We also conducted the same tests on a different Pentium 4 machine (0.5 Mb cache and 0.5 GB RAM using Matlab 7.0 R14 for Windows XP). Under this configuration, all qualitative results remained the same, but the performance ratios changed. For example, the A^{-1} -based method was “only” 15–20 times faster than the LU-based method using Intel’s MKL BLAS, and 8–12 faster under ATLAS BLAS.

2.3. Data Analysis

The simulation of Sec. 2.2 shows that, for sequentially solving the system and neglecting the preprocessing costs, the A^{-1} -based method is significantly faster than the LU-based method (per RHS). Since the number of arithmetic operations per RHS is the same for both methods (see

Efficient Solution of $Ax^{(k)} = b^{(k)}$ Using A^{-1}




Fig. 3. In addition average CPU time per RHS using a naive implementation for the LU-based method (solid) and the A^{-1} -based method (dashes).

Sec. 2.1), we conclude that the A^{-1} -based method is faster because the implementation of matrix multiplication in MKL BLAS or in ATLAS BLAS is faster than the implementation of a forward and backward substitution. Indeed, when we used a naive implementation, there was no performance difference between the two methods. In retrospect, the fact that the A^{-1} -based method is faster is not surprising since the memory structure and indexing of a full matrix is inherently simpler than that of a triangular matrix. Therefore, matrix multiplication can be more easily accelerated through software and hardware optimization. In other words, the performance difference between the two methods is due to the implementation of the numerical linear algebra software package. This conclusion is further supported by the observation that changing the BLAS package leads to a change in the performance ratio from 15–35 to 8–16, (see Fig. 3).

In the numerical experiments shown so far, we neglected the preprocessing cost and considered only the cost of the solving the linear systems (1.3) and (1.5). This corresponds to the case where $K \gg n$, since then the preprocessing costs are negligible. We now take into account the preprocessing cost and look at the overall cost of solving (1.1) in both methods. For example, setting $K = 10$ and comparing the overall cost (preprocessing+solving for 10 RHS) of the two methods, we obtain that the A^{-1} -based method is faster than the LU-based method for matrices of size $n < n_{\text{th}}$ and slower for $n > n_{\text{th}}$, where the threshold is $n_{\text{th}}(K = 10) = 700$ (see Fig. 4). Equivalently, for a given n , we can find a threshold $K_{\text{th}} = K_{\text{th}}(n)$, such the overall cost for solving (1.1) is faster for the A^{-1} -method when $K > K_{\text{th}}$ and slower when $K < K_{\text{th}}$.




Fig. 4. Overall cpu time of the LU-based method (dots) and the A^{-1} -based method (solid) with $K = 10$.

There is a linear relation between K_{th} and n . Indeed, the overall cost of the LU-based method is

$$T_{\text{lu}}^{\text{preprocessing}} + K \cdot T_{\text{lu}}^{\text{RHS}} = \alpha_{\text{lu}}^{\text{preprocessing}} n^3 + K \cdot \beta_{\text{lu}}^{\text{RHS}} n^2 + \mathcal{O}(Kn, n^2), \quad (2.2)$$

where $T_{\text{lu}}^{\text{preprocessing}}$ is the preprocessing cost of (1.2) and $T_{\text{lu}}^{\text{RHS}}$ is the cost of solving the linear systems (1.3) by forward and backward substitution. The overall cost of the A^{-1} -based method is

$$T_{\text{inverse}}^{\text{preprocessing}} + K \cdot T_{\text{inverse}}^{\text{RHS}} = \alpha_{\text{inverse}}^{\text{preprocessing}} n^3 + K \cdot \beta_{\text{inverse}}^{\text{RHS}} n^2 + \mathcal{O}(Kn, n^2), \quad (2.3)$$

where $T_{\text{inverse}}^{\text{preprocessing}}$ is the preprocessing cost (1.4) and $T_{\text{inverse}}^{\text{RHS}}$ is the cost of multiplication by A^{-1} , see (1.5). Hence,

$$\begin{aligned} K_{\text{th}} &= \frac{T_{\text{inverse}}^{\text{preprocessing}} - T_{\text{lu}}^{\text{preprocessing}}}{T_{\text{lu}}^{\text{RHS}} - T_{\text{inverse}}^{\text{RHS}}} \\ &= \frac{(\alpha_{\text{inverse}}^{\text{preprocessing}} - \alpha_{\text{lu}}^{\text{preprocessing}})n^3 + \mathcal{O}(n^2)}{(\beta_{\text{lu}}^{\text{RHS}} - \beta_{\text{inverse}}^{\text{RHS}})n^2 + \mathcal{O}(n)} \sim \frac{\alpha_{\text{inverse}}^{\text{preprocessing}} - \alpha_{\text{lu}}^{\text{preprocessing}}}{\beta_{\text{lu}}^{\text{RHS}} - \beta_{\text{inverse}}^{\text{RHS}}} n + \mathcal{O}(1). \end{aligned} \quad (2.4)$$

In Fig. 5 we plot the value of K_{th} as a function of n . The results show that $K_{\text{th}} \sim 0.0012n + 1.25$, in agreement with (2.4).

As noted, Observation 2.1 has been traditionally interpreted to imply that the LU-based method is faster than the A^{-1} -based method for

Efficient Solution of $Ax^{(k)} = b^{(k)}$ Using A^{-1}




Fig. 5. K_{th} as a function of n (solid). Dotted line is the fitted linear curve $K_{\text{th}} = 0.0124n + 1.25$.




Fig. 6. Average CPU time for solving $Ax = b^{(k)}$ using standard LU method (solid), LU with explicit pivot matrix (dash-dotted), LU and linsolve (dotted) and A^{-1} -method (dashed).

$K = \mathcal{O}(n)$ and that the run time of the two methods is the same for $K \gg n$. In contrast, in Fig. 6 we see that the A^{-1} -based method is faster for $K > cn$, where $c \approx 0.012$. The reason for this disagreement is the implicit assumption that the number of arithmetic operations is a good measure for comparison of run times. This assumption fails, however, to capture the large difference between $\beta_{\text{inverse}}^{\text{RHS}}$ and $\beta_{\text{lu}}^{\text{RHS}}$ that results from the difference in implementation of matrix-vector multiplication and of forward and backward substitutions.

2.4. Improving the LU Solver

In our numerical tests so far, we have used the most straightforward implementation of the LU algorithm for solving (1.1). We first calculated

the LU decomposition (1.2). Then, for each right-hand-side we solved the linear system using a forward substitution and a backward substitution, see equation (1.3), which are implemented by a backslash (`\`) or equivalently by the `mldivide` Matlab command.

As stated by the Matlab's help, the command

$$[L, U] = \text{lu}(A),$$

returns an upper-triangular matrix in U , and a “psychologically lower-triangular” matrix (i.e., a potentially permuted lower triangular matrix) in L . The “psychologically lower-triangular” form makes it harder to efficiently implement forward substitution. A faster way is, therefore, to enforce L to be a strict lower triangular matrix by using the command

$$[L, U, P] = \text{lu}(A), \quad (2.5)$$

where P is the permutation matrix. Then, the solution to $Ax = b$ can be obtained by

$$\begin{aligned} y &= L \setminus (P \cdot b), \\ x &= U \setminus y. \end{aligned} \quad (2.6)$$

The additional calculations in this modification (i.e., the calculation of $P \cdot b$) are $\mathcal{O}(n)$, hence are negligible compared with the time saved by forcing L to be triangular. Indeed, Figure 6 shows that using the strict lower-triangular matrix is about 2.5 times faster than the standard method.

The backslash command, e.g., $A \setminus b$, involves a preprocessing stage in which the properties of the matrix A are checked in order to pick the most appropriate solver for the problem. Thus, in the case of Eq. (2.6), the backslash command identifies that L and U are triangular matrices and uses backward or forward substitution to obtain the solution. Since it is known that L and U in (2.6) are lower- and upper-triangular matrices, it is possible to bypass the preprocessing stage by using the `linsolve` command instead of the backslash command. Therefore, we use (2.5) to obtain a strict lower-triangular matrix L , but replace (2.6) with

$$\begin{aligned} \text{optLT.LT} &= \text{true}; \quad \% \text{ lower-triangular property,} & (2.7) \\ \text{optUT.UT} &= \text{true}; \quad \% \text{ upper-triangular property,} \\ y &= \text{linsolve}(L, P \cdot b, \text{optLT}); \\ x &= \text{linsolve}(U, y, \text{optUT}). \end{aligned}$$

Figure 6 shows that using (2.5) and (2.7) is 4.6 times faster than the standard method, i.e. Eqs. (1.2) and (1.3), and 1.3 times faster than using

Efficient Solution of $Ax^{(k)} = b^{(k)}$ Using A^{-1}

a strict lower-triangular matrix without linsolve, i.e. (2.5) and (2.6). *Even after all these improvements, however, the A^{-1} method is still considerably faster.*

3. ACCURACY AND STABILITY

The second common argument in the literature against solving $Ax = b$ with A^{-1} is its numerical accuracy and stability (see citations in the Introduction). As noted, this argument is based on analysis which showed that the backward error in the A^{-1} -based method is larger than in the LU-based method, and that the ratio between the two increases with the condition number of A . While there has been no such analytical results for the forward error, it was implicitly assumed that it would behave similarly. Our simulations however suggest that this implicit assumption is wrong. Indeed, we observe that the *forward error is of the same magnitude for both the A^{-1} -based method and the LU-based method, regardless of the condition number of A .*

3.1. Forward Error

In the A^{-1} -based method,

$$\frac{\|\Delta x\|}{\|x\|} \leq \text{cond}(A) \frac{\|\Delta b\|}{\|b\|}.$$

In the LU-based method x is determined by solving the linear systems (1.3), hence,

$$\frac{\|\Delta x\|}{\|x\|} \leq \text{cond}(L)\text{cond}(U) \frac{\|\Delta b\|}{\|b\|}.$$

Since $\text{cond}(A) \leq \text{cond}(L)\text{cond}(U)$, the worst-case relative (forward) error $\frac{\|\Delta x\|}{\|x\|}$ is likely to be *larger* for the LU-based method than for the A^{-1} -based method. On the other hand, the roundoff error involved in matrix-vector multiplication is greater than the roundoff error of backward and forward substitution [4, p. 262].

We now show that, in practice, the relative error of the A^{-1} -based method is only slightly larger than the relative error of the LU-based method. To do so, we first solved linear systems of the form (1.1) using the LU-based and A^{-1} -based methods. For each $n = 100, 200, \dots, 4000$, we randomly choose 25 different $n \times n$ matrices and a solution vector x . For each matrix A , we first generate the vector $b = Ax$, then solve with each method $A\hat{x} = b$, calculate the relative forward error $\frac{\|\hat{x} - x\|_\infty}{\|x\|_\infty}$ and then




Fig. 7. (A) Relative forward error $\frac{\|\hat{x} - x\|_\infty}{\|x\|_\infty}$ for $\hat{x} = A^{-1}b$ (solid) and $\hat{x} = A \setminus b$ (dots), (B) Relative backward error for the same data.

average over all 25 matrices. Results in Fig. 7A show that the forward error of the LU-based is smaller by a factor of 2–4 than for the A^{-1} -based method. This 2–4 factor does not change significantly as n changes from 100 to 4000, and $\text{cond}(A)$ varies between 10^3 and 10^6 . In order to see whether larger differences would appear for ill-conditioned matrices, we perform the same test with Matlab’s matrix gallery of involutory ill-conditioned matrices (Fig. 7A). In this case, the relative forward error for the LU-based method is slightly below the error for the A^{-1} -based method, differing by a factor of 1–1.5, as $\text{cond}(A)$ varies between 1 and 10^{16} . Repeating this test for tridiagonal ill-conditioned matrices and asymmetric ill-conditioned matrices gave nearly identical results, with factors of 1–2 and 1.4–3.5, respectively.

3.2. Backward Error

The backward error $\|A\hat{x} - b\|$ is known to be significantly larger in the A^{-1} -based method than in the LU-based method when A is ill-conditioned [4, p. 262]. Indeed, Fig. 7B shows that the backward error of the LU-based method is smaller by a factor of 50–200 than for the A^{-1} -based method, for the same matrices used to generate Fig. 7A. Similarly, Fig. 7B shows that the backward error of the LU-based method becomes exponentially smaller than for the A^{-1} -based method, for the same ill-conditioned involutory matrices used to generate Fig. 8A.

These results, thus, confirm that the LU-based method is superior to the A^{-1} -based method in terms of backward error. However, comparison between Fig. 7A and B and between Fig. 8A and B shows, that the backward error is a poor prediction of the forward error.

Efficient Solution of $Ax^{(k)} = b^{(k)}$ Using A^{-1}




Fig. 8. (A) Forward error $\frac{\|x - \hat{x}\|_\infty}{\|\hat{x}\|_\infty}$ for $\hat{x} = A^{-1}b$ (solid) and for $\hat{x} = A \setminus b$ (dots). The two lines are almost indistinguishable, (B) Backward error $\frac{\|A\hat{x} - b\|_\infty}{\|b\|_\infty}$ for the same data.

Remark. If the matrix A can be diagonalized, the worst-case scenario is that b is the eigenvector of the maximal eigenvalue (in absolute value) and Δb is proportional to the eigenvector of the minimal eigenvalue. However, in most applications arising from PDE, this is not the case. The main ‘mass’ is concentrated in the lower modes, i.e. in the eigenvector correlated to the lower eigenvalues, while the ‘noise’ is related to the high-frequency modes. Therefore in such applications the worst-case scenario is unlikely to happen.

4. POTENTIAL APPLICATIONS

Any application that involves solving many linear systems with a constant *full* matrix A can benefit from adopting the A^{-1} -based approach over the LU-based method. We now list various applications which are potentially suitable for the A^{-1} -based method:

1. *Spectral differentiation methods for PDEs:* While finite difference methods require solving linear systems with a sparse banded matrix, spectral methods, in many cases, require solving linear systems with full matrices [5]. For example, when solving a parabolic PDE in an implicit method, each time step involves solving a linear system of the form

$$Dx^{(t+\Delta t)} = F(x^{(t)}),$$

where D is a full matrix which does not vary with time t . In this case, if the number of time steps $K > K_{\text{th}}(n)$, it is faster to adopt the A^{-1} -based approach, i.e., we invert D at the beginning of the simulation, and then at each step calculate

$$x^{(t+\Delta t)} = D^{-1}F(x^{(t)}).$$

2. *Modified Newton methods:* Newton methods are used to find the roots of nonlinear equations. They involve the iterations

$$x^{(k+1)} = x^{(k)} - [J^{(k)}]^{-1}F(x^{(k)}), \quad (4.1)$$

where $J^{(k)} = J(x^{(k)})$ is the Jacobian of the function F at the point $x^{(k)}$. Hence, $x^{(k+1)}$ is the solution of the linear system

$$J^{(k)}x^{(k+1)} = b^{(k)}, \quad b^{(k)} = J^{(k)}x^{(k)} - F(x^{(k)}). \quad (4.2)$$

The calculation of the Jacobian at every iteration is typically very expensive. To reduce the cost of the Jacobian calculation, in the modified Newton method [6], the Jacobian is calculated only once every several (K) iterations.

The common approach is to avoid the explicit calculation of J^{-1} , and find $x^{(k+1)}$ from (4.2). Our analysis shows that it is faster to use the A^{-1} -based approach and calculate $x^{(k+1)}$ from (4.1) for “small” matrices of size $n < n_{\text{th}}(K)$.

A similar algorithm which may benefit from applying the A^{-1} -based method is the Update Skipping BFGS algorithm used in optimization problems [7]. In this case, the iterations of the BFGS method are

$$B_k x^{(k+1)} = B_k x^{(k)} - \lambda J(x^{(k)}), \quad (4.3)$$

where B_k is an approximation to the Hessian of F at the point $x^{(k)}$ [8]. To reduce the cost of the Hessian calculation, in the Update Skipping BFGS algorithm, B_k is updated only once every K steps. In this case, Our analysis shows that it is faster to use the A^{-1} -based approach and calculate $x^{(k+1)}$ from

$$x^{(k+1)} = x^{(k)} - \lambda_k B_k^{-1} J(x^{(k)}) \quad (4.4)$$

for matrices of size $n < n_{\text{th}}(K)$.

3. *Inverse iteration:* The smallest eigenvalue of a full matrix A (and the corresponding eigenvector) can be found from the sequence $\frac{\|v^{(k+1)}\|}{\|v^{(k)}\|}$, where [4],

$$Av^{(k+1)} = v^{(k)}. \quad (4.5)$$

Table I. Comparison of CPU Run Time

Method	Equations	CPU time (seconds)
$[L, U]$	(1.2), (1.3)	19.76
$[L, U, P]$	(2.5), (2.6)	7.6
$[L, U, P]$ + linsolve	(2.5), (2.7)	4.62
A^{-1}	(1.4), (1.5)	1.88

The common approach is to calculate $v^{(k+1)}$ using the LU-based method, i.e.,

$$Ly = v^{(k)}, \quad Uv^{(k+1)} = y. \quad (4.6)$$

However, in this case, calculating $v^{(k+1)}$ using the A^{-1} -based approach

$$v^{(k+1)} = A^{-1}v^{(k)} \quad (4.7)$$

is faster if the number of iterations $K > K_{\text{th}}(n)$.

As a final example, we implement the inverse-power method (4.5) using both methods as follows. We randomly choose 250 random 100×100 matrices. For each matrix we measure the *overall* cpu-time (preprocessing and iterations) needed to calculate the smallest eigenvalue using the inverse-power method implemented with the A^{-1} -based method (4.7) and the LU-based method (4.6). We stop the iterations when the relative error is below 10^{-14} (with respect to the smallest eigenvalue found by Matlab's eig function).

In our simulations we observed that the A^{-1} -based method was about 11 times faster than the LU-based method (see Table I). The average number of iterations needed for the iterations to converge was 269 in A^{-1} -based method and 264 in the LU-based method. This 2% difference is probably due to the larger roundoff errors in the A^{-1} -based method. Therefore, in this case, the A^{-1} -based method is significantly faster than the LU-based method, and just as accurate. Using the ‘‘accelerated’’ LU algorithms (see Sect. 2.4) lead to a considerable reduction in the overall CPU time, but was still 2.5 times slower than using A^{-1} .

ACKNOWLEDGMENT

We thank Eli Turkel and Sivan Toledo for useful discussions.

REFERENCES

1. Golub, G., and van Loan, C. (1989). *Matrix Computations*, 2nd edn. The Johns Hopkins University Press, London.
2. Conte, S. D., and de Boor, C. (1980). *Elementary Numerical Analysis*, 3rd edn. McGraw-Hill book company.
3. Higham, N. J. (1996). *Accuracy and Stability of Numerical Algorithms*. SIAM, Philadelphia, PA.
4. Trefethen, N. L., and Bua, D. (1997). *Numerical Linear Algebra*. SIAM, Philadelphia.
5. Trefethen, N. L. (2000). *Spectral Methods in MATLAB*. SIAM, Philadelphia, PA.
6. Quarteroni, A., Sacco, R., and Saleri, F. (2000). *Numerical Mathematics*. Springer, New York.
7. Dennis, J. E., Schnabel, Jr., and Robert, B. (1983). *Numerical Methods for Unconstrained Optimization and Nonlinear Equations*. Prentice-Hall, New York.
8. Tamara, K. G., O'Leary, D. P., and Nazareth, L. (1998). BFGS with update skipping and varying memory. *SIAM J. Optim.* **8**(4), 1060–1083.