Bull. Nov. Comp. Center, Comp. Science, 19 (2003), 39–50

The Godunov Inverse Iteration algorithm for symmetric tridiagonal matrices

Anna M. Matsekh, Ella P. Shurina

We have developed a hybrid procedure based on the Godunov algorithm for computing eigenvectors of tridiagonal symmetric matrices and inverse iteration, which we call the Godunov Inverse Iteration algorithm. It employs the inverse iteration to improve the accuracy of eigenvectors computed according to the Godunov method with the embedded Modified Gram-Schmidt procedure to reorthogonalize eigenvectors corresponding to computationally coincidental eigenvalues and which may be missing a few digits of precision due to the round-off errors. We present some experimental results to illustrate that the new hybrid method produces results superior to both the Godunov method and standard implementations of the inverse iteration just on an iterative step. We also discuss some issues involved in the parallel implementation of the new method.

1. The Godunov method

Consider the fundamental algebraic eigenvalue problem, in which

$$Ax = \lambda x \tag{1}$$

for real symmetric matrices $A \in R^{n\times n}$. There always exists a real orthogonal transformation $X \in X^{n\times n}$ such that the matrix A is diagonalizable, that is,

$$X^T AX = \mathrm{diag}(\lambda_i), \tag{2}$$

where eigenvalues λ_i, $i = 1, \ldots, n$, are all real. The Godunov method employs the Rayleigh–Ritz scheme to solve problem (1) with a real symmetric matrix A in three steps:

(i) compute the orthonormal transformation Q such that matrix $T = Q^T AQ$ is tridiagonal;

(ii) solve the eigenproblem $Tu = \mu u$;

(iii) take (μ, Qu) as approximation to the eigenpair (λ, x).

The Godunov method [1] was designed to compute eigenvectors of unreduced symmetric tridiagonal matrices by a sequence of plane rotations on architectures that support an utmost precision and directed rounding. The Godunov eigenvector approximation u^i corresponding to the eigenvalue μ_i of the matrix T from the interval (α_i, β_i) such that

$$|\beta_i - \alpha_i| \le \varepsilon_{\text{mach}} \|T\|_2, \tag{3}$$

determined by a Sturm-sequence-based bisection procedure, is found recursively from the two-sided Sturm sequence $P_k(\mu_i)$, $i,k = 1,\dots,n$, where $P_{k-1}(\mu_i) = \operatorname{ctg}\theta_i$, setting

$$u_0^i = 1 \quad \text{and} \quad u_k^i = -u_{k-1}^i \frac{\operatorname{sign} b_{k-1}}{P_{k-1}(\mu_i)} \tag{4}$$

in just $O(n)$ operations per a normalized eigenvector. The two-sided Sturm sequence

$$P_0(\mu_i),\dots,P_{n-1}(\mu_i) \stackrel{\text{def}}{=} P_0^+(\alpha_i),\dots P_l^+(\alpha_i), P_{l+1}^-(\beta_i),\dots P_{n-1}^-(\beta_i) \tag{5}$$

is constructed from the left-sided and the right-sided Sturm sequences $P_k^+(\alpha_i)$, $k = 0,\dots,n-1$, and $P_k^-(\beta_i)$, $k = n-1,\dots,0$. The left-sided Sturm sequence $P_k^+(\alpha_i)$ is computed from minors of the matrix $T - \alpha_i I$ according to the formulas [1]

$$\begin{aligned} P_0^+(\alpha_i) &= \frac{|b_0|}{d_0 - \alpha_i}, \\ P_k^+(\alpha_i) &= \frac{|b_k|}{(d_k - \alpha_i - |b_{k-1}|)P_{k-1}^+(\alpha_i)}, \\ P_{n-1}^+(\alpha_i) &= \frac{1}{(d_{n-1} - \alpha_i - |b_{n-2}|)P_{n-2}^+(\alpha_i)}, \end{aligned} \tag{6}$$

while the right-sided Sturm sequence $P_k^-(\beta_i)$ is computed from minors of the matrix $T - \beta_i I$ as follows [1]

$$\begin{aligned} P_{n-1}^-(\beta_i) &= d_{n-1} - \beta_i, \\ P_k^-(\beta_i) &= \frac{d_k - \beta_i - |b_k|/P_{k+1}^-(\beta_i)}{|b_{k-1}|}, \\ P_0^-(\beta_i) &= d_0 - \beta_i - |b_0|/P_1^-(\beta_i). \end{aligned} \tag{7}$$

Although analytically equivalent, eigenvectors constructed in finite precision from the left-sided and the right-sided Sturm sequences for the same parameter λ are in general different. An eigenvector with an ensured accuracy is obtained when the left- and the right-hand sequences (5) are joint at an index l chosen according to the rule based on the Sturm theorem: for any real λ_0 the number of roots λ of the n-th principal minor of the matrix $T - \lambda I$, such that $\lambda < \lambda_0$ coincides with the number of non-positive values in the Sturm sequence $P_k(\lambda_i)$, $k = 1,\dots,n$. Let l^+ be the number of non-positive elements in the sequence $P_k^+(\alpha_i)$, $k = 0,\dots,n-1$, and $n - 1 - l^-$ be the number of non-negative elements in the sequence $P_k^-(\beta_i)$, $k = n-1,\dots,0$. Then the left and the right sequences (5) are joint at the index $l = l^+ = l^-$, for which the following condition is satisfied [1]:

$$(P_l^+(\alpha_i) - P_l^-(\beta_i))(1/P_{l+1}^-(\beta_i) - 1/P_{l+1}^+(\alpha_i)) \leq 0. \tag{8}$$

In our attempt to improve the Godunov method, we were motivated by the fact that it is a direct method, and due to the round-off errors in the finite precision, the error bound [1]

$$\|(T - \mu_k I)u_k\|_2 \leq 13\sqrt{3}\,\varepsilon_{\text{mach}}\|T\|_2\|u_k\|_2 \tag{9}$$

is not attained. At the same time, the two-sided Sturm sequence computations are susceptible to division by zero and overflow errors, while for computationally coincident and closely clustered interior eigenvalues it produces coincident or nearly coincident eigenvectors, taking no measures for reorthogonalization. In empirical studies, our implementation of the Godunov method consistently delivered residuals that were approximately two orders of magnitude larger than those of the eigenvectors computed according to the LAPACK version of the inverse iteration [2] and the EISPACK version of the inverse iteration TINVIT [3] for the same eigenvalue approximations. In addition, due to the round-off errors, the machine representation of the matrix T, that is, generally obtained either by the Householder or the Lanczos tridiagonalization, has the form $T_{\text{mach}} = T + G$ [4], where

$$\|G\| \leq k\sqrt{n}\,2^{-t} \tag{10}$$

and t is the number of mantissa bits in the machine representation of floating-point numbers. Therefore in the finite precision, the error bound (9) rather takes takes the following form:

$$\|(T_{\text{mach}} - \mu_k I)u_k\| \leq\ 13\sqrt{3}\,\varepsilon_{\text{mach}}\|T\|_2\|u_k\|_2 + k\sqrt{n}\,2^{-t}\|u_k\|. \tag{11}$$

2. The Godunov Inverse Iteration algorithm

The Godunov Inverse Iteration is a hybrid procedure for computing accurate approximations to the symmetric eigenvalue problem based on the algorithm by S.K. Godunov et al. [1] and on the inverse iteration. The Godunov Inverse Iteration algorithm was constructed to avoid common computational problems arising both in the Godunov method and in the inverse iteration method. It can be viewed as an algorithm that delivers the reorthogonalized iteratively improved the Godunov eigenvectors.

Instead of initiating an inverse iteration with a random vector, or solving a linear system to find a starting vector, as is customary in many implementations of the inverse iteration, we use the eigenvector, computed by the Godunov method in just $O(n)$ arithmetic operations as an extremely accurate starting vector in the Inverse Iteration. Before the inverse iteration is

applied, any non-numeric elements of the Godunov eigenvectors are substituted with random numbers. This semi-deterministic approach to finding initial vectors to the inverse iteration reduces the number of steps necessary for the convergence to a desired accuracy. In most cases the convergence is attained after one step of inverse iteration.

Typically, the inverse iteration breaks down when very accurate eigenvalue approximations are used as shifts, because the corresponding system in this case is nearly singular. To avoid such a breakdown, small disturbances are usually introduced into the eigenvalues used as shifts in the inverse iteration to ensure the convergence to the corresponding Ritz vectors. But even small arbitrary departures of the Ritz values from exact eigenvalues may produce significant departures of the Ritz vectors from actual eigenvectors. We solve this problem by using the right-hand bounds β_k of the intervals $\mu_i \in (\alpha_i, \beta_i)$, $i = 1, \ldots, n$, as shifts that are ensured to be within the error bounds (3). We apply the Modified Gram–Schmidt reorthogonalization for the eigenvector approximations corresponding to multiple eigenvalues or to the clustered eigenvalues with small relative gaps. We use Wilkinson's stopping criteria [4] $\|x_k\|_\infty \geq 2^t/(100n)$ to verify that the convergence is attained. Below we present a formal description of the Godunov Inverse Iteration algorithm.

Godunov Inverse Iteration algorithm Compute eigenvectors u^i, $i = 1, \ldots, n$, of the tridiagonal matrix $T = T^T \in R^{n \times n}$ with the main diagonal d and the codiagonal b on a processor with the machine precision $\varepsilon_{\text{mach}}$ and t mantissa bits.

```
godunov_inverse_iteration(d, b, n)

bisection(d, b, n)
  for (i=1, i<=n, i++)
    find eigenintervals (α_i, β_i) that contain eigenvalues μ_i ∈ (α_i, β_i)
    s.t. ||β_i − α_i| ≤ ε_mach ||T||_2, i = 1, ..., n.
end

godunov_eigenvector_method(d, b, n, α, β)
  for (i=1, i<=n, i++)
    compute the two-sided Sturm sequence P_n(μ_i):
      P_1^+(α_i) = |b_1|/(d_1 − α_i),
      P_k^+(α_i) = |b_k|/(d_k − α_i − |b_{k−1}|)P_{k−1}^+(α_i), k = 2, 3, ..., n − 1,
      P_n^+(α_i) = 1/(d_n − α_i − |b_{n−1}|)P_{n−1}^+(α_i),
      P_n^−(β_i) = d_n − β_i,
      P_k^−(β_i) = (d_k − β_i − |b_k|/P_{k+1}^−(β_i))/|b_{k−1}|, k = n − 1, n − 2, ..., 2,
      P_1^−(β_i) = d_1 − β_i − |b_1|/P_2^−(β_i)
```

find $l = l^+ = l^-$ s.t. $(P_l^+(\alpha_i) - P_l^-(\beta_i))(1/P_{l+1}^-(\beta_i) - 1/P_{l+1}^+(\alpha_i)) \le 0$

set $P_1(\mu_i), \ldots, P_n(\mu_i) \stackrel{\text{def}}{=} P_1^+(\alpha_i), \ldots P_l^+(\alpha_i), P_{l+1}^-(\beta_i), \ldots P_n^-(\beta_i)$

compute the Godunov eigenvector u^i: $u_1^i = 1$,

```
for (k=2, k<=n, k++)
```

$u_k^i = -u_{k-1}^i \operatorname{sign}(b_{i-1})/P_{i-1}(\mu_i)$

```
end
end
```

```
inverse_iteration(d, b, u, β)
  /* Preprocessing step */
  for (i=1, i<=n, i++)
    for (k=1, k<=n, k++)
      if u_k^i is not a machine number
      then set u_k^i to a random uniform number from (0,1)
    end
    use the right ends of the eigenintervals as shifts: γ_i = β_i
    perturb the computationally coincident eigenvalues:
    if (i > 0 ∩ |γ_i − γ_{i−1}| ≤ 10ε_mach|γ_i|)
    then γ_i = γ_{i−1} + 10ε_mach|γ_i|
  end
  /* Inverse iteration step */
  k = 0, δ = 2^t/(100n)
  do
    k = k + 1
    Solve (T − γ_i I)z = u^k
    The Modified Gram-Schmidt reorthogonalization step
    for (j=1, j < k, k++)
      if |γ_j − γ_k| ≤ ‖T‖_∞/1000
      then z = z − (z, u^j)u^j
    end
    u^k = z/‖z‖_2
  while (‖z‖_∞ > δ)
```

3. Experimental results

We have implemented and tested the Godunov method, the Godunov Inverse Iteration algorithm, the inverse iteration algorithm with random starting vectors which we call the Random Inverse Iteration algorithm (our implementation of the LAPACK procedure xSTEIN [1], and the inverse iteration algorithm with initial vectors found as direct solution to the eigenproblem, which we call the Direct Inverse Iteration algorithm (our implementation

of the EISPACK procedure TINVIT [3], in ANSI C (GNU C compiler) in IEEE double precision and tested these programs on an IntelR Xeon^{TM} CPU 1500MHz processor.

To make a fair comparison, we compute eigenvalue approximations to the symmetric eigenproblem

$$Ax_k = \lambda_k x_k, \ k = 1, 2, \ldots, n, \tag{12}$$

only once and use these eigenvalues to compute eigenvectors using four different routines, while in all the three inverse iteration implementations we use the same direct solver for systems of linear algebraic equations with tridiagonal symmetric matrices. We use the Householder tridiagonalization with dense matrices and restarted the Lanczos procedure with selective reorthogonalization with sparse matrices.

Following Godunov [1], we have implemented the bisection algorithm to find the intervals (α_i, β_i) containing the eigenvalues μ_i of the tridiagonal matrix $T = \tilde{Q}^T A \tilde{Q}$, with the ensured accuracy

$$|\beta_i - \alpha_i| \leq \varepsilon_{\text{mach}} \|T\|_2, \ i = 1, \ldots, n, \tag{13}$$

where $\varepsilon_{\text{mach}}$ is the unit round-off error. The bisection algorithm requires $O(tn^2)$ operations, where t is the number of bits of precision in a computer representation of floating point numbers.

In both the original Godunov method and in the new Godunov Inverse Iteration procedure we used eigenintervals (α_i, β_i) to compute the corresponding eigenvectors, while in the Random Inverse Iteration and the Direct Inverse Iteration versions of the inverse iteration algorithm we used $\mu_i = (\alpha_i + \beta_i)/2$, i.e., a typical choice for the i-th eigenvalue approximation in the inverse iteration implementations. If the analytical solution to a test eigenvalue problem is available, we report the accuracy of the computed eigenvalues as the maximum absolute deviation of the computed eigenvalue approximation $\mu_i = (\alpha_i + \beta_i)/2$ from its analytical counterpart. If the analytical solution is not available we report the ensured accuracy with which eigenintervals (α_i, β_i) were computed. To determine the orthogonality of the computed eigenvectors $\tilde{X}$, we consider the maximum euclidean vector norm

$$\text{err}_1 \stackrel{\text{def}}{=} \max_k \|(\tilde{X}^T \tilde{X} - I) e_k\|_2 \tag{14}$$

of the matrix $\tilde{X}^T \tilde{X} - I$, which represents a deviation of the computed basis $\tilde{X}$ from the standard orthonormal basis I comprised of the unit vectors e_k: $I = |e_k|$, $k = 1, 2, \ldots, n$.

To determine the accuracy of the computed eigendecomposition $\tilde{X}^T A \tilde{X} = \text{diag}(\tilde{\lambda}_i)$ of the eigenproblem we consider a maximum absolute residual error of the computed eigenpairs $(\tilde{\lambda}_k, \tilde{x}_k)$ in the euclidean norm:

$$\mathrm{err}_2 \stackrel{\mathrm{def}}{=} \frac{\max_k \|(A\tilde{x}_k - \tilde{x}_k\tilde{\lambda}_k)\|_2}{\max_k |\tilde{\lambda}_k|}. \tag{15}$$

In all the tests presented below, the Godunov Inverse Iteration procedure converged to desired accuracy in just one step, while the results were almost as accurate, and in a number of cases superior to the ones obtained with the Random Inverse Iteration and Direct Inverse Iteration.

Example 1. The tridiagonal symmetric eigenproblem $Rx = \lambda x$, $\lambda(R) = -\cos(k\pi/(n+1))$, $k = 1, \ldots, n$,

$$R = \begin{pmatrix} 0 & 0.5 & & \cdots & 0 \\ 0.5 & 0 & \ddots & & \vdots \\ & \ddots & \ddots & \ddots & \\ \vdots & & \ddots & \ddots & 0.5 \\ 0 & \cdots & & 0.5 & 0 \end{pmatrix}$$

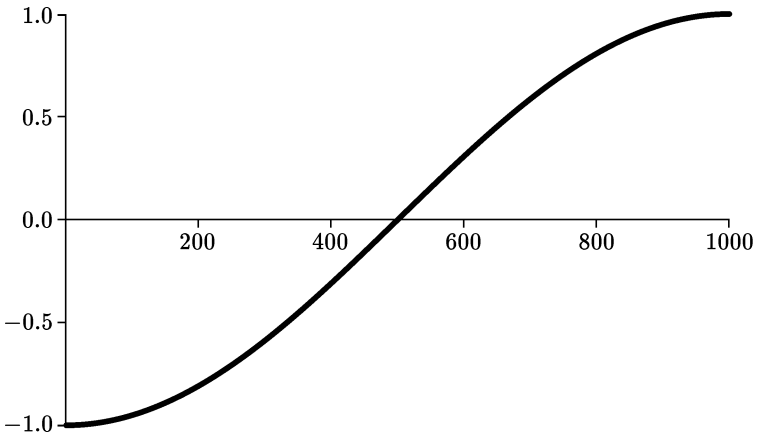

Figure 1. Eigenvalues $\tilde{\lambda}(R)$ computed by the bisection method for $n = 1000$

The problem of finding eigenvalues and eigenvectors of tridiagonal symmetric matrices with a zero main diagonal, in the so-called Golub–Kahan form [5], which arises in singular-value computations for bidiagonal matrices, and, generally, for non-symmetric matrices, presents a number of computational challenges. In Example 1, we compare the eigenvectors computed with the Godunov Inverse Iteration method against the eigenvectors computed according to the Godunov method, and Direct and Random Inverse Iteration algorithms for the same approximations of the eigenvalues of 1000×1000 tridiagonal matrix R which has zero diagonal entries and the entries equal 0.5 on the codiagonals. The eigenvalues of this matrix coincide with zeros of the Chebyshev polynomials of the second kind, and so we

were able to compare the analytical solution against eigenvalues computed with our bisection routine. Test results for this example are summarized in Table 1. For the eigenvalue approximations computed with a maximum absolute deviation of 3.3307e−16 from the analytical solution, some of the Direct Inverse Iteration and Random Inverse Iteration eigenvectors X did not converge and were set to zero, which is indicated by the fact that the orthogonality measure equals 1 in these tests.

Table 1. Error estimates for the eigenvectors of the matrix R ($n = 1000$), corresponding to the eigenvalues $\tilde{\lambda} = \tilde{\lambda}(R)$ computed with the specified maximum absolute deviation $\Delta(\lambda)$ from the exact eigenvalues $\lambda = \lambda(R)$

Algorithm	err_1	err_2	# iter.
	$\Delta(\lambda) = 3.3307\mathrm{e}{-16}$		
The Godunov Method	1.1535e−12	2.9458e−11	–
Direct Inverse Iteration	4.3392e−12	1.0000e+00	1
Random Inverse Iteration	2.3845e−16	1.0000e+00	3
Godunov Inverse Iteration	2.3461e−16	1.1138e−14	1
	$\Delta(\lambda) = 4.9960\mathrm{e}{-16}$		
The Godunov Method	2.0175e−12	4.8818e−11	–
Direct Inverse Iteration	5.9214e−12	8.2480e−11	1
Random Inverse Iteration	2.7557e−16	2.2042e−14	3
Godunov Inverse Iteration	2.6821e−16	1.0970e−14	1

The Godunov method and the Direct Inverse Iteration produced eigenvector approximations that were accurate only to 12 digits of machine precision, yet the Godunov eigenvectors satisfied the orthogonality measure to 11 digits of precision. In just one step of iterative improvement, the Godunov Inverse Iteration produced eigenvectors that satisfied the original problem to 16 digits of machine precision, just as the Random Inverse Iteration solution did after three iteration steps. In addition, the Godunov Inverse Iteration solution satisfied the orthogonality measure to 14 digits of machine precision.

When eigenvalues were computed with slightly lower precision (with a maximum absolute deviation of 4.9960e−16 from the analytical solution) all the three versions of the Inverse Iteration converged to a high accuracy, and again, the Godunov Inverse Iteration converged in only one step to virtually the same high accuracy as the Random Inverse Iteration in three steps. Clearly, the conventional Inverse Iteration implementations appear to be very sensitive to the accuracy with which eigenvalue approximations are computed, while the Godunov Inverse Iteration exhibits a robust behavior.

Example 2. Dense symmetric eigenproblem $Ux = \lambda x$:

$$U_{ij} = \begin{cases} 1/(i+j-1), & i = j, \\ -1/(i+j-1), & i \neq j, \end{cases}$$

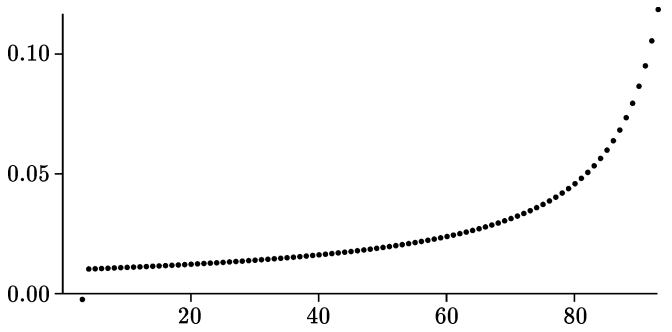

Figure 2. Eigenvalues $\tilde{\lambda}(U)$ computed by the bisection method for $n = 100$

In Example 2 we compare the Godunov method, the Godunov Inverse Iteration, the Random Inverse Iteration, and the Direct Inverse Iteration on an eigenvalue problem with a dense symmetric 100×100 matrix U which was derived from the Hilbert matrix: U has entries equal to $1/(i+j-1)$ on the main diagonal and off-diagonal elements equal to $-1/(i+j-1)$. This is an example of a matrix with clustered eigenvalues which makes a good demonstration of capabilities of both the Godunov and the Godunov Inverse Iteration algorithms. The Godunov method and the Inverse Iteration procedures were applied to the tridiagonal matrix obtained from the matrix U by the Householder reduction. Test results for this example are summarized in Table 2: here we see again that the Godunov Inverse Iteration solution satisfies the original problem and the orthogonality condition to 15 digits of machine precision in one step, virtually the same results as those produced by the Random Inverse Iteration in three steps. The Godunov method and the Direct Inverse Iteration delivered errors about two orders of magnitude higher than the Godunov Inverse Iteration and the Random Inverse Iteration.

Table 2. Error estimates for the eigenvectors of the matrix U ($n = 100$), corresponding to the eigenvalues $\tilde{\lambda} = \tilde{\lambda}(U)$ from the eigeninterval computed with ensured accuracy 8.8818e−16

Algorithm	err_1	err_2	# iter.
The Godunov Method	3.2466e−13	8.0953e−12	–
Direct Inverse Iteration	5.1916e−14	1.2718e−11	1
Random Inverse Iteration	9.1271e−16	2.8498e−15	3
Godunov Inverse Iteration	1.0023e−15	2.7849e−15	1

Example 3. The block-diagonal symmetric eigenproblem $Px = \lambda x$,

$$\lambda(P) = c_2 + 2c_0 \cos\frac{i\pi}{\sqrt{n}+1} + 2c_1 \cos\frac{j\pi}{\sqrt{n}+1}, \quad i,j = 1,\ldots,m,$$

$$P = \begin{pmatrix} c_2 & c_0 & & c_1 & & & \\ c_0 & c_2 & c_0 & & c_1 & & \\ & c_0 & c_2 & & & \ddots & \\ c_1 & & & \ddots & & & c_1 \\ & c_1 & & & c_2 & c_0 & \\ & & \ddots & & c_0 & c_2 & c_0 \\ & & & c_1 & & c_0 & c_2 \end{pmatrix}$$

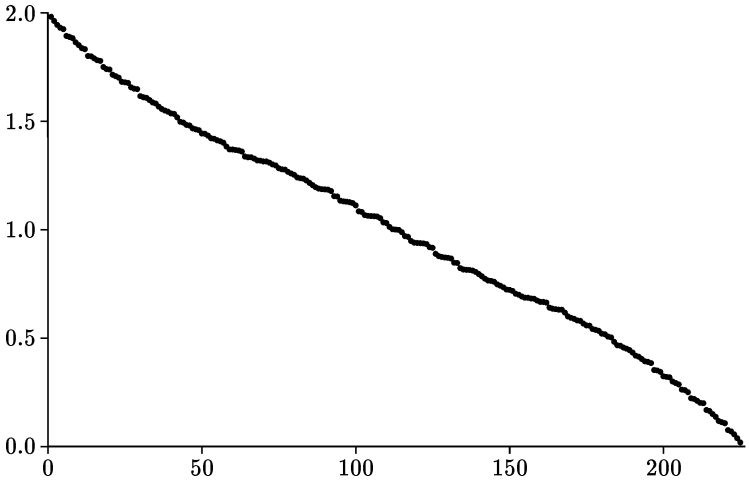

Figure 3. Eigenvalues $\tilde{\lambda}(P)$ computed by the bisection method for $c_2 = 1$, $c_0 = -0.33$, $c_1 = -0.17$, and $n = 225$

In Example 3 we provide the results of the Godunov Inverse Iteration procedure used to solve a sparse symmetric eigenproblem $Px = \lambda x$, transformed to the equivalent tridiagonal problem by the Lanczos method with selective reorthogonalization. The matrix $P \in R^{n \times n}$, $n = m^2$, is a version of the matrices arising in the finite difference approximations of the Laplacian on a rectangle. It has tridiagonal $m \times m$ blocks on the main diagonal and the codiagonals located m columns and m rows apart from the main diagonal. This is a very common test example with known analytical solution and makes a good illustration of the correctness of our routines. Test results for this example are summarized in Table 3: for both sets of parameters the original problem was satisfied to 16 digits and the orthogonality condition to 15 digits of precision.

Table 3. Error estimates for the eigenvectors of the matrix P with $c_0 = -0.33$, $c_1 = -0.17$, corresponding to the eigenvalues $\tilde{\lambda}$ computed with a maximum absolute deviation $\Delta(\lambda) =$ 6.6613e−16 from the exact eigenvalues λ by the Godunov Inverse Iteration

Parameters	err_1	err_2
$n = 225$, $c_2 = 1$	4.6356e−16	3.2918e−15
$n = 400$, $c_2 = 0.5$	5.2595e−16	4.8720e−15

4. The parallel implementation of the Godunov Inverse Iteration

The present-day implementations of eigensolvers for large sparse matrices are usually designed for distributed memory multicomputers which exchange data through message passing. The parallel implementation of the Godunov algorithm on such an architecture should not pose any challenges – the Godunov eigenvectors can be computed independently in only $11n$ floating point operations per eigenvector. By distributing a copy of two vectors containing, respectively, elements of the main diagonal and the codiagonal of the tridiagonal matrix to all p precesses and the sets of k eigenvalues for which $n = pk$ eigenvectors are desired, we can independently compute all the eigenvectors. The Godunov Inverse Iteration, on the other hand, requires a special treatment of the inverse iteration. On each step of the inverse iteration we need to solve a linear system of equations with tridiagonal matrix and perform the MGS reorthogonalization, which can bring up operation count to $O(n^3)$ in the worst case. Clearly, the inverse iteration is the most expensive step of the Godunov Inverse Iteration, and a parallel version of the algorithm will greatly depend on efficient implementation, as well as on optimal integration of the Godunov step. There are two ways to implement the Godunov Inverse Iteration – one is a truly parallel implementation that targets the efficient implementation of the inverse iteration when each processor gets a part of the matrix, and all eigenvalues, first making a contribution to the computation of a part of each of the Godunov's eigenvector, followed by the inverse iteration step, were again a tridiagonal solver and the MGS will act on the assigned part of the matrix in parallel. The other way is to sequentialize the algorithm, by distributing subsets of eigenvalues corresponding to the desired eigenvectors and vectors containing diagonals of the tridiagonal matrix to the processors. Neither scheme is superior for all the cases. One approach can be chosen on the other, depending on the size of a problem, the cost of having idle processes on the reorthogonalization step and the cost of data transposition in the sequentialized implementation on one hand, and the cost of extra communication overhead during the Sturm sequence and the Godunov vector computations in the truly parallel implementation.

5. Conclusions

The Godunov method for real symmetric matrices produces accurate eigenvector approximations, but usually these vectors have fewer digits of precision than eigenvectors computed according to some of the inverse iteration implementations. Designed for unreduced matrices, in the finite precision

the Godunov method produces almost colinear eigenvectors corresponding to closely clustered and computationally coincident eigenvalues. In the absence of directed rounding the Godunov method may produce non-numeric output. At the same time the choice of the initial vector in the Inverse Iteration algorithms does not ensure that the starting vector has a nontrivial component in the direction of the solution, and the algorithms do not always converge. The inverse iteration is very sensitive to the accuracy of the shift – we show that for eigenvalues computed by the bisection method with ensured accuracy in the order of machine precision the LINPACK and the EISPACK inverse iteration algorithms may break down. The Godunov Inverse Iteration algorithm was designed to solve these problems. Changing any non-numeric components of the Godunov eigenvectors to uniformly distributed random numbers, we apply the inverse iteration to those vectors, which usually attain the desired error bounds in one step, in contrast with other implementations of the inverse iteration algorithm which require a few more steps to achieve the same accuracy. The Godunov Inverse Iteration is fairly robust with respect to the choice of the inverse iteration shift – we use the right-hand bounds of the eigenvalue intervals computed by the bisection method as extremely accurate shifts in the Godunov Inverse Iteration. We resort to reorthogonalization within the iteration only in cases of computationally coincident or closely clustered eigenvalues. As the result, the Godunov Inverse Iteration algorithm produces accurate and robust solutions to the symmetric eigenvalue problem with higher accuracy than the Godunov method and in a fewer steps than the existing implementations of the inverse iteration algorithm.

References

[1] Godunov S.K., Antonov A.G., Kiriljuk O.P., Kostin V.I. Guaranteed Accuracy in the Solution of SLAE in Euclidean Spaces. – Novosibirsk: Nauka, 1988 (in Russian).

[2] Anderson E., Bai Z., Bischof C., et al. LAPACK Users' Guide: Second edition. – Philadelphia: SIAM, 1995.

[3] Smith B.T., Boyle J.M., Dongarra J.J., et al. Matrix eigensystem routines – EISPACK guide / Lecture Notes in Computer Science. – Berlin: Springer-Verlag, 1976. – Vol. 6.

[4] Wilkinson J.H. The Algebraic Eigenvalue Problem. – Oxford: University Press, 1965.

[5] Fernando K.V. Accurately counting singular values of bidiagonal matrices and eigenvalues of skew-symmetric tridiagonal matrices // SIAM J. on Matrix Analysis and Applications. – 1998. – Vol. 12, № 3. – P. 373–399.