Using the Field of Values for Pseudospectra Generation

Thierry Braconnier, R. Alan McCoy, Vincent Toumazou

CERFACS, 42 avenue G. Coriolis, 31057 Toulouse, France.

CERFACS Technical Report TR/PA/97/28

September 17, 1997

Keywords: Field of values, pseudospectra, Gershgorin disks

Abstract

The generation of spectral portraits requires the specification of a region in the complex plane which must be pre-determined. For cases where little is known a priori about the spectral properties of a matrix, an accurate and efficient means to compute the outer boundaries of the pseudospectra is desirable. It is shown that the field of values is an effective way to determine the extremities of the pseudospectrum; and that it provides advantages over other possible techniques such as Gershgorin disks, matrix norms, and eigenvalue solvers. Numerical comparisons of the different approaches are presented.

1 Introduction

Eigenvalues for nonnormal operators are known to be highly sensitive to perturbations. In some applications, the acute sensitivity of the eigenvalues may be a reflection of physical properties [12]. However, the numerical computation of eigenvalues for nonnormal matrices can be problematic; reliable eigenvalue methods may produce solutions significantly different from those obtainable in exact arithmetic [5]. Whether the perturbations come from finite precision arithmetic, physical approximations, or any other source, observing perturbation effects on eigenvalues can reveal much about the problem at hand.

One idea is to examine the eigenvalues of $A+E$, where E is a small perturbation on A. The ε-pseudoeigenvalue and ε-pseudospectrum are defined as:

- λ is an ε-pseudoeigenvalue of A if

$$\lambda \text{ is an eigenvalue of } A+E \text{ with } \|E\|_2 \leq \varepsilon \|A\|_2$$

- The ε-pseudospectrum of A is defined by

$$\Lambda_\varepsilon(A) = \{z \in \mathbb{C};\, z \text{ is an } \varepsilon\text{−pseudoeigenvalue of } A\}$$

For a fixed ε, the contour of $\Lambda_\varepsilon(A)$ can be defined as $\{z \in \mathbb{C};\ \|A\|_2 \|(A-zI)^{-1}\|_2 = \varepsilon^{-1}\}$. The graphical representation of

$$z \mapsto \log(\|A\|_2\ \|(A-zI)^{-1}\|_2) \tag{1}$$

is called the *spectral portrait* of A [6]. A logarithm is used to accommodate the wide range of values obtained near the eigenvalues of A, which are of course poles of the resolvent $(A-zI)^{-1}$. A spectral portrait thus provides an insightful visual representation of eigenvalue instability.

To generate a spectral portrait, the following steps can be taken:

1. Determine a region $\Re$ of interest in $\mathbb{C}$.
2. Discretise the region $\Re$.
3. For each point z of the discretisation, compute the value $\|(A - zI)^{-1}\|_2$.
4. Use a visualisation tool to display the computed values for $z \in \Re$.

We briefly mention that the computation of $\|(A - zI)^{-1}\|_2$ (Step 3) can be done in several ways. One could use a singular value decomposition, since $\|(A - zI)^{-1}\|_2 = \sigma_{\min}^{-1}(A-zI)$, where $\sigma_{\min}(A-zI)$ is the smallest singular value of $A-zI$. However, this approach is not suitable for large matrices. For such cases, continuation methods have been proposed [9, 3] ; while techniques using iterative eigensolvers (Lanczos type methods) have been developed [4, 11, 2].

Since all our definitions are taken relative (to ε), the level curves of (1) can be viewed as the level curves of the backward error associated with the computed eigenpairs and a spectral stability study can be done using the pseudospectra with small values of ε, i.e. small perturbations, compared to the machine precision ε_M. One of the consequences of this choice is that our pseudospectra code is implemented with a shift and invert technique and a shift equal to 0. It is more efficient for finding eigenvalues of largest magnitude. We then compute the largest eigenvalue of $[(A - zI)^*(A - zI)]^{-1}$ and thus the efficiency of this implementation increases when ε is small. In practice, it has been observed that due to the structure of the spectrum of $[(A - zI)^*(A - zI)]^{-1}$, the number of iterations for the Lanczos process increases with ε. The specifics we asked for the computed region $\Re$ (Step 1) are the following.

- $\Re$ should contain the eigenvalues of the matrix A.
- $\Re$ should contain the pseudospectra for small values of ε because they are more important in spectral stability studies.
- $\Re$ should not contain large regions in the complex plane where (1) is small, i.e., ε is large. The main reasons are that
  1. there regions are not interesting for spectral stability studies, and
  2. the shift and invert technique used will not be efficient on these nodes of the grid because the fact of shifting near 0 clusterizes the eigenvalues instead of separating them. Then, the code spends a lot of time determining values on areas which are not interesting.

The goal of this paper is to examine several methods for selecting the region which contains the values of interest for a spectral portrait. Among the approaches consisting of the QR algorithm (QRA), Gershgorin disks (GD), matrix norms (MN), or the field of values (FOV), we show that the region obtained by the FOV strategy provides the best overall solution in terms of pertinent global information and total computational costs. Using FOV as a first step, one obtains a region which contains all pseudospectral values of significant magnitude. If the application calls for it, the location of certain interesting eigenvalue subsets can thus be easily obtained. Furthermore, the region determined by FOV does not include extraneous areas where the pseudospectral values are small; this can lead to large computational savings.

We next describe the various approaches for determining an initial region in $\mathbb{C}$ for the spectral portrait. Numerical comparisons of each method are presented using matrices with different spectral properties in Section 4. Throughout the paper, $B(c, r)$ corresponds to the disk centered at c and radius r; and $p = p_{\min}\colon p_{\text{step}}\colon p_{\max}$ means that p takes values between $p_{\min}$ and $p_{\max}$ with a stepsize equal to p_{step}. If $p_{\text{step}} = 1$, we use the notation $p = p_{\min}\colon p_{\max}$. Finally, $\Lambda(A)$ denotes the exact spectrum of A.

2 Ways of Computing the Region $\Re$

In this section, we describe the four ways we have chosen for computing the region $\Re$ of the complex plane where we shall perform the computation of (1).

2.1 Eigenvalues Computation

The most obvious way to estimate the region is to compute the eigenvalues of the matrix A using a backward stable algorithm, for example the QR algorithm. There are two problems using this method. First, the QR algorithm is very time consuming for large matrices and does not use the possible sparsity of the original matrix A. And, if the matrix has a special pattern (for example it is upper triangular), the QR algorithm implemented with a multishift technique only uses the 1×1 or 2×2 diagonal blocks of A. This algorithm does not take into account the nonnormal part of the matrix given by the strictly upper triangular part of A [7] and only the exact spectrum is computed. For such matrices, we have made the following choice. The region is defined by

$$[\min_{i=1:n} \Re e(\lambda_i), \max_{i=1:n} \Re e(\lambda_i)] \times [-d, d]$$

where $d = \max_{i=1:n-1} |\lambda_{i+1} - \lambda_i|$. In this way, we can see if the two successive, more separated eigenvalues computed by QR are linked by finite precision computation. If the spectrum is not sensitive then the pseudospectra corresponding to either small or large values of ε are contained in the interval. And if the spectrum is sensitive then the pseudospectra corresponding to small values of ε are contained in the interval. Throughout the paper, the regions obtained by QR marked as $[a, b] \times [c, d]^*$ are computed using this choice.

2.2 Computation using $\|A\|_2$

Since $\rho(A) := \max_{i=1:n} |\lambda_i| \leq \|A\|_2$, the spectrum of the matrix A is contained in $B(0, \|A\|_2)$. $\|A\|_2$ can be easily computed using a Lanczos type algorithm. But, $B(0, \|A\|_2)$ can be much larger than the convex hull of the spectrum and $\Re$ can not satisfy our specifics.

2.3 Computation using Gershgorin Disks

The Gershgorin theorem tells us that

$$\Lambda(A) \subset \cup_{i=1:n} B(a_{ii}, \sum_{i \neq j} |a_{ij}|),$$

where $A = (a_{ij})_{i,j=1:n}$. A rectangle including this union of disks can be easily computed and this computation becomes more and more efficient with the sparsity of A. On the other hand, it can lead to a very rough estimate of the desired region because the rectangle can be much larger than the convex hull of the spectrum.

2.4 Field of Values

The field of values of a $n \times n$ complex matrix A is the set of all Rayleigh quotients

$$F(A) = \left\{ \frac{z^* A z}{z^* z} : 0 \neq z \in \mathbb{C}^n \right\}.$$

It can be shown that $F(A)$ is a convex set and when the matrix A is normal then it is the convex hull of the spectrum of A.

An algorithm has been proposed for computing the boundary of $F(A)$ [8]. It is based on the following fact. Any matrix A can be written as $A = A_H + A_K$, where $A_H = (A + A^*)/2$ is the Hermitian part of A and $A_K = (A - A^*)/2$ is the skew-Hermitian part of A. For any $z \in C$,

$$z^* A z = \underbrace{z^* A_H z}_{real} + \underbrace{z^* A_K z}_{pure\ imaginary} ,$$

which implies that

$$\min \{ \Re e(w) : w \in F(A) \} = \min_{z \neq 0} \frac{z^* A_H z}{z^* z} = \lambda_{\min}(A_H), \quad (2)$$

$$\max \{ \Re e(w) : w \in F(A) \} = \max_{z \neq 0} \frac{z^* A_H z}{z^* z} = \lambda_{\max}(A_H). \quad (3)$$

The intersection points of the vertical stripe defined by $\lambda_{\min}(A_H)$ and $\lambda_{\max}(A_H)$ and $F(A)$ are given by the Rayleigh quotients with A of the eigenvectors associated with $\lambda_{\min}(A_H)$ and $\lambda_{\max}(A_H)$. Since $F(e^{i\theta}A) = e^{i\theta}F(A)$, we can apply the same reasoning to the rotated matrix $A(\theta) = e^{i\theta}A$: by computing $\lambda_{\min}(A_H(\theta))$ and $\lambda_{\max}(A_H(\theta))$ and the Rayleigh quotients with A of the corresponding eigenvectors we determine two lines making an angle θ with the real axis between which $F(A)$ lies, together with the intersection points of the lines with $F(A)$. The complete computation can be done using a Lanczos type algorithm [2]. For determining $\Re$, since $F(A)$ is a convex set we only need to perform the previous computation using two angles: $\theta = 0$ and $\theta = \pi/2$. Thus, for our computation we need to compute the extremum eigenvectors of two Hermitian matrices.

Three remarks can be made. First, the computation of $F(A)$ is a unidimensional computation while the computation of pseudospectra is a two-dimensional one. Thus, it is less time consuming. Second, $F(A)$ gives *global* information about the numerical behavior of the spectrum of A while the pseudospectrum can give more *local* information because it can be computed only on a region containing a part of the spectrum of A. Finally, we have the following relationship between $F(A)$ and the pseudospectrum.

Theorem 2.1 *For all ϵ we have*

$$\Lambda_\epsilon(A) \subset F(A) + B(0, \epsilon\|A\|_2).$$

Proof: For all $\epsilon > 0$, $z \in \Lambda_\epsilon(A)$ means that there exists x such that

$$(A + E)x = zx \quad \text{with } \|E\|_2 \leq \epsilon\|A\|_2.$$

Without lost of generality, suppose $\|x\|_2 = 1$. Multiplying by x^* leads to

$$z = x^* A x + x^* E x.$$

Moreover, we have

$$|x^* E x| \leq \|x^*\|_2 \|Ex\|_2 \leq \|E\|_2 \leq \epsilon\|A\|_2. \quad \square$$

Theorem 2.1 means that all the pseudospectra are contained in $F(A)$ augmented with a stripe of length $\epsilon\|A\|_2$. Thus, for small values of ϵ the region enclosing $F(A)$ gives a good approximation of the region enclosing the corresponding pseudospectra. This result has motivated our choice of using FOV for computing $\Re$.

3 Test Matrices

The first test matrix is the upper triangular $n \times n$ real matrix defined by

$$W_1 = \begin{pmatrix} n & \varepsilon_M & \dots & \varepsilon_M & n \\ & n-1 & \varepsilon_M & \dots & \varepsilon_M \\ & & \ddots & & \vdots \\ & & & 2 & \varepsilon_M \\ & & & & 1 \end{pmatrix}$$

of size $n = 100$, where ε_M is the machine precision. Its spectrum is $\Lambda(W_1) = \{i\}_{i=1:n}$ and since its eigenbasis condition number [1] is equal to 1.86, the matrix W_1 has a small spectral sensitivity.

The second test matrix is the bidiagonal $n \times n$ real matrix defined by

$$W_2 = \begin{pmatrix} n & n & & & \\ & n-1 & \frac{n}{2} & & \\ & & \ddots & \ddots & \\ & & & 2 & \frac{n}{n-1} \\ & & & & 1 \end{pmatrix}$$

of size $n = 100$. Its spectrum is $\Lambda(W_2) = \{i\}_{i=1:n}$. In finite precision, the elements of the upper-diagonal have the effect of clustering the largest eigenvalues. This part of the spectrum becomes unstable and the eigenbasis condition number is equal to $1.75\ 10^{10}$.

The third test matrix is taken from [1]. The **Condif** matrix comes from the discretisation using centered finite differences on a $m \times m$ mesh of the two-dimensional model convection-diffusion problem

$$-\Delta u(x, y) + p_1 \nabla_x u(x, y) + p_2 \nabla_y u(x, y) - p_3 u(x, y)$$

on the unit square $[0, 1] \times [0, 1]$, with Dirichlet boundary conditions. The coefficients p_1, p_2 and p_3 are real numbers. This is a $n \times n$ non-symmetric matrix where $n = m^2$, with a block banded pattern and a bandwidth of size m and the eigenpairs are known explicitly. Let $h = 1/(m + 1)$, $\beta = p_1 h$, $\gamma = p_2 h$ and $\sigma = p_3 h^2$. Then, the eigenvalues of the matrices are given by

$$\lambda_{kl} = 4 - \sigma + 2\sqrt{1 - \beta^2} \cos\left(\frac{k\pi}{m+1}\right) + 2\sqrt{1 - \gamma^2} \cos\left(\frac{l\pi}{m+1}\right) \text{ for } l, k = 1{:}\, m.$$

Suppose we take[2] $\rho = p_1 = p_2 = p_3$. Figure 1 shows the exact spectrum (marked as plus) and the computed spectrum using the QR algorithm (marked as circles) of **Condif** for $n = 100$ and $\rho = 1$.

Then, the eigenvalues of the matrix $A(n, \rho)$ satisfy the following properties

1. if $\rho \le (\sqrt{n} + 1)$ then all the eigenvalues are real and contained within the interval $[0, 8]$, and
2. as the mesh-size h decreases, i.e., the size n increases, the separation of all the eigenvalues decreases.

[1] Let X (resp. Y) be the right (resp. left) eigenvectors of A and suppose $Y^*X = I$ and $||X||_2 = 1$, then the condition number is defined by $\kappa = ||Y||_2$ [5].

[2] For such a choice, the spectrum lies in the right half plane.

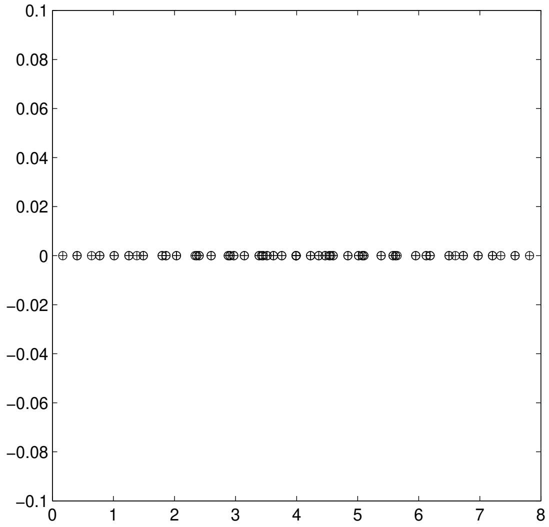

Figure 1: *Spectrum of* `Condif` *for* $n = 100$ *and* $\rho = 1$.

For the tests, we have used $n = 900$ and $\rho = 10$ which leads to a real spectrum and the corresponding matrix has an ill-conditioned eigenbasis because ρ is not very far from $\sqrt{n} + 1$. We shall see that the pseudospectra reveals the spectral sensitivity of `Condif`.

The last test matrix has been submitted by the laboratory LCR at Thomson-CSF and occurs in the simulation of nanometric recording. It is a complex symmetric (but non-Hermitian) matrix of size $n = 3886$ arising in a linear equation solved by the QMR method. As shown on Figure 2, the `Thomson3886` matrix has a banded structure and the computed spectrum using the QR algorithm is given by Figure 3.

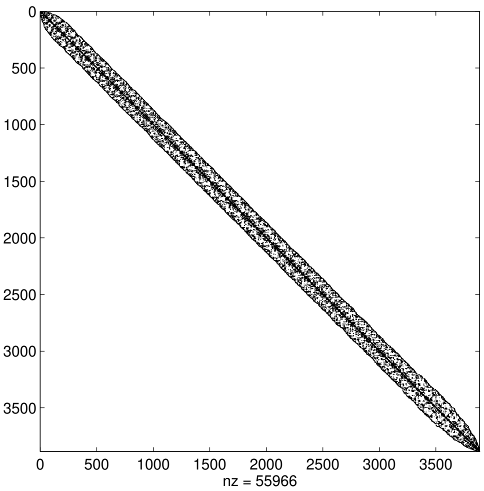

Figure 2: *Pattern of* `Thomson3886`.

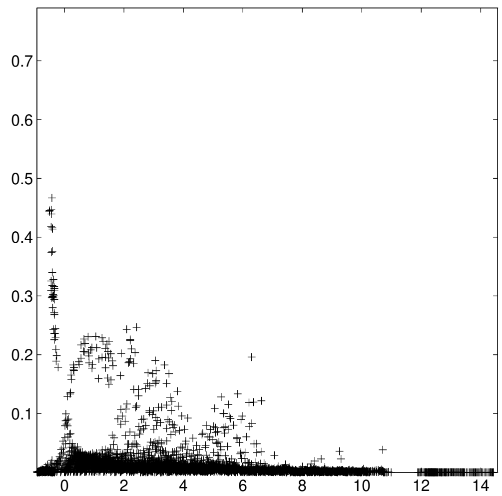

Figure 3: *Spectrum of* `Thomson3886`.

4 Numerical Results

In this section, we examine the results and compare the quality of the information provided by the pseudospectra as well as the corresponding computational time for the four strategies. The Figures show the level curves of the pseudospectra in a region computed by one of the above methods. The captions indicates the values of ε associated with the level curves where $\varepsilon = 10^{-p}$ and p varies in the specified range. Note that the axis of the plots, i.e., the limits of $\Re$, are given by the strategy used and can be very different from one strategy to another. We show that for a wide range of matrices, i.e.,

- matrix with small order having a small spectral sensitivity (W_1),
- matrix with small order having a large spectral sensitivity (W_2),
- large matrix having a large spectral sensitivity (`Condif`), and
- large matrix having a small spectral sensitivity (`Thomson3886`),

the region computed has the expected form and the pseudospectra computation is more efficient overall when FOV is used to determine $\Re$.

In all tables, the column labeled T_r gives the computational time for determining $\Re$ and the column labeled T gives the combined computational time (region+pseudospectra). The experiments on W_1, W_2 and `Condif` have been performed on a Silicon Graphic Power-challenge workstation using one $R10000$ processor and the experiments on `Thomson3886` have been done on a Sun Ultra-Sparc workstation. For both computers, the machine precision is $\varepsilon_M = 2.2\ 10^{-16}$.

4.1 Experiments on W_1

Table 1 gives the results and computational times for the pseudospectra computation of W_1 (on a 123×123 grid) using the four methods. The regions obtained for W_1 differ significantly for each method (see axis of Figures 4 to 7). The computation of $\Re$ using MN and GD takes the same time and are much faster than FOV. For the combined computation, QRA is very fast because no arithmetic operations are needed for reducing W_1 to Schur form. We can see that FOV and GD take the same time when QRA is much faster. MN is the slowest approach.

Method used	$\Re$	T_r (in sec)	T (in sec)
GD	$[-1.4\ 10^{-14}, 200] \times [0, 100]$	$1.3\ 10^{-3}$	176.1
FOV	$[-19.8, 120.8] \times [0, 50]$	2.62	198.3
MN	$[-141.4, 141.4] \times [-141.4, 141.4]$	$9.8\ 10^{-3}$	276.3
QRA	$[0.0, 101.0] \times [-1.0, 1.0]^*$	$2.0\ 10^{-1}$	166.4

Table 1: *Experiments on* W_1.

The resulting pseudospectra shown on Figure 4 to 7 allow us to understand such differences. We can see on Figure 5 that FOV gives a centered region on the interesting part of the complex plane. The area with large resolvent values is very small. GD determines a region larger than the one computed by FOV (see Figures 4 and 5). The computed region is larger using MN (see Figure 6). For this approach, many discritisation points give a large value of the resolvent. Since we use a shift and invert technique, the efficiency of the pseudospectra algorithm decreases for these points. For QRA, the combined computation is the fastest one. But the resulting pseudospectra

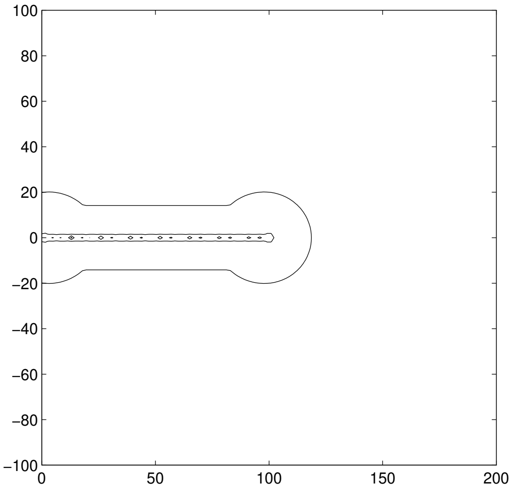

Figure 4: *GD.:* 10^{-p}, $p = 1{:}4$.

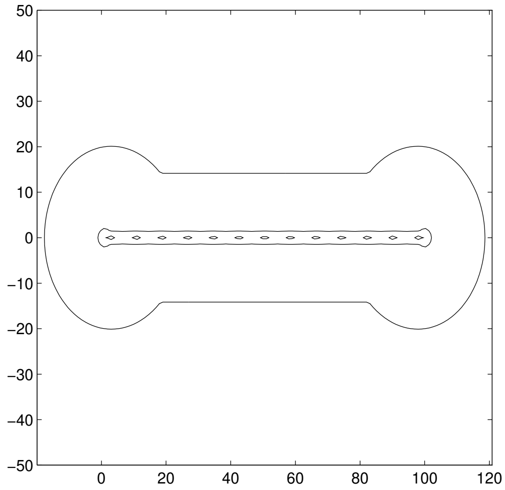

Figure 5: *FOV:* 10^{-p}, $p = 1{:}4$.

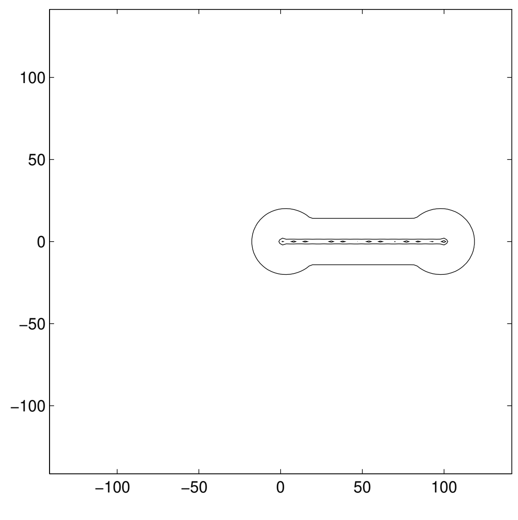

Figure 6: *MN:* 10^{-p}, $p = 1{:}4$.

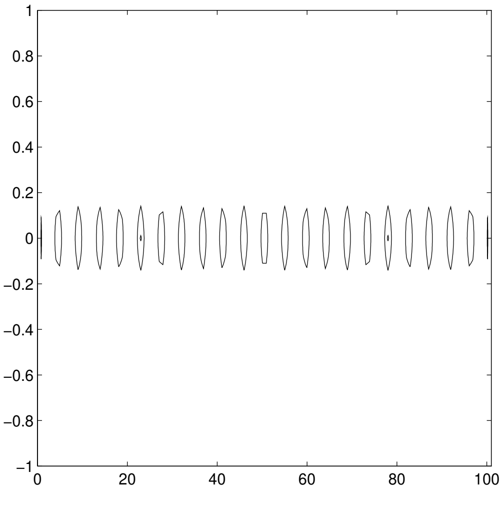

Figure 7: *QRA:* 10^{-p}, $p = 1{:}4$.

does not give all the information about the global instability of the spectra of W_1 as shown in Figure 7. Note that the pseudospectra (computed using FOV, GD and MN) confirm the small spectral sensitivity of W_1 because there is no region where the resolvent takes large values.

4.2 Experiments on W_2

Table 2 gives the results and computational times for the pseudospectra computation of W_2 (on a 123×123 grid) using the four methods.

Method used	$\Re$	T_r (in sec)	T (in sec)
GD	$[0, 200] \times [0, 100]$	$1.8\ 10^{-4}$	43.7
FOV	$[0.8, 155.7] \times [0, 56.4]$	$1.3\ 10^{-1}$	49.1
MN	$[-165.2, 165.2] \times [-165.2, 165.2]$	$1.5\ 10^{-3}$	110.9
QRA	$[0.0, 101.0] \times [-1.0, 1.0]^*$	$2.7\ 10^{-1}$	22.0

Table 2: *Experiments on* W_2.

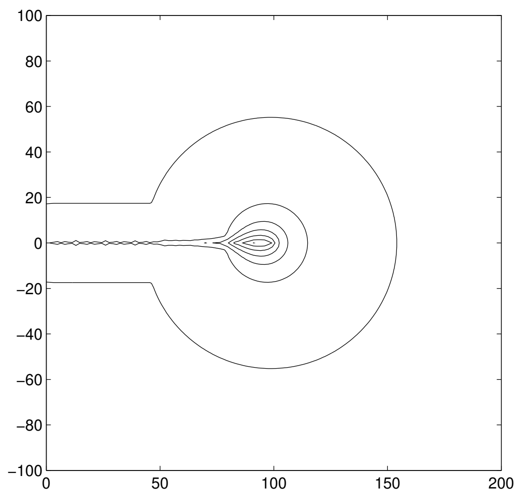

Figure 8: *GD:* 10^{-p}, $p = 1{:}2{:}13$.

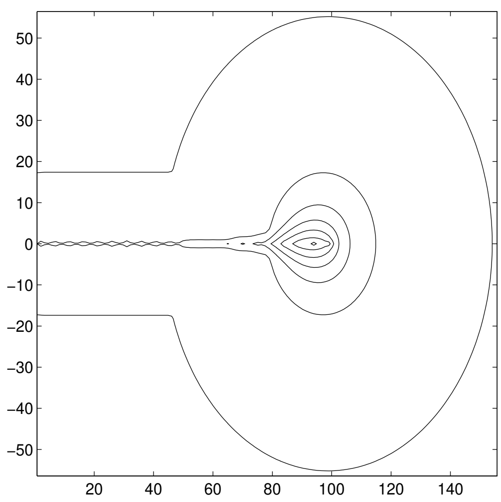

Figure 9: *FOV:* 10^{-p}, $p = 1{:}2{:}13$.

Concerning the shape of $\Re$ and the computational times, the same conclusions hold as for the experiments on W_1. Note that the pseudospectra (computed using FOV, GD and MN) confirm the large spectral sensitivity of W_2 because the resolvent takes large values on region far from the exact eigenvalues.

We have seen that for relatively small matrices having either a stable or an unstable spectra, GD and FOV approaches perform faster than the MN approach and slower than the QRA approach. But, FOV gives better information about the stability of the spectrum. Thus, FOV realizes the better trade-off between quality of information and combined computational time because the shift and invert technique is used. The next two test cases will show that for large matrices FOV performs better than the other approaches.

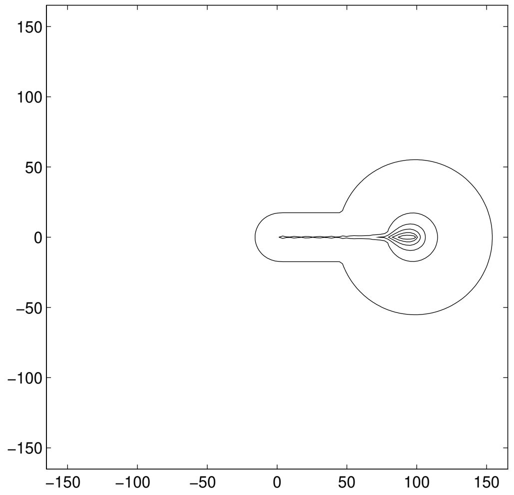

Figure 10: *MN.:* 10^{-p}, $p = 1{:}2{:}13$.

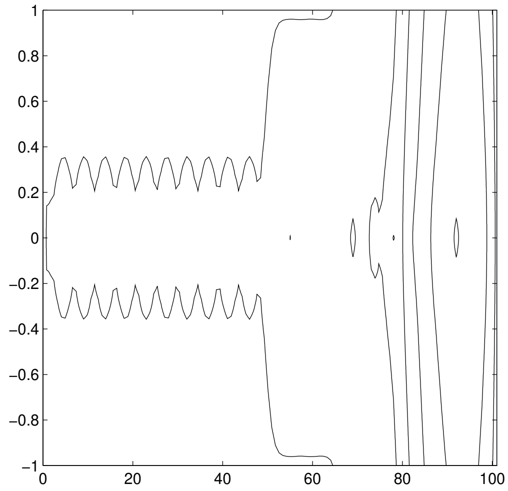

Figure 11: *QRA:* 10^{-p}, $p = 1{:}2{:}13$.

4.3 Experiments on `Condif`

Table 3 gives the results and computational times for the pseudospectra computation of **`Condif`** (on a 63×63 grid) using the four methods.

Method used	$\Re$	T_r (in sec)	T (in sec)
GD	$[-1.04\ 10^{-2}, 7.99] \times [0, 4.0]$	$2.03\ 10^{-3}$	655.9
FOV	$[1.01\ 10^{-2}, 7.94] \times [0, 1.28]$	1.28	521.4
MN	$[-7.97, 7.97] \times [-7.97, 7.97]$	$1.9\ 10^{-1}$	879.4
QRA	$[2.22\ 10^{-1}, 7.76] \times [-6.7\ 10^{-2}, 6.7\ 10^{-2}]^*$	613.97	1102.5

Table 3: *Experiments on* **`Condif`**.

Since **`Condif`** is not upper triangular, T_r for the QRA approach is now very large and T becomes prohibitive compared to the other methods. As with W_1 and W_2 for the combined computation, MN takes much more computational time and determines a region that is larger than FOV (see Figures 13 and 14). On this example, FOV performs better than GD because GD determines a region larger than FOV and on many points in this extra region the resolvent is large. Since the matrix is large, the shift and invert technique becomes very costly on these points.

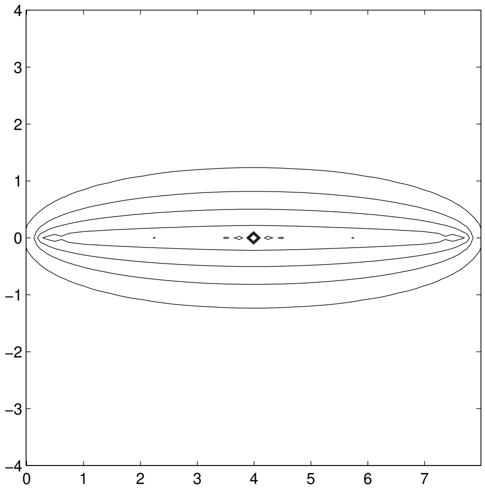

Figure 12: *GD:* 10^{-p}, $p = 2\colon 2\colon 10$.

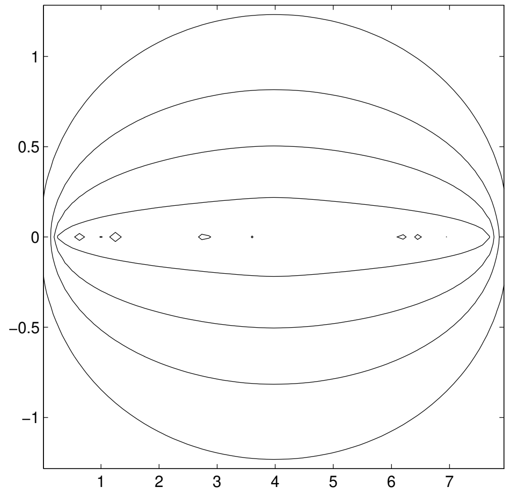

Figure 13: *FOV:* 10^{-p}, $p = 2\colon 2\colon 10$.

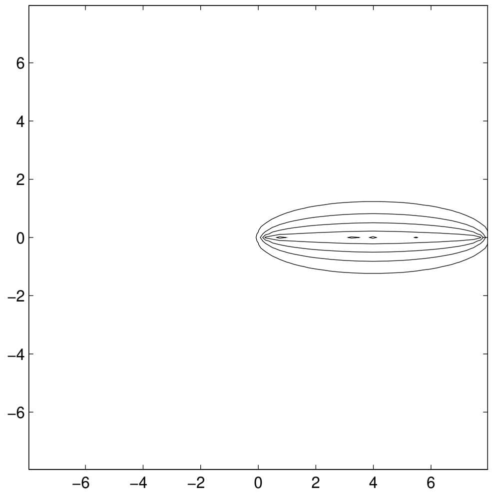

Figure 14: *MN:* 10^{-p}, $p = 2\colon 2\colon 10$.

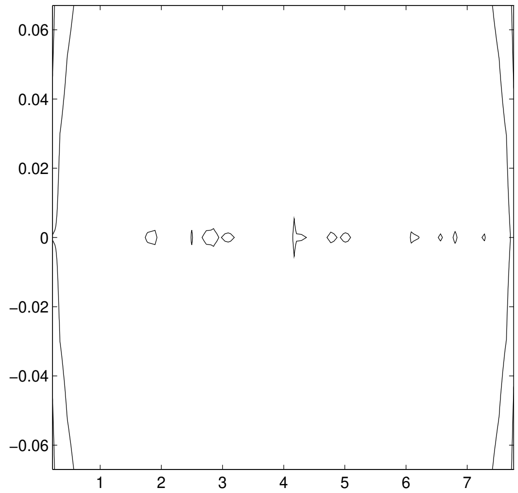

Figure 15: *QRA:* 10^{-p}, $p = 2\colon 2\colon 10$.

The pseusospectra computed using FOV and GD show the large sensitivity of `Condif`.

4.4 Experiments on `Thomson3886`

Table 4 gives the results and computational times for the pseudospectra computation of `Thomson3886` (on a 33×33 grid) using the four methods.

We can see in Figure 17 that the pseudospectra computed using FOV shows that the sensitive part of the spectrum of `Thomson3886` lies near the real axis. Once again, MN determines a region which is too large compared to FOV (see Figure 18) and the efficiency of the pseudospectra computation decreases when MN is used because of the shift and invert technique.

On this example, GD displays performance that is close to MN and that is much worse than using FOV. Thus, on such an example FOV gives better information about the spectral instability

Method used	$\Re$	T_r (in sec)	T (in sec)
GD	$[-9.76, 23.18] \times [-16.36, 16.36]$	$9.8\ 10^{-2}$	77294.6
FOV	$[-9.3\ 10^{-1}, 14.57] \times [1.4\ 10^{-9}, 7.9\ 10^{-1}]$	20.4	58581.1
MN	$[-14.6, 14.6] \times [-14.6, 14.6]$	1.45	85807.2
QRA	$[9.3\ 10^{-1}, 16.57] \times [7.0\ 10^{-7}, 4.7\ 10^{-1}]$	55316.3	114338.7

Table 4: *Experiments on* `Thomson3886`*.*

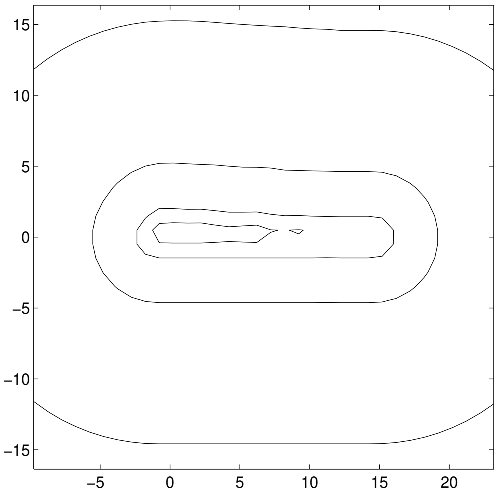

Figure 16: *GD:* 10^{-p}, $p = 0{:}0.5{:}1.5$.

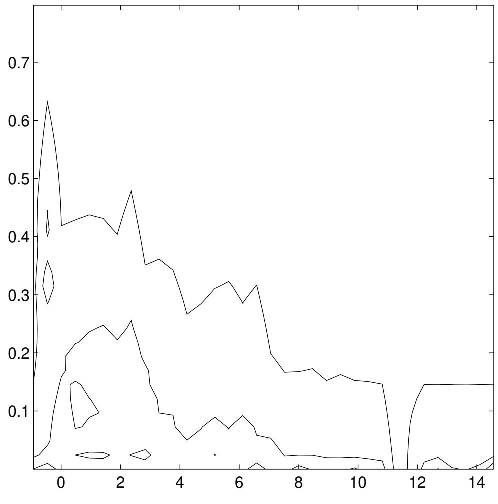

Figure 17: *FOV:* 10^{-p}, $p = 1{:}5$.

of `Thomson3886` and superior efficiency for the complete pseudospectra computation. On the other hand, as we can see in Figure 19, the eigenvalues computed by QR give a region very close to the one computed using FOV. But the complete computational cost is much higher than using FOV.

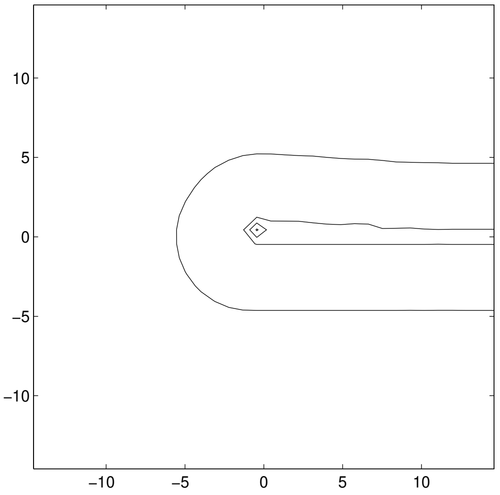

Figure 18: *MN:* 10^{-p}, $p = 0.5{:}1{:}3.5$.

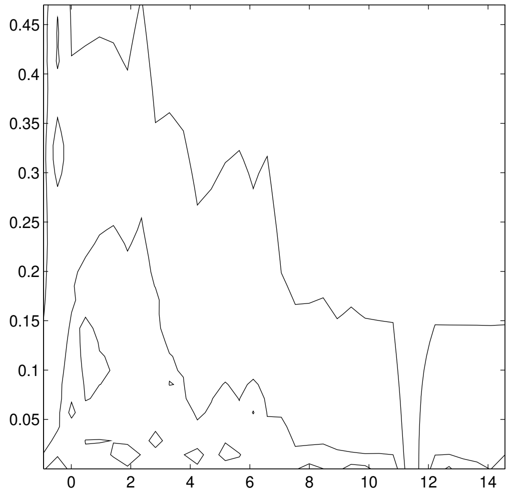

Figure 19: *QRA:* 10^{-p}, $p = 1{:}5$.

Note that $p_{\min}$ and $p_{\max}$ are different for Figures 16 to 19. For each method, these values are given by the range of computed values of $z \mapsto \log(\|A\|_2 \, \|(A - zI)^{-1}\|_2)$ for z in the computed region $\Re$. This example illustrates that the quality of the pseudospectra can be very different depending on the method used for computing $\Re$.

5 Conclusion

We have shown that for either relatively small or large matrices and that for matrices having either stable or unstable spectrum, the FOV approach offers the best trade-off between computational cost and quality of information for determining a region where the pseudospectra computation will take place. The region computed by the FOV approach does not include too many points where the resolvent is too small and then the Lanczos type method implemented with a shift and invert technique is much more efficient. If the user wants to refine the region where a spectral stability analysis will be carried out, FOV gives a good initial region where the user can zoom on the part of the spectrum in which he is interested.

Acknowledgments

This work has been partially supported by the HPCN european contract PINEAPL [10], where the FOV approach had been implemented in the pseudospectra code.

References

[1] Z. Bai. A collection of test matrices for the large scale nonsymmetric eigenvalue problem. Technical report, University of Kentucky, 1993.

[2] T. Braconnier and N. J. Higham. Computing the field of values and pseudospectra using the lanczos method with continuation. *BIT*, 36(3):422–440, 1996.

[3] M. Brühl. A curve tracing algorithm for computing the pseudospectrum. *BIT*, 36:3, 1996.

[4] J. F. Carpraux, J. Erhel, and M. Sadkane. Spectral portrait for non hermitian large sparse matrices. *Computing*, 53:301–310, 1994.

[5] F. Chaitin-Chatelin and V. Frayssé. *Lectures on Finite Precision Computations*. SIAM, Philadelphia, 1996.

[6] S. K. Godunov. Spectral portraits of matrices and criteria of spectrum dichotomy. In J. Herzberger and L. Atanassova, editors, *International symposium on computer arithmetic and scientific computation*, Oldenburg, Germany, October 1991. North-Holand.

[7] P. Henrici. Bounds for iterates, inverses, spectral variation and field of values of nonnormal matrices. *Numer. Math.*, 4:24–40, 1962.

[8] C. R. Johnson. Numerical determination of the field of values of a general complex matrix. *SIAM J. Num. Anal.*, 15(3):595–602, 1978.

[9] S. H. Lui. Computation of pseudospectra by continuation. *SISC*, 1995. To appear.

[10] PINEAPL. Parallel numerical applications and portable libraries, 1996. Fourth Framework Project: 20018, http://www.nag.co.uk/projects/PINEAPL.html.

[11] V. Toumazou. *Portraits spectraux de matrices : un outil d'analyse de la stabilité*. Ph. D. dissertation, Université Raymond Poincaré, Nancy I, CERFACS Tech. Rep. TH/PA/96/46, 1996.

[12] L. N. Trefethen, A. E. Trefethen, S. C. Reddy, and T. A. Driscoll. Hydrodynamic stability without eigenvalues. *Science*, 261:578–584, 1993.