Rational approximations of exp(x) for the calculation of rightmost eigenvalues

Gorik De Samblanx *Karl Meerbergen*
Adhemar Bultheel

Report TW 231, August 1995

Department of Computer Science, K.U. Leuven

Abstract

The article considers the finding of rightmost eigenvalues of a generalised eigenvalue problem $Ax = \lambda Bx$. A common technique is to apply Arnoldi's method or Subspace Iteration to the shift-invert transformation. This transformation implies the knowledge of a shift, which is in essence an approximation of the eigenvalue that is searched. To avoid the need for a shift, the article proposes the use of a rational approximation of the exponential function as a transformation for the problem matrices. This transformation has the advantage that it can enhance the convergence speed of the algorithms and that it is more robust compared to shift-invert with a possible inaccuratly chosen shift.
Keywords : sparse generalised eigenvalue problems, shift-invert, Chebyshev approximations, exponential function

1 Introduction

Suppose we want to solve the generalised eigenvalue problem

$$Ax = \lambda Bx,$$

where $A, B \in \mathbf{R}^{n \times n}$ are non-symmetric real matrices, and B is possibly singular; $\lambda \in \{\lambda_1, \ldots, \lambda_n\} \subset \mathbf{C}$ is an eigenvalue and $\mathrm{x} \in \{\mathrm{x}_1, \ldots, \mathrm{x}_n\} \subset \mathbf{C}^n$ is an associated eigenvector.

In many applications, for instance in stability analysis, we are not interested in the whole spectrum of eigenvalues, but only in the eigenvalue(s) with largest real parts, also called the 'dangerous' or leading eigenvalues. For these problems, A and B are sparse and n is large.

Common approaches to solve this problem are *shift-invert Arnoldi* [10] and *shift-invert Subspace Iteration* [7, 13]. These methods find eigenvalues near a shift point σ, where σ is an approximation for the rightmost eigenvalue. These methods are very sensitive to the choice of the shift. If the shift accidentally lies far from the rightmost eigenvalues, then it can happen that these eigenvalues are not computed. A practical shift selection strategy should take the possible loss of eigenvalues into account. This problem was studied by Meerbergen, Spence and Roose [9] for shift-invert subspace iteration.

Recently, other transformations have been proposed to find the rightmost eigenvalues of a matrix : the Cayley transform [6, 9] and the rational filter transform [8]. These approaches mainly use the shift-invert transformation as a building block for a more complex transformation.

In this paper, we present a rational transformation that maps the rightmost eigenvalues to well separated dominant eigenvalues and the other eigenvalues in a cluster near zero. The exponential function is such a transformation [3, 8]. We will show how a rational approximation of the exponential function can be used as a possible alternative for the shift-invert transformation.

It should be noticed that it is not the intention of this article to elaborate a theory on the approximation of e^A nor to derive a detailed error analysis, because we focus on the application of this function as a transformation for the eigenvalue problem. If the approximation error is small enough, i.e. reasonably smaller than the distance between the eigenvalues that must be distinguished, then the approximation will perform as good as an approximation that is even more accurate. Indeed, the convergence properties of the transformed problem depend on the approximated function, not on the approximation error.

The paper is built as follows : in Section 2, we show how a rational approximation of the exponential function can be used to find the rightmost eigenvalues. In Section 3 we propose three different approximations, each with different locality. Section 4 shows some experimental results that compare these methods with the shift-invert method. Section 5 concludes the paper with some general remarks.

2 The exponential transformation

Suppose that the eigenvalues λ_i are ordered by decreasing real part : $\mathrm{Re}(\lambda_i) \geq \mathrm{Re}(\lambda_{i+1})$ and suppose that the matrix B is non-singular. Consider the transformed eigenvalue problem

$$e^{\tau B^{-1}A}\mathrm{x} = \mu \mathrm{x}, \ 0 < \tau \in \mathbf{R}, \text{ where } \mu_i = e^{\tau \lambda_i}.$$

The matrix transformation $e^{\tau B^{-1}A}$ maps the eigenvalues with a large negative real part near zero and vertical lines in the complex plane to circles centered at the origin. Rightmost eigenvalues become dominant, independent of their imaginary part and their separation is in general guaranteed due to the exponential function. Unlike shift-invert, this property is independent of the precise location of the eigenvalues relative to a certain shift. Consequently, no knowledge about the spectrum is required in advance.

In our approach, we apply Arnoldi's method and subspace iteration to the matrix exponential. Since these iterative methods converge to well-separated dominant eigenvalues, the rightmost eigenvalues of $Ax = \lambda Bx$ are easily found by backtransformation of the dominant eigenvalues of $e^{\tau B^{-1}A}$.

The idea of using a rational approximation in different Krylov methods has also been proposed in [5, 12].

2.1 Use of the exponential transformation

It is clear that $e^{\tau B^{-1}A}$ should never be computed explicitly, because of the computational cost and because $e^{\tau B^{-1}A}$ is a full matrix, even for large sparse eigenvalue problems, for which the storage space is too expensive.

Actually, we do not need $e^{\tau B^{-1}A}$ in a closed form, as the solution methods only require the computation of the product $e^{\tau B^{-1}A} \cdot \mathrm{x}$ with $\mathrm{x} \in \mathbf{R}^n$ a vector..

To employ the exponential function as a transformation with a reasonable computational effort, we consider a rational approximation $r(\lambda)$ of $e^{\tau\lambda}$.

In practice, we can always find a region D that contains the rightmost eigenvalue. Such a region can e.g. be determined using Gerschgorin's theorem or any other way to bound the eigenvalues. Since we are not interested in eigenvalues outside D, we construct the rational function such that $|r(\lambda)|$ is small if $\lambda \notin D$.

So we try to find for a complex domain D, two polynomials $p(\lambda)$ and $q(\lambda)$ such that

$$e^{\tau\lambda} \approx r(\lambda) = \frac{p(\lambda)}{q(\lambda)}, \text{ for } \lambda \in D \subset \mathbf{C},$$

$$\text{with } \deg(p(\lambda)) = m, \ \deg(q(\lambda)) = n,$$

where

$$\epsilon(\lambda) = \left|e^{\tau\lambda} - r(\lambda)\right|, \ \lambda \in D \text{ and } 10^{\varepsilon} = \max\{\epsilon(\lambda) : \lambda \in D\}$$

are small. We only minimise the absolute size of the error value because otherwise for large domains D, the part where $\lambda \ll 0$ would be too much emphasised. Indeed, it is not possible to approximate $e^{\tau\lambda}$ by a rational function with a small relative error on an infinite subset of the complex plane. Moreover, we are only interested in accurate approximations for the rightmost eigenvalues.

We add two constraints to the approximation problem :

1. for $\lambda_i, \lambda_j \in D$ holds
$$\mathrm{Re}(\lambda_i) > \mathrm{Re}(\lambda_j) \Rightarrow |r(\lambda_i)| > |r(\lambda_j)|,$$
2. $\lim_{\lambda \to -\infty} r(\lambda) = 0 \Leftrightarrow n > m$.

Because $m < n$, and because we do not consider rational approximations with multiple poles, the approximation of $e^{\tau\lambda}$ has the partial fraction expansion

$$e^{\tau\lambda} = \sum_{i=1}^{n} \frac{\alpha_i}{\lambda - \beta_i}, \ \alpha_i, \beta_i \in \mathbf{C}$$

such that

$$e^{\tau B^{-1}A}\mathbf{x} = \sum_{i=1}^{n} \alpha_i (A - \beta_i B)^{-1} B\mathbf{x}, \ \alpha_i, \beta_i \in \mathbf{C}$$

can be computed as a weighted sum of n shift-invert operations.

Notes :

1. Because the exponential function is symmetric around the real axis, we can halve the amount of computational work. Indeed, if β_i is complex then

   $$\beta_i = \bar{\beta}_j \Leftrightarrow \alpha_i = \bar{\alpha}_j, \ 1 \leq i, j \leq n$$

   and

   $$\lambda \in \mathbf{R} \Rightarrow \frac{\alpha_i}{\lambda - \beta_i} + \frac{\bar{\alpha}_i}{\lambda - \bar{\beta}_i} = 2\mathrm{Re}\left(\frac{\alpha_i}{\lambda - \beta_i}\right)$$

   and similarly

   $$\alpha_i (A - \beta_i B)^{-1} B\mathbf{x} + \bar{\alpha}_i (A - \bar{\beta}_i B)^{-1} B\mathbf{x} = 2\mathrm{Re}(\alpha_i (A - \beta_i B)^{-1} B)\mathbf{x}$$

   which is a real vector if $\mathbf{x}$ is real.

   Hence from the set of non-real poles, we only have to consider the β_i with a positive imaginary part and we only have to factorise $(A - \beta_i B)$. We also can work with real vectors $\mathbf{x}$.

   This reduction of computational work can be seen as a trade off for the fact that using complex β_i makes full complex arithmetic necessary – this doubles the size of the problem.

2. The eigenvectors of the original problem are always eigenvectors of the rational matrix functions, but the inverse is not true. We expect that an eigenvector of $r(B^{-1}A)$ is an eigenvector of the original problem if the corresponding eigenvalue λ_i lies in D since

$$\lambda_i \neq \lambda_j \in D \Rightarrow r(\lambda_i) \simeq e^{\tau\lambda_i} \neq e^{\tau\lambda_j} \simeq r(\lambda_j).$$

So, if an approximate eigenpair $(\mu_i, \bar{\mathbf{x}}_i)$ of $r(B^{-1}A)$ is computed, then we can derive the original λ_i as

$$\lambda_i \approx \log \mu_i \text{ or } \lambda_i \approx \frac{\bar{\mathbf{x}}_i^T A \bar{\mathbf{x}}_i}{\bar{\mathbf{x}}_i^T B \bar{\mathbf{x}}_i},$$

where the latter is the Rayleigh quotient.

3. When B is singular, $A\mathbf{x} = \lambda B\mathbf{x}$ has an infinite eigenvalue with the nullspace of B as eigenspace. Since $\lim_{\lambda\to\infty} r(\lambda) = 0$, this eigenvalue is mapped to a zero eigenvalue of the transformed system and will not appear in the solution.

The computation of

$$\alpha_i(A - \beta_i B)^{-1}(B\mathbf{x})$$

does not imply the use of B^{-1}, so the possible singularity of B will not be a problem for the implementation. Moreover, the product $B\mathbf{x}$ will always filter out the part of $\mathbf{x}$ that points in the direction of a eigenvector associated with a zero eigenvalue of B [4, 11].

4. The transformation can only be computed if none of the β_i is an eigenvalue of $A\mathbf{x} = \lambda B\mathbf{x}$. If coincidentially β_i were an eigenvalue then the factorisation of $(A - \beta_i B)$ would raise an error.

2.2 An algorithm

The transformed Arnoldi procedure that we thus propose can be written as :

procedure Exponential_Arnoldi$(A, \mathbf{v}_1, \{\alpha_i, \beta_i \mid i = 1 \ldots n\}, N)$

for $k = 1$ **to** N **do**
  $\mathbf{v}_h \leftarrow \sum_{i=1}^{n} \alpha_i (A - \beta_i B)^{-1} B\mathbf{v}_k$
  Set $\bar{\mathbf{V}}_k = [\mathbf{v}_1, \cdots, \mathbf{v}_{k-1}, \mathbf{v}_h]$
  Compute $\bar{\mathbf{h}}_k \in \mathbf{R}^k$ such that if
  $\mathbf{V}_k \leftarrow [\mathbf{v}_1, \cdots, \mathbf{v}_{k-1}, \bar{\mathbf{V}}_k \cdot \mathbf{h}_k] \Rightarrow \mathbf{V}_k^T \mathbf{V}_k = I_k$
  Set $\mathbf{h}_k = [\bar{h}_{k,1} \cdots \bar{h}_{k,k}\ 0 \cdots 0]^T \in \mathbf{R}^N$
end for
Set $\mathbf{H} = [\mathbf{h}_1 \cdots \mathbf{h}_N]$
$\{\mu_1, \ldots, \mu_N\} \leftarrow$ *eigenvalues*$(\mathbf{H})$
$\mu = \mu_k$ where $|\mu_k| = \max\{|\mu_i|,\ i = 1 \ldots N\}$
$\lambda \leftarrow \log(\mu)$
$\mathbf{z} \leftarrow$ *corresponding_eigenvector*$(\mathbf{H}, \mu)$
$\mathbf{x} \leftarrow \mathbf{V}_N \cdot \mathbf{z}$
end procedure

3 Rational approximations to the exponential function

In this section we derive three different approximations of the product $e^{\tau B^{-1}A}\mathbf{x}$. First, we consider a rational Chebyshev approximation that has high local accuracy. Secondly, we propose an infinite Chebyshev approximation which spreads out this local quality over the real axis. The third approach tries to impose this property on the complex domain.

3.1 Chebyshev approximation on $[0,1]$

The first approximation that we propose is based on the well known Chebyshev approximation of e^{λ} on $[0,1]$. We compute

$$e^{\tau\lambda} \simeq r(\lambda)$$

where

$$r(\lambda) = \min_{r(\lambda)\in\Pi_{n-1,n}} \max_{\lambda\in[0,1]} \left|e^{\tau\lambda} - r(\lambda)\right| = 10^{\epsilon}.$$

$\Pi_{m,n}$ denotes the class of rational functions $p_m(\lambda)/q_n(\lambda)$ where $\deg p_m(\lambda) = m$, $\deg q_n(\lambda) = n$.

$r(\lambda)$ is called a *minimax approximation* : we minimise the maximal error of the approximation over an interval. A rational approximation that minimises the maximum error over a certain domain is called a *Chebyshev* approximation [2].

A side effect of this property is the *equi-oscillation* of the error, which means that each extremum of the error function reaches this minimal error value. We can expect that this minimum will be reached $m+n+2$ times, as the rational function has $m+n+1$ degrees of freedom, thus can interpolate $m+n+1$ times :

$$\left|e^{\tau\lambda} - r(\lambda)\right| \sim 10^{\epsilon} T_{m+n+1}(2\lambda - 1), x \in [0,1],$$

with $T_{m+n+1}(x)$ the Chebyshev polynomial of degree $m+n+1$.

Location of the poles The poles of all exponential approximations are located near a parabola that lies around the real axis :

$$\mathrm{Re}(x) = \beta - \alpha \mathrm{Im}^2(x).$$

To be able to compare the different approximations, we parameterise this curve with two parameters :

- α, which says how narrow the region is that is surrounded by the poles
- β, which says how close the poles lie near the imaginary axis.

It is not true that the exponential function is well approximated in the whole domain surrounded by the poles, but rightmost eigenvalues that are located closer to the poles will also be found as if it were a normal shift-invert transformation. On the other hand, the approximation can only hold in a subset of the region surrounded by the poles, so any wanted eigenvalue must be captured in that subset :

$$\mathrm{Re}(\lambda) < \beta' - \alpha' \mathrm{Im}^2(\lambda),\ \beta' < \beta,\ \alpha' < \alpha.$$

If the rightmost eigenvalue fulfils this condition, it will be found.

Both a small α and β indicate a good transformation, because they imply a broad domain close to the origin.

The error function of this Chebyshev approximation ($n = 2$) on the real axis is shown in Figure 1. The error behaviour can typically be divided in two regions :

- $x \in [0,1]$, where the quality of approximation is very high ($\epsilon(\lambda) \leq 10^{\epsilon}$), and where the minimax property holds,
- $x \in (-\infty, 0)$, where

$$\lim_{\lambda\to-\infty} e^{\tau\lambda} = 0 = \lim_{\lambda\to-\infty} r(\lambda),\ \text{but}$$

$$\text{since } \lim_{\lambda\to-\infty} \frac{e^{\tau\lambda}}{r(\lambda)} = 0,\ \lambda \ll 0 \Rightarrow e^{\tau\lambda} \ll r(\lambda).$$

We define a and b as

$$a = \max_{(-\infty,0)} \log_{10}\left|r(\lambda) - e^{\tau\lambda}\right| = \log_{10}\left|r(b) - e^{\tau b}\right|,\ a, b \in \mathbf{R}.$$

This 'large' approximation error a can be a disadvantage; but it often is not, as long as the rational transformation is small outside D. In many problems we focus on the eigenvalues close to the imaginary axis.

Table 1 shows some typical values of the five parameters defined for this approximation, in comparison with the other approximation methods.

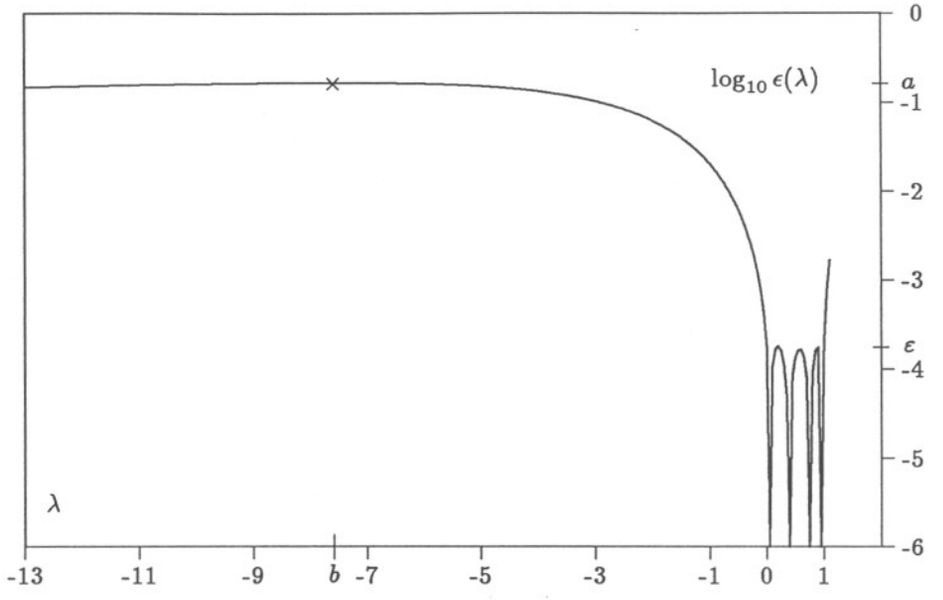


Figure 1: $\log_{10}\left|r(\lambda) - e^{\lambda}\right|$ for $n = 2$

	n	α	β	b	a	ε
$[0,1]$	3	0.10	4.2	-17.3	-1.0	-6.87
	4	0.07	5.5	-31.2	-1.1	-10.3
	5	0.06	6.8	-49.5	-1.2	-14.8
$(-\infty,0]$	3	0.20	1.4		-2.79	-2.79
	7	0.10	3.0		-6.66	-6.66
	10	0.07	4.2		-9.46	-9.46
$[0,1]\times i[-5,5]$	3	0.10	4.2	-17	-0.55	-4.0
	7	0.05	8.2	-71	-0.5	-10.6
	10	/	/	/	/	/

Table 1: Comparison between the transformations

3.2 Chebyshev Approximation on $(-\infty, 0]$

A possible drawback of the Chebyshev approximation on the interval $[0, 1]$ is that too much effort might be invested in approximating $e^{\tau\lambda}$ on $[0, 1]$ at the cost of a very bad approximation outside the unit interval. We can now try to shift and rescale that interval to an interval $[-c, 0]$ where $c > 1$. A big c will imply a larger approximation error ε inside the interval, but a better error a outside.

Moreover, we can chose $c = -\infty$ such that

$$a = \varepsilon \Rightarrow \epsilon(\lambda) \approx 10^{\varepsilon} T_{2n}\left(\frac{\lambda}{\lambda - 1}\right), \ \lambda \in [0, 1].$$

To achieve this, we use the rational Chebyshev approximation for the exponential function on the interval $(-\infty, 0]$:

$$10^{\varepsilon} = \max_{(-\infty,0)} \left| \frac{p_m(x)}{q_n(x)} - e^{\tau\lambda} \right| = \min_{\Pi_{m,n}} \max_{(-\infty,0)} \left| \frac{p_m(x)}{q_n(x)} - e^{\tau\lambda} \right| = \Lambda_{m,n}(e^{\tau\lambda}) \tag{1}$$

It is known that there is no solution in a closed form in the space $\mathrm{R}_{m,n}$ of rational functions that have no poles in the unit disk, but a nearby solution does exist [14].

The Chebyshev approximation over the negative part of the real axis for e^{λ} is very accurate. If n increases by 1 then the absolute error decreases by a factor of *at least* nine, which leads to the famous "1/9" conjecture in rational approximation theory.

Conjecture : ("1/9" Conjecture) [1]
Given $\Lambda_{m,n}$ defined by (1),

$$\lim_{n\to\infty} \Lambda_{n,n}^{1/n}(e^{\lambda}) = \frac{1}{9}.$$

The conjecture was proven to be wrong, the real value is about $1/9.289025\ldots$, implying that an approximation of degrees n will result in almost n correct digits. This is rather spectacular considering the fact that we approximate over an infinite interval.

Chebyshev Transformation on $(-\infty, 0]$: This Chebyshev approximation returns a very good result around the real axis without much loss of precision around the origin (compared to the Chebyshev approximation on $[0, 1]$, see Table 1).

The approximation domain lies closer to the imaginary axis, compared to the Chebyshev approximation on $[0, 1]$, but since α is larger, we would expect more difficulties in catching eigenvalues with a large imaginary part.

3.3 Chebyshev approximation on $D = [0, 1] \times i[-G, G]$

In the former subsection, we focussed on the amelioration of the approximation in the direction of the real axis. The result was good but still very restricted over the complex plane. This can be a problem when rightmost eigenvalues exist with large imaginary parts but small $\mathrm{Re}(\lambda_i) - \mathrm{Re}(\lambda_{i-1})$.

The exponential function is difficult to approximate on the imaginary axis, because the real and imaginary part of e^{λ} oscillate. We experienced that for each oscillation, at least two poles must be inserted :

$$e^{\lambda} = e^{Re(\lambda)}\left(\cos(\mathrm{Im}(\lambda)) + i\sin(\mathrm{Im}(\lambda))\right) \Rightarrow n > \frac{2G}{\pi} + 1 \Rightarrow n \text{ will be rather large.}$$

But if n gets too large, we get unwanted poles β_k to the left of the domain D which lie not on the parabolic curve. The maximum degree n to avoid these poles turns out do be around 7, depending on the value of G.

For our purpose, we do not need a closed form for the solution. We linearised the problem to the weighted $N \times (2n + 1)$ system

$$w(\mathrm{x})q(\mathrm{x})e^{\tau\mathbf{X}} - w(\mathrm{x})p(\mathrm{x}) = 0, \mathrm{x} \in D^N \text{ randomly chosen}, N \gg n$$

and solved it in a least square sense.

$w(\lambda)$ is a weight function chosen such that the error shows up some approximative equi-oscillatory behaviour :

$$\epsilon(\lambda) \sim 10^{\epsilon} T_{2n}(i\frac{\lambda}{G}).$$

This requirement was achieved by choosing

$$w(\lambda) = \frac{e^{-\lambda/2}}{\sqrt{\left(1 - \left(\frac{Im(\lambda)}{G}\right)^2\right)}}.$$

Complex Chebyshev Transformation Table 1 shows that this transformation can compete with the other transformations around the origin. In fact, the approximation in the imaginary direction is better, the parabola is wider (α small) but the poles lie on a relatively large distance (β larger) of the approximation domain. As can be expected, the approximation in the left half of the complex plane is not good at all.

3.4 Scaled exponential transformation

If we use these three transformations to approximate $r(\lambda) = e^{\lambda}$, we can always derive a rational approximation for $e^{\tau\lambda}$. If $\tau > 1$ then transformation gives faster convergence since $e^{\lambda_1 - \lambda_2} < e^{\tau(\lambda_1 - \lambda_2)}$. It appears that for large τ, the poles lie closer to the origin. The disadvantage is that the domain D shrinks. If $\tau = 1/|\lambda_i|$ then

$$\frac{\mu_i}{\mu_j} = e^{\tau(\lambda_i - \lambda_j)} \sim e^{-\lambda_j/\lambda_i}.$$

Thus scaling the function is useful when the rightmost eigenvalues are known to lie clustered (around the origin). The case where $1/\tau > 1$ is similar.

For the partial fraction expansion holds :

$$e^{\lambda} = \sum_{i=1}^{n} \frac{\alpha_i}{\lambda - \beta_i} \Rightarrow e^{\tau\lambda} = \sum_{i=1}^{n} \frac{\alpha_i}{\tau\lambda - \beta_i} = \sum_{i=1}^{n} \frac{\alpha_i/\tau}{\lambda - \beta_i/\tau},$$

so we only must divide the poles and coefficients by τ.

4 Some experimental results

Table 2 shows some experimental results on two rather small problems.

Example 1 The first problem is a Brusselator (dimension 200) with only a few complex eigenvalues which are the rightmost ones. They lie on a pitch-fork shaped curve that is shown in Figure 2.

The problem consists of the equations

$$\begin{aligned} \frac{\partial u}{\partial t} &= \frac{D_u}{L^2}\left[\frac{\partial^2 u}{\partial X^2} + \frac{\partial^2 u}{\partial Y^2}\right] - (B+1)u + u^2 v + C \\ \frac{\partial v}{\partial t} &= \frac{D_v}{L^2}\left[\frac{\partial^2 v}{\partial X^2} + \frac{\partial^2 v}{\partial Y^2}\right] - u^2 v + Bu \end{aligned}$$

for u and $v \in (0,1) \times (0,1)$ with homogeneous Dirichlet boundary conditions. They form a 2D reaction-diffusion model where u and v represent the concentrations of two reactants. The equations are discretised with central differences with grid-sizes $h_u = h_v = 1/(n+1)$ with $n = (N/2)^{1/2}$. For $x^T = [u_{1,1}, v_{1,1}, u_{1,2}, v_{1,2}, \ldots, u_{n,n}, v_{n,n}]$, the discretised equations can be written as $\dot{\lambda} = f(x)$. We wish to compute the rightmost eigenvalues μ_1 and μ_2 of the Jacobian matrix $A = \partial f/\partial x$ for $B = 5.45$, $C = 2$, $D_u = 0.004$, $D_v = 0.008$ and $L = 0.5$ for $u(X,Y) = B$ and $v(X,Y) = C/B$.

We applied a block Arnoldi process ($N = 10$ on a subspace $\mathbf{v}_1$ of dimension 2, the starting vectors were randomly chosen) combined with the shift-invert ($\sigma = 1.0$) transformation and with the three exponential transformations (of degree 2,4,8). The computations have been performed on a IBM SP2-processor.

In this case, the shift-invert transformation failed to find the correct rightmost eigenvalue ($0.1074 \pm 1.9017i$) – it found the second-rightmost eigenvalue. The result was correct when we iterated 20 times or when we changed the shift to $\sigma = 0.0$.

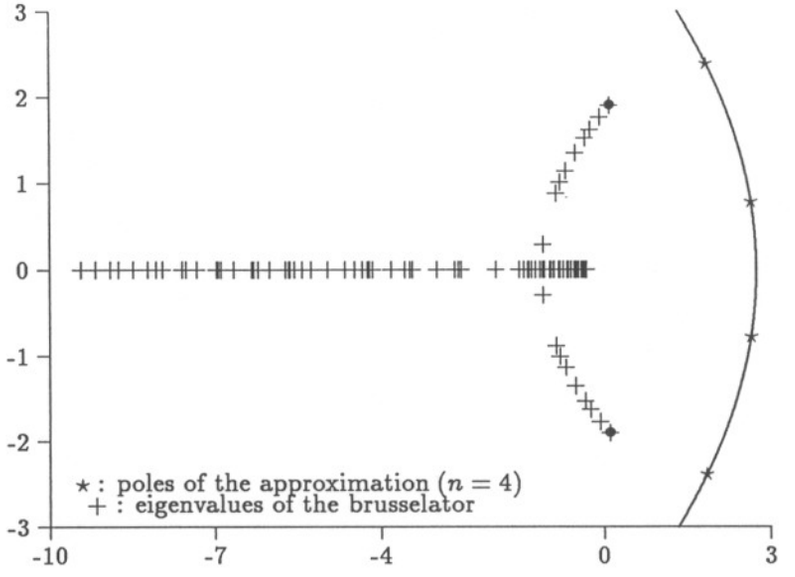


Figure 2: The Brusselator example

All the exponential approximations found the wanted eigenvalues all with about four digits. There was no increase in precision noticed when the degree of the rational function was raised. This is not surprising, since they all approximate the same function, they should return similar results.

Figure 2 shows for this problem the eigenvalues and the position of the poles of the Chebyshev approximation on $[0, 1]$ ($n = 4$).

Example 2 The second problem is an Olmstead model (dimension 100). The rightmost eigenvalues of this matrix lie close to the origin ($1.58\ 10^{-5} \pm 0.4470i$) but are not very well separated.

The model represents the flow of a layer of viscoelastic fluid heated from below. The equations are

$$\begin{cases} \dfrac{\partial u}{\partial t} &= (1-C)\dfrac{\partial^2 v}{\partial X^2} + C\dfrac{\partial^2 u}{\partial X^2} + Ru - u^3 \\ B\dfrac{\partial v}{\partial t} &= u - v \end{cases}$$

with boundary conditions $u(0) = u(1) = 0$ and $v(0) = v(1) = 0$. u represents the speed of the fluid and v is

		Brusselator $0.1074 \pm 1.917i$	Olmstead $1.58\ 10^{-5} \pm 0.4470i$
S.I.	$\sigma = 1.0$	$(0.2\ 10^{-3})$	$0.8\ 10^{-6}$
	$\sigma = 0.0$	$0.3\ 10^{-2}$	$0.5\ 10^{-12}$
$e^{2\lambda}$ on	$n = 2$	$0.5\ 10^{-4}$	$0.7\ 10^{-7}$
$[0, 1]$	$n = 4$	$0.3\ 10^{-7}$	
	$n = 8$	$0.9\ 10^{-5}$	
e^{λ} on	$n = 2$	$0.1\ 10^{-3}$	$0.1\ 10^{-6}$
$(-\infty, 0)$	$n = 4$	$0.1\ 10^{-2}$	
	$n = 8$	$0.1\ 10^{-3}$	
	$n = 2$	$0.1\ 10^{-4}$	$0.5\ 10^{-3}$ (scaled : $0.7\ 10^{-7}$)
e^{λ} on $[0, 1]\times$	$n = 4$	$0.1\ 10^{-4}$	
$i[-5, 5]$	$n = 8$	$0.3\ 10^{-4}$	

Table 2: Experimental results : displayed $\left|\bar{\lambda}_1 - \lambda_1\right|$

related to viscoelastic forces. The equation was discretised with central differences with gridsize $h = 1/(N/2)$. After discretisation the equation may be written as $\dot{\lambda} = f(x)$ with $x^T = [u_1, v_1, u_2, v_2, \ldots, u_{N/2}, v_{N/2}]$.

We applied here an Arnoldi process with $n = 10$ iterations and 1 starting vector $\mathbf{v}_1$.

The shift-invert transformation ($\sigma = 1.0$) easily found the solution, as did the exponential transformation. The exponential transformation reached a slightly higher level of accuracy. If we changed the shift to $\sigma = 0.0$ however, all three methods were outstripped by the shift-invert transformation.

In this example, the complex Chebyshev approximation turns out to converge very slowly because the poles lie much too far from the investigated area. If we scale the approximation with a factor 2, we find the solution with an accuracy corresponding to the other values in Table 2. This example shows that the exponential transformation also is sensitive to the distance between the eigenvalues and its poles. But this sensitivity is lower than the sensitivity to the position of only one shift.

5 Conclusion

Considering the fact that the exponential transformation roughly speaking doubles the computational work in comparison with the shift-invert transformation, the most important reason to use this transformation is the increased robustness. As we showed in the example, the exponential transformation searches the eigenvalues in a larger domain. In general, there is no need for a transformation with a very high degree ($n > 4$), unless we search for distinct eigenvalues with large imaginary parts. There is a need for more analysis on the approximation of $e^{\tau\lambda}$ in the complex plane if we want to make this transformation still more robust and more effective for the problems with such eigenvalues.

Acknowledgements

This research was supported by the National Fund for Scientific Research (NFWO), project Lanczos, grant #2.0042.93 and by the Human Capital and Mobility project ROLLS of the European Community under contract ERBCHRXCT930416.

References

[1] R. S. Varga A. J. Carpenter, A. Ruttan. Extended numerical computations on the "1/9" conjecture in rational approximation theory. *in Rational Approximation and Interpolation, Lecture Notes in Mathematics*, 1105:383–411, 1984.

[2] E.W. Cheney. *Introduction to approximation theory.* McGraw Hill, 1966.

[3] K.N. Christodoulou and L.E. Scriven. Finding leading modes of a viscous free surface flow : an asymmetric generalized eigenproblem. *J. Sci. Comput.*, 3:355–406, 1988.

[4] T. Ericsson. A generalised eigenvalue problem and the Lanczos algorithm. In J. Cullum and R.A. Willoughby, editors, *Large Scale Eigenvalue Problems*, pages 95–119. Elsevier Science Publishers BV, 1986.

[5] E. Gallopoulos and Y. Saad. Efficient solution of parabolic equations by Krylov approximation methods. *SIAM J. Sci. Statist. Comput.*, 13:1236–1264, 1992.

[6] T.J. Garratt, G. Moore, and A. Spence. Two methods for the numerical detection of Hopf bifurcations. In R. Seydel, F.W. Schneider, and H. Troger, editors, *Bifurcation and Chaos : Analysis, Algorithms, Applications*, pages 119–123. Birkhäuser, 1991.

[7] K. Meerbergen and D. Roose. Matrix transformations for computing rightmost eigenvalues of large sparse nonsymmetric matrices. Report TW206, Department of Computer Science, KU Leuven, 1994.

[8] K. Meerbergen and D. Roose. Finding eigenvalues to the right of a given line. In P.S. Vassilevski and S.D. Margenov, editors, *AMS Contemporary Mathematics Series.* IMACS Symposium on Iterative Methods in Linear Algebra, 1995. To appear.

[9] K. Meerbergen, A. Spence, and D. Roose. Shift-invert and Cayley transforms for detection of rightmost eigenvalues of nonsymmetric matrices. *BIT*, 34:409–423, 1994.

[10] R. Natarajan. An Arnoldi-based iterative scheme for nonsymmetric matrix pencils arising in finite element stability problems. *J. Comput. Phys.*, 100:128–142, 1992.

[11] B. Nour-Omid, B.N. Parlett, T. Ericsson, and P.S. Jensen. How to implement the Spectral Transformation. *Math. Comp.*, 48:663–673, 1987.

[12] Y. Saad. Analysis of some Krylov subspace approximations to matrix exponential operator. *SIAM J. Numer. Anal.*, 29:209–228, 1992.

[13] Y. Saad. *Numerical methods for large eigenvalue problems.* Algorithms and Architectures for Advanced Scientific Computing. Manchester University Press, Manchester, 1992.

[14] L. N. Trefethen. Rational Chebyshev approximation on the unit disk. *Numer. Math.*, 37:297–320, 1981.