Functions of structured matrices in numerical methods for ODEs

speaker: *Luciano Lopez*

Department of Mathematics
University of Bari, Italy

joint with: *N. Del Buono, R. Peluso, T. Politi, V. Simoncini*

Introduction

We are interested in numerical methods for evaluating

$$\exp(A)$$

and the product $\exp(\tau A)\, y$, when

$$A \in \mathbb{R}^{n\times n} \text{ is a square, real, sparse and large matrix,}$$

with a particular structure, i.e.

- skew-symmetric ($A = -A^T$);
- Hamiltonian ($A^T J = -JA$, with $J = \begin{pmatrix} 0 & I \\ -I & 0 \end{pmatrix}$);
- skew-symmetric and Hamiltonian.

$y \in \mathbb{R}^{n\times p}$ is a square (or rectangular) matrix which satisfies a geometric condition;

τ is a scaling factor which may be associated with the step size in a time integration method for ODEs.

We recall the following definitions:

The set of orthogonal matrices:

$$\mathcal{O}(n) = \left\{ Y \in \mathbb{R}^{n \times n} | Y \text{ non singular and } Y^T Y = I \right\}$$

The set of symplectic matrices:

$$\mathcal{S}(2n) = \left\{ Y \in \mathbb{R}^{2n \times 2n} | Y \text{ non singular and } Y^T J Y = J \right\}$$

The Stiefel manifold or the set of rectangular matrices with orthonormal columns:

$$\mathcal{S}(n,p) = \left\{ Y \in \mathbb{R}^{n \times p} | Y \text{ of rank } p \text{ and } Y^T Y = I_p \right\}.$$

Observe that:

- A skew-symmetric matrix $\Rightarrow \exp(A)$ orthogonal;
- A Hamiltonian matrix $\Rightarrow \exp(A)$ symplectic;
- A skew symmetric and Hamiltonian matrix $\Rightarrow \exp(A)$ ortho-symplectic;

Recall that the product of two orthogonal (resp. symplectic) matrices is again an orthogonal (resp. symplectic) matrix.

Main motivation of this study: construction of geometric numerical integrators for ODEs with invariants of orthogonal and symplectic type, for instance

ODEs evolving on the set of the orthogonal matrices;

ODEs evolving on the set of symplectic matrices;

ODEs evolving on the Stiefel manifold;

This kind of ODEs may arise, for instance, in

- the numerical computation of Lyapunov exponents of nonlinear dynamical systems;
- the numerical solution of advection-diffusion-reaction PDEs;
- the smooth QR decomposition of a matrix $A(t)$ depending on a parameter t.

Application to ODEs

Let $y(t)$ be the solution of the linear differential system

$$y' = A(t)y, \quad y(0) = y_0$$

.

Magnus's method provides

$$y(t) = \exp(\Omega(t))y_0.$$

where $\Omega(t)$ is a square matrix function satisfying a suitable ODE.

$A(t)$ skew-symmetric $\Rightarrow$ $\Omega(t)$ skew-symmetric $\Rightarrow$ $\exp(\Omega(t))$ orthogonal.

Then if $\quad y_0^T y_0 = I \Rightarrow y^T(t)y(t) = I$, for all $t > 0$.

Standard numerical methods are not structure-preserving.

Examples of structure-preserving methods are Magnus methods of 2nd and fourth order:

MG2	MG4
$A_n = A(t_n + \tau/2);$	$A_{n,1} = A(t_n + (\frac{1}{2} - \frac{\sqrt{3}}{6})\tau)$
$\omega_n = A_n;$	$A_{n,2} = A(t_n + (\frac{1}{2} + \frac{\sqrt{3}}{6})\tau)$
$y_{n+1} = \exp(\tau\omega_n)y_n$	$\omega_n = \frac{1}{2}(A_{n,1} + A_{n,2}) + \frac{\sqrt{3}}{12}\tau[A_{n,2}, A_{n,1}];$
	$y_{n+1} = \exp(\tau\omega_n)y_n$

with $t_n = t_0 + n\tau$.

The main computational requirement is that the numerical solution y_{n+1} must preserve the geometric behavior of the theoretical one.

This means that $\exp(\tau\omega_n)$ needs to be an orthogonal matrix at each n.

Methods in literature

Several methods may be found in literature to approximate the exponential matrix. Some of these may be splitted in

No structure-preserving methods:

- Padé and Chebyshev approximants;
- Arnoldi methods based on Krylov subspaces of dimension $m < n$ used to approximate $\exp(\tau A)y$ where y is a vector (see Hochbruk, Lubich, Moret, Simoncini);

Structure-preserving methods:

- Methods for approximating $\exp(\tau A)$ to a given order of accuracy with respect to τ. These methods are based on splitting techniques which exploit the structure of A (see Iserles and Celledoni).
- Methods based on the generalized polar decomposition of A. (see Iserles and Zanna, and Munthe-Kaas).
- All these methods have a cost of κn^3 flops where the constant κ increases with the order of the approximation.

The skew-symmetric case

We need to compute an approximation of

$$\exp(A)Y \text{ with } A = -A^T \text{ and } Y \text{ square orthogonal matrix.}$$

We may use a decomposition method based on two main steps:

- A is first reduced into a tridiagonal (and skew-symmetric) form H by using the tridiagonalization Lanczos process; at the end of this step we have $A = Q^T H Q$;
- then an effective Schur decomposition of H is obtained via the SVD of a bidiagonal matrix B of half size.

(see also Golub and van Loan book).

We observe that:

- The Lanczos process takes advantage from the possible sparsity of A due to the matrix-vector products involved.
- In floating point arithmetic, the Lanczos process provides H tridiagonal but the orthogonality of Q could be lost and a re-orthogonalization process could be required.
- For very large size problems, the storage of the columns of Q is the main drawback of this technique. In this case, the Lanczos process needs to be modified applying a storage procedure (see for instance Bergamaschi and Vianello).

The second step of the method

Suppose $A = Q^T H Q$ and n even integer.

In order to compute $\exp(H)$ we consider:

$$P = (e_1, e_3, \ldots, e_{n-1}, e_2, e_4, \ldots, e_n) \tag{1}$$

where $(e_1, e_2, \ldots, e_n)$ is the canonical basis of $\mathbb{R}^n$.

Then

$$P^T H P = \begin{pmatrix} 0 & -B \\ B^T & 0 \end{pmatrix}, \tag{2}$$

where B is a bi-diagonal square matrix of half size $w = \frac{n}{2}$.

Consider the SVD of B

$$B = U \Sigma V^T,$$

with $\Sigma = \mathrm{diag}(\sigma_1, \sigma_2, \ldots, \sigma_w)$ and $\sigma_1 > \sigma_2 > \ldots \sigma_w > 0$.

Then we can prove that:

$$\exp(A) = QPT(U,V,\Sigma)P^TQ^T$$

where:

$$T(U,V,\Sigma) = \begin{pmatrix} U\cos(\Sigma)U^T & -U\sin(\Sigma)V^T \\ V\sin(\Sigma)U^T & V\cos(\Sigma)V^T \end{pmatrix},$$

with

$$\cos(\Sigma) = \mathrm{diag}(\cos\sigma_1, \cos\sigma_2, \ldots, \cos\sigma_w),$$
$$\sin(\Sigma) = \mathrm{diag}(\sin\sigma_1, \sin\sigma_2, \ldots, \sin\sigma_w).$$

N. Del Buono & L. Lopez & R. Peluso, SISC 2005.

Flops count

When A is sparse, the main computational cost of this procedure is $\frac{35}{8}n^3$ flops, which should be compared with the ones of Matlab routines for matrix exponential which generally varies between $20n^3$ and $30n^3$ flops;

Instead, when A is a full matrix, the main computational cost is $\frac{51}{8}n^3$ flops.

Decay behavior

Although $\exp(A)$ is dense matrix, one can take computational advantages of the possible decay of entries of $\exp(H)$ away from the main diagonal. This behavior may be exploited in defining a banded approximation of $T(U, V, \Sigma)$.

Numerical comparisons

We have compared this approach (Matlab routine `AExp`) with the two Matlab functions

- `Expm` computing the exponential of A using a scaling and squaring algorithm with Padé approximations
- `Expm3` evaluating $\exp(A)$ via eigenvalues and eigenvectors decomposition.

Comparisons are done in terms of

- Flops (counted by Matlab 5.3 routine `flops`);
- Global error, defined as the 2-norm of the difference of `AExp` and `Expm`;
- Orthogonal error, defined as the distance of the computed exponential from the orthogonal manifold (i.e. $\|[\exp(A)]^T \exp(A) - I_n\|_F$, where $\|\cdot\|_F$ is the Frobenius norm on matrices)

Comparisons on sparse skew-symmetric matrices A of different dimensions n and entries randomly generated in $[-10, 10]$.

n	Method	Flops	Global error	Orthogonal error
50	`AExp`	901977	8.3755e-14	2.0214e-13
	`Expm`	2659508	-	1.7765e-14
	`Expm3`	5211760	-	2.6031e-14
100	`AExp`	7086541	8.2599e-14	1.9544e-13
	`Expm`	22970654	-	2.1794e-13
	`Expm3`	40787683	-	6.7479e-14
200	`AExp`	56085753	2.1696e-13	4.1762e-13
	`Expm`	182544884	-	8.7616e-14
	`Expm3`	318170737	-	1.8904e-13

Comparisons on full skew-symmetric matrices A of different dimensions n and entries randomly generated in $[-10, 10]$.

n	Method	Flops	Global error	Orthogonal error
50	`AExp`	1122393	8.0830e-14	1.6458e-13
	`Expm`	3409432	-	3.6161e-13
	`Expm3`	5244994	-	7.9323e-15
100	`AExp`	8863773	1.2236e-13	2.7711e-13
	`Expm`	26969568	-	5.9092e-13
	`Expm3`	41191787	-	1.6379e-14
200	`AExp`	70437373	3.0713e-13	6.2590e-13
	`Expm`	230548400	-	1.7157e-12
	`Expm3`	318310721	-	3.9922e-14

Hamiltonian and skew-symmetric matrices

Consider the case of $\mathcal{M}$ skew symmetric and Hamiltonian matrix :

$$\mathcal{M} = \begin{bmatrix} A & B \\ -B & A \end{bmatrix},$$

with $A \in \mathbb{R}^{n\times n}$ is a skew-symmetric matrix ($A^\top = -A$)

and $B \in \mathbb{R}^{n\times n}$ is symmetric (i.e., $B^\top = B$).

We start by analyzing the case in which $\mathcal{M}$ has the **special form**

$$\mathcal{M} = \begin{bmatrix} 0 & B \\ -B & 0 \end{bmatrix},$$

that is with $A = 0$.

The Schur decomposition of $\mathcal{M}$ may be derived by the decomposition

$$B = U\Lambda U^\top$$

U orthogonal and $\Lambda =$diag$(\lambda_1, \ldots, \lambda_n)$.

This decomposition may be obtained by the Lanczos process:

$$Q^\top BQ = T$$

with T symmetric tridiagonal matrix and Q orthogonal, then we may digonalize T

$$S^\top TS = \Lambda$$

with S orthogonal.

Finally computing the previous orthogonal matrix U as $U = QS$.

In floating-point arithmetic the columns of the matrix Q could progressively lose their orthogonality, hence a re-orthogonalization procedure could be required.

Hence,

$$\mathcal{M} = \begin{bmatrix} 0 & U\Lambda U^\top \\ -U\Lambda U^\top & 0 \end{bmatrix},$$

and we can show that:

$$\exp(\mathcal{M}) = \begin{bmatrix} U \ \cos(\Lambda) \ U^\top & U \ \sin(\Lambda) \ U^\top \\ -U \ \sin(\Lambda) \ U^\top & U \ \cos(\Lambda) \ U^\top \end{bmatrix}$$

where

- $\cos(\Lambda) = \mathrm{diag}(\cos(\lambda_1), \cos(\lambda_2), \ldots, \cos(\lambda_n))$
- $\sin(\Lambda) = \mathrm{diag}(\sin(\lambda_1), \sin(\lambda_2), \ldots, \sin(\lambda_n))$

If Y is ortho-symplectic then

$$Y = \begin{bmatrix} Y_1 & Y_2 \\ -Y_2 & Y_1 \end{bmatrix}$$

with the constrains:

$$Y_1^T Y_1 + Y_2^T Y_2 = I_n, \qquad Y_1^T Y_2 - Y_2^T Y_1 = 0.$$

If the matrix product

$$\exp(\mathcal{M})Y = \begin{bmatrix} U\cos(\Lambda)\,U^\top & U\sin(\Lambda)\,U^\top \\ -U\sin(\Lambda)\,U^\top & U\cos(\Lambda)\,U^\top \end{bmatrix} \begin{bmatrix} Y_1 & Y_2 \\ -Y_2 & Y_1 \end{bmatrix}$$

is required, then:

- We can avoid to compute the matrices $U\cos(\Lambda)\,U^\top$ and $U\sin(\Lambda)\,U^\top$ explicitly;
- Only the two blocks $(1,1)$ and $(1,2)$ in $\exp(\mathcal{M})Y$ need to be computed .

Splitting techniques

We now consider the general case:

$$\mathcal{M} = \begin{bmatrix} A & B \\ -B & A \end{bmatrix}, \qquad A \neq 0.$$

In the context of numerical methods for ODEs, splitting techniques are often used to reduce the cost of the exponential evaluation.

We may consider the following natural splitting

$$\begin{aligned} \mathcal{M} &= \quad \mathcal{M}_1 \quad + \quad \mathcal{M}_2 = \\ &= \begin{bmatrix} 0 & B \\ -B & 0 \end{bmatrix} + \begin{bmatrix} A & 0 \\ 0 & A \end{bmatrix} \end{aligned}$$

N. Del Buono & L. Lopez & T. Politi, to appear.

and to approximate the exponential map, we may apply:

- the first order accuracy approximation

$$\exp(\mathcal{M}) \cong \exp(\mathcal{M}_1)\exp(\mathcal{M}_2)$$

- or the Strang second order approximation scheme

$$\exp(\mathcal{M}) \cong \exp\left(\frac{1}{2}\mathcal{M}_2\right)\exp(\mathcal{M}_1)\exp\left(\frac{1}{2}\mathcal{M}_2\right).$$

- To compute $\exp(\mathcal{M}_2)$ effective methods for skew-symmetric matrices can be used;
- To compute $\exp(\mathcal{M}_1)$ the Schur decomposition method can be adopted;
- These splitting techniques preserve the geometric properties of the exponential, that is they provide matrices which are ortho-symplectic.

The general case

A general Hamiltonian and skew-symmetric matrix

$$\mathcal{M} = \begin{bmatrix} A & B \\ -B & A \end{bmatrix}$$

can be proved to be similar (by means of an ortho-symplectic matrix) to a canonical Hamiltonian and skew-symmetric matrix of the form

$$\begin{bmatrix} 0 & \Omega \\ -\Omega & 0 \end{bmatrix}$$

with Ω diagonal matrix.

However, this transformation method may be expensive and in the context of ODEs splitting techniques should be used.

Numerical Tests

Comparisons between the Matlab function `expm` and our procedure to compute $\exp(\mathcal{M})$ for matrices $\mathcal{M}$ with zero diagonal blocks (i.e., $A = 0$).

$2n$	Meth	Flops	Glob. err.	Orth. err.	Sympl. err.
50	O-Schur	486396	8.5197e-15	4.1587e-14	4.1587e-14
	`expm`	1651126	-	4.1412e-14	4.1412e-14
100	O-Schur	3555279	2.9592e-13	1.3000e-12	1.3000e-12
	`expm`	14950096	-	1.7453e-14	1.2249e-14
200	O-Schur	28459034	2.0328e-11	8.5270e-11	8.5270e-11
	`expm`	134396604	-	7.0984e-14	4.6904e-14
500	O-Schur	426743229	3.7416e-11	1.4595e-10	1.4595e-10
	`expm`	2.0898e+9	-	1.9512e-13	1.4210e-13

Computation of $\exp(\mathcal{M})$ in case of $\mathcal{M}$ in the general form (i.e. $A \neq 0$) by using splitting techniques, $(n = 200)$.

Meth	Flops	Glob. err.	Orth. err.	Sympl. err.
`Expm`	134548784	-	7.1416e-14	5.1095e-14
Splitting1	42828234	1.4855e-4	4.7570e-13	4.7570e-13
Splitting2	44827914	1.5752e-6	4.9231e-13	4.9231e-13

The rectangular orthogonal case

Suppose we need to compute an approximation of

$$Z = \exp(A)V$$

with A skew-symmetric;

V matrix of size $n \times p$ $(p << n)$ and with orthonormal columns.

We need a procedure which provides an approximation Z_m of Z with orthonormal columns.

Motived by the rectangular structure of V, we would like to apply Arnoldi approximations into Krylov subspces.

Case of $V = [v]$, $p = 1$ and $\|v\| = 1$.

An effective method is the Arnoldi approximation of $z = \exp(A)v$ using Krylov subspace:

$$\mathcal{K}_m \equiv \mathcal{K}_m(A, v) = \mathrm{span}\{v, Av, \dots, A^{m-1}v\}$$

V_m s.t. range$(V_m) = K_m(A, q)$ and $V_m^T V_m = I$

Arnoldi relation:

$$AV_m = V_m H_m + h_{m+1,m} v_{m+1} e_m^T$$

A common approach

$$\exp(A)v \;\approx\; z_m = V_m \exp(H_m) e_1, \qquad \|v\| = 1$$

and

$$\|v\| = 1 \Rightarrow \qquad \|z_m\| = 1.$$

Now, let $V = [v_1, \ldots, v_p]$ with orthonormal columns.

Regular Krylov subspaces $\mathcal{K}_m(A, v_i)$, $i = 1, \ldots, p$,

A skew-sym $\quad \Rightarrow \quad H_{m,i}$ skew-symmetric $\quad \Rightarrow \exp(H_{m,i})$ orthogonal

We may assume

$$\exp(A)v_i \quad \approx \quad z_{m,i} = V_{m,i}\exp(H_{m,i})e_1, \qquad i = 1, \ldots, p,$$

But it is not enough because

$\{z_{m,1}, \ldots, z_{m,p}\}$ are vectors of unit norm but not orthogonal vectors.

To preserve the orthonormal structure we need to use Block Krylov subspaces:

$$\mathcal{K}_m(A, V) = \mathrm{span}\{V, AV, \ldots, A^{m-1}V\}$$

A basis of $\mathcal{K}_m(A, V)$ is generated by the block Lanczos recursion:

$$A\mathcal{V}_m = \mathcal{V}_m\mathcal{H}_m + V_{m+1}h_{m+1,m}E_m^T$$

where:

- $\mathcal{V}_m = [V_1, \ldots, V_m] \in \mathbb{R}^{n\times mp}$ and $V_1 = V$,
- $\mathcal{H}_m$ is an $mp \times mp$ block tridiagonal and skew-symmetric matrix $\mathcal{H}_m = (h_{ij})$ with h_{ij} a $p \times p$ block,
- V_{m+1} is $n \times p$, $h_{m+1,m}$ is $p \times p$ and $E_m^T = [0, ..., 0, I_p]$.

L. Lopez & V. Simoncini, BIT 2006.

Then we have the following approximation

$$\exp(A)V \cong \mathcal{V}_m \exp(\mathcal{H}_m) E_1 \chi_0$$

where $\chi_0 \in \mathbb{R}^{p\times p}$ is such that $V = \mathcal{V}_m E_1 \chi_0$, and this approximation has orthonormal columns.

The rectangular symplectic case

Definition. *Let $Q \in \mathbb{R}^{2n \times 2p}$, we say that Q is a (rectangular) symplectic matrix if*

$$Q^T J Q = J_{2p}.$$

A Hamiltonian and Q symplectic $\Rightarrow$

$$Z = \exp(A)Q \text{ is still a rectangular symplectic matrix.}$$

We wish a symplectic approximation Z_m of Z.

In order to obtain Z_m we need a *symplectic* basis $\mathcal{V}_m$ of the subspace $\mathcal{K}_m(A, V)$ and a *Hamiltonian* representation $\mathcal{H}_m$ of A.

From Q we define the starting matrix V as

$$V = QP_1$$

with P_1 a suitable permutation matrix, so that

$$V^T JV = P_1 J_{2p} P_1^T, \tag{3}$$

Then V is symplectic upon permutation.

This permutation is commonly performed in the single vector case, i.e. for $p = 1$.

The algorithm proceeds by using the block Lanczos recurrence starting with V, that is

$$A\mathcal{V}_m = \mathcal{V}_m \mathcal{H}_m + V_{m+1} h_{m+1,m} E_m^T$$

and requiring the basis $\mathcal{V}_m$ to be symplectic upon permutation.

More precisely, the matrix $\mathcal{V}_m$ is constructed from the Lanczos recurrence with

$$(\mathcal{V}_m P_m)^T J (\mathcal{V}_m P_m) = J_{2mp}. \tag{4}$$

Moreover the matrix $P_m^T \mathcal{H}_m P_m$ will be Hamiltonian.

The approximation to $U = \exp(A)V$ is then given by

$$U_m = \mathcal{V}_m P_m \exp(P_m^T \mathcal{H}_m P_m)(\mathcal{V}_m P_m)^T JV,$$

which is equivalent to

$$U_m = \mathcal{V}_m \exp(\mathcal{H}_m)\mathcal{V}_m^T JV = \mathcal{V}_m \exp(\mathcal{H}_m) E_1 P_1 J_{2p} P_1^T,$$

and which is also symplectic upon permutation.

Stability problems and loss of symplecticity (or of rank) may destroy the Hamiltonian structure of $P_m^T \mathcal{H}_m P_m$ and some strategy should be used to avoid this problem.

L. Lopez & V. Simoncini, BIT 2006.

Linear Hamiltonian system:
$$\begin{cases} y' = Ay, & A = J^{-1}S \\ y(0) = y_0 \end{cases}$$

with $S \in \mathbb{R}^{400\times400}$ symmetric (eigs. in [1, 100])

Energy function: $E(y(t)) = y(t)^T S y(t)$ is constant for all $t > 0$.

Numerical symplectic integrator: starting with $y(0) = y_0$,

$$y_{n+1} = \exp(\tau A) y_n, \qquad n \geq 0 \qquad \tau = \frac{1}{40}$$

where y_n is the numerical approximation of $y(n\tau)$.

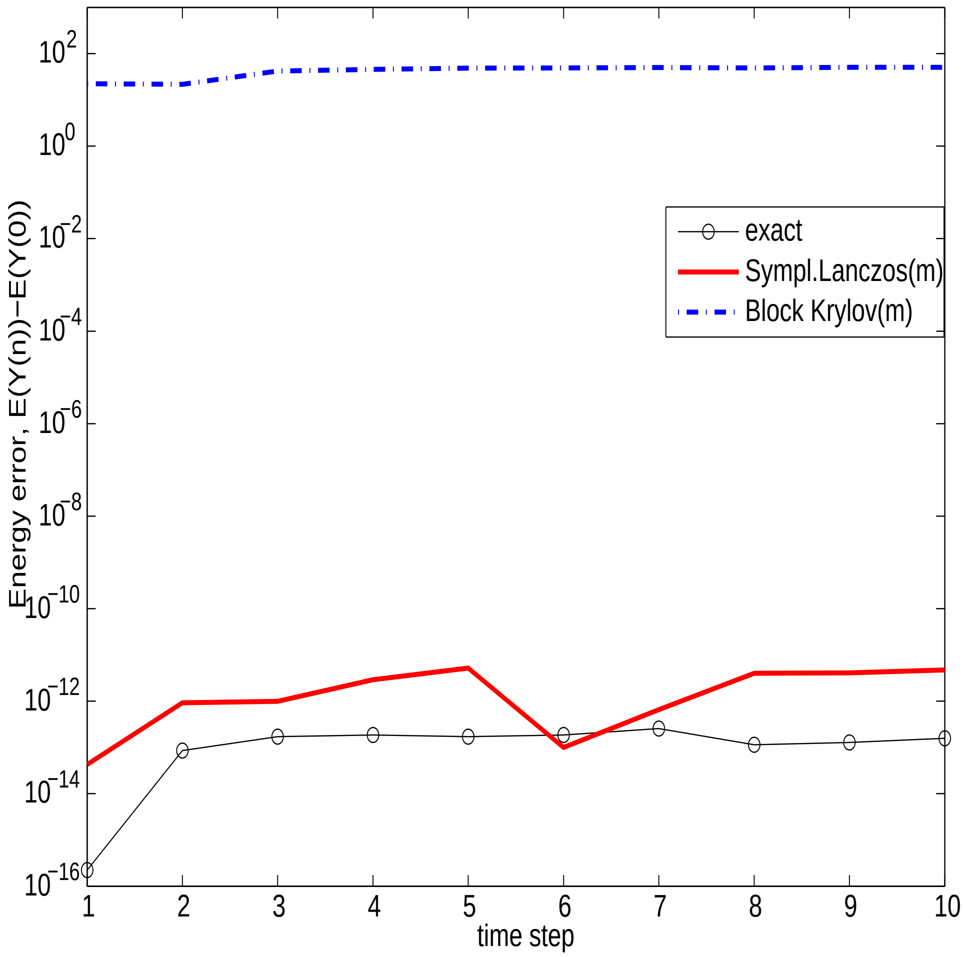

L. Lopez & V. Simoncini, BIT 2006.