Efficient computation of the Matrix Exponential by Generalized Polar Decomposition

Antonella Zanna
University of Bergen, Norway

email: anto@ii.uib.no
http://www.ii.uib.no/~anto

Overview

Motivations

- Integration methods using exponentials in GI need fast algorithms that approximate the matrix exponential
- Exact computation is not an issue – but it is crucial that the exponential approximation is in the Lie group.

Computation of the matrix exponential by GPDs

- A first approach (Z-MK)
- Fast methods (I-Z)

Generalized polar decompositions

Ingredients:

- A Lie group G
- An involutive automorphism σ

What we can do with them:

- We can factorize

$$z = xy \qquad \text{GPD}$$

  where

$$\sigma(x) = x^{-1}$$

  and

$$\sigma(y) = y$$

where x, y are appropriate group elements which are determined by z and σ.

In particular, if $z = \exp(tZ)$, it is true that

$$\exp(tZ) = \exp(X(t))\exp(Y(t))$$

and for t sufficiently small, the functions $X(t)$ and $Y(t)$ are uniquely determined.

Properties of the decomposition: at the group level

$$z = \exp(tZ) = \exp(X(t))\exp(Y(t)) = xy \qquad \text{GPD of } z$$

with $\sigma(x) = x^{-1}$, $\sigma(y) = y$.

Consider the sets

$$\begin{aligned} G^\sigma &= \{z \in G : \sigma(z) = z\} && \text{fixed points of } \sigma \\ G_\sigma &= \{z \in G : \sigma(z) = z^{-1}\} && \text{anti-fixed points of } \sigma \end{aligned}$$

- G^σ has the structure of a group:

$$z_1, z_2 \in G^\sigma \qquad \Rightarrow \qquad z_1 z_2 \in G^\sigma, \quad z_1^{-1} \in G^\sigma$$

- G_σ has the structure of a *symmetric space*,

$$z_1, z_2 \in G_\sigma \qquad \Rightarrow \qquad z_1 \star z_2 = z_1 z_2^{-1} z_1 \in G_\sigma.$$

At the algebra level...

Assume $z = \exp(tZ)$, where $Z \in \mathfrak{g}$, the Lie-algebra of G. The group automorphism σ induces a Lie-algebra map

$$\mathrm{d}\sigma(Z) = \frac{\mathrm{d}}{\mathrm{d}t}\Big|_{t=0} \sigma(\exp(tZ)), \qquad Z \in \mathfrak{g}$$

which is also an involutive automorphism since

$$\mathrm{d}\sigma([A,B]) = [\,\mathrm{d}\sigma(A), \mathrm{d}\sigma(B)], \quad A, B \in \mathfrak{g}, \qquad \mathrm{d}\sigma^2 = \mathrm{id},$$

We denote

$$\begin{array}{lcll} \mathfrak{k} & = & \{Z \in \mathfrak{g} : \mathrm{d}\sigma(Z) = Z\} & \text{subalgebra of } \mathfrak{g} \\ \mathfrak{p} & = & \{Z \in \mathfrak{g} : \mathrm{d}\sigma(Z) = -Z\} & \text{Lie triple system.} \end{array}$$

The subspaces $\mathfrak{p}$ and $\mathfrak{k}$ obey important inclusion properties:

$$\begin{array}{lcll} [\mathfrak{k}, \mathfrak{k}] & \subseteq & \mathfrak{k}, & (+1) \times (+1) = (+1) \\ [\mathfrak{k}, \mathfrak{p}] & \subseteq & \mathfrak{p}, & (+1) \times (-1) = (-1) \\ [\mathfrak{p}, \mathfrak{p}] & \subseteq & \mathfrak{k}, & (-1) \times (-1) = (+1). \end{array}$$

It is true that

$$\mathfrak{g} = \mathfrak{p} \oplus \mathfrak{k},$$

in other words, every $Z \in \mathfrak{g}$ can be uniquely written as

$$Z = P + K, \qquad \mathrm{d}\sigma(P) = -P, \quad \mathrm{d}\sigma(K) = K,$$

where

$$P = \frac{1}{2}(Z - \mathrm{d}\sigma(Z)), \qquad K = \frac{1}{2}(Z + \mathrm{d}\sigma(Z)).$$

In summary:

Group level	Algebra level
$G^\sigma := \{z \in G : \sigma(z) = z\}$, subgrp	$\mathfrak{k} = \{Z \in \mathfrak{g} : \mathrm{d}\sigma(Z) = Z\}$, subalg.
$G_\sigma := \{z \in G : \sigma(z) = z^{-1}\}$, symm. sp.	$\mathfrak{p} = \{Z \in \mathfrak{g} : \mathrm{d}\sigma(Z) = -Z\}$, LTS
$G = G_\sigma \cdot G^\sigma$	$\mathfrak{g} = \mathfrak{p} \oplus \mathfrak{k}$
$z = xy$	$Z = P + K$
$x = \exp(X(t)) \in G_\sigma$	$X(t) = \sum_{i=1}^\infty X_i t^i \in \mathfrak{p}, \quad X_i \in \mathfrak{p}$
$y = \exp(Y(t)) \in G^\sigma$	$Y(t) = \sum_{i=1}^\infty Y_i t^i \in \mathfrak{k}, \quad Y_i \in \mathfrak{k}$

The decomposition

$$Z = P + K$$

completely determines the functions $X(t)$ and $Y(t)$:

$$\begin{aligned}
X &= Pt - \tfrac{1}{2}[P,K]t^2 - \tfrac{1}{6}[K,[P,K]]t^3 \\
&\quad + \left(\tfrac{1}{24}[P,[P,[P,K]]] - \tfrac{1}{24}[K,[K,[P,K]]]\right)t^4 \\
&\quad + \left(\tfrac{7}{360}[K,[P,[P,[P,K]]]] - \tfrac{1}{120}[K,[K,[K,[P,K]]]] - \tfrac{1}{180}[[P,K],[P,[P,K]]]\right)t^5 \\
&\quad + \mathcal{O}(t^6),
\end{aligned}$$

$$\begin{aligned}
Y &= Kt - \tfrac{1}{12}[P,[P,K]]t^3 + \Big(\tfrac{1}{120}[P,[P,[P,[P,K]]]] \\
&\qquad + \tfrac{1}{720}[K,[K,[P,[P,K]]]] - \tfrac{1}{240}[[P,K],[K,[P,K]]]\Big)t^5 + \mathcal{O}(t^7).
\end{aligned}$$

In general, all the terms in the expansion of $X(t)$ and $Y(t)$ are obtained by explicit recurrence relations in terms of P and K.

Approximation of the matrix exponential

Recall that by GPD:

$$\exp(tZ) = \exp(X(t))\exp(Y(t)),$$

where

$$\begin{aligned}
X &= Pt - \tfrac{1}{2}[P,K]t^2 - \tfrac{1}{6}[K,[P,K]]t^3 \\
&\quad + \Big(\tfrac{1}{24}[P,[P,[P,K]]] - \tfrac{1}{24}[K,[K,[P,K]]]\Big)t^4 \\
&\quad + \Big(\tfrac{7}{360}[K,[P,[P,[P,K]]]] - \tfrac{1}{120}[K,[K,[K,[P,K]]]] - \tfrac{1}{180}[[P,K],[P,[P,K]]]\Big)t^5 \\
&\quad + \mathcal{O}(t^6),
\end{aligned}$$

$$\begin{aligned}
Y &= Kt - \tfrac{1}{12}[P,[P,K]]t^3 + \Big(\tfrac{1}{120}[P,[P,[P,[P,K]]]] \\
&\qquad + \tfrac{1}{720}[K,[K,[P,[P,K]]]] - \tfrac{1}{240}[[P,K],[K,[P,K]]]\Big)t^5 + \mathcal{O}(t^7).
\end{aligned}$$

and

$$Z = P + K, \qquad P = \frac{1}{2}(Z - \mathrm{d}\sigma(Z)), \qquad K = \frac{1}{2}(Z + \mathrm{d}\sigma(Z)).$$

A splitting method that approximates $\exp(tZ)$:

- Choose an appropriate σ.
- Split $Z = P + K$.
- Truncate the expansion

$$X(t) = Pt + \frac{1}{2}t^2[P,K] + \cdots, \quad Y(t) = Kt - \frac{1}{12}t^3[P,[P,K]] + \cdots$$

  to desired order.
- Compute the exponential of $X(t) \in \mathfrak{p}$
- Set $Z_1 = Y(t)$
- Repeat

In general, we iterate the procedure on the reduced space until we get a space of low dimension. At the end,

$$\exp(tZ) \approx \exp(X^{[1]})\exp(X^{[2]})\cdots\exp(X^{[m]})\exp(Y^{[m]}).$$

What are good choices of σ?

- The splitted factors should be easy to compute.
- Commutators should have a low complexity.
- Exponential/commutators of splitted parts should be easy to compute (either approximately or preferably exactly).

A good choice is splitting in matrices of low rank, for instance, borderded matrices: take

$$S = \left[\begin{array}{ccc|c} 1 & 0 & \cdots & 0 \\ 0 & 1 & \ddots & \vdots \\ \vdots & \ddots & \ddots & 0 \\ \hline 0 & \cdots & 0 & -1 \end{array}\right], \qquad \mathrm{d}\sigma(Z) = SZS,$$

then

$$P = \left[\begin{array}{ccc|c} 0 & \cdots & 0 & z_{1,n} \\ \vdots & 0 & \cdots & \vdots \\ 0 & \vdots & \ddots & z_{n-1,1} \\ \hline z_{n,1} & \cdots & z_{n,n-1} & 0 \end{array}\right], \qquad K = \left[\begin{array}{ccc|c} z_{1,1} & \cdots & z_{1,n-1} & 0 \\ \vdots & \cdots & \vdots & \\ z_{n-1,1} & \vdots & z_{n-1,n-1} & 0 \\ \hline 0 & \cdots & 0 & z_{n,n} \end{array}\right].$$

Such automorphisms work for $\mathrm{GL}(n), \mathrm{SL}(n), \mathrm{SO}(n)$.
Note that the commutators appearing in the expansion can be computed in $\mathcal{O}(n^2)$ computations (n^3 if the procedure is iterated for matrices of decreasing dimension)

An Euler–Rodrigues like formula for bordered matrices

The exponential of bordered matrices can be computed exactly by means of a formula similar to the Euler–Rodrigues formula for computing the exponential of a 3×3 skew-symmetric matrices.

Assume that $A \in \mathfrak{p}$ is of the form

$$A = \left[\begin{array}{c|c} O & \mathbf{a} \\ \hline \mathbf{b}^T & 0 \end{array}\right], \qquad \mathbf{a}, \mathbf{b} \in \mathbb{R}^n. \tag{4.1}$$

Then,

$$\exp(A) = \begin{cases} I + \frac{\sinh\theta}{\theta} A + \frac{1}{2}\left(\frac{\sinh(\theta/2)}{\theta/2}\right)^2 A^2, & \text{if } \mathbf{a}^T\mathbf{b} > 0, \theta = \sqrt{\mathbf{a}^T\mathbf{b}}, \\ I + A + \frac{1}{2}A^2, & \text{if } \mathbf{a}^T\mathbf{b} = 0, \\ I + \frac{\sin\theta}{\theta} A + \frac{1}{2}\left(\frac{\sin(\theta/2)}{\theta/2}\right)^2 A^2, & \text{if } \mathbf{a}^T\mathbf{b} < 0, \theta = \sqrt{-\mathbf{a}^T\mathbf{b}} \end{cases}$$

where

$$A^2 = \left[\begin{array}{c|c} \mathbf{a}\mathbf{b}^T & \mathbf{0} \\ \hline \mathbf{0}^T & \theta^2 \end{array}\right].$$

Note that $\exp(A)$ is never computed explicitely but always applied to a vector/matrix $\mathbf{v}$.

$$\begin{bmatrix} \mathbf{w}_k \\ w \end{bmatrix} = \exp(A)\mathbf{v} = \exp(A) \begin{bmatrix} \mathbf{v}_k \\ v \end{bmatrix} = \begin{bmatrix} \mathbf{v}_k + \zeta_1 \mathbf{a} \\ \zeta_2 \end{bmatrix},$$

where

$$\begin{aligned} \zeta_1 &= [\eta_1 v + \eta_2(\mathbf{b}^\top \mathbf{v}_k)], \\ \zeta_2 &= (1+\theta)v + (\mathbf{b}^\top \mathbf{v}_k). \end{aligned}$$

Cost of the computation (including both addition and multiplication) of the 'Euler–Rodrigues' exponential. The (k,k) column corresponds to the case when $\mathbf{a}, \mathbf{b}$ are full, the (k,p) corresponds to the case when $\mathbf{a}$ is full while only the last p components of $\mathbf{b}$ are nonzero and finally the (p,p) column corresponds to both $\mathbf{a}$ and $\mathbf{b}$ having only the last p components nonzero.

Cost of $\exp(A)$	(k,k)	(k,p)	(p,p)
$\mathbf{a}^\top \mathbf{b}$	$2k$	$2p$	$2p$
$\mathbf{b}^\top \mathbf{v}_k$	$2k$	$2p$	$2p$
$\zeta_1 \mathbf{a}$	k	k	p
$\mathbf{w}_k$	k	k	p
total, stage k	$6k$	$2k+4p$	$6p$
total, summing $1 \le k \le n$ (vector)	$3n^2$	n^2+4pn	$6pn$
matrix (n vectors)	$2n^3$	n^3+2pn^2	$4pn^2$

Back
Close

On the computation of commutators

Our first observation is that the involutions S are usually chosen so that $P = \Pi_{\mathfrak{p}}(Z)$ has low rank, hence only just a few nonzero eigenvalues.

- Use the theory of minimal polynomial, the least degree monic polynomial such that

$$p(\mathrm{ad}_A) = 0.$$

Lemma 4.1 *Consider the bordered matrix A in (4.1) with $\mathbf{a}\mathbf{b}^\top \neq O$. The minimal polynomial of ad_A is*

$$\begin{aligned} p(\lambda) &= \lambda(\lambda - 2\theta)(\lambda + 2\theta)(\lambda - \theta)(\lambda + \theta) \\ &= \lambda^5 - 5\mathbf{b}^\top\mathbf{a}\lambda^3 + 4(\mathbf{b}^\top\mathbf{a})^2\lambda, \end{aligned} \tag{4.2}$$

where $\theta = \sqrt{\mathbf{b}^\top\mathbf{a}}$. If $\mathbf{a}\mathbf{b}^\top = O$, and $\mathbf{a}$ and $\mathbf{b}$ are not both zero, then the minimal polynomial is

$$p(\lambda) = \lambda^3. \tag{4.3}$$

Proof. If A has distinct eigenvalues $\lambda_1, \lambda_2, \ldots, \lambda_m$ with algebraic multiplicities $r_1, r_2, \ldots, r_m$ respectively, the minimal polynomial of A has the form

$$q(\lambda) = \prod_{i=1}^{m} (\lambda - \lambda_i)^{g_i},$$

where g_i is the order of the largest Jordan block of A corresponding to the eigenvalue λ_i.

Let us assume first that $\mathbf{b}^\top \mathbf{a} \neq 0$. Assume that $\mathbf{v} = \begin{bmatrix} \mathbf{v}_1 \\ v_2 \end{bmatrix}$ is eigenvector of A corresponding to the eigenvalue λ. Imposing $A\mathbf{v} = \lambda \mathbf{v}$, we have

$$\begin{aligned} \mathbf{a} v_2 &= \lambda \mathbf{v}_1 \\ \mathbf{b}^\top \mathbf{v}_1 &= \lambda v_2 \end{aligned}$$

and deduce immediately that the eigenvalues of A are $\lambda = \pm\theta = \pm\sqrt{\mathbf{b}^\top \mathbf{a}}$ and $\lambda = 0$ with algebraic multiplicities one, one, and $n-2$ respectively. It is easily verified that these are also their geometric multiplicities: for $\lambda = \pm\theta$, eigenvectors are of the form $[\mathbf{a}, \pm 1]^\top$; for the zero eigenvalues, eigenvectors are of the form $[\mathbf{v}_1, 0]^\top$, $\mathbf{0} \neq \mathbf{v}_1 \in R^{n-1}$, satisfying $\mathbf{b}^\top \mathbf{v}_1 = 0$, furthermore, it is possible to find $n-2$ of those that are linearly independent.

Since the eigenvalues and eigenvectors of ad_A are the form $\lambda_i - \lambda_j$ and $\mathbf{y}_i^\top \mathbf{x}_j$ respectively, the λ_is being eigenvalues of A with left and right eigenvector $\mathbf{y}_i$ and $\mathbf{x}_i$ respectively, we deduce that ad_A has eigenvalues

$$\lambda = \pm 2\theta, \qquad \lambda = \pm\theta$$

with algebraic/geometric multiplicities one each, and

$$\lambda = 0$$

with algebraic and geometric multiplicity $n^2 - 4$. This implies that all Jordan blocks have size one, from which it follows directly that the minimal polynomial of ad_A is of the form (4.2).

Next, if $\theta = 0$ but $\mathbf{a}\mathbf{b}^\top \neq O$, namely $\mathbf{a}, \mathbf{b} \neq \mathbf{0}$, the eigenvalues of A, that we write as $[\mathbf{v}_1, v_2]^\top$, must obey the conditions

$$\begin{aligned} \mathbf{a} v_2 &= \mathbf{0} \\ \mathbf{b}^\top \mathbf{v}_1 &= 0. \end{aligned}$$

Since $\mathbf{a} \neq 0$, it must necessarily be $v_2 = 0$. Therefore eigenvalues must be of the form $[\mathbf{v}_1, 0]$. Recall that $\mathbf{v}_1$ has $n-1$ entries ($n-1$ free parameters) while the second equation $\mathbf{b}^\top \mathbf{v}_1 = 0$ gives only a linear constraint: This mean that we can find only $n-2$ linearly independent eigenvalues and two further linearly independent generalized eigenvalues. In terms of Jordan blocks, this means that A has a Jordan block of the form

$$J(0) = \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \end{bmatrix},$$

hence λ^3 is the minimal polynomial of A and, as a consequence, $A^3 = O$. Passing to the adjoint operator ad_A, recall that, for an arbitrary matrix C,

$$\mathrm{ad}_A^k C = \sum_{i=0}^{k} \binom{k}{i} (-1)^{k-i} A^i C A^{k-i}, \qquad k = 1, 2, \ldots. \tag{4.4}$$

Clearly, $\mathrm{ad}_A^5 C = O$ since in all terms there appears a power A^i with $i \geq 3$. For lower order powers, there are always terms of the type $A^i C A^{k-i}$ where $i, k-i \leq 2$. This means that it is always possible to find a matrix C for which at least one of terms does not vanish. Hence the minimal polynomial of ad_A is

$$p(\lambda) = \lambda^5.$$

Finally, in the case when either $\mathbf{a} = \mathbf{0}$ or $\mathbf{b} = \mathbf{0}$, by direct computation,

$$A^2 = O,$$

hence the minimal polynomial of A is λ^2. Insofar as ad_A is concerned, the first power to vanish in (4.4) is ad_A^3, and no lower power vanishes for arbitrary matrices C. Hence the minimal polynomial is

$$p(\lambda) = \lambda^3.$$

This completes the proof of the lemma. □

Theorem 4.2 *Assume that the matrix A is of the form* (4.1). *Then, for every $k = 1, 2, \ldots$, commutators by A can be computed as*

$$\mathrm{ad}_A^k = [C_1 + (-1)^k C_2] 2^k \theta^k + [C_3 + (-1)^k C_4] \theta^k, \qquad k = 1, 2, \ldots \tag{4.5}$$

when $\theta = \sqrt{\mathbf{b}^\top \mathbf{a}} \neq 0$, and

$$\begin{aligned} C_1 - C_2 &= \frac{1}{6}\left(-\frac{\mathrm{ad}_A}{\theta} + \frac{\mathrm{ad}_A^3}{\theta^3}\right) \\ C_3 - C_4 &= \frac{1}{3}\left(\frac{4\mathrm{ad}_A}{\theta} - \frac{\mathrm{ad}_P^3}{\theta^3}\right) \\ C_1 + C_2 &= \frac{1}{12}\left(-\frac{\mathrm{ad}_A^2}{\theta^2} + \frac{\mathrm{ad}_A^4}{\theta^4}\right) \\ C_3 + C_4 &= \frac{1}{3}\left(\frac{4\mathrm{ad}_A^2}{\theta^2} - \frac{\mathrm{ad}_A^4}{\theta^4}\right). \end{aligned} \tag{4.6}$$

If $\theta = 0$ but $\mathbf{ab}^\top \neq O$, then

$$\mathrm{ad}_A^k = O, \qquad k = 5, 6, 7, \ldots.$$

If $\theta = 0$ and either $\mathbf{a}$ or $\mathbf{b}$ is a zero vector, then

$$\mathrm{ad}_A^k = O, \qquad k = 3, 4, 5, \ldots.$$

Proof. It follows from the minimal polynomial (4.2). □

Back
Close

Complexity of the algorithms for full matrices:

Order 2	$\mathfrak{sl}(n), \mathfrak{so}(p,q)$ vector	$\mathfrak{sl}(n), \mathfrak{so}(p,q)$ matrix	$\mathfrak{so}(n)$ vector	$\mathfrak{so}(n)$ matrix
splitting	$1\frac{1}{3}n^3$	$1\frac{1}{3}n^3$	$\frac{2}{3}n^3$	$\frac{2}{3}n^3$
assembly exp	$3n^2$	$2n^3$	$3n^2$	$2n^3$
total	$1\frac{1}{3}n^3$	$3\frac{1}{3}n^3$	$\frac{2}{3}n^3$	$2\frac{2}{3}n^3$

Order 3(4)	$\mathfrak{sl}(n), \mathfrak{so}(p,q)$ vector	$\mathfrak{sl}(n), \mathfrak{so}(p,q)$ matrix	$\mathfrak{so}(n)$ vector	$\mathfrak{so}(n)$ matrix
splitting	$5(7)n^3$	$5(7)n^3$	$2\frac{1}{2}(4)n^3$	$2\frac{1}{2}(4)n^3$
assembly exp	$3n^2$	$2n^3$	$3n^2$	$2n^3$
total	$5(7)n^3$	$7(9)n^3$	$2\frac{1}{2}(4)n^3$	$4\frac{1}{2}(6)n^3$

These algorithms have a complexity that is comparable with other classical algorithms, like for instance diagonal Padé approximants.

Feasible algorithms are up to order 4, because for higher order the complexity becomes larger (although still $\mathcal{O}(n^3)$).

Back
Close

Faster algorithms

Combine GI and classical Linear Algebra techniques

The main difference with the approach presented before is that the matrix Z is preprocessed and reduced to a 'sparse' form stable under commutation, which is

- tridiagonal (for symmetric and skew-symmetric matrices),
- upper Hessenberg (for matrices in $\mathfrak{sl}(n)$),
- butterfly form (for symplectic matrices).

Again, we split rows and columns and start computing commutators.
In the following example, we consider a skew-symmetric tridiagonal matrix.

- 'red' for the $\mathfrak{p}$-part, 'blue' for the $\mathfrak{k}$-part
- updated elements are denoted with dots instead of crosses

Order 1:

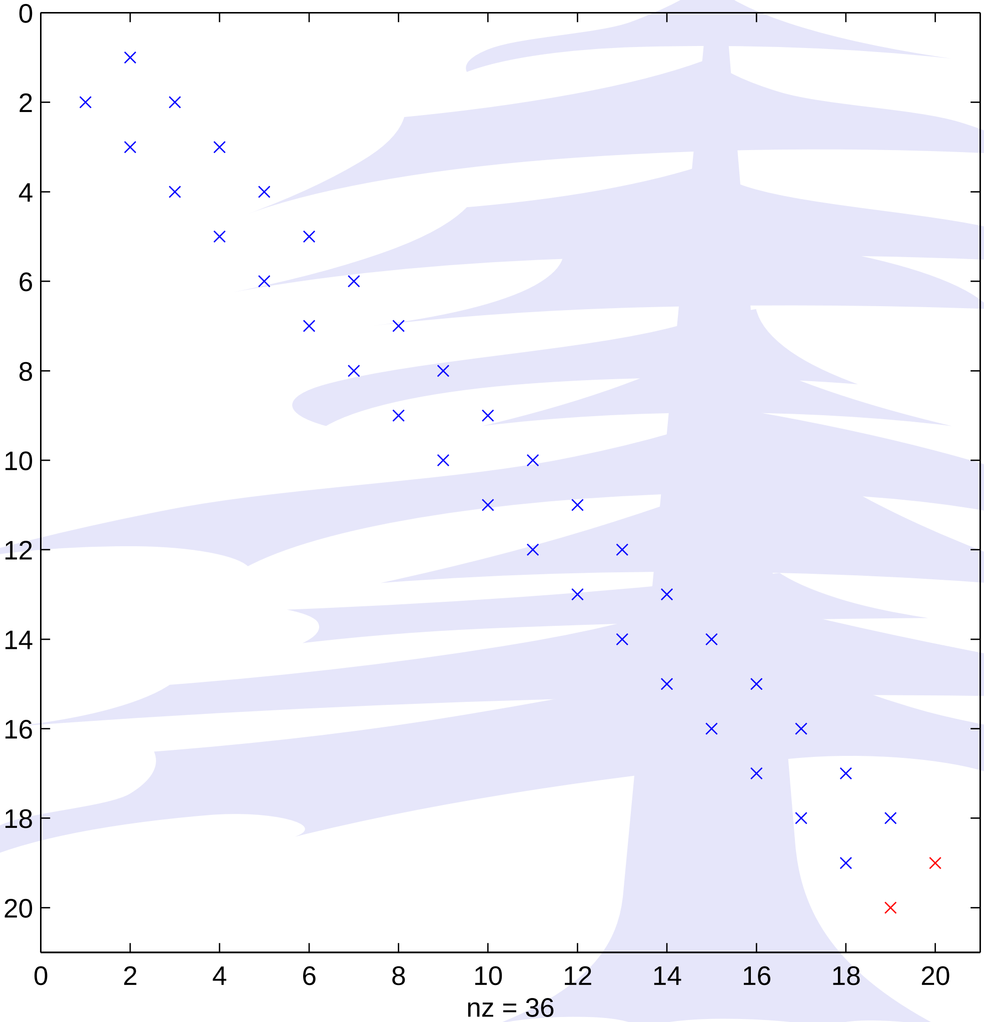

Back
Close

Order 2:

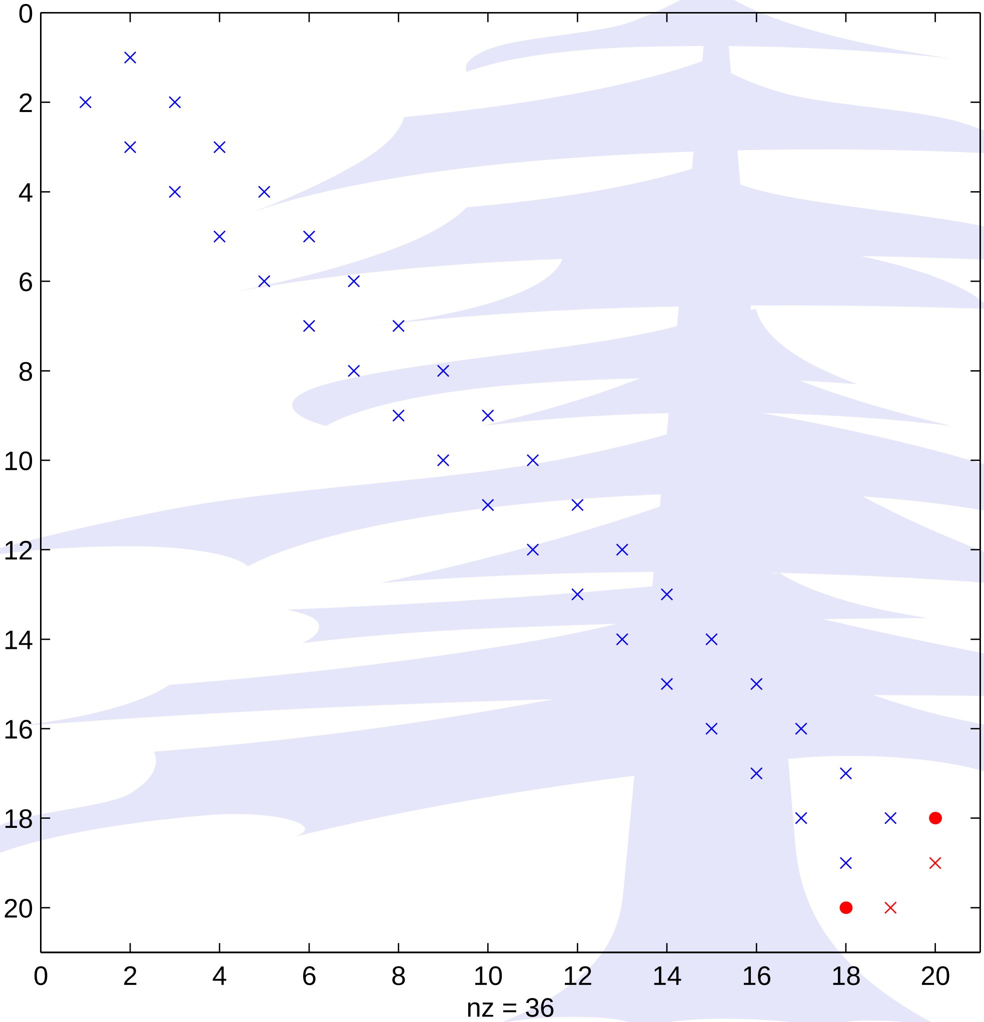

Order 3:

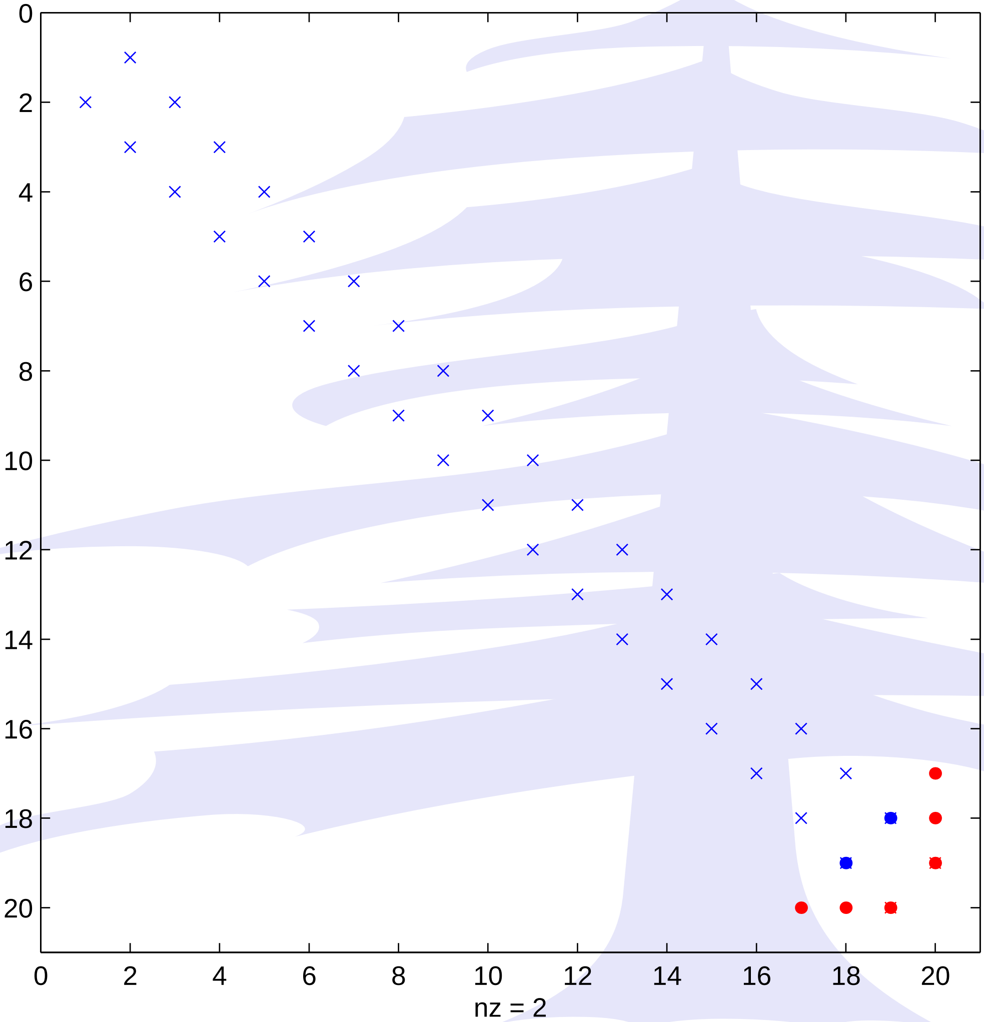

Order 4:

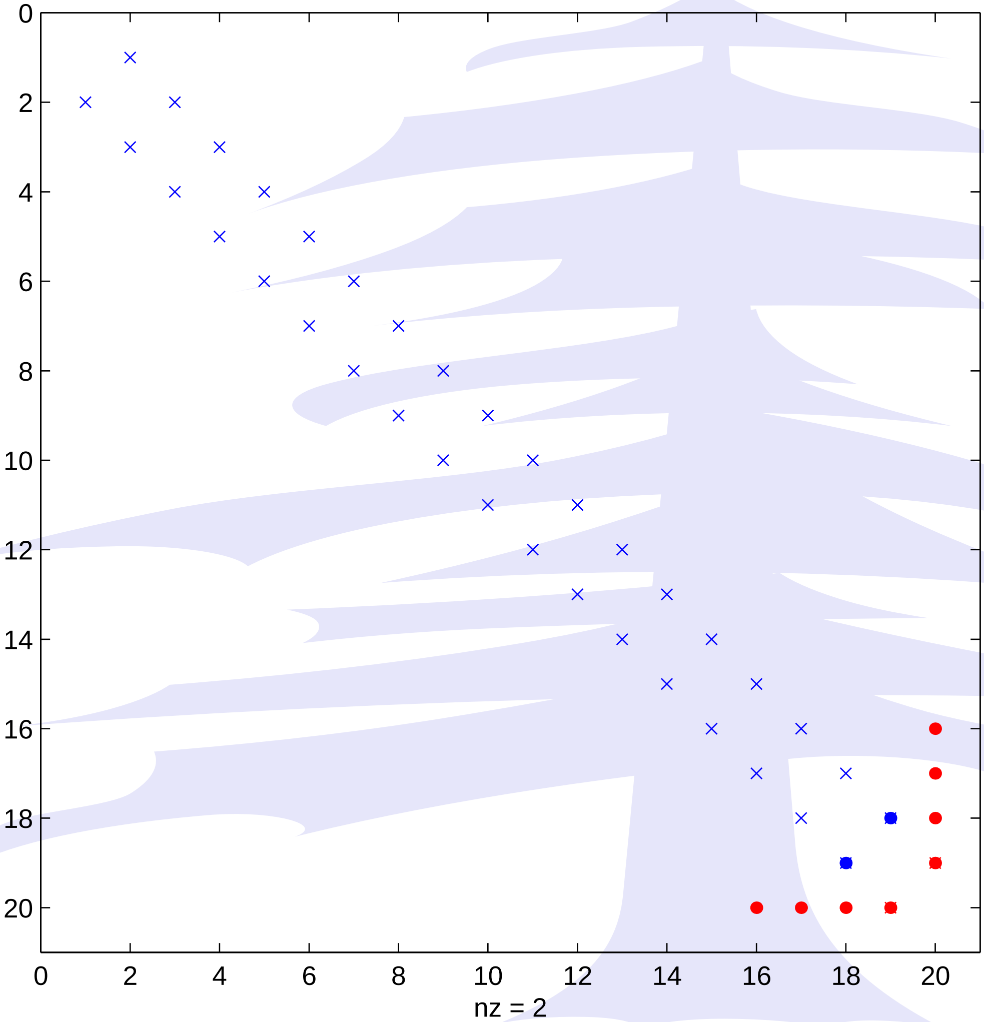

Order 5:

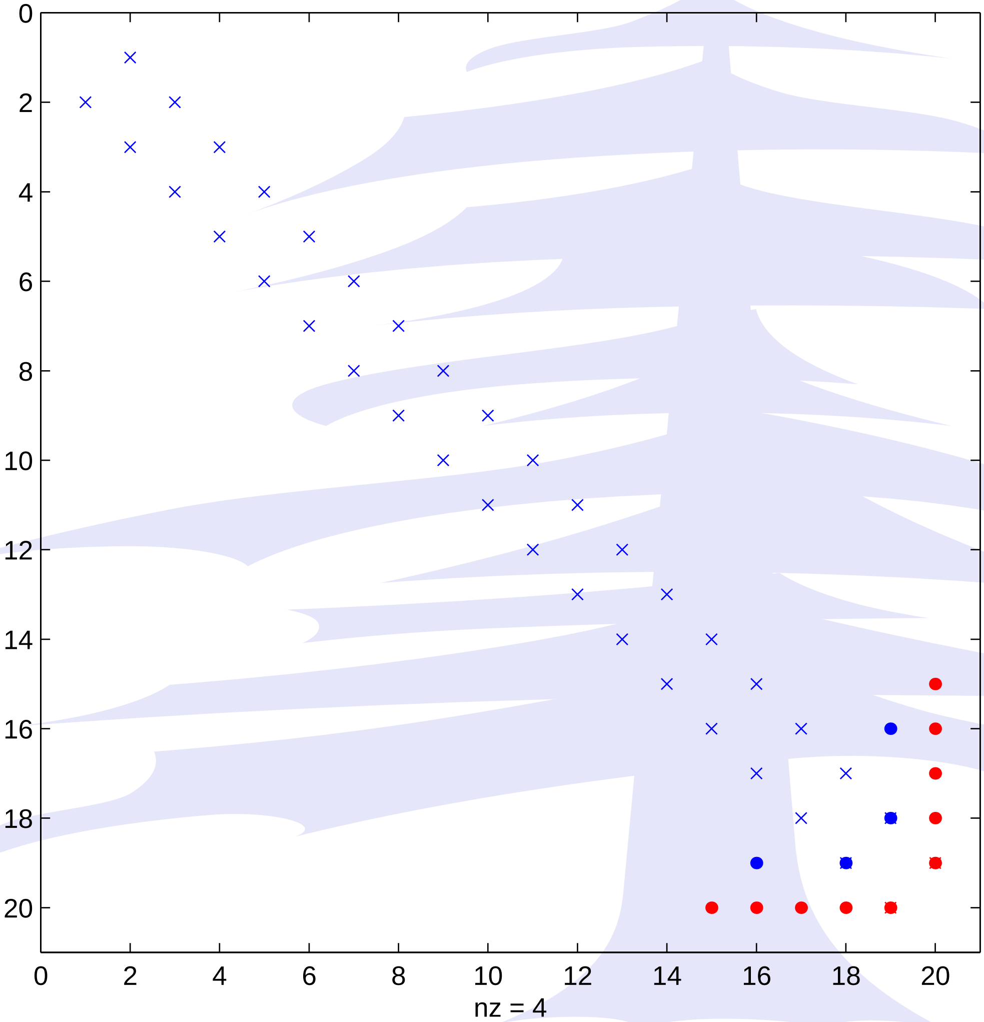

Main observations:

- Each extra order fills in two symmetric elements in the $\mathfrak{p}$-part.
- The fill-in in the $\mathfrak{k}$-part starts only at order 5.
- As long as the matrices P, K are tridiagonal, the commutators cost $\mathcal{O}(1)$.

The 'ugly' and the 'bad' fill-in

- The fill-in in the $\mathfrak{p}$-part is 'ugly' but not harmful: once once the $\mathfrak{p}$-term is computed up to desired order, one needs only compute the exponential.
- The fill-in in the $\mathfrak{k}$-part is much more dangerous: if not taken care of, it propagates and we lose the whole benefits of our tridiagonalization/reduction to Hessenberg

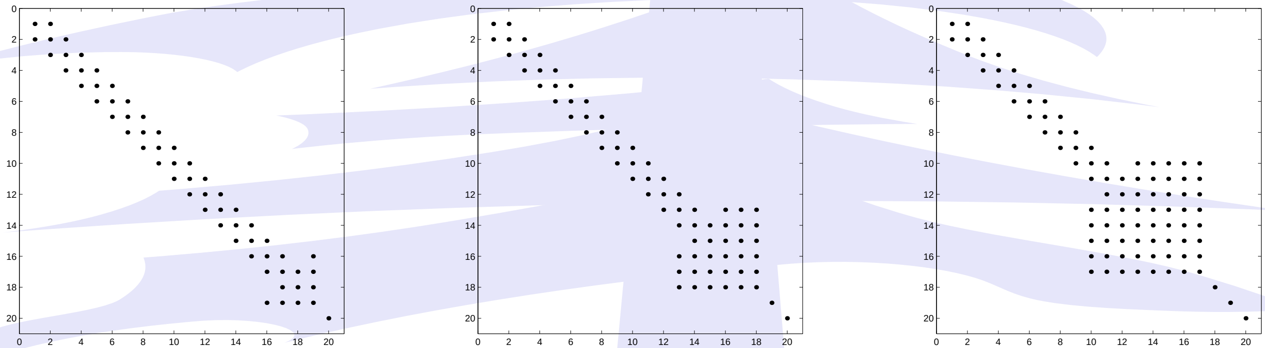

Therefore the fill-in elements in the $\mathfrak{k}$ part must be annihilated by, for instance, Givens rotations ($\mathcal{O}(1)$ computations)

Order 2	ZMK vector	ZMK matrix	IZ vector	IZ matrix
Tridiag.	–	–	n^3	n^3
order cond.	$\frac{2}{3}n^3$	$\frac{2}{3}n^3$	$\mathcal{O}(n)$	$\mathcal{O}(n)$
assembly exp	$3n^2$	$2n^3$	$6pn$	$4pn^2$
total	$\frac{2}{3}n^3$	$2\frac{2}{3}n^3$	$n^3+\mathcal{O}(pn)$	n^3+4pn^2

Order 3	ZMK vector	ZMK matrix	IZ vector	IZ matrix
Tridiag.	–	–	n^3	n^3
order cond.	$2\frac{1}{2}n^3$	$2\frac{1}{2}n^3$	$\mathcal{O}(n)$	$\mathcal{O}(n)$
assembly exp	$3n^2$	$2n^3$	$6pn$	$4pn^2$
total	$2\frac{1}{2}n^3$	$4\frac{1}{2}n^3$	$n^3+\mathcal{O}(pn)$	n^3+4pn^2

Order 4	ZMK vector	ZMK matrix	IZ vector	IZ matrix
Tridiag.	–	–	n^3	n^3
order cond.	$4n^3$	$4n^3$	$\mathcal{O}(n)$	$\mathcal{O}(n)$
assembly exp	$3n^2$	$2n^3$	$6pn$	$4pn^2$
total	$4n^3$	$6n^3$	$n^3+\mathcal{O}(pn)$	n^3+4pn^2

Comparison of cost of the approximation of the exponential without (ZMK) and with reduction to tridiagonal form (IZ) for splittings of order $2, 3, 4$. Only dominant terms are reported.

Matrices in $\mathrm{GL}(n), \mathrm{SL}(n)$

Order 5 terms for matrices in Hessenberg form:

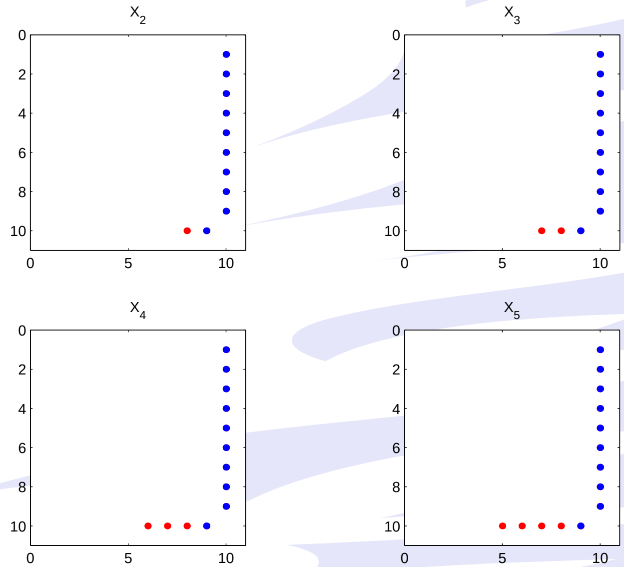

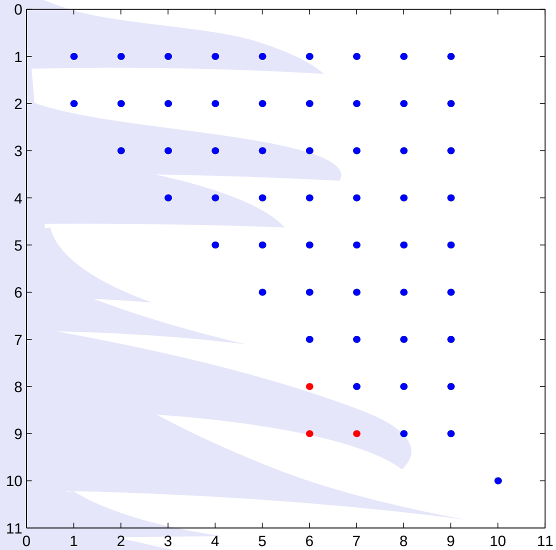

Back
Close

Symplectic matrices

Order 5 terms for matrices in butterfly form:

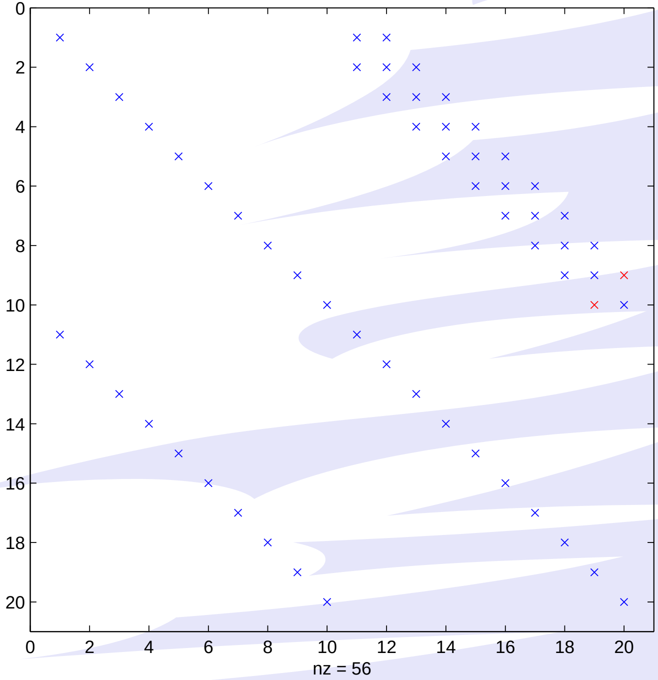

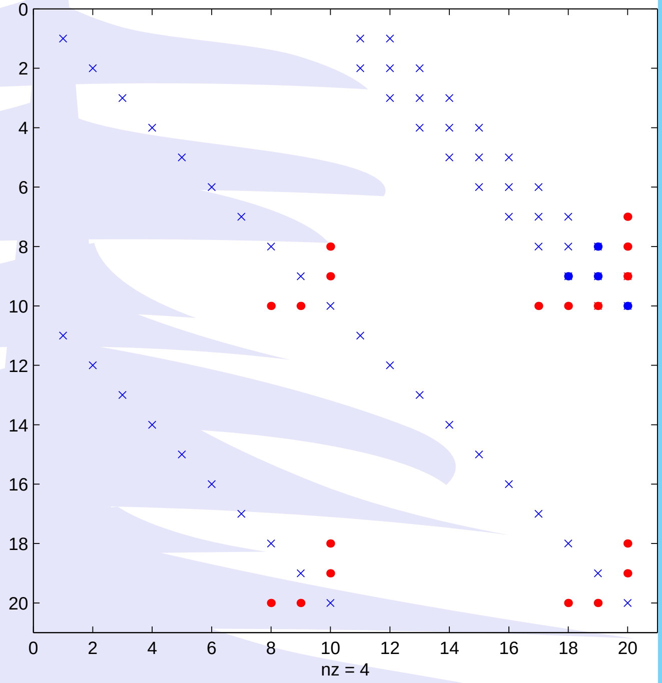

Reduction to tridiagonal/Hessenberg/ butterfly form

- For tridiagonalization/Hessenberg use Householder reflections:

$$H = I - \beta \mathbf{v}\mathbf{v}^T, \qquad \beta = \frac{2}{\|\mathbf{v}\|^2}$$

  Then,

$$HZH = Z - \beta \mathbf{v}\mathbf{v}^T Z - \beta Z \mathbf{v}\mathbf{v}^T + \beta^2 \underbrace{\mathbf{v}\mathbf{v}^T Z \mathbf{v}\mathbf{v}^T}_{0 \text{ if } Z \text{ is skew}}$$

  Cost:

  - n^3 for symmetric/skew-symmetric matrices
  - $\frac{10}{3}n^3$ for arbitrary matrices

Back
Close

- Reduction to butterfly form is done by means of symplectic transformations:
  - symplectic Givens/Householder (orthosymplectic)
  - symplectic Gauss transformations

  as proposed by Faßbender, Benner, Watkins (at the group level, for QR-like iterations).

  The basic idea of the algorithm is: at each step j,
  - bring the jth column of M into the desired form
  - bring the $(n+j)$th row of M into the desired form.

- **symplectic Givens transformations:,**

$$G = \left[\begin{array}{ccc|ccc} I_{k-1} & & & & & \\ & c & & & s & \\ & & I_{n-k} & & & \\ \hline & & & I_{k-1} & & \\ & -s & & & c & \\ & & & & & I_{n-k} \end{array}\right],$$

- **symplectic Householder transformations:**

$$H = \left[\begin{array}{cc|cc} I_{k-1} & & & \\ & Q & & \\ \hline & & I_{k-1} & \\ & & & Q \end{array}\right], \qquad Q = I_{n-k+1} - \beta \mathbf{v}\mathbf{v}^\top, \quad \beta = \frac{2}{\|v\|^2},$$

- **symplectic Gauss transformations:**

$$L = \left[\begin{array}{cccc|cccc} I_{k-2} & & & & & & & \\ & c & & & & & d & \\ & & c & & & d & & \\ & & & I_{n-k} & & & & \\ \hline & & & & I_{k-2} & & & \\ & & & & & c^{-1} & & \\ & & & & & & c^{-1} & \\ & & & & & & & I_{n-k} \end{array}\right].$$

Comparison of cost of the approximation of the exponential without (ZMK) and with reduction to Hessenberg form (IZ) for splittings of order $2, 3, 4$. Only dominant terms are reported.

Order	ZMK		IZ	
2	vector	matrix	vector	matrix
Hessenberg	–	–	$3\frac{1}{3}n^3$	$3\frac{1}{3}n^3$
order cond.	$1\frac{1}{3}n^3$	$1\frac{1}{3}n^3$	$\frac{1}{3}n^3$	$\frac{1}{3}n^3$
assembly exp	$3n^2$	$2n^3$	n^2	n^3
total	$1\frac{1}{3}n^3$	$2\frac{1}{3}n^3$	$3\frac{2}{3}n^3$	$4\frac{2}{3}n^3$

Order	ZMK		IZ	
3	vector	matrix	vector	matrix
Hessenberg	–	–	$3\frac{1}{3}n^3$	$3\frac{1}{3}n^3$
order cond.	$5n^3$	$5n^3$	$\frac{2}{3}n^3$	$\frac{2}{3}n^3$
assembly exp	$3n^2$	$2n^3$	n^2	n^3
total	$5n^3$	$7n^3$	$4n^3$	$5n^3$

Order	ZMK		IZ	
4	vector	matrix	vector	matrix
Hessenberg	–	–	$3\frac{1}{3}n^3$	$3\frac{1}{3}n^3$
order cond.	$7n^3$	$7n^3$	n^3	n^3
assembly exp	$3n^2$	$2n^3$	n^2	n^3
total	$7n^3$	$9n^3$	$4\frac{1}{3}n^3$	$5\frac{1}{3}n^3$

Some concluding remarks

What we got so far . . .

- Methods that stay on the correct Lie group
- They are cheap
- We can tackle most groups
- High order is easily achievable
- Easy scaling and squaring

Open issues

- A divide and conquer approach – is it suitable for large problems on parallel machines?
- Stiff problems? everything O.K. on $\mathrm{SO}(n)$, but what about other problems?
- Different choices of automorphisms that lead to other splittings
- Very large problems (for which $\mathcal{O}(n^3)$ is not feasible): compare with Krylov