UNIVERSITY OF CAMBRIDGE

CENTRE FOR MATHEMATICAL SCIENCES

DEPARTMENT OF APPLIED MATHEMATICS
& THEORETICAL PHYSICS

Lie-group methods in geometric numerical integration

Arieh Iserles

OUTLINE

1. Why geometric numerical integration?
2. Elements of differential geometry
3. The exponential and trivializations
4. Runge–Kutta–Munthe-Kaas methods
5. Magnus and Magnus-type expansions
6. Multivariate quadrature
7. Graded Lie algebras
8. The matrix exponential

Why geometric numerical integration?

A classical paradigm of applied mathematics:

Do rigorous, pure mathematics as much as possible. Determine qualitative features of your problem. And then, when exact analysis has reached its limits, resort to computation.

The problem: having spent great effort and ingenuity on finding precise qualitative information on the behaviour of our problem, which often has deep physical significance, we produce numerical solution that does not respect this qualitative information.

Invariants represent important qualitative information about the differential system: it is often advantageous to respect them under discretization. **This means designing numerical methods that share qualitative features of the differential equation(s).**

This is precisely the goal of ***geometric numerical integration (GNI)****.*

Major themes in GNI:

- Symplectic methods for DEs with Hamiltonian structure
  (Newton, Störmer, Verlet, de Vogelaerre, Feng Kang, Sanz-Serna, Scovel, Hairer & Lubich, Bennetin & Gorgili, . . .)

- Volume and energy conservation in DEs
  (Feng Kang, McLachlan & Quispel, . . .)

- Methods respecting Lie-Poisson structure
  (Marsden, Lewis & Simo, Ratiu, . . .)

- Methods replacing symplectic structure by a 'nearby' symplectic structure
  (Moser & Veselov, Marsden, Bridges & Reich, Hong, . . .)

- Methods for problems evolving on a differentiable manifold, in particular on a homogeneous manifold
  (Crouch, Munthe-Kaas, AI & Nørsett, Zanna, Owren, . . .)

Major techniques in GNI:

- Identification of 'classical' methods that retain geometric invariants
  (Sanz-Serna, Scovel, Skeel, . . .)
- Splitting and composition methods
  (Yošida, McLachlan & Quispel, Sanz-Serna & Murua, Blanes & Casas, . . .)
- Backward error analysis
  (Naishtadt, Bennetin & Gorgili, Reich, Hairer & Lubich, . . .)
- Clever asymptotic expansions
  (Cohen, Hairer & Lubich)
- Structure-preserving projection techniques
  (Hairer & Lubich, Calvo, AI & Zanna, . . .)
- Trivializations, group actions and Lie-algebraic techniques
  (Lewis & Simo, Crouch, Munthe-Kaas, AI, Nørsett, Zanna, P. Olver, Owren, . . .)
- Tricks from linear and abstract algebra
  (Munthe-Kaas & Owren, Munthe-Kaas, Quispel & Zanna, AI & Zanna, . . .)

Elements of differential geometry

A smooth manifold is a smooth domain which locally ‘looks like’ an Euclidean space: more formally, it can be covered by a smooth atlas made out of local coordinate charts. Conceptually, think of

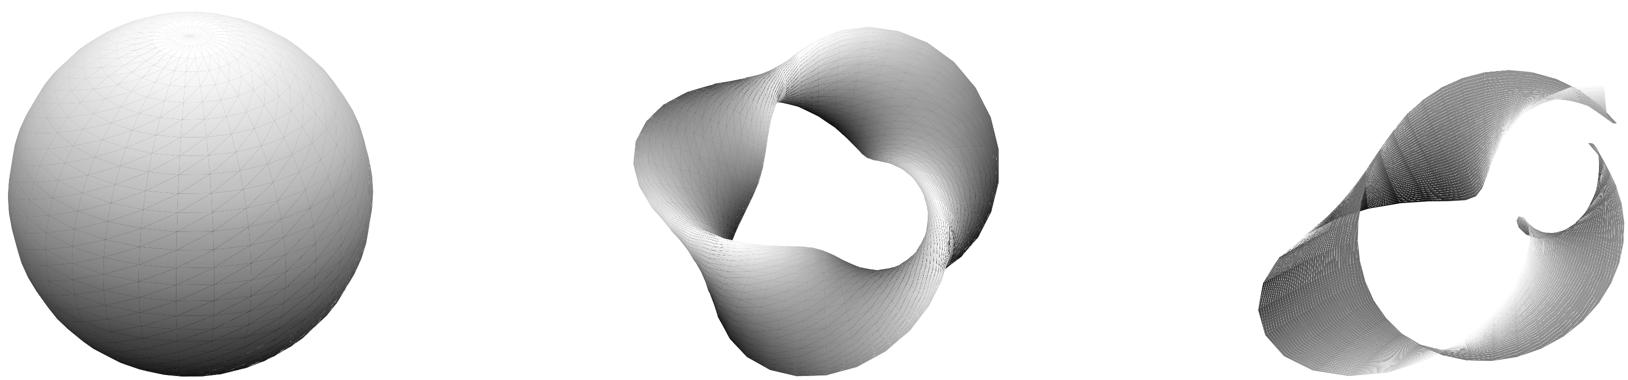

A tangent vector at $p \in \mathcal{M}$ is $\mathrm{d}\rho(t)/\mathrm{d}t|_{t=0}$, where $\rho(t) \in \mathcal{M}$ is a smooth curve s.t. $\rho(0) = p$.

The linear space of all tangents at p is the tangent space $\mathsf{T}\mathcal{M}|_p$, while

$$\mathsf{T}\mathcal{M} = \cup_{p \in \mathcal{M}} \mathsf{T}\mathcal{M}|_p$$

is the tangent bundle.

The cotangent space $\mathsf{T}^\star\mathcal{M}|_p$ consists of all linear functionals acting on elements of $\mathsf{T}\mathcal{M}|_p$.

Differential equations and tangents

A vector field on a manifold $\mathcal{M}$ is a smooth function $F(p) \in \mathsf{T}\mathcal{M}|_p$, $p \in \mathcal{M}$.

The set of all vector fields over $\mathcal{M}$ is denoted by $\mathfrak{X}(\mathcal{M})$ and, clearly, it is a linear space.

The differential equation

$$y' = F(y), \quad t \geq 0, \qquad y(0) = y_0 \in \mathcal{M},$$

where $F \in \mathfrak{X}(\mathcal{M})$, evolves on the manifold $\mathcal{M}$.

The flow of this DE is

$$y(t) = \Psi_{t,F}(y_0), \qquad t \geq 0.$$

Therefore

$$F(y) = \frac{\mathrm{d}}{\mathrm{d}t}\Psi_{t,F}(y_0)|_{t=0}.$$

In other words, F is the infinitisimal generator of the flow.

Noting that $\Psi_{\alpha,F} = \Psi_{1,\alpha F}$, we define

$$\Psi_{1,F} = \exp(F), \qquad \text{hence} \qquad \exp(tF) = \Psi_{t,F}.$$

This is the exponential map.

[This should not be confused with the concept of a semigroup in PDE theory, although at a deeper abstract level the two are actually quite similar!]

In general, flows fail to commute. Thus, let

$$\Phi_{s,t} = \exp(sF) \circ \exp(tG) \circ \exp(-sF) \circ (-tG).$$

If flows commute then $\Phi_{s,t}(y) = y$, but this isn't in general true.

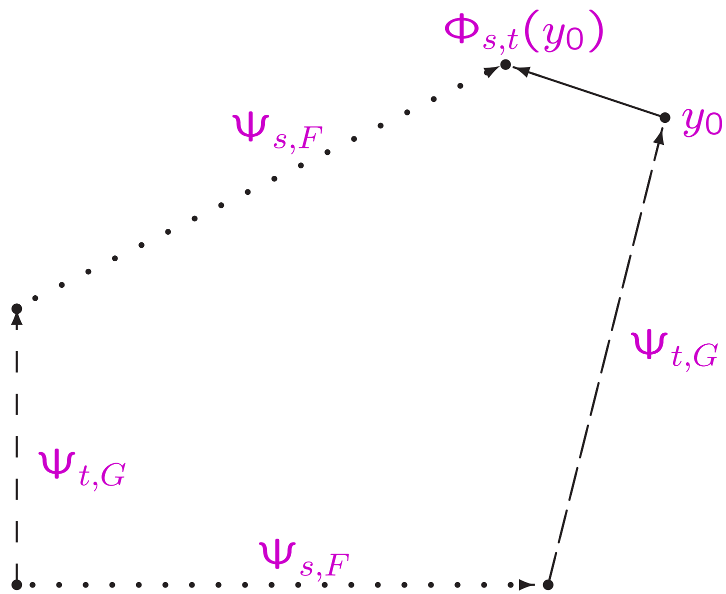

The local measure of lack of commutativity is the commutator $H = [F, G]$ where (translating to the standard Euclidean coordinates in $\mathbb{R}^n$)

$$H_i(y) = \sum_{j=1}^{n} \left\{ G_j(y) \frac{\partial F_i(y)}{\partial y_j} - F_j(y) \frac{\partial G_i(y)}{\partial y_j} \right\}.$$

Note that for every $F, G, H \in \mathfrak{X}(\mathcal{M})$ and scalar α we have the following features of the bracket operation:

Skew symmetry: $[F, G] = -[G, F]$;

Bilinearity:

$$[\alpha F, G] = \alpha[F, G],$$
$$[F + G, H] = [F, H] + [G, H];$$

The Jacobi identity

$$[F, [G, H]] + [G, [H, F]] + [H, [F, G]] = 0.$$

Therefore, $\mathfrak{X}(\mathcal{M})$ is a Lie algebra.

In the important case of linear DE $y' = ay$ (where a is a matrix) we have $\Psi_{t,F}(y) = \exp(ta)y_0$, with the familiar matrix exponential

$$e^{ta} = \sum_{m=0}^{\infty} \frac{1}{m!} t^m a^m.$$

In that case $[a, b] = ab - ba$.

Examples of manifolds:

- The sphere $\mathbb{S}_n \subset \mathbb{R}^n$: $\boldsymbol{x} \in \mathbb{R}^n$ s.t. $\|\boldsymbol{x}\| = 1$;
- The regular torus $\mathbb{T}_n \subset \mathbb{R}^n$: $\boldsymbol{x} \in \mathbb{R}^n$ which are 1-periodic in each coordinate;
- The orthogonal group $\mathrm{SO}(n)$ of $n \times n$ orthogonal matrices;
- The Grassmannian $\mathrm{G}_{n,m}$ of real $n \times m$ matrices, $m \leq n$, consisting of unit-length columns which are orthogonal to each other and equivalenced by orthogonal transformations;
- The isospectral orbit $\mathbb{I}_n(y_0)$ of all $n \times n$ symmetric matrices which are similar to the symmetric matrix y_0;
- The projective space $\mathbb{P}^n$ of all lines in $\mathbb{R}^n$ that pass through the origin.

Lie groups

Lie group $\mathcal{G}$ is smooth manifold, endowed with group structure which is continuous with respect to the topology of the manifold.

$\mathsf{O}(n, \mathbb{R})$**:** real $n \times n$ orthogonal matrices (*orthogonal group*);

$\mathsf{SL}(n, \mathbb{R})$**:** real $n \times n$ matrices with unit determinant (the *special linear group*);

$\mathsf{SU}(n, \mathbb{C})$**:** complex $n \times n$ unitary matrices with unit determinant (the *special unitary group*);

$\mathsf{SO}(n, m, \mathbb{R})$**:** real $n \times n$ matrices a s.t. $apa^\top = p$, $p = \mathsf{diag}\,[\mathbf{1}_m, -\mathbf{1}_{n-m}]$, with unit determinant. ($n = 4$, $m = 1$: the Lorenz group);

$\mathsf{E}(n, \mathbb{R})$**:** The Euclidean group of all translations and length-preserving linear transformations in $\mathbb{R}^n$;

$\mathsf{A}(n, \mathbb{R})$**:** The affine group of all translations and area-preserving linear transformations in $\mathbb{R}^n$.

Why are Lie groups useful?

Lie groups are important since they provide an appropriate formalism to investigate symmetries, invariants and qualitative behaviour of differential equations.

There are many other good reasons why Lie groups are important. We'll see one (group actions on manifolds) but there are many others, not least in number theory.

Many physical laws are conveniently formulated with built-in Lie-group symmetries.

For example, laws of motion have $\mathrm{SO}(3)$ symmetry, equations of special relativity evolve in $\mathrm{SO}_{4,1}$ and theory of superstrings can be formulated in $\mathrm{SU}(32)$ and in $\mathrm{E}(8) \times \mathrm{E}(8)$.

A finite-dimensional Lie group can be usually (but not always!) represented as a subgroup of the set of $n \times n$ nonsingular square matrices, $\mathrm{GL}_n[\mathbb{F}]$, for some $n \geq 1$ and a field $\mathbb{F}$. Such groups are called *matrix Lie groups.*

The tangent space of a Lie group

Let $\mathcal{G}$ be a Lie group and $\mathfrak{g} = \mathsf{T}\mathcal{G}|_I$ the tangent space at identity. Since $\mathcal{G}$ is a group, it follows at once that

$$\mathsf{T}\mathcal{G}|_X = \mathfrak{g}X \qquad \text{for all} \qquad X \in \mathcal{G}.$$

Since (as we have already seen) there is a natural isomorphism between (finite-dimensional) linear vector fields and square matrices, we deduce that

- g is a Lie algebra (cf. next slide)

- The exponential map is the classical matrix exponential.

- Let ρ and σ be smooth curves on $\mathcal{G}$ s.t. $\rho(t) = I + ta + \cdots$ and $\sigma(t) = I + tb + \cdots$. Then $a, b \in \mathfrak{g}$ and

$$[a, b] = \frac{\partial^2}{\partial t \partial s} \rho(s)\sigma(t)\rho(-s)|_{t=s=0} = ab - ba.$$

Lie algebras

We say that the abstract set $\mathfrak{g}$ is a Lie algebra if it is a linear space, which in addition is closed under the binary operation $[\,\cdot\,,\,\cdot\,] : \mathfrak{g} \times \mathfrak{g} \to \mathfrak{g}$ which obeys the following axioms:

1. Linearity: For every $a, b, c \in \mathfrak{g}$ and scalars α, β it is true that
$$[\alpha a + \beta b, c] = \alpha[a, c] + \beta[b, c].$$

2. Skew-symmetry: For every $a, b \in \mathfrak{g}$
$$[a, b] = -[b, a].$$

3. The Jacobi identity: For every $a, b, c \in \mathfrak{g}$
$$[[a, b], c] + [[b, c], a] + [[c, a], b] = 0.$$

Lie algebras can be fairly strange objects and not all of them originate as tangent spaces of Lie groups. A nice example is the modular group $\mathrm{SL}(n, \mathbb{Z})$ of integer matrices with unit determinant, which is important in number theory.

Examples of Lie algebras

1. Lie group: $\mathrm{O}(n, \mathbb{R}), \mathrm{SO}(n, \mathbb{R})$
   Lie algebra: The set $\mathfrak{so}(n, \mathbb{R})$ of real $n \times n$ skew symmetric matrices;

2. Lie group: $\mathrm{SL}(n, \mathbb{R})$:
   Lie algebra: The set $\mathfrak{sl}(n, \mathbb{R})$ of real $n \times n$ matrices with zero trace;

3. Lie group: $\mathrm{SO}(n, m, \mathbb{R})$:
   Lie algebra: The set $\mathfrak{so}(n, m, \mathbb{R})$ of real $n \times n$ matrices a such that $ap + pa^\top = 0$, where $p = \mathrm{diag}[\mathbf{1}_m, -\mathbf{1}_{n-m}]$. With greater generality, for any nonsingular symmetric p and the quadratic Lie group

   $$\{x \in \mathrm{GL}(n, \mathbb{R}) \ : \ xpx^\top = p\}$$

   the corresponding Lie algebra is

   $$\{a \in \mathfrak{gl}(n, \mathbb{R}) \ : \ ap + pa^\top = 0.$$

Group actions

An action of a Lie group $\mathcal{G}$ on a manifold $\mathcal{M}$ is a smooth map $\Lambda : \mathcal{G} \times \mathcal{M} \to \mathcal{M}$ s.t.

$$\Lambda(I, y) = y, \qquad y \in \mathcal{M},$$
$$\Lambda(p, \Lambda(q, y)) = \Lambda(pq, y), \qquad p, q \in \mathcal{G},\ y \in \mathcal{M}.$$

- Each group acts on itself;
- $\mathrm{O}(n, \mathbb{R})$ acts on on the sphere $\mathbb{S}_n$, $\Lambda(p, y) = py$;
- $\mathrm{SO}(n, \mathbb{R})$ acts on the Grassmannian

$$\mathbb{G}(n, m) = \mathrm{S}(n)/(\mathrm{SO}(m) \times \mathrm{SO}(n - m))$$

- ...and on the isospectral orbit $\mathbb{I}_n(y_0)$ via

$$\Lambda(p, y) = pyp^{-1}.$$

As a matter of fact, all previous examples of manifolds are subject to transitive group action: every point in $\mathcal{M}$ is reachable from any other point via the group action. In that case M is a homogeneous manifold.

Differential equations and group actions

Given a homogeneous manifold $\mathcal{M}$ and a group action Λ we define $\lambda_* : \mathfrak{g} \to \mathfrak{X}(\mathcal{M})$ as

$$\lambda_*(a)(y) = \frac{\mathrm{d}}{\mathrm{d}s}\Lambda(\rho(s), y)|_{s=0},$$

where $\rho(s) = I + at + \cdots$ is a smooth curve in $\mathcal{G}$.

Suppose that $\mathcal{G}$ is a matrix group. Then, for $a \in \mathfrak{g}$ the flow of λ_*, i.e.

$$y' = \lambda_*(a)(y), \quad t \geq 0, \qquad y(0) = y_0 \in \mathcal{M},$$

can be expressed in the form

$$y(t) = \Lambda(s(t), y_0), \qquad t \geq 0,$$

where

$$s' = as, \quad t \geq 0, \qquad s(0) = I.$$

This can be generalized from fixed $a \in \mathfrak{g}$ to a sufficiently smooth function $a : \mathbb{R}_+ \times \mathcal{M} \to \mathfrak{g}$: The solution of the differential equation

$$y' \lambda_*(a(t, y))(y), \quad t \geq 0, \qquad y(0) = y_0 \in \mathcal{M},$$

can be represented as $y(t) = \Lambda(s(t), y_0)$, where

$$s' = a(t, \Lambda(s, y_0))s, \quad t \geq 0, \qquad s(0) = I,$$

evolves in $\mathcal{G}$.

Therefore,

instead of solving the equation in $\mathcal{M}$, i.e. finding y_{N+1} given y_N, say, find a group action that takes y_N to y_{N+1}.

We conclude that

if we can devise a numerical method that respects a Lie-group structure then it can be extended to respect every homogeneous manifold structure acted upon by that group.

The main paradigm of Lie-group methods

Instead of a group action, we can consider an algebra action. Specifically, if $\mathcal{G} \ni x = \mathrm{e}^{v}$, where $v \in \mathfrak{g}$, then $\mu(v, y) = \Lambda(x, y)$. Note that this is less general than a group action: it is entirely possible that there is no single $v \in \mathfrak{g}$ s.t. $x = \mathrm{e}^{v}$, although in a finite-dimensional group there always exist $v_1, v_2, \ldots, v_r$ s.t. $x = \mathrm{e}^{v_1} \cdots \mathrm{e}^{v_r}$.

For example, each element of $\mathrm{SL}(2)$ can be obtained as a product of two exponentials of elements from $\mathfrak{sl}(2)$.

We follow the pattern

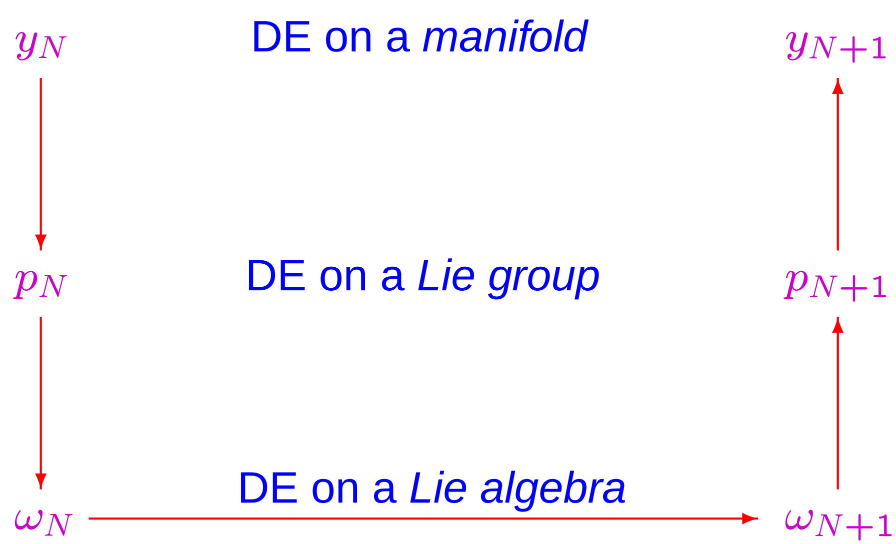

Some Lie-group methods do not follow this pattern:

- The method of Crouch & Grossman represents the equation in the form $y' = \sum_{l=1}^{d} \alpha_l(t) q_l$, using a rigid frame $q_l \in \mathrm{T}\mathcal{M}|_y$, $l = 1, \ldots, d$, and composes the solution from one-dimensional 'steps'.

- The method of McLachlan, Quispel & Robidoux represents the equation in the skew-gradient form $y' = S(t, y)\nabla g(y)$, where S is skew-symmetric, and suitably discretizes the gradient.

- The approach of cannonical coordinates of the second kind (Marthinsen & Owren) writes

$$\mathcal{G} \ni y(t) = \mathrm{e}^{\theta_1(t)\phi_1} \mathrm{e}^{\theta_2(t)\phi_2} \ldots \mathrm{e}^{\theta_d(t)\phi_d} y_0,$$

  where $\dim \mathfrak{g} = d$, $\{\phi_1, \ldots, \phi_d\}$ is a basis of $\mathfrak{g}$ and $\theta_1, \ldots, \theta_d$ are scalar functions. It is then possible to derive differential equations for the unknowns θ_k. The CCSK approach is particularly useful when the basis corresponds to the root space decomposition of g, since this simplifies the equations a great deal.

The exponential and trivializations

An equation evolving on a homogeneous manifold $\mathcal{M}$ can be written in the form

$$y' = \mu_*(a(t,y))(y), \quad t \geq 0, \qquad y(0) = y_0 \in \mathcal{M},$$

where the function $\mu_* : \mathfrak{g} \to \mathfrak{X}(\mathcal{M})$ is a Lie-algebra homomorphism, $\mathfrak{X}(\mathcal{M})$ is the set of vector fields on $\mathcal{M}$ and $a : [0,\infty) \times \mathcal{M} \to \mathfrak{g}$.

Many such ODEs occur in applications:

$\mathrm{O}(n,\mathbb{R})$ *(Orthogonal group):* mechanical systems, robotics, computer vision, computation of Lyapunov exponents, isospectral flows, numerical linear algebra;

$\mathrm{SL}(n,\mathbb{R})$ *(Special linear group):* conservation of volume, Riccati systems, Sturm–Liouville problems, image processing;

$\mathrm{Sp}(n,\mathbb{R})$ *(Symplectic group):* Hamiltonian and Lie–Poisson systems;

$\mathrm{SO}(1,3)$ *(Lorenz group):* relativity theory (also $\mathrm{SO}(2,5)$);

$\mathrm{SU}(n,\mathbb{C})$ *(Unitary group):* quantum mechanics.

Example 1: 'Classical' Lie-group equations:

$$y' = a(t,y)y, \quad t \geq 0, \qquad y(0) = y_0 \in \mathcal{G},$$

where $a : [0,\infty) \times \mathcal{G} \to \mathfrak{g}$.

Example 2: Isospectral flows:

$$y' = [b(t,y), y], \quad t \geq 0,$$

where $y(0) = y_0 \in \mathrm{Sym}_n[\mathbb{R}]$ and $b : [0,\infty) \times \mathbb{I}[y_0] \to \mathfrak{so}_n[\mathbb{R}]$.

Example 3: Equations on a sphere:

$$\boldsymbol{y}' = \boldsymbol{a}(t,\boldsymbol{y}) \times \boldsymbol{y}, \quad t \geq 0, \qquad \boldsymbol{y}(0) = \boldsymbol{y}_0 \in \mathbb{S}_2,$$

where $\boldsymbol{a} : [0,\infty) \times \mathbb{S}_2 \to \mathbb{R}^3$ and $\times$ is a vector product. E.g., the Lagrange top equations of a rigid body in body coordinates are

$$\boldsymbol{\Pi}_b' = \boldsymbol{\Pi}_b \times \boldsymbol{\Omega}_b + M g\, l\, \boldsymbol{\Gamma}_b \times \boldsymbol{\chi}, \quad \boldsymbol{\Gamma}_b = \boldsymbol{\Gamma}_b \times \boldsymbol{\Omega}_b,$$

where $\boldsymbol{\Pi}_b$ is angular momentum, $\boldsymbol{\Omega}_b$ the angular velocity ($\boldsymbol{\Omega}_b = I_b^{-1}\boldsymbol{\Pi}_b$, I_s the inertia tensor) and $\boldsymbol{\Gamma}_s$ the gravity, while M is the mass, g the gravitational constant, l the distance between the centre of mass and the centre of frame of reference and $\boldsymbol{\chi}$ the unit vector on the body axis.

Trivialization

We solve homogeneous-manifold equations by five conceptual steps, to be implemented in every time step,

1. Transform the equation from $\mathcal{M}$ to $\mathcal{G}$;
2. Transform the equation from $\mathcal{G}$ to $\mathfrak{g}$;
3. Discretize the equation in $\mathfrak{g}$;
4. Transform the outcome from $\mathfrak{g}$ to $\mathcal{G}$;
5. Transform the outcome from $\mathcal{G}$ to $\mathcal{M}$.

The advantage of this approach is that

while in general $\mathcal{M}$ is a *nonlinear* surface, the Lie algebra $\mathfrak{g}$ is a *linear space.* As long as we discretize there with just linear operations and commutators, we are bound to respect its structure!

Steps 1 and 5 are already clear from our discussion, at present we focus on steps 2–4.

A trivialization is a smooth mapping $\Phi : \mathfrak{g} \to \mathcal{G}$ such that $\Phi(0) = I$. The main idea is write the solution p of a Lie-group equation in the form

$$p(t) = \Phi(\omega(t)), \qquad t \geq 0$$

and replace the equation for p by the Lie-algebraic equation for ω.

Formally,

$$\begin{aligned} & p' = a(t,p)p = \mathrm{d}\Phi(\omega, \omega') \\ \Rightarrow \qquad & \omega' = \mathrm{d}\Phi_\omega^{-1} a(t,p) \\ \Rightarrow \qquad & \omega' = \mathrm{d}\Phi_\omega^{-1} a(t, \Phi(\omega)), \end{aligned}$$

with the initial condition $\omega(0) = 0$.

Note that a trivialization might be valid only in (hopefully, large) neighbourhood of the origin: this represents a possible restriction on the step size of the discretization method.

An obvious (and most useful) trivialization is exponentiation,

$$p(t) = \mathrm{e}^{\omega(t)} p_N.$$

The outcome is the *dexpinv equation*

$$\omega = \mathrm{dexp}_{\omega}^{-1} a = \sum_{m=0}^{\infty} \frac{\mathrm{B}_m}{m!} \mathrm{ad}_{\omega}^{m} a, \qquad \omega(t_N) = 0$$

(Hausdorff), where

$$\mathrm{ad}_{\omega}^{m} a = \overbrace{[\omega, [\omega, \ldots, [\omega}^{m \text{ times}}, a] \cdots]]$$

and $\{\mathrm{B}_m\}_{m \geq 0}$ are *Bernoulli numbers,*

$$\sum_{m=0}^{\infty} \frac{\mathrm{B}_m}{m!} z^m = \frac{z}{\mathrm{e}^z - 1} \approx 1 - \tfrac{1}{2} z + \tfrac{1}{12} z^2 - \tfrac{1}{720} z^4 + \tfrac{1}{30240} z^6.$$

Note that the dexpinv equation is always nonlinear and consists of an infinite sum. However, infinite sums can be truncated and nonlinearity is a worthwhile price to pay.

Specifically, the dexpinv equation is

$$\omega' = a - \frac{1}{2}[\omega, a] + \frac{1}{12}[\omega, [\omega, a]] - \frac{1}{720}[\omega, [\omega, [\omega, [\omega, a]]]] + \frac{1}{30240}[\omega, [\omega, [\omega, [\omega, [\omega, [\omega, a]]]]]] + \cdots .$$

For a quadratic Lie group, i.e. $G = \{p : pqp^\top = q\}$, where $q \in \mathrm{Sym}_n[\mathbb{R}]$ is nonsingular, an interesting alternative is the Cayley trivialization

$$\Phi(z) = \frac{1 + \frac{1}{2}z}{1 - \frac{1}{2}z}.$$

It results in the dcayinv equation

$$\omega' = a - \frac{1}{2}[\omega, a] + \frac{1}{4}\omega a \omega.$$

The Cayley trivialization has two advantages *vis-á-vis* the standard exponential: the equation is much simpler and the evaluation of $\Phi(\omega)$ *much* cheaper, in particular for large n. On the other hand, of course, it applies only to quadratic Lie groups.

Runge–Kutta–Munthe-Kaas methods

We apply an RK method to the dexpinv equation, in place of the original Lie-group equation. E.g., instead of the familiar 3rd-order scheme

$$\begin{aligned} k_1 &= a(t_N, y_N)y_N \\ k_2 &= a(t_{N+1/2}, y_N + \tfrac{1}{2}hk_1)(y_N + \tfrac{1}{2}hk_1) \\ k_3 &= a(t_{N+1}, y_N - hk_1 + 2hk_2)(y_N - hk_1 + 2hk_2) \\ \Delta &= h(\tfrac{1}{6}k_1 + \tfrac{2}{3}k_2 + \tfrac{1}{6}k_3) \\ y_{N+1} &= y_N + \Delta, \end{aligned}$$

we use

$$\begin{aligned} k_1 &= a(t_N, y_N) \\ k_2 &= a(t_{N+1/2}, \mathrm{e}^{\frac{1}{2}hk_1} y_N) \\ k_3 &= a(t_{N+1}, \mathrm{e}^{-hk_1+2hk_2} y_N) \\ \Delta &= h(\tfrac{1}{6}k_1 + \tfrac{2}{3}k_2 + \tfrac{1}{6}k_3) \\ y_{N+1} &= \mathrm{e}^{\Delta + \frac{1}{6}h[\Delta, k_1]} y_N. \end{aligned}$$

This is again a 3rd-order method, except that it is guaranteed to evolve on $\mathcal{G}$.

- The beauty of this approach is that we can take any RK method and convert it into a Lie-group method without any major effort.

- Each step of an ν-stage RK–MK method requires ν function evaluations and ν computations of the matrix exponential. The latter can be fairly expensive for high ν.

- The number of commutators increases rapidly with order (in particular if we want really high order methods!). This can be alleviated by techniques from graded Lie algebras, which will be considered later.

- Implicit RK–MK methods can be used, but they require the solution of underlying nonlinear algebraic equations. Owren & Welfert have presented an extension of Newton's method that respects Lie group structure but, again, it ain't cheap.

- Similar approach is valid for other trivializations, in particular for the dcay-inv equations.

Magnus and Magnus-type expansions

Consider (for simplicity) the linear equation

$$y' = a(t)y, \quad t \geq 0, \qquad y(0) \in \mathcal{G}, \quad a(t) \in \mathfrak{g}.$$

Recall that $y(t) = \mathrm{e}^{\omega(t)} y_0$, where ω was given by the dexpinv equation.

Wilhelm Magnus showed that

$$\omega(t) = \int_0^t a(\xi)\mathrm{d}\xi - \frac{1}{2}\int_0^t\int_0^{\xi_1}[a(\xi_2), a(\xi_1)]\mathrm{d}\xi_2\mathrm{d}\xi_1$$

$$+ \frac{1}{4}\int_0^t\int_0^{\xi_1}\int_0^{\xi_2}[[a(\xi_3), a(\xi_2)], a(\xi_1)]\mathrm{d}\xi_3\mathrm{d}\xi_2\mathrm{d}\xi_1$$

$$+ \frac{1}{12}\int_0^t\int_0^{\xi_1}\int_0^{\xi_1}[a(\xi_3), [a(\xi_2), a(\xi_1)]]\mathrm{d}\xi_3\mathrm{d}\xi_2\mathrm{d}\xi_1$$

$$+ \cdots .$$

Extensive use of Magnus expansions in theoretical physics, quantum chemistry, control theory, stochastic DEs, ...

Systematic theoretical analysis and numerical implementation of the Magnus expansion by AI & Nørsett, using graph theory.

- Graph: The pair $\boldsymbol{G} = \langle \boldsymbol{V}, \boldsymbol{E} \rangle$, where $\boldsymbol{V} = \{v_1, \ldots, v_r\}$ are vertices and $\boldsymbol{E} \subseteq \boldsymbol{V} \times \boldsymbol{V}$ the edges;
- Path from v_i to v_j: Ordered set $\{(v_{s_l}, v_{s_{l+1}})\}_{l=1}^{q}$ of distinct vertices s.t. $s_1 = i$ and $s_{q+1} = j$;
- A graph is connected if all vertices connected by a path and it is a tree if such path is unique;
- Rooted tree: The pair $\boldsymbol{T} = (\boldsymbol{G}, w)$, where $\boldsymbol{G}$ is a tree and w (the root) is one of its vertices.
- A rooted tree admits a natural partial ordering of ancestor/successor and parent/child. Note that root has no parent. Vertices with no children are called leaves.

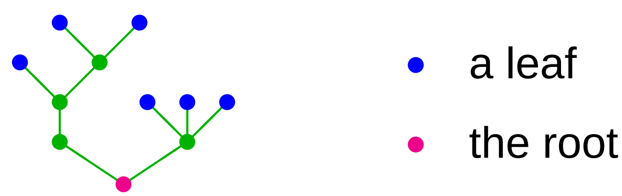

It is easy to see that all terms in a Magnus expansion of ω' are of the form

$$\left[\int \text{simpler term}, \text{another simpler term} \right].$$

This can be expressed by recursion (the proof follows by Picard iteration):

1. We commence from $a(t)$;

2. If $H_1(t)$ and $H_2(t)$ already feature in the expansion, so does

$$\left[\int_0^t H_1(\xi) \mathrm{d}\xi, H_2(t) \right].$$

We model this with *rooted binary trees,*

1. The tree • corresponds to $a(t)$;

2. If $H_{\tau_1} \rightsquigarrow \tau_1$, $H_{\tau_2} \rightsquigarrow \tau_2$ have been already derived, then

$$\left[\int_0^t H_{\tau_1}(\xi) \mathrm{d}\xi, H_{\tau_2}(t) \right] \rightsquigarrow$$

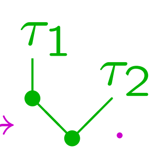

Magnus in terms of trees

Let $\mathbb{T}_m$ be the set of trees with m vertical lines (i.e., corresponding to terms with m integrals). Then

$$\omega(t) = \sum_{m=0}^{\infty} \sum_{\tau \in \mathbb{T}_m} \alpha(\tau) \int_0^t H_\tau(\xi) \mathrm{d}\xi,$$

where the constants $\alpha(\tau)$ are obtained recursively: Any $\tau \in \mathbb{T}_m$, $m \geq 1$, can be written uniquely in the form

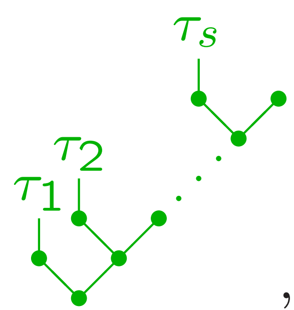

,

where

$$\alpha(\bullet) = 1,$$

$$\alpha(\tau) = \frac{\mathrm{B}_s}{s!} \prod_{l=1}^{s} \alpha(\tau_l).$$

Al, Nørsett & Rasmussen: Appropriately truncated "by power", a Magnus expansion $y_{n+1} = \Upsilon_h y_n$ is time symmetric, i.e. $\Upsilon_h \circ \Upsilon_{-h} = \mathrm{Id}$. Therefore $\Upsilon_h = \mathrm{e}^{\Psi_h}$, where Ψ_h is an odd function and it follows that such an expansion is always of an even order – if we truncate it to produce odd order, we gain for free an extra unit of order.

Sixth-order Magnus expansion:

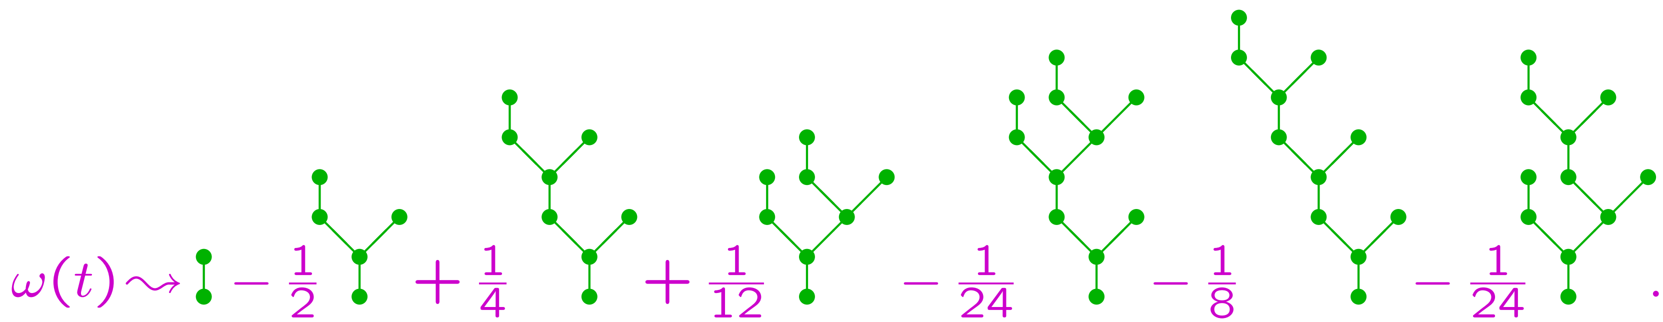

Note that we truncate the expansion by throwing away high-order terms (truncation "by power"), rather than trees of sufficiently high m. This results in much smaller number of surviving trees.

Convergence of the Magnus expansion

Theorem *The Magnus expansion converges if*

$$\int_0^t \|a(\xi)\| \mathrm{d}\xi \le \int_0^{2\pi} \frac{\mathrm{d}\xi}{4 + \xi(1 - \cot\frac{\xi}{2})} \approx 1.086869.$$

Original proof by Blanes, Casas, Oteo & Ros. Shorter proof by Moan: Integration & triangle inequality imply that

$$\|\omega(t)\| \le \int_0^t \|\mathrm{dexp}^{-1}_{\omega(\xi)} a(\xi)\| \mathrm{d}\xi \le \int_0^t \sum_{k=0}^{\infty} \frac{|\mathrm{B}_k|}{k!} (2\|\omega(\xi)\|)^k a(\xi) \mathrm{d}\xi$$
$$= \int_0^t g(2\|\omega(\xi)\|) \|a(\xi)\| \mathrm{d}\xi,$$

where $g(x) = 2 + \frac{x}{2}(1 - \cot\frac{x}{2})$. Bihari-type inequality: Let $h, g, v \in \mathrm{C}[0, t^*)$ positive, g nondecreasing. Then $h(t) \le \int_0^t g(h(\xi)) v(\xi) \mathrm{d}\xi$ implies

$$h(t) \le \tilde{g}^{-1}\left(\int_0^t v(\xi) \mathrm{d}\xi\right), \qquad \tilde{g}(x) = \int_0^x \frac{\mathrm{d}\xi}{g(\xi)}.$$

Letting $h(t) = 2\|\omega(t)\|$, $v(t) = \|a(t)\|$ completes the proof. □

The Fer expansion

We approximate the solution of $y' = a(t)y$ in the form

$$y(t) = \exp\left[\int_0^t a(\xi)\mathrm{d}\xi\right] v(t).$$

Then

$$v' = \mathrm{fer}_{\int_0^t a(\xi)\mathrm{d}\xi} a(t)\, v, \quad t \geq 0, \qquad v(0) = y(0),$$

where

$$\begin{aligned}\mathrm{fer}_b a &= \frac{(I + \mathrm{ad}_b)\mathrm{e}^{-\mathrm{ad}_b} - I}{\mathrm{ad}_b} a \\ &= \sum_{k=1}^{\infty} (-1)^k \frac{k}{(k+1)!} \mathrm{ad}_b^k a.\end{aligned}$$

This procedure can be iterated: the outcome is the Fer expansion

$$y(t) = \mathrm{e}^{\int_0^t \omega_0(\xi)\mathrm{d}\xi} \mathrm{e}^{\int_0^t \omega_1(\xi)\mathrm{d}\xi} \cdots y(0), \qquad t \geq 0.$$

AI: $\omega_m(t) = \mathcal{O}\left(t^{2^{m+2}-2}\right)$, hence the order increases very fast.

Also the Fer expansion can be expanded in rooted trees. Thus, for order 6, just two terms are required and

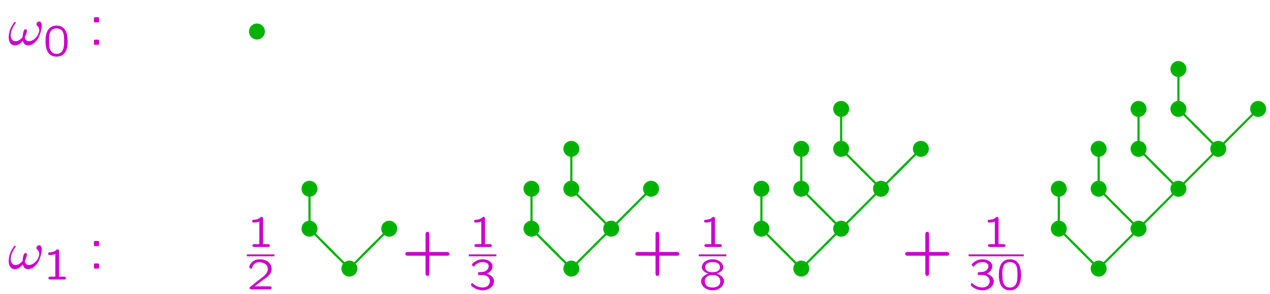

Detailed complexity analysis of Lie-group methods (Celledoni, AI, Nørsett & Orel) demonstrates that in general the Magnus expansion is always cheaper than the Fer expansion. A major reason is that Fer requires more evaluations of the exponential. However, Fer might be useful for ‘niche’ computations.

The Cayley expansion

For a quadratic Lie group with the Cayley trivialization we let

$$y(t) = \mathrm{cay}_{\omega(t)} y_0 = \frac{I + \frac{1}{2}\omega(t)}{I - \frac{1}{2}\omega(t)} y_0, \qquad t \ge 0,$$

consequently

$$\omega' = a - \tfrac{1}{2}[\omega, a] + \tfrac{1}{4}\omega a \omega, \quad t \ge 0, \qquad \omega(0) = 0.$$

In line with the Magnus expansion, we can show that

$$\begin{aligned} \omega(t) = & \int_0^t a(\xi)\mathrm{d}\xi - \tfrac{1}{2}\int_0^t \int_0^{\xi_1} [a(\xi_2), a(\xi_1)]\mathrm{d}\xi_2 \mathrm{d}\xi_1 \\ & + \tfrac{1}{4}\int_0^t \int_0^{\xi_1} \int_0^{\xi_2} [[a(\xi_3), a(\xi_2)], a(\xi_1)]\mathrm{d}\xi_3 \mathrm{d}\xi_2 \mathrm{d}\xi_1 \\ & - \tfrac{1}{4}\int_0^t \int_0^{\xi_1} a(\xi_2)\mathrm{d}\xi_2 a(\xi_1) \int_0^{\xi_1} a(\xi_2)\mathrm{d}\xi_2 \mathrm{d}\xi_1 + \cdots . \end{aligned}$$

This Cayley expansion can be also expanded in rooted trees, except that we require slightly more complicated structures.

AI: Specifically, we have the following composition rules:

1. $\int_0^t a(\xi)\mathrm{d}\xi \rightsquigarrow$  is a term.

2. If $H \rightsquigarrow \tau$ is a term, then so is

$$\int_0^t [H(\xi), a(\xi)]\mathrm{d}\xi \rightsquigarrow$$ 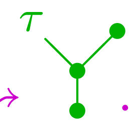

.

3. If $H_1 \rightsquigarrow \tau_1$ and $H_2 \rightsquigarrow \tau_2$ are terms, then so is

$$\int_0^t H_1(\xi)a(\xi)H_2(\xi)\mathrm{d}\xi \rightsquigarrow$$ 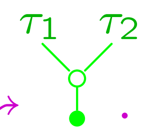

.

We obtain

$\omega(t) \rightsquigarrow$  $- \frac{1}{2}$ 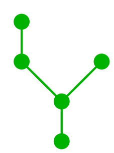 $+ \frac{1}{4}$ 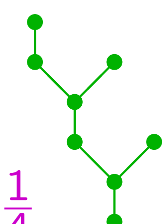 $- \frac{1}{4}$ 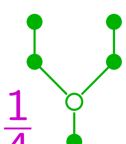 $+ \cdots .$

Multivariate quadrature

Each term in a Magnus, Fer or Cayley expansion corresponds to a multivariate integral over a different polytope. In principle, multivariate integration is very expensive: we need many function evaluations for *every* polytope. Fortunately, everything simplifies!

All integrals in question are of the form

$$\mathscr{I}(h) = \int_{h\mathcal{S}} \mathcal{L}(a(\xi_1), a(\xi_2), \ldots, a(\xi_r)) \mathrm{d}\boldsymbol{\xi},$$

where $\mathcal{L}$ is a multilinear form and $\mathcal{S}$ is a polytope,

$$\mathcal{S} = \{\boldsymbol{x} \in \mathbb{R}^r \,:\, 0 \le x_k \le x_{i_k}, \quad k = 1, 2, \ldots, r\},$$

with $x_0 = 1$ and $i_k \le k - 1$, $x_0 = 1$, $k = 1, 2, \ldots, r$. Specifically, for the first few integrals we have the following polytopes $\mathcal{S}$ and multilinear forms $\mathcal{L}$:

Tree	r	$\mathcal{L}(z_1,\ldots,z_r)$	$\mathcal{S}$
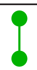	1	z_1	
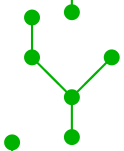	2	$[z_2,z_1]$	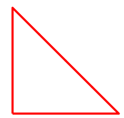
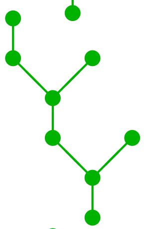	3	$[[z_3,z_2],z_1]$	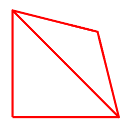
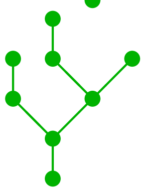	3	$[z_3,[z_2,z_1]]$	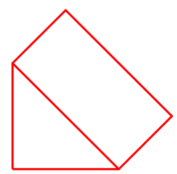

Let $c_1,\ldots,c_\nu \in [0,1]$ be distinct quadrature points and $a_k = a(t_N + c_k h)$. We approximate $\mathscr{I}(h)$ with

$$\mathscr{Q}(h) = h^r \sum_{\boldsymbol{i}\in C_\nu^r} \beta_{\boldsymbol{i}} \mathcal{L}(a_{i_1}, a_{i_2}, \ldots, a_{i_r}),$$

where $C_\nu^r = [1,2,\ldots,\nu]^r$. Note that we repeatedly recycle just ν function evaluations of a!

Theorem *(AI & Nørsett) The order of quadrature is* $\nu + s$*, where*

$$\int_0^1 \zeta^{i-1} q(\zeta) \mathrm{d}\zeta = 0, \qquad i = 1, 2, \ldots, s,$$

and $q(x) = \prod_{k=1}^{\nu} (x - c_k)$.

The proof is long and complicated. We'll prove instead a simpler statement. For this (and for much future use) we require the

Alekseev–Gröbner Lemma *Given the ODE* $y' = f(t, y)$*, suppose that* u *is any* C^1 *function s.t.* $u(t_N) = y(t_N)$*. Then*

$$u(t) - y(t) = \int_{t_N}^{t} \Phi(t - x)[f(x, u(x)) - u'(x)] \mathrm{d}x,$$

where Φ *is the solution of the variational equation* $\Phi' = \frac{\partial f}{\partial y} \Phi$*,* $\Phi(0) = I$*.*

Theorem *(Zanna) The order of the approximation to truncated Magnus expansion with the above quadrature is* $\nu + s$*.*

To prove this, we replace the matrix a by its interpolating polynomial $\hat{a}$ s.t. $\hat{a}(t_N + c_k h) = a(t_N + c_k h)$, $k = 1, \ldots, \nu$. Let $u' = \hat{a}u$, $u(t_N) = y(t_N)$. Therefore, by the A–G Lemma

$$u(t_{N+1}) = y(t_{N+1}) + \int_{t_N}^{t_{N+1}} \Phi(t_{N+1} - x)[\hat{a}(x) - a(x)]u(x)\mathrm{d}x.$$

The integrand vanishes at $c_1, \ldots, c_\nu$. Therefore, once we discretize the integral by the quadrature formula, all the is left is the quadrature error, i.e. $\mathcal{O}\left(h^{\nu+s+1}\right)$. □

Corollary *If* $c_1, \ldots, c_\nu$ *are Gauss–Legendre points in* $[0, 1]$ *then the order is* 2ν.

Example: A 4th-order approximantion

$$\int_0^h \int_0^{\xi_1} [a(\xi_2), a(\xi_1)]\mathrm{d}\boldsymbol{\xi} \approx \frac{\sqrt{3}}{6} h^2 [a_1, a_2],$$

where $c_1 = \frac{1}{2} - \frac{\sqrt{3}}{6}$, $c_2 = \frac{1}{2} + \frac{\sqrt{3}}{6}$.

To construct a quadrature formula in greater detail, let $\ell_k \in \mathbb{P}_{\nu-1}$ be the kth cardinal polynomial of Lagrangian interpolation,

$$\ell_k(c_k) = 1, \qquad \ell_k(c_j) = 0, \quad j \neq k.$$

Then

$$\beta_{\boldsymbol{i}} = \int_{\mathcal{S}} \prod_{k=1}^{r} \ell_{i_k}(\xi_k) \mathrm{d}\boldsymbol{\xi}.$$

GOOD NEWS

It takes just ν evaluations of the matrix a to compute all the integrals in the expansion to order 2ν.

BAD NEWS

The numbers of commutators increases very fast, in line with the exponential growth in the number of combinations in C_ν^r. Therefore, the volume of linear algebra required in every time step becomes prohibitive for large orders.

Graded Lie algebras

The number of commutators can be vastly reduced by an idea due to Munthe-Kaas & Owren.

Replace $ha_1, \ldots, ha_\nu$ by $b_0, b_1, \ldots, b_{\nu-1}$, where

$$\sum_{l=0}^{\nu-1} \frac{1}{l!}(c_k - \tfrac{1}{2})^l b_l = ha_k, \qquad k = 1, 2, \ldots, \nu.$$

Then $b_l \approx h^{l+1} a^{(l)}(\frac{1}{2}h) = \mathcal{O}(h^{l+1})$: the term b_l is of grade $l+1$. The grades propagate naturally under commutation: if the grade of x_i is w_i for $i = 1, 2$, then the grade of $[x_1, x_2]$ is $w_1 + w_2$. Therefore, the term

$$\mathcal{L}(b_{i_1}, b_{i_2}, \ldots, b_{i_r})$$

is of grade $|\boldsymbol{i}| = \sum i_l$. We obtain the order-$(2\nu)$ quadrature

$$\widetilde{\mathcal{Q}}(h) = \sum_{|\boldsymbol{i}| \le 2\nu} \tilde{\beta}_{\boldsymbol{i}} \mathcal{L}(b_{i_1}, b_{i_2}, \ldots, b_{i_r}),$$

where

$$\tilde{\beta}_{\boldsymbol{i}} = \int_{\mathcal{S}} \prod_{l=1}^{r} (\xi_l - \tfrac{1}{2})^{i_l} \mathrm{d}\xi_r \cdots \mathrm{d}\xi_1.$$

Free Lie algebras

For future use we require the concept of a Free Lie algebra. Given generators $\Phi = \phi_1, \phi_2, \ldots, \phi_\nu$, we say that

$$\mathfrak{F} = \mathsf{FLA}(\phi_1, \phi_2, \ldots, \phi_\nu)$$

is the free Lie algebra generated by Φ if it is the closure of the generators under commutation and linear combinations. Trivially, it is a Lie algebra.

We attach to each ϕ_l the grade $g(\phi_l) = \kappa_l \in \mathbb{N}$ and let grades propagate by commutation:

$$g(t_1) = \mu_1,\; g(t_2) = \mu_2 \qquad \Rightarrow \qquad g([t_1, t_2]) = g(t_1) + g(t_2).$$

We denote by $\mathscr{K}_m^\nu$ the linear space of all terms in the FLA of grade m. Of course $\mathfrak{F} = \mathsf{Sp}\{\mathscr{K}_m^\nu \,:\, m \geq 1\}$, although we will not make much use of it.

The surprisingly small dimension of $\mathscr{K}_m^\nu$ is the key to our goal, to reduce the number of commutators.

The great commutator throw-away

Three mechanisms allow us to get rid of most terms:

1. All terms of grade $\geq 2\nu + 1$ are not required and can be thrown out.
2. It is possible to show that the discrete expansion is time symmetric. Therefore sums of terms of even grade vanish and we don't need to compute them.
3. The classical Witt–Birkhoff formula gives the (surprisingly small) dimension of $\mathscr{K}_m^\nu$ for unit grades $\kappa_l \equiv 1$ (corresponding to $x_l = h a_l$). Munthe-Kaas & Owren have extended the Witt–Birkhoff formula to arbitrary grades. Specifically, let λ_k, $k = 1, 2, \ldots, s = \max \kappa_l$, be the zeros of the polynomial

$$1 - \sum_{i=1}^{\nu} z^{\kappa_i}.$$

Then

$$\dim \mathscr{K}_m^\nu = \frac{1}{m} \sum_{d|m} \mu(d) \sum_{i=1}^{s} \lambda_i^{m/d},$$

where μ is the Möbius function,

$$\mu(d) = \begin{cases} 1, & d = 1, \\ (-1)^q, & n_i = 1,\ i = 1, 2, \ldots, q, \\ 0, & \text{otherwise}, \end{cases}$$

where $d = p_1^{n_1} p_2^{n_2} \cdots p_q^{n_q}$ is the prime-number decomposition of d.
They also gave a recursive procedure for the formation of a basis of $\mathscr{K}_m^\nu$ (a generalisation of Hall and Lyndon bases).

The number of required terms:

order:	2	4	6	8	10
Naive	1	5	80	3304	1256567
Clever	1	2	7	22	73

Is this the optimal Magnus expansion?

Further economy is possible by clever algorithmic tricks and aggregating terms. This has been worked out for even orders 4–8 by Blanes, Casas & Ros.

For example, for order 6, we compute

$$
\begin{aligned}
&c_1 = \tfrac{1}{2} - \tfrac{\sqrt{15}}{10}, \quad c_2 = \tfrac{1}{2}, \quad c_3 = \tfrac{1}{2} + \tfrac{\sqrt{15}}{10}, \\
&a_k = ha(t_N + c_k), \quad k = 1, 2, 3, \\
&b_1 = ha_2, \qquad b_2 = \tfrac{\sqrt{15}}{3}(a_3 - a_1), \\
&b_3 = \tfrac{10}{3}(a_3 - 2a_1 + a_1), \\
&v_1 = [b_1, b_2], \\
&v_2 = [b_1, 2b_3 + v_1], \\
&v_3 = [-20b_1 - b_3 + v_1, b_2 - \tfrac{1}{60}v_2], \\
&\omega_N = b_1 + \tfrac{1}{12}b_3 + \tfrac{1}{240}v_3, \\
&y_{N+1} = \mathrm{e}^{\omega_N} y_N.
\end{aligned}
$$

Thus,

just three function evaluations, three commutators and one exponential!

The matrix exponential

It is not enough to "approximate" the exponential: we must do so while mapping an element from $\mathfrak{g}$ to $\mathcal{G}$! Otherwise all the hard work in designing Lie-group methods would have been in vain.

Standard methods for the calculation of the matrix exponential:

- Eigenvalue/eigenvector decomposition:

  Exceedingly expensive (in particular for large systems) and ill conditioned.

- Polynomial and rational approximations:

  We cannot expect the result to live in $\mathcal{G}$, with one exception: when the Lie group is quadratic and we use diagonal Padé approximations (Feng Kang). On the other hand, no rational approximation can map $\mathfrak{sl}(n)$ to $\mathrm{SL}(n)$ for $n \geq 3$.

- Krylov subspace approximation:

  Very effective means for large matrices (Hochbruck & Lubich). However, there is no reason for them to map a Lie algebra to 'its' Lie group.

Splitting methods (Celledoni & AI)

Given $a \in \mathfrak{g}$, we split its exponential,

$$\mathrm{e}^{ta} \approx \mathrm{e}^{tb_1}\mathrm{e}^{tb_2}\cdots\mathrm{e}^{tb_s}$$

such that

1. $\sum_k b_k = a$, each b_k lives in $\mathfrak{g}$;
2. It is easy to evaluate each e^{tb_k} exactly; and
3. The overall error is consistent with the order of the Lie-group method

Note that necessarily $\mathrm{e}^{tA} \in \mathcal{G}$.

If $b_k = b_{s+1-k}\ \forall k$ then the approximation is time symmetric. In that case the order can be further enhanced with the *Yošida device.*

Practical splittings decompose a into low-rank matrices, using the following explicit formula for the matrix exponential. Suppose that u and v are both $n \times r$ matrices, $r \ll n$, and $c = uv^\top$. Then

$$\mathrm{e}^{tc} = I + tud^{-1}(\mathrm{e}^{td} - I)v^\top,$$

where $d = v^\top u$ is $r \times r$, hence 'small'.

For example, for $\mathfrak{g} = \mathfrak{so}_n$ we let $r = 2$, $s = n - 1$, set

$$b^{[0]} = a = [\boldsymbol{b}_1^{[0]}, \boldsymbol{b}_2^{[0]}, \ldots, \boldsymbol{b}_n^{[0]}]$$

and choose

$$b_1 = \boldsymbol{b}_1^{[0]}\boldsymbol{e}_1^\top - \boldsymbol{e}_1\boldsymbol{b}_1^{[0]\top} \in \mathfrak{so}(n)$$

and $b^{[1]} = b^{[0]} - b_1$. Observe that the first row and column of $b^{[1]}$ vanish. We continue in this manner, similar to LU factorization, letting $b^{[i]} = b^{[i-1]} - b_i$ and

$$b_i = \boldsymbol{b}_i^{[i-1]}\boldsymbol{e}_i^\top - \boldsymbol{e}_i\boldsymbol{b}_i^{[i-1]\top} \in \mathfrak{so}_n, \quad i = 1, \ldots, n-1.$$

Note that $\mathrm{rank}\, b_i = 2$.

The Yošida device

Let Φ_h be a time-reversible numerical method of order p for the solution of $y' = f(t, y)$.

Time reversibility implies that p is even.

We consider a new method,

$$\Psi_h = \Phi_{\alpha h} \circ \Phi_{\beta h} \circ \Phi_{\alpha h}.$$

1. The method Ψ_h is time reversible.

2. Suppose that

$$\alpha = \frac{1}{2 - 2^{1/(p+1)}}, \qquad \beta = -\frac{2^{1/(p+1)}}{2 - 2^{1/(p+1)}}.$$

   Then Ψ_h is of order $p + 2$.

This approach (originally developed by Haruo Yošida for symplectic methods) can be used to improve the order of our approximation of the exponential.

Coordinates of the second kind (Celledoni & AI)

Let $d = \dim \mathfrak{g}$ and let $\{\phi_1, \phi_2, \ldots, \phi_d\}$ be its basis. We seek scalar functions $\theta_1, \theta_2, \ldots, \theta_m$ s.t.

$$\mathrm{e}^{ta} = \mathrm{e}^{\theta_1(t)\phi_1} \mathrm{e}^{\theta_2(t)\phi_2} \ldots \mathrm{e}^{\theta_d(t)\phi_d}.$$

Such functions always exist but their practical evaluation in closed form is virtually impossible for $d \geq 3$ (Wei & Norman). Alternatively, we approximate the functions θ_k to requisite order.

The structure constants of $\mathfrak{g}$ are numbers

$$c^j_{k,l}, \quad k, l, j = 1, \ldots, d, \qquad \text{such that} \qquad [\phi_k, \phi_l] = \sum_{j=1}^{d} c^j_{k,l} \phi_j.$$

The Taylor expansion of each θ_k can be derived explicitly in terms of the structure constants (which, of course, depend on the choice of the basis). Let

$$a = \sum_{k=1}^{d} \alpha_k \phi_k.$$

We have

$$
\begin{aligned}
\theta_k(0) &= 0, \\
\theta_k'(0) &= \alpha_k, \\
\theta_k''(0) &= \sum_{l=1}^{d} \sum_{j=1}^{l-1} \alpha_l c_{l,j}^k \alpha_j, \\
\theta_k'''(0) &= 2 \sum_{l=1}^{d} \sum_{j=1}^{l-1} c_{l,j}^k [\theta_l''(0)\alpha_j + \theta_j''(0)\alpha_l] \\
&\quad + 2 \sum_{l=1}^{d} \sum_{j=1}^{l-1} \sum_{i=1}^{j-1} \sum_{m=1}^{d} c_{l,j}^m c_{i,m}^k \alpha_l \alpha_j \alpha_i + \sum_{l=1}^{d} \sum_{j=1}^{d} \sum_{i=1}^{l-1} c_{l,i}^j c_{i,j}^k \alpha_l \alpha_i^2
\end{aligned}
$$

and so on. More and more summation!

Assume, though, that the structure constants are sparse: they are almost all zero. In that case the cost reduces a very great deal! This can be accomplished by choosing as our basis a root space decomposition of $\mathfrak{g}$: the number of summations typically drops by a factor of two and the cost of a second-order approximation in $\mathfrak{so}(n)$ is just $\mathcal{O}(n)$.

Generalized polar decomposition (Munthe-Kaas, Quispel & Zanna)

An involutory automorphism of a Lie group $\mathcal{G}$ is a one-to-one map $\sigma : \mathcal{G} \to \mathcal{G}$ s.t.

$$\begin{aligned} &\sigma(x \cdot y) = \sigma(x) \cdot \sigma(y), \qquad x, y \in \mathcal{G}, \\ &\sigma(\sigma(x)) = x, \qquad x \in \mathcal{G}. \end{aligned}$$

Each $\mathcal{G}$-automorphism can be lifted to the underlying Lie algebra, resulting in a $\mathfrak{g}$-automorphism,

$$\mathrm{d}\sigma(a) = \frac{\mathrm{d}}{\mathrm{d}t} \sigma(\mathrm{e}^{ta})|_{t=0}.$$

We use $\mathrm{d}\sigma$ to define the sets

$$\begin{aligned} \mathfrak{k} &= \{a \in \mathfrak{g} \,:\, \mathrm{d}\sigma(a) = a\}, \\ \mathfrak{p} &= \{a \in \mathfrak{g} \,:\, \mathrm{d}\sigma(a) = -a\}. \end{aligned}$$

While $\mathfrak{k}$ is a Lie subalgebra, $\mathfrak{p}$ is a Lie triple system: a linear space s.t.

$$a, b, c \in \mathfrak{g} \qquad \Rightarrow \qquad [a, [b, c]] \in \mathfrak{g}.$$

We have

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

Specifically, $\mathfrak{g} \ni a = p + k$, where

$$p = \tfrac{1}{2}[a - \mathrm{d}\sigma(a)], \qquad k = \tfrac{1}{2}[a + \mathrm{d}\sigma(a)].$$

Note that

$$\begin{aligned} a, b \in \mathfrak{k} : &\qquad [a,b] \in \mathfrak{k}, \\ a, b \in \mathfrak{p} : &\qquad [a,b] \in \mathfrak{k}, \\ a \in \mathfrak{k}, b \in \mathfrak{p} : &\qquad [a,b] \in \mathfrak{p}, \\ a \in \mathfrak{p}, b \in \mathfrak{k} : &\qquad [a,b] \in \mathfrak{p}. \end{aligned}$$

It is possible to prove that the Lie group can be factorized into generalized polar decomposition

$$\mathcal{G} \ni z = xy, \qquad \sigma(y) = y, \quad \sigma(x) = x^{-1}.$$

At the algebra level, this corresponds to

$$\mathcal{G} \ni \mathrm{e}^{c} = \mathrm{e}^{a}\mathrm{e}^{b}, \qquad a \in \mathfrak{p}, \quad b \in \mathfrak{k}.$$

The main idea is to obtain the leading terms of a, b from p, k.

We have

$$
\begin{aligned}
a = & \, pt - \tfrac{1}{2}[p,k]t^2 - \tfrac{1}{6}[k,[p,k]]t^3 \\
& + (\tfrac{1}{24}[p,[p,[p,k]]] - \tfrac{1}{24}[k,[k,[p,k]]])t^4 + \mathcal{O}\left(t^5\right), \\
b = & \, kt - \tfrac{1}{12}[p,[p,k]]t^3 + \mathcal{O}\left(t^5\right).
\end{aligned}
$$

Suppose now that $\dim \mathfrak{p}$ is very small. In that case e^a is very cheap to compute. This is not the case with e^b, but recall that $\mathfrak{k}$ is a subalgebra! Hence, we can go on splitting it!

We thus obtain a sequence of algebra authomorphisms and low-dimensional triple Lie systems $\mathfrak{p}_k$ for $k = 1, 2, \ldots, m$, such that

$$
\mathrm{e}^a = \mathrm{e}^{p_1}\mathrm{e}^{p_2}\cdots\mathrm{e}^{p_m}, \qquad p_k \in \mathfrak{p}_k.
$$

A convenient way of generating such $\mathrm{d}\sigma_k$ is through involutory inner authomorphisms

$$
\sigma(x) = sxs^\top, \qquad \mathrm{d}\sigma(a) = sas^\top,
$$

where $s \in \mathcal{G} \cap \mathrm{O}(n)$.

Let $s_j = \mathrm{diag}(1 - 2\mathrm{e}_j)$. Then elements in each $\mathfrak{p}_j$ are nonzero just along the j row and column, hence $\mathrm{rank}\,\mathfrak{p}_j \equiv 2$. We are 'peeling' the matrix a from the top down.

Al & Zanna: Suppose that we have brought a to an upper Hessenberg form, $a = q\tilde{a}q^\top$, where $q \in \mathrm{O}(n)$. The above algorithm does not respect this form. But suppose that we 'peel' the matrix from the bottom, not from the top! In that case, moving to each subsequent subalgebra $\mathfrak{k}_j$ 'contaminates' just few elements under the first subdiagonal, at the bottom of the matrix, which can be 'cleaned' with fre Givens rotations.

This yields an algorithm which is competitive with classical methods to compute the matrix exponential, even when the conservation of Lie-group structure is not at issue!