

UNIVERSITY OF CAMBRIDGE

Numerical Analysis Reports

**Approximately preserving symmetries in
numerical integration**

Arieh Iserles, Robert McLachlan & Antonella Zanna

DAMTP 1997/NA22
December 1997

Department of Applied Mathematics and Theoretical Physics
Silver Street
Cambridge
England CB3 9EW

Approximately preserving symmetries in numerical integration

Arieh Iserles,^{*} Robert McLachlan[†] and Antonella Zanna[‡]

December 5, 1997

Abstract

We present a general procedure for recursively improving the invariance of a numerical integrator φ under a symmetry group. If h is a symmetry, we construct the *adjoint method* $h^{-1}\varphi h$. In each time step we apply either the original method or the adjoint method, according to a prescription based on the *Thue-Morse sequence*. The outcome is a solution sequence which displays progressively smaller symmetry errors, to any desired order in the time step. The method can also be used to force the solution to stay close to a desired submanifold of phase space, while retaining structural properties of the original method.

Keywords: symmetry, differential equations, numerical integration, geometric integration, composition methods

1 Orthogonal flows

A recent study of the numerical integration of matrix differential equations (Calvo, Iserles & Zanna 1996, Calvo, Iserles & Zanna 1997*b*, Dieci, Russell & van Vleck 1994, Iserles 1997), explored methods that keep the solution in a desired solution space: for example, that render the computed matrix orthogonal whenever the exact solution is orthogonal. This is expensive to do exactly, but a method of alternating between the original equation and its adjoint was discovered which preserves orthogonality not exactly, but up to any given order in the time step.

An *orthogonal flow* is a matrix differential equation of the form

$$\dot{y} = a(t, y)y, \quad t \geq 0, \quad y(0) = y_0, \quad (1.1)$$

where $y_0 \in \mathrm{O}(n)$, the orthogonal group, and a maps $\mathbb{R}^+ \times \mathrm{O}(n)$ to $\mathfrak{so}(n)$, the set of $n \times n$ skew-symmetric matrices (the Lie algebra of $\mathrm{O}(n)$). It is well-known that $y(t) \in \mathrm{O}(n)$ for all $t \geq 0$ and the conservation of this feature under discretization is

^{*}Department of Applied Mathematics and Theoretical Physics, Silver Street, Cambridge CB3 9EW, UK.

[†]Department of Mathematics, Massey University, Private Bag 11-222, Palmerston North, New Zealand.

[‡]Newnham College, University of Cambridge, Cambridge CB3 9DF, UK.

important to a wide range of applications, not least to the calculation of Lyapunov exponents, isospectral flows and problems in numerical algebra (Dieci et al. 1994).

It is known that certain Runge–Kutta methods can integrate (1.1) whilst conserving orthogonality, but all such methods are necessarily implicit, hence excessively costly (Calvo, Iserles & Zanna 1997a, Dieci et al. 1994). An alternative approach has been explored in (Calvo et al. 1996). Given an arbitrary Runge–Kutta method of order p and solution sequence $\{y_m\}_{m \in \mathbb{Z}^+}$, where $y_m \approx y(m\Delta t)$, the *orthogonality error* $e_m := \|y_m^T y_m - I\|$ is (locally) $\mathcal{O}((\Delta t)^{p+1})$. Suppose, however, that we use (1.1) to advance from y_{2m} , whilst employing the *adjoint equation*

$$\dot{x} = -xa(t, x^{-1}) \quad (1.2)$$

to advance from $x_{2m+1} = y_{2m+1}^{-1}$, setting $y_{2m+2} = x_{2m+2}^{-1}$. It is instructive to consider the function a in (1.2) as belonging to the left Lie algebra of $O(n)$, while the same function in (1.1) is in the right Lie algebra (cf. (Olver 1995) for Lie-group terminology).

Although the Runge–Kutta method alternating between (1.1) and (1.2) is of the same order p , it has been proved in (Calvo et al. 1996) that $e_{2m} = \mathcal{O}((\Delta t)^{p+2})$. It has been further suggested in (Calvo et al. 1997b) that this approach can be used to increase progressively the order of the orthogonality error. Let us denote a single step by the original orthogonal flow (1.1) by ‘0’ and a step of the adjoint equation (1.2) by ‘1’. We have already mentioned that the sequence ‘000000 ...’ results in $e_m = \mathcal{O}((\Delta t)^{p+1})$, while the outcome of ‘010101 ...’ is $e_{2m} = \mathcal{O}((\Delta t)^{p+2})$. By the same token, ‘01100110 ...’ yields $e_{4m} = \mathcal{O}((\Delta t)^{p+3})$, ‘01101001 ...’ results in $e_{8m} = \mathcal{O}((\Delta t)^{p+4})$ and so on. Numerical experiments in (Iserles 1997) affirm that, after initial perturbations, the orthogonality error is rapidly decreasing and often is driven down to machine accuracy.

2 Automorphisms of numerical integrators

In the present paper we interpret and explore this approach in the language of symmetries of differential equations, and generalise it to a technique which preserves (again, up to any order in the time step) any group of symmetries of a differential equation. A consequence is that the fixed set of the symmetries, which should be invariant under the dynamics, is kept invariant up to any order in the time step. Orthogonality preservation is an example of such a fixed set. Three other examples considered here are symmetry preservation in isospectral flows such as the Toda lattice; $f \leftrightarrow g$ symmetry in splitting methods for the equation $\dot{y} = f + g$; and spatial symmetries of finite-difference operators for PDEs.

Since we aim to make our methods and analysis as general as possible, our discussion must somehow involve the space of all numerical methods. Human ingenuity being unlimited, this space is difficult to handle; it must be delimited somehow. One approach that has been used fruitfully was presented in (McLachlan, Quispel & Turner 1997), and we adopt it here as well. (Indeed, we view the present paper as a further advertisement for the approach.)

We let M be a manifold and consider differential equations

$$\dot{y} = f(y), \quad y \in M \quad (2.1)$$

with phase space M . (Since we're computing here, we shall often have $M = \mathbb{R}^n$.) Let F be a set of such vector fields, and let Φ be a set of functions

$$\varphi : F \times [-T, T] \rightarrow \text{Diff}(M)$$

$$\varphi_{\Delta t}(f) : y \rightarrow \tilde{y} \quad (f \in F, \Delta t \in [-T, T]).$$

The desirable members of Φ are the ones for which $\tilde{y}$ is close to the time- Δt flow of f , although not all members have this property. The functional dependence of the diffeomorphism φ on its arguments Δt and f are what we think of as a ‘numerical method’; however, one can also freeze these arguments and study the resulting map (the ‘dynamics of numerics’ point of view, see (Stuart & Humphries 1996)), or study its image point $\tilde{y} = \varphi_{\Delta t}(f)y$. A flexible perspective is needed.

The final step is to let the set of maps

$$\{\varphi_{\Delta t}(f) : \varphi \in \Phi, f \in F, \Delta t \in \mathbb{R}\}$$

generate a group, G , where the group operation is composition of diffeomorphisms. Notice that G has elements in which different time steps and vector fields are mixed. G being a group of diffeomorphisms, it has a Lie algebra $\mathfrak{g}$ of vector fields.¹

In many of our examples, the phase space M is *itself* a Lie group, such as $O(n)$ or $GL(n)$. It is important to keep this group conceptually separate from G itself, which is a group of diffeomorphisms of M .

Many properties of numerical methods can be related to properties of the group G . Moreover, many of *those* are related to particular automorphisms and antiautomorphisms of G . Given a group G and a function $\mathcal{A} : G \rightarrow G$, we let

$$\text{fix } \mathcal{A} = \{\varphi \in G : \varphi = \mathcal{A}\varphi\}$$

stand for the group of the *fixed points* of $\mathcal{A}$. A function $\mathcal{A}_+ : G \rightarrow G$ is said to be an *automorphism* if

$$\mathcal{A}_+(\varphi_1\varphi_2) = (\mathcal{A}_+\varphi_1)(\mathcal{A}_+\varphi_2), \quad \varphi_1, \varphi_2 \in G,$$

while $\mathcal{A}_- : G \rightarrow G$ is an *antiautomorphism* if

$$\mathcal{A}_-(\varphi_1\varphi_2) = (\mathcal{A}_-\varphi_2)(\mathcal{A}_-\varphi_1), \quad \varphi_1, \varphi_2 \in G.$$

Fixing $\psi \in G$, standard examples of automorphisms and antiautomorphisms are the *conjugation* $\mathcal{A}_+\varphi = \psi^{-1}\varphi\psi$ and the *inversion* $\mathcal{A}_-\varphi = \varphi^{-1}$ respectively. Noting that (anti)automorphisms have a natural group operation, namely composition, we write $\text{Aut}_{\pm}(G)$ for a group of (anti)automorphisms acting on G . If at least one member is an antiautomorphism, then the automorphisms $\text{Aut}_+(G)$ form a normal subgroup of index 2. The antiautomorphisms $\text{Aut}_-(G)$ do not form a subgroup. An important property is that

$$\varphi \in \text{fix } \text{Aut}_+(G) \quad \Rightarrow \quad \varphi\mathcal{A}_-\varphi \in \text{fix } \text{Aut}_{\pm}(G) \quad (2.2)$$

¹Actually, making this statement precise rapidly becomes extremely technical. All we need, however, is that the diffeomorphisms can be arbitrarily well approximated by flows, to within any power of their distance from the identity. See (Reich 1997).

for an arbitrary $\mathcal{A}_- \in \text{Aut}_-(G)$. Here

$$\text{fix Aut}(G) = \bigcap_{\mathcal{A} \in \text{Aut}(G)} \text{fix } \mathcal{A}.$$

This feature, a way of constructing group elements that are invariant under all of $\text{Aut}_\pm(G)$, is known as the *generalized Scovel projection* and was used in (McLachlan et al. 1997) to construct numerical integrators that preserve the reversing symmetry group of a differential equation.

In the application to numerical methods, important operations are the antiautomorphism ‘inverse’, $\mathcal{I}\varphi := \varphi^{-1}$, and the automorphism ‘reverse’, $\mathcal{R}\varphi_{\Delta t} := \varphi_{-\Delta t}$. (Note that $\varphi \in \text{fix } \mathcal{IR} \Rightarrow \varphi_{\Delta t}\varphi_{-\Delta t} = \mathbb{1}$, that is, the method is ‘time-symmetric’ or ‘self-adjoint’ (Hairer & Wanner 1991, Yoshida 1990).) Another is ‘ h -conjugation’, $\mathcal{N}_h\varphi := h^{-1}\varphi h$, where $h \in \text{Diff}(M)$ is a given diffeomorphism (for example, a change of variables). The fixed set of $\mathcal{N}_h$ are the equivariant maps $\varphi h = h\varphi$; in other words, they have h as a symmetry.

Notice that until now, we have not acted on the differential equation itself, the argument f of the functions φ . One way to do this is to let a diffeomorphism h act on f (by push-forward, or change of variables): $\mathcal{V}_h\varphi(f) := \varphi((dhf)h^{-1})$. (Usually, $\mathcal{V}_h\varphi(f)$ is not an integrator of f —it’s acting on the wrong differential equation.) Now, instead of changing variables in the differential equation, we should perhaps have changed them in the map as well; so we need the ‘coordinate transformation’ automorphism

$$\mathcal{K}_h\varphi(f) := \mathcal{N}_h\mathcal{V}_h\varphi(f) = h^{-1}\varphi((dhf)h^{-1})h$$

as well. Intuitively, changing variables, integrating in the new variables, then changing back, should give the same answer as integrating in the original variables (that is, we should have $\varphi \in \text{fix } \mathcal{K}_h$), since this is true for the exact flow of a differential equation. However, this is not true in general for numerical integrators.

If H is a group of transformations, the automorphisms $\mathcal{K}_h$, $\mathcal{V}_h$, and $\mathcal{N}_h$ have the further automorphic property that, e.g., $\mathcal{N}_{gh} = \mathcal{N}_g\mathcal{N}_h \ \forall g, h \in H$. We say that φ is *H -covariant* if $\varphi \in \text{fix } \mathcal{K}_h$ for all $h \in H$. Different methods can then be classified by their covariance group. For example, Runge–Kutta methods are covariant with respect to the affine group $H = \text{A}(n) := \{x \mapsto Ax + b : A \in \text{GL}(n, \mathbb{R}), b \in \mathbb{R}^n\}$, which justifies their classification as ‘linear’ methods. It’s when we start playing around with the set of candidate vector fields F , the phase space M , and the group G that problems arise. For example, splitting methods are a good way of constructing φ ’s that are symplectic or volume-preserving; but they are not linearly covariant.

Why is covariance desirable in the first place? The full picture is not yet clear, except in the special case that H is the symmetry group of the differential equation f . Then $(dhf)h^{-1} = f \ \forall h \in H$, so $\mathcal{K}_h\varphi(f) = \mathcal{N}_h\varphi(f)$ and H -covariance implies H -symmetry of the integrator φ . (The advantage of considering covariance instead of just symmetry is that it can be studied independently of any special vector fields f ; it’s a ‘one level higher’ property.) Thus, Runge–Kutta methods preserve any linear or affine symmetries that any differential equation may happen to have. However, more general methods must be considered as well, because (i) symmetries may not

be globally linearisable, such as the matrix symmetry $A \mapsto A^{-T}$ considered below in Section 4; (ii) if a system has a group of symmetries, they may not be simultaneously linearisable; (iii) even a linear symmetry group may not be preserved by some methods, such as splitting methods, which are useful for their other conservation properties.

If H is a symmetry group of f , then $\varphi(f)$ has the same symmetry group if it is H -covariant (equivalently, if it is H -equivariant). Thus, it is desirable to construct φ s which are fixed under given automorphism groups $\text{Aut}_+(G)$. Indeed, this is necessary even to apply the generalized Scovel projection (2.2)—note that (2.2) makes no mention of how to ensure that $\varphi \in \text{fix } \text{Aut}_+$. There is one trick which is sometimes useful: If φ is in a linear space and the action of $\text{Aut}_+(G)$ is linear, one can sum over the group:

$$\overline{\varphi} := \sum_{\mathcal{A} \in \text{Aut}_+(G)} \mathcal{A}\varphi$$

satisfies $\overline{\varphi} = \text{Aut}_+(G)\overline{\varphi}$. However, the spaces of symplectic maps, isospectral maps, orthogonal matrices and so on are not linear.

3 Iterative error reduction

The general question is, for a group G and automorphism group $\text{Aut}_+(G)$, how can one systematically construct group elements $\varphi \in \text{fix } \text{Aut}_+(G)$? Even in the simplest case of a single automorphism $\mathcal{A}$ satisfying $\mathcal{A}^2 = \mathbb{1}$ (i.e., $\mathcal{A}$ being an involution) one immediately observes that the analogue of the generalized Scovel projection, namely $\psi = \varphi\mathcal{A}\varphi$, will not do, since $\mathcal{A}\psi = (\mathcal{A}\varphi)\varphi \neq \psi$. But to discard ψ entirely would be a big mistake.

The central observation of this paper is the following: *if φ is near the identity, then $\psi := \varphi\mathcal{A}\varphi$ is more nearly fixed by $\mathcal{A}$ than φ is.* We measure this by the *discrepancy*

$$d(\varphi) := d(\varphi, \mathcal{A}) := \log[\varphi(\mathcal{A}\varphi)^{-1}].$$

Proposition 1 *Let $\mathcal{A}$ be an involutory automorphism and assume that $\varphi \in G$ is close to the identity, i.e. $\varphi = e^{\mathcal{O}(\varepsilon)}$ for sufficiently small ε . Letting $\psi = \varphi\mathcal{A}\varphi$, it is true that*

$$d(\psi) = \mathcal{O}(\varepsilon d(\varphi)).$$

In particular, $d(\varphi) = \mathcal{O}(\varepsilon)$ implies that $d(\psi) = \mathcal{O}(\varepsilon^2)$.

Proof Let $\mathfrak{g}$ be the Lie algebra of G . We can write

$$\varphi = e^{\varepsilon S + \delta N},$$

where $S \in \mathfrak{g}$ is chosen so that $e^{\varepsilon S} \in \text{fix } \mathcal{A}$ and $\delta = \mathcal{O}(\varepsilon)$. It follows that

$$\mathcal{A}\varphi = e^{\varepsilon S + \delta M}$$

for some $M \in \mathfrak{g}$. Note that, using the Baker–Campbell–Hausdorff (BCH) formula (Yoshida 1990)

$$\begin{aligned} d(\varphi) &= \log[\varphi(\mathcal{A}\varphi)^{-1}] = \log[e^{\varepsilon S + \delta N} e^{-(\varepsilon S + \delta M)}] \\ &= \delta(N - M) - \frac{1}{2}[\varepsilon S + \delta N, \varepsilon S + \delta M] + \mathcal{O}(\varepsilon^3) \\ &= \delta(N - M) - \frac{1}{2}\varepsilon\delta[S, N + M] + \mathcal{O}(\varepsilon^3) = \mathcal{O}(\varepsilon), \end{aligned}$$

where $[\cdot, \cdot] : \mathfrak{g} \times \mathfrak{g} \rightarrow \mathfrak{g}$ is the Lie bracket of vector fields. On the other hand, $\mathcal{A}$ being an involutory automorphism, it is true that $\mathcal{A}(\varphi\mathcal{A}\varphi) = \mathcal{A}\varphi\mathcal{A}(\mathcal{A}\varphi) = (\mathcal{A}\varphi)\varphi$, hence

$$e^{d(\psi)} = (\varphi\mathcal{A}\varphi)(\mathcal{A}(\varphi\mathcal{A}\varphi))^{-1} = (\varphi\mathcal{A}\varphi)((\mathcal{A}\varphi)\varphi)^{-1} = \varphi(\mathcal{A}\varphi)\varphi^{-1}(\mathcal{A}\varphi)^{-1},$$

a group commutator. However, if $x, y = \mathcal{O}(\varepsilon)$, then according to the BCH formula, $\log(e^x e^y e^{-x} e^{-y}) = [x, y] + \mathcal{O}(\varepsilon^3)$, so

$$d(\psi) = [\varepsilon S + \delta N, \varepsilon S + \delta M] + \mathcal{O}(\varepsilon^3) = \varepsilon\delta[S, M - N] + \delta^2[N, M] + \mathcal{O}(\varepsilon^3) = \mathcal{O}(\varepsilon\delta),$$

thereby proving the proposition. $\square$

Let us consider the following chain: $\psi_0 = \varphi$ and $\psi_{m+1} = \psi_m \mathcal{A} \psi_m$, $m \in \mathbb{Z}^+$. Given that $\psi_0 = e^{\mathcal{O}(\varepsilon)}$, we have just deduced that $d(\psi_1) = \mathcal{O}(\varepsilon d(\psi_0))$. It follows from the method of proof of the proposition that $\psi_1 = e^{\mathcal{O}(\varepsilon)}$ and this can be extended to all ψ_m . Therefore $d(\psi_{m+1}) = \mathcal{O}(\varepsilon d(\psi_m))$ and we readily deduce by induction that

$$d(\psi_m) = \mathcal{O}(\varepsilon^m d(\varphi)), \quad m \in \mathbb{Z}^+. \quad (3.1)$$

Note that each ψ_m entails 2^m applications of φ . Moreover, $\mathcal{A}$ being an involutory automorphism, it is true that

$$\begin{aligned} \psi_0 &= \varphi = \mathcal{A}^0 \varphi, \\ \psi_1 &= \varphi \mathcal{A} \varphi = (\mathcal{A}^0 \varphi)(\mathcal{A}^1 \varphi), \\ \psi_2 &= \psi_1 \mathcal{A} \psi_1 = (\mathcal{A}^0 \varphi)(\mathcal{A}^1 \varphi) \mathcal{A}((\mathcal{A}^0 \varphi)(\mathcal{A}^1 \varphi)) = (\mathcal{A}^0 \varphi)(\mathcal{A}^1 \varphi)(\mathcal{A}^1 \varphi)(\mathcal{A}^0 \varphi). \end{aligned}$$

In general, suppose that

$$\psi_m = \prod_{i=1}^{2^m} \mathcal{A}^{l_{m,i}} \varphi,$$

where $l_{m,i} \in \{0, 1\}$, $i = 1, 2, \dots, 2^m$. Then

$$\psi_{m+1} = \left(\prod_{i=1}^{2^m} \mathcal{A}^{l_{m,i}} \varphi \right) \mathcal{A} \left(\prod_{i=1}^{2^m} \mathcal{A}^{l_{m,i}} \varphi \right) = \left(\prod_{i=1}^{2^m} \mathcal{A}^{l_{m,i}} \varphi \right) \left(\prod_{i=1}^{2^m} \mathcal{A}^{l_{m,i}+1} \varphi \right).$$

However, $\mathcal{A}$ being an involution, $\mathcal{A}^{l_{m,i}+1} = \mathcal{A}^{l'_{m,i}}$, where $l'_{m,i}$ is the binary complement of $l_{m,i}$. Therefore

$$l_{m+1,i} = l_{m,i}, \quad l_{m+1,2^m+i} = l'_{m,i}, \quad i = 1, 2, \dots, 2^m, \quad (3.2)$$

allowing us to identify the set $\{l_{m,i}\}$.

The ordered set $\{\sigma_i\}_{i \in \mathbb{N}}$ is said to be the *Thue–Morse sequence* if $\sigma_1 = 0$ and each subsequence $\sigma_1, \sigma_2, \dots, \sigma_{2^m}$ therein is followed by its binary complement: hence, 0, 1, 1, 0, 1, 0, 0, 1, $\dots$ (Lothaire 1983). It follows from (3.2) that $l_{m,i} = \sigma_i$, $i = 1, 2, \dots, 2^m$.

An alternative definition of the Thue–Morse sequence is that $\sigma_i = 0$ whenever the number of ones in the binary representation of the integer $(i - 1)$ is even, $\sigma_i = 1$ otherwise. This is an easy consequence of (3.2), since the number of ones in the binary representation of $(2^m + i - 1)$ differs in parity from the number of ones in the binary representation of $(i - 1)$.

Our discussion can be readily generalized from a single automorphism to a finite group Γ of automorphisms. Letting $\Gamma = \{\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_r\}$, we extend the concept of discrepancy in an obvious manner,

$$d(\varphi, \Gamma) = \max_{l=1,2,\dots,r} d(\varphi, \mathcal{A}_l), \quad \varphi \in G.$$

Proposition 2 *Let φ be close to the identity and set*

$$\psi = \prod_{l=1}^r \mathcal{A}_l \varphi \tag{3.3}$$

Then

$$d(\psi, \Gamma) = \mathcal{O}(\varepsilon d(\varphi, \Gamma)).$$

Proof Choosing a nonzero $S \in \mathfrak{g}$ such that $e^{\varepsilon S} \in \text{fix } \Gamma$, it is true that $\mathcal{A}_l \varphi = e^{\varepsilon S + \delta M_l}$, $M_l \in \mathfrak{g}$, $l = 1, 2, \dots, r$. Let $\mathcal{B} \in \Gamma$. Then $\mathcal{B}\mathcal{A}_l = \mathcal{A}_{m(l)}$, $l = 1, 2, \dots, r$, hence $\mathcal{B}\mathcal{A}_l \varphi = e^{\varepsilon S + \delta M_{m(l)}}$. We deduce that

$$\begin{aligned} \psi(\mathcal{B}\psi)^{-1} &= (\mathcal{A}_1 \varphi)(\mathcal{A}_2 \varphi) \cdots (\mathcal{A}_r \varphi) \{ \mathcal{B}[(\mathcal{A}_1 \varphi)(\mathcal{A}_2 \varphi) \cdots (\mathcal{A}_r \varphi)] \}^{-1} \\ &= (\mathcal{A}_1 \varphi)(\mathcal{A}_2 \varphi) \cdots (\mathcal{A}_r \varphi) [(\mathcal{B}\mathcal{A}_1 \varphi)(\mathcal{B}\mathcal{A}_2 \varphi) \cdots (\mathcal{B}\mathcal{A}_r \varphi)]^{-1} \\ &= (\mathcal{A}_1 \varphi)(\mathcal{A}_2 \varphi) \cdots (\mathcal{A}_r \varphi) (\mathcal{A}_{m(r)} \varphi)^{-1} \cdots (\mathcal{A}_{m(2)} \varphi)^{-1} (\mathcal{A}_{m(1)} \varphi)^{-1} \\ &= e^{\varepsilon S + \delta M_1} e^{\varepsilon S + \delta M_2} \cdots e^{\varepsilon S + \delta M_r} e^{-\varepsilon S - \delta M_{m(r)}} \cdots e^{-\varepsilon S - \delta M_{m(2)}} e^{-\varepsilon S - \delta M_{m(1)}}. \end{aligned}$$

It is now trivial to prove, inductively using the BCH formula, that

$$\psi(\mathcal{B}\psi)^{-1} = \exp \left[\delta \sum_{l=1}^r (M_l - M_{m(l)}) + \mathcal{O}(\varepsilon \delta) \right].$$

Γ being a group, the mapping $m : \{1, 2, \dots, r\} \rightarrow \{1, 2, \dots, r\}$ is one-to-one. Therefore

$$\sum_{l=1}^r (M_l - M_{m(l)}) = 0,$$

completing the proof of the proposition. $\square$

For example, if $\Gamma \cong \mathbb{Z}_3$ we can represent ψ by the sequence 012 and ψ_2 by the sequence 012 120 201, and so on. In this context the Thue–Morse sequence is just the

simplest of a set of similar sequences associated to each finite group, it being the one associated to $\mathbb{Z}_2$.

If φ were more nearly fixed by some elements of Γ than by others, then those could be omitted from ψ .

Often one can compute not just the single element $\varphi = e^{\varepsilon S + \mathcal{O}(\varepsilon^2)}$, but $\varphi_\varepsilon = e^{\varepsilon S + \mathcal{O}(\varepsilon^2)}$ for any $\varepsilon > 0$. This is true for numerical integrators, where ε is related to the time step. In this case one can ask whether some specially chosen composition of $\mathcal{A}^j \varphi_{c_i \varepsilon}$ might not achieve a similar symmetry error with fewer stages than the Thue–Morse sequence. In fact, this is possible, but, unlike the Thue–Morse sequence, the compositions now depend on the relative size of δ and ε and on the group Γ . We analyse a few such cases in Section 7.

4 Application to symmetries

In the application to differential equations with a symmetry group H , we have $\Gamma = \{\mathcal{N}_h : h \in H\}$, and $\mathcal{N}_h \varphi(f) = h^{-1} \varphi(f) h$, therefore

$$\psi(f) = \prod_{h \in H} h^{-1} \varphi(f) h.$$

Because in this case the flow of f is in $\text{fix } \mathcal{N}_h$, ψ has the same order as φ .

For example, if H is the cyclic group of order n generated by the element h , hence $H = \{1, h, h^2, \dots, h^{n-1}\}$, then it is readily affirmed that

$$\psi = (h^{-(n-1)} \varphi h^{n-1}) \dots (h^{-2} \varphi h^2) (h^{-1} \varphi h) \varphi = (h \varphi)^n.$$

In other words, one simply ‘twists’ x a little each step, and after n steps, the symmetry error is reduced by $\mathcal{O}(\varepsilon)$. Using $\psi_2 = (h(h \varphi)^n)^n$, one twists twice every n steps.

We have left open the question of whether ψ_n should be considered the map approximating one time step (i.e., we take $\varepsilon = 2^{-n}(\Delta t)$) or 2^n time steps (i.e., $\varepsilon = \Delta t$). Naturally the latter is preferable, but may not be possible for stability reasons. Let φ and h be linear maps. Although φ and $h^{-1} \varphi h$ have the same eigenvalues, their composition may not. For example, if

$$\varphi : x \mapsto \begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix} x, \quad h : x \mapsto \begin{bmatrix} \alpha & 1 - \alpha^2 \\ 1 & -\alpha \end{bmatrix} x$$

(where h is an involution), then φ has trace 0 and is elliptic, but $\varphi h^{-1} \varphi h$ has trace $2 + \alpha^4$ and is hyperbolic for $\alpha \neq 0$. In this example, it is clear that φ is not near enough to the identity. How near is ‘near enough’ is a stability question.

Further to the question of choice of time step, consider ψ_n as a method with 2^n stages. It contains ψ_j for $j < n$ as sub-methods. If δ is the symmetry error in the basic method φ , the local symmetry error after 2^j stages is $\mathcal{O}(\varepsilon^j \delta)$. The *global* symmetry error is one power of ε less throughout, or $\mathcal{O}(\varepsilon^{\min(j, n-1)} \delta)$. Thus, we can take output every 2^{n-1} stages and still have global errors of $\mathcal{O}(\varepsilon^{n-1} \delta)$.

Example Let $M_n[\mathbb{R}]$ be the set of real $n \times n$ matrices, $l \in M_n[\mathbb{R}]$, $b : M_n[\mathbb{R}] \rightarrow M_n[\mathbb{R}]$ and consider the matrix differential equation

$$\dot{l} = [b(l), l] = b(l)l - lb(l) \quad (4.1)$$

whose flow is isospectral, i.e., the spectrum of $l(t)$ is constant in time. Numerical integrators have been developed that also preserve the spectrum of l (Calvo et al. 1997a, Calvo et al. 1996, Calvo et al. 1997b). Indeed, the very idea of alternating between (1.1) and (1.2) has originated in this context. Here we are just interested in symmetries, since the most important applications of isospectral flows (4.1), not least to Toda lattice equations (Toda 1981) and to problems in numerical linear algebra (Chu 1994) pertain to symmetric matrices. If $b(l) = -b(l^T)^T$, then $T : l \mapsto l^T$ is a symmetry of (4.1). Then, if the initial condition $l(0)$ is symmetric, it remains so for all time. This symmetry is an involution, so the hierarchy of methods is

$$\psi_0 = \varphi, \quad \psi_1 = \varphi T \varphi T, \quad \psi_2 = \varphi T \varphi T \varphi, \quad \dots$$

This example is continued and tested in the next example, on orthogonal flows.

Having looked at symmetries, which entail the invariance of the fixed set $\text{fix } h = \{y \in M : y = h(y)\}$, we now turn the question around and consider directly requiring $\text{fix } h$ to be invariant. The two are not necessarily equivalent. For a differential equation, invariance of $\text{fix } h$ is equivalent to $f = (dhf)h^{-1}$ for all $y \in \text{fix } h$; for a map, it is equivalent to $\varphi = h^{-1}\varphi h$ for all $y \in \text{fix } h$ (i.e., *not* for all $y \in M$.) To take a simple example, let $y \in \text{GL}(n)$, $b : \text{GL}(n) \rightarrow \text{O}(n)$, and $h : y \mapsto y^{-T}$, so that $\text{fix } h = \text{O}(n)$. Then $\varphi : y \mapsto b(y)y$ *always* leaves $\text{O}(n)$ invariant (just a statement of the group property of $\text{O}(n)$) but it only has h as a symmetry if $b(y^{-T}) = b(y)$ for all $y \in \text{GL}(n)$.

The beauty of the method described above—namely, $\psi_{n+1} = \psi_n \mathcal{N}_h \psi_n$ —is that *it still works!* (This is shown formally below.) This may seem paradoxical at first. Indeed, we first thought that one should use $\mathcal{K}_h \varphi$, because this is always has the same order as φ as an integrator. But methods in $\text{fix } \mathcal{K}_h$ do not necessarily leave $\text{fix } h$ invariant, as one can easily check.

Roughly, trying to drive φ towards $\mathcal{N}_h \varphi$ helps to preserve the invariance of $\text{fix } h$ even when h is not a symmetry of the differential equation. Of course, it is dangerous to enforce something in a numerical integrator which is not even satisfied by the flow of the differential equation. The following propositions state that in the present case this course of action is justified. Proposition 3 shows that the symmetry error is under control, and Proposition 4 shows that the truncation error is under control.

It is convenient to consider the case when phase space M is a Lie group, as we then have a natural measure of closeness in M : x is close to y if $\log(y^{-1}x) \ll 1$. Of course, this includes the important case when M is the abelian group $\mathbb{R}^n$, in which case $\log(y^{-1}x) = x - y$.

Proposition 3 *Let the phase space M be a Lie group and let $\varphi \in G$ be nearly invariant under $\mathcal{N}_h$, i.e. $d(\varphi, \mathcal{N}_h) = \mathcal{O}(\delta)$ where $\delta \ll 1$. Let the initial condition y be nearly invariant under h , i.e. $\eta := \log(y^{-1}h(y)) \ll 1$. Then the discrepancy satisfies $\tilde{\eta} := \log(\tilde{y}^{-1}h(\tilde{y})) = \mathcal{O}(\eta, \delta)$, where $\tilde{y}$ is the value of the flow at the next time step.*

Proof Let $\Delta = \exp(d(\varphi, \mathcal{N}_h)) = \varphi h^{-1} \varphi^{-1} h$, $y \in M$, and $E = y^{-1} h(y)$. Δ and E are both near to the identity in their respective spaces. From these definitions we have

$$\begin{aligned}\varphi(yE) &= \varphi(hy) = \Delta^{-1}(h\varphi(y)) = \Delta^{-1}(h\tilde{y}) = \Delta^{-1}(\tilde{y}\tilde{E}) \\ &\Rightarrow \tilde{E} = \tilde{y}^{-1} \Delta(\varphi(yE))\end{aligned}$$

Now if $m \in \text{Diff}(M)$, then $\log(m(y)^{-1} m(yE)) = \mathcal{O}(\log E)$ —a statement of the differentiability of m . Equivalently, $m(yE) = m(y)E_1$ where $\log E_1 = \mathcal{O}(\log E)$. So

$$\tilde{E} = \tilde{y}^{-1} \Delta(\varphi(y)E_1) = (\tilde{y}^{-1} \Delta(\tilde{y}))(E_2) := \tilde{\Delta}(E_2)$$

where $\log(E_1), \log(E_2) = \mathcal{O}(\log(E))$. Since $\tilde{y}$ is also an element of $\text{Diff}(M)$ (acting by left multiplication), we see that $\tilde{\Delta}$ is conjugate to Δ and therefore $\log(\tilde{\Delta}) = \mathcal{O}(\log(\Delta))$. Taking logs of both sides proves the proposition. $\square$

Proposition 4 *Let f be a vector field with $(dhf)h^{-1} = f$ when $y = h(y)$. Let $\varphi_{\Delta t}(f)$ be an integrator for f with local truncation error $\mathcal{O}((\Delta t)^{p+1})$. Let $\eta = \log(y^{-1} h(y))$. Then $\mathcal{N}_h \varphi$ is an integrator for f of order $\mathcal{O}((\Delta t)^{p+1}, \eta)$.*

Proof We have

$$\varphi(f) = e^{\Delta t f + (\Delta t)^{p+1} F + \mathcal{O}(\Delta t)^{p+2}} = e^{(\Delta t)^{p+1} F} e^{\Delta t f} e^{\mathcal{O}((\Delta t)^{p+2})}$$

where $\text{fix } h$ is invariant under the flow of f , but not necessarily under the flow of F . Then

$$\mathcal{N}_h \varphi(f) = h^{-1} \varphi(f) h = (h^{-1} e^{(\Delta t)^{p+1} F} h) (h^{-1} e^{\Delta t f} h) e^{\mathcal{O}((\Delta t)^{p+2})}$$

The first factor is $e^{\mathcal{O}((\Delta t)^{p+1})}$, and when applied to the initial condition y , the second factor is $e^{\Delta t f} e^{\mathcal{O}(\eta)}$. This proves the result. $\square$

After time $\mathcal{O}(1)$, the global symmetry error η is $\mathcal{O}(\delta/\Delta t)$. If $\delta = \mathcal{O}((\Delta t)^{p+1})$ (i.e., the leading term in the truncation error is not symmetric), this term dominates the local truncation error, leading to global truncation errors of $\mathcal{O}((\Delta t)^p + \delta/(\Delta t)^2)$. Thus there would be an order reduction if we tried to use the method $\mathcal{N}_h \varphi$ by itself, as one would expect. Fortunately, this order reduction does not appear if the leading term in the truncation error is symmetric, as it is in the method $\psi = \varphi \mathcal{N}_h \varphi$. This method has $\delta = (\Delta t)^{p+2}$ or smaller.

Corollary 5 *If φ has order p , and $\eta = 0$ initially, then $\psi = \varphi \mathcal{N}_h \varphi$ has order p .*

Proof Locally, ψ is the composition of two methods of order p . Globally, symmetry errors do not dominate over times $\mathcal{O}(1)$. $\square$

Of course, there *is* a problem if η is actually larger than the truncation error. This can only happen if the differential equation did not leave $\text{fix } h$ invariant in the first

place. In this curious situation the methods ψ_n still succeed in staying on fix h with the expected error, but they are not approximating the flow of f any longer.

Note that the application of a Thue–Morse-type solver in (Calvo et al. 1997b) is different from the above: the isospectral manifold therein is considered as a homogeneous space under the action of the special orthogonal group $\mathrm{SO}(n)$, whereby the problem of solving (4.1) is reduced to an equation of type (1.1). This guarantees the recovery of isospectrality but, with most practical methods, leads to breakdown in symmetry. It is the latter phenomenon that motivated the original discussion of this type of ‘alternating solvers’.

Example We consider again the differential equation (1.1):

$$\dot{y} = a(t, y)y, \quad y(0) = y_0, \quad (4.2)$$

except that we no longer assume that $y_0 \in \mathrm{O}(n)$ and that $a : \mathbb{R}^+ \times \mathrm{O}(n) \rightarrow \mathfrak{o}(n)$. If H is a matrix Lie group with Lie algebra $\mathfrak{h}$, given $y_0 \in H$ and $a : \mathbb{R}^+ \times H \rightarrow \mathfrak{h}$, it is well-known that, in this case, $y(t) \in H$ for all $t \geq 0$. For the purposes of numerical integration we extend the phase space from H to all of $\mathrm{GL}(n)$.

We concern ourselves in the present example with *quadratic* Lie groups $\{y \in \mathrm{GL}(n) : ypy^T = y\}$, which are the fixed set of the involution

$$h(y) = py^{-T}p^{-1}, \quad y \in \mathrm{GL}(n),$$

H will be invariant under (4.2) if $a(t, y) \in \mathfrak{h}$ for $y \in H$. However, to integrate in GL the function a needs to be defined on all of GL . Different lifts of this type will in principle lead to different results. However, we do not find that requiring, say, $a(t, y) \in \mathfrak{h}$ for all $y \in \mathrm{GL}$, or for (4.2) to have h as a symmetry, necessarily leads to better results.

Since $h(h(y)) = pp^{-T}yp^Tp^{-1}$, the function h is an involution if and only if $pp^{-T} = \pm I$. Hence for convenience we restrict p to either a symmetric or a skew-symmetric matrix.

If we change variables to $x = h(y)$, then we find that x obeys the differential equation

$$\dot{x} = (dhf)h^{-1} = -pa(t, h^{-1}(x))^T p^{-1}x, \quad x(0) = h(y_0). \quad (4.3)$$

If also $a \in \mathfrak{h}$, i.e., $pa^T = -ap$, then

$$\dot{x} = a(t, h^{-1}(x))x, \quad x(0) = h(y_0). \quad (4.4)$$

If also $a(t, h^{-1}(x)) = a(t, x)$ then h is a symmetry of (4.2) and we have

$$\dot{x} = a(t, x)x, \quad x(0) = h(y_0). \quad (4.5)$$

We can now implement the Thue–Morse technique and explore different variants of the ‘adjoint method’ $\mathcal{A}\varphi$. Stepping with (4.3) is what we have called $\mathcal{K}_h\varphi$. If a is chosen appropriately, this is equal to stepping with (4.5), i.e., mapping y before and after the step, but using the same differential equation. The recommended method is to step with (4.5), *regardless* of the function a .




Figure 1: The discrepancy for $\psi_3 = \text{'01101001'}$ applied to Eq. (4.6), where method 0 is φ and method 1 is $\mathcal{N}_h \varphi$. Here $\Delta t = \frac{1}{200}$. Each eighth step (i.e., a complete cycle) is denoted by a cross.

Note that all of these methods are different from (1.2), which was the method we set out to ‘explain’! The difference is related to the partial covariance of the methods used with (1.2), and is explained in the Appendix. Eq. (1.2) turns out to be equivalent to using $\mathcal{K}_h \varphi$.

The two most important applications of the technique are when $p = I$, i.e. $H = O(n)$, and when, for even n ,

$$p = \begin{bmatrix} O & I \\ -I & O \end{bmatrix}, \quad p^2 = -I,$$

whereby H becomes $\text{Sp}(n)$, the *symplectic group*.

Figs. 1–3 illustrate the results on these methods in one case. We first consider the system (4.2) with

$$b(y) = l_+ - l_-, \quad l = y l_0 y^{-1}, \quad (4.6)$$

where l_+ (l_-) is the upper (lower) triangular part of l , and l_0 is a fixed symmetric matrix. We have chosen this example because if l is tridiagonal, then it is the well-known Toda lattice. This system can be regarded as a lift of the Toda lattice to GL. This system leaves $O(n)$ invariant, and $y l_0 y^{-1}$ tends to a diagonal matrix with the




Figure 2: Methods ψ_0 (Euler), $\psi_1 = \psi_0 \mathcal{N}_h \psi_0$, $\psi_2, \dots, \psi_4$ compared for (4.6). Dotted lines connect sample points. Solid lines show predicted slope (symmetry order) of the method, namely η (the symmetry error after a fixed time $t = 4$) = $\mathcal{O}((\Delta t)^{n+1})$ for method ψ_n .

eigenvalues of l_0 on the diagonal. Furthermore, $-b(y^{-T})^T = b(y)$, so $h(y) = y^{-T}$ is a symmetry of this system.

We took $n = 3$, $y_0 = I$ and used Euler as the basic method,

$$\varphi : \tilde{y} = y + \Delta t b(y)y,$$

so that $\delta = \varepsilon^2$. First we confirmed numerically that $\mathcal{K}_h \varphi$ and $\mathcal{N}_h \varphi$ are indeed equivalent in this case, as explained above. (They are both equivalent to (1.2), see the Appendix.) We then explored the hierarchy of methods $\psi_0, \dots, \psi_4$, taking $\varepsilon = \Delta t$, measuring the discrepancy $\log_{10}[yy^T]$. Figure 1 shows the results for $\psi_3 = \text{'01101001'}$. The dramatic error reduction as the Thue-Morse sequence completes each ‘sweep’ of length 2^j is evident. One can also see how, after about 32 time steps—just four complete cycles—the errors after $4n$ and after $8n$ steps are about equal, so one could usefully take output every four steps, instead of eight as might be expected. This is the ‘double sampling’ phenomenon explained above. As Δt is decreased, the number of time steps required to reach this balance increases.

Figure 2 compares all five methods, showing that ψ_n has a global symmetry error of $\mathcal{O}((\Delta t)^{n+1})$, as expected.




Figure 3: As in Fig. 2: methods $\psi_0, \dots, \psi_4$ compared for $\mathcal{A} = \mathcal{N}_h$ (crosses) and $\mathcal{A} = \mathcal{K}_h$ (circles) for a nonsymmetric system. Only $\mathcal{N}_h$ has the good order behaviour.

We then studied a different system which did not have h as a symmetry. Let l_j be the j th diagonal of l . Take $b(y) = l_1 - l_1^T$ still leaves $O(n)$ invariant, but does not have h as a symmetry. Figure 3 compares the two choices $\mathcal{A} = \mathcal{N}_h$ and $\mathcal{A} = \mathcal{K}_h$. It shows that $\mathcal{N}_h$ works as expected. However, $\mathcal{K}_h$ still works quite well, and is indistinguishable for ψ_1 and ψ_2 . It seems to gain one or two powers of ε easily and then ‘saturate’. We have not explained this observation, although we saw it in many different systems.

Example Let $H = \text{SL}(n)$, the *special linear group* of $n \times n$ real matrices with unit determinant. We set

$$h(y) = \frac{1}{(\det y)^{2/n}} y, \quad y \in \text{GL}(n).$$

It is trivial to observe that h is an involutory automorphism with fixed set $\pm H$. For good discretization schemes the departure from the Lie group is likely to be small, hence it is highly unlikely for the time-stepping scheme to cross into the wrong component of the Lie group $\{y \in \text{GL}(n) : \det y = \pm 1\}$. This is an important observation, since many Lie groups of practical interest, not least $O(n)$, consist of a number of disconnected components, one of which is itself a Lie group. Our remark is valid in this, more general, situation.

It is fortunate that our method $\varphi \mathcal{N}_h \varphi$ does not require changing variables in the differential equation, for $(dhf)h^{-1}$ is extremely complicated in this example.




Figure 4: The discrepancy for (4.7) with a step-size $\frac{1}{50}$. The top figures display the discrepancy for each solution sequence at points $m = 2^l$, $l = 0, 1, 2$, while the bottom figure depicts $\log_{10} d(y_n)$ for the three sequences.

Let us consider the system

$$\begin{aligned} \dot{y}_{1,1} &= y_{1,1}^2 y_{2,2}, & \dot{y}_{1,2} &= y_{1,1} y_{1,2} y_{2,2}, \\ \dot{y}_{2,1} &= -y_{2,1} - y_{1,2} y_{2,1}^2, & \dot{y}_{2,2} &= -y_{2,2} - y_{1,2} y_{2,1} y_{2,2}, \end{aligned} \quad (4.7)$$

with $y_0 \in \text{SL}(2)$. It is an easy matter to verify that (4.7) can be written in the form (1.1), such that $\text{tr } a(y) = 0$ for $y \in \text{SL}(2)$. However, $\text{tr } a(y) \neq 0$ in general, and (4.7) is very far from having h as a symmetry! Figures 4 and 5 display the error in the determinant when (4.7) is solved by the forward Euler method

$$y_{m+1} = [1 + ha(y_m)]y_m, \quad m \in \mathbb{Z}^+,$$

with stepsizes $h = \frac{1}{50}$ and $h = \frac{1}{100}$ respectively. We have displayed the discrepancy $\delta_l = d(y_n)$ of the ‘naive’ solution sequence 000000... ($l = 0$), the sequence 010101... ($l = 1$) and 01100110... ($l = 2$). The behaviour of discrepancies is consistent with our analysis.

5 Splitting methods

One can consider applying the underlying idea of Section 2 not to ordinary symmetries, but to symmetries of *pairs* of differential equations, as arising in splitting methods. Let $\varphi(f, g)$ be a method applied to the pair of vector fields f and g . The flow of $\dot{x} = f + g$ has the symmetry $\varphi(f, g) = \varphi(g, f)$ and one can consider requiring this of numerical integrators as well. For example, in a linear space one can use $e^{\varepsilon f} e^{\varepsilon g} + e^{\varepsilon g} e^{\varepsilon f}$, which has this symmetry. The symmetry has no dynamical significance in general, although it may do in particular cases, for example if f and g represent alternating directions of a partial differential equation. Most commonly-used composition methods do not have this symmetry.

To apply the symmetrising composition, let $\mathcal{A}\varphi(f, g) := \varphi(g, f)$, an involutory automorphism. Each application of $\mathcal{A}$ increases the order of symmetry by one. If $\varphi(f, g) = e^{\varepsilon f}$, then $\psi_1 = \varphi\mathcal{A}\varphi = e^{\varepsilon f} e^{\varepsilon g}$ is symmetric to first order in ε , $\psi_2 = e^{\varepsilon f} e^{2\varepsilon g} e^{\varepsilon f}$ is symmetric to second order (it is the leapfrog method), and so on. Let us take various second-order compositions, scale step lengths so that all give $e^{\varepsilon(f+g)}$ at first order, and expand them in series as $\exp(\varepsilon(f+g) + \varepsilon^3(c_1[f, [f, g]] + c_2[g, [g, f]])) + \dots$ (see (McLachlan 1995) for a survey of similar composition methods).

method	c_1	c_2	stages	scaled error	stability limit
ψ_2	$-1/24$	$1/12 \approx 0.0833$	3	0.7500	2
ψ_3	$1/192$	$1/192 \approx 0.0052$	6	0.1875	2.472
ρ	0.0061	0.0061	5	0.1520	2.544
ξ	0.0101	0.0101	4	0.1616	2.264

Table 1: Symmetrised compositions: coefficients, scaled error and stability limits.

Table 1 displays a number of interesting symmetrised compositions. Here

$$\rho = e^{z\varepsilon f} e^{1/2\varepsilon g} e^{(1-2z)\varepsilon f} e^{1/2\varepsilon g} e^{z\varepsilon f}, \quad z = \frac{6 - \sqrt{5}}{8} \approx 0.1910$$

is the five-stage, second-order composition with least error; note that the error is minimised when the method is $\mathcal{A}$ -symmetric to third order, which indicates another usefulness of this symmetry. It also has the ‘time symmetry’ $\rho(\varepsilon)\rho(-\varepsilon) = 1$. The method

$$\xi = e^{w\varepsilon f} e^{(1-w)\varepsilon g} e^{(1-w)\varepsilon f} e^{w\varepsilon g}, \quad w = \frac{1}{\sqrt{2}},$$

is the four-stage, second-order composition with smallest error. It is not time symmetric, but it can be factored into $\zeta\mathcal{IR}\mathcal{A}\zeta$, where $\mathcal{IR}\zeta(\varepsilon) = \zeta(-\varepsilon)^{-1}$ is an antiautomorphism under which flows are invariant. The ‘scaled error’ of an m -stage method, which takes into account the required amount of work, is now $\max\{c_1, c_2\}m^2$. Subsequent methods ψ_n for $n > 3$ all have the same scaled errors as ψ_3 , since symmetrisation affects only higher order terms.

We note from Table 1 that all three alternative methods beat leapfrog, insofar as symmetry and scaled errors are concerned, and that the method ψ_3 is competitive

with ρ and ξ , despite these methods having been optimised. Our next remark arises from a general procedure, rather than a specific case: suppose that φ is a method of *any* order. Then φ^2 of course has the same scaled error as φ . But the leading order error term of $\varphi\mathcal{A}\varphi$ is the sum of the errors of φ and $\mathcal{A}\varphi$. Thus terms a and b corresponding to commutators g and $\mathcal{A}g$ are both replaced by $a + b$. The worst case is when $a = b$ already; otherwise there is some cancellation.

Proposition 6 *The method $\varphi\mathcal{A}\varphi$ has scaled error less than or equal to that of φ , with equality if and only if φ is already $\mathcal{A}$ -symmetric in its leading error term.*

Unfortunately, symmetrisation can affect stability. If φ is strongly stable (if it is a stable method for a well-posed, dissipative partial differential equation, for instance) then there is no problem. But in neutrally-stable problems one may have to halve the time step when moving from ψ_n to ψ_{n+1} . In Table 1 we have also listed the maximum stable time step for each method applied to the harmonic oscillator with phase space (q, p) and $f = \partial/\partial p$, $g = -\partial/\partial q$. The stability limit for method ψ_n tends to π as $n \rightarrow \infty$: the point at which the pair of eigenvalues on the unit circle meet at -1 in the exact solution. There is no way to prevent them bifurcating off the unit circle with loss of stability for the method. By contrast, the maximum stability limit possible with six stages is 6, but this can only be achieved by tailoring the method to have this property; it does not seem to sit easily with the symmetrisation procedure.

6 Other symmetrising compositions

The product (3.3) over a group Γ of automorphisms uses r stages to improve the symmetry error by a single power of ε , $r = |\Gamma|$ being the order of the group. By iterating, one can improve the symmetry error by $\mathcal{O}(\varepsilon^j)$ in r^j stages, *for all* $\delta = \mathcal{O}(\varepsilon)$ *and all* Γ . In fact, using more complicated products can give similar improvements in considerably fewer stages. However, the precise product then depends on the relative size of the symmetry error δ to the time step ε , and on the Γ . The case $\Gamma \cong \mathbb{Z}_2$ is naturally the easiest and we treat it first.

Let $\delta = \varepsilon^p$. This would be the case if $\varphi(f)$ were a method of order $p - 1$ and f had symmetry group Γ . We cannot group all symmetric and all nonsymmetric terms together, as we did in Proposition 1 by writing $\varphi = e^{\varepsilon S + \delta N}$, because when the time step is changed, S and N will change too. We must take

$$\varphi = e^{\varepsilon S + \delta(N_0 + \varepsilon N_1 + \varepsilon^2 N_2 + \dots)}.$$

Then $\mathcal{A}\varphi$ is the same but with M_i replacing N_i . By changing the time step we can also evaluate

$$\varphi(\alpha) = e^{\alpha \varepsilon S + \alpha^p \delta(N_0 + \alpha \varepsilon N_1 + \alpha^2 \varepsilon^2 N_2 + \dots)}.$$

Composing these maps leads to very complicated product formula. To make some progress we consider the simpler case of $N_i = M_i = 0$ for $i \geq 1$, e.g. when $\varphi = e^{\varepsilon f}$. The resulting methods also apply in other cases:

- When terms corresponding to odd powers of ε are symmetric, as they are when, e.g., $\mathcal{A}\varphi = \mathcal{I}\mathcal{R}\varphi$, such as when $\varphi = e^{\varepsilon f} e^{\varepsilon g}$;

- When $\varphi = \mathcal{IR}\varphi$, i.e., φ is time-symmetric—then only odd powers of ε appear in the above expansion and all new commutators are higher order.

Case 1. $\Gamma \cong \mathbb{Z}_2$, **gain** $\mathcal{O}(\varepsilon)$: This is the main case treated above, in (3.3). With two stages, $\varphi\mathcal{A}\varphi$, the symmetry error is multiplied by $\mathcal{O}(\varepsilon)$.

Case 2. $\Gamma \cong \mathbb{Z}_2$, **gain** $\mathcal{O}(\varepsilon^2)$: Let $\delta = \varepsilon^p$. This would be the case if $\varphi(f)$ were a method of order $p - 1$ and f had symmetry group Γ . As in Proposition 1, we have $\varphi = e^{\varepsilon S + \delta N}$ and $\mathcal{A}\varphi = e^{\varepsilon S + \delta M}$. By changing the time step we can also evaluate $\varphi(\alpha) = e^{\alpha\varepsilon S + \alpha^p\delta N}$. The symmetry error terms to be eliminated are those of order δ (due to N and M), $\varepsilon\delta$ (due to $[S, N]$ and $[S, M]$) and, if $p = 1$, of order δ^2 (due to $[N, M]$). However, if ψ is our method, the coefficients of N and M appearing in $\psi(\mathcal{A}\psi)^{-1}$ are always equal, and by taking a symmetric product, all single commutators are eliminated. Therefore we try

$$\psi = \varphi(\alpha)\mathcal{A}\varphi(\beta)\varphi(\alpha) = e^{(2\alpha+\beta)\varepsilon S + 2\alpha^p\delta N + \beta^p\delta M + \mathcal{O}(\varepsilon^3)} \quad (6.1)$$

using the *symmetric Baker–Campbell–Hausdorff formula* (Yoshida 1990)

$$e^x e^y e^x = e^{2x+y+([y,[y,x]]-[x,[x,y]])/6+\dots}. \quad (6.2)$$

We now compare ψ to $\mathcal{A}\psi$, which is just (6.1) with N and M swapped. Thus we have symmetry error $\mathcal{O}(\varepsilon^3)$ and time step ε if

$$2\alpha + \beta = 1, \quad 2\alpha^p = \beta^p \quad (6.3)$$

or

$$\alpha = 1/(2 + 2^{1/p}), \quad \beta = 1 - 2\alpha. \quad (6.4)$$

For $p = 1, 2, 3, 4$ and ∞ we have $\alpha = 0.25, 0.2929, 0.3068, 0.3136$ and $\frac{1}{3}$ respectively—note that the stage lengths become equal as $p \rightarrow \infty$. This family (6.1), (6.4) of product formulae with is a beautiful counterpart to the Suzuki–Yoshida family (Yoshida 1990, Suzuki 1990) developed for the case when φ has order $p - 1$, ψ has order $p + 1$, and $\mathcal{A} = \mathbb{1}$. That family has $\alpha = 1/(2 - 2^{1/p}) = 1.7071, 1.3512, 1.2334$ for $p = 2, 3, 4$; consequently $\beta < -1$ and the large backward steps make that family inaccurate (and, for some problems, unstable). Another difference is that (6.1) starts at $p = 1$ instead of $p = 2$.

If the product (6.1) is iterated ($\psi_2 = \psi(\alpha)\mathcal{A}\psi(\beta)\psi(\alpha)$, etc.) one gains order $\mathcal{O}(\varepsilon^{2j})$ in 3^j stages, roughly 1.732^j stages per power, compared to 2^j stages in case 1.

Case 3. $\Gamma \cong \mathbb{Z}_2$, **gain** $\mathcal{O}(\varepsilon^3)$: We first try four factors, in the form

$$\psi = \varphi(\alpha)\mathcal{A}\varphi(\beta)\varphi(\beta)\mathcal{A}\varphi(\alpha) \quad (6.5)$$

Now the N and M errors are zero by symmetry, and the $[S, N]$ and $[S, M]$ errors are equal by symmetry, as are the $[S, [S, N]]$ and $[S, [S, M]]$ errors. That seems to leave three equations for α and β . However, on expanding the above product with the BCH formula, we find that all of these are equivalent to the single equation

$$\alpha^{p+1} + 2\alpha^p\beta - \beta^{p+1} = 0. \quad (6.6)$$

Together with $2\alpha + 2\beta = 1$, we have $\alpha^p(1 - \alpha) = (\frac{1}{2} - \alpha)^{p+1}$. For $p = 1, 2, 3, 4$ the real solutions near $\frac{1}{4}$ are $(2 - \sqrt{2})/4 \approx 0.1464$, $(6 - 2\sqrt{5})/8 \approx 0.1910$, 0.2087 and 0.2183 . The $\mathcal{O}(\delta^2)$ term due to $[N, M]$ is $\alpha^{2p} + 2\alpha^p\beta^p - \beta^{2p}$, and is nonzero on the above solutions for $p > 1$. However, for $p = 1$ (precisely the case when this term is dominant and is required to vanish) this coefficient matches (6.6) exactly and hence is zero. Thus, (6.5) improves symmetry errors by $\mathcal{O}(\varepsilon^3)$ for all δ .

When applied to a general φ , these methods provide a gain of $\mathcal{O}(\varepsilon^2)$ instead of $\mathcal{O}(\varepsilon^3)$. (There are no N_1 or N_2 errors, by symmetry, but there are new $[S, N_1]$ ones.) However, in numerical tests, this four-stage method did somewhat better than the four-stage method of Section 3, $\varphi(\frac{1}{4})\mathcal{A}\varphi(\frac{1}{4})\mathcal{A}\varphi(\frac{1}{4})\varphi(\frac{1}{4})$, so it deserves to be kept in mind.

Case 4. $\Gamma \cong \mathbb{Z}_2$, **gain** $\mathcal{O}(\varepsilon^4)$: We take five factors, in the form

$$\psi = \varphi(\alpha)\mathcal{A}\varphi(\beta)\varphi(\gamma)\mathcal{A}\varphi(\beta)\varphi(\alpha)$$

The calculation is very similar to Case 2, but longer. From the symmetry, there is one equation at order δ , namely $2\alpha^p + \gamma^p = \beta^p$ (as in (6.3)); no equations at order $\varepsilon^i\delta^j$ when $i + j$ is even; and a single equation at order $\varepsilon^2\delta$, which turns out to be

$$\begin{aligned} & [-\alpha^p(8, 12, 3, 18, 9, 12) + \beta^p(6, 8, 3, 12, 6, 9) - \gamma^p(3, 3, 1, 6, 3, 3)] \\ & \times (\alpha^2, \beta^2, \gamma^2, \alpha\beta, \alpha\gamma, \beta\gamma) = 0 \end{aligned}$$

As before, there is a solution with nearly equal steps: for $p = 1$ it is $(\alpha, \beta) = ((3 - \sqrt{5})/8, 1/4)$. For $p = 2$ the solution is $(0.1339745962155614, 0.2320508075688791)$; for $p = 3$ it is $(0.1516218764988578, 0.2238296732545284)$. For $p = 1$ it might appear that all triple commutators of S , N , and M are dominant (there are eight of these); but as in Case 3, these are all zero on the above solution. This is because when $p = 1$, we are composing $e^{\varepsilon(S+N)}$ and $e^{\varepsilon(S+M)}$, so there are only two independent quantities, namely $S + N$ and $S + M$, to begin with.

When iterated, this family improves the symmetry error by $\mathcal{O}(\varepsilon^{4j})$ in 5^j stages, or roughly 1.495^j stages per order.

We note that the optimised second order methods given in the previous section (see Table 1) for approximating $e^{\varepsilon(f+g)}$ are examples of these symmetrisation procedures. We start with $\varphi(f, g) = e^{\varepsilon f}$, which does not preserve the symmetry $\mathcal{A}\varphi(f, g) = \varphi(g, f)$ at all, so it has $p = 1$. The 4-stage method in Case 3 raises this to $p = 4$, and the 5-stage method in case 4 raises it to $p = 5$. Alternatively, starting with $\varphi(f, g) = e^{\varepsilon f}e^{\varepsilon g}$ (with $p = 2$) or $\varphi(f, g) = e^{\frac{1}{2}\varepsilon f}e^{\varepsilon g}e^{\frac{1}{2}\varepsilon f}$ (with $p = 3$) gives other methods analogous to those in Table 1.

Case 5. Γ arbitrary, **gain** $\mathcal{O}(\varepsilon^2)$: For a general group Γ , we need to find a product ψ such that $\mathcal{A}\psi$ is close to ψ for all $\mathcal{A} \in \Gamma$. Applying $\mathcal{A}$ has the effect of permuting the product according to the group law, which has a complicated effect on the product according to the Baker–Campbell–Hausdorff theorem. Also, there may be new determining equations for each $\mathcal{A} \in \Gamma$. This approach seems hopeless. The best we can do in general is to apply the symmetry used in Cases 2 and 4 above.

Clearly the product must run over the whole group, or applying some $\mathcal{A} \in \Gamma$ will

introduce new terms. So let

$$\psi = \prod_{\mathcal{A} \in \Gamma} \mathcal{A}\varphi \prod'_{\mathcal{A} \in \Gamma} \mathcal{A}\varphi \quad (6.7)$$

where in $\prod'$ the product is taken in reverse order to that in $\prod$. From the symmetric BCH formula (6.2), ψ has no terms of order $\varepsilon\delta$ or δ^2 . The same is true for $\mathcal{A}\varphi$ for any $\mathcal{A} \in \Gamma$. Thus it only remains to check the terms of order δ . If the time steps are all equal as in (6.7), then $\mathcal{A}\psi$ is a permutation of the product in (6.7), hence all terms of order δ match.

Thus (6.7) has a symmetry error of $\mathcal{O}(\varepsilon^2\delta, \varepsilon\delta^2)$, and contains only $2n$ factors. Using (3.3) twice would use n^2 factors.

Note that for $n = 2$, this method is indistinguishable from two iterations of the standard symmetrisation $\psi_2 = \psi\mathcal{A}\psi = \varphi\mathcal{A}\varphi\mathcal{A}\varphi$.

As a closing example, we mention finite difference methods for PDE's. Strang's seminal paper on splitting (Strang 1964) decomposes a two-dimensional spatial operator into two one-dimensional ones, say f and g . To retain second-order accuracy in time he recommended the composition $e^{\frac{1}{2}\varepsilon f}e^{\varepsilon g}e^{\frac{1}{2}\varepsilon f}$. Homogeneous PDEs on square grids will have $(f, g) \mapsto (g, f)$, as a symmetry and this can be (approximately) retained to any desired order using the methods of this section. Similarly, operators on triangular lattices lead to splittings with symmetries of order three. Finally, one can imagine getting simpler finite difference stencils by deliberately breaking a required spatial symmetry, such as translation invariance, only to recover it 'on the fly' during the time integration. We leave a detailed investigation of these ideas for the future.

Acknowledgements. R.M. is grateful to the Isaac Newton Institute for the Mathematical Sciences, Cambridge, UK, for their hospitality during 1996, and to Lottery Science New Zealand and the Marsden Fund of the RSNZ for their financial support.

References

- Calvo, M. P., Iserles, A. & Zanna, A. (1996), 'Runge–Kutta methods for orthogonal and isospectral flows', *Appl. Num. Maths* **22**, 153–163.
- Calvo, M. P., Iserles, A. & Zanna, A. (1997*a*), 'Numerical solution of isospectral flows', *Maths Comput.* to appear.
- Calvo, M. P., Iserles, A. & Zanna, A. (1997*b*), Semi-explicit methods for isospectral flows, Technical Report 1997/NA16, University of Cambridge.
- Chu, M. T. (1994), A list of matrix flows with applications, in A. Bloch, ed., 'Hamiltonian and Gradient Flows', Fields Institute Communications, Amer. Maths Soc., Providence, RI, pp. 87–97.
- Dieci, L., Russell, R. D. & van Vleck, E. S. (1994), 'Unitary integrators and applications to continuous orthonormalization techniques', *SIAM J. Numer. Anal.* **31**, 261–281.

- Hairer, E. & Wanner, G. (1991), *Solving Ordinary Differential Equations II: Stiff Problems and Differential-Algebraic Equations*, Springer-Verlag, Berlin.
- Iserles, A. (1997), Beyond the classical theory of computational ordinary differential equations, in I. Duff & G. Watson, eds, ‘State of the Art in Numerical Analysis’, Oxford University Press, Oxford, pp. 171–192.
- Lothaire, M. (1983), *Combinatorics on Words*, Addison-Wesley, Reading, MA.
- McLachlan, R. I. (1995), ‘On the numerical integration of ordinary differential equations by symmetric composition methods’, *SIAM J. Sci. Comp.* **16**, 151–168.
- McLachlan, R. I., Quispel, G. R. W. & Turner, G. S. (1997), ‘Numerical integrators that preserve symmetries and reversing symmetries’, *SIAM J. Numer. Anal.* to appear.
- Olver, P. J. (1995), *Equivalence, Invariants, and Symmetry*, Cambridge University Press, Cambridge.
- Reich, S. (1997), ‘Dynamical systems, numerical integration, and exponentially small estimates’, Habilitationsschrift, Free University of Berlin.
- Strang, G. (1964), ‘Accurate partial difference methods ii: nonlinear problems’, *Numer. Math.* **6**, 37.
- Stuart, A. M. & Humphries, A. R. (1996), *Dynamical Systems and Numerical Analysis*, Cambridge Monographs on Applied and Computational Mathematics, Cambridge University Press, Cambridge.
- Suzuki, M. (1990), ‘Fractal decomposition of exponential operators with applications to many-body theories and Monte-Carlo simulations’, *Phys. Letters A* **146**, 319–323.
- Toda, M. (1981), *Theory of Nonlinear Lattices*, Springer-Verlag, Berlin.
- Yoshida, H. (1990), ‘Construction of higher order symplectic integrators’, *Phys. Letters A* **150**, 262–268.

A Alternating methods in the current formalism

The alternating methods that inspired this paper appear to use different adjoints from our $\mathcal{N}_h\varphi$. Compare, for example, (1.2) with (4.5). Here we examine this difference. Suppose the target diffeomorphism h factors into $h = h_1h_2$, as for example $y^{-T} = (y^{-1})^T$ does. Suppose the basic method φ is h_2 -covariant but not h_1 -covariant. For example, Runge–Kutta methods are covariant with respect to linear functions such as $y \mapsto y^T$. Then

$$\mathcal{K}_h\varphi = \mathcal{K}_{h_1}\mathcal{K}_{h_2}\varphi = \mathcal{K}_{h_1}\varphi$$

This is the method (1.2), which changes variables by ‘half’ of h .




Figure 5: The discrepancy for (4.7) with a step-size $\frac{1}{100}$. The meaning of different figures is identical to that in Figure 4.

$\mathcal{K}_h\varphi$ is not the recommended adjoint method because it does not always leave h invariant. But we do have

$$\mathcal{N}_h\varphi = \mathcal{N}_{h_1}\mathcal{N}_{h_2}\varphi = \mathcal{N}_{h_1}\mathcal{V}_{h_2}^{-1}\mathcal{K}_{h_2}\varphi = \mathcal{N}_{h_1}\mathcal{V}_{h_2}^{-1}\varphi.$$

(Recall that $\mathcal{V}_h$ changes variables just in the differential equation, not in the initial conditions.) Applied to (1.1), for example, this expresses the adjoint method as

$$\dot{x} = -xa(t, x^{-1}), \quad x(0) = y_0^T.$$

B Antiautomorphisms and the Thue–Morse technique

It is often more natural to express the discrepancy of an element in terms of an antiautomorphism, rather than an automorphism. For example, it makes more sense to let $d(\varphi) = \log[\varphi\varphi^T]$ in $O(n)$ directly, by letting $\mathcal{A}\varphi = \varphi^T$, rather than in the roundabout manner of Section 2. As a matter of fact, this is precisely the approach of (Calvo et al. 1997b) and we generalise it in this section to general Lie groups G .

Given an involutory antiautomorphism $\mathcal{A}$, we redefine the discrepancy as $d(\varphi) =$

$\log \varphi \mathcal{A} \varphi$ and let $\psi = \varphi(\mathcal{A} \varphi)^{-1}$. Supposing that φ is near identity, we again let

$$\varphi = e^{\varepsilon S + \delta N},$$

where $\mathcal{A} e^{\varepsilon S} = e^{\varepsilon S}$ and $\delta = \mathcal{O}(\varepsilon)$, and observe that

$$\mathcal{A} \varphi = e^{\varepsilon S + \delta M}$$

for some M . Therefore, by the BCH formula,

$$e^{d(\varphi)} = e^{\varepsilon S + \delta N} e^{\varepsilon S + \delta M} = e^{2\varepsilon S + \delta(N+M) + \mathcal{O}(\varepsilon^2)} = I + \mathcal{O}(\varepsilon).$$

The automorphism $\mathcal{A}$ being involutory, it is true that $(\mathcal{A} \varphi)^{-1} = \mathcal{A} \varphi^{-1}$, therefore

$$\begin{aligned} e^{d(\psi)} &= \varphi(\mathcal{A} \varphi)^{-1} \mathcal{A}[\varphi(\mathcal{A} \varphi)^{-1}] = \varphi(\mathcal{A} \varphi)^{-1} \varphi^{-1}(\mathcal{A} \varphi) \\ &= e^{\varepsilon S + \delta N} e^{-\varepsilon S - \delta M} e^{-\varepsilon S - \delta N} e^{\varepsilon S + \delta M} \\ &= e^{\delta(N-M) - \frac{1}{2}[\varepsilon S + \delta N, \varepsilon S + \delta M] + \mathcal{O}(\varepsilon^3)} e^{-\delta(N-M) - \frac{1}{2}[\varepsilon S + \delta N, \varepsilon S + \delta M] + \mathcal{O}(\varepsilon^3)} \\ &= e^{\delta(N-M) + \frac{1}{2}\varepsilon\delta[S, N-M] + \mathcal{O}(\varepsilon^3)} e^{-\delta(N-M) + \frac{1}{2}\varepsilon\delta[S, N-M] + \mathcal{O}(\varepsilon^3)} \\ &= e^{\varepsilon\delta[S, N-M] + \mathcal{O}(\varepsilon^3)} = I + \mathcal{O}(\varepsilon d(\varphi)). \end{aligned}$$

We deduce that the statement of Proposition 1 remains valid in the present setting.

The extension to the Thue–Morse formalism, *à la* Section 2, is straightforward. Thus, let again $\psi_0 = \varphi$ and $\psi_{m+1} = \psi_m(\mathcal{A} \psi_m)^{-1}$, $m \in \mathbb{Z}^+$. It follows at once from our analysis that

$$d(\psi_m) = \mathcal{O}(\varepsilon^m d(\varphi)).$$

Moreover, $\mathcal{A}$ being an involution,

$$\begin{aligned} \psi_0 &= \varphi, \\ \psi_1 &= \varphi(\mathcal{A} \varphi)^{-1}, \\ \psi_2 &= \varphi(\mathcal{A} \varphi)^{-1} \{ \mathcal{A}[\varphi(\mathcal{A} \varphi)^{-1}] \}^{-1} = \varphi(\mathcal{A} \varphi)^{-1} [\varphi^{-1}(\mathcal{A} \varphi)]^{-1} = \varphi(\mathcal{A} \varphi)^{-1} (\mathcal{A} \varphi)^{-1} \varphi, \\ \psi_3 &= \varphi(\mathcal{A} \varphi)^{-1} (\mathcal{A} \varphi)^{-1} \varphi \{ \mathcal{A}[\varphi(\mathcal{A} \varphi)^{-1} (\mathcal{A} \varphi)^{-1}] \}^{-1} \\ &= \varphi(\mathcal{A} \varphi)^{-1} (\mathcal{A} \varphi)^{-1} \varphi(\mathcal{A} \varphi)^{-1} \varphi \varphi(\mathcal{A} \varphi)^{-1} \end{aligned}$$

and so on. We again obtain a sequence of two elements—this time φ and $(\mathcal{A} \varphi)^{-1}$ —alternating according to the Thue–Morse prescription.

As an example, we revisit the case of *quadratic Lie groups* from Section 2. Thus, $G = \{y \in \mathrm{GL}(n) : ypy^T = p\}$, where $p \in \mathrm{GL}(n)$ is given. We set

$$\mathcal{A} \varphi = p \varphi^T p^{-1}, \quad \varphi \in \mathrm{GL}(n),$$

therefore $d(\varphi) = \log \varphi p \varphi^T p^{-1}$ vanishes if and only if $\varphi \in G$. As in Section 2, we require that p is either symmetric or skew-symmetric, since this ensures that $\mathcal{A}$ is an involution.

We set $x = \mathcal{A} y$, where y is the solution of (1.1). Therefore

$$\begin{aligned} \dot{x} &= p \dot{y}^T p^{-1} = p y^T a(t, y)^T p^{-1} \\ &= x p a(t, p^T x^T p^{-T})^T p^{-1}. \end{aligned}$$

Like in the first example in Section 4, we can replace pa^T with $-ap$. The outcome is our adjoint equation

$$\dot{x} = -xa(t, p^T x^T p^{-T}), \quad t \geq 0, \quad x(0) = py_0^T p^{-1}. \quad (\text{B.1})$$

Note that (1.2) is a special case of (B.1) when $p = I$ and (1.1) is a flow in $O(n)$. Therefore, our construction is a genuine generalisation of the work in (Calvo et al. 1997b).