PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A
MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

Geometric integration using discrete gradients

Robert I. McLachlan, G. R. W. Quispel and Nicolas Robidoux

Phil. Trans. R. Soc. Lond. A 1999 **357**, 1021-1045
doi: 10.1098/rsta.1999.0363

THE ROYAL SOCIETY

Geometric integration using discrete gradients

By Robert I. McLachlan[1], G. R. W. Quispel[2] and Nicolas Robidoux[1]

[1] *Mathematics Department, Massey University, Palmerston North, New Zealand*
[2] *Faculty of Science, LaTrobe University, Bundoora, Melbourne 3083, Australia*

This paper discusses the discrete analogue of the gradient of a function and shows how discrete gradients can be used in the numerical integration of ordinary differential equations (ODEs). Given an ODE and one or more first integrals (i.e. constants of the motion) and/or Lyapunov functions, it is shown that the ODE can be rewritten as a 'linear-gradient system'. Discrete gradients are used to construct discrete approximations to the ODE which preserve the first integrals and Lyapunov functions exactly. The method applies to all Hamiltonian, Poisson and gradient systems, and also to many dissipative systems (those with a known first integral or Lyapunov function).

Keywords: geometric integration; integrals; gradient systems; Lyapunov functions; Hamiltonian systems; discrete gradients

1. Introduction

Consider the Duffing oscillator without forcing (Guckenheimer & Holmes 1983):

$$\dot{x} = y,$$
$$\dot{y} = x - x^3 - ay.$$

This is not a Hamiltonian system, but the 'energy' $V = \frac{1}{2}y^2 - \frac{1}{2}x^2 + \frac{1}{4}x^4$ does play a distinguished role, because $\dot{V} = -ay^2 = 0$ for $a = 0$ (the undamped Hamiltonian case), and $\dot{V} \leqslant 0$ for $a > 0$ (the damped case (Hale & Kocak 1991)). Observe that this system can be written in the form

$$\begin{pmatrix} \dot{x} \\ \dot{y} \end{pmatrix} = L\nabla V, \qquad L = \begin{pmatrix} 0 & 1 \\ -1 & -a \end{pmatrix}.$$

When $a = 0$, L is antisymmetric and energy is preserved; when $a > 0$, L is negative definite and energy decreases. In this paper all systems that preserve or decrease a function V are written in the form $L\nabla V$, and this form is used to construct geometric integrators that likewise preserve or decrease V.

Geometric integrators are numerical methods for differential equations which preserve structural properties like symplectic structure (Sanz-Serna & Calvo 1994), phase-space volume (Kang & Wang 1994), integrals (Gonzalez 1996; Quispel & Capel 1996), symmetries (McLachlan & Quispel 1998; Iserles *et al.* 1999), reversing symmetries (McLachlan *et al.* 1998; McLachlan & Quispel 1998), isospectrality (Iserles & Zanna 1996), Lie group integrators (McLachlan 1995) or orthonormality (Dieci *et al.* 1994). For an elementary introduction, see Quispel & Dyt (1997). Further examples are found in Stuart & Humphries (1996), which considers systems with linear decay, dissipativity or contractivity, as well as gradient systems, and studies the preservation of fixed points, (quasi-)periodic orbits and attractors.

One can broadly distinguish between general-purpose integrators that 'happen' to preserve a property when present, and methods expressly designed to enforce the property. An example of a method which happens to automatically preserve quadratic integrals is the implicit midpoint rule; an example of a method custom built for systems with a Hamiltonian of the form

$$H(q, p) = T(p) + V(q)$$

is the symplectic splitting method

$$q' = q + \tau \nabla T(p), \qquad p' = p - \tau \nabla V(q'). \tag{1.1}$$

One can also distinguish geometric integrators whose applicability depends on dynamically significant features from those integrators whose applicability does not. An example of the latter is (1.1): the subclass of separable Hamiltonians has no dynamical significance and is only considered because it allows the explicit method (1.1). Finally, custom-built geometric integrators themselves fall into two main categories, depending on whether or not they require the equation to be expressed in a special form (which frequently makes explicit the structural property).

Although one expects geometric integrators to need more analytic information about the differential equation than general-purpose integrators, custom-built integrators requiring a special form may still seem very specialized. Nonetheless, several interesting new methods—e.g. volume-preserving integrators (Kang & Wang 1994; McLachlan & Quispel 1998), Lie group- and isospectrality-preserving integrators (Iserles & Zanna 1996)—are of this type and it is likely that the study of such integrators will lead to settling outstanding questions in the field of numerical integration. Such non-conventional methods explore the space of computable solutions and stretch the notion of a 'numerical method'.

A review of several of these special forms, which may be viewed as special coordinates on the class of systems at hand, is given in McLachlan & Quispel (1998). For example, the versatile 'splitting' technique relies on such special coordinates. Sometimes it is possible for one coordinate system to cover the whole class of systems; sometimes many sets of coordinates are needed, so that the representation is not unique. A disadvantage of any approach relying on special coordinates is that if a system has more than one special property, the desired coordinates may conflict; one is then forced to study the intersection of the relevant spaces.

The present article concerns itself with custom-built geometric integration methods based on rewriting the ordinary differential equation (ODE) in 'linear-gradient form', which is now introduced. Although all the relevant systems can be written in this form, the representation is far from unique.

The conservation of the energy H in Hamiltonian and non-canonical (Poisson) systems,

$$\dot{x} = J(x)\nabla H(x), \qquad (J(x) \text{ antisymmetric}), \tag{1.2}$$

is a consequence of the antisymmetry of the matrix J, since

$$\dot{H} = \nabla H \cdot \dot{x} = (\nabla H)^{\mathrm{T}} J \nabla H = 0.$$

Conversely, and more generally, a first integral $V(x)$ of the ODE $\dot{x} = f(x)$ can be treated as an 'energy function' and the ODE rewritten in a form analogous to (1.2), that is, there exists an antisymmetric matrix $A(x)$ such that

$$\dot{x} = f(x) = A(x)\nabla V(x), \qquad (A(x) \text{ antisymmetric}).$$

Elementary linear algebra yields an instance of such a matrix (Quispel & Capel 1996): under the assumption that ∇V is non-vanishing, the matrix

$$\frac{1}{|\nabla V|^2} f(\nabla V)^{\mathrm{T}}$$

maps ∇V to f. Because V is conserved,

$$0 = \dot{V} = \dot{x}^{\mathrm{T}}\nabla V = f^{\mathrm{T}}\nabla V = 0,$$

so that the antisymmetrizing term $\nabla V f^{\mathrm{T}}$ maps ∇V to 0. Consequently,

$$A = \frac{1}{|\nabla V|^2}(f(\nabla V)^{\mathrm{T}} - (\nabla V)f^{\mathrm{T}}) \tag{1.3}$$

is antisymmetric and maps ∇V to f.

Writing the system $\dot{x} = f(x)$ in the *skew-gradient form*,

$$\dot{x} = A(x)\nabla V(x) \qquad (A(x) \text{ antisymmetric}), \tag{1.4}$$

lends itself to the construction of integrators that have V as an integral, and puts arbitrary integrals on a par with energy functions.

The 'skew-gradient' form (1.4) turns out to be powerful and versatile; this paper can be considered as an essay on its many variations. We consider the following.

(i) Systems with a first integral V, i.e. for which $f \cdot \nabla V = 0$ for all x. The motion stays on the level set $\Sigma_c := \{x : V(x) = c\}$. Preserving this property leads to good nonlinear stability, especially if Σ_c is compact.

(ii) Systems with a weak integral V, i.e. for which there exists a value c such that $f \cdot \nabla V = 0$ for all $x \in \Sigma_c$. Now only the particular level set Σ_c is preserved, which may be stable or unstable. If Σ_c has an interior, then preserving V as a weak integral means that interior orbits cannot escape, a form of stability.

(iii) Gradient systems $\dot{x} = f = -\nabla V$. Here V is not conserved but obeys $\dot{V} = -|\nabla V|^2 \leqslant 0$, that is, V decreases and, subject to some further conditions on V, all orbits tend to fixed points.

(iv) Systems with a Lyapunov function V, for example systems with a forward invariant region (Kloeden & Lorenz 1986). The Lyapunov function can be 'weak' (i.e. $\dot{V} = f \cdot \nabla V \leqslant 0$) or 'strong' ($\dot{V} = f \cdot \nabla V < 0$ except at fixed points; also known as a 'strict' Lyapunov function) (Hirsch & Smale 1974). Preserving these properties is important because they are equivalent to the existence of a stable, respectively asymptotically stable, attracting set. The Lyapunov property can also be local ($\dot{V} \leqslant 0$ for trajectories starting in a neighbourhood of some invariant set, for example a fixed point) or global ($\dot{V} \leqslant 0$ for all trajectories), corresponding to local versus global basins of attraction.

(v) Systems with multiple features as above. Suitable combinations may lead to greater stability (Lakshmikantham *et al.* 1991). For example, even if the level set of each (integral or Lyapunov) function is non-compact, their intersection may be compact. (Multiple Lyapunov functions arise frequently in control theory (Lakshmikantham *et al.* 1991).)

The methods presented in this article are not substantially affected if f, but not V, is time dependent.

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

When dealing with an ODE $\dot{x} = f(x)$ which admits a 'monitor' function V, it is the sign of

$$\alpha(x) := f(x) \cdot \nabla V(x)$$

which distinguishes the different cases above. No matter which case applies, the methods presented in the present article require V and, in some cases, the sign (negative, positive or zero) of α to be known. (Compare with the implicit midpoint rule, which preserves quadratic integrals even when everybody is oblivious to their existence!)

There is considerable overlap between the topic of the present article and the excellent work by Stuart & Humphries (1996). There the structural assumptions on f are quite severe (e.g. dissipativity, contractivity) which allows the integration methods to be fairly general (e.g. Runge–Kutta linear multistep methods). In contrast, this article concerns itself with 'looser' structural properties, but has to specifically construct methods which preserve them. For example, one of the properties considered in Stuart & Humphries (1996) is *monotonicity:* $f \cdot x < 0$ for all $x \neq 0$, a consequence of which is that the function $V(x) = |x|^2$ decreases in $D = \mathbb{R} \setminus \{0\}$. We generalize to almost any function $V(x)$, and almost any domain D.

Many variations and refinements of the class of Lyapunov functions have been used, either for proving stability in different cases (see, for example, Rouche *et al.* 1977), or for ensuring that the Lyapunov function can be preserved under discretization, as in Stuart & Humphries (1996). The key point for us is whether the Lyapunov function can be written as a factor of a linear-gradient formulation of the ODE. We do not need to make any further assumptions on V, because V is not altered in any way in the time discretization.

It is worth remembering that closed properties, such as integrals or weak Lyapunov functions, are easily destroyed by discretization, whereas open properties, such as strong Lyapunov functions, can be preserved under time-discretization by any method for a sufficiently small time-step.

For any of the above continuous time systems, we systematically write them in *linear-gradient* form:

$$\dot{x} = L(x)\nabla V(x),$$

where L is a matrix-valued function. (Note that the term 'linear-gradient' has nothing to do with whether or not the ODE is linear.) The corresponding approximating discrete map $x \mapsto x'$ will have the form

$$\frac{x' - x}{\tau} = \tilde{L}(x, x', \tau)\overline{\nabla} V(x, x')$$

where, for consistency, it is required that $\tilde{L}(x, x, 0) = L(x)$ and $\overline{\nabla} V(x, x) = \nabla V(x)$. Here, $\overline{\nabla} V$ is a *discrete gradient,* that is

$$\overline{\nabla} V(x, x') \cdot (x' - x) = V(x') - V(x), \qquad \forall x, x'. \tag{1.5}$$

The properties of discrete gradients are at the core of the present article and are studied in § 3. In one form or another, (1.5) has appeared several times in the literature. Proving that a specific scheme is energy conserving often relies on an analogue of the discrete gradient's defining property (1.5); some early references are Chorin *et al.* (1978), Gotusso (1985), Itoh & Abe (1988) and Kriksin (1993). In the present context, the first time it appears as an axiom is in Gonzalez (1996), where it is called a 'discrete derivative'. In Quispel & Turner (1996), the method was extended to non-Hamiltonian systems with an integral. The method is now further extended to systems with a Lyapunov function.

Earlier, Roe (1981) introduced requirements similar to (1.5) when formulating linearized Riemann problems in the context of numerically solving systems of hyperbolic partial differential equations of conservation type. Perhaps not surprisingly, many solutions and tricks found in the geometric integration literature had a previous life as tools for the numerical integration of conservation laws; we study this connection in § 4. Discrete gradients are the second theme of this paper, and we will explore their properties and their applications in these different fields.

We first introduce some notation and terminology. We define a *positive definite matrix* L as one such that, for all non-zero real vectors x, $x^{\mathrm{T}}Lx > 0$; L is *positive semidefinite* if $x^{\mathrm{T}}Lx \geqslant 0$ for all real vectors x; *negative definiteness* is defined analogously. If, in addition, L is symmetric, we say so explicitly (unlike others who define definite matrices only in the context of symmetry).

We use the simplest notation from multivariate calculus: ∇ for the Euclidean gradient, $|v|$ for the Euclidean norm, $u \cdot v$ for the Euclidean inner product and $u \wedge v$ for the wedge product $(u \wedge v)_{ij} = u_i v_j - u_j v_i$. The argument x of vector fields $f(x)$, $v(x)$, etc., is often suppressed, as much of our work reduces to linear algebra at each point x. Throughout, A indicates an antisymmetric matrix or skew-symmetric tensor, and S a symmetric matrix or tensor. We often, though not always, use the so-called Einstein summation convention that repeated indices are summed over all meaningful values.

2. Gradient formulations of dynamical systems

Although the linear-gradient form is the most general, we develop it through some special cases:

(i) systems with an integral V can be written as $\dot{x} = A(x)\nabla V(x)$ with A an antisymmetric matrix (§ 2 *a*);

(ii) systems with a strong Lyapunov function V can be written as $\dot{x} = S(x)\nabla V(x)$, where S is symmetric negative definite (§ 2 *b*);

(iii) systems with an integral or a weak or strong Lyapunov function can be written $\dot{x} = L(x)\nabla V(x)$, where L is suitably antisymmetric, negative semidefinite or negative definite (§ 2 *c*). (In fact, we can take the symmetric part of L to be a multiple of the identity.)

Although some systems do of course naturally occur in form (ii), the advantage of (iii) is that it merges all cases smoothly.

(*a*) *Skew-gradient form for systems with integrals*

As pointed out earlier, if $f = A\nabla V$, where A is antisymmetric, then V is an integral of f. The converse also holds.

Proposition 2.1. *Let $f \in C^r(\mathbb{R}^n, \mathbb{R}^n)$, $r \geqslant 1$, $n > 1$, be a vector field and $V \in C^{r+1}(\mathbb{R}^n, \mathbb{R})$ be an integral of f so that $f \cdot \nabla V = 0$ for all x. Then there exists an antisymmetric matrix function A, C^r on the domain $\{x : \nabla V \neq 0\}$, such that $f = A\nabla V$. Moreover, A can be chosen so as to be bounded near every non-degenerate critical point, so that A is locally bounded if V is a Morse function, that is, a smooth function all of whose critical points are non-degenerate.*

Proof. Effectively, we want to solve $A\nabla V = f$ for the antisymmetric matrix A: n linear equations in $\frac{1}{2}n(n-1)$ unknowns. A particular solution, C^r on $\{x : \nabla V \neq 0\}$, is

$$A = \frac{1}{|\nabla V|^2} f \wedge \nabla V \tag{2.1}$$

(where $\wedge$ is the wedge product $(u \wedge v)_{ij} = u_i v_j - u_j v_i$), so that (2.1) is a rewriting of (1.3)).

We now study the behaviour of this A near points where $\nabla V = 0$. The Morse lemma (Abraham & Marsden 1978; Abraham *et al.* 1988) states that there is a coordinate chart about any non-degenerate critical point of V in which

$$V(x) = V(0) + \tfrac{1}{2}x^{\mathrm{T}}Bx,$$

where B, the Hessian of V at $x = 0$, is non-degenerate.

Let x be arbitrary. For λ small enough,

$$\lambda x^{\mathrm{T}} B f(\lambda x) = (\nabla V(\lambda x))^{\mathrm{T}} f(\lambda x) \equiv 0.$$

This implies that

$$x^{\mathrm{T}} B f(0) = \lim_{\lambda \to 0} x^{\mathrm{T}} B f(\lambda x) = \lim_{\lambda \to 0} \frac{1}{\lambda} \lambda x^{\mathrm{T}} B f(\lambda x) = 0,$$

so that the non-degeneracy of B implies that $f(0) = 0$.

Because f vanishes at 0, $|f(x)|/|x|$ is locally bounded. Because B is non-degenerate, $|x|/|\nabla V(x)|$ is also locally bounded. Multiplying the two bounded quotients, we see that $|f(x)|/|\nabla V(x)|$ is locally bounded. Since $|A_{ij}| \leqslant 2|f(x)|/|\nabla V(x)|$, the matrix A is bounded in the neighbourhood of any non-degenerate critical point of V. ■

Near a degenerate critical point of V, A may be unbounded: consider, for example, $f = 4(y, -x)^{\mathrm{T}}$ and $V = (x^2+y^2)^2$ in $\mathbb{R}^2$. When $n = 2$, the matrix A satisfying $Av = f$ for a given non-zero f and v is unique; here A is the unbounded

$$1/(x^2+y^2)\begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}.$$

However, the condition that V is a Morse function is not necessary: for example, if $V = x^4 + y^4$, then A must be bounded if f is bounded. For simplicity, in the remainder of this paper it will be assumed that V is a Morse function.

We call the matrix A in equation (2.1) *canonical*. If $n > 2$, there are other solutions. For example, if g is a non-vanishing vector field then we may take $A = (f \wedge g)/(g \cdot \nabla V)$ (where $v = \nabla V$); or, if $B(x)$ is a non-singular symmetric matrix, we can take $g = Bv$ so that $A = (f \wedge Bv)/(v^{\mathrm{T}}Bv)$. Whatever particular solution is chosen, one can add a homogeneous solution.

Proposition 2.2. *Given a non-zero vector v, the general homogeneous solution of $Av = 0$ is $A_{ij} = \boldsymbol{A}_{ijk} v_k$, where $\boldsymbol{A}$ is any completely skew-symmetric 3-tensor.*

Proof. Given any homogeneous solution A, take

$$\boldsymbol{A}_{ijk} = \frac{1}{|v|^2}(A_{ij}v_k + A_{jk}v_i + A_{ki}v_j).$$

■

In general, the homogeneous solutions and totally skew-symmetric 3-tensors are not in one-to-one correspondence: the dimensionality of the space of totally skew-symmetric 3-tensors is $\frac{1}{6}n(n-1)(n-2)$ while that of the homogeneous solution is $\frac{1}{2}(n-1)(n-2)$, as can be seen from the following argument. Take a basis in which $v = e_1$. Then $Av = 0$ is equivalent to $A_{i1} = 0$, $i = 1, n$, which, taking into account the antisymmetry of A, yields the count. So, for $n > 3$, there are more suitable 3-tensors than homogeneous solution 2-tensors.

Other representations of the homogeneous solutions are

$$A = y \wedge z,$$

where y and z are any vectors orthogonal to v (giving exactly the $\frac{1}{2}(n-1)(n-2)$ independent solutions) and, projecting arbitrary vectors y and z onto $v^{\perp}$,

$$A = (v \cdot v)y \wedge z + (y \cdot v)z \wedge v - (z \cdot v)y \wedge v,$$

which, however, does not have the convenient property of being linear in v like $\boldsymbol{A}_{ijk}v_k$. Another representation of the homogeneous solution is given in Quispel & Capel (1996).

So, in general, many A will satisfy $A\nabla V = f$. For given f and V it may, or may not, be possible to find a particularly nice A: constant or smooth or with many zero entries. Even if A is given at the outset, for example when dealing with a Hamiltonian system, it may be that adding a cleverly chosen homogeneous solution, in conjunction with an appropriate discrete gradient, leads to a better integrator. Of course there cannot be a function $A(f, v)$ such that $A(A^*v, v) = A^*$ for all A^*, because we are only given information about the product $A^*v = f$.

One can build some A that have a certain pattern of zeros, but these will be singular in general. For example, as $Av = 0$ represents $n-1$ independent equations, suppose we take $A_{ij} \neq 0$ only for $|j - i| = 1$. Then

$$A_{12} = \frac{f_2}{v_1}, \qquad A_{i,i+1} = \frac{f_{i+1}}{v_i} + A_{i-1,i}v_{i-1},$$

which is generally singular on the $(n-1)$-manifolds $v_i = 0$. Although the product Av is still non-singular, when we use this formulation to build numerical integrators, the product $A\tilde{v}$, where $\tilde{v}$ is only approximately equal to v, will arise, and in *this* product we do not want singularities.

The canonical A is distinguished by being, in some sense, minimal.

Proposition 2.3. *Given f and v non-zero orthogonal vectors in $\mathbb{R}^n$, let $\tilde{A}$ be an antisymmetric matrix solution of $\tilde{A}v = f$, and let*

$$A := \frac{1}{|v|^2} f \wedge v$$

be the canonical solution. Then

(i) *the rank of A is 2; the rank of $\tilde{A}$ is at least 2;*

(ii) *the codimension of the kernel of A is 2; the codimension of the kernel of $\tilde{A}$ is at least 2; and*

(iii) *if the kernel of $\tilde{A}$ has codimension 2 and the range of $\tilde{A}$ is contained in* $\mathrm{span}(f, v)$, *then $\tilde{A} = A$.*

Proof. Obvious. ∎

Now, f is at least as smooth as A and ∇V; the tricky issue has to do with the converse: given smooth f and V, how smooth can one make A? We suspect that there is always an A which is as smooth as f and ∇V. An indication of this is given by the following proposition.

Proposition 2.4. *Let f be an analytic vector field and V be an analytic function such that $f \cdot \nabla V = 0$. Then, in a neighbourhood of a non-degenerate critical point of V, there exists a real analytic antisymmetric matrix function A such that $A\nabla V = f$, which can be extended to a globally C^∞ solution if V is a Morse function.*

Proof. In the neighbourhood of a non-degenerate critical point, we can take coordinates such that $V = \frac{1}{2}x^{\mathrm{T}}Bx$ with B constant and non-singular. Change variables to $\tilde{x} = Bx$, so that $\nabla_x V = \tilde{x}$. Dropping the tildes, we have to solve $Ax = f$ given that $f \cdot x = 0$.

Expand the f_i in Taylor series. Because $\sum_{j=1}^{n} x_j f_j = f \cdot x = 0$, f_i has no term containing only x_i, for otherwise the leading such term would dominate near the origin on the ith coordinate axis.

We construct a real analytic matrix A row by row.

Suppose the first $k-1$ rows have been found. We proceed to complete the kth row. (For $k = 1$, the whole of the first row is constructed by the following procedure.) By antisymmetry, A_{ki} is already known for $k \leqslant i$. We show that these match all terms in f_k containing only $x_1, \ldots, x_k$. This is equivalent to showing that, when $x_{k+1} = \cdots = x_n = 0$, $\sum_{j=1}^{k} A_{kj} x_j = f_k$. By assumption, $\sum_{i=1}^{k} x_i f_i = 0$ and $\sum_{j=1}^{k} A_{ij} x_j = f_i$ for $i = 1, \ldots, k-1$, so

$$0 = \sum_{i,j=1}^{k} x_i A_{ij} x_j = \sum_{i=1}^{k-1} x_i f_i + x_k \sum_{j=1}^{k} A_{kj} x_j = x_k \left(-f_k + \sum_{j=1}^{k} A_{kj} x_j \right).$$

Thus, all remaining terms in f_k contain a factor x_j for some $j > k$. These can be matched by assigning the A_{kj} $(j > k)$ in any way, e.g. by letting $A_{k,k+1} x_{k+1}$ match all terms with a factor x_{k+1}, then letting $A_{k,k+2} x_{k+2}$ match all remaining terms with a factor x_{k+2}, and so on.

The only operations that have been applied to the (absolutely convergent) Taylor series for the f_i are selecting a subsequence and cancelling a factor x_j. Hence, A is given by a convergent Taylor series and is analytic.

Suppose now that V is a Morse function, so that its critical points are non-degenerate. Away from critical points, the canonical A (2.1) is analytic. As the equation $A\nabla V = f$ is linear in A, a C^∞ partition of unity can be used to blend the canonical A and the above into a C^∞ partition. ∎

Example 2.5. Suppose that $f_1 = x_1 x_2 + x_2 x_3$. We may take $A_{12} = x_1 + x_3$ and $A_{1j} = 0$ for $j > 2$ (other choices are possible.) Then, to satisfy $x_1 f_1 + x_2 f_2 \cdots = 0$, f_2 must contain a term $-x_1^2$; no other f_i could cancel the term $x_1^2 x_2$. $-A_{12} x_1$ cancels this term, leaving a function of $x_2, \ldots, x_n$ only, which can be matched using A_{2j}, $j > 2$.

Example 2.6. At a non-degenerate critical point, consider the lowest order non-vanishing terms. Let $V = \frac{1}{2}x^{\mathrm{T}}Bx$ and $f = Dx$ where $f \cdot \nabla V = 0$, that is, BD is antisymmetric. Then the unique constant antisymmetric A satisfying $A\nabla V = f$ is $A = DB^{-1}$, and is the one constructed by the above algorithm. At the next order, the linear terms in A are not unique, for $\boldsymbol{A}_{ijk}x_k$ can be added for any constant skew tensor $\boldsymbol{A}$.

We suspect, but have not proved, that a C^{r-1} A always exists when f is C^r and V is C^{r+1}.

The above proof also shows that, if V is quadratic and f is a polynomial, then there is a polynomial A such that $A\nabla V = f$. However, it does not follow that this is true for all polynomial V, for the change of variables which we used to make V quadratic generally exists only in the neighbourhood of a critical point.

Example 2.7. Let $n = 2$, $V = p(x)^2q(y)$ for some polynomials p, q and

$$A = \begin{pmatrix} 0 & 1/p \\ -1/p & 0 \end{pmatrix}$$

so that $f = (pq', -2p'q)^{\mathrm{T}}$ is a polynomial. For this f, A is unique, but it is not a polynomial. However, if p and q' have no zeros then V has no critical points, and A is analytic.

(*b*) *Gradient form for systems with Lyapunov functions*

The similarity between gradient systems

$$\dot{x} = -\nabla V$$

and skew-gradient systems (systems with an integral) will now be apparent. More generally, one can consider systems of the form

$$\dot{x} = S(x)\nabla V, \tag{2.2}$$

where $S(x)$ is symmetric and negative definite. Then

$$\dot{V} = \nabla V^{\mathrm{T}} S \nabla V \leqslant 0,$$

with equality only where $\nabla V = 0$, that is, at fixed points. A discussion of the dynamics of such systems may be found in Hirsch & Smale (1974).

A partial converse also holds: systems which admit a decreasing function V can be written in the form (2.2); that is, systems with a Lyapunov function are *generalized gradient systems* (with S defining a metric on $\mathbb{R}^n$ and an associated generalized gradient (Abraham *et al.* 1988; Hirsch & Smale 1974)). (Stuart & Humphries (1996) call such systems 'gradient systems' because of the example $\dot{x} = -\nabla V$; we see now that their terminology could not be more appropriate!)

Proposition 2.8. *Let f be a vector field and V a Lyapunov function, that is, $f \cdot \nabla V \leqslant 0$. Then there exists a symmetric negative definite matrix function S with domain $D = \{x : f \cdot \nabla V(x) \neq 0\}$, such that $S\nabla V = f$, which is smooth if V and f are smooth.*

Proof. Write $v := \nabla V$. Consider points where $f \cdot v \neq 0$. Let $(v, y_1, \ldots, y_{n-1})$ be an orthogonal basis for $\mathbb{R}^n$, which can be chosen to be smooth if V is. We form S from a particular and a homogeneous solution to $Sv = f$, namely

$$S = S^{\text{part}} + S^{\text{homog}} = \frac{1}{f \cdot v} f f^{\mathrm{T}} + c_i y_i y_i^{\mathrm{T}}, \qquad c_i < 0, \qquad i = 1, \ldots, n-1.$$

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

Clearly, $Sv = S^{\text{part}}v = f$ and

$$z^{\mathrm{T}}Sz = \frac{1}{f \cdot v}(f \cdot z)^2 + c_i(y_i \cdot z)^2 \leqslant 0,$$

with equality only if $z = 0$. Furthermore, any solution to $S^{\text{homog}}v = 0$ can be written as $S^{\text{homog}} = \sum c_i y_i y_i^{\mathrm{T}}$ ($\{y_i\}$ and v are the eigenvectors of S, which are orthogonal since S is symmetric); but it is complicated to determine the values of the c_i that make $S = S^{\text{part}} + S^{\text{homog}}$ negative definite. ∎

A caveat is that S may blow up if the angle between f and ∇V approaches $\frac{1}{2}\pi$; for example, if $\dot{V} = 0$ away from fixed points. With V a strong Lyapunov function, S may still blow up, although only at fixed points. This can happen even under quite stringent assumptions on V and f and even if S is allowed to be only *semi*definite.

Proposition 2.9. *With notation as in proposition 2.8, let f be C^1 and the Morse function V be a Lyapunov function for f. If the angle between f and ∇V is locally bounded away from $\frac{1}{2}\pi$, then S^{part} is locally bounded on D. Otherwise, there may be no bounded negative semidefinite solution to $S\nabla V = f$, even if V and f are analytic and V is a strong Lyapunov function.*

Proof. Suppose that the angle between f and ∇V is locally bounded away from $\frac{1}{2}\pi$, so that $f \cdot \nabla V$ vanishes only if f or ∇V vanishes. We wish to show that $S^{\text{part}} := 1/(f \cdot \nabla V)ff^{\mathrm{T}}$ is locally bounded.

Because $S^{\text{part}} = 0$ if f, but not ∇V, vanishes it is sufficient to consider critical points of V.

We proceed to show that f vanishes at every critical point of the Lyapunov Morse function V; for future reference, note that this does not rely on the angle hypothesis. Since V is Morse, all of its critical points are non-degenerate. Consequently, there exists a coordinate chart about any critical point in which

$$V(x) = V(0) + \tfrac{1}{2}x^{\mathrm{T}}Bx,$$

where B, the Hessian of V at $x = 0$, is non-degenerate.

Let x be arbitrary. For λ small enough,

$$\lambda x^{\mathrm{T}}Bf(\lambda x) = (\nabla V(\lambda x))^{\mathrm{T}}f(\lambda x) \leqslant 0.$$

This implies that

$$x^{\mathrm{T}}Bf(0) = \lim_{\lambda\to 0+} x^{\mathrm{T}}Bf(\lambda x) = \lim_{\lambda\to 0+} \frac{1}{\lambda}\lambda x^{\mathrm{T}}Bf(\lambda x) \leqslant 0.$$

Replacing λ by $-\lambda$ in the above yields that $x^{\mathrm{T}}Bf(0) \geqslant 0$. Consequently, $x^{\mathrm{T}}Bf(0) = 0$. The non-degeneracy of B now implies that $f(0) = 0$.

The 2-norm of S^{part},

$$|S| = \frac{|f|^2}{|f \cdot \nabla V|} = \frac{|f|}{|\nabla V|\,|\cos\theta|}$$

is now seen to be bounded since $\cos\theta$ is bounded away from 0, and since both $|x|/|\nabla V| = |x|/|Bx|$ and $|f|/|x|$ are bounded, by virtue of B being non-singular and f vanishing smoothly.

Now consider

$$f = \begin{pmatrix} -y \\ x \end{pmatrix} - (x^2+y^2)\begin{pmatrix} x \\ y \end{pmatrix}, \qquad V = \tfrac{1}{2}(x^2+y^2), \qquad \nabla V = \begin{pmatrix} x \\ y \end{pmatrix}.$$

Here f and the Morse function V are analytic and V is a strong Lyapunov function for f. Let

$$S = -\begin{pmatrix} a & b \\ b & c \end{pmatrix}$$

be a negative semidefinite solution of $S\nabla V = f$. We show that S blows up as one approaches the origin from any direction. The rotational symmetry makes the positive x-axis typical; assume $y = 0$ without loss of generality. $S\nabla V = f$ is equivalent to $a = x^2$ and $b = 1$. On the other hand, negative semidefiniteness is equivalent to a and c being non-negative together with $b^2 \leqslant ac$. Consequently, we must have $c \geqslant 1/x^2$, so that S blows up at the origin. ∎

This suggests the following construction, which builds the symmetric matrix out of the component of f parallel to ∇V—fixing the angle between ∇V and this component to 0 or π, depending on whether V is a non-decreasing or non-increasing function—and handles the component of f normal to ∇V with an antisymmetric matrix. This leads to a form which includes both integral-preserving ($\alpha := f \cdot \nabla V = 0$) and Lyapunov-decreasing ($\alpha < 0$) cases, and smoothly merges the two. Our view is that this is the interesting case, for $\alpha < 0$ is an open property and can be preserved under time discretization by any method for sufficiently small time-step, whereas $\alpha \leqslant 0$ is not an open property and is more difficult to preserve.

(*c*) *Linear-gradient form for general systems*

Instead of using an antisymmetric matrix, suitable for $\alpha = 0$, or a symmetric negative definite matrix, suitable for $\alpha < 0$, the key idea is to use them both. That is, we write our system in the form $\dot{x} = f = Lv$, with L suitably definite but not necessarily symmetric. For such systems, $\dot{V} = v^{\mathrm{T}} f = v^{\mathrm{T}} Lv < 0$ at points where L is negative definite, yet boundedness is retained.

Proposition 2.10. *Let f be a C^r, $r \geqslant 1$, vector field and V a C^{r+1} function. Let $v = \nabla V$, $D = \{x : v(x) \neq 0\}$, and $\alpha = f \cdot v$. Then for all $x \in D$ there exists a C^r matrix L, such that $f = Lv$ and L is negative definite, antisymmetric or positive definite when α is negative, zero, or positive, respectively. If in addition V is a Morse function which is a Lyapunov function for f, then L can be chosen so as to be locally bounded.*

Proof. At each point x, decompose f into its components along and perpendicular to v: $f = u + (\alpha/|v|^2)v$ where $u \in v^{\perp}$. From proposition 2.1 there is a C^r antisymmetric matrix A such that $Av = u$. We take any such A and let $L = A + (\alpha/|v|^2)I$. Then, $z^{\mathrm{T}} Lz = |z|^2\alpha/|v|^2$, and therefore its sign is the same as that of α. ∎

Of course, since we know nothing about L apart from its action in the direction v, this solution is far from unique.

This formulation covers classes (i)–(iv) of the systems listed in the introduction, i.e. those with a single function V. If the system has a weak integral, i.e. $\dot{V} = 0$ on

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

$\Sigma_c = \{x : V(x) = c\}$, then $\alpha = 0$ for $x \in \Sigma_c$ and we recover the skew-gradient form there; off Σ_c, we need no information about the sign of α. If the system has a local or global Lyapunov function, this is reflected by the domain in which $\alpha < 0$.

If the system has an invariant region, we must find a function V that has the boundary of that region as a level set. Then a linear-gradient system can be constructed using this V. For example, the proposition allows one to construct all systems which leave the sphere $|x|^2 = 1$ totally invariant as

$$\dot{x} = \{A(x) + (|x|^2 - 1)\beta(x)I\}x,$$

where A is any antisymmetric matrix function and β is any real function. If the system has a non-empty compact uniformly asymptotically stable region, then such a V is guaranteed to exist (Kloeden & Lorenz 1986).

As in § 2 *a*, one may be concerned to find the smoothest possible matrix L. Taking the symmetric part of L to be a multiple of the identity is then not a good idea, for the scalar $(f \cdot v)/|v|^2$ is usually not analytic at $v = 0$; more general representations $f = (A + S)\nabla V$ must be considered. In the case of strong Lyapunov functions, L can be made to be smooth near critical points.

The following is the 'smooth' version of proposition 2.9; we believe that its conclusion holds under considerably weaker hypotheses, since in effect all we need to do is control the signs of the n eigenvalues of the symmetric part of L, and after satisfying $L(x)\nabla V(x) = f(x)$ we are left with $n^2 - n$ degrees of freedom in L.

Proposition 2.11. *Let f be a C^r (resp. analytic) vector field and let V be a C^{r+1} function with a non-degenerate critical point in a neighbourhood of which $f \cdot \nabla V > 0$ (except, of course, at the critical point itself). Assume further that the angle between f and ∇V is locally bounded away from $\frac{1}{2}\pi$, and that Df is non-singular at the critical point. Then there exists a C^{r-1} (resp. analytic) matrix function L, positive definite in a neighbourhood of the critical point, such that $L\nabla V = f$.*

Proof. The proof of proposition 2.9 establishes that f vanishes at the non-degenerate critical point. As in the proof of proposition 2.4, after a change of variables we have to solve $Lx = f$ with the sign of L equal to the sign of $f \cdot x$. A solution is

$$L(x) := \int_0^1 Df(\xi x)\,\mathrm{d}\xi,$$

that is, $L(x)$ is the average of the derivative of f over the straight line segment running from the origin to x. (Note the similarity with the mean-value discrete gradient, discussed later in the paper.)

That L is C^{r-1} (resp. analytic) is a consequence of the Leibniz and chain rules (see, for example, Dieudonné 1978).

To show that $Lx = f$, note that

$$L(x)x = \int_0^1 Df(\xi x)x\,\mathrm{d}\xi = \int_{[0,x]} Df(\chi)\,\mathrm{d}\chi = f(x) - f(0) = f(x).$$

Finally, we need to show that $L(x)$ is positive definite in a neighbourhood of the critical point. Because the angle between f and $\nabla V = x$ is locally bounded away from $\frac{1}{2}\pi$, there exists a constant $b > 0$ such that

$$f(x) \cdot x \geqslant b|f(x)||x|$$

near the origin. Because Df is invertible in a neighbourhood of 0—since $Df(0)$ is non-singular—and smooth, there exists a constant $B > 0$ such that

$$|f(x)| \geqslant B|x|$$

near the origin. Consequently,

$$\frac{1}{|x|}x \cdot L(x)\frac{1}{|x|}x = \frac{1}{|x|^2}x \cdot f(x) \geqslant b\frac{|f(x)|}{|x|} \geqslant bB,$$

which, since x was arbitrary in its direction, implies that

$$y \cdot L(0)y \geqslant bB|y|^2$$

for all y. The smoothness of L now implies that this last inequality extends to an analogous inequality in a neighbourhood of the critical point, so that $L(x)$ is positive definite there. ■

It may not be possible to choose L analytic everywhere.

Example 2.12. Let $f = (x + x^2, y + y^2)^{\mathrm{T}}$, $V = \frac{1}{2}(x^2 + y^2)$. Because the variables decouple, the unique analytic L satisfying $L\nabla V = f$ is

$$L = \begin{pmatrix} 1+x & 0 \\ 0 & 1+y \end{pmatrix}.$$

But the sign of L does not equal the sign of $f \cdot \nabla V = x^2 + x^3 + y^2 + y^3$ everywhere, for example, at $x = 1$, $y = -\frac{3}{2}$.

The worst case is when the set of points where $\alpha(x) = 0$ passes through a critical point and we also demand that L be antisymmetric on that set, as in proposition 2.10. Then there may be no differentiable L.

Example 2.13. Let $f = (x, 0)$ and $V = \frac{1}{2}(x^2 + y^2)$, so that $f \cdot \nabla V = x^2$. If L is differentiable, then

$$L(0) = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix},$$

which is not antisymmetric on the line $x = 0$.

So we are restricted to our original intention, of making L antisymmetric where $\alpha = 0$ away from critical points, for instance, on manifolds corresponding to weak integrals.

(d) *Multilinear-gradient systems*

We omit the development of systems with several integrals and systems with several Lyapunov functions and consider the general (mixed) case directly.

Proposition 2.14. *Let f be a vector field and let $V^1, \ldots, V^p$ be p functions. Let $v^i = \nabla V^i$, $\alpha^j = f \cdot v^j$ and $B_{ij} = v^i \cdot v^j$, so that B is symmetric. At points where the v^i are linearly independent so that B is non-singular, there exists a completely skew-symmetric $(p+1)$-tensor A such that*

$$f_i = A_{ij_1 \cdots j_p} v^1_{j_1} \cdots v^p_{j_p} + B^{-1}_{jk} \alpha^k v^j_i. \tag{2.3}$$

Proof. Assuming the v^i to be linearly independent, the matrix B is positive definite and thus non-singular. Split f into its component in the linear span of the v^j and component (call it u) perpendicular to the v^j, so that $f = u + \beta^j v^j$.

Let us solve the linear equations

$$A_{ij_1\cdots j_p} v^1_{j_1} \cdots v^p_{j_p} = u_i$$

for the skew-symmetric tensor A, or, writing '$\cdot$' for inner product, $A\cdot(v^1, \ldots, v^p) = u$. (See Darling (1994) for the multilinear algebra; we work in a Euclidean basis of $\mathbb{R}^n$ and identify vectors and 1-forms.) Because $f \cdot v^j = 0$ for all j, we have

$$(f \wedge v^1 \wedge \cdots \wedge v^p) \cdot (v^1, \ldots, v^p) = f \det B,$$

so a particular solution is

$$A = \frac{1}{\det B} f \wedge v^1 \wedge \cdots \wedge v^p.$$

(See below for the homogeneous solution.)

It remains to show that $B^{-1}_{jk} \alpha^k v^j_i$ equals $\beta^j v^j_i$, which added to u_i is f_i. This is a consequence of the invertibility of B together with

$$\alpha^k = f \cdot v^k = v^j \cdot v^k \beta^j = B_{jk} \beta^j.$$

■

We illustrate the significance of this formulation in the case of $p = 2$ integrals or Lyapunov functions. First, if $\dot{V}^1 = \dot{V}^2 = 0$, we have

$$f = A_{ijk} v^1_j v^2_k$$

as a representation of *all* systems with two integrals. This form was originally suggested to us by the example of Nambu dynamics (Nambu 1973; Takhtajan 1994), which has a similar form except that the 'Nambu tensor' A must satisfy an additional differential identity, analogous to the Jacobi identity of Hamiltonian dynamics, which further characterizes the motion. The most important such Nambu tensor has $n = p + 1$ and A is the alternating n-tensor defined by $A_{1\cdots n} = 1$. The important feature is that the integrals appear explicitly in the formulation. It was shown by Quispel & Capel (1999) that all systems with integrals have a Nambu-like form.

To illustrate the constructions geometrically, consider the case of two Lyapunov functions. Now $\det B = |v^1|^2|v^2|^2 - |v^1 \cdot v^2|^2 > 0$ by the Schwarz inequality and we have

$$\begin{aligned} f &= u + (\det B)^{-1}((B_{22}\alpha^1 - B_{12}\alpha^2)v^1 + (B_{11}\alpha^2 - B_{12}\alpha^1)v^2) \\ &= u + (\det B)^{-1}(\alpha^1(B_{22}v^1 - B_{12}v^2) + \alpha^2(B_{12}v^1 - B_{11}v^2)), \end{aligned}$$

showing how the component of f in the (v^1, v^2) plane has components of the appropriate sign along directions perpendicular to each v^i. (For example, $B_{22}v^1 - B_{12}v^2$ is perpendicular to v^2.) The signs of the α^j indicate the sign of each $\dot{V}^j$, and as $\alpha^j \to 0$ we smoothly recover the completely skew form appropriate to systems with two integrals.

We now consider the homogeneous solution for the skew part.

Proposition 2.15. *The system*

$$A_{ij_1\cdots j_p} v^1_{j_1} \cdots v^p_{j_p} = 0 \tag{2.4}$$

in the unknown skew $(p+1)$*-tensor* A *with* p *given linearly independent vectors* $v_i \in \mathbb{R}^n$*, has rank* $n-p$*, and has*

$$\binom{n}{p+1} - (n-p)$$

linearly independent solutions. The general solution can be written as

$$A_{ij_1\cdots j_p} = C^l_{ij_1\cdots j_p k} v^l_k, \tag{2.5}$$

where the C^l*,* $l = 1, \ldots, p$*, are arbitrary skew* $(p+2)$*-tensors.*

Proof. First note that, being a skew tensor, A has

$$\binom{n}{p+1}$$

independent components. Temporarily taking a coordinate system in which $v^1 = e^1, \ldots, v^p = e_p$, the equations read $A_{i,1,2,\ldots,p} = 0$ for all i. But these components are zero already for $i \leqslant p$, so only $(n-p)$ of the equations are independent. Unfortunately, working in this basis does not lead to a convenient form for the general solution. With A as in (2.5),

$$A \cdot (v^1, \ldots, v^p) = C^l \cdot (v^l, v^1, \ldots, v^p) = 0$$

from the skew symmetry of each C^l. This form gives

$$p \binom{n}{p+2}$$

solutions: for $n \leqslant p+1$ it gives $A = 0$ only, which is the only solution; for $n = p+2$, C^l is proportional to the alternating tensor, so it gives p independent solutions, which is all there are; for $n > p+2$ we have

$$p \binom{n}{p+2} > \binom{n}{p+1} - (n-p),$$

which is more than enough.

We must show that any A satisfying (2.4) can be written as $A = C^l \cdot v^l$ for appropriate skew $(p+2)$-tensors C^l. To isolate C^1, say, we contract with v^k for $k > 1$:

$$\tilde{A} := A \cdot (v^2, \ldots, v^p) = C^1 \cdot (v^1, \ldots, v^p).$$

Because $\tilde{A} \cdot v^j = 0$ for all j,

$$(\tilde{A} \wedge v^1 \wedge \cdots \wedge v^p) \cdot (v^1, \ldots, v^p) = \tilde{A} \det B,$$

so one solution for C^1 is

$$C^1 = \frac{1}{\det B} \tilde{A} \wedge v^1 \wedge \cdots \wedge v^p.$$

■

Another way of stating the proposition is to say that a basis for the skew $(p+1)$-tensors satisfying $A\cdot(v^1,\dots,v^p)=0$ are the skew $(p+1)$-tensors satisfying $A\cdot v^k=0$ for some k. For, just as in proposition 2.2, these can be written as $A=C\cdot v^k$ for some skew $(p+2)$-tensor C, by taking $C=(A\wedge v^k)/|v^k|^2$.

As in the case of $p=1$ integral, systems may have much simpler A tensors than those given above. For example, we might take the $(n-p)$ independent components to be $A_{i,i+1,\dots,i+p}$, with all other components zero. For a given f, this A will usually be singular, but it does provide a nice way of constructing simple examples of vector fields having a given list of integrals.

3. Discrete gradients and the discretization of linear-gradient systems

(a) *Definitions and characterizations*

Definition 3.1 (Gonzalez 1996). Let V be a differentiable function. Then $\overline{\nabla}V$ is a discrete gradient of V if it is continuous and

$$\begin{cases}\overline{\nabla}V(x,x')\cdot(x'-x)=V(x')-V(x),\\ \overline{\nabla}V(x,x)=\nabla V(x).\end{cases} \tag{3.1}$$

Proposition 3.2. *$\overline{\nabla}V$ is a discrete gradient if it is continuous and*

$$\overline{\nabla}V(x,x')=\frac{V(x')-V(x)}{|x'-x|^2}(x'-x)+w(x,x'),\qquad (x\neq x'), \tag{3.2}$$

where $w(x,x')$ is a vector-valued function such that

$$\begin{cases}w(x,x')\cdot(x'-x)=0,\qquad (x\neq x'),\\ \lim_{x'\to x}\{w(x,x')-\pi_{(x'-x)^\perp}\nabla V(x)\}=0,\end{cases} \tag{3.3}$$

where $\pi_{(x'-x)^\perp}$ is the projection on the component perpendicular to $x'-x$.

Proof. Suppose that $\overline{\nabla}V(x,x')$ is a discrete gradient. Let w be defined by (3.2). To show that w satisfies (3.3), note that $w(x,x')$ is orthogonal to $x'-x$ because $\overline{\nabla}V(x,x')\cdot(x'-x)=V(x')-V(x)$, and that

$$\begin{aligned}w(x,x')-\pi_{(x'-x)^\perp}\nabla V(x)\\ &=\overline{\nabla}V(x,x')-\frac{V(x')-V(x)}{|x'-x|^2}(x'-x)-\nabla V(x)+\frac{\nabla V(x)\cdot(x'-x)}{|x'-x|^2}(x'-x)\\ &=\{\overline{\nabla}V(x,x')-\nabla V(x)\}-\frac{1}{|x'-x|}\{V(x')-V(x)-\nabla V(x)\cdot(x'-x)\}\frac{(x'-x)}{|x'-x|},\end{aligned}$$

which tends to zero as $x'\to x$.

Conversely, take $\overline{\nabla}V$ to be defined by (3.2) and let t be a unit vector. Then

$$\begin{aligned}(\nabla V(x)-\overline{\nabla}V(x,x))\cdot t&=\lim_{s\to 0}\{\nabla V(x)-\overline{\nabla}V(x,x+st)\}\cdot t\\ &=\lim_{s\to 0}\left\{\frac{1}{s}(\nabla V(x)\cdot st-\overline{\nabla}V(x,x+st)\cdot st)\right\}\\ &=\lim_{s\to 0}\left\{\frac{1}{s}(\nabla V(x)\cdot st-(V(x+st)-V(x)))\right\}=0.\end{aligned}$$

Consequently, $\overline{\nabla}V(x,x)=\nabla V(x)$. ∎

Because $(V(x') - V(x))/|x' - x|$, the component of *any* discrete gradient in the direction $x' - x$, by proposition 3.2, is equal to the average of the component of the gradient in that direction by the potential theorem, we chose the symbol $\overline{\nabla}$ to denote a discrete gradient, since it suggests 'average gradient', which, in at least one direction, is just what it is.

In one dimension, the unique discrete gradient is the difference quotient $(V(x') - V(x))/(x' - x)$. In higher dimensions, there are many possible discrete gradients since only the component along $x' - x$ is tightly constrained, being set to the difference quotient $(V(x') - V(x))/|x' - x|$. It is worth noting that this constraint limits the ability of discrete gradients to approximate point values of continuum gradients to second order.

An analogous definition for vector fields allows us to study the analogue of the chain rule.

Definition 3.3. Let f be a differentiable vector field. Then, $\overline{D}f$ is a *discrete derivative* if it is continuous and

$$\begin{cases} \overline{D}f(x, x') \cdot (x' - x) = f(x') - f(x), \\ \overline{D}f(x, x) = Df(x). \end{cases} \tag{3.4}$$

Clearly, $\overline{D}f$ is a discrete derivative if and only if each of its rows is a discrete gradient for the corresponding component, so that a discrete gradient is a one-dimensional discrete derivative.

Proposition 3.4 (chain rule property). *Let f and g be vector fields, and $\overline{D}f$ and $\overline{D}g$ be corresponding discrete Jacobian matrices. Then, $\overline{D}f(g(x), g(x'))\overline{D}g(x, x')$ is a discrete derivative for $f \circ g$.*

Proof. Omitting the arguments of $\overline{D}f$ and $\overline{D}g$,

$$\begin{aligned} (\overline{D}f\overline{D}g) \cdot (x' - x) &= \overline{D}f \cdot (\overline{D}g \cdot (x' - x)) \\ &= \overline{D}f \cdot (g(x') - g(x)) \\ &= f(g(x')) - f(g(x)). \end{aligned}$$

The other necessary property is a consequence of the (continuum) chain rule. ■

The unique discrete derivative of a curve $x(t)$ is the difference quotient $(x(t') - x(t))/(t' - t)$. Thus, the discrete derivative axiom *itself* is a statement of a special case of the chain-rule property. Essentially, it is this property that makes discrete gradient integrators work.

(*b*) *Some discrete gradients*

The *mean value discrete gradient* (Harten *et al.* 1983) is

$$\overline{\nabla}_1 V(x, x') := \int_0^1 \nabla V((1 - \xi)x + \xi x')\, \mathrm{d}\xi, \qquad (x \neq x');$$

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

that is, $\overline{\nabla}_1 V(x,x')$ is the average of the gradient of V on $[x,x']$, the segment joining x and x'. That it is a discrete gradient is a consequence of the potential theorem:

$$\begin{aligned}\overline{\nabla}_1 V(x,x')\cdot(x'-x) &= \int_0^1 \nabla V((1-\xi)x+\xi x')\cdot(x'-x)\,\mathrm{d}\xi\\ &= \int_{[x,x']}\nabla V(\chi)\cdot\mathrm{d}\chi = \int_{[x,x']}\mathrm{d}V\\ &= V(x')-V(x).\end{aligned}$$

Introduced by Gonzalez (1996), the *midpoint discrete gradient* is

$$\overline{\nabla}_2 V(x,x') := \frac{V(x')-V(x)}{|x'-x|^2}(x'-x)+\pi_{(x'-x)^{\perp}}\nabla V(\tfrac{1}{2}(x+x')),\qquad (x\neq x'). \quad (3.5)$$

In Gonzalez (1996), it is given by the following equivalent formula:

$$\begin{aligned}\overline{\nabla}_2 V(x,x') := {}& \nabla V(\tfrac{1}{2}(x'+x))\\ &+\frac{V(x')-V(x)-\nabla V(\tfrac{1}{2}(x'+x))\cdot(x'-x)}{|x'-x|^2}(x'-x),\qquad (x\neq x').\end{aligned}$$

Both the mean value and the midpoint discrete gradient are second-order approximations to the value of the gradient at the midpoint of $[x,x']$, being exact on linearly varying ∇V.

Fitting somewhat between the mean value and the midpoint are discrete gradients which rely on higher-order quadratures to compute approximations to the average of the component of ∇V *normal* to $x'-x$ along the segment $[x,x']$, instead of the second-order $\pi_{(x'-x)^{\perp}}\nabla V(\frac{1}{2}(x+x'))$. However, the mean value and midpoint discrete gradients capture the mean of the gradient along $[x,x']$ and the value of the gradient at the midpoint of this segment as well as a discrete gradient can, given the stiff constraint along $[x,x']$.

The following discrete gradient differs from the above in that it is associated with a piecewise linear path joining x and x', each piece parallel to one of the coordinate axes, rather than along the segment $[x,x']$. Introduced in Itoh & Abe (1988), the *coordinate increment discrete gradient* is defined as follows: choose an ordering of the coordinates x_i; for the sake of exposition assume that this ordering is $x_1, x_2, x_3, \ldots, x_n$. Then,

$$\overline{\nabla}_3 V(x,x') := \begin{pmatrix} \dfrac{V(x_1',x_2,x_3,\ldots,x_n)-V(x_1,x_2,x_3,\ldots,x_n)}{x_1'-x_1}\\ \dfrac{V(x_1',x_2',x_3,\ldots,x_n)-V(x_1',x_2,x_3,\ldots,x_n)}{x_2'-x_2}\\ \vdots\\ \dfrac{V(x_1',\ldots,x_{n-2}',x_{n-1}',x_n)-V(x_1',\ldots,x_{n-2}',x_{n-1},x_n)}{x_{n-1}'-x_{n-1}}\\ \dfrac{V(x_1',\ldots,x_{n-2}',x_{n-1}',x_n')-V(x_1',\ldots,x_{n-2}',x_{n-1}',x_n)}{x_n'-x_n}\end{pmatrix},$$

in which $0/0$ is understood to be $\partial V/\partial x_i(x)$.

$\overline{\nabla}_3 V$ is a discrete gradient because

$$\overline{\nabla}_3 V(x,x')\cdot(0,\ldots,0,x_i'-x_i,0,\ldots,0) \\ = V(x_1',\ldots,x_{i-1}',x_i',x_{i+1},\ldots,x_n) - V(x_1',\ldots,x_{i-1}',x_i,x_{i+1},\ldots,x_n),$$

so that $\overline{\nabla}_3 V(x,x')\cdot(x'-x)$ is a collapsing sum which reduces to $V(x')-V(x)$, and because in the limit $\overline{\nabla}_3 V(x,x')$ is the vector of partial derivatives.

The potential theorem implies that $\overline{\nabla}_3 V(x,x')$ is the vector of the means of the tangential components of the gradient of V along each of the n segments of the path joining x and x' by incrementing the coordinates one at a time, that is

$$(\overline{\nabla}_3 V(x,x'))_i = \int_0^1 \frac{\partial V}{\partial x_i}(x_1',\ldots,x_{i-1}',x_i+\xi(x_i'-x_i),x_{i+1},\ldots,x_n)\,\mathrm{d}\xi.$$

This last formulation has the advantage of not requiring an interpretation of 0/0.

The coordinate-increment discrete gradient is only a first-order approximation of the gradient at the midpoint of the interval $[x,x']$.

Finally, for *some* functions, one can form a discrete gradient with a chain of simple midpoint evaluations. It has long been known that the implicit midpoint rule preserves quadratic integrals (Cooper 1987). One quick way to see this is to check that it is a discrete gradient in this case: if $V(x)=x^{\mathrm{T}}Bx$,

$$\begin{aligned}\nabla V(\tfrac{1}{2}(x+x'))\cdot(x'-x) &= (x'-x)^{\mathrm{T}}B(x+x') \\ &= x'^{\mathrm{T}}Bx' - x^{\mathrm{T}}Bx \\ &= V(x')-V(x).\end{aligned}$$

Now suppose $V(x)$ is not quadratic, but can be made so by a quadratic change of variables $y=h(x)$. Then by the chain rule property (proposition 3.4), $\overline{D}h(x)^{\mathrm{T}}\overline{\nabla}\tilde{V}(y)$ is a discrete gradient of $V(x)$, where $\tilde{V}(y)=V(x)$. From the assumptions that $\tilde{V}(y)$ and $h(x)$ are quadratic, we can use the midpoint rule for the discrete derivatives, giving the *quadratic midpoint discrete gradient* (with $\overline{x}=\frac{1}{2}(x+x')$):

$$\overline{\nabla}_4 V(x,x') = (Dh)(\overline{x})^{\mathrm{T}}\nabla\tilde{V}(\overline{y}). \tag{3.6}$$

This is further generalized by considering a sequence of changes of variables, each of which is either quadratic or has a quadratic inverse—if $h^{-1}(y)$ is quadratic, then $(D(h^{-1})(\overline{y}))^{-T}\nabla\tilde{V}(\overline{y})$ is a discrete gradient of $V(x)$. This construction was first used in solving hyperbolic conservation laws, where it is known as 'Roe averaging' (see § 4). This is an example of a method which depends on non-dynamically significant features of the equations for its applicability.

(*c*) *The integration method*

For the ODE $\dot{x}=L(x)\nabla V(x)$, we consider integration methods of the form

$$\frac{x'-x}{\tau} = \tilde{L}(x,x',\tau)\overline{\nabla}V(x,x'), \tag{3.7}$$

where we now assume that L, $\tilde{L}$, and V are continuously differentiable, and that $\tilde{L}$ is an approximation of L, i.e. $\tilde{L}(x,x,0)=L(x)$.

If L and $\overline{\nabla}V$ are differentiable, then (3.7) is a consistent method. If we also have the symmetry $\tilde{L}(x,x',\tau)=\tilde{L}(x',x,-\tau)$, then the map is time symmetric (McLachlan

& Quispel 1998) and hence second-order accurate. This can be ensured by taking, e.g. $\tilde{L}(x, x', \tau) = L(\frac{1}{2}(x + x'))$ and $\overline{\nabla} V(x, x') = \overline{\nabla} V(x', x)$, as occurs in the mean value and midpoint discrete gradients.

Gonzalez (1996) carefully discusses the well-posedness and convergence properties of one-step implicit schemes like (3.7).

This method keeps V constant if $\tilde{L}$ is skew (Gonzalez 1996; Quispel & Turner 1996), and decreases V if $\tilde{L}$ is negative definite. Moreover, all such maps can be written in the form (3.7) for any choice of discrete gradient.

Definition 3.5. The function V is an integral of the map $\phi : x \mapsto x'$ if $V(\phi(x)) = V(x)$ for all x. It is a weak Lyapunov function if $V(\phi(x)) \leqslant V(x)$ for all x; it is a strong Lyapunov function if also $V(\phi(x)) = V(x)$ if and only if $x = \phi(x)$.

Proposition 3.6. *If $\tilde{L}(x, x', \tau)$ is antisymmetric for all x, x' and τ, the map (3.7) has V as an integral. Conversely, for any map $x \mapsto x'$ with integral V, at pairs (x, x') such that $\overline{\nabla} V(x, x') \neq 0$ there is an antisymmetric matrix $\tilde{L}$ such that (3.7) holds.*

If $\tilde{L}(x, x', \tau)$ is negative definite (resp. semidefinite) for all x, x', the map (3.7) has V as a strong (resp. weak) Lyapunov function. Conversely, for any map $x \mapsto x'$ with strong (resp. weak) Lyapunov function V, at pairs (x, x') such that $\overline{\nabla} V(x, x') \neq 0$ there is a negative definite (resp. semidefinite) matrix $\tilde{L}$ such that (3.7) holds.

Proof. Assuming the map is given by (3.7), we have

$$V(x') - V(x) = \overline{\nabla} V(x, x') \cdot (x' - x) = \tau \overline{\nabla} V(x, x') \cdot \tilde{L}(x, x', \tau) \overline{\nabla} V(x, x'),$$

which is zero, negative or non-positive when L is skew, negative definite or negative semidefinite, respectively.

For the converses, τ is irrelevant. Let the map be $x' = \phi(x)$. In the first case, we are given that $\overline{\nabla} V$ and $x' - x$ are orthogonal when $x' = \phi(x)$. Therefore, we can take, for example, $\tilde{L} = |\overline{\nabla} V|^{-2} \overline{\nabla} V \wedge (x' - x)$ on these points, and any smooth extension elsewhere. The second case is similar, with $\tilde{L}$ as given in proposition 2.8 or 2.10. ∎

When dealing with weak integrals, one can only assume that L is antisymmetric *for x on the proper level set of V*; indeed, without care, (3.7) may very well move one out of a fixed point. This can happen, for example, if $L(x, x')$ is taken to be $L(\frac{1}{2}(x + x'))$ and the invariant set is not convex. The following proposition shows that, in some situations, all is not lost.

Proposition 3.7. *Let $L(x)$ be a matrix function and $\overline{\nabla} V(x, x')$ be a discrete gradient for V, such that $L(x)$ is antisymmetric for all x on some level set of V, say the zero set of V. Then, this level set is invariant under*

$$\frac{x' - x}{\tau} = L(x) \overline{\nabla} V(x, x').$$

Proof. If x is in the zero set of V, then

$$V(x') - V(x) = \overline{\nabla}(x, x') \cdot (x' - x) = \overline{\nabla}(x, x') \cdot \tau L(x) \overline{\nabla} V(x, x') = 0.$$

∎

So, in the general case, one has to be careful to evaluate L only at points where it will have the correct sign.

A popular way to increase the order of simple methods is by composing several steps with carefully chosen time-steps (Sanz-Serna & Calvo 1994; McLachlan 1995). In the integral-preserving case, this is fine, because the integral is still preserved by each step. However, in the case of a Lyapunov function, to obtain order more than two requires steps with negative time steps. These steps will *increase* V and there is no way to ensure that the total method still has V as a Lyapunov function. We leave the development of Lyapunov-preserving methods with order greater than two as an interesting problem for future research.

We now investigate whether discrete gradient methods can preserve symmetries of the original differential equation. Runge–Kutta methods, for example, preserve all linear symmetries automatically (see, for example, McLachlan *et al.* 1998), but geometric methods tend to break this property. Here the choice of the matrix $\tilde{L}$ and the discrete gradient strongly affect symmetry preservation.

Proposition 3.8. *Let $\dot{x} = L\nabla V$ have a symmetry h. We call it* a special symmetry *if it is also a symmetry of L, $\tilde{L}$ and V, i.e. $\mathrm{d}hL\mathrm{d}h^{\mathrm{T}} = L \circ h$ and $V \circ h = V$. The discrete gradient method $x' - x - \tau\tilde{L}(x, x', \tau)\overline{\nabla}V(x, x') = 0$ has h as a symmetry if*

(i) *h is linear, special and $\overline{\nabla}$ is the mean value discrete gradient;*

(ii) *h is orthogonal, special and $\overline{\nabla}$ is the midpoint discrete gradient;*

(iii) *h is linear diagonal, special and $\overline{\nabla}$ is the coordinate increment discrete gradient; or*

(iv) *h is linear, $\tilde{L}$ is constant and $\overline{\nabla}$ is the mean value discrete gradient.*

Proof. The condition $\mathrm{d}hL\mathrm{d}h^{\mathrm{T}} = L \circ h$ is analogous to the definition of h being a Poisson map in the case that L is a Poisson tensor. It means that under $y = h(x)$ the differential equation $\dot{x} = L\nabla_x V$ transforms to $\dot{y} = L\nabla_y V$. It is a necessary assumption because of the way the discrete gradient breaks up the L and V components. If L has this property, then it is easy to also get it for $\tilde{L}$, by taking $\tilde{L} = L(\frac{1}{2}(x + x'))$ for example.

The condition $V \circ h = V$ implies $\mathrm{d}h^{\mathrm{T}}\nabla V \circ h = \nabla V$. We show that if this is also true for $\overline{\nabla}$, then the symmetry is preserved.

The map defined implicitly by $\psi(x, x') = 0$ has h as a symmetry if $\psi = 0 \Rightarrow \psi \circ h = 0$. For linear special symmetries $x \mapsto Hx$,

$$\begin{aligned} 0 &= x' - x - \tau\tilde{L}(x, x', \tau)\overline{\nabla}V(x, x') \\ \Rightarrow 0 &= Hx' - Hx - \tau H\tilde{L}(x, x', \tau)\overline{\nabla}V(x, x') \\ \Rightarrow 0 &= Hx' - Hx - \tau\tilde{L}(Hx, Hx', \tau)H^{-T}\overline{\nabla}V(x, x'), \end{aligned}$$

so $\psi \circ h = 0$ if $H^{\mathrm{T}}\overline{\nabla}V(Hx, Hx') = \overline{\nabla}V(x, x')$. The mean value discrete gradient is a linear function of ∇V, so in case (i) this follows immediately. For the midpoint discrete gradient (setting $\overline{x} = \frac{1}{2}(x + x')$ and $\Delta x = x' - x$),

$$\begin{aligned} \overline{\nabla}V(x, x') &= \nabla V(\overline{x}) - \frac{\nabla V(\overline{x})^{\mathrm{T}}\Delta x + V(x') - V(x)}{\Delta x^{\mathrm{T}}\Delta x}\Delta x \\ &= H^{\mathrm{T}}\nabla V(H\overline{x}) - \frac{\nabla V(H\overline{x})^{\mathrm{T}}HH^{-1}\Delta Hx + V(Hx') - V(Hx)}{\Delta(Hx)^{\mathrm{T}}H^{-T}H^{-1}\Delta Hx}H^{-1}\Delta Hx \\ &= H^{\mathrm{T}}\overline{\nabla}V(Hx, Hx') \end{aligned}$$

if $H^{-1} = H^{\mathrm{T}}$, i.e. if H is orthogonal. Case (iii) is obvious; the coordinate increment discrete gradient can preserve only diagonal symmetries because any coupling of the variables completely destroys its structure. For case (iv) (the only case in which we can preserve non-special symmetries; note that the system is necessarily Poisson in this case) the method reduces to

$$\frac{x'-x}{\tau} = \int_0^1 f(\xi x + (1-\xi)x')\,\mathrm{d}\xi,$$

which is linear. ■

For systems with multiple integrals and Lyapunov functions, a discrete version of the formulation (2.3) works if gradients are replaced by discrete gradients in just the right places.

Proposition 3.9. *Let $\tilde{A}(x, x', \tau)$ be a skew $(p+1)$-tensor, let $\overline{\nabla} V^j$ be discrete gradients of the p functions V^j and let $\tilde{B}_{ij} = \overline{\nabla} V^i \cdot \overline{\nabla} V^j$. Then the map defined by*

$$\frac{x'-x}{\tau} = \tilde{A}\cdot(\overline{\nabla} V^1, \dots, \overline{\nabla} V^p) + \tilde{B}_{ij}^{-1}\alpha^j \overline{\nabla} V^i$$

obeys

$$V^j(x') - V^j(x) = \alpha^j.$$

Proof. Immediate, on using the discrete gradient property (3.1) and contracting with $\overline{\nabla} V^j$. ■

Notice that we have to take $\alpha^j = f\cdot\nabla V^j$ to get the right constancy, increase or decrease of V^j, but still replace ∇V^j by $\overline{\nabla} V^j$ in the matrix B. With the mean value or midpoint discrete gradients, for example, we might use $\alpha^j(\frac{1}{2}(x+x'))$. Then if V^j is an integral or weak or strong Lyapunov function for f, it also is for the map.

4. The Roe method for PDEs

There is an interesting parallel between the use of discrete gradients for ODEs and their occurrence in the Roe method for treating systems of hyperbolic conservation laws in one space dimension. Consider the system of PDEs

$$\dot{u} = \frac{\partial}{\partial x} f(u) = J(u)\frac{\partial u}{\partial x}, \tag{4.1}$$

where $u(x,t) \in \mathbb{R}^n$, $x \in \mathbb{R}$, and J is the Jacobian of f. For any f, this system has n integrals $\int u\,\mathrm{d}x$. It is important to preserve discrete analogues of these in numerical integration—this is necessary to make shocks propagate at the right speed, for example. Such a scheme is called conservative.

Our presentation follows LeVeque (1992) and Harten *et al.* (1983). In Godunov-type schemes, $u(x,0)$ is approximated by a piecewise constant function, $u(x,0) = u_j$ for $x_{j-(1/2)} < x < x_{j+(1/2)}$. At each time step, for each cell, one solves a Riemann problem, namely equation (4.1) with initial data $u = u_{j-1}$ for $x < x_{j-(1/2)}$, $u = u_j$ for $x > x_{j-(1/2)}$. This does lead to a conservative scheme, but because of its expense, approximate Riemann solvers are often used instead. This approximation can break the conservative property.

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

Roe (1981) considered replacing (4.1) by the constant-coefficient equation

$$\dot{u} = \tilde{J}(u_{j-1}, u_j)\frac{\partial u}{\partial x}, \tag{4.2}$$

where

$$\left.\begin{aligned} &\text{(i)} \quad \tilde{J}(u,v)(v-u) = f(v) - f(u), \\ &\text{(ii)} \quad \bar{J}(u,v) \text{ is diagonalizable with real eigenvalues,} \\ &\text{(iii)} \quad \tilde{J}(u,v) \to J(u) \text{ smoothly as } v \to u. \end{aligned}\right\} \tag{4.3}$$

Axiom (ii) ensures that the linearized problem is hyperbolic and solvable. Of interest to us is that axioms (i) and (iii) make $\tilde{J}$ into a discrete derivative. These, in turn, ensure that the Godunov-type scheme constructed from (4.2) is conservative. (Axiom (i) of equation (4.3) also means that the solution of the exact Riemann problem and of (4.2) are equal in the case of a single shock.) Thus, in this presentation of the Roe method, there is a clear link between the discrete gradient axiom and preservation of a conservation law.

Another entirely different presentation is possible, however. Letting $u_j^n \approx u(j\Delta x, n\Delta t)$, the fully discrete version

$$\frac{u_j^{n+1} - u_j^n}{\Delta t} = \frac{f(u_{j+1}^n, u_j^n) - f(u_j^n, u_{j-1}^n)}{\Delta x} \tag{4.4}$$

has $\sum_j u_j^{n+1} = \sum_j u_j^n$ for any choice of the 'numerical flux function' $f(u,v)$, i.e. it preserves discrete analogues of the integrals $\int u \, \mathrm{d}x$. (Linear integrals are preserved by *any* consistent time-integration scheme. The importance of (4.4) is in its assumed form of the spatial discretization; this is easy to find because of the simplicity of the desired integral.) 'Upwind' schemes are obtained by taking

$$f(u,v) = \tfrac{1}{2}(f(u) + f(v) - d(u,v)), \qquad d(u,v) = |J(\tfrac{1}{2}(u+v))|(v-u) + o(v-u).$$

The Roe method is equivalent to taking $d(u,v) = |\tilde{J}(u,v)|$. It is just one of many possible second-order conservative upwind schemes. One could say that the axioms (4.3) are now responsible for making it into a Godunov-type scheme, which is appealing on physical grounds. However, there is then no apparent connection to the desired conservation laws.

Thus, perhaps the popularity of the Roe method is partly due to its lying in the intersection of the two main classes of methods for systems of hyperbolic conservation laws.

(Good numerical methods for these problems should also satisfy an 'entropy condition', essentially that a scalar function $V(u)$ should obey $\dot{V} \leqslant 0$. Entropy is a Lyapunov function, and it would be interesting to study the use of discrete gradient methods to ensure that it does not increase. In this connection, it is intriguing that in Harten *et al.* (1983) it is noted that any conservative scheme can be fixed to not increase entropy by modifying its numerical flux function to $f(u,v) - \beta(u,v)(v-u)$, where β is a form of limiting projection involving f and V—the comparison with the general discrete gradient (3.2) is striking.)

What methods are used to construct Roe linearizations $\tilde{J}(u,v)$? Harten *et al.* (1983) give what we have called the mean value discrete gradient, and show that it also satisfies axiom (ii) in (4.3). They presented it primarily to show that a $\bar{J}$ always exists; the general solution (3.2) and the Gonzalez discrete gradient (3.5) not having arisen in this field. We suggest that they could also be very useful here.

The linearization which is most often used is the quadratic-midpoint discrete derivative (3.6). Many interesting equations, such as the isothermal equations of gas dynamics and the full Euler equations, have flux functions f which can be made quadratic after a single change of variables and whose inverse is itself quadratic. We suspect that this is, in turn, due to their geometric characterization as 'Euler systems' (Lie–Poisson systems with quadratic kinetic energy) (Marsden & Ratiu 1994), which are quadratic.

Even though, presumably, one cannot characterize the f which can be made quadratic, such systems usually have few components so using the Gonzalez method will be just as practical as the midpoint approach. This extends the Roe method to all systems.

The authors thank the Marsden Fund of the Royal Society of New Zealand, and Lottery Science New Zealand, for their financial support. G.R.W.Q. is indebted to the Massey University Mathematics Department for their hospitality during his visit to Massey, at which time this paper was conceived.

The authors also thank J. A. G. Roberts for pointing out the reference by Hale & Kocak (1991), and B. van Brunt for discussions which led to the proof of proposition 2.4.

References

Abraham, R. & Marsden, J. E. 1978 *Foundations of mechanics*, 2nd edn. Reading, MA: Benjamin/Cummings.

Abraham, R., Marsden, J. E. & Ratiu, T. 1988 *Manifolds, tensor analysis and applications*, 2nd edn. New York: Springer.

Chorin, A. J., Hughes, T. J. R., Marsden, J. E. & McCracken, M. 1978 Product formulas and numerical algorithms. *Commun. Pure Appl. Math.* **31**, 205–256.

Cooper, G. J. 1987 Stability of Runge–Kutta methods for trajectory problems. *IMA J. Numer. Analysis* **7**, 1–13.

Darling, R. W. R. 1994 *Differential forms and connections*, Cambridge University Press.

Dieci, L., Russell, R. D. & Van Vleck, E. S. 1994 Unitary integrators and applications to orthonormalization techniques. *SIAM J. Numer. Analysis* **31**, 261–281.

Dieudonné, J. 1978 *Foundations of modern analysis*. New York: Academic.

Gonzalez, O. 1996 Time integration and discrete Hamiltonian systems. *J. Nonlinear Sci.* **6**, 449–467.

Gotusso, L. 1985 On the energy theorem for the Lagrange equations in the discrete case. *Appl. Math. Comput.* **17**, 129–136.

Guckenheimer, H. & Holmes, P. 1983 *Nonlinear oscillations, dynamical systems and bifurcations of vector fields*. New York: Springer.

Hale, J. & Kocak, H. 1991 *Dynamics and bifurcations*. New York: Springer.

Harten, A., Lax, P. D. & van Leer, B. 1983 On upstream differencing and Godunov-type schemes for hyperbolic conservation laws. *SIAM Rev.* **25**, 35–61.

Hirsch, M. & Smale, S. 1974 *Differential equations, dynamical systems and linear algebra*. New York: Academic.

Iserles, A. & Zanna, A. 1996 A scalpel, not a sledgehammer: qualitative approach to numerical mathematics. DAMTP 1996/NA07, University of Cambridge.

Iserles, A., McLachlan, R. I. & Zanna, A. 1999 Approximately preserving symmetries in numerical integration. *Eur. J. Appl. Math.* (In the press.)

Itoh, T. & Abe, K. 1988 Hamiltonian-conserving discrete canonical equations based on variational difference quotients. *J. Comput. Phys.* **77**, 85.

Kang, F. & Wang, D. L. 1994 *Computational mathematics in China* (ed. Z. Shi & C. Yang), vol. 163. Contemporary Mathematics Series. AMS

Kloeden, P. E. & Lorenz, J. 1986 Stable attracting sets in dynamical systems and their one-step discretizations. *SIAM J. Numer. Analysis* **23**, 986–995.

Kriksin, Yu. A. 1993 A conservative difference scheme for a system of Hamilton's equations with external action. *Zh. Vychisl. Mat. Fiz.* **33**, 206–218.

Lakshmikantham, V., Matrosov, V. M. & Sivasundaram, S. 1991 *Vector Lyapunov functions and stability analysis of nonlinear systems*. Dordrecht: Kluwer.

LeVeque, R. J. 1992 *Numerical methods for conservation laws*. Basel: Birkhäuser.

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. 1998 Generating functions for dynamical systems with symmetries, integrals, and differential invariants. *Physica* D **112**, 298–309.

McLachlan, R. I., Quispel, G. R. W. & Turner, G. S. 1998 Numerical integrators that preserve symmetries and reversing symmetries. *SIAM J. Numer. Analysis* **35**, 586–599.

Marsden, J. E. & Ratiu, T. 1994 *Introduction to mechanics and symmetry*. New York: Springer.

Nambu, Y. 1973 Generalized Hamiltonian dynamics. *Phys. Rev.* D **7**, 2405–2412.

Quispel, G. R. W. & Capel, H. W. 1996 Solving ODEs numerically while preserving a first integral. *Phys. Lett.* A **218**, 223–228.

Quispel, G. R. W. & Capel, H. W. 1999 Solving ODEs numerically while preserving all first integrals. *SIAM J. Sci. Comp.* (Submitted.)

Quispel, G. R. W. & Dyt, C. 1997 Solving ODEs numerically while preserving symmetries, Hamiltonian structure, phase space volume, or first integrals. In *Proc. 15th IMACS World Congress* (ed. A. Sydow), vol. 2, pp. 601–607. Berlin: Wissenschaft und Technik.

Quispel, G. R. W. & Turner, G. S. 1996 Discrete gradient methods for solving ODEs numerically while preserving a first integral. *J. Phys.* A **29**, L341–L349.

Roe, P. L. 1981 Approximate Riemann solvers, parameter vectors and difference schemes. *J. Comput. Phys.* **43**, 357–372.

Rouche, N., Habets, P. & LaLoy, M. 1977 *Stability theory by Liapunov's direct method*. New York: Springer.

Sanz-Serna, J. M. & Calvo, M. P. 1994 *Numerical Hamiltonian problems*. Cambridge: Chapman & Hall.

Stuart, A. M. & Humphries, A. R. 1996 *Dynamical systems and numerical analysis*. Cambridge University Press.

Takhtajan, L. 1994 On foundation of the generalized Nambu mechanics. *Commun. Math. Phys.* **160**, 295–315.

MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES
THE ROYAL SOCIETY A
PHILOSOPHICAL TRANSACTIONS OF

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES