SKEW-ADJOINT FINITE DIFFERENCE METHODS ON NONUNIFORM GRIDS

ADRIAN KITSON, ROBERT I. MCLACHLAN, AND NICOLAS ROBIDOUX

ABSTRACT. We study differentiation matrices that are skew-adjoint with respect to a diagonal discrete inner product, and show that these can be used to construct conservative and hence nonlinearly stable semidiscretizations of certain PDEs. Families of such finite differences of any order are constructed on arbitrary, periodic, one-dimensional grids and are studied for their performance with respect to the roughness of the grid for linear and nonlinear wave equations.

1. INTRODUCTION

Robust numerical methods should be stable and accurate on a wide range of equations, grids, and boundary conditions—a set of design goals which are often in conflict. Many researchers agree that stability is a paramount goal and have focused attention on methods for which the stability proof for the numerical method is directly analogous to the stability proof for the continuous problem. The canonical example is the periodic one-way wave equation $u_t = u_x$, $u(x+2\pi,t) = u(x,t)$, with C^1 initial data. The energy $\mathcal{H} = \int_0^{2\pi} \frac{1}{2}u^2\,dx$ obeys

$$\frac{d\mathcal{H}}{dt} = \int_0^{2\pi} uu_t\,dx = \int_0^{2\pi} uu_x\,dx = -\int_0^{2\pi} (\frac{1}{2}u^2)_x\,dx = 0.$$

The equation is stable with respect to the L_2 inner product.

On the grid $\{x_j\}$ with data values $u_j \approx u(x_j)$, $\mathbf{u} = (u_j)$, when is the semidiscretization $\mathbf{u}_t = D\mathbf{u}$ stable? If we introduce a discrete symmetric bilinear form $\langle \mathbf{u}, \mathbf{v}\rangle_S := \mathbf{u}^t S\mathbf{v}$ (where S is a symmetric matrix), then

$$\frac{d}{dt}\frac{1}{2}\langle \mathbf{u}, \mathbf{u}\rangle_S = \mathbf{u}^t(SD + D^tS)\mathbf{u},$$

which is zero for all $\mathbf{u} \in \mathbb{R}^n$ if and only if D is skew-adjoint with respect to $\langle \mathbf{u}, \mathbf{v}\rangle_S$. Further, the set $\{\mathbf{u} : \langle \mathbf{u}, \mathbf{u}\rangle_S = \text{const.}\}$ is bounded if and only if S is positive definite. Keeping the solution on a bounded level set of an energy function limits its growth and keeps it from blowing up; errors may not be small, but at least they are bounded.

Definition 1. *The differentiation matrix D is* skew-adjoint with respect to the symmetric bilinear form $\langle \mathbf{u}, \mathbf{v}\rangle_S := \mathbf{u}^t S\mathbf{v}$ *(or* with respect to S*) if*

$$\langle \mathbf{u}, D\mathbf{v}\rangle_S + \langle D\mathbf{u}, \mathbf{v}\rangle_S = 0 \quad \forall \mathbf{u},\ \mathbf{v} \in \mathbb{R}^n;$$

equivalently, if $D = S^{-1}A$ for some antisymmetric matrix A. If S is diagonal then we say that D is diagonally skew-adjoint.

1991 *Mathematics Subject Classification.* 65M06.

Such differentiation matrices D are said to satisfy a *summation by parts* identity and were introduced for finite difference methods by Kreiss and Scherer [8]. Recent work [2, 3, 12, 15] has focused on stable boundary closures, different boundary conditions, and different grids; in this paper we study the effect of nonlinearities and rough grids on periodic 1D problems.

An analogy with finite element methods should be pointed out here. The Galerkin finite element approximation of u_x with elements $\{f_i\}$ automatically has the form $S^{-1}A\mathbf{u}$, where

$$S_{ij} = \int_0^{2\pi} f_i f_j \, dx \quad \text{and} \quad A_{ij} = \int_0^{2\pi} f_i f_j' \, dx. \tag{1}$$

The matrix A is antisymmetric if the elements are continuous. That is, finite element differentiation matrices easily inherit a skew-adjoint structure directly from the continuous problem, even on arbitrary grids. Can finite difference methods do this too?

(Most of the results of this paper on discretizing first derivatives have analogs for second derivatives, with skew-adjointness replaced by self-adjointness; see Section 3.2.)

Skew-adjointness is easy to achieve on a uniform grid, e.g., by central differences. Skew-adjointness is harder to achieve on more general grids, but it is particularly on rough grids that it is important to maintain stability. We also wish to extend the use of the energy method for linear PDEs [1, 2, 3, 5, 6, 12, 14, 15] to energy functionals other than $\int u^2 \, dx$ and to nonlinear PDEs. In this paper we show that:

(1) skew-adjoint finite differences can be used to control *nonlinear* as well as linear stability;
(2) there exist skew-adjoint differentiation matrices of arbitrary order on arbitrary grids, and stable skew-adjoint ones on most grids; and
(3) they have a reasonable cost overhead compared to standard methods.

However, we shall see that the differentiation matrices cannot be *locally* determined by the grid (Propositions 2 and 3 of Section 3), but must be determined globally by solving some linear equations (Proposition 4). These equations have nonunique solutions, so in our numerical tests in Section 5 we address the question of how to select a good solution. The best methods are compared to the more conventional 'mapping method' (developed in Section 4), in which the nonuniform grid is mapped to a uniform grid before taking finite differences.

The relationship between a finite difference matrix D, its eigenvalues, and its factorizations $D = S^{-1}A$ is somewhat subtle. Not all matrices have such factorizations, while for those that do, the factorization is not unique. The sign of S does not in general determine the signs of the real parts of the eigenvalues of D. However, we do have the following. (We pass to the complex case for convenience.)

Proposition 1. *$D \in \mathbb{C}^{n\times n}$ is diagonalizable and has imaginary eigenvalues if and only if there exists a positive definite hermitian matrix S and a skew hermitian matrix A such that $D = S^{-1}A$.*

Proof. Suppose that D is diagonalizable and its spectrum is purely imaginary, so that $D = X\Lambda X^{-1}$ with Λ diagonal and imaginary. Then (writing $K^{-\dagger} := (K^\dagger)^{-1}$),

$$D = (XX^\dagger)(X^{-\dagger}\Lambda X^{-1}).$$

The first factor is hermitian positive definite, the second, skew-hermitian.

Conversely, suppose that D is skew-adjoint. Because S is positive definite hermitian, it has a positive definite hermitian square root K. Consider

$$\widetilde{D} := K^{-1}AK^{-1}.$$

Note that $D = K^{-1}\widetilde{D}K$. Because diagonalizability and spectrum are invariant under similarity transformations, it suffices to consider $\widetilde{D}$. $\widetilde{D}$ is normal, and consequently diagonalizable, since

$$\begin{aligned}\widetilde{D}\widetilde{D}^\dagger - \widetilde{D}^\dagger\widetilde{D} &= K^{-1}AK^{-1}K^{-\dagger}A^\dagger K^{-\dagger} - K^{-\dagger}A^\dagger K^{-\dagger}K^{-1}AK^{-1} \\ &= -K^{-1}AS^{-1}AK^{-1} + K^{-1}AS^{-1}AK^{-1} \\ &= 0.\end{aligned}$$

$\widetilde{D}$ has purely imaginary spectrum because it is skew-hermitian. □

When they exist, factorizations $D = S^{-1}A$, with S hermitian and A skew-hermitian, are not unique: If T is an invertible real diagonal matrix, and $S^{-1}A$ is such a factorization of $D = X\Lambda X^{-1}$ with Λ diagonal and imaginary, then $\hat{S} := (XTX^\dagger)^{-1}$ is hermitian and $\hat{A} := \hat{S}D = X^{-\dagger}\Lambda X^{-1}$ is skew-hermitian, so that $D = \hat{S}^{-1}\hat{A}$ is also such a factorization. (If, additionally, S and T are positive definite, then so is $\hat{S}$.) This "rescaling of the eigenvectors of D" does not exhaust all possible factorizations if D has eigenspaces of dimension greater than one. (For example, for $D = 0$, SD is antisymmetric for any S.)

Note that (i) $\hat{S}$ is not positive definite for all T; and (ii) for general D, there is not enough freedom in T to make $\hat{S}$ diagonal.

Example 1. *For $S = \left(\begin{smallmatrix} 0 & 1 \\ 1 & 0 \end{smallmatrix}\right)$ and $\Lambda = \left(\begin{smallmatrix} 1+i & 0 \\ 0 & -1+i \end{smallmatrix}\right)$, the product $S\Lambda$ is skew-hermitian, and S is hermitian, even though Λ does not have imaginary eigenvalues. This is possible because S is not positive definite.*

Example 2. *For any real diagonal S and imaginary diagonal Λ, $S\Lambda$ is skew-hermitian: S is not required to be positive definite to make $D = S^{-1}A$ have imaginary eigenvalues.*

Therefore, if we want D's that factor as $D = S^{-1}A$ with S diagonal and A antisymmetric, they will have to be deliberately constructed. We examine several possible ways of doing this in Sections 3 and 4, and test their performance in Section 5.

Another treatment of the same problem is possible, using the support operator method [4], which hinges on discretizing the formally adjoint gradient and divergence operators in such a way that their matrix representations are formally adjoint with respect to inner products induced by usually diagonal, preferably positive definite, matrices. Here, we concern ourselves primarily with discretizing skew-adjoint operators, and for this reason only one "differentiation" matrix and one "inner product" matrix is involved, instead of two. Furthermore, when we discuss the discretization of self-adjoint operators (the second derivative), we do not factor it (as in $\nabla^2 = \nabla \cdot \nabla$), but rather discretize it "all at once," focusing on self-adjointness and ignoring the duality of factoring.

2. Nonlinear stability by conserving a non-quadratic energy

The conserved quantity $\int u^2\,dx$ can be approximated in many sensible ways by the quadratic form $\mathbf{u}^t S\mathbf{u}$; stability in the sense that $\{\mathbf{u} : \mathbf{u}^t S\mathbf{u} = const.\}$ be

bounded only requires that S be positive definite. The situation for more general conserved quantities $\int h(u)\,dx$ is different. If we stick to linear quadrature methods then the only reasonable approximation is $\sum h(u_i)s_i$, where the s_i are quadrature weights. In the quadratic case this reduces to $\mathbf{u}^t S\mathbf{u}$ where S is diagonal. We shall now show that, in one general method of constructing conservative discretizations, the differentiation matrices that appear should be skew/self-adjoint with respect to the same S that appears in the quadrature formula. That is, S should be diagonal.

For convenience, we consider $u \in \mathfrak{F} := C^\infty(\mathbb{R})$. We consider PDEs of the form

$$u_t = \mathcal{D}(u)\frac{\delta \mathcal{H}}{\delta u}, \tag{2}$$

where $\mathcal{D}(u)$ is a linear differential operator (that may depend nonlinearly on u), $H : \mathfrak{F} \to \mathbb{R}$ is a conserved quantity such as energy, $\frac{\delta \mathcal{H}}{\delta u}$ is the variational derivative defined by

$$\int \frac{\delta \mathcal{H}}{\delta u} v\,dx = \frac{d}{d\epsilon}\mathcal{H}(u+\epsilon v)|_{\epsilon=0}$$

for all u, $v \in \mathfrak{F}$, and $\mathcal{D}(u)$ is skew-adjoint with respect to L_2, i.e.

$$\int v\mathcal{D}(u)w\,dx + \int w\mathcal{D}(u)v = 0$$

for all u, v, $w \in \mathfrak{F}$. These conditions ensure that $d\mathcal{H}/dt = 0$. Hamiltonian PDEs are naturally in the form (2) with, e.g., $\mathcal{D} = \left(\begin{smallmatrix} 0 & I \\ -I & 0 \end{smallmatrix}\right)$ or $\mathcal{D} = \partial_x$. For further details see [11, 13].

We seek a semidiscretization of (2) of the analogous discrete form, namely an ODE of "skew-gradient" form [10] with conserved quantity $H(\mathbf{u})$:

$$\mathbf{u}_t = L(\mathbf{u})\nabla H(\mathbf{u}), \tag{3}$$

where $L(\mathbf{u}) = -L^t(\mathbf{u})$. Of course we want an accurate discretization of the PDE, but we do *not* insist that $H(\mathbf{u}) \approx \mathcal{H}(u)$ when $u_i \approx u(x_i)$.

Consider first the case $\mathcal{H}(u) = \int h(u)\,dx$ with discretization $H(\mathbf{u}) = h'(\mathbf{u})^t\mathbf{s}$ (where we have made the natural extension of h from $\mathbb{R}$ to $\mathbb{R}^n$). Then $\frac{\delta \mathcal{H}}{\delta u} = h'(u)$ and $\nabla H = Sh'(\mathbf{u})$, where $\mathbf{s}$ is the vector of quadrature weights and S is the diagonal matrix with diagonal $\mathbf{s}$. Eq. (3) will be a good approximation of the PDE if $L(\mathbf{u})S \approx \mathcal{D}(u)$. Since we also need $L(\mathbf{u}) = -L^t(\mathbf{u})$, the discretizations of any derivatives appearing in $\mathcal{D}(u)$ should be skew/self-adjoint with respect to S.

For example, if $\mathcal{D} = i$ then we can take $L = iS^{-1}$ (which is skew-hermitian) to get the discretization $\mathbf{u}_t = (iS^{-1})\nabla H(\mathbf{u}) = (iS^{-1})(Sh'(\mathbf{u})) = ih'(\mathbf{u})$.

More significantly, if $\mathcal{D} = \partial_x$ then we can take $L = DS^{-1}$ with $D = S^{-1}A$ to get

$$\mathbf{u}_t = Dh'(\mathbf{u}) = S^{-1}AS^{-1}\nabla H(\mathbf{u}) = L\nabla H(\mathbf{u}).$$

This discretization conserves $H(\mathbf{u})$ and is a Hamiltonian system.

These conclusions carry over to more general conserved quantities. Let $h : \mathbb{R}^4 \to \mathbb{R}$ and let $\mathcal{H}(u)$ be the differential function

$$\mathcal{H}(u) = \int h(x, u, u_x, u_{xx})\,dx$$

with variational derivative

$$\frac{\delta \mathcal{H}}{\delta u} = h_2 - \partial_x h_3 + \partial_{xx} h_4.$$

For the discretization

$$H(\mathbf{u}) = \sum_i h(x_i, u_i, (D_1\mathbf{u})_i, (D_2\mathbf{u})_i)s_i$$

we have (using the natural extension of h to $(\mathbb{R}^4)^n$)

$$S^{-1}\nabla H(\mathbf{u}) = h_2 + S^{-1}D_1^t S h_3 + S^{-1}D_2^t S h_4.$$

Thus we require that $S^{-1}D_1^t S h_3 \approx -D_1 h_3$, $S^{-1}D_2^t S h_4 \approx D_2 h_4$, *and* that the discretization of $\mathcal{D}(u)$ is skew-adjoint with respect to S. Although various solutions are possible, the most obvious discretization which satisfies these requirements is

$$D_1 = S^{-1}A_1 \quad (A_1^t = -A_1),$$
$$D_2 = S^{-1}A_2 \quad (A_2^t = A_2),$$

with all first derivatives in $\mathcal{D}$ discretized by D_1 and all second derivatives discretized by D_2.

This conclusion, that all discrete derivatives should be skew-adjoint or self-adjoint with respect to the same, diagonal, bilinear form, only applies if one wants a systematic discretization procedure for all equations of the form (2). In particular, if h is quadratic then nondiagonal bilinear forms are possible; in this and other cases we have shown further [11] that one can use discrete derivatives of the form $S^{-t}A$ where S is not even symmetric. However, for a nonlinear Hamiltonian, such as the sine-Gordon Hamiltonian $\int(\frac{1}{2}p^2 + \frac{1}{2}qq_{xx} - \cos q)\,dx$, it is difficult to see how using different bilinear forms on the different terms could be an advantage.

3. Construction of skew-adjoint finite differences: exact-on-polynomials methods

3.1. **First derivatives.** We explore the construction of differentiation matrices of the form $D = S^{-1}A$, skew-adjoint with respect to the diagonal bilinear form defined by S, i.e., A is antisymmetric. The bandwidths of D and A are then equal. Let $S_{ij} = s_i\delta_{ij}$. Let the grid points be x_i with spacings $h_i = x_{i+1} - x_i$.

The standard finite difference approximation of $d^n u/dx^n$, formed by differentiating the degree d-polynomial interpolant at $d+1$ arbitrary points, has order at least $d+1-n$ when u is C^{d+1} [7]. We will therefore seek differentiation matrices D which are exact on polynomials of various degrees d. Without some extra control over D (e.g. on its norm, or on the grid spacing ratios h_i/h_{i+1}), we cannot expect to attain the same order as the standard method, but this criterion does serve as a useful guide. We shall construct S and A simultaneously so that D is exact on polynomials of various degrees, considering firstly approximations of the first derivative.

If A is tridiagonal, then D can be exact on linears: the unique solution is $A_{i,i+1} = 1$, $s_i = x_{i+1} - x_{i-1}$.

If A has bandwidth 5, then A and S together contain three unknowns per grid point. If the finite difference is to be exact on quadratics, there are three equations to satisfy per grid point. This suggests that such a finite difference might exist. However, we shall show (Propositions 2 and 3) that the requirement of antisymmetry globally couples the entries of A. The finite difference stencil cannot, unfortunately, be a simple function of a few neighbouring grid spacings. Since A and S depend on the entire grid, boundary conditions play a role. We proceed to study the case of periodic boundary conditions, both theoretically (showing the

existence of suitable A and S in Proposition 4) and in practise (Section 5). The entries of A lie in the null space of a certain matrix (called L below), which must be computed in advance to determine the finite difference. The choice of A then determines S explicitly.

Proposition 2. *There is no finite difference approximation of the first derivative which has bandwidth 5, is invariant under the translations $x_i \mapsto x_i + c$, is diagonally skew-adjoint, is exact on quadratics, and whose coefficients at x_i depend only on $x_{i-2}, \dots, x_{i+2}$.*

Proof. We first show that the coefficients cannot depend on all of $x_{i-2}, \dots, x_{i+2}$. Let $A_{i,i+1} = a(x_{i-2}, \dots, x_{i+2})$ and $A_{i,i+2} = b(x_{i-2}, \dots, x_{i+2})$. By antisymmetry, $A_{i+2,i} = -A_{i,i+2} = -b(x_{i-2}, \dots, x_{i+2})$, which is assumed to depend only on $x_i, \dots, x_{i+4}$. Thus $A_{i,i+2} = b(x_i, \dots, x_{i+2})$. Similarly, we find $A_{i,i+1} = a(x_{i-1}, \dots, x_{i+2})$. By invariance under translations, these functions must depend only on the grid spacings, so we have the finite difference

$$\begin{aligned}(S^{-1}A\mathbf{u})_i = s_i^{-1}\big(&- b(h_{i-2}, h_{i-1})u_{i-2} + a(h_{i-2}, h_{i-1}, h_i)u_{i-1} \\ &+ a(h_{i-1}, h_i, h_{i+1})u_{i+1} + b(h_i, h_{i+1})u_{i+2}\big).\end{aligned}$$

We require that this finite difference be exact on quadratics. Take the test functions $u_j = (x_i - x_j)^\nu$, $\nu = 0, 1, 2$. Their derivatives at $x = x_i$ are 0, 1, and 0, respectively. The equation for $\nu = 1$ can be taken to fix the value of s_i, leaving the equations for $\nu = 0$ and $\nu = 2$. Specializing to $i = 2$ for convenience, these are

$$-b(h_0, h_1) - a(h_0, h_1, h_2) + a(h_1, h_2, h_3) + b(h_2, h_3) = 0 \tag{4}$$

and

$$-b(h_0, h_1)(h_0+h_1)^2 - a(h_0, h_1, h_2)h_1^2 + a(h_1, h_2, h_3)h_2^2 + b(h_2, h_3)(h_2+h_3)^2 = 0. \tag{5}$$

Considering first Eq. (4), holding h_0, h_1, and h_2 fixed and varying h_3 shows that $a(h_1, h_2, h_3) + b(h_2, h_3)$ must be independent of h_3; so let

$$a(h_1, h_2, h_3) = g(h_1, h_2) - b(h_2, h_3).$$

Substituting and varying h_2 then shows that $b(h_1, h_2) + g(h_1, h_2)$ is a function of h_1 only, so let

$$g(h_1, h_2) = f(h_1) - b(h_1, h_2).$$

Substituting again leaves $f(h_1) - f(h_0) = 0$ for all h_0 and h_1, i.e., $f(h) \equiv F$. That is, $a(h_1, h_2, h_3) = F - b(h_1, h_2) - b(h_2, h_3)$.

Substituting this value for a into Eq. (5) and varying first h_0 and then h_3, gives the two results

$$b(h_0, h_1) = \frac{d(h_1)}{h_0(h_0 + 2h_1)},$$

$$b(h_2, h_3) = \frac{e(h_2)}{h_3(h_3 + 2h_2)}.$$

But these equations are incompatible, for they imply either $d(h) = e(h) = 0$, in which case Eq. (5) is not satisfied, or

$$\frac{d(h_3)}{e(h_2)} = \frac{h_2(h_2 + 2h_3)}{h_3(h_3 + 2h_2)},$$

yet the right hand side is not separable. □

A similar result is presumably true for any bandwidth and any local dependence of a_{ij} on the h_j. Merely being exact on constants indicates this:

Proposition 3. *A local skew-symmetric finite difference, exact on constants, obeys*

$$\sum_{j\le i, k>i} a_{jk} = C$$

for some C, for all i.

Proof. We take the difference of the two sums for two adjacent values of i, say $i-1$ and i.

$$\begin{aligned}\sum_{j\le i,k>i} a_{jk} - \sum_{j\le i-1,k>i-1} a_{jk} &= \sum_{k>i} a_{ik} - \sum_{j<i} a_{ji} \\ &= \sum_{k>i} a_{ik} + \sum_{k<i} a_{ik} \\ &= 0\end{aligned}$$

since the row sums of A must be zero if the finite difference is exact on constants. □

So either C is a universal constant independent of the h_i—which seems problematic—or C depends on all of the h_i, i.e., the entries of A are globally coupled.

Since locally determined finite differences are impossible, we consider globally determined finite differences. To be specific, we consider periodic grids of n points. Naive counting suggests that with bandwidth $2d+1$ we have $n(d+1)$ unknowns in A and (the diagonal of) S, so it should be possible to choose these to make the finite difference exact on polynomials of degree $\le d$. This is indeed the case.

Proposition 4. *Let S be diagonal, let A be antisymmetric with half bandwidth d, and let the differentiation matrix $D = S^{-1}A$ be exact on functions which are locally polynomials of degree $\le d$ on a given one-dimensional periodic grid of n points. There is a linear space of dimension at least $\lfloor d/2 \rfloor + 1$ of such matrices A, each A having an associated matrix S. Each 1-dimensional subspace of such A's corresponds to the same D, so the manifold of such D's contains the projective space of dimension $\lfloor d/2 \rfloor$.*

Proof. The equations to be solved are

$$\begin{aligned}\sum_j A_{ij}(x-x_j)^\nu|_{x=x_i} &= S_{ii}\nu(x-x_i)^{\nu-1}|_{x=x_i} \\ &= \begin{cases} 0, & \nu \ne 1 \\ S_{ii}, & \nu = 1 \end{cases}\end{aligned} \tag{6}$$

for $i = 1,\dots,n$, $\nu = 0,\dots,d$, and $x_j = \sum_{i=1}^{j} h_i$. The S_{ii} can be determined from the equations for $\nu = 1$, leaving nd homogeneous linear equations in the A_{ij}. We represent these equations as

$$L_\nu \mathbf{w} = 0,\ \nu \ne 1; \quad \mathbf{s} = L_1 \mathbf{w} \tag{7}$$

where $\mathbf{w}$ contains the upper diagonals of A, $\mathbf{s} = \mathrm{diag}(S)$, and $L_\nu \in \mathbb{R}^{n\times nd}$. We denote the combined matrix of coefficients of the determining equations by

$$L := [L_0; L_2; \dots; L_d] \in \mathbb{R}^{nd\times nd}.$$

Now, using antisymmetry of A, one can check that the sum of the n equations corresponding to each even value of ν is zero. Thus, L has corank at least $m =$

$\lfloor d/2 \rfloor + 1$. Each matrix A in the m-dimensional subspace of solutions determines a matrix S and hence (barring a measure zero space of grids for which some $s_i = 0$ and S is not invertible; we shall want to stay well away from $s_i = 0$ in any case) a finite difference matrix D. If, however, $A^{(1)} = rA^{(2)}$, then $S^{(1)} = r^{-1}S^{(2)}$, so $D^{(1)} = D^{(2)}$. Thus the space of such differentiation matrices, exact on polynomials of degree $\leq d$, contains the projective space $\mathbb{RP}^{\lfloor d/2 \rfloor}$ of dimension $\lfloor d/2 \rfloor$. □

Thus the generic situation is to get, for example, an isolated solution which is exact on linears, a one-dimensional space of finite differences exact on polynomials of degrees ≤ 2 and ≤ 3, and a two-dimensional space of finite differences exact on polynomials of degrees ≤ 4 and ≤ 5. From these spaces of solutions it will be necessary to choose a particular solution, a point we consider in Section 5. We consider just one condition here, the positivity of S.

It will be crucial, for most applications, that S be positive definite, i.e., that $S_{ii} > 0$ for all i. This ensures that D has imaginary eigenvalues and that $\{\mathbf{u} : \mathbf{u}^t S \mathbf{u} = 1\}$ is compact. However, we will see from the geometric condition given below in Proposition 5 that not all the solutions will have S positive definite. Moreover, there exist extremely rough grids for which *none* of the solutions are positive definite.

Example 3. *For $n = 5$ (the smallest sensible case), one fairly smooth grid for which this happens is $\mathbf{h} \propto (3, 6, 10, 3, 6)^t$. The solutions for $\mathbf{s}$ which are exact on polynomials of degree ≤ 2 are*

$$\mathbf{s} = \begin{pmatrix} -0.5837 & 0.6416 \\ 0.0864 & -0.6453 \\ -0.5238 & 0.3977 \\ -0.6013 & 0.0262 \\ 0.1259 & -0.1143 \end{pmatrix} \begin{pmatrix} c_1 \\ c_2 \end{pmatrix}.$$

(The coefficient matrix here is computed as $S = L_1 \operatorname{null}(L_0; L_2)$.) One can check that there is no choice of c_1, c_2 which makes all $s_i > 0$ (see Proposition 5). We examine this phenomenon for more realistic grids in Section 5.

The set of finite differences for which S is positive definite can be determined geometrically as follows.

Proposition 5. *Let a basis for the solutions* $\operatorname{diag}(S)$ *be $\mathbf{s}^1, \ldots, \mathbf{s}^m$, $m = \lfloor d/2 \rfloor + 1$. Let $\mathbf{s}^i = (s^i)_j$ and let $\mathbf{p}^1, \ldots, \mathbf{p}^n \in \mathbb{R}^m$ be defined by $(p^j)_i = s^i_j$. There is a positive definite solution S iff the origin lies outside the convex hull of the $\mathbf{p}^j$.*

Proof. The general solution is $\mathbf{s} = \sum c_i \mathbf{s}^i$. Since solutions related by a scale factor are equivalent, we can take $\|\mathbf{c}\| = 1$. Let Q be any orthogonal matrix whose first row is $\mathbf{c}$. Then $s_j = (Q\mathbf{p}^j)_1$. Thus, the problem of determining a positive S becomes that of finding a common rotation which sends the n points $\mathbf{p}^j$ into the halfspace $p_1 > 0$; or equivalently, of determining if all of the $\mathbf{p}^j$ lie to one side of some hyperplane through the origin. This is true if the origin is not contained in the convex hull of the $\mathbf{p}^j$. □

Thus, if it were necessary to completely automate the process, one could actually construct the convex hull (in time $\mathcal{O}(n \log n)$). In practise, for $d = 2$ and $d = 3$, we have $m = 2$ and one can merely search amongst planar rotations for the one which, e.g., maximizes $\min s_j$.

It could be argued that finding these finite differences is difficult because we are asking so much of them, namely, high order on arbitrarily rough grids. Consider the (nonlocal) Galerkin approximation constructed using piecewise linear elements. It is only $\mathcal{O}(h)$ accurate on rough grids, yet globally $\mathcal{O}(h^2)$ accurate on smooth grids, and even $\mathcal{O}(h^4)$ accurate at the grid points on a uniform grid. Perhaps it is natural to seek this kind of order behaviour. We shall see below that both exact-on-polynomials and mapping methods do show such a phenomenon of superconvergence on smooth grids.

3.2. **Second derivatives.** The theory and construction of matrices D_2 representing second derivatives is very similar. The analog of Prop. 1 holds, namely there exists a positive definite hermitian matrix S and a *hermitian* matrix A such that $D_2 = S^{-1}A$ if and only if D_2 is diagonalizable and has *real* eigenvalues. However, thinking of the test equation $u_{tt} = u_{xx}$ with semidiscretization $\mathbf{u}_{tt} = D_2\mathbf{u}$, we should require that D_2 has real *nonpositive* eigenvalues, since this implies $\frac{d}{dt}\langle \mathbf{u}, \mathbf{u}\rangle_S \leq 0$. This cannot be determined by looking at the factors S and A separately, as we did for first derivatives.

Example 4.

$$S^{-1}A = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$$

has eigenvalues ± 1*, despite being self-adjoint with respect to a diagonal, positive definite bilinear form; while*

$$S^{-1}A = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} \begin{pmatrix} -6 & 1 \\ 1 & 1 \end{pmatrix}$$

does have negative real eigenvalues $\lambda \approx -5.8,\ -1.2$*, despite being self-adjoint with respect to an indefinite bilinear form.*

The analogue of Proposition 2 holds for second derivatives. Being exact on constants fixes the diagonal of A; begin exact on quadratics fixes S. By analogy with finite differences based on interpolation, to get a second order method we impose that the finite difference be exact on cubics.

The equations analogous to Eq. (6) are

$$\begin{aligned} \sum_j A_{ij}(x - x_j)^\nu|_{x=x_i} &= S_{ii}\nu(\nu-1)(x-x_i)^{\nu-2}|_{x=x_i} \\ &= \begin{cases} 0, & \nu = 1, 3, 4, \ldots, d+1 \\ 2S_{ii}, & \nu = 2. \end{cases} \end{aligned} \tag{8}$$

Since A is symmetric, the sum of the equations corresponding to *even* ν is zero, so the dimension of the solution space is still $\lfloor d/2 \rfloor$ with bandwidth $2d+1$, as in Prop. 4. Note that the S matrices constructed in this way will generally be different from those in the first derivatives, a drawback if both first and second derivatives appear in the equation. We could force the same S to appear in both, but this would increase the bandwidth of A.

4. Construction of skew-adjoint finite differences: the mapping method

The mapping method is one of the standard methods of constructing finite differences on nonuniform grids; it also has very attractive adjointness properties. It

is simple to implement and amenable to analysis, and can yield diagonal matrices S which are the same for first and second derivatives and positive on all but extremely distorted grids. There is even a second derivative which is superconvergent, i.e. has order 1 on rough grids instead of the expected 0. The main disadvantage of the mapping method is that its order on rough grids is limited to, at best, 1.

To derive the finite differences we need a precise definition of the 'smoothness' of the grid.

Definition 2. *A family $(x_i^n)_{i\in\mathbb{Z}}^{n>0}$ of grids, periodic in the sense that $x_{i+n}^n = x_i + 2\pi$, is said to be of smoothness class C^r if the function*

$$x(\xi) := \lim_{n\to\infty} x^n_{\lfloor n\xi/2\pi\rfloor}$$

exists and is C^r. The family is rough *if it is C^0 but not C^1.*

We usually omit the superscript and write, e.g., x_i. For example, a grid with $\lim_{n\to\infty} h_i = 0$ but $\lim_{n\to\infty} |h_i/h_{i+1}| \neq 1$ is rough.

For the mapping method we take grid points at $x_i = x(\xi_i) = x(ih)$ where $x(\xi)$ is a given function—i.e., the points are equally spaced in ξ. Then the chain rule gives

$$u_x = (x_\xi)^{-1} u_\xi,$$

which has the natural discretization

$$D\mathbf{u} = S^{-1}A\mathbf{u}$$

where $S = \mathrm{diag}(\mathbf{s})$, $s_i \approx x_\xi(\xi_i)$, and $(A\mathbf{u})_i \approx (u_\xi)(x_i)$, the advantage being that A is a differentiation matrix on a *uniform* grid and can easily be chosen to be antisymmetric. For better performance on rough grids, it is desirable to choose $\mathbf{s} = A\mathbf{x}$ instead of $s_i = x_\xi(\xi_i)$; this makes the approximation exact on linears and hence at least first order. Using central differences at the point ξ_0 gives the tridiagonal finite difference (writing $u_{(n)} = \partial^n u/\partial x^n$, $x_{(n)} = \partial^n x/\partial \xi^n$, and evaluating derivatives at $\xi = 0$,)

$$\frac{u_1 - u_{-1}}{x_1 - x_{-1}} = u_x + \frac{1}{6}h^2\big(3u_{(2)}x_{(2)} + u_{(3)}{x_{(1)}}^2\big) + \mathcal{O}(h^4) \tag{9}$$

if x and u are smooth functions. The discretization is robust in the sense that $s_i = (x_{i+1} - x_{i-1})/(2h) > 0$ on all grids.

The Taylor series expansion (9) also gives a rough idea of the method's performance on nonsmooth grids. If $h_i = hU(a,b)$ where $U(a,b)$ is a random variable uniform on $[a,b]$, then $x_\xi \sim 1$ and $x_{\xi\xi} \sim h^{-1}$; the error drops from $\mathcal{O}(h^2)$ to $\mathcal{O}(h)$. This is confirmed by computing directly that if u is C^2,

$$\left|\frac{u_1 - u_{-1}}{x_1 - x_{-1}} - u_x\right| \leq \max(x_1 - x_0, x_0 - x_{-1}) \max_{x_{-1}\leq x\leq x_1} |u_{xx}(x)|.$$

On C^1 grids, however, $x_{\xi\xi} \sim 1$ and the error remains $\mathcal{O}(h^2)$ (as can also be checked directly from the finite difference).

The analogous fourth order finite difference is

$$\begin{aligned}\frac{-u_2+8u_1-8u_{-1}+u_{-2}}{-x_2+8x_1-8x_{-1}+x_{-2}} = u_x - \frac{1}{30}h^4\Big(&u_{(2)}\big(10\frac{x_{(2)}x_{(3)}}{x_{(1)}}+5x_{(4)}\big)\\ &+u_{(3)}\big(15{x_{(2)}}^2+10x_{(1)}x_{(3)}\big)\\ &+u_{(4)}\big(10{x_{(1)}}^2x_{(2)}\big)\\ &+u_{(5)}\big({x_{(1)}}^4\big)\Big)+\mathcal{O}(h^6),\end{aligned}$$

which indicates that the method is $\mathcal{O}(h)$ on rough (C^0) grids, $\mathcal{O}(h^2)$ on C^1 grids, and $\mathcal{O}(h^4)$ on C^3 grids, *providing* that x_ξ is bounded away from zero. If we take $x(\xi)\sim\xi^2$ as $\xi\to 0$, perhaps to resolve a singularity of u at $x=0$, the method need not even converge as $h\to 0$. Also, $s_i=(-x_{i+2}+8x_{i+1}-8x_{i-1}+x_{i-2})/(12h)$ is not always positive; this requires the mild restriction $\frac{1}{7}<h_i/h_{i+1}<7$.

For second derivatives the pictures is a little different: better for second order methods, worse for fourth order. The chain rule gives

$$x_\xi u_{xx}=(u_\xi/x_\xi)_\xi,$$

but taking $D\mathbf{u}=S^{-1}AS^{-1}A\mathbf{u}$, i.e. using the square of the first derivative approximation, is not a good idea as this leads to large bandwidth approximations with odd and even points decoupled; it does not reproduce the compact central difference on uniform grids.[1] Instead, an analogy with staggered grids suggests using

$$D\mathbf{u}=S^{-1}(-A^tR^{-1}A)\mathbf{u}$$

where $(A\mathbf{u})_i\approx u_\xi(x_{i+1/2})$, $R=\operatorname{diag}(\mathbf{r})$, $r_i\approx x_\xi(\xi_{i+1/2})$, and $S=\operatorname{diag}(\mathbf{s})$, $s_i\approx x_\xi(\xi_i)$.

Using central differences throughout ensures that the approximation is exact on linears and gives the familiar tridiagonal approximation

$$Q:=\frac{\dfrac{u_1-u_0}{x_1-x_0}-\dfrac{u_0-u_{-1}}{x_0-x_{-1}}}{(x_1-x_{-1})/2}=u_{xx}+\frac{1}{12}h^2\big(4u_{(3)}x_{(2)}+u_{(4)}{x_{(1)}}^2\big)+\mathcal{O}(h^4)$$

if x and u are smooth functions. The error looks better than expected, since unlike Eq. (9), no $x_{(3)}$ term appears to leading order. By a happy chance, the finite difference is in fact exact on quadratics, and hence at least first order on any grid. (In a sense it's even better than the analogous first derivative approximation.) It coincides with the approximation obtained by differentiating the local quadratic interpolant, which is surprising, since such approximations do not usually have known adjointness properties. Specifically, if u is C^4 and the grid is C^r, we have

$$\begin{aligned}|Q-u_{xx}| &\le \frac{1}{3}\,|(h_0-h_{-1})u_{xxx}(x_0)|+\frac{1}{6}\max(h_0^2,h_{-1}^2)\max_{x_{-1}\le x\le x_1}|u_{xxxx}(x)|\\ &=\mathcal{O}(h^{1+\min(r,1)}),\end{aligned}$$

while if u is C^3,

$$|Q-u_{xx}|\le\frac{2}{3}\max(h_0,h_{-1})\max_{x_{-1}\le x\le x_1}|u_{xxx}(x)|=\mathcal{O}(h).$$

[1]The use of $D=(S^{-1}A)^2$ was found to be satisfactory in some cases in the cyclo-difference method [3].

Table 1: Finite differences using the mapping method

operator	bandwidth	order
∂_x	tridiagonal	$\mathcal{O}(h)$ on C^0 grids, $\mathcal{O}(h^2)$ on C^1 grids
∂_{xx}	tridiagonal	$\mathcal{O}(h)$ on C^0 grids, $\mathcal{O}(h^2)$ on C^1 grids
∂_x	5-diagonal	$\mathcal{O}(h)$ on C^0 grids, $\mathcal{O}(h^4)$ on C^3 grids, $x_\xi > \delta > 0$
∂_{xx}	7-diagonal	$\mathcal{O}(1)$ on C^0 grids, $\mathcal{O}(h^4)$ on C^4 grids, $x_\xi > \delta > 0$

Table 2: Definition of the test grids

grid name	smoothness	generation equation
smooth	C^∞	$h_i = 1 + \alpha \sin^2(\frac{i\pi}{n})$
sawtooth	C^1	$h_i = \begin{cases} 1 + \alpha(\frac{i}{m}), & \text{when } 1 \le i \le m \\ 1 + \alpha(\frac{m+1+i}{m}), & \text{when } m+1 \le i \le n \end{cases}$
binary	C^0	$h_i = \begin{cases} 1, & \text{when } 1 \le i \le m \\ 1 + \alpha, & \text{when } m+1 \le i \le n \end{cases}$
random	C^0	$h_i = U(1, 1+\alpha)$

At fourth order no such luck obtains. The best we can do is to take

$$(A\mathbf{u})_i = (-u_{i+2} + 27u_{i+1} - 27u_i + u_{i-1})/(24h),$$

$R = \mathrm{diag}(A\mathbf{x})$, and S as for the first derivative. This increases the bandwidth to 7 and gives an approximation equal to

$$u_{xx} + h^4 \mathcal{O}\left(u_{(2)} \frac{{x_{(2)}}^2 x_{(3)}}{{x_{(1)}}^3}, u_{(2)} \frac{x_{(5)}}{x_{(1)}}, \ldots, u_{(6)} {x_{(1)}}^4 \right)$$

which is only $\mathcal{O}(1)$ on rough grids and behaves disastrously if $x_\xi \to 0$.

These results are summarized in Table 1.

5. Selection and performance of the finite differences

Although in Section 3 we established the existence of finite differences, skew-adjoint with respect to diagonal bilinear forms, exact on polynomials of any degree, many issues could not be resolved since the finite differences are determined implicitly. In this section we carry out some numerical tests with the goals of

(1) Validating the order of the exact-on-polynomials methods;
(2) Exploring whether the choice of method is important, and exploring some heuristics for making this choice;
(3) Determining whether there are schemes with accuracy comparable to mapping method schemes on smooth grids, but which degrade more gracefully on rough grids;
(4) Measuring any deterioration of accuracy as a result of imposing skew-adjointness; and
(5) Testing the performance of high-bandwidth, skew-adjoint schemes when solving nonlinear wave equations for long times.

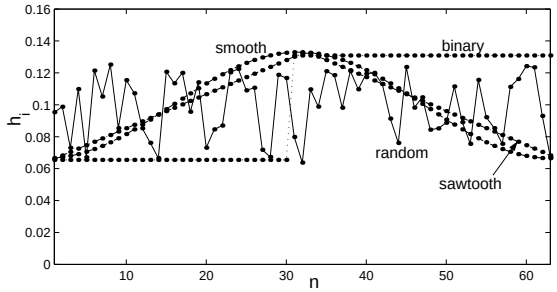

Figure 1: The four test grids for $n = 63$ points and roughness $\alpha = 1$.

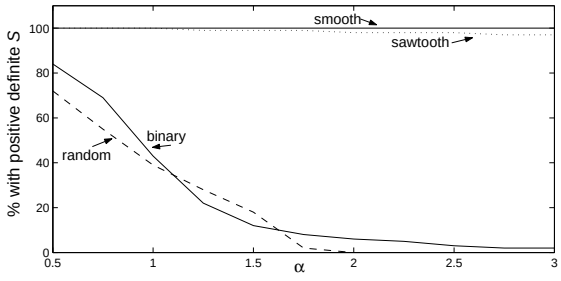

Figure 2: The proportion of second order skew-adjoint finite differences which are stable (i.e. for which S is positive definite) depends on the grid roughness α.

We have used four families of periodic grids, which we call the *smooth*, *sawtooth*, *binary*, and *random* grids (see Table 2 and Figure 1). Each is defined using a roughness parameter α such that $h_i \in [1, 1+\alpha]$: when $\alpha = 0$, each grid is uniform. Each h_i is then scaled by $2\pi / \sum_i h_i$ to produce a grid with period 2π. Because of the known problems of solving wave equations with an even number of grid points [9], we only use an odd number $n := 2m + 1$ of grid points.

Recall the order conditions (7), $L\mathbf{w} = 0$. In the case $d = 2$ (i.e., D is 5-diagonal and exact on quadratics), Proposition 4 shows that $[\mathbf{w}_1; \mathbf{w}_2] := \mathrm{null}(L_0; L_2)$ is 2-dimensional, so we let $\mathbf{w}(\theta) = \mathbf{w}_1 \cos\theta + \mathbf{w}_2 \sin\theta$ and $\mathbf{s}(\theta) = L_1\mathbf{w}(\theta)$ determine the differentiation matrix $D(\theta)$ for $\theta \in (0, \pi]$. We have not performed precise optimizations with respect to θ but simply evaluated $D(\theta_j)$ for $\theta_j = \pi j/100$, $j = 1, \ldots, 100$. 'Method number j' in the figures refers to the differentiation matrix $D(\theta_j)$, where the basis of the null space is chosen as in point 2 below.

Our testing procedure is as follows.

(1) Firstly, if the s_i change sign then we exclude the method from further consideration, since S is not positive definite. That is, we require $s_i/s_1 > 0$ for all i. The fraction of the solutions with positive definite S depend on how rough the grid is and reaches 0 at a finite value of α, as shown in Figure 2. This value is extremely large for the smooth and sawtooth grids, but only 2–4 for the C^0 grids. A possible conclusion is that the requirements of

skew-adjointness with respect to a diagonal bilinear form and order greater than 1 on extremely rough grids are simply incompatible.

(2) Given a family of methods of a certain order, a standard procedure is to choose the one that minimizes the residuals of the order conditions of the next highest order. This is easily done in our case since the order conditions are homogeneous. Let $\mathbf{w} = W\mathbf{c} = \text{null}(L)$ be the general solution of the order conditions and let $P\Lambda Q^t$ be the singular value decomposition of $L_{d+1}W$. We want to

$$\begin{aligned}\min_{\|\mathbf{w}\|=1} \|L_{d+1}\mathbf{w}\|^2 &= \min_{\|\mathbf{c}\|=1} \mathbf{c}^t W^t L_{d+1}^t L_{d+1} W\mathbf{c} \\ &= \min_{\|\mathbf{c}\|=1} \mathbf{c}^t Q\Lambda^t P^t P\Lambda Q^t \mathbf{c} \\ &= \min_{\|\mathbf{y}\|=1} \mathbf{y}^t \Lambda^t \Lambda^t \mathbf{y} \\ &= \min_{\|\mathbf{y}\|=1} \sum \lambda_i^2 y_i^2\end{aligned}$$

where $\mathbf{y} = Q^t\mathbf{c}$. The minimum and maximum errors are given by the minimum and maximum singular values of $L_{d+1}W$. Henceforth we use a basis for W adjusted to this decomposition. In the case $d = 2$, this means that $D(\theta_{50})$ has the minimum 3rd order error and $D(\theta_{100})$ has the maximum 3rd order error.

(3) Very small weights s_i may lead to a form of instability as the ellipsoid $\langle \mathbf{u}, \mathbf{u} \rangle_S$ becomes very elongated. This suggests maximizing $\min s_i / \max s_i$.

(4) When S is positive, D has pure imaginary eigenvalues; its largest in magnitude, $\max|\lambda(D)|$, will limit the time step of any explicit scheme used to integrate $\mathbf{u}_t = D\mathbf{u}$. This suggests minimizing $\max|\lambda(D)|$ to get the "most stable" scheme. Since D is skew-adjoint, we have

$$\max|\lambda(D)| = \|D\|_S$$

so this criterion also ensures that D itself is "small".

To test the methods in practise and to validate the above heuristics requires a test function. Strictly speaking, we should use the whole family of test functions $u = \sin(kx)$, $k = 1, 2, \ldots$; but the best method will depend on k, n, and the grid. We make a compromise choice of $u = e^{2\sin x}$ which includes a range of frequencies, remembering that it will look smoother on finer grids which will favour certain methods. To get a consistent error measure independent on the choice of method, but respecting the grid, we define

$$\|\mathbf{u}\|_h^2 := \sum_{i=1}^{n} u_i^2 \frac{h_i + h_{i-1}}{2} \approx \int u(x)^2\, dx.$$

(5) We measure the error in the derivative,

$$e_1 := \|D\mathbf{u} - \mathbf{u}'\|_h / \|\mathbf{u}'\|_h,$$

where $(\mathbf{u}')_i = u'(x_i)$.

(6) We consider the solution of the test equation $u_t = u_x$ discretized as $\mathbf{u}_t = D\mathbf{u}$ after one period; this is long enough time for high frequency errors to manifest themselves but not so long that the solution disperses entirely. It

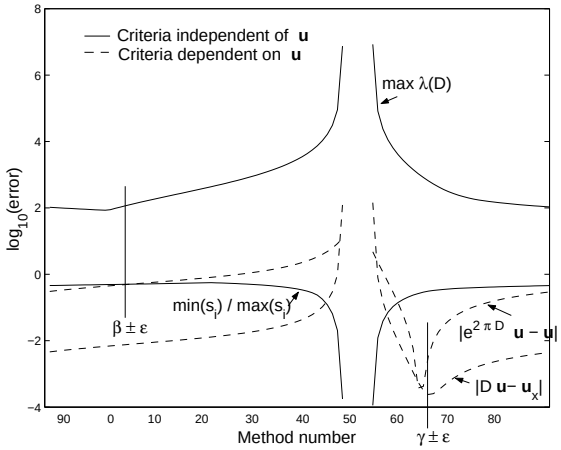

Figure 3: Comparison of four heuristics for choosing the 'best' finite difference matrix D from a 1-parameter family of methods. Note the gap where S is not positive definite.

also means each part of u is exposed to each part of the grid. The error after one period is

$$e_2 := \|e^{2\pi D}\mathbf{u} - \mathbf{u}\|_h / \|\mathbf{u}\|_h.$$

Comparison of the four heuristics is shown in Figure 3 for a typical case. The methods which optimize $\max|\lambda(D)|$ and $(\min s_i / \max s_i)$ are very similar (marked $\beta \pm \epsilon$) and the methods which minimize the errors e_1 and e_2 for the test function $u = e^{2\sin x}$ are also very similar (marked $\gamma \pm \epsilon$). However, the two methods β and γ are very far apart, and while the errors e_1 and e_2 are reasonable for method β, they are far from being as small as possible. The two sets of conditions are essentially independent; this trade-off between time-stability and accuracy is common in finite difference schemes. We have not found a better criterion which is independent of the test function u. Since we want to know if it is *possible* for the high-order exact-on-polynomials methods to beat other approaches, we restrict our attention from now on to the error measure e_2 (relative error after one period for the test function $e^{2\sin x}$).

We can now study the performance of the methods as the grid becomes rougher. Firstly, consider the methods which minimize (method 50) or maximize (method 100) the third order error. Figures 4–5 show that minimizing this error is a good choice provided the grid is not too rough: it is close to optimum for all α on the sawtooth C^1 grid, for $\alpha \leq 0.75$ on the binary grid, and for $\alpha \leq 0.25$ on the random grid. This suggests that it would be a good (and easy to determine) choice on grids which are C^1 in most places with a few small C^0 dislocations, as happens a lot in practise. Curiously, *maximizing* the third order error gives a fairly robust method for a good range of α; but for α large enough it is necessary to optimize to get a good method.

In Figure 6 we compare the best methods to two other approaches. Firstly, the tridiagonal mapping method, discussed in Section 4, is diagonally skew-adjoint with respect to a positive definite S on all grids, but only first order on rough grids

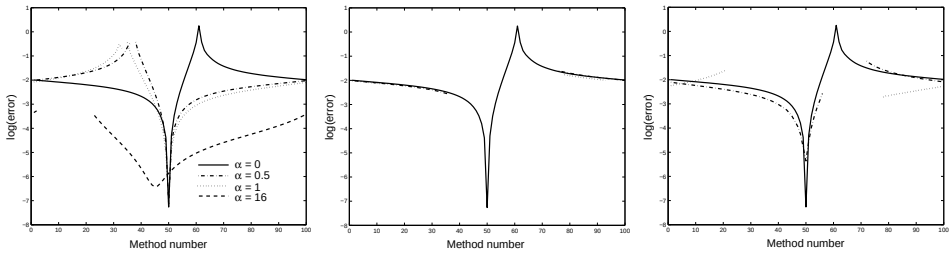

Figure 4: Left: Dependence of the error measure e_2 on the choice of method for the sawtooth (C^1) grid for $n = 63$ grid points. Minimizing the third order error (method 50) always selects a good method—in fact it gives order 4 on C^1 grids. Unstable methods are not shown. All logs are base 10. Centre: For the random (C^0) grid. Right: For the binary (C^0) grid. As α increases, method 50 becomes unstable.

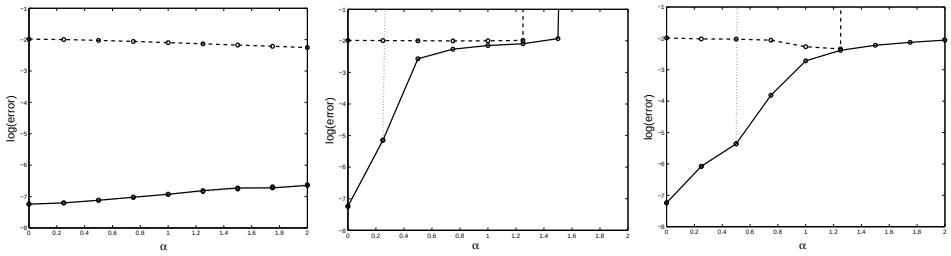

Figure 5: Left: Error e_2 on the sawtooth grid with $n = 63$ for three methods—the smallest (dotted line) and largest (dashed line) third order error, and the smallest error overall (solid line). Centre: for the random grid. On moderately rough grids, *maximizing* the third order error seems to give a fairly good method. Right: For the binary grid.

and second order on smooth grids. The exact-on-polynomial methods can therefore beat it on the sawtooth grid (Figure 6, left) and on C^0 grids (Figure 6, right and centre) provided α is small enough or n is large enough. If the grid is too rough, we see that forcing the method to be exact on quadratics works to its detriment, and the (lower order) mapping method is superior.

Secondly, there is a simple, standard, tridiagonal method given by differentiating the quadratic interpolant through x_{i-1}, x_i, and x_{i+1}. It appears to have pure imaginary eigenvalues for all grids we have tested. However, we have no proof of this, nor can we find a diagonal S with respect to which it is skew-adjoint. It performs very similarly to the mapping method in practise (see Fig. 6). Being exact on polynomials of one higher degree than the mapping method doesn't seem to help it or harm it.

We find that the behaviours described above for the exact-on-quadratics methods persist for higher orders. The theory (Proposition 4) gives orders 2, 3, and 4 for bandwidth 5, 7 and 9 respectively, even on rough grids, and we confirm this numerically (see Table 3). We again observe higher orders for smooth grids. Note the lower than expected orders for the C^∞ grids with the bandwidth 9 method, this

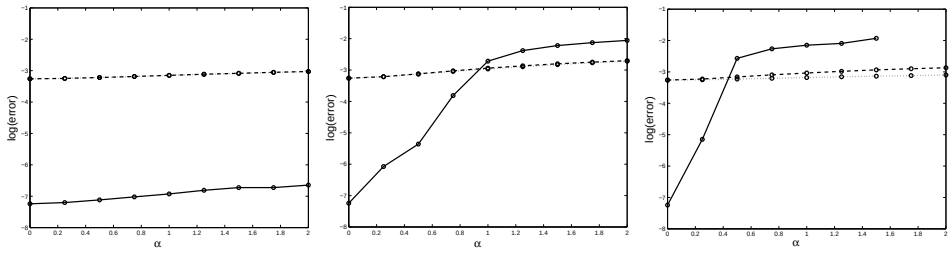

Figure 6: Left: Comparison of the best skew-adjoint method exact on quadratics (solid line) with the (skew) mapping method (dashed line) and the standard exact-on-quadratics method (dotted line, indistinguishable), for the sawtooth grid. Center: For the random grid. Right: For the binary grid.

Table 3: Observed order of the skew-adjoint exact-on-polynomials methods on the four grids for different bandwidths, determined from the best method for $n = 37$, 67, 97, and 127.

grid	bandwidth		
	5	7	9
smooth ($\alpha = 0.5$) C^∞	3.9	6.0	4.6
sawtooth ($\alpha = 0.5$) C^1	3.9	5.9	5.4
binary ($\alpha = 0.5$) C^0	2.2	3.8	3.9
random ($\alpha = 0.5$) C^0	2.0	3.0	4.0

is due to the difficulty in finding the 'best' linear combination of three independent methods.

Finally, we will look at how well the skew-adjoint exact-on-polynomials method performs when solving a nonlinear PDE. Following [9], we will simulate the two-soliton solution of the sine-Gordon equation $u_{tt} = u_{xx} - \sin u$ in a periodic domain with length 40. The one-kink solution with spatial period 10 is

$$u(x,t) = 2 \arcsin \operatorname{sn}\left(\frac{x - ct}{\kappa\sqrt{1-c^2}}; \kappa\right).$$

We have taken $\kappa = 1 - 10^{-5}$ (for which the speed $c = 0.677295615572506$) and placed one kink in the interval $[0, 10]$ and an anti-kink $-u$ in the interval $[30, 40]$.

Following Section 2, this system is Hamiltonian with energy

$$\mathcal{H} = \int \frac{1}{2}\left(p^2 + (q_x)^2\right) + \cos q \, dx,$$

and the discretization

$$\mathbf{q}_t = \mathbf{p}, \quad \mathbf{p}_t = D_2\mathbf{q} - \sin \mathbf{q}$$

is Hamiltonian with energy

$$H = \sum_i \left(\frac{1}{2}\left(p_i^2 + q_i(D_2\mathbf{q})_i\right) + \cos q_i\right) s_i.$$

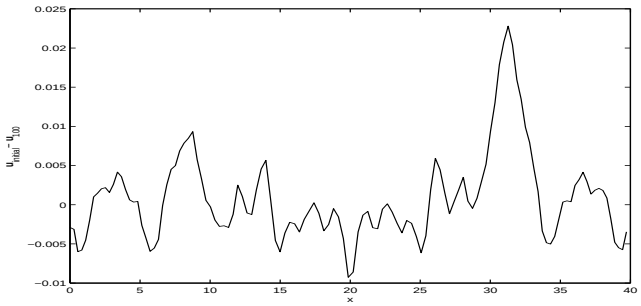

Figure 7: The difference between the initial condition and the solution after 100 collisions for the sine-Gordon equation run on a smooth grid with $\alpha = 8$.

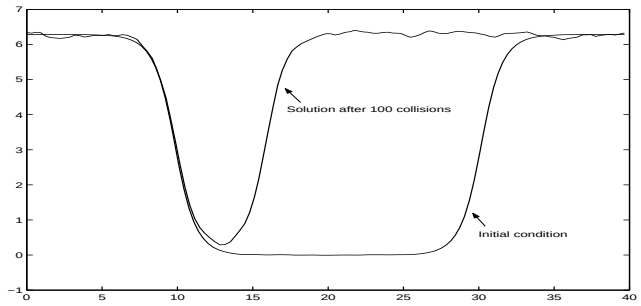

Figure 8: The initial condition and the solution after 100 collisions for the sine-Gordon equation run on a random grid with $\alpha = 0.5$.

Here we use the exact-on-polynomials method to generate a suitable skew-adjoint D_2, exact on cubics.

Using the symplectic leapfrog (Störmer-Verlet) method for the time integration, a typical energy error is shown in Figure 9. The spikes correspond to the collisions of the wavefronts, but there is no drift in the energy. This is what we expect from a symplectic algorithm. Note that because the bandwidth is greater than 3, this method is probably not multisymplectic in the sense of [9], but we still get excellent long-time stability on this equation.

We tested all four families of grids. When the grid was smooth (small α) and C^∞ the method performed well and the wavefront after 100 collisions looked almost exactly the same as the initial condition; for this reason Figure 7 shows the difference between the initial condition and the solution after 100 collisions. On the other hand, in Figure 8 we see the wavefront of the solution on a rough grid after 100 collisions together with its initial condition. The two solitons are still well-formed but there has been some phase drift.

While the wiggles do increase as the roughness α increases they remain bounded except at high α where the energy starts to drift after several collisions.

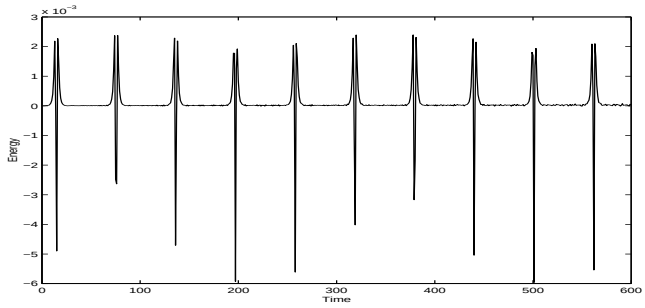

Figure 9: The energy error for the two-soliton solution of the sine-Gordon equation for several collisions run on a random grid ($\alpha = 0.25$).

6. Conclusions

We have established the existence of finite difference methods which are skew-adjoint with respect to diagonal bilinear forms. To get this property, and have order greater than 1 (the limit for the standard mapping method) on rough grids, one has to pay a price: the coefficients of the finite differences are globally coupled, and the bandwidth is increased. (Although we note that the bandwidth 5 methods appear to improve to order 4 on C^1 grids, the order of a standard bandwidth 5 method). As the grid becomes rougher or less differentiable the methods remain stable and their accuracy degrades gracefully. However, for very rough grids, e.g. if h_i/h_{i+1} is too small, even these methods eventually lose stability, in the sense that *none* of the possible diagonal S matrices are positive.

We have mostly considered periodic boundary conditions; other boundary conditions remain an important problem for the future. They change the counting in the order conditions (Eq. (6)) but not the overall approach. If the bandwidth of A is large, then one loses unknowns in A near an (e.g.) Dirichlet boundary, and the local order of approximation has to be reduced appropriately at such grid points.

Another possibility is to consider differences formed from narrow stencils, e.g. using $u_{i+1} - u_i$ to approximate $hu'(x_{i+1/2})$. As in the uniform grid case, these are likely to be more accurate that the wide stencils.

Relaxing the requirement that S be diagonal would give much more scope, indeed this is common in the summation-by-parts literature, but, as shown in Section 2, it does rule out the direct application to nonlinear equations, if one wants to preserve, e.g., energy. It remains to be seen if the present approach is practical for higher-dimensional equations such as fluid and reaction-diffusion equations on rough grids.

Acknowledgement. We are grateful to the Marsden Fund of the Royal Society of New Zealand for financial support.

References

[1] M. H. Carpenter and D. Gottlieb. Spectral methods on arbitrary grids. *J. Comput. Phys.*, 129:74–86, 1996.

[2] M. H. Carpenter, D. Gottlieb, and S. Abarbanel. Time-stable boundary conditions for finite-difference schemes solving hyperbolic systems: methodology and application to high-order compact schemes. *J. Comput. Phys.*, 111:220–236, 1994.

[3] M. H. Carpenter and J. Otto. High-order "cyclo-difference" techniques: An alternative to finite differences. *J. Comput. Phys.*, 118:242–260, 1995.
[4] J. E. Castillo, J. M. Hyman, M. J. Shashkov, and S. Steinberg. The sensitivity and accuracy of fourth order finite-difference schemes on nonuniform grids in one dimension. *Comput. Math. Appl.*, 30(8):41–55, 1995.
[5] D. Gottlieb and S. A. Orszag. *Numerical analysis of spectral methods: theory and applications*. Number 26 in CBMS-NSF Regional Conference Series in Applied Mathematics. SIAM, Philadelphia, 1977.
[6] A. Harten, P. D. Lax, and B. van Leer. On upstream differencing and Godunov-type schemes for hyperbolic conservation laws. *SIAM Rev.*, 25(1):35–61, 1983.
[7] E. Isaacson and H. B. Keller. *Analysis of Numerical Methods*. Dover, New York, 1994.
[8] H. O. Kreiss and G. Scherer. Finite element and finite difference methods for hyperbolic partial differential equations. In C. de Boor, editor, *Mathematical Aspects of Finite Elements in Partial Differential Equations: Symposium 1974*, pages 195–212. Academic Press, 1974.
[9] J. E. Marsden, G. W. Patrick, and S. Shkoller. Multisymplectic geometry, variational integrators, and nonlinear pdes. *Comm. Math. Phys.*, 199(2):351–395, 1998.
[10] R. I. McLachlan, G.R.W. Quispel, and N. Robidoux. Geometric integration using discrete gradients. *R. Soc. Lond. Philos. Trans. Ser. A Math. Phys. Eng. Sci.*, 357:1021–1046, 1999.
[11] R. I. McLachlan and N. Robidoux. Antisymmetry, pseudospectral methods, and conservative PDEs. In B. Fiedler, K. Gröger, and J. Sprekels, editors, *Proc. Int. Conf. EQUADIFF 99*, pages 994–999. World Scientific, 2000.
[12] P. Olsson. Summation by parts, projections, and stability. I. *Math. Comput.*, 64:1035–1065, S23–S26, 1995.
[13] P. J. Olver. *Applications of Lie Groups to Differential Equations*, volume 107 of *Graduate Texts in Mathematics*. Springer-Verlag New York, New York, second edition, 1993.
[14] M. Shashkov. *Conservative Finite-Difference Methods on General Grids*. CRC Press, Boca Raton, 1996.
[15] B. Strand. Summation by parts for finite difference approximations for d/dx. *J. Comput. Phys.*, 110:47–67, 1994.

Institute of Fundamental Sciences, Massey University, Private Bag 11–222, Palmerston North, New Zealand
E-mail address: R.McLachlan@massey.ac.nz, N.Robidoux@massey.ac.nz