



Parallelization, initialization, and boundary treatments for the diamond scheme

Stephen Marsland¹ · Robert I. McLachlan² · Matthew C. Wilkins²

Received: 6 March 2018 / Accepted: 5 July 2019 / Published online: 17 July 2019
© Springer Science+Business Media, LLC, part of Springer Nature 2019

Abstract

We study a class of general purpose linear multisymplectic integrators for Hamiltonian wave equations based on a diamond-shaped mesh. On each diamond, the PDE is discretized by a symplectic Runge–Kutta method. The scheme advances in time by filling in each diamond locally. We demonstrate that this leads to greater efficiency and parallelization and easier treatment of boundary conditions compared with methods based on rectangular meshes. We develop a variety of initial and boundary value treatments and present numerical evidence of their performance. In all cases, the observed order of convergence is equal to or greater than the number of stages of the underlying Runge–Kutta method.

Keywords Multisymplectic integrators · Multi-Hamiltonian PDE · Geometric numerical integration

Mathematics Subject Classification (2010) 37M15 · 37K05 · 65P10

1 Introduction

The diamond scheme is a family of fully discrete numerical methods for first-order hyperbolic PDEs introduced in [10]. It is based on the diamond grid shown in Fig. 1. The family is parameterized by its number of stages, r . The dependent variables are associated with each r node on each edge of the grid; from data on the lower two edges of a diamond, data on the top two edges can be computed *locally within a single diamond*. This feature is unique among schemes of such broad applicability

✉ Robert I. McLachlan
r.mclachlan@massey.ac.nz

¹ School of Mathematics and Statistics, Victoria University of Wellington, Wellington, New Zealand

² School of Fundamental Sciences, Massey University, Palmerston North, New Zealand

and motivates its further exploration. We argued in [10] that the local nature of the scheme indicates that it is highly parallelizable and amenable to local boundary and initialization treatments. In this paper, we present numerical evidence supporting this view.

Following an introduction to the diamond scheme in the context of Hamiltonian PDEs in Section 2, we move on to investigating the parallelization and initialization details of the method. In particular, in Section 3, we present both serial and parallel implementations of the diamond scheme for a nonlinear wave equation. The experimental results are exceptionally good, showing high convergence orders and almost perfect speedup with only $\gtrsim 5$ diamonds per processor. In Section 4, we develop two novel initialization methods and compare their performance with a reference method in which the diamonds are initialized with the exact solution of the differential equation. The observed orders of convergence are at least as good (and sometimes better) than those of the reference method. Dealing with this unconventional aspect of the scheme, in which initial values are not determined by the problem data, is crucial. In Section 5, we develop local boundary treatments for nonhomogeneous Dirichlet and Neumann boundary conditions. These are tested on a variety of linear and nonlinear wave equations for extremely long integration times. In all the tests in all sections, the observed order of convergence of the r -stage method is at least r , although often it exceeds this. The method is marked by its particular theoretical advantages and by its observed performance on a range of tests.

2 Overview of the diamond scheme

We describe the method as it applies to the family of Hamiltonian PDEs:

$$K\mathbf{z}_t + L\mathbf{z}_x = \nabla S(\mathbf{z}), \quad (1)$$

where K and L are constant $n \times n$ real skew-symmetric matrices, $\mathbf{z}: \Omega \rightarrow \mathbb{R}^n$, $\Omega \subset \mathbb{R}^2$, and $S: \mathbb{R}^n \rightarrow \mathbb{R}$. Any solutions to (1) satisfy the *multisymplectic conservation law*:

$$\omega_t + \kappa_x = 0, \quad (2)$$

where $\omega = \frac{1}{2}(d\mathbf{z} \wedge K d\mathbf{z})$ and $\kappa = \frac{1}{2}(d\mathbf{z} \wedge L d\mathbf{z})$ [2]. (That is, if $\mathbf{u}_1, \mathbf{u}_2$ are solutions to the variational equation $K\mathbf{u}_t + L\mathbf{u}_x = S''(\mathbf{z})\mathbf{u}$, then $\omega(\mathbf{u}_1, \mathbf{u}_2) = \mathbf{u}_1^T K \mathbf{u}_2$ and $\kappa(\mathbf{u}_1, \mathbf{u}_2) = \mathbf{u}_1^T L \mathbf{u}_2$.) A numerical method that satisfies a discrete version of (2) is called a *multisymplectic integrator* (see [2, 7] for reviews of multisymplectic integration).

For ODEs, there are effective symplectic integrators—such as symplectic Runge–Kutta methods—that apply to the entire class of ODEs $K\mathbf{z}_t = \nabla S(\mathbf{z})$, and have excellent numerical properties, including symplecticity, arbitrary order, small error constants, unconditional stability, and linear equivariance. One generalization of these methods to the PDEs (1) is to apply high-order Runge–Kutta methods in space and in time on a rectangular mesh [14]. This approach inherits some of the good features listed; some of its properties, including dispersion and order behavior, are studied in [9]. However, the scheme does have some drawbacks: it is fully implicit,

and it leads to discrete equations without a solution for periodic boundary conditions unless r and N (the number of cells in space) are *both* odd [9, 15]. Solvability of the discrete equations is also affected by the boundary conditions, and no general effective treatment of boundary conditions is known for this method.

The first two issues, implicitness and boundary treatment, are related. They can be avoided for some PDEs, like the nonlinear wave equation, by applying suitably partitioned Runge–Kutta methods [8, 15–17]. When they apply, they lead to explicit ODEs amenable to explicit time-stepping, can have high order, and can deal with general boundary conditions. However, the partitioning means that they are not linear methods.

The diamond scheme is a different generalization of symplectic Runge–Kutta methods from ODEs to PDEs. It provides an approach that is multisymplectic, applies to all PDEs of the form (1), is linear in z , and is locally well defined for any number of stages. We first give the definition of the class of diamond schemes that we will consider.

Definition 1 A *diamond scheme* for the PDE (1) is a quadrilateral mesh in space–time together with a mapping of each quadrilateral to a square to which a Runge–Kutta method is applied in each dimension, together with initial data specified at sufficient edge points such that the solution can be propagated forward in time by locally solving for pairs of adjacent edges.

We only consider Runge–Kutta methods in which the nodes c_i lie in $[0, 1]$.

It is convenient to first map each diamond in Fig. 1 to a unit square using the linear transformation T defined by (omitting unnecessary additive constants):

$$T: \quad \tilde{x} = \frac{1}{\Delta x}x + \frac{1}{\Delta t}t, \quad \tilde{t} = -\frac{1}{\Delta x}x + \frac{1}{\Delta t}t. \quad (3)$$




Fig. 1 Schematic of the diamond scheme for periodic boundary conditions. Information flows upwards within individual diamonds, as indicated by the blue arrows. The solution, z , is stored at the Runge–Kutta nodes on the edges, indicated by black dots (shown here for the 3-stage Gauss method), and initialized at the dots on the blue zigzag. A step of the diamond scheme consists of two half-steps. The first half-step calculates z at the dots on the green zigzag; with periodic boundary conditions, information from the right-hand triangle is shared with the left-hand triangle. The second half-step uses z at the dots on the green zigzag to calculate z at the dots on the red zigzag

By the chain rule, (1) transforms to:

$$\tilde{K}\tilde{z}_{\tilde{t}} + \tilde{L}\tilde{z}_{\tilde{x}} = \nabla S(\tilde{z}), \quad (4)$$

where

$$\tilde{K} = \frac{1}{\Delta t}K - \frac{1}{\Delta x}L, \quad \tilde{L} = \frac{1}{\Delta t}K + \frac{1}{\Delta x}L, \quad (5)$$

and $\tilde{z}(\tilde{x}, \tilde{t}) = z(x, t)$. As the PDE (4) is still of the same class as (1), we may apply the multisymplectic Runge–Kutta collocation method given by Reich [14] to (4) within a single unit square; applying T^{-1} yields a method on the diamond lattice.

Let (A, b) be the parameters of an r -stage Runge–Kutta method. In what follows, we will take the method to be the Gauss Runge–Kutta method. Figure 2 shows a diamond with $r = 3$, and its transformation to the unit square. The square contains $r \times r$ internal grid points, as determined by the Runge–Kutta coefficients c , and internal stages Z_i^j , which are analogous to the usual internal grid points and stages in a Runge–Kutta method. The internal stages also carry the variables X_i^j and T_i^j , which approximate z_x and z_t , respectively, at the internal stages.




Fig. 2 The diamond transformed by a linear transformation, T , to the unit square. The square contains $r \times r$ ($r = 3$ in this example) internal stages Z_i^j . The solution is known along the bottom and left-hand sides. The method proceeds as two sets of r Gauss Runge–Kutta r -stage methods: internal stage values, Z_i^j , X_i^j , T_i^j , are calculated, then the right and top updated

The dependent variables of the method are the values of $\mathbf{z}$ at the edge grid points. The Runge–Kutta discretization is:

$$\mathbf{Z}_i^j = \tilde{\mathbf{z}}_\ell^j + \sum_{k=1}^r a_{ik} \mathbf{X}_k^j, \quad (6)$$

$$\mathbf{Z}_i^j = \tilde{\mathbf{z}}_i^b + \sum_{k=1}^r a_{jk} \mathbf{T}_i^k, \quad (7)$$

$$\nabla S(\mathbf{Z}_i^j) = \tilde{K} \mathbf{T}_i^j + \tilde{L} \mathbf{X}_i^j, \quad (8)$$

together with the update equations:

$$\tilde{\mathbf{z}}_r^i = \tilde{\mathbf{z}}_\ell^i + \sum_{k=1}^r b_k \mathbf{X}_k^i, \quad (9)$$

$$\tilde{\mathbf{z}}_i^t = \tilde{\mathbf{z}}_i^b + \sum_{k=1}^r b_k \mathbf{T}_i^k, \quad (10)$$

for $i, j = 1, \dots, r$. Here, $\tilde{\mathbf{z}}_\ell^j$ and $\tilde{\mathbf{z}}_i^b$ are known. Equations (6)–(8) are first solved for the internal stage values $\mathbf{Z}_i^j$, $\mathbf{X}_i^j$, and $\mathbf{T}_i^j$, then (9) and (10) are used to calculate $\tilde{\mathbf{z}}_i^t$ and $\tilde{\mathbf{z}}_r^i$. Equations (6)–(8) are $3r^2$ equations in $3r^2$ unknowns $\mathbf{Z}_i^j$, $\mathbf{X}_i^j$, and $\mathbf{T}_i^j$. Equations (6) and (7) are linear in $\mathbf{X}_i^j$ and $\mathbf{T}_i^j$; these variables may be explicitly eliminated. Thus, in practice, the method requires solving a set of r^2n equations for $\mathbf{Z}_i^j$ in each diamond (where n is the number of dimensions).

The method does not use values at the corners. If necessary, solutions at the corners can be obtained using Runge–Kutta update equations along the edges, combined with averages where two edges meet.

The following basic properties of the diamond scheme are established in [10]. The conservation law in Theorem 2 is a discretization of the integral of $\omega_t + \kappa_x = 0$ over a single diamond, transferred to the boundary of the diamond using Stokes' theorem, and discretized by Gauss quadrature.

Theorem 2 *The diamond scheme satisfies the discrete symplectic conservation law:*

$$\frac{1}{\Delta t} \sum_{i=1}^r b_i \left(\omega_i^t + \omega_r^i - \left(\omega_\ell^i + \omega_i^b \right) \right) + \frac{1}{\Delta x} \sum_{i=1}^r b_i \left(\kappa_r^i + \kappa_i^b - \left(\kappa_i^t + \kappa_\ell^i \right) \right) = 0,$$

where $\omega_n^m = \frac{1}{2} d\mathbf{z}_n^m \wedge K d\mathbf{z}_n^m$, $\kappa_n^m = \frac{1}{2} d\mathbf{z}_n^m \wedge L d\mathbf{z}_n^m$.

We also recall the result of [9] that the Reich collocation scheme with the r -stage Gauss method, when applied on a rectangular mesh to the hyperbolic PDE (1) with initial conditions on $t = 0$ and periodic boundary conditions, has global errors of order at least r . In some cases, order $r + 1$ is observed, which is the stage order of the method. Therefore, we expect convergence of order at least r from the diamond scheme in cases where it is stable.

Table 1 Sample problems

Name	Equation	Range	Solution
Esin	$u_{tt} - u_{xx} = 0$	$0 \leq x \leq 2\pi$	$e^{2 \sin(x-t-3)}$
Sincos	$u_{tt} - u_{xx} = 0$	$0 \leq x \leq 2\pi$	$\sin(x) \cos(t)$
Coscoss	$u_{tt} - u_{xx} = -u$	$0 \leq x \leq \pi$	$\cos(2x) \cos(\sqrt{5}t)$
sine-Gordon	$u_{tt} - u_{xx} = -\sin(u)$	$-30 \leq x \leq 30$	$4 \arctan \left(\frac{\sin(\frac{t}{\sqrt{2}})}{\cosh(\frac{x}{\sqrt{2}})} \right)$

One major difference between the diamond and rectangular meshes is that the diamond mesh means that the method is effectively multistage. Indeed, it is the extra initial data, at different time levels, that allow the diamonds to be filled in independently. In the ODE case $L = 0$, the diamond scheme yields an r -step integrator whose underlying 1-step method is the original Gauss method.

3 Diamond implementation

In this section, we describe both serial and parallel implementations of the diamond scheme. The aim is to demonstrate the numerical stability of the method as well as to evaluate the speedup resulting from the parallelization.

We test the diamond scheme numerically on several different wave equations $u_{tt} - u_{xx} = f(u)$ on the domain $x \in [a, b]$, $t \in [0, T]$; these are listed in Table 1. The initial conditions are Cauchy, i.e., $u(x, 0)$ and $u_t(x, 0)$ are given. In this section, we use periodic boundary conditions at $x = a$ and $x = b$, but in Section 5 we demonstrate methods to impose Dirichlet and Neumann boundary conditions and demonstrate them numerically.

Note that there are other methods (in particular, Lobatto IIIA–IIIB in space and an explicit symplectic splitting method in time) that perform outstandingly well on nonlinear wave equations. We adopt this equation simply as a first test: if the method fails here, it almost certainly fails overall.

The wave equation has several formulations of the form (1). Here, we use the formulation with $v = u_t$, $w = u_x$, $\mathbf{z} = (u, v, w)^T$:

$$K = \begin{pmatrix} 0 & -1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}, \quad L = \begin{pmatrix} 0 & 0 & 1 \\ 0 & 0 & 0 \\ -1 & 0 & 0 \end{pmatrix}, \quad (11)$$

$$S(\mathbf{z}) = -V(u) + \frac{1}{2}v^2 - \frac{1}{2}w^2, \text{ and } f(u) = V'(u). \quad (12)$$

Then, we have the following stability result.

Theorem 3 [9] *The diamond scheme with $r = 1$ applied to the wave equation $u_{tt} = u_{xx}$ with periodic boundary conditions is linearly stable when $\lambda = \frac{\Delta t}{\Delta x} \leq 1$.*

Indeed, from Fig. 1, there will always be a CFL stability condition that a characteristic originating in the bottom left edge must not intersect the bottom right edge. For the wave equation, we conclude that $\frac{\Delta t}{\Delta x} \leq 1$ is necessary for stability regardless of the underlying Runge–Kutta scheme.

Three implementations of the diamond scheme were prepared: serial implementations in Python and C and a parallel version in C. Although the scheme is most naturally expressed using rank-3 tensors, for the implementations, the tensors were flattened to matrices. Solving the nonlinear equations on each diamond requires a nonlinear solver. The Python code used the SciPy [5] routine `fsolve`, which is a wrapper around the MINPACK [11] `hydrd` and `hybrj` algorithms, which are based on the Powell hybrid method [13]. The C codes used the GNU Scientific Library [4] routines `gsl_multiroot_fsolver_hybrids` and `gsl_multiroot_fdfsolver_gnewton` which again are wrappers around the MINPACK `hydrd` and `hybrj` algorithms. The `hydrd` algorithm approximates the Jacobian, whereas `hybrj` requires the exact Jacobian; both versions were used and compared.

3.1 Convergence of the serial implementation

We first test the convergence of the serial implementation on a nonlinear wave equation. The diamond scheme with varying r was used to solve the sine-Gordon equation $u_{tt} - u_{xx} = -\sin(u)$. An exact solution is the so-called breather (see also Table 1):

$$u(x, t) = 4 \arctan \left(\frac{\sin \left(\frac{t}{\sqrt{2}} \right)}{\cosh \left(\frac{x}{\sqrt{2}} \right)} \right). \quad (13)$$

The domain is taken significantly large, $[a, b] = [-30, 30]$, so the solution can be assumed to be periodic. (Note that $u_x(-30, \sqrt{2}) \approx 3 \times 10^{-9}$.) The error is the discrete 2-norm of u :

$$E^2 = \frac{b-a}{N} \sum_{i=1}^N (u_i - u(a + i\Delta x, T))^2. \quad (14)$$

The number of diamonds at each time level is $N = 40, 80, \dots, 1280$, and the integration time, T , is twice the largest time step. The Courant number $\frac{\Delta t}{\Delta x} = \frac{1}{2}$ is held fixed. The $2rN$ initial values of $z = (u, u_t, u_x)$ needed at the bottom edge of the first row of diamonds are provided by the exact solution; we will consider other methods of initialization in Section 4.

The results for the global error are shown in Fig. 3. It is apparent that for this problem the method converges, and the order appears to be r when r is odd and $r + 1$ when r is even.

3.2 Parallel implementation and speedup

In the parallel implementation of the diamond scheme, the domain is divided into strips, finite in width in the x direction and potentially infinite in the t direction. The width L of the domain is a multiple of Δx . There are $N = \frac{L}{\Delta x}$ diamonds in each




Fig. 3 The error of the diamond scheme with varying r applied to the multisymplectic Hamiltonian PDE arising from the sine-Gordon equation. The true solution was the so-called breather on the domain $[-30, 30]$. The Courant number is fixed at $\frac{1}{2}$ as Δt is decreased. The order of the method appears to be r when r is odd and $r + 1$ when r is even (reproduced from [10])

row, and p processors. The N diamonds are divided as equally as possible into p contiguous regions. If N is a multiple of p , then each processor will get an equal number of diamonds to work on. Otherwise, $p - n$ processors get k diamonds, and n processors get $k + 1$ diamonds, where $n = N - pk$ and $k = \left\lfloor \frac{N}{p} \right\rfloor$.

Each processor calculates the solution on its strip of diamonds. After the first half time step (see Fig. 1), the solution at the right edge of the strip must be passed to the right-hand neighbor and received from the left. This results in p transmits of vectors of length $n(r + 1)$ in length. At the next half time step, the left-hand edge must be passed to the left and received from the right, another p transmits. If the solution is to be output, each processor must send a number of values—twice the number of diamonds in its sub-domain—to one processor. This communication was achieved with an MPI gather. Since this processor has very slightly more work to perform, it is desirable to ensure it is one of the $p - n$ processors that receive k diamonds to work on.

To test how much the diamond scheme could benefit from running in parallel, it was run on the New Zealand eScience Infrastructure's (NeSI's) Pan Intel Linux cluster physically located at the University of Auckland, New Zealand. At the time of use, the cluster had approximately 6000 cores each running somewhere between 2 and 3 GHz with most around 2.7 GHz or 2.8 GHz. Most cores had at least 10 GB of RAM available, far more than we required. Due to the busyness, of the cluster it was impractical to request specific CPUs for each job; thus, there was a certain variability in timing tests simply because of the different speeds of the CPUs available. However, the uni-processor jobs, arguably the most important while testing parallel speedups, did run on the most common 2.7-GHz or 2.8-GHz processors.

The diamond scheme was initialized with the triangle initialization method detailed in Section 4.1, r was set to 5, the number of time steps was 1000, $\Delta t = 0.05$, and the periodic sine-Gordon problem from Table 1 was used. Despite the cluster having approximately 6000 cores, by trial and error, it was apparent that only about a maximum of 300 or 400 cores could be readily available on demand. So each trial consisted of nine runs with the number of cores being 1, 3, 7, 20, 56, 100, 150, 300, and 350. For each run, the wall-clock time was recorded using the Unix `date` command, the program run, and then the wall-clock time checked again. Each trial (set of nine runs) was performed twice with a couple of days in between each trial, and the time results averaged.

According to Amdahl's law [1], for a particular problem size, if n is the number of cores, and $B \in [0, 1]$ is the fraction of the algorithm that is strictly serial, then the theoretical time $T(n)$ for the algorithm to run on n cores is:

$$T(n) = T(1) \left(B + \frac{1}{n}(1 - B) \right).$$

Thus, the theoretical speedup $S(n)$ is:

$$S(n) = \frac{T(1)}{T(n)} = \frac{1}{B + \frac{1}{n}(1 - B)}. \quad (15)$$

Letting $n \rightarrow \infty$ gives a theoretical maximum speedup of $\frac{1}{B}$. By increasing n until the speedup begins to tail off, it is possible to estimate B . For a perfectly parallelizable algorithm, the speedup should be equal to the number of cores used. In the first trial, shown on the left in Fig. 4, Δx was such that there were 4000 diamonds across the domain.




Fig. 4 Speedup of the diamond scheme versus the number of cores for the code running on the NeSI Pan cluster. Code that was perfectly parallelizable would have the speedup equal to the number of cores (the blue line). As the ratio of the number of cores to the amount of work (number of diamonds across the domain) increases, one would expect the speedup to deviate from the perfect blue line. On the left, with at least $4000/350 > 11$ diamonds per core, the speedup is very good. On the right, with as few as 2 diamonds per core, the performance deteriorates. The speedup was calculated from a single run (not an average of two runs) for the 450 and 499 number of cores runs

As the ratio of the number of cores to the amount of work (number of diamonds across the domain) increases, one would expect the speedup to diverge from the perfect speedup line. This is because the overhead in communication will gradually swamp the gains in computation time. In this first trial, the speedup is very good and it is impossible to estimate B , the fraction of the algorithm that is strictly serial. Ideally, the number of cores would be increased until a deviation from the perfectly parallelizable line could be reliably detected; however, no more cores were easily available.

So, instead of increasing the number of cores, the number of diamonds was decreased. Figure 4 (right) shows the second trial where the number of diamonds across the domain was decreased to 1000. One of the trials included two extra runs with $n = 450$ and $n = 499$. Because it took many days for these runs to begin executing, the second trial did not include these large runs, and no averaging could take place. This figure shows the speedup reaching approximately 250 before beginning to tail off. So for this size of problem, from (15) this equates to $B \approx 0.4\%$, which is remarkably low. The conclusion is that the diamond scheme is exceptionally parallelizable.

4 Diamond scheme initialization

The diamond mesh creates an issue for initialization which is not present for rectangular meshes. In this section, we develop two initialization methods, the *triangle initialization* and the *phantom diamond initialization* methods, and compare them with the reference “method” in which initial values are taken from the exact solution (as used in the previous section).

4.1 Triangle initialization method

For the nonlinear wave equations in Table 1, initial conditions for u and u_t are specified at $t = 0$. Differentiating with respect to x , we can assume that u_x , and hence all components of $\mathbf{z}$, are known at $t = 0$.

The x -axis cuts the first row of diamonds in half, yielding a row of triangles. In the triangle initialization method, each triangle is mapped to the unit square, then the usual set of (6)–(10) solved, giving values for $\mathbf{z}$ on the top left and top right edges of the triangle.

The transformation, $x = \tilde{x} - \tilde{t}$, $t = \tilde{x}\tilde{t}$, illustrated in Fig. 5, takes the unit square to the triangle $(-1, 0)$, $(1, 0)$, $(0, 1)$.

Adding a translation and scaling results in the map:

$$x = \frac{\Delta x}{2}(\tilde{x} - \tilde{t}) + b \quad (16)$$

$$t = \frac{\Delta t}{2}\tilde{x}\tilde{t} \quad (17)$$




Fig. 5 The unit square under the map $x = \tilde{x} - \tilde{t}$, $t = \tilde{x}\tilde{t}$

which takes the unit square to the triangle $(b - \frac{\Delta x}{2}, 0)$, $(b + \frac{\Delta x}{2}, 0)$, $(b, \frac{\Delta t}{2})$. Recall that the transformed K and L are given by:

$$\begin{aligned}\tilde{K} &= g_t K + g_x L \\ \tilde{L} &= f_t K + f_x L,\end{aligned}$$

where $(\tilde{x}, \tilde{t}) = (f(x, t), g(x, t))$. This yields:

$$\begin{aligned}\tilde{K} &= \frac{2}{\Delta x \Delta t (\tilde{x} + \tilde{t})} (\Delta x K - \tilde{t} \Delta t L) \\ \tilde{L} &= \frac{2}{\Delta x \Delta t (\tilde{x} + \tilde{t})} (\Delta x K + \tilde{x} \Delta t L).\end{aligned}$$

In summary, the triangle initialization method consists of four steps:

1. The analytical solution for $\mathbf{z} = (u, u_t, u_x)$ is known along the whole edge DAB in Fig. 5.
2. This provides exact values for $\tilde{\mathbf{z}}$ on the Runge–Kutta nodes on the edges $\tilde{D}\tilde{A}$ and $\tilde{A}\tilde{B}$.
3. The square $\tilde{A}\tilde{B}\tilde{C}\tilde{D}$ is completed (using the existing diamond method algorithm and code) which provides values of $\tilde{\mathbf{z}}$ at the Runge–Kutta nodes on the edges $\tilde{D}\tilde{C}$ and $\tilde{C}\tilde{B}$.
4. These values map directly to the values of $\mathbf{z}$ at the Runge–Kutta nodes on the edges DC and CB .

Note that the map in step 4 is the identity, because $\tilde{\mathbf{z}}(\tilde{x}, \tilde{t}) = \mathbf{z}(x, t)$ and the locations of the Runge–Kutta nodes correspond under the map (16) which is linear on the edges.

The initialization method cannot affect the stability. However, it potentially affects the accuracy, particularly in view of the singularity in (16) at $(x, t) = (0, 0)$. We therefore study the convergence of the method numerically.

Figure 6 shows the error of exact and triangle initialization as Δt is reduced while keeping the Courant number $\frac{\Delta t}{\Delta x} = \frac{1}{2}$. The integration time $T = 1.5$ is twice the largest time step. It is apparent that for this problem the triangle initialization has errors equal to or less than that of exact initialization. Table 2 shows the observed convergence order of the two initialization methods.

At first sight, this observation is surprising. However, any error result requires some mapping from the infinite-dimensional function spaces used for the initial data




Fig. 6 The error of the diamond scheme initialized using the exact and triangle methods (Section 4.1) applied to the sine-Gordon equation. The exact solution is a so-called breather on $[-30, 30]$. The Courant number is fixed at $\frac{1}{2}$ as Δt is decreased. The error is the 2-norm of the global error at the final time (see (14)). For this problem, the triangle initialization is as accurate or better than exact initialization

to the finite-dimensional spaces of the numerical method. Methods that use “processing” [9] increase the order of convergence by choosing this mapping carefully. In this sense, the exact initialization, which uses the mapping that restricts the exact solution to the nodes, is just one choice among many.

We next apply the triangle initialization method to the four different wave equations with different forcings and initial conditions given in Table 1. In each case, the boundary conditions are periodic. The number of diamonds at each time level is

Table 2 Observed order of convergence of the diamond scheme initialized with the exact solution and with the triangle initialization method

r	Order	
	Exact	Triangle
1	1	2
2	3	3
3	3	5
4	5	5
5	5	6

(Data obtained from Fig. 6.) It is apparent that for this problem the triangle initialization performs as well, or better, than the exact initialization. The order appears to be $r + 1$ for most r (for $r = 4$ the order is $r + 2$ for some reason), whereas for the exact initialization the order is r (r odd) and $r + 1$ (r even)




Fig. 7 The error of the diamond scheme with varying numbers of stages r applied to the Esin, Sincos, Cocos, and sine-Gordon problems with periodic boundary conditions and the triangle initialization method (see Table 1). The Courant number is fixed at $\frac{1}{2}$ as Δt is decreased. Table 3 summarizes the observed convergence order given by the slope of these lines

$N = 10, 20, \dots, 1280$, and the integration time $T = 1.5$ is twice the largest time step.

The computed global errors are shown in Fig. 7. Table 3 shows the observed convergence order of the diamond scheme for these problems. It is apparent that for this problem, the order is at least r .

Table 3 Observed convergence order of the problems given in Table 1 (data from Fig. 7)

r	Order			
	Esin	Sincos	Cocos	sine-Gordon
1	2.0	2.1	2.1	1.8
2	1.9	1.9	1.9	4.1
3	4.0	3.4	3.3	5.6
4	3.8	3.7	3.8	5.4
5	6.2	4.8	4.6	6.8

The observed convergence order of the r stage method is at least r

4.2 Phantom diamond initialization method

In this method, a phantom diamond is constructed about the $t = 0$ axis as illustrated in Fig. 8. The bottom left and bottom right edges are missing values of $\mathbf{z}$; these will now become free variables, whose values will be determined by the Runge–Kutta update equations. To compensate, more information will be specified at the stages on $t = 0$. This information will be obtained from the initial conditions. For any underlying r -stage Runge–Kutta method, r of the r^2 internal stages lies on $t = 0$.

We illustrate the procedure for the nonlinear wave (11) with initial conditions $u(x, 0) = h(x)$, $u_t(x, 0) = g(x)$. Each edge node contains 3 degrees of freedom. The bottom left and bottom right edges together contain $6r$ degrees of freedom. Therefore, additional $6r$ values must be specified at the r internal stages on $t = 0$, or 6 additional values per stage. Each of these stages has 3 degrees of freedom for each of the variables $\mathbf{Z}$ ($\approx \mathbf{z} = (u, u_t, u_x)$), $\mathbf{X}$ ($\approx \mathbf{z}_x$), and $\mathbf{T}$ ($\approx \mathbf{z}_t$). We proceed as follows. In each case the given data is evaluated, and the resulting condition imposed, at the Runge–Kutta nodes on $t = 0$.

1. The initial conditions specify u and $u_t (= v)$, providing the data $Z_1 = h(x)$ and $Z_2 = g(x)$.
2. The third component of the PDE, namely $w = u_x$, provides the data $Z_3 = h'(x)$.
3. Differentiating the second and third components of the PDE with respect to x gives $v_x = u_{tx}$ and $w_x = u_{xx}$. These provide the data $X_2 = g'(x)$ and $X_3 = h''(x)$, respectively.
4. The equation $u_{tx} = u_{xt} = w_t$ provides the 6th and final required piece of data, namely $T_3 = g'(x)$.

In total, we now have 6 out of the 9 components of $\mathbf{Z}$, $\mathbf{X}$, and $\mathbf{T}$ specified at each $t = 0$ internal stage. The Runge–Kutta equations now yield a closed system that can be solved locally within each phantom diamond separately, and the update equations yield the values of $\mathbf{z}$ on the top left and top right edges.

The results are shown in Fig. 9. The phantom diamond initialization method does not appear to adversely affect the order of the scheme: it is still at least r . In addition,




Fig. 8 An initial phantom diamond in the $r = 2$ scheme. The solution $\mathbf{z}$ is not known on the bottom left or bottom right edges. Internally, nothing is known at the stars, and some information is known at the dashed circles. There needs to be enough known at the dashed circles to match the missing information on the bottom left and right stars




Fig. 9 The error of the diamond scheme initialized with the triangle method (Section 4.1) and the phantom diamond method (Section 4.2), with varying numbers of stages r applied to the Coscos D-N problem (see Table 4). The Courant number is fixed at $\frac{1}{2}$ as Δt is decreased. The phantom diamond initialization method is as good as, or better than, the triangle initialization method

the errors are either close to (when r is odd) or much small than (when r is even) those obtained with the triangle initialization method.

Although the phantom diamond initialization method works well in this example, it relies on being able to compute enough information from the initial conditions. We do not know if this is possible for all PDEs in all formulations of the form (1).


5 Dirichlet and Neumann boundary conditions

The construction of stable, high-order methods for hyperbolic initial-boundary value problems is not easy.¹ Nevertheless, the compactness of the diamond scheme indicates that we might hope to construct entirely local boundary treatments, systematically for all r .

We have developed and tested many approaches. A basic requirement is that the resulting method should be solvable and stable subject to a CFL condition, ideally $\Delta t/\Delta x < 1$. In addition, we found that some methods that worked well on a simple test problem (e.g., the linear wave equation with homogeneous Dirichlet boundary conditions) did not work on more complicated problems. Thus, extensive testing was required. For example, the following methods were not robustly able to solve wave equations with a variety of boundary conditions: (i) specifying some components of

¹For example, for wide-stencil finite difference methods, a significant development effort over many years has yielded stable methods using the *summation by parts* and *simultaneous approximation term* approaches [3, 6, 12, 19]. These finite difference operators approximate u_x (resp. u_{xx}) at all points, using different finite differences near the boundary. Stability is achieved by requiring that the finite difference is skew- (resp. self-) adjoint with respect to an inner product, designed along with the method.

Fig. 10 A left-hand boundary phantom diamond in the $r = 2$ scheme. The solution $\mathbf{z}$, is known at the circles on the bottom right edge. For an internal diamond, $\mathbf{z}$ would also be known at the stars on the bottom left edge, but not in this case. Internally, nothing is known at the stars, and some information is known at the dashed circles. There needs to be enough known at the dashed circles to match the missing information on the bottom left stars



$\mathbf{z}$ at more points on the boundary than just the r stages; (ii) using extra information from an adjacent interior diamond; (iii) mapping the boundary triangle to a square, as in the triangle initialization method; and (iv) combinations of these.

The method that was ultimately successful is the following *boundary scheme* that we describe for Dirichlet and Neumann boundary conditions for the nonlinear wave equation. In both cases, we use the phantom diamond as shown in Fig. 10, with $3r$ conditions imposed at the internal stages on the left-hand boundary $x = a$, to compensate for the $3r$ missing values at the bottom left edge. The entire set of (8) is then solved simultaneously, and the top right edge values filled in using the update (9), (10). The approach is similar to that used in the phantom diamond initialization method, although the counting is different as there is more information available at the edges and fewer boundary conditions.

The $3r$ new conditions are equations in the $9r$ internal dependent variables $\mathbf{Z}$ ($\approx \mathbf{z} = (u, v, w)$), $\mathbf{X}$ ($\approx \mathbf{z}_x$), and $\mathbf{T}$ ($\approx \mathbf{z}_t$). These equations come from the boundary conditions, their derivatives with respect to t , and the derivatives of the PDE with respect to x and t . However, note that we cannot use the PDE itself as it is already imposed at the internal stages.

Table 4 Sample nonperiodic problems

Name	Domain	Left boundary	Right boundary
Esin DD	$0.2 \leq x \leq \frac{\pi}{3}$	Dirichlet	Dirichlet
Sincos DD	$0.2 \leq x \leq \frac{\pi}{3}$	Dirichlet	Dirichlet
Sincos DN	$0.2 \leq x \leq \frac{\pi}{3}$	Dirichlet	Neumann
Coscoss DD	$0.2 \leq x \leq \frac{\pi}{3}$	Dirichlet	Dirichlet
Coscoss DN	$0.2 \leq x \leq \frac{\pi}{3}$	Dirichlet	Neumann
sine-Gordon DD	$-2 \leq x \leq 2$	Dirichlet	Dirichlet

See Table 1 for the equations and exact solutions. The Dirichlet boundary conditions are found using the exact solution, and the Neumann conditions by differentiating the exact solution with respect to x

1. For the Dirichlet boundary condition $u(a, t) = g(t)$, the 3 new conditions at each of the r stages at $x = a$ are obtained in the following way:
 - (a) From $u(a, t) = g(t)$, we get $Z_1 = g(t)$;
 - (b) By differentiating the boundary condition, we get $u_t(a, t) = g'(t)$, yielding $Z_2 = g'(t)$; and




Fig. 11 The error of the diamond scheme with varying r applied to the Esin DD (Dirichlet/Dirichlet boundary conditions), Sincos DD, Sincos DN (Dirichlet/Neumann), Coscos DD, Coscos DN, and sine-Gordon DD problems (see Table 4). The Courant number is fixed at $\frac{1}{2}$ as Δt is decreased. Table 5 summarizes the observed convergence order given by the slope of these lines

Table 5 Observed convergence orders for the initial-boundary value problems given in Table 4

r	Order					
	Esin DD	Sincos DD	Sincos DN	Coscoss DD	Coscoss DN	sine-Gordon DD
1	2.3	2.4	2.2	2.3	2.2	2.0
2	2.2	2.4	2.1	2.4	2.1	3.2
3	3.7	3.8	3.6	4.0	4.1	4.4
4	4.2	4.5	3.8	4.3	4.1	4.5
5	5.2	5.4	5.4	5.4	5.3	6.3

To one significant figure, the order appears to be r , although in some cases it exceeds this

- (c) Differentiating again gives $u_{tt}(a, t) = g''(t)$, which together with the PDE $u_{tt} - u_{xx} = f(u)$ gives $w_x(a, t) = g''(t) - f'(g(t))$, yielding $X_3 = g''(t) - f'(g(t))$.
2. For the Neumann boundary condition $u_x(a, t) = h(t)$, the 3 new conditions at each of the r stages at $x = a$ are obtained in the following way:
 - (a) From $u_x(a, t) = h(t)$ and $u_x = w = Z_3$, we get $Z_3 = h(t)$;
 - (b) By differentiating the boundary condition, with respect to time, we get $w_t = h'(t)$, yielding $T_3 = h'(t)$; and
 - (c) Equality of mixed partial derivatives $v_x = w_t$ yields $X_2 = h'(t)$.

This boundary scheme is now applied to the four sample problems given in Table 1 with a mix of Dirichlet and Neumann boundary conditions. Table 4 summarizes the problems. In each case, 10^6 – 10^7 time steps were computed to ensure that the equations were solvable at each step and the solutions remained bounded.

Because the exact solution is known for all the sample problems, it is easy to impose whatever boundary condition is desired on any spatial domain. The domains were chosen so the solutions were not periodic or symmetric in any way, because while testing other potential methods, it became apparent that using “easy” problems gave false confidence in the method. To ensure only one thing was tested at a time, the triangle initialization scheme was used. As a comparison, the phantom diamond initialization method was used for one problem (Coscoss DN). Because the domains are smaller than the periodic tests, a smaller number of diamonds was used, $N = 2, 4, \dots, 128$. Figure 11 shows the error for the various test problems as Δt is decreased. From this data, the observed order of convergence is summarized in Table 5.

6 Conclusions

The novel diamond mesh introduced in [10] raised hopes that it would be suitable for parallelization and for a wide variety of initial-boundary value problems, with local boundary closures not affecting the interior part of the scheme. We have presented

numerical evidence that this is indeed possible. It still remains to prove stability and the observed orders of convergence. That this may be possible is suggested by the known results that the Reich method (for a wide class of equations of the form (1), (4)) has convergence order at least r , and that it also preserves discrete forms of arbitrary quadratic conservation laws, whenever the PDE has any such [18].

Funding information This research was supported by the Marsden Fund of the Royal Society Te Apārangi and by a Massey University PhD Scholarship.

References

1. Amdahl, G.M.: Validity of the single-processor approach to achieving large scale computing capabilities. AFIPS Conference Proceedings **30**, 493–485 (1967)
2. Bridges, T.J., Reich, S.: Numerical methods for Hamiltonian PDEs. J. Phys. A **39**, 5287–5320 (2006)
3. Carpenter, M., Gottlieb, D., Abarbanel, S.: 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)
4. Gough, B., Galassi, M.: GNU Scientific Library Reference Manual, Network Theory Ltd 3 Ed. (2009)
5. Jones, E., Oliphant, T., Peterson, P.: SciPy: Open source scientific tools for Python. <http://www.scipy.org/> (2001)
6. Kreiss, H., Scherer, G.: On the Existence of Energy Estimates for Difference Approximations for Hyperbolic Systems, Tech. Report, Technical Report, Dept. of Scientific Computing, Uppsala University (1977)
7. Leimkuhler, B., Reich, S.: Simulating Hamiltonian Dynamics, Cambridge Monogr. Appl. Comput. Math. Cambridge University Press, Cambridge (2004)
8. McLachlan, R., Sun, Y., Tse, P.: Linear stability of partitioned runge–Kutta methods. SIAM J. Numer. Anal. **49**, 232–263 (2011)
9. McLachlan, R.I., Sun, Y., Ryland, B.: High order multisymplectic runge–Kutta methods. SIAM J. Sci. Comput. **36**, A2199–A2226 (2014)
10. McLachlan, R.I., Wilkins, M.: The multisymplectic diamond scheme. SIAM J. Sci. Comput. **37**, A369–A390 (2015)
11. Moré, J.J., Garbow, B.S., Hillstom, K.E.: User Guide for MINPACK-1, ANL-80-74 Argonne national laboratory (1980)
12. Olsson, P.: Summation by parts, projections, and stability. I. Math. Comput. **64**, 1035–1065 (1995)
13. Powell, M.J.D.: An efficient method for finding the minimum of a function of several variables without calculating derivatives. Comput. J. **7**, 155–162 (1964)
14. Reich, S.: Multi-symplectic runge–Kutta collocation methods for Hamiltonian wave equations. J. Comput. Phys. **157**, 473–499 (2000)
15. Ryland, B.: Multisymplectic integration, PhD thesis, Massey University, New Zealand (2007)
16. Ryland, B.N., McLachlan, R.I.: On multisymplecticity of partitioned runge–Kutta methods. SIAM J. Sci. Comput. **30**, 1318–1340 (2008)
17. Ryland, B.N., McLachlan, R.I., Frank, J.: On the multisymplecticity of partitioned runge–Kutta and splitting methods. Int. J. Comput. Math. **84**, 847–869 (2007)
18. Sun, Y.: Quadratic invariants and multi-symplecticity of partitioned runge–Kutta methods for Hamiltonian PDEs. Numer. Math. **106**, 691–715 (2007)
19. Svård, M., Nordström, J.: Review of summation-by-parts schemes for initial-boundary-value problems. J. Comput. Phys. **268**, 17–38 (2014)

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.