INTERNATIONAL JOURNAL FOR NUMERICAL METHODS IN FLUIDS
Int. J. Numer. Meth. Fluids 0000; **00**:1–29
Published online in Wiley InterScience (www.interscience.wiley.com). DOI: 10.1002/fld

Multirate Time Stepping for Accelerating Explicit Discontinuous Galerkin Computations with Application to Geophysical Flows

B. Seny[1*], J. Lambrechts[1], R. Comblen[1], V. Legat [1], J.-F. Remacle[1]

[1] *Université catholique de Louvain, Institute of Mechanics, Materials and Civil Engineering (iMMC), Bâtiment Euler, Avenue Georges Lemaître 4, 1348 Louvain-la-Neuve, Belgium*

SUMMARY

This paper presents multirate explicit time stepping schemes for solving Partial Differential Equations with Discontinuous Galerkin (DG) elements in the framework of large-scale marine flows. It addresses the variability of the local stable time steps by gathering the mesh elements in appropriate groups. The real challenge consists to develop methods exhibiting mass conservation and consistency. Two multirate approaches, based on standard Explicit Runge-Kutta methods, are analyzed. They are well suited and optimized for the DG framework. The significant speedups observed for the hydrodynamic application of the Great Barrier Reef confirm the theoretical expectations.

Received . . .

KEY WORDS: Discontinuous Galerkin; High Performance Computing; Multirate time stepping; Explicit Runge-Kutta; Shallow water equations; Great Barrier Reef

1. INTRODUCTION

The development of suitable and fast time integration methods for ocean modeling constitutes an important challenge. It is indeed impossible to use one single time-discretization scheme that is effective for all physical processes in a complex marine model, as different subsystems have widely different characteristics in terms of time scales, dynamic behaviour, and accuracy requirements. The primitive equations for ocean flows allow for the existence of phenomena exhibiting a wide spectrum of propagation speeds. Typically, external gravity waves propagate at $10 - 100\ ms^{-1}$ and internal waves at a few meters per second, whereas advection is characterized by speeds ranging from 10^{-3} to $1\ ms^{-1}$. Large- and small-scale processes have significant interactions so that it is essential to simulate them simultaneously. It seems impossible today to reproduce all scales with structured uniform grids since the computational cost can become very crippling with the high resolution that is required. Therefore variable resolution is needed both temporally and spatially.

*Correspondence to: bruno.seny@uclouvain.be

Unstructured grids are well suited to capture complex topography and also allow the representation of a wide spectrum of time and length scales in a single model. The finite volumes and the finite elements are the two main methods that make use of unstructured grids. Many groups are now developing finite volume codes for coastal applications like FVCOM (Finite Volume Community Ocean Model) [1] and others [2, 3]. Finite elements methods are also widely used in the area of large-scale ocean modeling by communities such as FEOM (Finite Element Ocean Model) [4], ICOM (Imperial College Ocean Model) [5, 6] and [7, 8]. Our research team is developing the **S**econd-generation **L**ouvain-la-Neuve **I**ce-ocean **M**odel (SLIM)[†] [9, 10, 11, 12] which is a discontinuous Galerkin-based finite element model.

The variable resolution and the complexity of unstructured mesh generation processes generally lead to grids with an important dispersion of element sizes. More specifically, while finite element meshers are able to control the average element size of a mesh, they are usually unable to control the smallest element size. In this context, the classical conservative explicit time discretization methods are limited due to stability requirements. The Courant-Friedrichs-Lewy (CFL) condition, that combines the finest cell size and the highest wave velocity, may highly restrict the global allowable explicit time step. Accordingly, the computational efficiency of explicit time-stepping methods may be drastically low.

For instance, consider the case of a typical mesh of the Great Barrier Reef (GBR), illustrated by Figure 1, made up of about 1 million triangles. This mesh is built by means of the open source software GMSH[‡] [13]. Element sizes were determined in order to capture the relevant bathymetric and topographic features, and the associated hydrodynamic processes, such as eddies and tidal jets [10]. For the mesh and bathymetry presented in Figure 1, the estimated minimum and maximum stable time steps are $0.154\ s$ and $7.972\ s$, respectively. To run a 24 hours simulation with a classical explicit method, $561\ 039$ time steps would have to be performed on almost 1 million elements. One possibility to reduce these expensive computations is to adapt the time steps under local stability conditions.

Multirate schemes represent a class of methods that use various time steps on different grid cells. The strategy consists in gathering the grid cells in different groups that satisfy the local CFL stability conditions for a certain range of time steps. Standard Explicit Runge-Kutta (ERK) methods are applied on bulk groups with a local time step in such a way that the total computational efforts are drastically reduced. Buffer groups are introduced, with adapted ERK methods, in order to accommodate the transition between the different bulk groups. However, the development of such methods is still challenging since convergence and conservation properties should remain satisfied during the communication between the groups. In this context, two multirate approaches that attempt to partly solve the transitions issues are explored. The first one, introduced by Constantinescu and Sandu, [14], preserves the system invariants but is at most second order accurate. On the other hand, Schlegel et al. [15] have proposed a method that borrows some ideas of the implicit-explicit splitting scheme [16, 17]. It can be proved that a third order multirate scheme can be achieved with an appropriate base ERK method. Unfortunately, this method turns out to be non-conservative.

[†] http://www.climate.be/SLIM
[‡] http://www.geuz.org/gmsh

DOI: 10.1002/fld

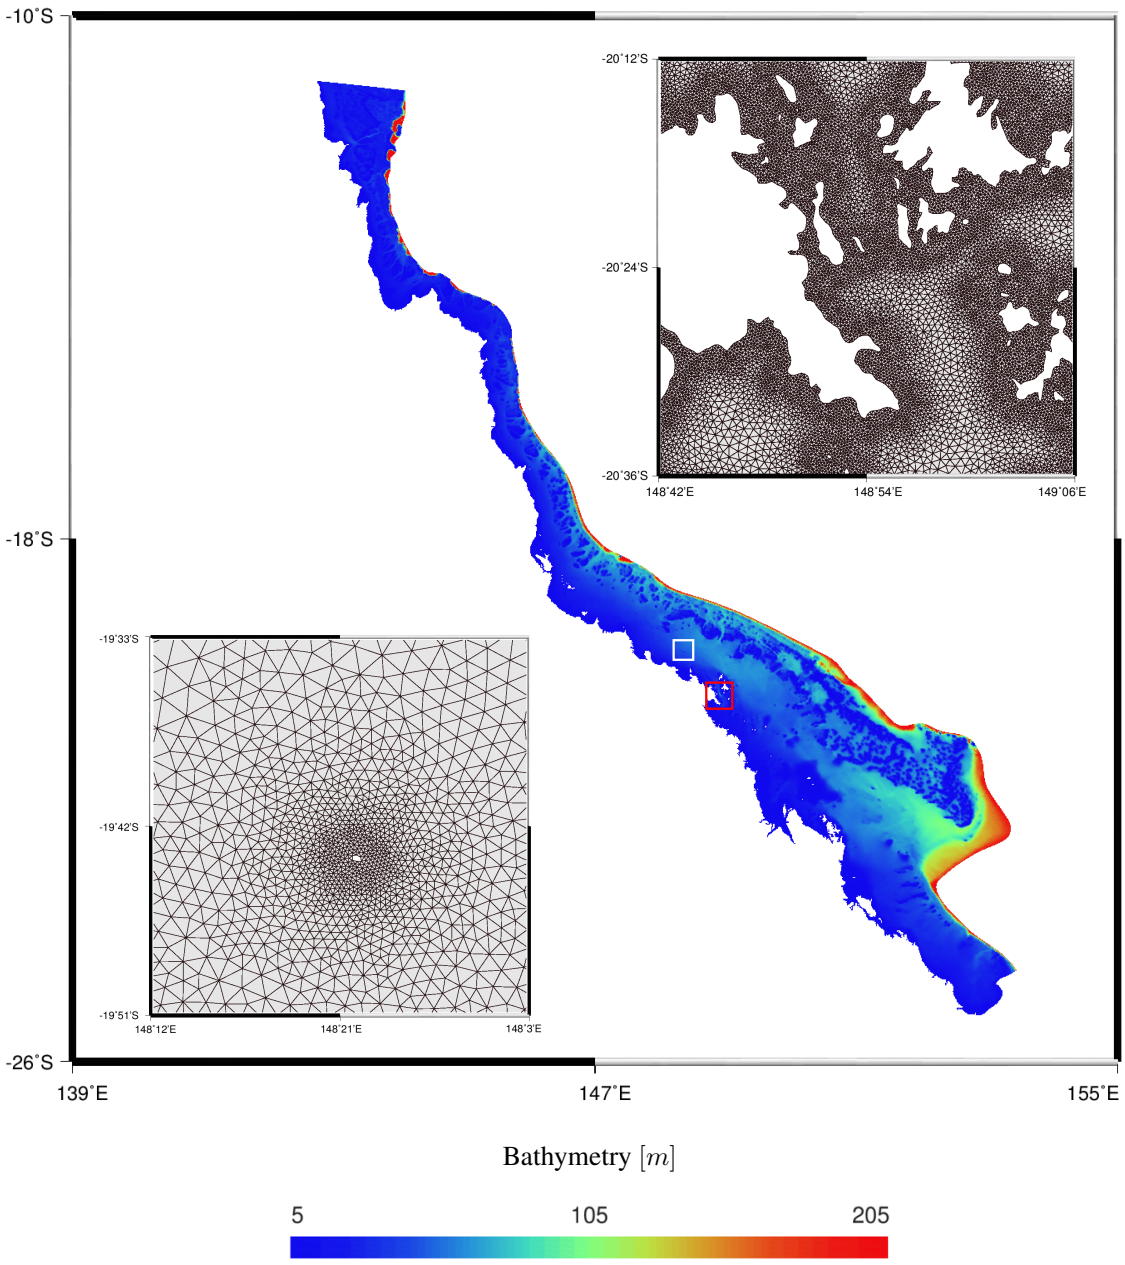

Figure 1. Bathymetry and mesh of the GBR with a first zoom (bottom left, white) on the Holbourne Island and a second one (upper right, red) on the Withsundays Inslands Archipelago. The mesh is made up of 909,185 triangles, with inner radii comprised between *circa* 29 m and 1.3 km, and 444,598 nodes.

The aim of this paper is to develop and adapt these multirate methods to large unstructured meshes in the framework of the Dicontinuous Galerkin Method (DGM). The standard ERK methods and their time step restrictions are described in Section 2. Two multirate approaches [14, 15], with different features, are introduced and analyzed in the DGM framework, Section 3. The construction of multirate groups for multiple levels of refinement and a way to optimize the speedup is addressed in Section 4. Finally numerical experiments will be shown and discussed for the GBR in Section 5.

DOI: 10.1002/fld

2. EXPLICIT TIME INTEGRATION

When solving time dependent partial differential equations (PDEs), it is a common practice to first discretize the spatial variables in order to obtain a semi-discrete method of lines (MOL) by leaving the time variable continuous. In other words the spatial and temporal discretizations are then independent. The advantage of this procedure is that the problem reduces to a system of ordinary differential equations (ODEs) to which a numerical method for initial value ordinary equations can be applied.

As an illustration, consider the case of a one-dimensional scalar advection equation written in a conservative form

$$\frac{\partial u}{\partial t} + \frac{\partial}{\partial x}(cu) \quad = \quad 0 \tag{1}$$

with initial condition

$$u(x,0) \quad = \quad u^0(x) \tag{2}$$

and appropriate boundary conditions. Assume a conservative DGM spatial discretization of (1) which can be represented by a function $f(u_i, t)$. It contains the volume and interface parts of the steady-state residual of the problem multiplied by the inverse of the mass matrix associated the element. By notational convenience, we define $u_i \triangleq (u_{i,1} \ \cdots \ u_{i,n})$ as the set off all local discrete values defined in element Ω_i. The semi-discrete DG approximation can be written, for each grid element Ω_i, as the following Cauchy problem

$$\frac{du_i}{dt}(t) = f(u_j(t), t), \qquad u_i(0) = u_i^0 \tag{3}$$

which needs to be solved in time. A class of numerical methods, among many, to integrate the solution in time is the family of Explicit Runge-Kutta (ERK) schemes. The MOL approach may be extended to any conservation law with larger dimensions and/or multiple unknown fields. Other discretization techniques may also be used to approximate the spatial terms of the PDEs.

2.1. Explicit Runge-Kutta schemes

ERK methods are among the most popular time stepping schemes [18, 19]. They are self-starting meaning that they give the solution at the next time step only in terms of the current solution. Therefore only the initial condition is needed to start the time integration. Other well-known schemes like the Adam-Bashfort methods are multistep and use solutions at different previous time steps. ERK methods have also the property of being relatively flexible. For instance the time step may be changed at each iteration of the scheme. These explicit schemes may also be developed up to high orders of accuracy with the constraint that a r^{th} order RK method needs $s \geq r$ inner stages, i.e., evaluations of the steady-state residual and multiplication by the inverse mass matrix. A widely used ERK method is the classical 4 stages, 4^{th} order scheme (RK44). Flux limiters, strongly recommended for hyperbolic conservation laws with DG discretization in space, may be applied to the solution in a simple manner. They ensure that non-oscillatory properties are achieved for strong

Prepared using fldauth.cls
DOI: 10.1002/fld

shocks. An s-stage ERK method computes the next step solution u^{n+1}, at time $t^{n+1} = t^n + \Delta t$, using the current solution u^n at t^n by applying the following algorithm:

- For k=1 : s do

$$u^{(k)} = u^n + \Delta t \sum_{l=1}^{k-1} A_{kl} K^l \tag{4}$$

$$K^k = f\left(u^{(k)},\, t^n + c_k \Delta t\right) \tag{5}$$

- Compute u^{n+1} as

$$u^{n+1} = u^n + \Delta t \sum_{l=1}^{s} b_l K^l \tag{6}$$

Butcher tableaus conveniently represent this family of ERK methods. They are defined by a matrix $A \in \Re^{s \times s}$ and two vectors $b, c \in \Re^s$ [20]:

$$\begin{array}{c|c} c & A \\ \hline & b^T \end{array}$$

where A is strictly lower triangular. For consistency, it is required that $\sum_{l=1}^{k-1} A_{kl} = c_k$. The order of the method is related to the constraints imposed on A_{kl}, b_l and c_k. The $u^{(k)}$ variable represents the solution at stage k of the method, that corresponds to the intermediate time $\tilde{t}^k = t^n + c_k \Delta t$. At each stage of the ERK method K^k is computed, i.e., the steady-state residual evaluated for $u^{(k)}$ and multiplied by the inverse of the mass matrix. Afterwards, the next step solution u^{n+1} is obtained by summing u^n with a linear combination of the K^k. These methods can be rewritten as a convex combination of Euler steps [21]. Therefore they may be categorized in the family of strong stability preserving (SSP) time stepping schemes. This property ensures that a certain norm, like the total variational (TV) norm [21], of the solution does not increase in time. SSP numerical methods are often required for problems with discontinuous solutions, such as shock waves in hyperbolic problems. Non-physical behaviors like spurious oscillations can be avoided in this way. Gottlieb et al. [22] discussed in detail Runge-Kutta SSP schemes and several examples of these methods can be found in [21]. As an example, consider the 2-stage, second order method, RK2a, defined by the Butcher tableau represented in Table I.

0		
1	1	
	1/2	1/2

$$\begin{aligned} u^{(1)} &= u^n, & K^1 &= f(u^{(1)}, t^n) \\ u^{(2)} &= u^n + \Delta t K^1, & K^2 &= f(u^{(2)}, t^n + \Delta t) \\ u^{n+1} &= u^n + \tfrac{1}{2}\Delta t(K^1 + K^2) \end{aligned}$$

Table I. Butcher tableau and development of RK2a method

DOI: 10.1002/fld

2.2. *Time-step restrictions*

Even if ERK time integration methods are known to be very efficient for solving several types of PDEs, they have a major drawback due to their stability requirements. Indeed, it is well-attested that the global time step should be taken below a critical value, determined by the CFL condition. For advection dominated advection-diffusion equations, the CFL constraint on the time step can be expressed as a ratio of the grid spacing Δx and the amplitude of the wave/advective velocity c. But in almost all realistic scenarios, the CFL condition is not constant both spatially and temporally.

On the one hand, unstructured meshes have elements with a wide spectrum of sizes. Several numerical applications require that some regions of the domain are examined more closely. Local refinement is often needed to capture the topography of complex geometry and/or some specific physical behaviors. On the other hand, even for structured meshes, the wave speed may vary considerably across the entire domain. As an example, consider the case of the two-dimensional shallow water equations, where the wave speed is defined as $\sqrt{gH}$ with g defined as the gravity and $H(x, y)$ a strong varying water depth depending on the local horizontal coordinates. The global time step is then determined by the element where $H(x, y)$ reaches its maximum.

For problems with unstructured meshes, made up of N grid elements, and non-constant wave velocities, the CFL condition can be written as:

$$\max_{\Omega_i} \left(\frac{|c|}{\Delta x} \right) \Delta t \quad < \quad C, \tag{7}$$

where Ω_i represents the element number $i \in [1, ..., N]$ of the mesh. The constant C depends on the particular PDE and on the ERK scheme that defines the shape of the stability zone [23]. This is a severe restriction on the time step in order to guarantee overall stability. In the case of the GBR mesh of Figure 1, the global time step is critically smaller than the one required for most elements.

Using fully implicit time integration schemes constitutes a way to avoid the restriction mentioned before. In such strategies the only limitations on the time step come from accuracy purposes. The drawback, however, is that implicit methods require solving large (non-) linear systems of equations. Indeed, the dimension of the systems to solve increases with the number of degrees of freedom.

Another alternative, while using explicit schemes, is to use a multirate approach. The main idea is to consider different regions in the discretized spatial domain where the CFL condition is locally satisfied. Mesh elements are sorted, by their own characteristic stable time step, in different groups, specified by a maximum time step, for which ERK methods are stable and achieve the target accuracy. With locally adapted time steps the computational efforts of the global algorithm could be considerably reduced.

To illustrate this, consider a one-dimensional mesh, represented on Figure 2, where elements Ω_i have a size equal to h except Ω_0 that is twice smaller. If equation (1) is to be solved on this mesh, assuming a constant wave speed c, one can determine the stable time steps for both kinds of cells. If $\frac{1}{2}\Delta t$ is stable for Ω_0 then Δt may be assumed stable for $\Omega_{i>0}$. Consider that we are able to apply the same s-stage ERK method with a time step $\frac{1}{2}\Delta t$ on Ω_0 and Δt on the other elements. For a large value of N, the number of elements in the mesh, one can show that a speedup of 2 is obtained compared to the same s-stage ERK method applied with the same time step $\frac{\Delta t}{2}$ everywhere.

DOI: 10.1002/fld

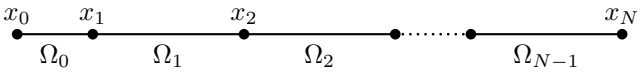

Figure 2. One-dimensional unstructured mesh

But, in such a strategy , there exists an inconsistency at the interface between the small and large element since they use a different time step. Therefore a coherent transition should be ensured at the interface between groups of elements. In particular, convergence and conservation properties should be fulfilled at the interfaces. This constitutes the major difficulty in developing multirate schemes.

3. MULTIRATE TIME INTEGRATION

Multirate schemes for conservation laws have been reported in the literature since the early 1980s, but they were either locally inconsistent or not mass-conservative. Mass preserving multirate schemes were developed by Osher and Sanders [24] as well as by Dawson and Kirby [25] but it turns out that the timestepping accuracy of the overall method is only first order due to the treatment of the interfaces. Tang and Warnecke [26] proposed multirate schemes, based on standard 2-stage ERK methods, that achieve second order consistency in time. The drawback, however, is that the resulting schemes are not mass preserving. Hundsdorfer et al. [27] discuss, within the framework of partitioned Runge-Kutta methods, the defects of multirate methods of order 1 and 2 due to either the local inconsistency or the lack of mass conservation. They give a particular attention to monoticity properties of the considered multirate schemes.

Multirate methods have also been developed in the framework of self-adjusting strategies. Savcenco et al. [28] consider implicit time stepping methods suitable for stiff ODEs. Those methods use an error estimator to determine if smaller time steps are required to keep the error below given tolerance for all components. The aim is to minimize the execution time without loosing accuracy. At the interfaces it is necessary to interpolate the solutions associated to different times. In this context Hundsdorfer et al. [29] studied a particular multirate scheme: the θ-method with one level of temporal local refinement. Stability, local accuracy and propagation of interpolation errors are analyzed in detail.

A first step, when developing multirate schemes, is undeniably to ensure that the different local time steps are well synchronized. A solution for a coherent time progression is to combine different time steps th ofat are integer multiples of each other. For the sake of simplicity the analysis will be restricted to groups that have time step ratios of 2. If a reference time step Δt_* is assumed for the group with the largest stable time step, the other partitions will be time-integrated with stable time steps $\Delta t_*/2^z,\ z = 1, \cdots z^*$, with $z^* + 1$ the number of groups and z the multirate exponent of the group. Multirate schemes may also be developed for arbitrary integer time step ratios κ where the different groups would have stable time steps defined as $\Delta t_*/\kappa^z$.

The difficulty, when developing multirate strategies, is to manage the transition between groups of elements that use a different stable time step. Indeed, if a final time is to be reached, the number of stages of the respective ERK methods won't be the same on two neighboring elements that belong to different multirate groups. Some information is thus missing in order to ensure a

DOI: 10.1002/fld

coherent transition. This problem of communication reveals two underlying issues: conservation of the fluxes and accuracy of a multirate method. Both Constantinescu [14] and Schlegel [15] proposed a solution by introducing buffer groups. In these regions, adapted ERK methods are applied to bridge the transition between bulk groups where the standard ERK methods are used. Both multirate approaches are based on Partitioned Runge-Kutta (PRK) schemes that are used to solve problems with two different ERK methods [30, 18].

3.1. Second order conservative Multirate Runge-Kutta schemes

The idea of Constantinescu and Sandu [14] is to extend singlerate ERK methods to multirate ERK methods. They developed a general systematic approach, based on PRK methods [30], to construct a family of second order multirate PRK schemes (MPRK-2). For the sake of convenience, the following notation is used: $\mathcal{RK}^x[\Delta t_*]$ means that a given ERK method x is used with a Δt_* time step and its associated Butcher tableau. The methodology consists in choosing a second order accurate s-stage ERK base method $\mathcal{RK}^b = [A^b,\ b^b,\ c^b]$ and extend it to a $2s$-stage ERK method in the buffer region. It should ensure that the transition between two partitions that have a time step ratio of 2. Constantinescu [14] has shown that, if the ERK base method belongs to the family of the SSP time discretization methods, the corresponding multirate scheme will maintain this property.

For the sake of simplicity, the development of the multirate approach of Constantinescu is detailed through a basic one-dimensional example on a mesh similar to Figure 2. This will clarify the features and the size of the buffer regions.

3.1.1. Introductory example First of all let's define some conventions and notations. Assume that the right-hand side f_i, in equation (3), computed on element Ω_i, can be split into the volume contribution $f_{i,i}$ and the left and right interface contributions $f_{i-1,i}$ and $f_{i+1,i}$, depicted in Figure 3, that use information from the neighboring elements:

$$f_i(u_{i-1}, u_i, u_{i+1}, t) \quad = f_{i,i} + f_{i-1,i} + f_{i+1,i} \tag{8}$$

The volume and interface terms for each element at each stage k of an ERK method may be defined as follows:

$$f_{i,i}^{(k)} = f_{i,i}(u_i^{(k)}, t^n + c_k\Delta t), \qquad f_{i-1,i}^{(k)} = f_{i-1,i}(u_{i-1}^{(k)}, u_i^{(k)}, t^n + c_k\Delta t) \tag{9}$$

At each stage of the ERK method $K_i^k = f_{i,i}^{(k)} + f_{i-1,i}^{(k)} + f_{i+1,i}^{(k)}$. For a classical singlerate ERK method $f_{i-1,i}^{(k)} = -f_{i,i-1}^{(k)}$ since the normals at the interface between two neighboring elements are opposite each other. This property ensures global conservation of the fluxes after each iteration of the method.

Consider the second order accurate SSP ERK base method RK2a represented in Table II (a) and the mesh depicted in Figure 2. The key idea is to extend the RK2a method, $\mathcal{RK}^b[\Delta t_*]$, to a 4-stage adapted Butcher tableau $\mathcal{RK}^a[\Delta t_*]$, Table II (b), where the base method is repeated twice on the same time interval. Actually, $\mathcal{RK}^a$ is strictly equivalent to the base method $\mathcal{RK}^b$ if it is used on all elements. The Butcher tableau shown in Table II (c) corresponds to the base method applied twice successively with the same time step $\frac{1}{2}\Delta t_*$. Actually this Butcher tableau contains implicitly the

DOI: 10.1002/fld

update for $u^{n+\frac{1}{2}}$. In other words RK2a is applied a first time to u^n to get $u^{n+\frac{1}{2}}$, that corresponds to time $t^n + \frac{1}{2}\Delta t_*$, and then again to $u^{n+\frac{1}{2}}$ in order to compute u^{n+1}. The methods $\mathcal{RK}^a[\Delta t_*]$ and $\mathcal{RK}^{2b}[\Delta t_*]$ have now the same number of stages and therefore it enables the transition between the two types of elements.

0		
1	1	
	1/2	1/2

(a) $\mathcal{RK}^b = [A^b, b^b, c^b]$

0				
1	1			
0	0	0		
1	0	0	1	
	1/4	1/4	1/4	1/4

(b) $\mathcal{RK}^a = [A^a, b^a, c^a]$

0				
1/2	1/2			
1/2	1/4	1/4		
1	1/4	1/4	1/2	
	1/4	1/4	1/4	1/4

(c) $\mathcal{RK}^{2b} = [A^{2b}, b^{2b}, c^{2b}]$

Table II. Butcher tableaus corresponding to (a) the RK2a base method, (b) the buffer adapted method and (c) RK2a applied twice successively with a 2 times smaller time step

Consider the setup of Figure 3: we apply $\mathcal{RK}^{2b}$ to element Ω_0 and $\mathcal{RK}^a$ to elements $\Omega_{1,2,3}$. Using the Butcher tableaus of Table II and the definition of the volume and interface contributions (8, 9) on each element, it is now possible to develop the computations related to each stage of the method. In order to distinguish the intermediate times, defined by the c vectors of the Butcher tableau's, notice that $\tilde{t}_i^k$ represents the current inner time used at stage k on element Ω_i.

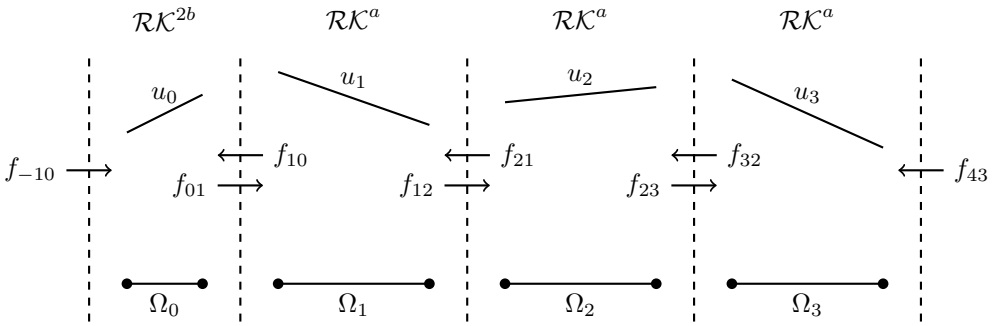

Figure 3. One-dimensional unstructured mesh with interface fluxes, $f_{i,i-1}$ and $f_{i-1,i}$, between neighboring elements Ω_{i-1} and Ω_i. $\mathcal{RK}^{2b}$ is applied on element Ω_0 while $\mathcal{RK}^a$ is used for elements $\Omega_{1,2,3}$.

At the first stage of the coupled methods there are no ambiguities. It is identical to apply the same base method RK2a everywhere. The K_i^1 are all computed at the same intermediate time level $\tilde{t}_i^1 = t^n$. Virtual incoming fluxes, f_{-10} and f_{43}, are supplied at the boundary of the domain as represented on Figure 3. It is assumed that the virtual element Ω_{-1} (resp. Ω_4) is of the same type – same size, same ERK method applied – as its neighbor Ω_0 (resp. Ω_3).

$$u_0^{(1)} = u_0^n, \qquad u_1^{(1)} = u_1^n, \qquad u_2^{(1)} = u_2^n, \qquad u_3^{(1)} = u_3^n$$

$$K_0^1 = f_{-10}^{(1)} + f_{00}^{(1)} + f_{10}^{(1)}, \quad K_1^1 = f_{01}^{(1)} + f_{11}^{(1)} + f_{21}^{(1)}, \quad K_2^1 = f_{12}^{(1)} + f_{22}^{(1)} + f_{32}^{(1)}, \quad K_3^1 = f_{23}^{(1)} + f_{33}^{(1)} + f_{43}^{(1)}$$

DOI: 10.1002/fld

At the second stage the intermediate time levels are not the same on each element, i.e., $\tilde{t}_0^2 = t^n + \frac{1}{2}\Delta t_*$ while $\tilde{t}_i^2 = t^n + \Delta t_*$ for $i = 1, 2, 3$.

$$u_0^{(2)} = u_0^n + \tfrac{1}{2}\Delta t_* K_0^1, \qquad u_1^{(2)} = u_1^n + \Delta t_* K_1^1, \qquad u_2^{(2)} = u_2^n + \Delta t_* K_2^1, \qquad u_3^{(2)} = u_3^n + \Delta t_* K_3^1$$

$$K_0^2 = f_{-10}^{(2)} + f_{00}^{(2)} + f_{10}^{(2)}, \quad K_1^2 = f_{01}^{(2)} + f_{11}^{(2)} + f_{21}^{(2)}, \quad K_2^2 = f_{12}^{(2)} + f_{22}^{(2)} + f_{32}^{(2)}, \quad K_3^2 = f_{23}^{(2)} + f_{33}^{(2)} + f_{43}^{(2)}$$

Several simplifications can be done at the third stage and are highlighted with bold characters. Every entry of the third row in Table II (b) is zero and therefore $u_i^{(3)} = u_i^{(1)}$ for $i = 1, 2, 3$. It follows that $f_{i,j}^{(3)} = f_{i,j}^{(1)}$ if and only if i and j belong to the set $\{1, 2, 3, 4\}$ implying that $K_2^3 = K_2^1$ and $K_3^3 = K_3^1$. At this stage the intermediate times are different: $\tilde{t}_0^3 = t^n + \frac{1}{2}\Delta t_*$ while $\tilde{t}_i^3 = t^n$ for $i = 1, 2, 3$.

$$u_0^{(3)} = u_0^n + \tfrac{1}{4}\Delta t_*(K_0^1 + K_0^2), \qquad \mathbf{u_1^{(3)}} = u_1^n = \mathbf{u_1^{(1)}}, \qquad \mathbf{u_2^{(3)}} = u_2^n = \mathbf{u_2^{(1)}}, \qquad \mathbf{u_3^{(3)}} = u_3^n = \mathbf{u_3^{(1)}}$$

$$K_0^3 = f_{-10}^{(3)} + f_{00}^{(3)} + f_{10}^{(3)}, \quad K_1^3 = f_{01}^{(3)} + f_{11}^{(3)} + f_{21}^{(3)} = f_{01}^{(3)} + f_{11}^{(1)} + f_{21}^{(1)}, \quad \mathbf{K_2^3} = f_{12}^{(3)} + f_{22}^{(3)} + f_{32}^{(3)} = f_{12}^{(1)} + f_{22}^{(1)} + f_{32}^{(1)} = \mathbf{K_2^1}, \quad \mathbf{K_3^3} = f_{23}^{(3)} + f_{33}^{(3)} + f_{43}^{(3)} = f_{23}^{(1)} + f_{33}^{(1)} + f_{43}^{(1)} = \mathbf{K_3^1}$$

At the fourth and last stage it can be deduced from the previous simplifications that $u_2^{(4)} = u_2^{(2)}$ and that $u_3^{(4)} = u_3^{(2)}$. It follows that $f_{i,j}^{(4)} = f_{i,j}^{(2)}$ for i and j belonging to the set $[2, 3, 4]$. The unique simplification that can be performed at this level is thus: $K_3^4 = K_3^2$. The intermediate times, $\tilde{t}_i^4$, are, logically, all equal to $t^n + \Delta t_*$ at this last stage.

$$u_0^{(4)} = u_0^n + \tfrac{1}{4}\Delta t_*(K_0^1 + K_0^2 + 2K_0^3), \qquad u_1^{(4)} = u_1^n + \Delta t_* K_1^3, \qquad \mathbf{u_2^{(4)}} = u_2^n + \Delta t_* K_2^3 = u_2^n + \Delta t_* K_2^1 = \mathbf{u_2^{(2)}}, \qquad \mathbf{u_3^{(4)}} = u_3^n + \Delta t_* K_3^3 = u_3^n + \Delta t_* K_3^1 = \mathbf{u_3^{(2)}}$$

$$K_0^4 = f_{-10}^{(4)} + f_{00}^{(4)} + f_{10}^{(4)}, \quad K_1^4 = f_{01}^{(4)} + f_{11}^{(4)} + f_{21}^{(4)} = f_{01}^{(4)} + f_{11}^{(4)} + f_{21}^{(4)}, \quad K_2^4 = f_{12}^{(4)} + f_{22}^{(4)} + f_{32}^{(4)} = f_{12}^{(4)} + f_{22}^{(2)} + f_{32}^{(2)}, \quad \mathbf{K_3^4} = f_{23}^{(4)} + f_{33}^{(4)} + f_{43}^{(4)} = f_{23}^{(2)} + f_{33}^{(2)} + f_{43}^{(2)} = \mathbf{K_3^2}$$

The final operation is the update where the next step solution u^{n+1} is computed by using equation (6). From the above simplifications, highlighted in bold characters, it follows that:

$$u_0^{n+1} = u_0^n + \frac{1}{4}\Delta t_*(K_0^1 + K_0^2 + K_0^3 + K_0^4), \tag{10}$$

$$u_1^{n+1} = u_1^n + \frac{1}{4}\Delta t_*(K_1^1 + K_1^2 + K_1^3 + K_1^4), \tag{11}$$

$$u_2^{n+1} = u_2^n + \frac{1}{4}\Delta t_*(2K_2^1 + K_2^2 + K_2^4), \tag{12}$$

$$\mathbf{u_3^{n+1}} = u_3^n + \frac{1}{4}\Delta t_*(2K_3^1 + 2K_3^2) = \mathbf{u_3^n} + \frac{\mathbf{1}}{\mathbf{2}}\mathbf{\Delta t_*}(\mathbf{K_3^1} + \mathbf{K_3^2}) \tag{13}$$

DOI: 10.1002/fld

Equation (13) is thus equivalent to equation (6) for the base method $\mathcal{RK}^b[\Delta t_*]$ in Ω_3. This means that the 4-stage adapted method collapses into the original 2-stage base method RK2a if and only if the particular element is at a minimum distance of 2 connected elements from Ω_0. In other words: applying RK2a with $\frac{1}{2}\Delta t_*$ on Ω_0 only has an influence on the integration scheme used on the two next elements, i.e., Ω_1 and Ω_2. Therefore a buffer region of size 2 is needed between the two bulk groups that are integrated with $\mathcal{RK}^b[\Delta t_*]$ and $\mathcal{RK}^{2b}[\Delta t_*]$ respectively. Elements Ω_1 and Ω_2 are stable for Δt_* but require twice more computations than Ω_3. Despite that, this multirate approach requires less computation than the classical singlerate method.

About the conservation of the fluxes at the interfaces between elements, we can check that $f^{(k)}_{i-1,i} = -f^{(k)}_{i,i-1}$ for $i = 1, 2, 3$ and $k = 1, 2, 3, 4$. Since the b vectors of the Butcher tableau are equal, $b^a = b^{2b}$, for all elements, the sum of the fluxes cancels at each interface:

$$\sum_{k=1}^{4} b^a_k f^{(k)}_{i-1,i} + \sum_{k=1}^{4} b^{2b}_k f^{(k)}_{i,i-1} \quad = \quad \sum_{k=1}^{4} \frac{1}{4}(f^{(k)}_{i-1,i} - f^{(k)}_{i-1,i}) = 0, \quad \text{for } i = 1, 2, 3 \tag{14}$$

The so called first and second order conditions are verified for the two methods considered separately [14]. At the critical interface between Ω_0 and Ω_1 the order of the coupling between $\mathcal{RK}^a$ and $\mathcal{RK}^{2b}$ has to be considered. The first order coupling conditions are implicitly satisfied. It can be verified that the second order PRK coupling conditions, i.e.,

$$\sum_{k=1}^{4} b^{2b}_k c^a_k = \frac{1}{2}, \qquad \sum_{k=1}^{4} b^a_k c^{2b}_k = \frac{1}{2} \tag{15}$$

are satisfied [30]. The RK2a multirate method of Constantinescu is thus globally second order accurate. Indeed, for PRK methods, the global order is defined as the minimum among the orders of the two methods considered separately and the order of their coupling [30].

3.1.2. Generalization The strategy of Constantinescu [14] may be used to manage different integer time step ratios. A time step ratio $\kappa = 2$ between the different multirate groups seems to be sufficient for our target applications. Stable time steps of two neighboring cells are assumed to be relatively close for the vast majority of the mesh elements. This multirate approach may be extended, not only to any s-stage ERK base method, as shown in Table III , but also to multiple levels of refinement. But for an s-stage base method, a buffer region of at least s connected elements is required between two bulk groups. It is only at that distance that it is possible to collapse the adapted method into the base method. This general property can be proved using the same arguments as in the above introductory example. Imbricated multirate groups for buffers of size 2, 3 and 4 are illustrated around the Holbourne island in Figures 13(a), 13(b) and 13(c).

In the DGM formulation, elements are connected through their interfaces (nodes in 1D, segments in 2D and faces in 3D). In the context of multirate methods, this is a major advantage compared to the standard continuous Finite Element Method (FEM) where elements are connected through nodes. Indeed, buffer regions are generally larger and thus more elements, compared to the discontinuous case, need twice as many operations as required by their stable time step. Hence, the efficiency of the multirate methods is lower. Another issue to take into account is the handling of the mass matrix

DOI: 10.1002/fld

that is not block diagonal and thus couples the whole solution. This would complicate the use of several time steps on different multirate groups. However, we did not investigate in practice the multirate approach for continuous finite elements.

c	A
	b^T

(a) $\mathcal{RK}^b$

c	A	
c		A
	$\frac{1}{2}b^T$	$\frac{1}{2}b^T$

(b) $\mathcal{RK}^a$

$\frac{1}{2}c$	$\frac{1}{2}A$	
$\frac{1}{2}\mathbb{1} + \frac{1}{2}c$	$\frac{1}{2}\mathbb{1}b^T$	$\frac{1}{2}A$
	$\frac{1}{2}b^T$	$\frac{1}{2}b^T$

(c) $\mathcal{RK}^{2b}$

Table III. Butcher tableaus for (a) the arbitrary s-stage ERK base method, (b) the adapted buffer method and (c) the base method with half of the time step applied twice successively.

Since this multirate strategy is based on the PRK method, the order of the coupled method can be obtained as the minimum among the base methods used and the order of their coupling [30, 31]. Constantinescu [14] has shown that the MPRK-2 schemes, defined by the Butcher tableaus in Table III, *are (1) second order accurate if the base method is at least second order accurate* and *have (2) at most a second order accurate coupling regardless of the order of the base method*. The third order coupling conditions are never all satisfied for this multirate strategy [14]. It is actually at each critical interface, between a buffer group and a more constrained bulk group, that the coupling reduces to second order accuracy.

In spite of the order restrictions, the MPRK-2 schemes present the advantage of being conservative. It is shown in [14] that *any partitioned Runge-Kutta method with the same weights ($b^a = b^{2b}$) is conservative. In particular MPRK-2 (described by Table III) is conservative.*

Multiple levels of refinements can be defined recursively on nested multirate groups. The base method, $\mathcal{RK}^b$, and the associated adapted method, $\mathcal{RK}^a$, are applied successively to the different buffer and bulk groups. Consider an arbitrary problem with a succession of multiple multirate groups. Bulk groups, Ω_z^b, are integrated with a $\Delta t_*/2^z$ stable time step as well as their neighboring buffer groups, Ω_z^a. This procedure is illustrated in Figure 4 for a general case. The overall speedup that this technique would yield compared to a classical singlerate ERK method strongly depends on the amount of elements that are allocated to each multirate group.

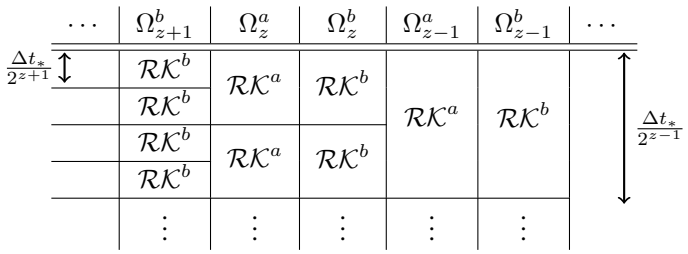

Figure 4. Multiple levels of refinement. The multirate exponent of a group is z.

DOI: 10.1002/fld

3.2. *Recursive Flux Splitting Multirate*

Knoth and Wolke [17] developed implicit-explicit (IMEX) integration methods in the context of advection-diffusion equations in air pollution. An efficient solution is expected by splitting the right-hand side of the differential equation (16) in a non-stiff advection part, the $\mathcal{G}$ term, and a stiff diffusion part, the $\mathcal{F}$ term:

$$\frac{du}{dt} = \mathcal{F}(u) + \mathcal{G}(u), \qquad u(0) = u_0, \tag{16}$$

which are solved with an explicit and an implicit method respectively. The IMEX method should ensure that the cumulative integration interval for $\mathcal{F}$ equals the explicit time step used for $\mathcal{G}$. The key idea of Schlegel et al. [15] is to consider that $\mathcal{F}$ is non-stiff like $\mathcal{G}$ but is restricted by a smaller time step, then solve them together with an inner method for $\mathcal{F}$ and an outer method for $\mathcal{G}$ that are both ERK methods. The same base method can either be used for the two parts or two different schemes ca be mixed. These choices strongly depend on the stability requirements with respect to $\mathcal{F}$. An imbricated system with $s \times q$ stages is obtained by combining a s-stage outer method $\mathcal{RK}^O$ with a q-stage inner method $\mathcal{RK}^I$. For a complete explanation about the construction of the new Butcher tableaus, see [15]. The resulting method is called the Recursive Flux Splitting Multirate (RFSMR) and may be written in a PRK form:

$$u^{n+1} = u^n + \Delta t_* \sum_{k=1}^{s\times q} b_k^{\mathcal{F}} \mathcal{F}(u^{(k)}, t^n + c_k^{\mathcal{F}} \Delta t_*) + \Delta t_* \sum_{k=1}^{s\times q} b_k^{\mathcal{G}} \mathcal{G}(u^{(k)}, t^n + c_k^{\mathcal{G}} \Delta t_*), \tag{17}$$

$$u^{(k)} = u^n + \Delta t_* \sum_{l=1}^{k-1} A_{kl}^{\mathcal{F}} \mathcal{F}(u^{(l)}, t^n + c_l^{\mathcal{F}} \Delta t_*) + \Delta t_* \sum_{l=1}^{k-1} A_{kl}^{\mathcal{G}} \mathcal{G}(u^{(l)}, t^n + c_l^{\mathcal{G}} \Delta t_*) \tag{18}$$

for $k = 1, ..., s$ and where $A^{\mathcal{F}}, b^{\mathcal{F}}, c^{\mathcal{F}}$ and $A^{\mathcal{G}}, b^{\mathcal{G}}, c^{\mathcal{G}}$ are the Runge-Kutta parameters of the resulting method. They are obtained by combinations of the original inner and outer methods parameters: A^O, b^O, c^O and A^I, b^I, c^I.

Order conditions can be established for these mixed schemes. They consist in both the classic order conditions for the base ERK methods and additional coupling conditions [31]. It can be shown that the resulting methods for $\mathcal{G}$ and $\mathcal{F}$ as well as their coupling are second order accurate if and only if the underlying base method is at least second order accurate. Knoth and Wolke [17] have derived an additional third order consistency condition that, when satisfied by the base method, leads to a third order accurate multirate scheme experimentally. Yet, the theoretical proof of this property remains an open question. In particular, the RK43 scheme represented in Table IV, used as inner and outer method, fulfills this condition and leads to a third order accurate multirate scheme. The

0				
1/2	1/2			
1/2	-1/6	2/3		
1	1/3	-1/3	1	
	1/6	1/3	1/3	1/6

Table IV. RK43 Butcher tableau

two resulting schemes, $\mathcal{RK}^{\mathcal{G}}$ and $\mathcal{RK}^{\mathcal{F}}$, that have both $s^2 = 16$ stages, may be constructed [15].

DOI: 10.1002/fld

The 10-stages methods, represented in Tables V and VI, are obtained by eliminating redundant rows and columns in the resulting Butcher tableaus. More explanations about the RFSMR method and its properties can be found in [15]. Our analysis is restricted to the interpretation of the resulting schemes and their effective application in a multirate approach.

0										
1/4	1/4									
1/4	1/4	0								
1/2	1/2	0	0							
1/2	**1/2**	0	0	0						
1/2	**-1/6**	0	0	0	**2/3**					
3/4	1/12	0	0	0	1/6	1/2				
3/4	1/12	0	0	0	1/6	1/2	0			
1	1/3	0	0	0	-1/3	1	0	0		
1	**1/3**	0	0	0	**-1/3**	**1**	0	0	0	
	1/6	0	0	0	**1/3**	**1/3**	0	0	0	**1/6**

Table V. $\mathcal{RK}^{\mathcal{G}}$ - Outer Buffer

0										
1/4	**1/4**									
1/4	**-1/12**	**1/3**								
1/2	**1/6**	**-1/6**	**1/2**							
1/2	1/12	1/6	1/6	1/12						
1/2	**1/12**	**1/6**	**1/6**	**1/12**	0					
3/4	**1/12**	**1/6**	**1/6**	**1/12**	0	**1/4**				
3/4	**1/12**	**1/6**	**1/6**	**1/12**	0	**-1/12**	**1/3**			
1	**1/12**	**1/6**	**1/6**	**1/12**	0	**1/6**	**-1/6**	**1/2**		
1	1/12	1/6	1/6	1/12	0	1/12	1/6	1/6	1/12	
	1/12	**1/6**	**1/6**	**1/12**	0	**1/12**	**1/6**	**1/6**	**1/12**	0

Table VI. $\mathcal{RK}^{\mathcal{F}}$ - Inner Buffer

The $\mathcal{RK}^{\mathcal{G}}$ (resp. $\mathcal{RK}^{\mathcal{F}}$) method is strictly equivalent to the base method RK43 applied once with a time step Δt_* (resp. twice successively with a time step $\Delta t_*/2$), if only this method is used for all the domain variables. Indeed, if the bold entries of table V (resp. VI) are gathered, by eliminating rows and columns that are redundant when the method is considered independently, the Butcher tableau of method RK43 (resp. $2\times$ RK43 with half of the time step) is obtained. These methods are used in an inner and outer buffer group that accommodate the transition between two bulk groups that have a time step ratio of 2.

The critical interface is now located between the inner and the outer buffer group. Since $c^{\mathcal{G}} = c^{\mathcal{F}}$, the solutions at each inner stage of the method are all computed at the same intermediate time steps $\tilde{t}_i^k$. This is not sufficient to draw a conclusion about the order of the coupled method but it simplifies a lot of the third order conditions that have to be satisfied. A major drawback of the method is that there is no conservation of the fluxes at the critical interface since $b^{\mathcal{G}} \neq b^{\mathcal{F}}$.

The buffer groups have a different meaning compared to the multirate approach of Constantinescu [14]. The total buffer has again a size of two connected elements, but it isn't necessary that it separates two bulk groups. An inner buffer either separates a bulk group and an outer buffer group or two outer buffer groups that have a different stable time step. This property can easily be verified by developing the Butcher tableaus presented above. Construction of appropriate multirate groups

DOI: 10.1002/fld

for the approach of Schlegel et al. [15] will be detailed in the next section. Figure 13(d) illustrates the different multirate groups for the method of Schlegel around the Holbourne island.

The Butcher tableaus have 10 stages in the buffer regions. This is more than the 8 stages needed when the base method RK43 is applied twice successively. However, Table V indicates that only K^1, K^5, K^6 and K^{10} have to be computed while in Table VI K^5 and K^{10} are superfluous. Indeed, some columns in the corresponding Butcher tableaus are equal to zero everywhere. This has to be taken into account when implementing this method but the 10 intermediate solutions $u^{(k)}$ are effectively needed to ensure a coherent transition between the multirate groups.

To have an idea about how the two buffers communicate with their respective bulk neighbors, we just need to copy Tables V and VI and replace all the non-bold entries by zeros. The two resulting Butcher tableaus are then respectively equivalent to RK43 applied with a Δt_* time step and RK43 applied twice successively with a $\Delta t_*/2$ time step.

4. MULTIRATE GROUPS

The key idea in order to achieve the best speedup is to take advantage of the multiple levels of refinements. But the speedup that can be reached with such multirate strategies strongly depends on the distribution of the characteristic stable time steps among the elements of the mesh. In particular, the gap between the minimum and the maximum stable time steps as well as the amount of elements present in each multirate group has a significant influence on the computational efficiency of the methods. Therefore the mesh elements are to be organized in an optimized way.

The major difficulty, when implementing multirate methods, is to manage the different groups and the communication between them. In the 2D-DGM framework, we propose to gather elements in groups that share the same multirate characteristics and than treat them one by one. Inside each group of elements 3 types of groups of interfaces are distinguished: (i) interfaces that are common to elements of the same element group, (ii) interfaces that are common to 2 different element groups and (iii) interfaces that are part of a physical boundary. The multirate groups communicate through the interface groups of type (ii).

A generic way to construct these multirate groups is developed in section 4.1. Afterwards, section 4.2,two efficiency issues are addressed in: the influence of the reference time on the speedup and the duplicate computations of interface residuals.

4.1. Construction of multirate groups

Consider that the stable time steps may be computed for each element of a mesh. The minimum and maximum stable time steps are noted Δt_m and Δt_M. A reference time step $\Delta t_* < \Delta t_M$ is fixed. By using a time step ratio of 2, the different ranges to which the elements have to belong may be defined recursively. The maximum multirate exponent z^* is determined as follows:

$$z^* = \left\lceil \log_2 \frac{\Delta t_*}{\Delta t_m} \right\rceil \tag{19}$$

DOI: 10.1002/fld

such that $\frac{\Delta t_*}{2^N} \leq \Delta t_m$. Elements may now be sorted in $z^* + 1$ groups. Indeed, the stable time steps of each element in the mesh belongs to one of the following sets:

$$\Omega_{3(z^*-z)} = \begin{cases} \left[\Delta t_*, \Delta t_M\right] & \text{if } z = 0 \\ \left[\frac{\Delta t_*}{2^{z+1}}, \frac{\Delta t_*}{2^z}\right[& \text{if } z = 1, ..., z^* \end{cases} \tag{20}$$

where z stands for the multirate exponent of the group. Since buffer groups have to be inserted, a tag θ is attributed to each multirate group Ω_θ:

$$\theta = 3(z^* - z) + \sigma, \qquad \sigma = 0, 1, 2 \tag{21}$$

where the integer σ defines whether the group is characterized as a bulk, $\sigma = 0$, an inner buffer, $\sigma = 1$, or an outer buffer, $\sigma = 2$. This is a general notation that is adapted to manage the two multirate strategies. For an s-stage MPRK-2 method, the inner buffer groups are always empty sets and the outer buffer groups have a size s. Note that two bulk groups that are integrated with two different time steps never have neighboring elements. The building procedure is illustrated step by step, for the two multirate approaches, on a simple mesh represented by Figure 5(a). For the method of Constantinescu, the illustration is limited to the multirate scheme which uses a 2-stage base method but it can be extended to a buffer of any size.

The first step, common to both methods, is to assign a bulk tag, defined by (21), to each element depending on its characteristic time step. Buffer groups are neglected at this level. As illustrated in Figure 5(a) there are 4 initial groups: Ω_0, Ω_3, Ω_6 and Ω_9. The transition between them is then ensured by introducing the buffer groups.

The procedure is quite simple for the method of Constantinescu. Since there are no inner buffer elements, the tags are either equal to $3(z^* - z)$ or $3(z^* - z) + 2$. The buffers have a size of two connected elements, since the base method, RK2aC, has 2 stages. The group with the smallest multirate tag, Ω_0, remains the same and then successively the buffer elements are introduced. At the same time it is ensured that two neighboring elements have neighboring tags. In other words the tags of the elements are smoothed, i.e., two connected elements have either the same tag, if they are both members of the same buffer or bulk group, or two successive tags. Figure 5(b) shows the distribution of the mesh elements in their respective multirate groups.

The smoothing procedure is more complex for the method of Schlegel because two types of buffers are introduced. It is divided in two steps. The first one, shown in Figure 5(c), introduces the outer buffer elements. The technique is identical than for the previous method but with an outer buffer of size 1. Afterwards, as illustrated by Figure 5(d), inner buffer elements are introduced by changing the tags of the elements that are still in the current bulk groups but that have an interface in common with an upstream outer buffer group. Bulk group may be empty for the multirate method of Schlegel.

As expected, introducing buffer groups has a significant influence on the repartition of the elements. Indeed, many element are attributed to groups that have a smaller time step than prescribed

DOI: 10.1002/fld

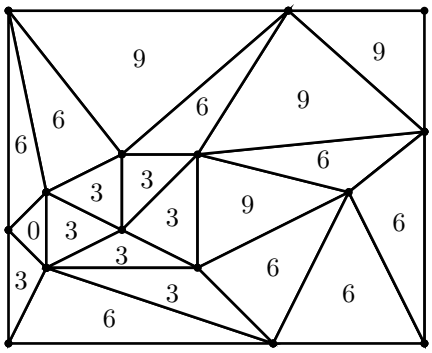

(a) First step (Common) - Assign bulk tags: $\theta = 3(z^* - z)$.

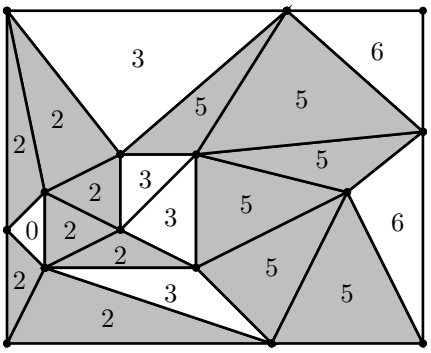

(b) Second step (Constantinescu) - Introduce buffer groups of size 2: $\theta = 3(z^* - z) + 2$, dark gray elements.

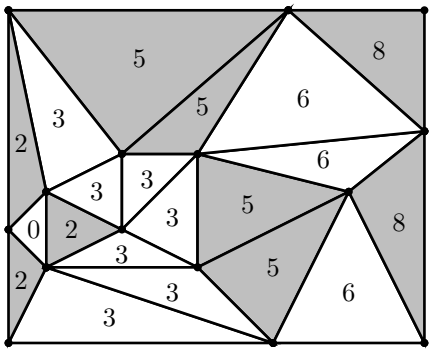

(c) Second step (Schlegel) - Introduce outer buffer groups of size 1: $\theta = 3(z^* - z) + 2$, dark gray elements.

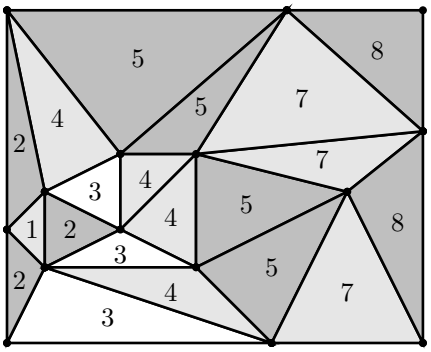

(d) Third step (Schlegel) - Introduce inner buffer groups of size 1: $\theta = 3(z^* - z) + 1$, light gray elements.

Figure 5. Construction of multirate groups for discontinuous elements. Inner buffer groups are empty for the method of Constantinescu (b). Inner buffer groups may recover a whole bulk group for the method of Schlegel, i.e., tags 1 and 7 (d).

a priori. Initially present multirate tags may disappear due to the inserted buffer groups and this leads to a loss in efficiency. In the example of Figure 5 no elements remains in group Ω_9. Nevertheless, it is a necessary condition to construct a coherent multirate scheme. If continuous elements are used for the spatial discretization the efficiency would be worse. As shown in Figures 6(a) and 6(b), the impact of introducing buffer groups is much more severe than in the discontinuous case. Elements are not anymore connected by faces but by nodes and therefore the size of the buffer regions drastically increases.

By looking to the algorithm that constructs the groups it seems difficult to determine *a priori* the effective distribution of the elements in multirate groups. Therefore it seems inevitable to build the groups in order to compute the theoretical speedup. Furthermore, it will be shown in the next section that the choice of the reference time step Δt_* has a significant influence on the repartition of the elements and therefore on the theoretical speedup.

4.2. *Efficiency*

4.2.1. Choice of the reference time step In practice, the effective speedup of multirate versus singlerate is determined by taking the ratio of the two corresponding CPU times. However, it is

Prepared using fldauth.cls
DOI: 10.1002/fld

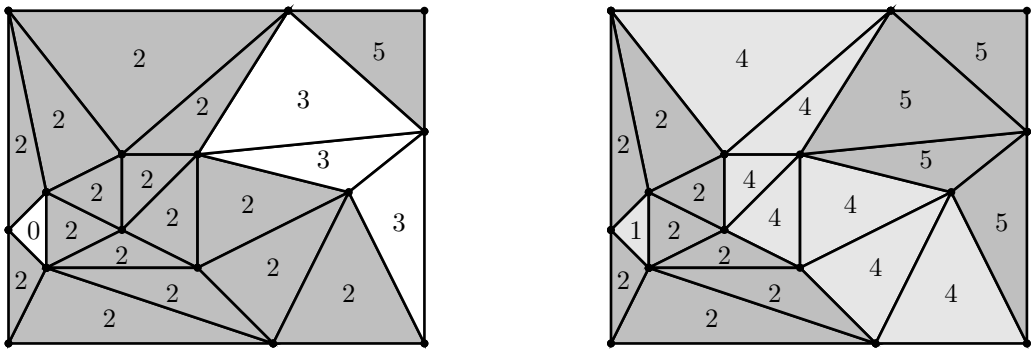

(a) Method of Constantinescu. (b) Method of Schlegel

Figure 6. Construction of the multirate groups for continuous elements. The buffer elements a drastically larger part of the domain than in the discontinuous case.

worth to have an *a priori* estimation about the theoretical speedup that could be achieved. The work performed at each stage of an ERK method for each individual element may be approximated as constant. The total work is then determined as the sum of the number of elements in each multirate group times the number of ERK stages that have to be performed in this group. The theoretical speedup, S_{th}, of a multirate method versus its singlerate equivalent can then be expressed as the ratio of their respective work:

$$S_{th} = \frac{\Delta t_m^{-1} s N}{\Delta t_*^{-1} s \sum_{i=1}^{N} 2^{\gamma(i)}} = \frac{\Delta t_* N}{\Delta t_m \sum_{i=1}^{N} 2^{\gamma(i)}} \tag{22}$$

where N stands for the total number of elements in the mesh. The discrete repartition function $\gamma(i)$ defines the effective multirate exponent associated to element i. This means that $2^{\gamma(i)s}$ stages have to be performed on element i in order to achieve a Δt_* time step. The function $\gamma(i)$ does not only depend on the stable time step of the element but also on its corresponding multirate group. This is because, in buffer regions, more stages have to be performed than prescribed *a priori* by their effective stable time steps.

Assume a fixed mesh with specific Courant numbers per element that do not vary in time. Two factors may then influence the theoretical multirate speedup (22): the reference multirate time step Δt_* and the function $\gamma(e)$. But for each Δt_* there exists an optimal configuration of the multirate groups. Actually the repartition function γ is entirely dependent on Δt_*. For a given Δt_* the number of multirate groups may vary as well as their organization and the number of elements present in each of them. Variations of Δt_* may therefore have a significant influence on the theoretical speedup. The analysis can be reduced to a fixed range of Δt_*.

Firstly, all Δt_* such that $2\Delta t_* < \Delta t_M$ have to be proscribed. In this situation, $S_{th}(2\Delta t_*) > S_{th}(\Delta t_*)$ because the reference time step is twice larger and that γ only differs for the elements with the largest characteristic time steps where it has a larger value. The reference time step Δt_* has to be chosen in a range defined as follows:

$$\Delta t_* = \max_z(\alpha 2^z \Delta t_m), \quad \text{such that} \quad 2^z \alpha \Delta t_m \leq \Delta t_M \tag{23}$$

DOI: 10.1002/fld

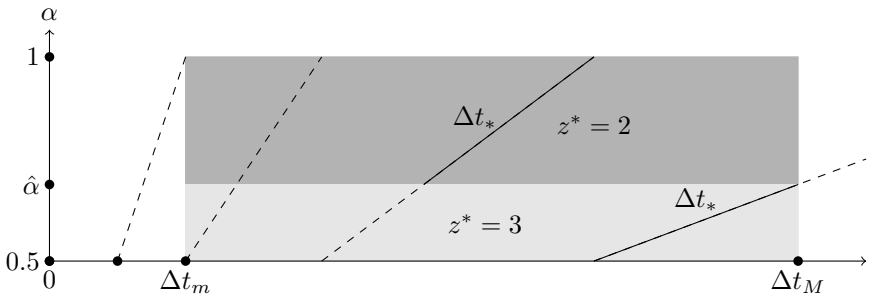

Figure 7. Influence of α on Δt_* and the number of multirate groups. For a certain value $\hat{\alpha}$ an additional multirate group appears.

where α is a factor that determines the time step, $\alpha\Delta t_m$, used for the multirate groups with the largest multirate exponent. When $\alpha = 1$ this group is simply integrated with Δt_m that is the maximum time step authorized for global stability. For $\alpha \leq \frac{1}{2}$ the multirate groups are exactly the same than for an α belonging to the set $]\frac{1}{2}, 1]$ and therefore $]\frac{\alpha}{2}\Delta t_m, \alpha\Delta t_m[$ is an empty set. The objective is to determine the maximum of equation (22) with the constraint defined by (23). This can be summarized as:

$$\max_{\alpha} S_{th}(\alpha), \quad \tfrac{1}{2} < \alpha \leq 1 \tag{24}$$

Multirate groups are generally almost the same for very small variations of α but it seems very difficult to predict the maximum of the objective function $S_{th}(\alpha)$. In practice, the theoretical speedup can be computed for each α and optimization techniques can be used to determine the maximum α^*. Figure 12 shows the evolution of the reference time step Δt_* and the corresponding multirate subdivision in function of α for an arbitrary Δt_m and Δt_M. For a certain $\hat{\alpha}$, the reference time step, Δt_*, jumps to a new curve starting at Δt_M. Consequently a new multirate subdivision appears and z^* jumps from 2 to 3, which means that for $\alpha > \hat{\alpha}$ (resp. $\alpha < \hat{\alpha}$) there will be 3 (resp. 4) levels of refinement. This phenomenon partly shows the complexity of predicting the theoretical speedup depending on the parameter α.

4.2.2. Avoid duplicate computations Two types of operations are performed during ERK time stepping: summing up vectors when the current solution is computed at an inner stage of a method, equation (4), or when the next step solution is updated, equation (6), and computing the steady-state residuals, equation (5). Almost all the computational efforts are contained in the operations of the second type and can be split into an interface and a volume term. It was shown in the previous sections that the effective computational gain, when using multirate methods, relies on the amount of computations that are avoided compared to a singlerate method. In the DGM formulation, the interface contributions of the steady-state residual have to be computed only once at each stage of the method. Indeed for two neighboring elements the interface fluxes are opposite each other. Ideally, the interface terms of type (ii), at the boundary between two multirate groups, should only be computed once. However, this is not simple to implement since each group is treated separately and runs with a different time step. In our implementation, this superfluous computation of interface

DOI: 10.1002/fld

terms is avoided for the methods of Constantinescu but not for the method of Schlegel. Consequently the effective speedup will be lower than the theoretical one for the method of Schlegel.

5. NUMERICAL EXPERIMENTS AND RESULTS

The two multirate ERK methods are applied for the temporal integration in the framework of ocean modeling. A depth-averaged barotropic 2D model is used to compute the mean horizontal velocity vector $\mathbf{u}$ and the free-surface elevation η for shallow waters. Consider the non-conservative shallow water equations:

$$\frac{\partial \eta}{\partial t} + \boldsymbol{\nabla} \cdot \big((h+\eta)\mathbf{u}\big) = 0, \tag{25}$$

$$\frac{\partial \mathbf{u}}{\partial t} + \mathbf{u} \cdot (\boldsymbol{\nabla} \boldsymbol{u}) + f\mathbf{k} \times \mathbf{u} + g\boldsymbol{\nabla}\eta = \frac{1}{H}\boldsymbol{\nabla} \cdot \big(H\nu(\boldsymbol{\nabla}\boldsymbol{u})\big) + \frac{\boldsymbol{\tau}^s - \boldsymbol{\tau}^b}{\rho H}, \tag{26}$$

where f, g, ν and ρ are respectively the Coriolis parameter, the gravitational acceleration, the horizontal eddy viscosity and the mean water density. The actual water depth is $H = h + \eta$, where h is the reference water depth below the mean sea level. The bottom and wind stresses are parametrized as $\boldsymbol{\tau^b}$ and $\boldsymbol{\tau^s}$ respectively. The equations are discretized with P_1^{DG} discontinuous finite elements for both elevation and velocity fields. Three methods of Constantinescu, RK2aC, RK33C and RK44C, based on the corresponding base methods, and the method of Schlegel RK43S are compared. The convergence and the performance of the methods are first analyzed on a simple shallow water test case. The step-by-step procedure of the different multirate strategies is then illustrated on a more realistic application, the GBR.

5.1. Convergence and performance of the multirate methods: an island in a rectangular basin

We introduce a simple shallow water test case in order to compare the different multirate methods in terms of efficiency, convergence and conservation properties. Consider the water circulation in a rectangular closed basin defined on the domain $[-W, W] \times [-L, L]$ where $W = 350\ m$ and $L = 75\ m$ with an elliptic island in the middle. Figure 8(a) represents this domain with an assiocated bathymetry that varies between $10\ m$ and $5\ m$ (around the island). The bottom stress is a quadratic dissipation term that depends on the bathymetry. A Coriolis force is also acting on the system. The initial condition corresponds to an exponential elevation where $0\ m \leq\ \eta(x, y)\ \leq 0.05\ m$ as represented on Figure 8(b).

We use this example to compare the different time-stepping schemes through 3 experiments. The first experiment compares the efficiency of the methods by measuring the integrated L_2 errors for both elevation and velocities and the CPU time after 25 seconds of simulation on different meshes. We consider 4 meshes that are obtained by successive refinements of the original mesh of Figure 8(c). The number of mesh elements, the reference time steps and the theoretical speedups associated with each mesh and each multirate method are listed in Table VII. The maximum multirate exponent z^* is 4 for all meshes and methods. This means that the time step

DOI: 10.1002/fld

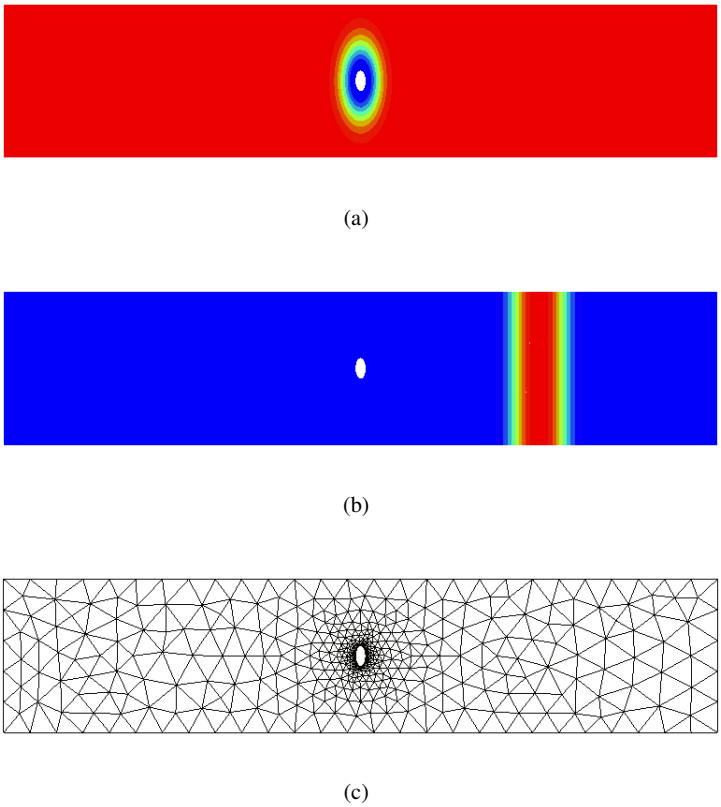

(a)

(b)

(c)

Figure 8. Illustrations related to the simple shallow water testcase: an island in a rectangular basin. (a) Bathymetry. (b) Initial Condition. (c) Coarse Mesh.

for the singlerate methods is $\Delta t_{min} = \Delta t_* / 2^4$. The stability requirements of the explicit temporal discretization limit the time step to very small values. Therefore the temporal error is much smaller than the spatial error and the total error is expected to scale as the spatial error [32]. In this case we use a P_1^{DG} discretization and we expect a second order convergence for all fields when the mesh and the corresponding time step are refined. Indeed, as represented on Figure 9, a convergence of order 2 is observed for both elevation and velocities regardless of the time-stepping method that is used. Multirate methods have thus no adverse effect on the global space-time error. Moreover, all the multirate methods give a better ratio than the singlerate ones between CPU time and error because they need less operations. Figure 9 reveals that the RK2aC method is the most efficient. This is because it yields the best effective speedup and uses only 2 stages. The RK43S method needs 4 stages and has an effective speedup that is lower than the theoretical one, for the reasons mentioned in section 4.2.

Figures 10(a) and 10(b) give the L_2 error in function of the CPU time for both elevation and velocities after 0.5 seconds of simulation. The same mesh, represented by Figure 8(c), is used for all computations but with different time steps. The original time step associated with the mesh is divided successively by a factor 2 such that the pure temporal error is visible. The expected convergence rates are observed for all time-stepping schemes. However the multirate methods produce larger temporal errors than their singlerate counterparts. For multirate methods, the error

Prepared using fldauth.cls
DOI: 10.1002/fld

	h	$h/2$	$h/4$	$h/8$
# elements	892	3568	14272	57088
$\Delta t_*[sec]$	0.1084	0.0541	0.0270	0.0135
RK2aC	2.6283	2.8370	3.0110	3.1232
RK33C	2.4189	2.7022	2.9133	3.0638
RK44C	2.2296	2.5789	2.8239	3.0072
RK43S	3.2451	3.2484	3.2542	3.2553

Table VII. Theoretical speedups ($\alpha = 1$) corresponding to the 4 multirate methods evaluated for 4 meshes obtained by successive refinements of the original mesh of Figure 8(c).

associated with the largest time step Δt_* propagates to all mesh elements after a certain time. The global temporal error is of the order of the largest time step. In our case, the largest time step of the multirate RK2aC method is 16 times bigger than the time step of the singlerate RK2a method. Both methods being of quadratic precision the error is 16^2 times larger. This may be verified on Figure 10 by comparing the 2 blue convergence curves. The third order accurate RK43S method gives the best precision for a fixed CPU time. After 3 refinements of the original time step it beats the RK2a method. All the methods of Constantinescu achieve second order accuracy but RK33C is slightly more performant than RK2aC andSecond RK44C.

Finally we compare the conservation properties of the 4 selected multirate schemes. The total water volume at a time t is computed as follows:

$$V(t) \quad = \quad \int_\Omega \eta(t, \boldsymbol{x}) + H(\boldsymbol{x}) d\boldsymbol{x} \tag{27}$$

We evaluate the conservation defect of a particular method as:

$$D(t, \text{method}) \quad = \quad \frac{V(t, \text{method}) - V(0)}{V(0)} \tag{28}$$

Figure 11 shows the conservation defects for the selected multirate methods for a simulation of 200 seconds. The experiment confirms the theory since there is a perfect conservation of the total water volume for the schemes of Constantinescu. The scheme of Schlegel does not preserve the water volume.

From those three experiments, the RK2aC scheme seems the most appropriate for application with P_1^{DG} spatial discretization. It delivers the best total speedup (lowest number of stages, minimum buffer size) and respects an important property in oceanography: mass conservation. Note that first order time-stepping schemes are too dissipative and therefore inappropriate. However, the other schemes may present some advantages for other applications e.g. when the temporal error is dominant.

5.2. Hydrodynamics of the Great Barrier Reef

Consider the unstructured mesh of the Great Barrier Reef depicted in Figure 1 on which the two-dimensional shallow water equations (25), (26) are solved. The parametrization of the equations as well as multiple details about the model and the mesh can be found in [10]. Bathymetry, wind stress

DOI: 10.1002/fld

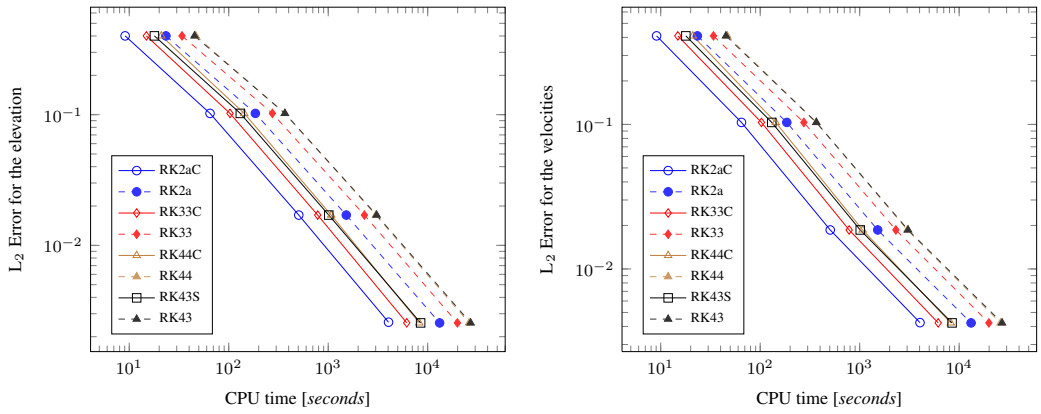

(a) (b)

Figure 9. Integrated L_2 errors for the elevation (a) and velocities (b) as a function of the CPU time for the 4 selected multirate schemes and their singlerate equivalents. The errors are computed after 25 seconds of simulation. The first mark of each curve corresponds to the reference mesh of Figure 8(c) with an element size h. The 3 next marks are associated with 3 successive refinements of the original mesh: $h/2$, $h/4$ and $h/8$. The maximum multirate exponent is 4 for all the multirate schemes applied on all meshes (with $\alpha = 1$). Second order convergence is observed for all the schemes as expected.

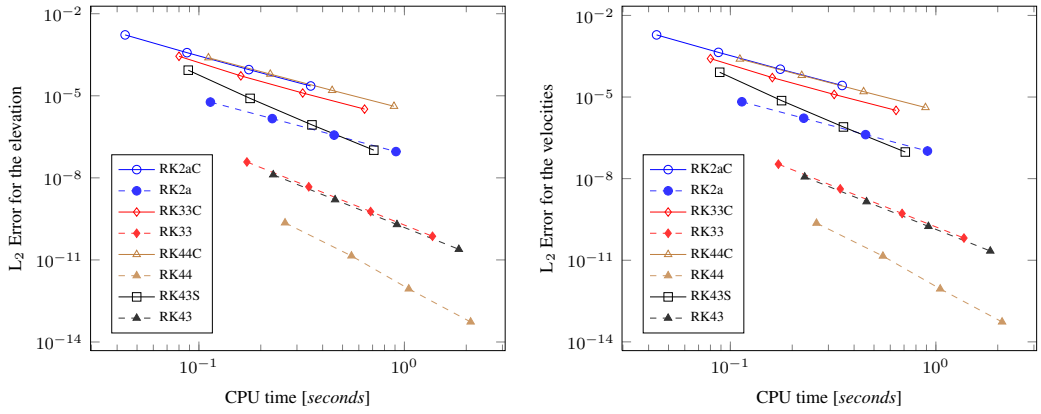

(a) (b)

Figure 10. Integrated L_2 errors for the elevation (a) and velocities (b) on the original mesh of Figure 8(c) as a function of the CPU time for the 4 selected multirate schemes and their singlerate equivalents. The errors are computed after 0.5 seconds of simulation. Each mark is associated to a ratio of the original time step: Δt_{min}, $\Delta t_{min}/2$, $\Delta t_{min}/4$ and $\Delta t_{min}/8$ for the singlerate methods and Δt_*, $\Delta t_*/2$, $\Delta t_*/4$ and $\Delta t_*/8$ for the multirate methods. The expected convergence rates are observed for the 8 schemes. Second order convergence for the 3 schemes of Constantinescu and third order for the scheme of Schlegel.

and open sea boundary conditions are obtained from terrain data or measured data. A zero mass flux and a tangential momentum proportional to the mean tangential velocity are imposed along the impermeable boundaries (coasts and islands). The parametrization of Smagorinsky [33] for the kinematic viscosity ν is used to incorporate unresolved features.

The 4 multirate approaches have been tested. Figure 12 shows the theoretical speedup depending on parameter α. As expected we have that $S_{th}(\frac{1}{2}) = S_{th}(1)$. The 3 methods of Constantinescu yield a curve of almost the same shape but a shift of the maximum is observed. Since the number of elements in the buffer groups increases with the number of stages of the method, the speedup

DOI: 10.1002/fld

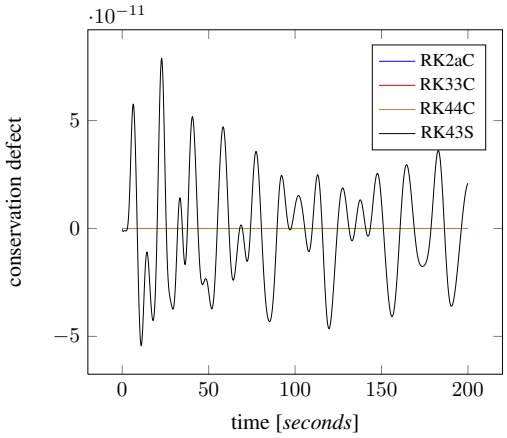

Figure 11. Conservation defects evaluated as the mean water volume per cubic meter that is added or removed from the original total water volume as a function of time. As expected the methods of Constantinescu are conservative. An oscillation of the relative mass is observed for the one of Schlegel.

declines. The method of Schlegel, RK43S, achieves a significantly higher speedup due to buffers that do not perform more expensive operations than actually needed.

The optimal values for α associated with the reference time step and the maximum multirate exponent z^* as well as the corresponding theoretical speedups are listed in Table VIII for the 4 different multirate methods. An illustration of the corresponding multirate groups is shown in Figure 13 for a zoom around the Holbourne Island. Observe that for RK2aC, RK33C and RK44C the size of the buffer groups is increasing with the number of stages of the base method. For RK43S a distinction can be made between the inner and outer buffer groups that are both of size 1. Figure 14 shows the multirate groups on the whole GBR for method RK2aC. The global percentage of inner and outer buffer elements for the whole GBR mesh is given in Table VIII. The optimal α yield a maximum multirate exponent $z^* = 6$ for the 3 methods of Constantinescu while z^* stays at 5 for the method of Schlegel.

The 4 selected methods were used to run the GBR test case and the CPU times have been measured. The same runs have also been performed with the corresponding singlerate methods where the global time step is simply the minimum among all. The experimental speedups, listed in Table VIII, are obtained by taking the ratio of the singlerate and multirate CPU times. The theoretical and experimental speedups are relatively close for the 3 methods of Constantinescu while the one of Schlegel yields a worser experimental speedup. As mentioned in section 4.2, this lack in efficiency is due to the computational overhead caused by the superfluous operations performed at the interfaces between multirate groups. Recall that the speedup strongly depends on the kind of mesh that is used. Significantly higher speedups may be obtained for the same problem with other meshes where the average element size is drastically larger than the smallest one.

The RK2aC scheme is used to perform a 24 hour simulation on the mesh presented in Figure 1 with data corresponding to the first of march 2000. A plot of the velocity vectors and the sea surface elevation is presented on Figure 15 corresponding to time 21:51:23. Tidal jets and eddies due to the interaction of the flow with the topography near the open-sea boundary are clearly visible.

DOI: 10.1002/fld

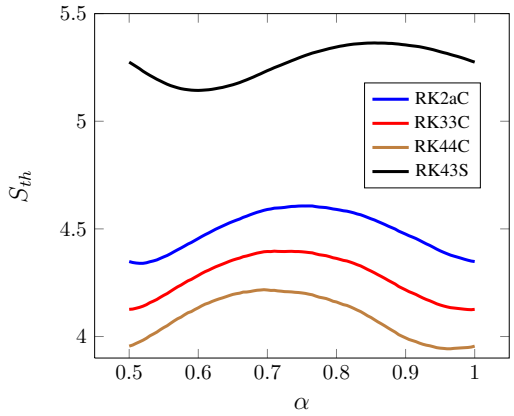

Figure 12. Theoretical speedup as a function of parameter $\alpha \in]\frac{1}{2}, 1]$ for the GBR.

Multirate Method	α^*	Δt_* [sec]	z^*	% Outer buffer	% Inner buffer	Theoretical speedup	Experimental speedup
RK2aC	0.75	7.381	6	20.0 %	0 %	4.606	4.461
RK33C	0.735	7.233	6	26.3 %	0 %	4.396	4.327
RK44C	0.695	6.840	6	31.3 %	0 %	4.218	4.183
RK43S	0.855	4.207	5	12.9 %	13.0 %	5.364	4.343

Table VIII. Comparison of the 4 selected multirate schemes.

6. CONCLUSION AND FUTURE WORK

In this paper, two ERK multirate approaches have been implemented, in the DGM framework, for solving large-scale problems with different time steps. The first strategy is conservative and reaches second order accuracy while the second one is not conservative but is third order accurate. Even if the multirate methods are more complex to implement than their singlerate equivalents, they inherit a lot of properties that makes them particularly adapted to multi-scale simulations. A significant speedup, for a well chosen reference time step, has been observed for the two kind of multirate methods on an unstructured mesh of the Great Barrier Reef. However, the speedup turns out to be highly dependent on the nature of the mesh. Furthermore other parameters, like the choice of the reference time step, have a significant impact on the speedup.

Large-scale applications such as the Great Barrier Reef require the use of parallel computers. Some kind of load balancing strategy has to be supplied to accomodate multirate schemes. Indeed, small elements have a higher cost than large elements in such a strategy and will require more frequent updates at inter-processor interfaces. The key idea consists in creating an optimized mesh partition such that the amount of grid cells of the different multirate groups is ideally the same on each computer core. However, a compromise should also be found between the effective work on each processor and the amount of communications between them.

Until now we have not considered that some parameters related to the local stability condition may change in time. The meshes could be readapted at some time steps and the multirate groups would have to be consequently changed. A more physical constraint is that the wave/advective velocity changes considerably in time and could cause the solution to blow up after a certain number of

DOI: 10.1002/fld

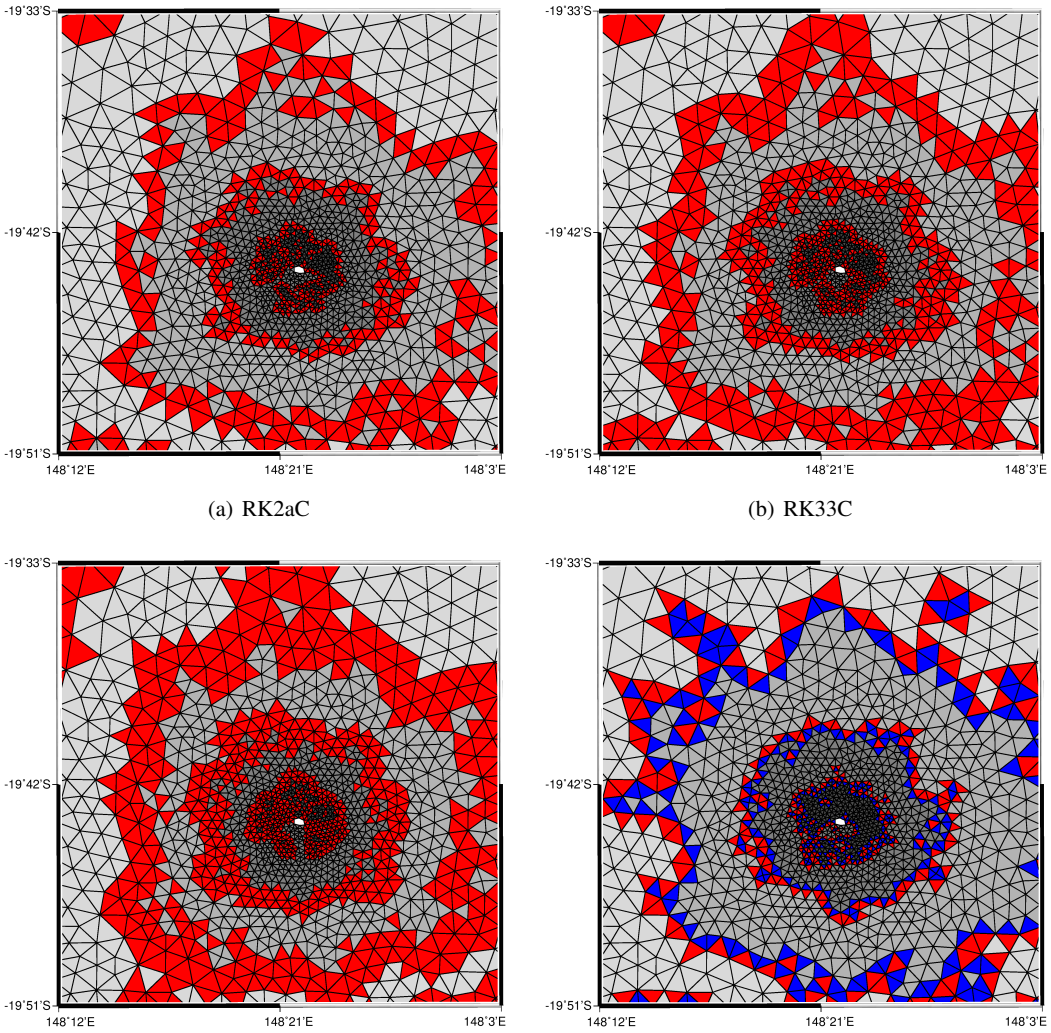

(a) RK2aC (b) RK33C

(c) RK44C (d) RK43S

Figure 13. Multirate groups around the Holbourne island for the 4 multirate strategies obtained for the optimal Δt_*. Outer buffer groups are colored in red and inner buffer groups in blue. Bulk groups have colors that vary from light gray to dark gray depending on the size of their time step.

iterations. A criterion could eventually be found to determine whether it is worth or not to recompute the multirate groups at a certain moment in order to stay stable all along the simulation.

ACKNOWLEDGEMENTS

Bruno Seny is a Research fellow with the Belgian Fund for Research in Industry and Agriculture (FRIA). The present study was carried out within the scope of the project "Taking up the challenges of multi-scale marine modeling", which is funded by the Communauté Française de Belgique, as Actions de Recherche Concertées, under contract ARC 10-15/028.

DOI: 10.1002/fld

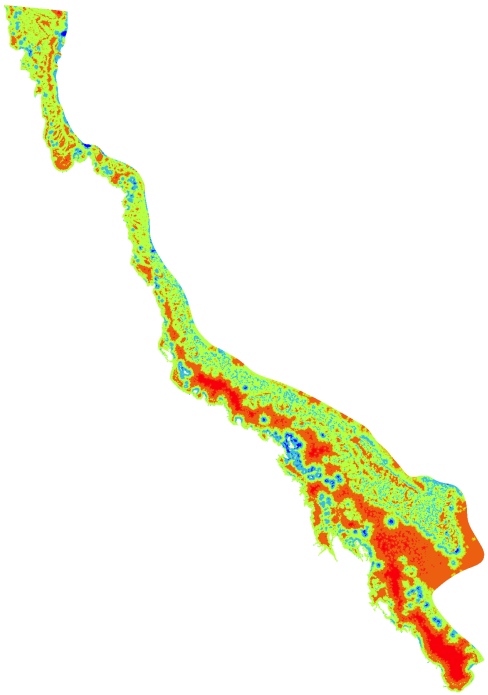

Figure 14. Multirate groups for the RK2aC method on the whole GBR. Elements have colors that depend on their multirate groups. Small (resp. large) time steps are used on blue (resp. red) elements.

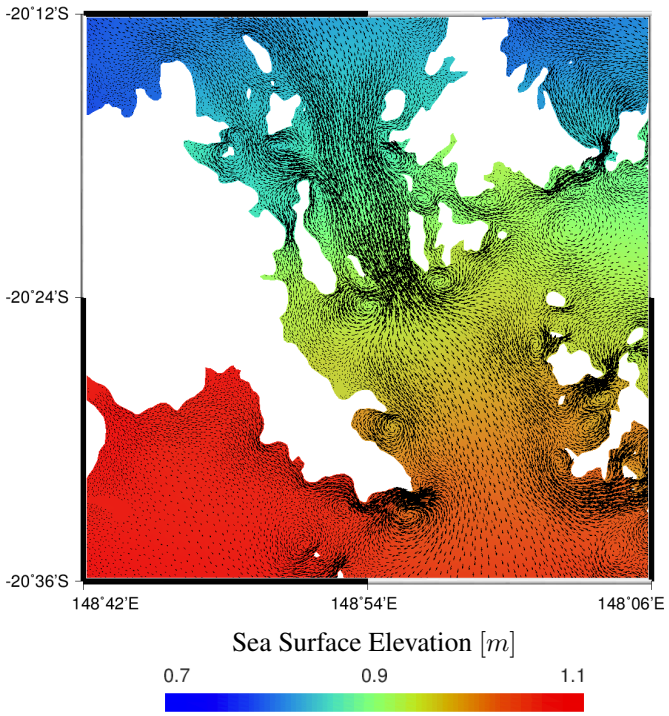

Figure 15. Sea surface elevation (color levels) and bidimensional velocity field (arrows) around the Withsunday Island Archipelago. Velocity vectors have a norm that varies between 0 and $0.822\ m/s$.

DOI: 10.1002/fld

REFERENCES

1. Chen C, Liu H, Beardsley RC. An unstructured grid, finite-volume, three-dimensional, primitive equations ocean model: Applications to coastal ocean and estuaries. *Journal of Atmospheric and Oceanic Technology* 2003; **20**:159–186.
2. Fringer OB, Gerritsen M, Street RL. An unstructured-grid, finite-volume, nonhydrostatic, parallel coastal ocean simulator. *Ocean Modelling* 2006; **14**:139–173, doi:10.1016/j.ocemod.2006.03.006.
3. Casulli V, Walters RA. An unstructured grid, three-dimensional model based on the shallow water equations. *International Journal for Numerical Methods in Fluids* 2000; **32**:331–348.
4. Wang Q, Danilov S, Schröter J. Finite element ocean circulation model based on triangular prismatic elements, with application in studying the effect of topography representation. *Journal of Geophysical Research* 2008; **113**, doi:10.1029/2007JC004482.
5. Ford R, Pain CC, Piggott M, Goddard A, de Oliveira CR, Umpleby A. A non-hydrostatic finite element model for three-dimensional stratified oceanic flows, Part I: Model formulation. *Monthly Weather Review* 2004; **132**:2816–2831, doi:10.1175/MWR2824.1.
6. Piggott MD, Gorman GJ, Pain CC, Allison PA, Candy AS, Martin BT, Wells MR. A new computational framework for multi-scale ocean modelling basid on adapting unstructured meshes. *International Journal For Numerical Methods In Fluids* 2008; **56**:1003–1015, doi:10.1002/fld.1663.
7. Luettich RA, Westerink JJ, Scheffner NW. Adcirc: an advanced three-dimensional circulation model for shelves, coasts and estuaries 1991; .
8. Lyard F, Lefevre F, Letellier T, Francis O. Modelling the global ocean tides: modern insights from FES2004. *Ocean Dynamics* 2006; **56**:394–415.
9. Comblen R, Lambrechts J, Remacle JF, Legat V. Practical evaluation of five part-discontinuous finite element pairs for the non-conservative shallow water equations. *International Journal for Numerical Methods in Fluids (in press)* 2009; doi:10.1002/fld.2094.
10. Lambrechts J, Hanert E, Deleersnijder E, Bernard PE, Legat V, Remacle JF, Wolanski E. A multiscale model of the whole Great Barrier Reef hydrodynamics. *Estuarine, Coastal and Shelf Science* 2008; **79**:143–151, doi: 10.1016/j.ecss.2008.03.016.
11. Lambrechts J, Comblen R, Legat V, Geuzaine C, Remacle JF. Multiscale mesh generation on the sphere. *Ocean Dynamics* 2008; **58**(5):461–473, doi:10.1007/s10236-008-0148-3.
12. de Brye B, de Brauwere A, Gourgue O, Kärnä T, Lambrechts J, Comblen R, Deleersnijder E. A finite-element, multi-scale model of the Scheldt tributaries, river, estuary and ROFI. *Coastal Engineering (accepted)* 2010; .
13. Geuzaine C, Remacle JF. Gmsh: a three-dimensional finite element mesh generator with buit-in pre- and post-processing facilities. *International Journal for Numerical Methods in Engineering* 2009; **79**:1309–1331, doi: 10.1002/nme.2579.
14. Constantinescu EM, Sandu A. Multirate timestepping methods for hyperbolic conservation laws. *Journal of Scientific Computing* 2007; **33**:239–278, doi:10.1007/s10915-007-9151-y.
15. Schlegel M, Knoth O, Arnold M, Wolke R. Multirate Runge-Kutta schemes for advection equations. *Journal of Computational and Applied Mathematics* 2009; **226**:345–357, doi:10.1016/j.cam.2008.08.009.
16. Ascher UM, Ruuth SJ, Wetton BT. Implicit-explicit methods for time-dependent partial differential equations. *SIAM Journal on Numerical Analysis* 1995; **32**:797–823, doi:10.1137/0732037.
17. Knoth O, Wolke R. Implicit-explicit runge-kutta methods for computing atmospheric reactiv flows. *Applied Numerical Mathematics* 1998; **28**(2-4):327–341.
18. Hairer E, Nørsett SP, Wanner G. *Solving Ordinary Differential Equations I: Nonstiff Problems*. 2 edn., Springer Series in Computational Mathematics, Springer: Berlin, 2000.
19. Hairer E, Nørsett SP, Wanner G. *Solving Ordinary Differential Equations II: Stiff and Differential-Algebraic Problems*. 2 edn., Springer Series in Computational Mathematics, Springer: Berlin, 2002.
20. Butcher J. Numerical methods for ordinary differential equations in the 20th century. *Journal of Computational and Applied Mathematics* 2000; **125**:1–29.
21. Shu CW, Osher S. Efficient implementation of essentially non-oscillatory shock-capturing schemes,ii. *J. Comput. Phys.* 1989; **83**(1):32–78, doi:http://dx.doi.org/10.1016/0021-9991(89)90222-2.
22. Gottlieb S, Shu CW, Tadmor E. Strong stability-preserving high-order time discretization methods. *SIAM Rev* 2001; **43**:89–112.
23. Kubatko EJ, Dawson C, Westerink JJ. Time step restrictions for runge-kutta discontinuous galerkin methods on triangular grids. *J. Comput. Phys.* December 2008; **227**:9697–9710.
24. Osher S, Sanders R. Numerical approximations to nonlinear conservation laws with locally varying time and space grids. *Mathematics of Computation* 1983; **41**(164):321–336.

DOI: 10.1002/fld

25. Dawson C, Kirby R. High resolution schemes for conservation laws with locally varying time steps. *SIAM J. Sci. Comput.* 2000; **22**(6):2256–2281, doi:http://dx.doi.org/10.1137/S1064827500367737.
26. Tang H, Warnecke G. A class of high resolution difference schemes for nonlinear hamilton–jacobi equations with varying time and space grids. *SIAM Journal on Scientific Computing* 2005; **26**(4):1415–1431.
27. Hundsdorfer W, Mozartova A, Savcenco V. Analysis Of Explicit Multirate And Partitioned Runge-Kutta Schemes For Conservation Laws. *Technical Report MAS-E0715* 2007.
28. Savcenco V, Hundsdorfer W, Verwer JG. A Multirate Time Stepping Strategy For Stiff ODEs. *BIT : Numerical Mathematics* January 2007; **47**(1):137 – 155.
29. Hundsdorfer W, Savcenco V. Analysis Of A Multirate Theta-Method For Stiff ODEs. *Applied Numerical Mathematics* 2009; **59**:693 – 706.
30. Hairer E. Order conditions for numerical methods for partitioned ordinary differential equations. *Numerische Mathematik* 1981; **36**:431–445.
31. Jackiewicz Z, Vermiglio R. Order conditions for partitioned runge-kutta methods. *Applications of Mathematics* 2000; **45**:301–316(16).
32. Chevaugeon N, Hillewaert K, Gallez X, Ploumhans P, Remacle JF. Optimal numerical parametrization of discontinuous Galerkin method applied to wave propagation problems. *Journal of Computational Physics* 2007; **223**:188–207.
33. Smagorinsky J. General circulation experiments with the primitive equations. *Monthly Weather Review* 1963; **91**:99–164.

DOI: 10.1002/fld