

Emmanuelle Itam · Stephen Wornom · Bruno
Koobus · Alain Dervieux

A Volume-agglomeration multirate time advancing for high Reynolds number flow simulation

Abstract

A frequent configuration in computational fluid mechanics combines an explicit time advancing scheme for accuracy purposes and a computational grid with a very small portion of much smaller elements than in the remaining mesh. Examples of such situations are the travel of a discontinuity followed by a moving mesh, and the large eddy simulation of high Reynolds number flows around bluff bodies where together very thin boundary layers and vortices of much more important size need to be captured. For such configurations, explicit time advancing schemes with global time stepping are very costly. In order to reduce this problem, the multirate time stepping approach represents an interesting improvement. The objective of such schemes, which allow to use *different time steps* in the computational domain, is to avoid penalizing the computational cost of the time advancement of unsteady solutions which can become large due to the use of small global time steps imposed by the smallest elements such as those constituting the boundary layers. In the present work, a new multirate scheme based on control volume agglomeration is proposed for the solution of the compressible Navier-Stokes equations possibly equipped with turbulence models. The method relies on a prediction step where large time steps are performed with an evaluation of the fluxes on macro-cells for the smaller elements for stability purpose, and on a correction step in which small time steps are employed only for the smaller elements. The accuracy and efficiency of the proposed method are evaluated on several benchmarks flows: the problem of a moving contact discontinuity (inviscid flow), the computation with a hybrid turbulence model of flows around bluff bodies like a tandem cylinders at Reynolds number 1.66×10^5 , a circular

E. Itam and B. Koobus
IMAG, Université de Montpellier, Place Eugène Bataillon, 34095 Montpellier cedex, France
E-mail: itamemmanuelle@gmail.com, koobus@umontpellier.fr

S. Wornom
LEMMA, 2000 Route des Lucioles, 06902 Sophia-Antipolis, France
E-mail: Stephen.Wornom@inria.fr

A. Dervieux
INRIA, 2004 Route des Lucioles, 06902 Sophia-Antipolis, France
E-mail: alain.dervieux@sophia.inria.fr

cylinder at Reynolds number 8.4×10^6 , and a flow around a space probe model at Reynolds number 10^6 .

Keywords: computational fluid dynamics, multirate time advancing, explicit scheme, volume agglomeration, unstructured grid, hybrid turbulence model.

1 INTRODUCTION

A frequent configuration in CFD calculations combines an explicit time advancing scheme for accuracy purpose and a computational grid with a very small portion of much smaller elements than in the remaining mesh. Two typical examples are the following: A first example is the hybrid RANS/LES simulation of high Reynolds number flows around bluff bodies. In that case, very thin boundary layers need be addressed with extremely small cells. When applying explicit time advancing, the computation is penalized by the very small time-step to be applied (CFL number of order 1). But this is not the only interesting region of the computational domain. An important part of the meshing effort is devoted to large regions of medium cell size in which the motion of vortices need be accurately captured. For these vortices, the efficient and accurate time-step is of order of the ratio of local mesh size by vortex velocity. We can apply an implicit scheme with such a time-step, which would produce a local CFL of order 1 for the vortices and a local CFL of order hundreds for the boundary layer. But the implicit scheme used with a too small time step will be of larger cost than the explicit one, with a cost increasing importantly with formal accuracy, think of implicit RK schemes. For higher time lapse, implicit schemes like BDF1/2 show much more dissipation than explicit schemes and are much less accurate.

Our second example concerns an important complexity issue in unsteady mesh adaptation. Indeed, unsteady mesh adaptive calculations are penalized by the very small time-step imposed by accuracy requirements on regions involving small space-time scales. This small time step is for example an important computational penalty for mesh adaptive methods of AMR type [5]. This is also the case for unsteady fixed-point mesh-adaptive methods as in [4]. Conversely, using an implicit scheme and large time steps will strongly degrade the accuracy. In that latter method, the loss of efficiency is even more crucial when the anisotropic mesh is strongly locally strongly stretched. In [4], this loss is evaluated as limiting the numerical convergence order for discontinuities to $8/5$ instead of second-order convergence.

This limitation also applies to mesh adaptation by mesh motion. Our second example will concentrate on the computation of an isolated traveling discontinuity. The discontinuity needs to be followed by the mesh, preferably in a mesh-adaptive mode. Except if the adaptation works in a purely Lagrangian mode, an implicit scheme will smear the discontinuity of the solution. An explicit scheme will applied a costly very small time step on the whole computational domain.

In order to overcome these problems, the multirate time stepping approach represents an interesting alternative. A part of the computational domain is advanced in time with the small time-step imposed by accuracy and stability constraints. Another part is advanced with the larger time-step giving a good compromise between accuracy and efficiency. Many works have been published on multirate methods in the field of ODE, see for example [24, 1, 9, 29] and [25, 2, 15, 8, 7, 14, 13, 27], but fewer works were conducted on multirate time advancing schemes for the solution of PDE and hyperbolic conservation laws [6, 26, 23, 18, 28, 21], and rare applications were performed in Computational Fluid Dynamics (CFD), for shock propa-

gation in [21] and for shallow water computations in [28]. Therefore, there is still much work to do to provide viable multirate methods for CFD applications.

In this work, we propose a new multirate scheme based on control volume agglomeration which is at the same time very simple and well suited to a large class of finite volume approximations. The agglomeration produces macro-cells by grouping together several neighboring cells of the initial mesh. The method relies on a prediction step where large time steps are used with an evaluation of the fluxes performed on the macro-cells for the region of smallest cells, and on a correction step advancing solely the region of small cells, this time with a small time step.

We demonstrate the method in a numerical framework using a vertex centered approximation, the mixed finite volume/finite element formulation.

Target applications are three-dimensional unsteady flows modeled by the compressible Navier-Stokes equations equipped with turbulence models and discretized on unstructured possibly deformable meshes. The numerical illustration involves the two above examples.

The proposed algorithm is described in Section 2. Section 3 provides some motivations of this construction. Section 4 gives several examples of applications.

2 Multirate time advancing by volume agglomeration

2.1 Finite-Volume Navier-Stokes

The multirate time advancing scheme based on volume agglomeration is developed for the solution of the three-dimensional compressible Navier-Stokes equations. The main assumption is that the computational domain is split into computational finite volume cells such that cells intersect only by their boundaries and cover the whole computational domain. The discrete Navier-Stokes system is assembled by into a flux summation Ψ_i summing convective and diffusive fluxes evaluated at all the interfaces separating two cells. More precisely, the finite-volume spatial discretization combined with an explicit forward-Euler time-advancing writes for the Navier-Stokes equations possibly equipped with a $k - \varepsilon$ model:

$$vol_i w_i^{n+1} = vol_i w_i^n + \Delta t \Psi_i, \quad \forall i = 1, \dots, ncell,$$

where vol_i is the volume of $cell_i$, Δt the time step,

$$w_i^n = (\rho_i^n, (\rho u)_i^n, (\rho v)_i^n, (\rho w)_i^n, E_i^n, (\rho k)_i^n, (\rho \varepsilon)_i^n)$$

are as usually the density, moments, total energy, turbulent energy and turbulent dissipation at $cell_i$ and time level t^n , and $ncell$ the total number of cells in the mesh.

Given an explicit -conditionally stable- time advancing, we assume that we can define a maximal stable time step (*local time step*) $\Delta t_i, i = 1, \dots, ncell$ on each node. The stable local time step is classically defined by the combination of a viscous stability limit and an advective one according to the following formula:

$$\Delta t_i \leq \frac{CFL \times \Delta l_i^2}{\Delta l(\|\mathbf{u}_i\| + c_i) + 2 \frac{\gamma}{\rho_i} \left(\frac{\mu_i}{Pr} + \frac{\mu_{t_i}}{Pr_t} \right)} \quad (1)$$

where Δl_i is a local characteristic mesh size, \mathbf{u}_i the local velocity, c_i the sound celerity, γ the ration of specific heats, ρ_i the density, $\frac{\mu_i}{Pr} + \frac{\mu_{t_i}}{Pr_t}$ the sum of local viscosity to Prandtl ratio, laminar and turbulent, and CFL a parameter depending of the time advancing scheme,

of the order of unity. Using the local time step Δt_i leads to a stable but not consistent time advancing.

A consistent and stable time advancing should use a *global/uniform time step* defined by:

$$\Delta t = \min_{1,ncell} \Delta t_i.$$

2.2 Inner and outer zones

We first define the inner zone and the outer zone, the coarse grid, and the construction of the fluxes on the coarse grid, ingredients on which our multirate time advancing scheme is based. For this splitting into two zones, the user is supposed to choose a (integer) *time step factor* $K > 1$.

– Definition of the Inner and Outer zones :

- We define the **outer zone** as the set of cells i for which the explicit scheme is stable for a time step $K \Delta t$

$$\Delta t_i \geq K \Delta t,$$

- the **inner zone** is the set of cells for which

$$\Delta t_i < K \Delta t.$$

– Definition of the coarse grid :

– Objective :

- Advancement in time is performed with time step $K \Delta t$
- Advancement in time preserves accuracy in the outer zone
- Advancement in time is consistent in the inner zone

In the example given below, the accuracy of the initial scheme can be defined as a third-order spatial accuracy on smooth meshes, through the use of a MUSCL-type upwind-biased finite volume, combined with a fourth-order time accuracy through the use of the standard Runge-Kutta scheme, see [19] for details.

- A **coarse grid** is defined on the inner zone by applying cell agglomeration in such a way that on each macro-cell, the maximal local stable time step is at least $K \Delta t$. Agglomeration consists in considering each cell and aggregating to it neighboring cells which are not yet aggregated to an other one (Figure 1). Agglomeration into macro-cell is re-iterated until macro-cells with maximal time step smaller then $K \Delta t$ have disappeared.

– Time-advancing on the macro-cells :

- We advance in time the chosen explicit scheme on the coarse grid with $K \Delta t$ as time step.

– Construction of the flux on the coarse grid

- The nodal fluxes Ψ_i are assembled on the fine cells (as usual)
- Fluxes are summed on the macro-cells I (inner zone) :

$$\Psi^I = \sum_{k \in I} \Psi_k \tag{2}$$

Remark 1: Stability can be reinforced by adding a smoothing of the coarse flux (inner zone) :

$$\Psi^I = \left(\sum_{K \in \mathcal{V}(I)} \Psi^K \text{vol}^K \right) / \left(\sum_{K \in \mathcal{V}(I)} \text{vol}^K \right). \quad (3)$$

We did not need to apply this flux-smoothing. \square

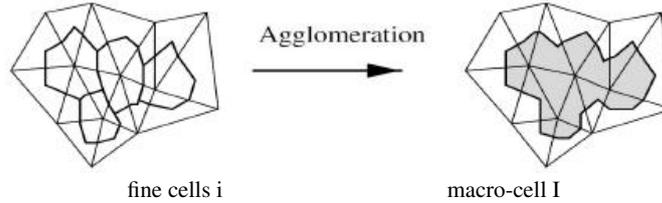


Fig. 1 Sketch (in 2D) of the agglomeration of 4 cells into a macro-cell. Cells are dual cells of triangles, bounded by sections of triangle medians.

2.3 Multirate time advancing

The multirate algorithm is then based on a **prediction step** and a **correction step** as defined hereafter :

Step 1 (prediction step) :

The solution is advanced in time with time step $K \Delta t$, on the macro-cells in the inner zone and on the fine cells in the outer zone :

For $\alpha = 1$, *RK step*

$$\text{outer zone : } \quad \text{vol}_i w_i^{(\alpha)} = \text{vol}_i w_i^{(0)} + b_\alpha K \Delta t \Psi_i^{(\alpha-1)} \quad (4)$$

$$\text{inner zone : } \quad \text{vol}^I w^{I,(\alpha)} = \text{vol}^I w^{I,(0)} + b_\alpha K \Delta t \Psi^{I,(\alpha-1)} \quad (5)$$

$$w_i^{(\alpha)} = w^{I,(\alpha)} \quad \text{for } i \in I \quad (6)$$

EndFor α .

where b_α denote the Runge-Kutta parameters, and vol^I the volume of macro-cell I.

Step 2 (correction step) :

- The unknowns are frozen in the outer zone at level $t^n + K \Delta t$.
- The outer unknowns near the boundary of the outer zone which are necessary for advancing the inner zone are interpolated in time.
- In the inner zone, using these interpolated values, the solution is advanced in time with the chosen explicit scheme and time step Δt .

This time advancing writes:

For $kt = 1, K$
 For $\alpha = 1, RK\ step$

$$\text{inner zone : } \quad vol_i w_i^{(\alpha)} = vol_i w_i^{(0)} + b_\alpha \Delta t \Psi_i^{(\alpha-1)} \quad (7)$$

EndFor α .
 EndFor kt .

Remark 2: The complexity, proportional to the number of points in the inner zone, is therefore mastered. \square

3 Elements of analysis

3.1 Stability

The central question concerning the coarse grid is the stability resulting from its use in the computation.

Considering (1), we expect that the viscous stability limit will improve by a factor four for a twice larger cell. The viscous stability limit can therefore be considered as more easily addressed by our coarsening. For the advective stability limit, we can be a little more precise. The coarse mesh is an unstructured partition of the domain in which cells are polyhedra. Analyses of time advancing schemes on unstructured meshes are available in L^2 norm for unstructured meshes, see [3, 11, 10]. Here we solely propose a L^∞ analysis of the first order advection scheme. The gain in L^∞ stability can be analysed for a first-order upwind advection scheme. We get the following (obvious) lemma:

Lemma : The upwind advection scheme is positive on the mesh made of macro-cells as soon as for all macro-cell I :

$$\Delta t \|V_I\| < \left[\sum_{J \in \mathcal{N}(I)} \int_{\partial cell(I) \cap \partial cell(J)} d\Sigma \right]^{-1} \int_{cell(I)} dx$$

where $\mathcal{N}(I)$ holds for the neighbouring macro-cells of I . \square

Thanks to the application of an adequate neighboring-cell agglomeration, macro-cell shapes will not degrade and the ratio in RHS will be increased, producing finally a K -times larger stability limit.

3.2 Accuracy

In contrast to more sophisticated multirate algorithms, the proposed method has not a strong control of the accuracy. Let us first remark that the generic situation involves variable-size meshes, which limits the unsteady accuracy on small scales, already before applying the multirate method.

However the two following remarks tend to show that the scheme accuracy - on the coarser grid- is conserved:

- the predictor step involves a sum of the fluxes and is at least as accurate as an equivalent coarse-grid approximation,
- the corrector step will improve the result in a way which depends on mesh smoothness, *i.e.* in better extent if the transition from small cells to larger cells is a smooth one.

3.3 Efficiency

The proposed two-level multirate depends on only one parameter, the ratio K between the large and small time step. Considering a mesh with N vertices, a short loop on the mesh will produce the function $K \mapsto N^{small}(K) \leq N$ which gives the number of cells in the inner region for K .

If $CPU_{ExpNode}(\Delta t)$ denotes the CPU per node and per time step Δt of the underlying explicit scheme, a model for the multirate cpu per Δt would be

$$CPU_{MR(K)}(\Delta t) = \left(\frac{N}{K} + N^{small}(K) \right) \times CPU_{ExpNode}(\Delta t)$$

to be compared with the explicit case:

$$CPU_{Expli}(\Delta t) = N \times CPU_{ExpNode}(\Delta t).$$

Once we have evaluated $K \mapsto N^{small}(K)$ for a given mesh it is possible to predict a theoretical optimum K_{opt} for minimising the CPU time in scalar execution.

3.4 Parallelism

The proposed method is experimented with a parallel MPI software which relies on mesh partitioning. Cell-agglomeration is applied inside partitions to save communications.

In first study, we have not optimised the MPI partitioning in order to adapt it to the two different steps of the multirate algorithm, predictor over the whole mesh, corrector on the inner cells only. The efficiency with respect to non-parallel calculation may then deteriorate. If the inner zone is large, mesh partition can be optimised by splitting uniformly the inner zone between the processors. We have not applied this strategy because, in the experiments we performed, we have observed that the inner zone can be of very small size, and, on the other hand, we have preferred to keep the present domain decomposition which optimizes the (many) intercore communications that occur during a computation. A consequence is a poor parallel efficiency for the correction step, related to a defavourable ratio between computation and communication, and not related to unbalanced load when a small inner zone is involved.

In a second study, we use a special partition, the multi-constrained partition of Karypis, [17] which is incorporated in the METIS library. The principle is apply the partition optimizer by minimizing the communications while imposing the two constraints that both the whole domain and the inner zone are balanced by the partition. In practice, we have applied the Karypis algorithm to a vertex partition and then transposed to elements for running our CFD code.

If the partition is ideally perfectly balanced for the inner nodes and for all nodes, then for $nproc$ processors, both $CPU_{MR(K)}(\Delta t)$ and $CPU_{expli}(\Delta t)$ will be $nproc$ smaller and the

theoretical gain is as for the scalar case.

Remark 3 : Another idea, not implemented in this work, is to organize concurrency between the corrector step and a part of the predictor. Indeed, the corrector only needs the predicted values of an intermediate zone in the coarse cells region located near the interface with the fine cells region. Then a possible improvement of parallel efficiency could be to first compute the predictor step in the intermediate zone, then communicate in order to allow a concurrent computation of (the rest of, with some overlapping) the predictor step together with the corrector step (to be concentrated on a very few processors). \square

4 Applications

The multirate algorithm is implemented into the parallel (MPI) CFD code AIRONUM shared by INRIA Sophia-Antipolis, LEMMA company and University of Montpellier. A description of this tool, which solves with a mixed element/volume method on unstructured meshes the compressible Euler and Navier-Stokes equations possibly equipped with a turbulence model, can be found in [22] and [16].

4.1 Contact discontinuity

In this first example, we consider the case of a moving contact discontinuity. For this purpose, the compressible Euler equations are solved in a rectangular parallelepiped as computational domain where the density is initially discontinuous at its middle (see Figure 2) while velocity and pressure are uniform.

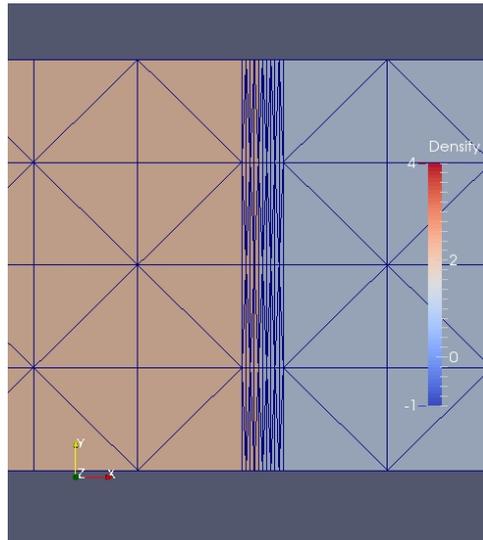


Fig. 2 ALE calculation of a traveling contact discontinuity. Instantaneous mesh with mesh concentration in the middle of zoom and corresponding advected discontinuous fluid density.

The uniform velocity is a purely horizontal one. As can be seen in Figure 2, small cells are present on either side of the discontinuity. The mesh moves during the computation in such a way that the nodes located at the discontinuity are still the same, and that the number of small cells are equally balanced on either side of the discontinuity. An Arbitrary

K	CPU explicit (s/ Δt /node)	$N^{small}(K)/N$ (%)	Expected gain (scalar)	CPU pred. phase (s/ $K\Delta t$)	CPU correc. phase (s/ $K\Delta t$)	Measured gain (parallel)
5	$4.96 \cdot 10^{-6}$	1.3	4.7	0.124	0.244	1.7
10	$4.96 \cdot 10^{-6}$	1.3	8.8	0.124	0.482	2.0
15	$4.96 \cdot 10^{-6}$	1.3	12.5	0.124	0.729	2.2

Table 1 ALE propagation of a contact discontinuity: Time step factor K , CPU of the explicit scheme per explicit time-step Δt and per node, percentage of nodes in the inner region, theoretical gain in scalar mode, CPU of the prediction phase per time-step $K\Delta t$, CPU of the correction phase per time-step $K\Delta t$, and measured parallel gain.

Lagrangian-Eulerian formulation is then used to solve the Euler equations on the resulting deforming mesh. Our long term objective is to combine the multirate time advancing with a mesh adaptation algorithm in such a way that the small time steps imposed by the necessary good resolution of the discontinuity remain of weak impact on the global computational time.

The mesh used in this simulation contains 25000 nodes and 96000 tetrahedra. The computational domain is decomposed into 2 subdomains. When integer K , used for the definition of the inner and outer zones, is set to 5, 10 and 15, the percentage of nodes located in the inner zone is always 1.3%, which corresponds to the vertices of the small cells located on either side of the discontinuity.

The multirate scheme with the aforementioned values of K , as well as a 4-stage Runge-Kutta method, are used for the computation. Each simulation was run on 2 cores of a Bullx B720 cluster. In Table 1, CPU times (prediction phase / correction phase) are given for the multirate approach and different time step factors K . An improvement in the efficiency of about 1.7, 2.0 and 2.2 is observed when K is set to 5, 10 and 15, respectively. Though a gain up to 2.2 is reached, the small number of inner nodes, leading to a poor computation/communication ratio in the correction phase, can explain that the obtained gain is rather moderate when compared to the scalar theoretical case.

4.2 Tandem cylinders

Our main study concerning the application of multirate is the calculation of a flow around a tandem cylinders at Reynolds number 1.66×10^5 . This was a test case of an AIAA workshop, see [20]. It is a challenging computation since several complex flow features need to be captured around multiple bodies (stagnation zones, boundary layers, shear layers, separations, laminar-turbulent transition, recirculations, vortex sheddings, wakes). Furthermore, small cells are necessary for a proper prediction of the very thin boundary layers, which implies very small global time steps so that classical explicit calculations become very costly. The application of our multirate scheme to the tandem cylinders benchmark is also made more difficult by the fact that we use a hybrid turbulence model based on RANS and VMS-LES approaches [16], so that additional equations associated with turbulent variables need to be advanced in time.

In order to illustrate the quality of resolution, the Q-criterion isosurfaces are shown in Figure 5. It shows the complex flow features and the very small structures that need to be captured by the numerical model and the turbulence model, which renders this simulation particularly challenging. Further information concerning the comparison between computation and experiments are available in [16].

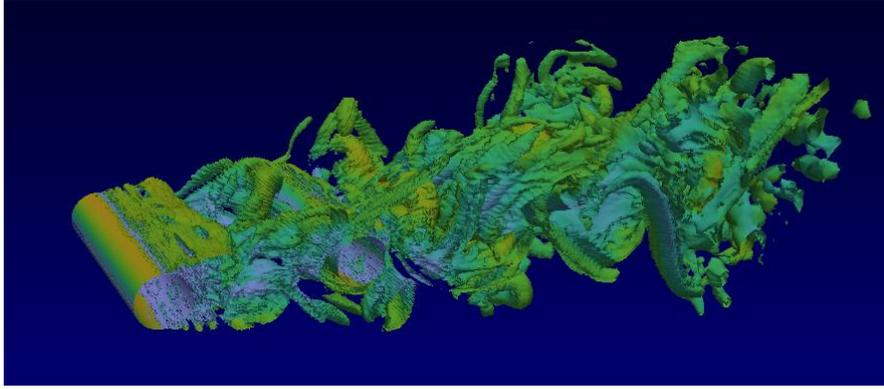


Fig. 3 Tandem cylinders - coarse mesh : instantaneous Q-criterion isosurfaces.

Two meshes were used for this study : a coarse mesh which contains 2.6 million nodes and 15 million tetrahedra, and a fine mesh with 16 million nodes and 96 million tetraedra. For both meshes, the smallest cell thickness is $1.2 \cdot 10^{-4}$.

– Coarse mesh

The computational domain is decomposed into 192 subdomains. The CFL number is set to 0.5 for the explicit and multirate computations. When integer K , used for the definition of the inner and outer zones, is set to 2, 5 and 10, the percentage of nodes located in the inner zone is 4%, 16% and 25%, respectively.

CPU times (prediction phase / correction phase) are given in Table 2 for the multirate approach and different time step factors K . For this test case, the multirate scheme is not very efficient due to a too costly correction phase (a large number of inner nodes not equally distributed among the subdomains). It is also true that the theoretical scalar gain is rather small. One can notice that with an implicit simulation and a CFL number set to 30, the gain is large compared to the explicit option. However, the accuracy is degraded with the implicit approach compared to the multirate option (see the relative error in Table 2 and Figure 4).

K	CPU explicit (s/ Δt /node)	$N^{small}(K)/N$ (%)	Expected gain (scalar)	CPU pred. phase (s/ $K\Delta t$)	CPU cor. phase (s/ $K\Delta t$)	Measured gain (parallel)	Error (%)
2	$3.5 \cdot 10^{-7}$	4	1.85	0.92	0.86	1.03	$1.67 \cdot 10^{-5}$
5	$3.5 \cdot 10^{-7}$	16	2.77	0.92	4.10	0.92	$2.57 \cdot 10^{-5}$
10	$3.5 \cdot 10^{-7}$	25	2.86	0.92	8.48	0.98	$7.67 \cdot 10^{-4}$
Implicit						53	$2.5 \cdot 10^{-1}$

Table 2 Tandem cylinder - coarse mesh: Time step factor K , CPU of the explicit scheme per explicit time-step Δt and per node, percentage of nodes in the inner region, theoretical gain in scalar mode, CPU of the prediction phase per time-step $K \Delta t$, CPU of the correction phase per time-step $K \Delta t$, measured parallel gain, and relative error.

– Fine mesh

The computational domain is decomposed into 768 subdomains, and as many cores on a Bullx cluster were used to perform these computations. When integer K , used for the definition of the inner and outer zones, is set to 5, 10 and 20, the percentage of nodes located in the inner zone is 18%, 24% and 35%, respectively (see Table 3).

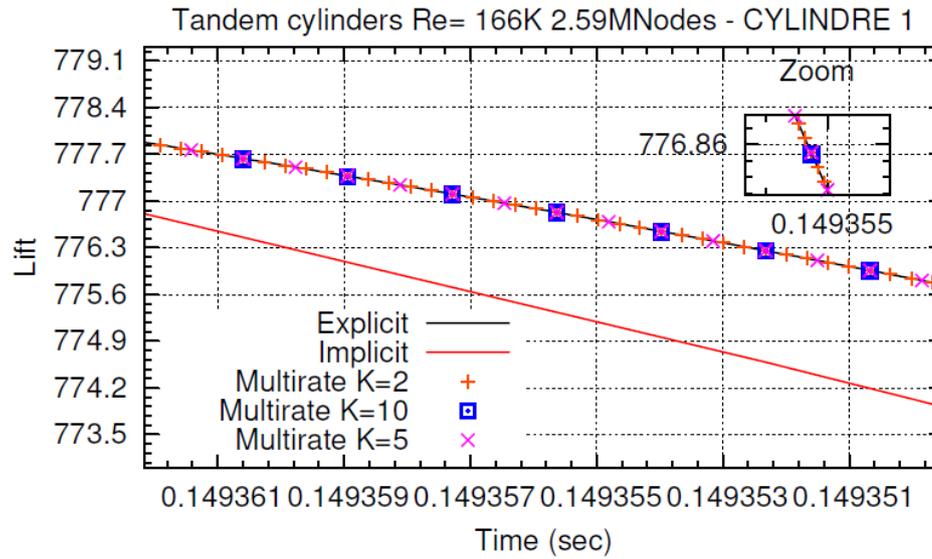


Fig. 4 Tandem cylinder - coarse mesh : zoom of the lift curves obtained with explicit, implicit and multirate schemes for the first cylinder.

The CPU times for the explicit and multirate schemes are shown in Table 3. As for the coarse mesh and for the same reason, the multirate option turns out to be not very efficient. Again, the theoretical scalar gain is rather small in this case (see Table 3).

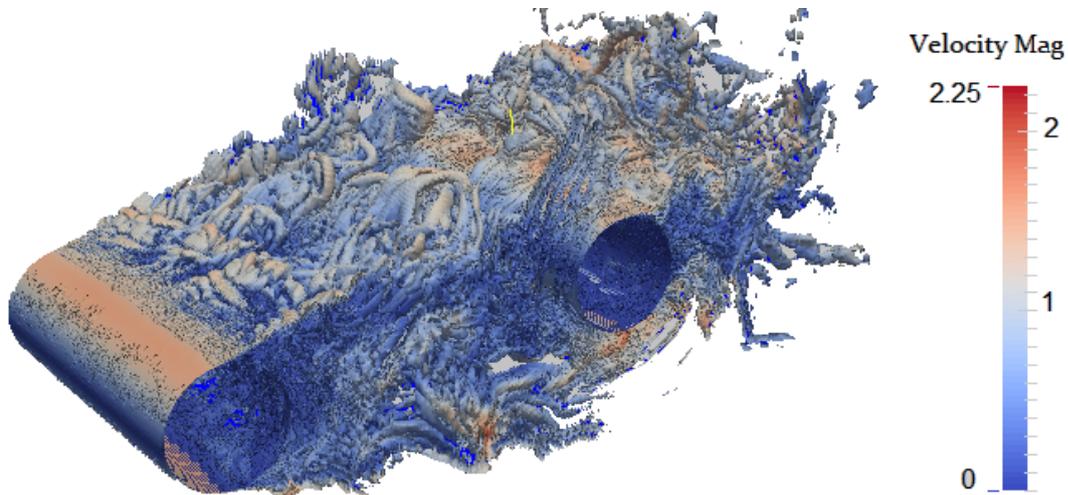


Fig. 5 Tandem cylinders - fine mesh : instantaneous Q-criterion isosurfaces. (coloured with velocity modulus).

4.3 Circular cylinder at very high Reynolds number

The third application concerns the simulation of the flow around a circular cylinder at Reynolds number 8.4×10^6 . As for the previous benchmark, the computational domain is made of small cells around the body in order to allow a proper representation of the very thin boundary layer that occurs at such a high Reynolds number. On the other hand, the same hybrid RANS/VMS-LES model as that of the previous benchmark is used to compute this flow, which implies again that both the fluid and turbulent variables need to be advanced by the time integration

K	CPU explicit (s/ Δt /node)	$N^{small}(K)/N$ (%)	Expected gain (scalar)	CPU pred. phase (s/ $K\Delta t$)	CPU cor. phase (s/ $K\Delta t$)	Measured gain (parallel)
5	10^{-7}	18	2.63	1.55	6.93	0.91
10	10^{-7}	24	2.94	1.52	14.15	0.99
20	10^{-7}	35	2.50	1.53	28.94	1.02

Table 3 Tandem cylinder - fine mesh: Time step factor K , CPU of the explicit scheme per explicit time-step Δt and per node, percentage of nodes in the inner region, theoretical gain in scalar mode, CPU of the prediction phase per time-step $K\Delta t$, CPU of the correction phase per time-step $K\Delta t$, and measured parallel gain.

scheme, and therefore also the multirate method. Figure 6 depicts the Q-criterion isosurfaces and shows the very small and complex structures that need to be captured by the numerical and the turbulence models, which renders this simulation very challenging.

The mesh used in this simulation contains 4.3 million nodes and 25 million tetrahedra. The smallest cell thickness is $2.5 \cdot 10^{-6}$. The computational domain is decomposed into 768 subdomains. When integer K , used for the definition of the inner and outer zones, is set to 5, 10 and 20, the percentage of nodes located in the inner zone is 15%, 19% and 24%, respectively (see Table 4).

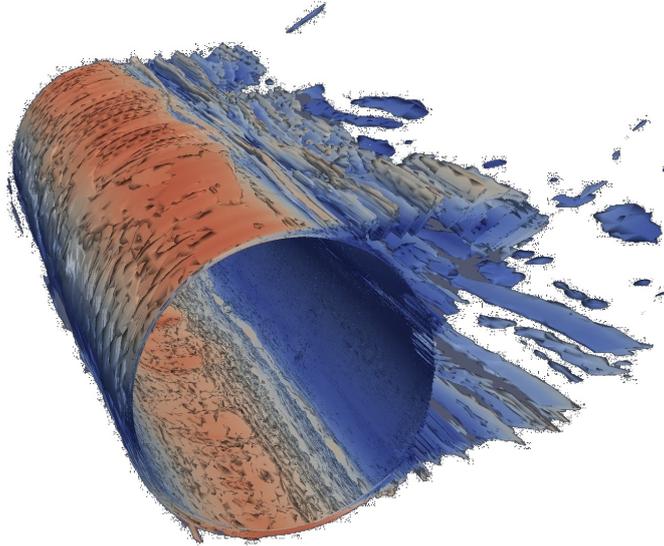


Fig. 6 Circular cylinder at Reynolds number 8.4×10^6 : instantaneous Q-criterion isosurfaces (coloured with velocity modulus).

The explicit scheme is the 4-stage Runge-Kutta method. For each simulation, 768 cores were used on a Bullx B720 cluster, and the CFL number was set to 0.5. CPU times for the explicit and multirate schemes with different values of K are given in Table 4. One can observe that the efficiency of the multirate approach is rather moderate, probably less than 1%. The cost of the correction phase is indeed relatively high compared to the prediction phase. This is certainly due to an important number of inner nodes (which implies also a moderate theoretical scalar gain) and a non uniform distribution of these nodes among the computational cores. An implicit simulation, with a CFL number set to 30, was also performed. An impor-

K	CPU explicit ($s/\Delta t/\text{node}$)	$N^{\text{small}}(K)/N$ (%)	Expected gain (scalar)	CPU pred. phase ($s/K\Delta t$)	CPU cor. phase ($s/K\Delta t$)	Measured gain (parallel)	Error (%)
5	$8.4 \cdot 10^{-8}$	15	2.86	0.39	1.53	1.02	$4.4 \cdot 10^{-4}$
10	$8.4 \cdot 10^{-8}$	19	3.45	0.39	3.12	1.11	$7.8 \cdot 10^{-4}$
20	$8.4 \cdot 10^{-8}$	24	3.45	0.39	6.24	1.18	$2.6 \cdot 10^{-3}$
Implicit						12.12	1.0

Table 4 Circular cylinder: Time step factor K , CPU of the explicit scheme per explicit time-step Δt and per node, percentage of nodes in the inner region, theoretical gain in scalar mode, CPU of the prediction phase per time-step $K\Delta t$, CPU of the correction phase per time-step $K\Delta t$, measured parallel gain, and relative error.

tant gain is observed compared to the multirate case, but at the cost of a notable degradation of the accuracy (see Table 4 and Figure 7).

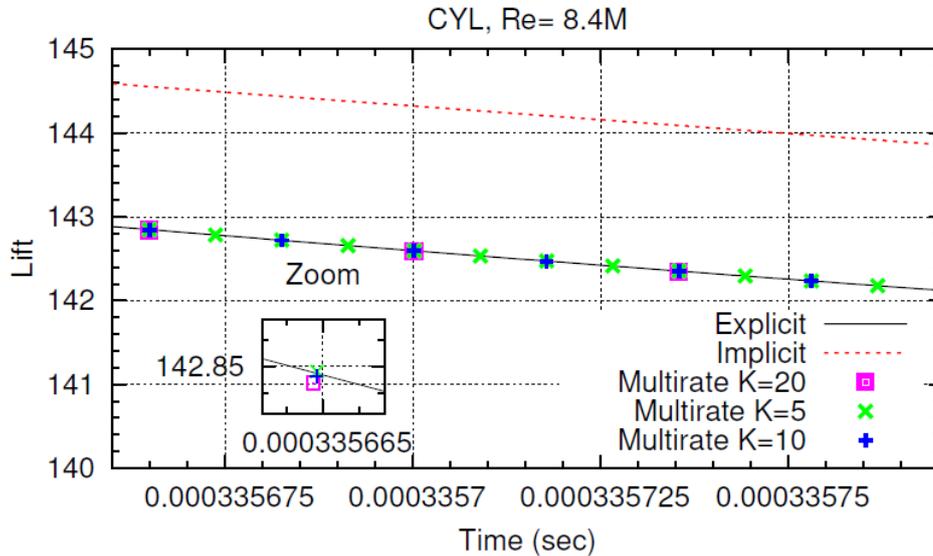


Fig. 7 Circular cylinder at Reynolds number 8.4×10^6 : zoom of the lift curves obtained with explicit, implicit and multirate schemes.

4.4 Spatial probe

The last case is the supersonic flow around a probe model for Exomars (see for example [12]). The Reynolds number is 1 million with respect to probe diameter. Delicate features in this simulation are a separation arising on a highly curved wall and relatively large recirculation zone at afterbody. Hybrid RANS-LES calculation brings more information than pure RANS does. The mesh involves 4.38 million nodes, 26 million tetrahedra, and the smallest mesh thickness is 2.10^{-5} . A sketch of this flow is presented in Figure 8. The gain in efficiency varies from 2.93 with $K = 10$ and 56 cells in the inner zone to a maximum of 3.82 with $K = 40$ and 151 cells in the inner zone (located in the high curvature region of the boundary layer where few very small cells are present). For $K = 40$, only 15 subdomains have inner cells (mean number 15 inner cells). On the other hand, the impact of the multirate approach on accuracy is very low compared to the implicit case, even for $K = 40$ which corresponds to an effective CFL number of 20 (see Table 4.4 and Figure 9).

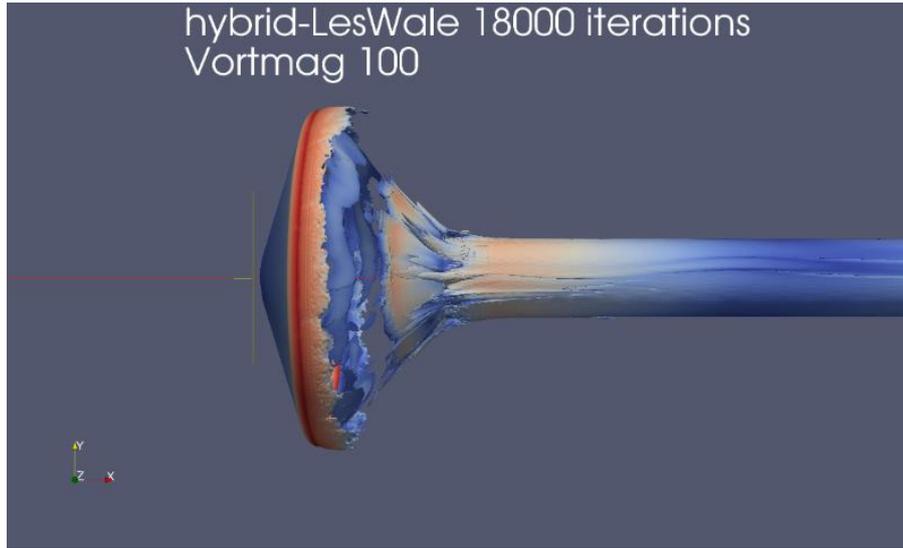


Fig. 8 Flow around a probe model at Reynolds number 1 million. Q criterion.

K	CPU explicit ($s/\Delta t/\text{node}$)	$N^{small}(K)/N$ (%)	Expected gain (scalar)	CPU pred. phase ($s/K\Delta t$)	CPU cor. phase ($s/K\Delta t$)	Measured gain (parallel)	Error (%)
10	$4.13 \cdot 10^{-7}$	0.015	8.69	1.81	4.36	2.93	$1 \cdot 10^{-5}$
40	$4.13 \cdot 10^{-7}$	0.040	15.38	1.83	17.35	3.82	$1.6 \cdot 10^{-4}$
Implicit						36.88	$2 \cdot 10^{-2}$

Table 5 Spatial probe: Time step factor K , CPU of the explicit scheme per explicit time-step Δt and per node, percentage of nodes in the inner region, theoretical gain in scalar mode, CPU of the prediction phase per time-step $K\Delta t$, CPU of the correction phase per time-step $K\Delta t$, measured parallel gain, and relative error.

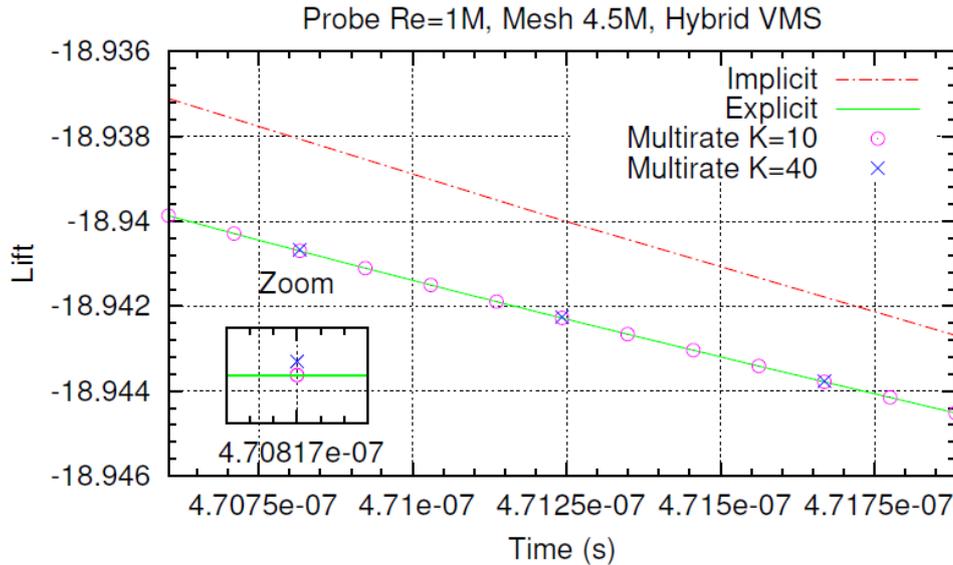


Fig. 9 Spatial probe : zoom of the lift curves obtained with explicit, implicit and multirate schemes.

4.5 Numerical experiments with a constrained partition

In this section we use a multi-constrained partitioning described in [17] and included in the METIS library.

The motivation of multi-constrained partitioning is to produce a partition which is quasi op-

Test case	Mesh size (vertices)	nproc	K	Measured gain (usual partition)	Measured gain (MCP)
Single cylinder Re=8.4M	4.3M	384	20	1.18	1.51
Tandem cylinder Re=166K	2.59M	192	20	1.2	1.29
Tandem cylinder Re=166K	16M	192	20	not computed	1.77
Tandem cylinder Re=166K	16M	768	20	1.02	2.0

Table 6 Gains obtained by parallel multirate simulations using usual partitions and parallel multirate simulations using multi-constrained partitions (noted MCP), when compared to the parallel explicit simulations using the same partitions.

timal for the several steps of the parallel MPI multirate algorithm of the considered code.

The two constraints in our case are:

- balance the predictor phase applying to the whole mesh,
- balance the corrector phase applying to the inner-nodes submesh

New multi-constrained partition have been built for the cylinder and the tandem problems using the METIS library. Table 6 shortly presents the main output of this new series of computations. Better parallel gains, when compared to the explicit scheme on the same partitions, are observed with the multirate method on multi-constrained partitions. However, it should be noted that we could not obtain a perfect partition on the considered meshes which would have allowed for an equally shared workload between the computational cores for both the predictor step and the correction step.

5 Conclusion

A new simplified multirate strategy for unstructured finite volume CFD is proposed in this work. The motivation of this research is two folds. The first motivation is related to the computation of high Reynolds number flows. Hybrid turbulence simulations with Reynolds numbers ranging from hundreds of thousands to millions can be computed with implicit time advancing methods in order to maintain a reasonable computational cost. But in many cases this is achieved with an important degradation of the accuracy with respect to computations based on smaller time steps on the same mesh. The second motivation is related to the arising of novel anisotropic mesh adaptation methods. The complexity of computations with large and very small cells in a mesh needs to be mastered with new methods. Now, since the average hybrid turbulence algorithm is already very complex, it is important to develop multirate methods which are simple and for which the programming into existing CFD is only slightly intrusive. The proposed method is based on control volume agglomeration, and relies on a prediction step where large time steps are used and where the fluxes for the smaller elements are evaluated on macro-cells for stability purpose. A correction step follows in which only the smaller elements are advanced in time with a small time step. The modifications in an existing explicit unstructured code, required by the implementation of such a multirate algorithm, demand little effort. Preliminary interesting results are given. They show that the proposed multirate strategy can be applied to complex unsteady CFD problems such as the prediction of three-dimensional flows around bluff bodies with an hybrid RANS/LES turbulence model.

Simulations, representative of problems that can be encountered in industrial applications, up to a Reynolds number as high as 8.4 million, are performed. For the considered flow calculations, the fully explicit option is still usable but of high computational cost. We observed that the proposed multirate strategy offers a really superior efficiency when the number of inner nodes, associated with very small cells in the mesh, is rather moderate. Thanks to the use of an explicit Runge- Kutta time advancing, the time accuracy of the multirate scheme remains high and the dissipation remains low, by comparison with an implicit computation. Only very small time-scales are lost with respect to a pure explicit computation.

The case of a simplified mesh adaptive calculation is also studied. Due to its simplicity, the proposed method can be easily extended to several multirate layers corresponding to different time step stability regions in order to separate, for example in the case of three layers, very small scales from intermediate ones, and intermediate scales from larger ones.

In a first series of numerical experiments, CPU gains are obtained with a usual mesh partition. We also performed calculations using a particular partition constrained by the balancing of both global mesh and inner nodes submesh. Ideally, the workload should be equally shared for both Step1 (prediction phase) and Step2 (correction phase) of the multirate algorithm. However, we could not obtain such an ideal partition with our meshes. Further work on partitioning is probably necessary to get perfectly balanced workloads without too large communication times.

In summary, the proposed multirate method is easy to program into a complex CFD code, is very stable in practice and the loss of accuracy with respect to an explicit scheme is very low, in contrast to implicit BDF2-based calculations, although we applied the implicit scheme with a CFL of 30, not much larger than with the multirate calculation (CFL= 10 for $K = 20$ in our simulations). Implicit accuracy is limited not only by the intrinsic scheme accuracy but also by the conditions required to achieve greater efficiency which involve a sufficiently large time-step and a short, parameter dependant, convergence of the linear solver performed in the time advancing step. In contrast, explicit and multirate computations are parameter safe, and the accuracy of the multirate method is optimal in regions complementary to the inner zone, that is, in our vortex shedding flow simulations, where it can be necessary to propagate accurately vortices, for example from the first cylinder to the second one in the case of the tandem cylinders. A future work will assess the effect of the better accuracy provided by the multirate scheme on vortex propagation in LES and hybrid turbulence simulations.

References

1. J.F. Andrus. Numerical solution of systems of ordinary differential equations separated into subsystems. *SIAM J. Numerical Analysis*, 16(4), 1979.
2. J.F. Andrus. Stability of a multi-rate method for numerical integration of ODE's. *Computers Math. Applic.*, 25, No 2:3–14, 1993.
3. F. Angrand and A. Dervieux. Some explicit triangular finite element schemes for the Euler equations. *Int. J. for Numer. Methods in Fluids*, 4:749–764, 1984.
4. A. Belme, A. Dervieux, and F. Alauzet. Time accurate anisotropic goal-oriented mesh adaptation for unsteady flows. *J. Comput. Phys.*, 239:6323–6348, 2012.
5. M. J. Berger and P. Colella. Local adaptive mesh refinement for shock hydrodynamics. *J. Comput. Phys.*, 82:64–84, 1989.
6. E. Constantinescu and A. Sandu. Multirate timestepping methods for hyperbolic conservation laws. *J. sci. Comp*, 33(3):239–278, 2007.

7. C. Engstler and C. Lubich. Multirate extrapolation methods for differential equations with different time scales. *Computing*, 58:173–185, 1997.
8. C. Engstler and C. Lubich. Mur8: a multirate extension of the eight-order Dormer-Prince method. *Applied Numerical Mathematics*, 25:185–192, 1997.
9. C.W. Gear and D.R. Wells. Multirate linear multistep methods. *BIT*, pages 484–502, 1984.
10. M.B. Giles. Energy stability analysis of multi-step methods on unstructured meshes. *CFDL Report 87-1, MIT Dept. of Aero. and Astro.*, 1987.
11. M.B. Giles. Stability analysis of a Galerkin/Runge-Kutta Navier-Stokes discretisation on unstructured tetrahedral grids. *Journal of Computational Physics*, 132(2):201–214, 1997.
12. A. Guelhan, J. Klevanski, and S. Willems. Experimental study of the dynamic stability of the exomars capsule. In: *Proceedings of 7th European Symposium on Aerothermodynamics. ESA Communications. 7th European Symposium on Aerothermodynamics, 9.-12. Mai 2011, Brugge, Belgium*, 2011.
13. M. Günther, A. Kvaerno, and P. Rentrop. Multirate partitioned Runge-Kutta methods. *BIT*, 41(3):504–514, 2001.
14. M. Günther, A. Kvaerno, P. Rentrop, A. Guelhan, J. Klevanski, and S. Willems. Multirate partitioned Runge-Kutta methods. *BIT*, 38(2):101–104, 1998.
15. M. Günther and P. Rentrop. Multirate row methods and latency of electric circuits. *Applied Numerical Mathematics*, 13:83–102, 1993.
16. E. Itam, S. Wornom, B. Koobus, B. Sainte-Rose, and A. Dervieux. Simulation of multiple blunt-body flows with a hybrid variational multiscale model. *Conference on Modelling Fluid Flow (CMFF 15) The 16th International Conference on Fluid Flow Technologies Budapest, Hungary, September 1-4, 2015*.
17. G. Karypis and V. Kumar. Multilevel algorithms for multi-constraint graph partitioning. Internal Report 98-019, University of Minnesota, Department of Computer Science / Army HPC Research Center Minneapolis, MN 55455, 1998.
18. R. Kirby. On the convergence of high resolution methods with multiple time scales for hyperbolic laws. *Mathematics of Computation*, 72(243):129–1250, 2002.
19. B. Koobus, F. Alauzet, and A. Dervieux. Some compressible numerical models for unstructured meshes. In F. Magoulès, editor, *CFD Handbook*. CRC Press, Boca Raton, London, New York, Washington, 2011.
20. D. Lockard. Summary of the tandem cylinder solutions from the benchmark problems for airframe noise computations-I. *Proceedings of Workshop AIAA-2011-35*, 2011.
21. R. Löhner, K. Morgan, and O.C. Zienkiewicz. The use of domain splitting with an explicit hyperbolic solver. *Comput. Methods Appl. Mech. Engrg.*, 45:313–329, 1984.
22. C. Moussaed, M.V. Salvetti, S. Wornom, B. Koobus, and A. Dervieux. Simulation of the flow past a circular cylinder in the supercritical regime by blending rans and variational-multiscale les models. *Journal of Fluids and Structures*, 47:114–123, 2014.
23. P.R. Mugg. Construction and analysis of multi-rate partitioned Runge-Kutta methods. *Thesis, Naval Postgraduate School, Monterey, California, June 2012*.
24. J. R. Rice. Split Runge-Kutta method for simultaneous equations. *Journal of Research of the National Bureau of Standards - B. Mathematics and Mathematical Physics*, 64B(3), 1960.
25. J. Sand and S. Skelboe. Stability of backward Euler multirate methods and convergence of waveform relaxation. *BIT*, 32:350–366, 1992.
26. A. Sandu and E. Constantinescu. Multirate explicit Adams methods for time integration of conservation laws. *J. Sci. Comp*, 38:229–249, 2009.

27. V. Savcenco, W. Hundsdorfer, and J.G. Verwer. A multirate time stepping strategy for stiff ordinary differential equations. *BIT*, 47:137–155, 2007.
28. B. Seny, J. Lambrechts, T. Toulorge, V. Legat, and J.-F. Remacle. An efficient parallel implementation of explicit multirate Runge-Kutta schemes for discontinuous Galerkin computations. *J. of Comp. Phys.*, 256:135–160, 2014.
29. S. Skelboe. Stability properties of backward differentiation multirate formulas. *Applied Numerical Mathematics*, 5:151–160, 1989.