

Multi-scale problems, high performance computing and hybrid numerical methods

G. Balarac, G.-H. Cottet, J.-M. Etancelin, J.-B. Lagaert, F. Perignon and C. Picard

Abstract The turbulent transport of a passive scalar is an important and challenging problem in many applications in fluid mechanics. It involves different range of scales in the fluid and in the scalar and requires important computational resources. In this work we show how hybrid numerical methods, combining Eulerian and Lagrangian schemes, are natural tools to address this multi-scale problem. One in particular shows that in homogeneous turbulence experiments at various Schmidt numbers these methods allow to recover the theoretical predictions of universal scaling at a minimal cost. We also outline how hybrid methods can take advantage of heterogeneous platforms combining CPU and GPU processors.

1 Introduction

Numerical simulations have become a routine tool to develop, prototype and/or validate products and processes in industry. Applications encompass virtually all sectors of activity from Aeronautics, Automotive industry and Oil exploration to Circuit design, Biomechanics and Animations studios, to name a few. With the need to perform more and more realistic simulations and the advent of supercomputers, available in national or regional centers, the field of High Performance Computing

G. Balarac

LEGI, CNRS and Université de Grenoble, BP 53, 38041 Grenoble Cedex 9, France, guillaume.balarac@grenoble-inp.fr

G.-H. Cottet, J.-M. Etancelin, F. Perignon and C. Picard

Laboratoire Jean Kuntzmann, CNRS and Université de Grenoble, BP 53, 38041 Grenoble Cedex 9, France, georges-henri.cottet@imag.fr

J.-B. Lagaert

Laboratoire de Mathématiques, Université Paris 11, 91405 Orsay Cedex, France, jean-baptiste.lagaert@math.u-psud.fr

(HPC) is not anymore restricted to academia and scientific grand challenges but starts to reach SMEs.

HPC requires easy and flexible access to HPC facilities, obviously, and to master the appropriate programming language, but also to question the numerical methods and algorithms that are used in the simulations. These methods and algorithms should be adapted both to the physics of the problems to solve and to the architecture of the simulation platforms. Moreover since these platforms are often of an hybrid nature, that is combine different type of processors, typically CPU and GPU processors, one may also wish to develop or use methods which couple different types of algorithms that can be optimally distributed to different types of processors.

This is particularly desirable if the problem to solve is multi-level by nature. In that case the different scales that are to be represented can also be resolved on different types of processors. This is hybrid computing. In some sense the nature of the problem and of the hardware inspires the type of mathematical and numerical models that should be used for optimal efficiency.

The purpose of this paper is to describe ongoing work in our group towards hybrid computing for applications in turbulent transport of passive scalar. In the next section we briefly describe the physical context of this work. In section 3 we describe a hybrid method coupling a semi-Lagrangian method for the scalar transport and a spectral method for incompressible flows and we show some results obtained with this method. Section 4 is devoted to the implementation of scalar transport on GPU processors.

2 Universal scaling in turbulent transport

The prediction of the dynamics of a scalar advected by a turbulent flow is an important challenge in many applications. Some of these applications are illustrated in Figures 1 to 3. Figure 1 shows the dynamics of a pollutant ejected by a sewer in the Los Angeles bay at two different times. Figure 2 shows the atomization of a jet. In that case the transported quantity is the interface water-air interface [1]. Figure 3 shows a similar experiment but in the context of combustion. In this case the transported quantities are concentrations of chemical species [2]. All these illustrations share a common feature, namely that very small scales spontaneously appear and need to be captured if accurate predictions are needed for the location of the pollutant, the size of the droplets or the combustion efficiency, respectively, are sought.

The production of small scales in an advected scalar indeed reflects some fundamental turbulence properties and is driven by the value of the Schmidt number, Sc , the ratio between the viscosity of the fluid and the diffusivity of the scalar. if $Sc > 1$, the so-called Batchelor scale η_B which measures the size of the smallest scalar fluctuations is smaller than the smallest length scales of the turbulent flow (the Kolmogorov scale η_K). These scales are related by $\eta_B = \eta_K / \sqrt{Sc}$. More precisely, for $Sc > 1$, Batchelor [3] reports that the classical Corrsin-Obukhov cascade

associated with a $k^{-5/3}$ law (where k is the wave number) for the scalar variance spectrum [4, 5] is followed by a viscous-convective range with a k^{-1} power law. This viscous-convective range is followed by the dissipation range, where various theoretical scalings have been proposed for the spectrum [3, 6]. A direct consequence of this fact is that, for $Sc > 1$, in numerical simulations the scalar is more demanding, in terms of grid resolution, than the flow itself. It is therefore natural to envision numerical approaches which use different grid resolutions for the scalar and the momentum.

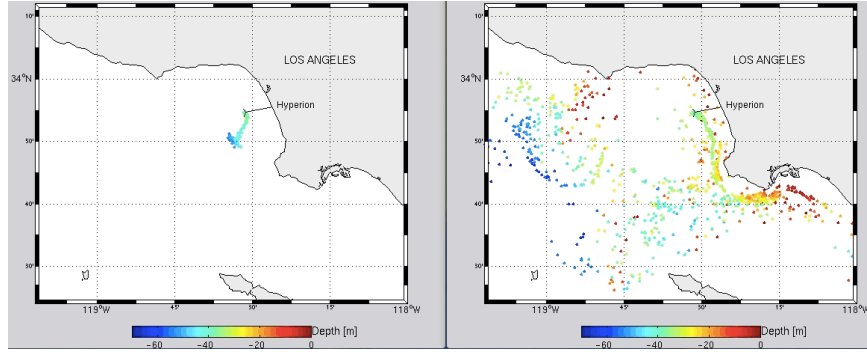


Fig. 1 Transport of a pollutant in the bay of Los Angeles at two different times. Courtesy of E. Blayo (Université Joseph Fourier, Grenoble)




Fig. 2 Atomization of a jet (Courtesy of S. Zaleski, Université Pierre et Marie Curie, Paris).




Fig. 3 Reacting jet . Courtesy of L. Vervisch, INSA Rouen.

3 Hybrid particle-spectral method

We consider in the following scalar equation

$$\frac{\partial \theta}{\partial t} + \mathbf{u} \cdot \nabla \theta = \nabla \cdot (\kappa \nabla \theta) \quad (1)$$

coupled with the incompressible Navier-Stokes equation

$$\frac{\partial \mathbf{u}}{\partial t} + \mathbf{u} \cdot \nabla \mathbf{u} = \nabla \cdot (\nu \nabla \mathbf{u}) - \nabla p, \quad \nabla \cdot \mathbf{u} = 0, \quad (2)$$

in a periodic box. κ is the molecular scalar diffusivity, ν the flow viscosity and \mathbf{u} the flow velocity field. Using different grid resolutions for the scalar and the flow has already been considered for instance in [7, 8]. In the latter reference a compact finite-difference method was used for the scalar and a pseudo-spectral method for the flow. A significant speed-up over a pure spectral solver with high resolution for

both the momentum and the scalar was obtained. Our choice here is to combine a particle method for the scalar and a spectral method for the flow.

Our motivation to choose a particle method for the scalar advection comes from the fact that, for large Schmidt numbers, the scalar dynamics is essentially driven by advection, a regime for which Lagrangian or semi-Lagrangian methods are well suited. Moreover in such method, the stability limits for the time-step are governed by the amount of strain in the flow, and not by the grid size. In the present context where high resolutions of the scalar are desired, this is definitely a feature that is expected lead to an important speed up.

More precisely, the method we use for the scalar is a semi-Lagrangian (or remeshed) particle method, where at every time step particles carrying the scalar values are moved along the streamlines of the velocity then remeshed on a regular cartesian grid. Remeshing particles on a regular grid is a way to guarantee the accuracy of particle methods. This approach has been systematically used and validated in a number of simulation of vortex flows [9, 10, 11, 12, 13] or in combination with level set methods for interface capturing [14, 15, 16]. Remeshing particles at every time-step also allows to easily couple the method to grid based methods, in particular when velocity values are computed on a grid. These methods can be summarized by the following formula

$$\theta_i^{n+1} = \sum_j \theta_j^n \Gamma \left(\frac{x_j^{n+1} - x_i}{\Delta x^\theta} \right). \quad (3)$$

where θ_j^n denotes the value of the scalar at the grid point x_j and at time $t_n = n\Delta t^\theta$, x_j^{n+1} is the location after one advection step of the particle initialized at time t_n on the grid point x_j , and Δx^θ and Δt^θ denote the grid size and the time-step. In the above formula Γ is an interpolating kernel, the smoothness and the moment properties of which govern the spatial overall accuracy of the method [17]. In this work we chose the following kernel second order kernel

$$\Gamma(x) = \begin{cases} \frac{1}{12} (1 - |x|) \left(25|x|^4 - 38|x|^3 - 3|x|^2 + 12|x| + 12 \right) & \text{if } 0 \leq |x| < 1 \\ \frac{1}{24} (|x| - 1) (|x| - 2) \left(25|x|^3 - 114|x|^2 + 153|x| - 48 \right) & \text{if } 1 \leq |x| < 2 \\ \frac{1}{24} (3 - |x|)^3 (5|x| - 8) (|x| - 2) & \text{if } 2 \leq |x| < 3 \\ 0 & \text{if } 3 \leq |x|. \end{cases}$$

The scalar time-step is given by $\Delta t^\theta = (|\nabla \mathbf{u}|_{\max})^{-1}$. As already mentioned this value does not depend on the scalar grid size.

For the momentum equation we use a classical pseudo-spectral method, with the 3/2 rule to de-alias inertial terms and a second-order Runge-Kutta scheme is used both for the time-stepping of the spectral method and to advect particles. Precise descriptions of the methods and of the experimental set up are given in [18, 17].

Figure 4 shows a comparison of the scalar spectra obtained by the present coupling method and pure spectral method in an experiment of decaying homogeneous

turbulence. In the hybrid method two different resolutions were used for the scalar. This experiment shows that, provided the particle method is used with slightly more grid points than needed by the spectral method, the scalar values are well recovered all the way to the dissipation scale. In this Direct Numerical Simulation, the Schmidt number was equal to 50 and the momentum equation was solved with 256 modes in each direction.




Fig. 4 Spectra of the scalar variance $E_\theta(k, t)$ at two two different times for $Sc = 50$. Right picture is a zoom of left picture on the smallest scales.

To evaluate the efficiency of the hybrid method, we show in Table 1 CPU times for the full spectral method and the hybrid method for $Sc = 50$. All runs correspond to fully resolved simulations for the Navier-Stokes equations. One can see that, because it can use much larger time-steps, the hybrid method, even when it uses slightly more points to accurately resolve the finest scales, leads to significant savings over the pure spectral method. Additional validation and diagnostics are give in [18].

The computational efficiency of the hybrid method allows to address more challenging cases and to investigate in a systematic fashion the universal scaling laws in the case of forced homogeneous turbulence. Table 2 summarizes the simulation set up corresponding to two values of the Reynolds number and several Schmidt numbers. For the highest Schmidt number, $Sc = 128$, the simulation used 3064^3 compu-

Method	N^u	N^θ	$\Delta t^u (\times 10^{-4})$	$\Delta t^\theta (\times 10^{-4})$	total CPU time
Spectral	1024^3	1024^3	2.5	2.5	43 590s
Spectral	256^3	1024^3	2.5	2.5	16 671s
Hybrid	256^3	1024^3	10	100	1 139s
Hybrid	256^3	1280^3	10	100	1 328s

Table 1 Numerical efficiency of the different methods on a decaying turbulence experiment- Runs are performed on 2048 cores of a Blue Gene Q. N^u, N^θ denotes the spatial resolution for velocity and scalar and $\Delta t^u, \Delta t^\theta$ are the numerical time steps for momentum and scalar equations. CPU times correspond to the simulation time $t = 6$.

tational elements for the scalar equation, on a IBM Blue Gene supercomputer. The ratio between the time-step used in the present simulation and that which would have been required in a comparable spectral simulation is almost equal to 100. Figure 5 shows the compensated spectra of the scalar for a Reynolds number based on the Taylor micro-scale R_λ [21] equal to 130. These spectra do exhibit a k^{-1} decay on a range which increases with the Schmidt number. Beyond this viscous-convective range, the spectra follow an exponential decay coinciding with the scaling law proposed by Kraichnan [6].

R_λ	N^u	Δt^u	Sc	N^θ	Δt^θ	$\Delta t_{\text{spec}}^\theta$
130	256^3	1.2e^{-2}	0.7	512^3	8.6e^{-2}	6e^{-3}
			4	1024^3		3e^{-3}
			8	1024^3		3e^{-3}
			16	1536^3		2e^{-3}
			32	1536^3		2e^{-3}
			64	2048^3		1.5e^{-3}
			128	3064^3		1e^{-3}
210	512^3	3e^{-3}	0.7	770^3	2e^{-2}	2e^{-3}
			4	1024^3		1.5e^{-3}

Table 2 Setup of simulations performed in forced homogeneous turbulence. Δt^u is the time step used to solve the Navier-Stokes equation with a pseudo-spectral solver. Δt^θ is the time step used to solve the scalar transport equation with the particle method. $\Delta t_{\text{spec}}^\theta$ is the time step which would be needed if a pseudo-spectral method was used for the same number of scalar grid points [18]

Figure 6 shows vorticity and scalar contours in a cross section of the computational box for the simulation corresponding $R_\lambda \approx 130$ and $Sc = 128$. It illustrates the scale separation between the flow and the scalar for these parameters. The extension of the hybrid method to the coupling of particle methods with finite-volume methods to address engineering configurations is under way.




Fig. 5 Compensated spectra for the scalar variance at $R_\lambda \approx 130$. The arrow shows the direction of increasing Schmidt numbers. In the dissipative region, the circles show the law proposed by Kraichnan and the squares show the law proposed by Batchelor in the dissipative scales. The vertical axis shows the spectra compensated by the Batchelor law predicting a k^{-1} decay in the intermediate scale.




Fig. 6 Cross-section colored by the vorticity magnitude (left, blue regions are for the lowest vorticity values and red regions are for the highest vorticity values) and by the passive scalar (middle, blue regions are for the lowest scalar values and red regions are for the highest scalar values) for $R_\lambda \approx 130$ and $Sc = 128$. The zooms (right) for the vorticity magnitude (top) and the scalar (bottom) correspond to the white box with a length approximately equal to the Kolmogorov scale.

4 Towards hybrid computing

As already mentioned the multi-scale nature of turbulent transport makes natural the idea of hybrid computing methodologies where different part of the problems are distributed to different types of hardware. To be able to implement hybrid algorithms on hybrid architectures, one needs to develop frameworks and libraries with a high level description which allows to distribute different solvers and grids

to different parts of the clusters in a seamless fashion. Both particle advection and particle remeshing are local operations. This limits the communications between computational elements and makes semi-Lagrangian particle methods well suited to parallel implementations [20]. Such an implementation is described in [19] for the 2D Navier-Stokes equations and in [17] for 3D linear transport equations. In [17], to achieve good portability, the computational frameworks are written using OpenCL.

The efficiency of GPU algorithms is very much conditioned by memory access strategies. To minimize the resulting computational overhead, we use a directional splitting where particles are pushed and remeshed successively along each direction. This allows to send a given number of independent particle lines on a single workgroup. This strategy requires to transpose data after each direction has been processed. However on modern GPU cards, transpositions can be achieved at a cost close to that of a simple copy operation. Figure 5 shows the computational cost of our GPU implementations [17] in double precision arithmetics for different remeshing kernels, for 2D and 3D experiments using about 16 million points. In these experiments a second order splitting was used to alternate one dimensional particle advection and remeshing. The number of points in the kernel stencils in each direction varied from 4 to 8 [17]. These calculations were done on a NVIDIA Tesla K20m. These performances reached between 20 and 50 % of the peak performance of the GPU, depending on the size of the stencil and represented a speed up of about 25 over a multi-threaded MPI implementation running on 8 Xeon E5-2640 cores

Hybrid computing would consist of combining the above implementation of scalar transport at high resolution with flow calculations on CPU processors. Based on timing obtained in our CPU and GPU implementations, in the case when the full scalar grid fits on a single GPU, up to resolutions of 512^3 , a target toy configuration where computational times on CPU and GPU would be similar, consist of a 128^3 flow resolution running on 8 CPUs together with a 512^3 scalar resolution running on the GPU, for an overall computational time of about 1s per iteration. To obtain such performance it is essential that communications between velocity data processed on the CPUs and the GPU are processed in an optimal way. This is the object of ongoing research.

5 Conclusion

Combining high order semi-Lagrangian and Eulerian methods is an efficient strategy to address turbulent transport problems. It allows to describe accurately the viscous-convective range and dissipation scales of the scalar at a minimal cost. This is due to the fact that semi-Lagrangian methods are not subject to CFL conditions. The local nature of particle methods naturally opens the way to hybrid computations using heterogeneous hardware.




Fig. 7 Code profiling of one time step for different problem sizes in double precision.

Acknowledgments

This work was partially supported by the Agence Nationale pour la Recherche (ANR) under Contracts No. ANR-2010-JCJC-091601 and ANR-2010-COSI-0009. G.-H. C. is also grateful for the support from Institut Universitaire de France. Computations reported in section 3 were performed using HPC resources from GENCI-IDRIS (Grant 2012-020611).

References

1. G. Tryggvason, R. Scardovelli, S. Zaleski, Direct Numerical Simulations of Gas-Liquid Multiphase Flows, Cambridge University Press, 2011.
2. G. Lodato, P. Domingo, L. Vervisch, Three-dimensional boundary conditions for direct and large-eddy simulation of compressible viscous flows, *Journal of Computational Physics* 227 (1) (2008) 5105–5143.
3. G. K. Batchelor, Small-scale variation of convected quantities like temperature in turbulent fluid part 1. general discussion and the case of small conductivity, *Journal of Fluid Mechanics* 5 (01) (1959) 113–133.
4. S. Corrsin, On the spectrum of isotropic temperature fluctuations in an isotropic turbulence, *Journal of Applied Physics* 22 (4) (1951) 469–473.
5. A. M. Obukhov, The structure of the temperature field in a turbulent flow, *Dokl. Akad. Navk. SSSR* 39 (1949) 391.
6. R. Kraichnan, Small-scale structure of a scalar field convected by turbulence, *Phys. Fluids* 11 (1968) 945–953.
7. G.-H. Cottet, G. Balarac, M. Coquerelle, Subgrid particle resolution for the turbulent transport of a passive scalar, in: *Advances in Turbulence XII, Proceedings of the 12th EUROMECH European Turbulence Conference, September, Vol. 132, 2009*, pp. 779–782.
8. T. Gotoh, S. Hatanaka, H. Miura, Spectral compact difference hybrid computation of passive scalar in isotropic turbulence, *Journal of Computational Physics* 231 (21) (2012) 7398–7414.
9. P. Koumoutsakos, A. Leonard, High-resolution simulations of the flow around an impulsively started cylinder using vortex methods, *Journal of Fluid Mechanics* 296 (1995) 1–38.

10. G.-H. Cottet, B. Michaux, S. Ossia, G. Vanderlinden, A comparison of spectral and vortex methods in three-dimensional incompressible flows, *J. Comput. Phys.* 175 (2) (2002) 702–712.
11. G.-H. Cottet, P. Poncet, Advances in direct numerical simulations of 3d wall-bounded flows by vortex-in-cell methods, *Journal of Computational Physics* 193 (1) (2004) 136–158.
12. P. Ploumhans, G. Winckelmans, J. Salmon, A. Leonard, M. Warren, Vortex methods for direct numerical simulation of three-dimensional bluff body flows: Application to the sphere at $Re=300, 500$ and 1000 , *Journal of Computational Physics* 178 (2) (2002) 427–463.
13. P. Poncet, Topological aspects of the three-dimensional wake behind rotary oscillating circular cylinder, *J. Fluid Mech.* 517 (2004) 27–53.
14. S. E. Hieber, P. Koumoutsakos, A lagrangian particle level set method, *Journal of Computational Physics* 210 (1) (2005) 342–367.
15. M. Bergdorf, P. K. Koumoutsakos, A Lagrangian particle-wavelet method, *Multiscale Modeling and Simulation: A SIAM Interdisciplinary Journal* 5 (2006) 980–995.
16. A. Magni, G. Cottet, Accurate, non-oscillatory, remeshing schemes for particle methods, *Journal of Computational Physics* 231 (1) (2012) 152–172.
17. G.-H. Cottet, J.-M. Etancelin, F. Perignon, C. Picard, High-order semi-lagrangian particle methods for transport equation: numerical analysis and implementation issues, to appear in *ESAIM: Mathematical Modelling and Numerical Analysis*.
18. J.-B. Lagaert, G. Balarac, G.-H. Cottet, Hybrid spectral-particle method for the turbulent transport of a passive scalar, *Journal of Computational Physics* 260 (1) (2014) 127–142.
19. D. Rossinelli, M. Bergdorf, G.-H. Cottet and P. Koumoutsakos, GPU accelerated simulations of bluff body flows using vortex methods, *Journal of Computational Physics*, 229 (9), 33163333, 2010.
20. I.F Sbalzarini, J.H. Walther, M. Bergdorf, S.E. Hieber, E.M. Kotsalis, P. Koumoutsakos, PPM - A highly efficient parallel particle-mesh library for the simulation of continuum systems. *J. Computational Physics*, 215:566-588, 2006
21. M. Lesieur, *Turbulence in fluids, Fluid mechanics and its applications*, Springer, Dordrecht, 2008.