Grid Design and the Fate of Eddies in External Flows

Philippe R. Spalart[1], Mikhail Kh. Strelets[2], and Andrey Garbaruk[2]

[1]Boeing Commercial Airplanes, Seattle, WA 98124, USA
philippe.r.spalart@boeing.com

[2]New Technologies and Services, St.-Petersburg 197198, Russia
strelets@mail.rcom.ru

Summary. A few years ago we introduced a nomenclature for grid regions in Detached-Eddy Simulations (DES) of external flows: the RANS, Focus, Euler, and Departure Regions [1]. These may help organize grid generation and error estimation/reduction. We have four purposes here: to point out that most of these concepts apply to LES, if not to DNS, and are not truly limited to external flows; to discuss what really defines a DNS for external flows, and even for internal flows with inflow-outflow conditions; to ask whether the SGS model should be sensitized to increasing grid spacing along the flow, hoping to achieve stable simulations of "real-life flows" with energy-conserving differencing, which is definitely not the norm today; and finally to present an initial study of an explicit definition of the boundary between Focus Region and Departure Region. The long-term goal is to extend automatic grid adaptation to the turbulence-resolving approaches.

Key words: Grid Generation, Accuracy, Turbulence Modelling

1 Introduction

Large-Eddy Simulation and its hybrids with RANS are rapidly expanding towards complex applications, which is very gratifying but can also produce some chaos, and even a backlash when work that is not careful enough is publicized or introduced in designs which then prove faulty. This concern is central to the organisation of QLES (although complex flows are not the only ones considered), but curiously the list of topics either for 2007 or 2009 does not contain grid design, which is however our primary focus here. Fewer than ten out of thirty-three papers this week highlight grid issues, and only two have adaptation. In our opinion, the grid is just as important as the solver, and is arguably more important than the SGS model, except very near walls.

A factor in how research topics are chosen is that SGS research is more publishable and less dependent on heavy software than grid research. Grid-generation technology and practices have much control both over the accuracy of research simulations and over the practicality and reliability of LES in engineering practice. Once

M.V. Salvetti et al. (eds.), *Quality and Reliability of Large-Eddy Simulations II*, ERCOFTAC Series 16, DOI 10.1007/978-94-007-0231-8_24,

the geometry has any complexity, and even for simple shapes when they cause massive separation, there is much latitude on where to "spend" the grid points, and sharp differences in the true resolution requirements at different points. As a result, grids generated with too little regard for these facts are wasteful.

In addition, the long-term goal of the CFD community must be automatic generation of grids with little waste, grids which will necessarily depend on the flow and not just the geometry; in other words, be adaptive. For this to develop, the thinking needs to be organized as well as possible, and to attract some degree of consensus. Note that as of 2009, adaptation is not the norm even for steady RANS solutions, especially not in commercial solvers. The benefits in time savings *and* in accuracy will be extensive. It is an active research area, but an arduous one, with deep open questions including the balance between solver tolerance for very uneven grids and smoothness in grid generation, and the option between heuristic adaptation strategies and more systematic ones, often based on adjoint formulations. Of course, since many flows contain boundary layers, thin shear layers, and possibly shocks, isotropic adaptation cannot be sufficient.

A dominant interest of the authors is in hybrid RANS-LES methods and especially DES, but significant overlap with grid design in LES and DNS has become clear. Similarly, some of the thinking which is unavoidable in external flows because of the unbounded domain proved to have some logic in internal flows as well. This little-explored territory led to a paper containing more "vision/speculation" than "finished products," with three sections made up of remarks with strategic or "philosophical" aspects, and a fourth with a much more concrete nature and a proposal which may contribute in due time to grid adaptation.

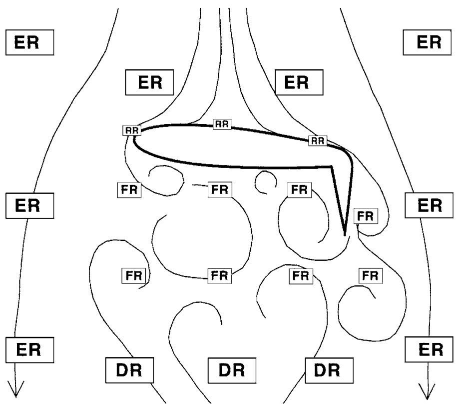

Fig. 1. Sketch of grid regions in the DES of an airfoil, with flow coming from above.

2 Grid Regions in DES, LES, and DNS

The "Young Person's Guide to DES Grids" is easily available as a NASA report [1], and the method contained in it will be called the "YPG approach." Figure 1 is a sketch from that report, and we use here the same nomenclature of RANS Region or RR, Focus Region or FR, Departure Region or DR, and Euler Region or ER, to which it will help to add VR for Vortical Region (this creates a conflict with "VR" for Viscous Region in the YPG, but the Viscous Region is not discussed here). At any instant, the VR is simply the complement of the ER, of course. However in the design of a stationary grid, the VR must be defined as the region which is visited by vorticity at any time during the simulation, and then the ER as its complement. This is because the resolution needs of VR are more demanding than those of the ER. This will be illustrated.

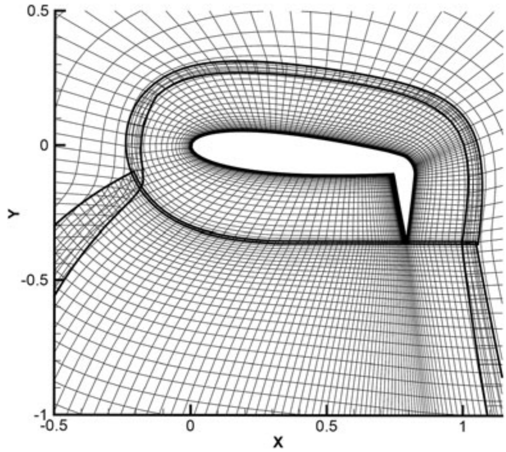

Fig. 2. Block-structured grid generated for airfoil [2].

Figure 2 shows an early example of grid design in DES, accomplished within the constraints of a block-structured solver with overset capability. Recall that the flow comes from above; this is the wing of a tilt-rotor aircraft during hover, with down-wash from the rotors. The Euler Region is most visible, in being coarser; the spacing in the z direction was also larger than in the other regions. The wasteful aspect that is not seen, of course, is that the entire simulation used a single time step, far shorter than the ER would need (however, the number of points in the ER is not dominant). The block in contact with the airfoil, with an unusual "snail" topology, covers the RANS region and part of the Focus Region. It extends into the ER above the airfoil, much farther than necessary for accuracy, but the advantage was to contain the number of blocks. The lower block covers part of the FR, and the Departure Region. Gradual coarsening is seen in the DR.

If an LES is considered, there is no pure RANS Region the way there is in a DES, where it includes the attached boundary layer on the upstream side of the airfoil. Still, if the LES has wall modelling, the very-near-wall layer will have a model with RANS character, and the grid will usually have RANS proportions, i.e., large cell aspect ratios. In a DNS, there is no such region at all; the cell aspect ratios are much larger than 1, but not unlimited nor Reynolds-number dependent. On the other hand, all simulations have an Euler Region, devoid of small eddies, and the most economical grid will have fairly coarse spacing in the ER.

The less obvious border is between the FR and DR. Recall that the DR is a region downstream of the bodies in which we accept poorer and poorer resolution, even though it contains turbulence. Having the grid spacing prevent the resolution of the smallest eddies is fundamental in LES, but not in DNS. Yet, all DNS of external flows has a DR, under another name or no name. The alternative to this would be to have a domain short enough to afford full DNS resolution everywhere, with an outflow boundary condition placed closely enough to keep the grid count under control. However, such conditions are crude and interfere with vortex shedding and similar phenomena, and the consensus seems clear that it is better to have a large domain with gradually expanding grid cells, and to push the outflow condition far away from the body. Most probably, in the DNS to date, the DR has effectively operated in Implicit LES (ILES) mode, that is, some feature of the numerics stabilized the simulation, as opposed to viscosity acting at the Kolmogorov length scale η which is not resolved any more. It could have been logical to declare that the DR was an LES region appended to a DNS, and activated an SGS model only in that region. Hoffman clearly recognized this concept, using the name DNS/LES, and counting on numerical stabilization in the effective DR [3].

Similar grid-design decisions face the scientist for some internal flows. Consider a backward-facing step flow, and the position of the downstream boundary of the simulation. If we insist on "DNS quality" over the entire domain, in simple terms on resolving η and the wall viscous scale ν/u_τ, much effort will be spent on the region, well past reattachment, where it is intuitively clear the small-scale turbulence cannot influence the step region. Again, one option is to end the domain close to the reattachment line with an outflow condition; the other option is to let the grid coarsen, arguing that only the larger eddies have a "dialogue" with the region upstream. The developing inaccuracy in skin friction, due to under-resolving the near-wall layer, will hardly be felt. Therefore, DNS resolution is not needed, and the best use of resources is to declare this region as a DR. This could also be viewed as a combination of DNS and LES. The resolution is testable in both regions: comparing results with different enough grid spacings provides meaningful indications of the residual errors (purely numerical in the DNS region, mixed numerical/SGS modelling in the LES region). Therefore, a thorough demonstration of accuracy is possible. This is although the order of convergence has no reasons to be the same in the two regions; in fact, to our knowledge, convergence according to a theoretical order of accuracy has never been demonstrated in a DNS or LES complex enough to be interesting. To begin with, it is usually impossible to reduce the scatter due to a finite time sample to levels such that grid convergence could be observed at all.

3 Requirements for a Simulation to be a DNS

DNS first developed primarily in homogeneous turbulence, with periodic conditions and uniform grids. This kept the criteria for "DNS Quality" very simple: the grid spacing (a single number) and time step had to be small enough in Kolmogorov viscous units. The assumption was not of an exact numerical solution, but that numerical errors were negligible. When DNS progressed to channels, the attitude was and still is much the same. There is agreement that *a DNS contains no empirical modelling, and numerical errors are negligible*, which is verified by running at least one case with a resolution finer than the standard, in all four directions. After that, usually when progressing to higher Reynolds numbers, guidelines are applied either in wall units or Kolmogorov units; in short, in viscous units. The simulation is highly accurate everywhere. Yet we have just argued that this attitude is not sustainable in external flows, and in some internal but inflow-outflow cases, and suggested that appending an LES region to a DNS is not illogical, at least downstream. It has not been formally applied to our knowledge; we rather presume that numerical dissipation has been at play. A proper study of that type would include grid-refinement studies in both regions, or for the DNS region the respect of established guidelines in viscous units, and also varying the DNS-LES interface position.

While planning a large DNS of pipe flow (for which they did not find the CPU time), the authors had reflections which are related to the ones just applied to external and inflow-outflow cases, even though the DNS was to be periodic in the streamwise direction. The overwhelming interest was in the Law of the Wall and Karman constant κ, so that the core region (say $r/R < 0.6$ with obvious notation) only needed to support the Reynolds shear stress that coupled with the wall layer. Therefore, resolving η in the core region of the pipe was not essential to the understood purpose of the study. Was it needed, then? Or could a genuine "DNS" aimed only at the wall layer be assembled with classical DNS in the wall region, and LES in the core region, thus reducing the cost? Boundary-layer DNS could lead to similar reflections: is the viscous superlayer resolved accurately, and does it need to be? This depends on the exact purpose of the study. Since some estimates make the viscous super-layer as thin as the viscous sub-layer, it certainly has not been resolved at any interesting Reynolds number.

Even in the law-of-the-wall region of the pipe, the dissipation range was not needed in itself; only the Reynolds shear stress, which is of course dominated by large eddies. If so, it made sense to let the grid spacing Δ grow proportionally to y, the distance from the wall, whereas η grows closely to $y^{1/4}$. At the very large y^+ values which were hoped for, this made some difference. Again, this does not deny the need for grid-refinement studies, but the difference between a systematic grid refinement with $\Delta \propto y$ and one with $\Delta \propto y^{1/4}$ is deep. This DNS may yet be conducted, although recent trends in DNS directed at the Karman constant are very discouraging; the Reynolds numbers needed are far higher than expected even a few years ago.

A similar and less controversial example is as follows. Some DNS studies have focused on the mean velocity and Reynolds-stress tensor. Others have instead focused on predictions of the dissipation tensor, to provide data for Reynolds-stress-transport models. It is clear that calculating these two tensors to the same accuracy does not require quite the same resolution; the dissipation tensor is more demanding. Thus, no DNS is exact, and the appropriate resolution depends on the purpose of the study, including the region of interest, the quantities of interest, and probably other features.

4 Sub-Grid-Scale Modelling in Coarsening Resolution

We have seen that resolved turbulence routinely needs to be convected into coarser grid regions, certainly for external flows and also for some internal ones. This demands the trouble-free destruction of the eddies too small for the grid spacing. The same applies to sound waves; it is not practical to track all such waves all the way to the outer boundaries, which gives great value to the Ffowcs-Williams-Hawkings equation, applied with permeable surfaces made rather tight around the turbulent region. The irrotational region outside the FWH surface can be viewed as a DR for waves, and jumps or even rapid increases in Δ pose a real danger of reflections. A study of the acoustic energy propagating into the outer layers of the grid revealed a steep decline, albeit for a wave with relatively short wavelength [4].

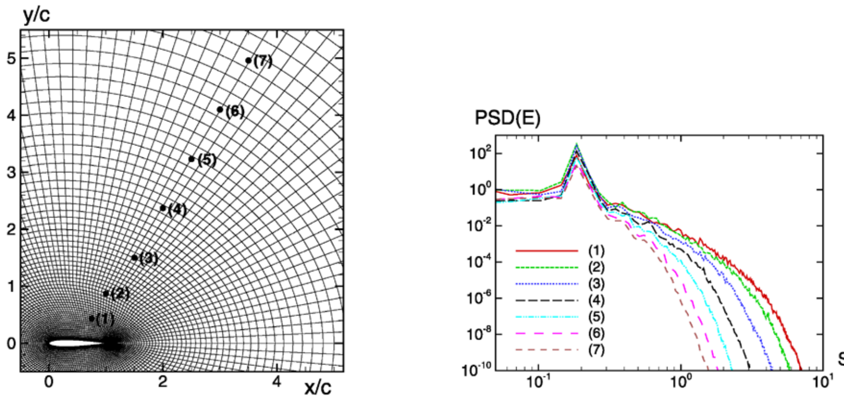

Fig. 3. Airfoil at 60° angle of attack. a) location of points in the wake, and b) corresponding spectra.

Vanella, Piomelli & Balaras carefully addressed this phenomenon, but for step changes in resolution, and that between regions each having a uniform grid [5]. They also considered refinement as well as coarsening; not surprisingly, refinement was trouble-free, as the energy cascade simply extended itself when the grid spacing and accordingly the eddy viscosity decreased. Their position was that, when coarsening, the SGS model ideally should be adjusted to make the energy removal uneventful. In some versions of the model, they used knowledge of the step in Δ a short distance upstream of that change, a similar idea to that of using the derivative of the grid spacing, as proposed here. At first sight, their use of centered differencing schemes left little room for error, but in the end they focused on accuracy and did not report any problems with numerical stability.

Figure 3a shows a line of points in the FR and DR behind a stalled airfoil, at which frequency spectra were calculated. For the present qualitative purposes, Taylor's hypothesis is sufficient to illustrate the behaviour of wave-number spectra, especially since the velocity does not vary widely. Still, flow visualisations show that these points are bathed in turbulence for most of the time, but not all: they spend brief periods in irrotational fluid. This intermittency means that the signals are not homogeneous in time; as a result, even results from very fine grids and high Reynolds numbers may not contain an inertial range. Here, the spectra for the first few points are actually rather consistent with a $-5/3$ slope.

The comparison in Figure 3b displays the gradual "shortening" of the spectrum as the grid coarsens farther from the airfoil, as well as the absence of energy pile-up. In this region the differencing scheme is very close to centered: a weighted average of 4th-order centered and at most 10% of the 5th-order upwind one. From the first to the last point, Δ increases by a ratio of about 7.5, and the frequency of the apparent steep spectral drop drops by a factor of about 5, which is not a perfect fit, but makes the connection plausible. The Nyquist frequency associated with the time step is $St \approx 17$, therefore well above the value where the spectrum drops; therefore, time-integration errors are not a factor in this region (the first spectrum reaches $St = 7$ in the figure, but the spectral density is down by 12 orders of magnitude).

This simulation was continued with a fully centered scheme, and rapidly deteriorated, with spurious short waves invading the wake and expanding into the irrotational region, as seen in figure 4b. This is in spite of implicit time integration. This confirms that the weighted scheme used is close to optimal [6].

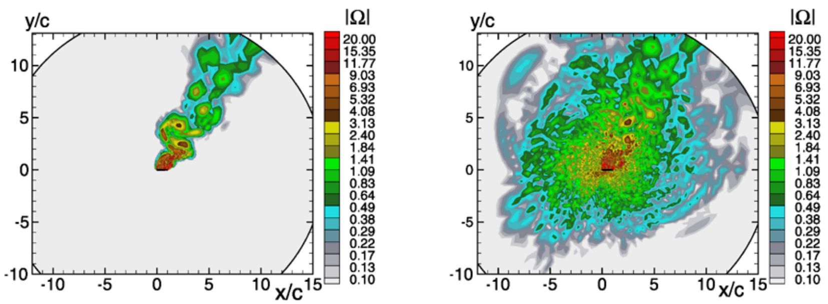

Fig. 4. Vorticity in: a) normal simulation, and b) continuation with centered differencing.

A revealing way to measure numerical dissipation may be as follows. Consider the simple convection of a wave in one dimension, and the effect of upwind differencing on its amplitude after it travels by one wave-length (in other words, after one period). This depends on the resolution, and we pick a resolution with 10 points per wave-length as a compromise. In the past, we have indicated that using 7 points provides useful accuracy with our schemes; on the other hand, some careful studies with other schemes have concluded to the need for 30 points, if not more. Some of these studies were based on very sensitive problems, such as the growth of near-neutral TS waves, and may therefore be unduly pessimistic.

For the first-order upwind scheme, the damping after one period is by 85% in amplitude. It is extreme. For the energy, that is 97.75%, but we will use amplitude. For the second-order upwind scheme, the damping is 31%, which is still intense. The third-order and fifth-order upwind-biased schemes give 12% and 1%, respectively. When a weighted average with any centered scheme is used, the damping can become almost negligible. Using a 10% weight for the fifth-order upwind-biased scheme, which has been common and stable enough for the focus region in our studies, leads to only 0.1% damping (and 0.5% with 7 points per wave). Thus the scheme is not rigorously energy-conserving, but comes extremely close to it for all but the shortest waves. For instance, a wave with only 3 points in it will have 38% damping; thus, the high order of accuracy makes the damping extremely selective.

Returning to the second-order upwind scheme, using 30 points gives a small damping of 1.5%, but comparing with 10, the number of grid points in 3D has increased by a huge ratio. The same second-order scheme with only a 10% weight gives only 3% damping with 10 points per wave. Thus, it is unfortunate that very few codes with the LES option offer the possibility of blending centered and upwind schemes. We recognize the extreme difficulty in obtaining better than second-order accuracy in unstructured codes, but blending two schemes does not appear as challenging. On the other hand, organizing the different blending weights in different regions of the domain is not trivial. Therefore, even to experiment with such an option requires access to the source code, and adds to the complexity of designing the LES study.

This transport into coarser grid has served as a justification for the use of numerical schemes, either monotone or upwind-biased, which do not conserve energy. This refers to the spatial discretisation, and the effects of temporal discretisation are left for future work, although we just remarked that the time step is effectively very short in the DR. It is very common for external-flow turbulence-resolving simulations not to be stable with energy-conserving schemes. The simulations which are stable then take on some of the nature of ILES, if only in some regions. Codes with switches between an upwind-biased scheme in the RR and ER and a much-less biased scheme in the FR and DR are in routine use and very valuable as just dicussed [6]. On the other hand, it would be conceptually closer to perfection to have an energy-conserving scheme, and task the SGS model with the energy removal (the amount removed would therefore be accessible to calculation, which is not the case with numerical dissipation). We will now attempt an estimate of the magnitude of this effect, in an inertial range.

In this concept, in simple terms of energy, the SGS model accounts both for the energy cascade, as usual, and the needed removal of short waves. Ideally, the first role represents interaction with non-resolved waves; it is fairly localized near the high end of the spectrum, but not in a simple way. Ideally, the second role would be very localized, acting only at the extreme end of the spectrum. Applied in Fourier space, it could be a “top-hat” filter, trimming the end of the spectrum. This implies that an eddy viscosity would not be very truthful, since it is far being local in Fourier space.

The SGS model in its first role consumes energy at the rate ε (we are assuming a large cell turbulent Reynolds number, so that the resolved dissipation is negligible). For the second role, let Δ be the grid spacing. If the grid cells are far from isotropic, let it be the larger dimension. The resolved turbulence is responding to a positive value of $D\Delta/Dt$, the Lagrangian derivative of the grid spacing (for waves rather than eddies, the rate of change of Δ at a speed $U \pm c$ would be relevant). We assume that the shape of the end of the resolved spectrum is invariant; in other words, the energy distribution in and beyond the inertial range is $C\varepsilon^{2/3}k^{-5/3}f(k/k_{\max})$ where C is the Kolmogorov constant, $k_{\max}$ is the cut-off wave-number and f a function over $[0,1]$ with $f(x) \to 1$ when $x \to 0$. Its shape is not trivial, because it can be controlled both by spatial and temporal discretisation, not to mention the Smagorinsky or similar constant. Figure 3b does not contradict the idea of an invariant shape, but a larger sample and a firm inertial range would be needed to be precise.

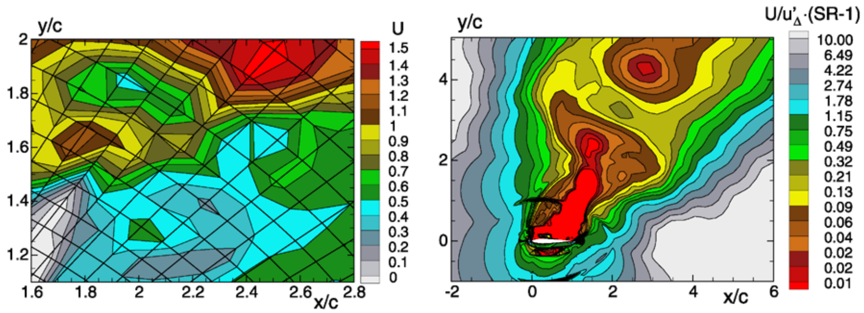

Fig. 5. Flow past an airfoil. a) Velocity contours, showing level of activity within grid cells; b) estimate of coarsening-related energy removal, normalized by true dissipation (equ. 3).

The grid-related energy-removal rate then results from integrating the spectrum up to $k_{\max}$, which is spatially-varying and inversely proportional to Δ so that $D(\log k_{\max})/Dt = -D(\log \Delta)/Dt$. We then have

$$\frac{Dk_{res}}{Dt} = -C' \, \varepsilon^{2/3} \, \Delta^{-1/3} \, \frac{D\Delta}{Dt} \tag{1}$$

where C' is a constant of order 1, related to the Kolmogorov constant and to the LES spectrum shape f. This equation does not give the total derivative of k_{res}; only the rate related to grid spacing. The ratio of the two rates is

$$\frac{1}{\varepsilon} \, \frac{Dk_{res}}{Dt} = C' \, \varepsilon^{-1/3} \, \Delta^{-1/3} \, \frac{D\Delta}{Dt} \tag{2}$$

and will be used to get a sense of the relative intensity of the two phenomena.

We first note that this ratio decreases if the grid is refined by a uniform ratio, because it is proportional to $D(\Delta^{2/3})/Dt$. This is reassuring: the issue vanishes during grid convergence. On the other hand, ε is not easy to estimate. We have $\nu_t S_{ij} S_{ij}$ where ν_{SGS} is the sub-grid viscosity and S_{ij} the strain tensor, but another measure may be more revealing and accessible. The velocity variations over one grid cell, which we will denote by u'_Δ, are proportional to $\varepsilon^{1/3}\Delta^{1/3}$, so that we have for the ratio in (2):

$$\frac{1}{\varepsilon}\,\frac{Dk_{res}}{Dt} = \frac{C''}{u'_\Delta}\,\frac{D\Delta}{Dt} = C''\frac{U}{u'_\Delta}\,(SR-1), \tag{3}$$

where C'' is also a constant of order 1, and SR is the stretching ratio from one cell to the next in the direction of the velocity U. A common value for SR in the wall layer, for the wall-normal spacing, is 1.25, but good grids in the FR often require values below 1.1, if not below 1.05 [7].

The values of u'_Δ, tentatively defined in each cell as $u_{\max} - u_{\min} + v_{\max} - v_{\min} + w_{\max} - w_{\min}$, as not as small as expected; see figure 5a, in which the quantity is often 1/2 of the freestream velocity. As a consequence, the values of $U/u'_\Delta\ (SR-1)$ are generally much below 1, as seen in figure 5b. If u'_Δ is defined as the rms within the cell instead of the $\max - \min$ definition, the values are of course smaller, but still not $\ll U$. If the constant C'' is of order unity as we presumed, this shows that on the present grid, the energy-removal effect is of small magnitude relative to the dissipation, and therefore the need for stabilized differencing schemes may *not* be attributable to it. It should be easily compensated by a moderate turn-up of the end of the spectrum, which raises the dissipation of the resolved field (since both ν_t and $S_{ij}S_{ij}$ increase). This thought is consistent with the findings of Vanella *et al*. This issue may deserve more precise estimates, but the determination of f and C'' requires an inertial range, which is rarely found except in simulations of homogeneous turbulence, which in turn are done with centered differencing and uniform grids.

At this stage, we have not tried and are not proposing a model for this grid-coarsening effect, merely suggesting it for reflection. We consider that an ILES character is very acceptable in the DR, whereas ideally we would prefer to have no numerical dissipation at all in the FR.

5 Explicit Determination of Grid Regions: an Exercise

Automatic grid adaptation will, sooner or later, apply to LES and other turbulence-resolving approaches, like it has with great usefulness to the full-potential equations and is slowly developing for steady Navier-Stokes work. Adaptation shortens the user time and reduces the dependence of the results on the user's skills and patience. It can also improve accuracy for a given computing cost. However, as of today, we know of very few examples of automatic adaptation in DNS or LES [3]. We are only presenting a concept which may be one component of such a system. It relates to the YPG Approach.

The goal is to determine the Departure Region quantitatively. This information can then be used to allow gentle but sustained grid stretching in the DR, while the grid in the FR has a much more uniform target spacing. A plausible definition of the FR is that it contains turbulent eddies which may now or in the future interact with the region of interest. Here, this “region of interest” will be taken to be the immediate vicinity of the airfoil. For a backward-facing step, the analogous region could be the near-wall layer in the recirculation bubble, which is the sensitive region; it could also include a few lines of points across the separated region. The definition depends on the purpose of the study.

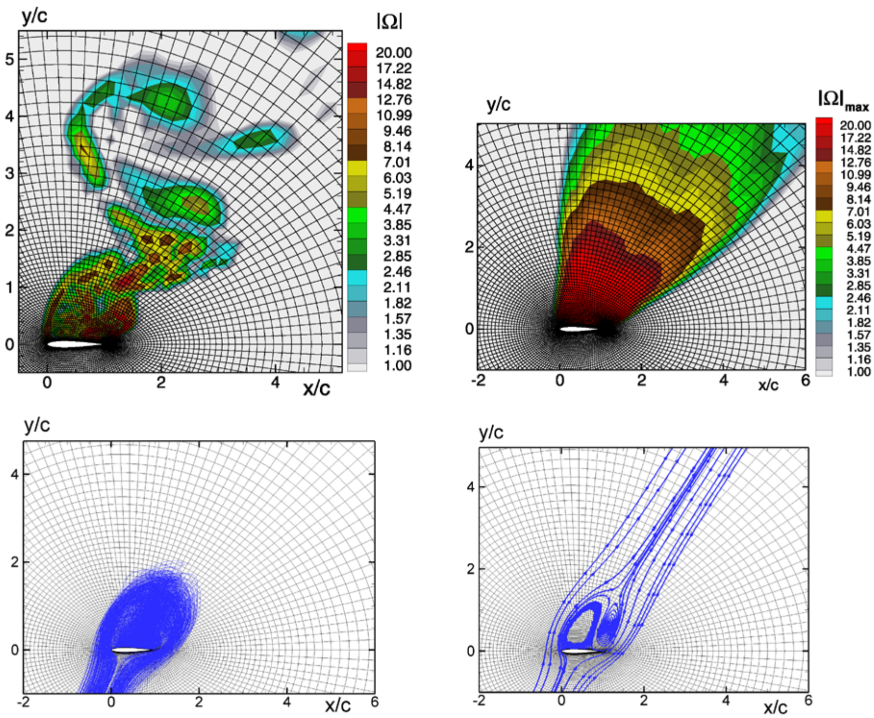

Fig. 6. Flow past an airfoil. a) Instantaneous vorticity; b) maximum vorticity during run; c) reverse trajectories; d) streamlines of the average flow field.

The idea, then, is to calculate reverse trajectories from that region of interest. In simple terms, together these trajectories will outline the region from which particles (eddies) may travel to impact the region of interest. We call this the Contact Region (CR). It is stationary; we are not envisioning time-varying grids yet. It is calculated *after* the simulation and is expensive in terms of storage, because a time-dependent 3D field is stored. As a result, it is practical only on a grid coarser than the one on which the primary simulation is conducted. This will be a precursor simulation. In the present one, the grid is 142*101*35 and the time step $0.03c/U_\infty$. Running such simulations is very good practice in any case, even if the grid generation is manual, and grid generation is rapidly reaching a capability such that little user time will be wasted in doing this (and of course little CPU time, the resolution being coarser in all four directions).

The CR is not a new concept added to the YPG list. It is a mathematical rather than a descriptive definition. Similarly, we are here seeking a mathematical description of the ER. The use of trajectories and of the concept of "contact" appears simplistic, knowing that pressure interactions occur at a distance, but recall that we are after the issue of suppressing small eddies, in other words of displacing vorticity by a small distance. Thus, vorticity is most representative of the physics at play, and it is a transported quantity.

Here, we are considering a flow which is homogeneous in the z direction, which makes it easier to outline the CR without a huge time sample. It is defined in the x-y plane and extruded. In this paper, it is only visual, but a workable definition in a code will be to tag every grid cell which has been visited by a trajectory, and then fill in the likely gaps to arrive at a contiguous area or volume.

Another step in the YPG Approach is to identify the Euler Region. Its elementary definition is simple: the region never visited by non-zero vorticity. For this, the maximum vorticity magnitude $\omega_{\max}$ during the simulation is recorded, and the ER defined by a low threshold for $\omega_{\max}$. Again, here, the max is taken both over t and z. This is illustrated in figures 6a and 6b. The peak vorticity $\omega_{\max}$ behaves as expected, and setting a threshold around $1 \times U_\infty/c$ is effective.

Figure 6c shows the reverse trajectories from a thin layer on top of the airfoil, located at $y = 0.007c$. The starting points are distributed uniformly spanwise, and the trajectories cover a time interval of $35c/U_\infty$.

The reverse trajectories naturally extend to ∞ upstream and therefore enter the Euler Region. Therefore, the Contact Region gives somewhat of an over-estimate of the extent of the Focus Region. In addition, they are seen to form a bundle, over one chord wide, upstream. The reason is engulfment of irrotational fluid by the turbulent region. This makes the CR based on trajectories overlap with the ER; the true FR is smaller than the CR. The indications it provides are still very concrete regarding the shape and extent of the FR: convincingly, it has an oblong shape truncated by the airfoil, and ends near $y/c = 2$.

When treating this flow with the YPG Approach, the RANS Region is easily outlined along the wall, and $\omega_{\max}$ helps pick its thickness on the lower side of the airfoil. On the upper side, the RR blends into the FR, which is the intersection of the CR and the Vortical Region. The DR is then the intersection of the VR and the complement of the CR.

The exercise is convincing visually, which was helped by the homogeneity in z, but simple recipes to use the definitions in an automatic system were given. Naturally, such a system will need many decisions on the grid spacing itself, starting with the target FR spacing Δ_0, the first wall-normal spacing, stretching ratios particularly in the FR and DR, and so on. If using a flexible grid-generation package, the spacing in

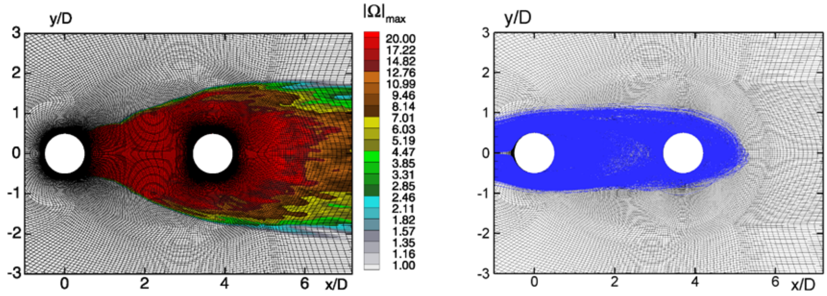

Fig. 7. Flow past tandem cylinders. a) Vortical and Euler regions; b) reverse trajectories.

the ER can be significantly larger than in the VR, including in z. Only the time step is fully shared.

The relationship between the CR and the more established concept of Recirculation Region is studied with figures 6c and 6d. The latter shows the streamlines of the Reynolds-averaged LES flow field. This mean flow being two-dimensional, a closed region or "bubble" dominated by two large circulation regions exists. Its extent is not too different from that of the CR, although it is smaller as could be expected, and the strictly-defined "bubble" (delimited by streamlines which make contact with the body) does not have the upstream "bundle." It also fails to include edge regions which are visited by vorticity and turbulence.

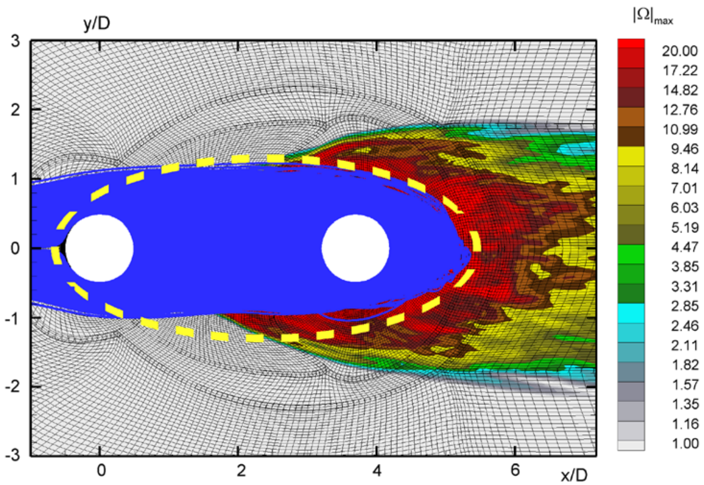

Fig. 8. Flow past tandem cylinders. Vortical region, reverse trajectories, and dashed line tentative Focus Region.

The CR concept also accounts for interference between bodies, as illustrated in figure 7 for tandem cylinders. Both cylinders were declared as Regions of Interest. The second cylinder then "emits" a fairly wide CR which bridges the two bodies. However, the CR is narrower than the VR, which is quite wide past the second cylinder; this will help save on the number of grid points. Note that the mean streamlines as in figure 6d would fail to bridge the two cylinders; thus the CR concept is demonstrating its potential, and plausibility relative to intuition.

Some day, simulations of an airplane will be aimed at whether the turbulence created by a wing spoiler causes buffet on the horizontal tail, or the wake of the nose landing gear disturbs an air inlet or ram-air turbine, and it is possible the present CR concept will be helpful. Other examples would be multiple road vehicles.

6 Outlook

We have provided material which has some preliminary and inconclusive aspects, but we hope is somewhat stimulating. We are fairly confident in this for the QLES community, but not as much about the wider turbulence/CFD community. It appears that only about 100 copies of the bound proceedings are sold outside of the QLES members, which is small. Naturally, this does not account for the "soft" copies of articles disseminated by internet, but presumably these reach only readers which are already connected to a QLES contributor. A valid question is how best to boost the quality of the turbulent CFD in the widest circles we can.

Journal articles derived from the work of QLES contributors must have a readership much larger than 100. Personal exchanges at QLES meeting are also often substantial and productive, thus contributing to research quality. Yet, informal surveys suggest that much of the "turbulence knowledge" CFD users outside academia apply actually travels with the CFD codes they use rather than literature they read. The users extensively consult the support staff of the CFD vendors or the government scientists who write and maintain the codes, for instance at DLR and NASA. The commercial companies have strong incentives to point at the ease of use of their products, which include the turbulence treatments. This creates a danger of overselling these treatments, which we all know to be shaky in many respects, turbulence being "the chief outstanding difficulty of our subject" according to Bradshaw's 1994 article, still fully applicable today [8]. One task of the turbulence research specialists is to decrease the incidence of poor "plug and play" practices, without discouraging the use of CFD or even the use of cutting-edge approaches.

Publishing accessible overview articles, with impartial and candid discussions of the state-of-the-art, is not so easy and is a good path. Contributing to the documentation which travels with the codes could be another. We find that most CFD companies have retained and grown some very competent turbulent-CFD experts, blending genuine scientific attitudes with commercial urges. The same is true for government code suppliers. QLES-type specialists should use every opportunity to collaborate with them, and possibly to directly contribute to the documentation they attach to their products. This appropriately goes beyond the intellectual pleasure and the prestige of publishing at high level, to educate “in real time” the ever-widening ranks of turbulent CFD practitioners.

References

1. Spalart P R (2001) “Young person’s guide to Detached-Eddy Simulation grids.” NASA CR-2001-211032. http://techreports.larc.nasa.gov/ltrs/PDF/2001/cr/NASA-2001-cr211032.pdf
2. Spalart P R, Hedges L S, Shur M, Travin A (2003) Simulation of active flow control on a stalled airfoil. Flow, Turb. and Comb., 361–373.
3. Hoffman J (2006) Adaptive simulation of the subcritical flow past a sphere. J. Fluid Mechanics Vol. 568, 77–88.
4. Spalart P R, Shur M L, Strelets M Kh (2009) Added sound sources in jets; theory and simulation. Int. J. Aeroacoustics, to appear.
5. Vanella M, Piomelli U, Balaras E (2008) Effect of local grid refinement on large-eddy simulation statistics and flow-fields. ETMM7.
6. Travin A, Shur M L, Strelets M Kh, Spalart P R Physical and Numerical Upgrades in the Detached-Eddy Simulation of Complex Turbulent Flows. In: Fluid Mechanics and Its Applications: Advances in LES of Complex Flows, pp.239–254. Proceedings of the Euromech Colloquium 412. R. Friedrich and W. Rodi, (eds.).
7. Shur M L, Spalart P R, Strelets M (2005) Noise prediction for increasingly complex jets, Part I: Methods and tests. International J. of Aeroacoustics, 4(3+4), 213–246.
8. Bradshaw P (1994) Turbulence: the chief outstanding difficulty of our subject. Expts. in Fluids, 16(3–4), 203–216.