



Rendiconti
Accademia Nazionale delle Scienze detta dei XL
Memorie di Scienze Fisiche e Naturali
122° (2004), Vol. XXVIII, pp. 315-329

C.M. CASCIOLA - F. SALVADORE - R. PIVA*

Large scale simulations in turbulence and combustion

Introduction

Fluid dynamics is one of the prominent fields where intensive computations have been systematically developed and employed. A simple reason is given by the relevance of gases and liquids for a number of applications. However the spreading of fluid motion in different areas of science and technology alone cannot explain the role of large-scale computations in the fluid-dynamics community. The real explanation lays in the huge variety of non-linear phenomenologies which are generated by the relatively simple mathematical model given by the Navier-Stokes equations. Such complexity is even increased in presence of an active micro-structure, e.g. in the case of dilute solutions of long-chain polymers, micro-bubbles, solid particles or even reactive scalars as can be found in combustion problems.

In these conditions one of the few feasible ways of understanding the basic physics of the fluid motion is provided by numerical simulations, which may reach and easily exceed the limits of available computational power. Hence the need for efficient algorithms for large-scale computations and the adoption of fast, massively parallel computers, either in shared or distributed memory architecture. A real challenge is provided by cooperative or grid computing, where different huge machines can be combined through an external net to contribute different parts of a single simulation.

Purpose of the present paper is to discuss the main difficulties associated with large scale simulations in the fluid dynamics of complex fluids. Here we shall focus on combustion, where chemical reactions, diffusion processes and fluid-turbulence lead to a very challenging problem. In perspective, our aim would be to achieve the direct numerical simulation (DNS) of the propagation of a simple turbulent flame.

* Dipartimento di Meccanica e Aeronautica, Università degli Studi di Roma.

A brief presentation of the mathematical model is provided first. Then an efficient algorithm is described, based on a high order discretization of the relevant system of partial differential equations. Finally an outlook is given to the basic ideas for the extension of the algorithm in the framework of grid computing.

Mathematical Model

The basic set of equations for fluid dynamics is provided by the Navier-Stokes equation in three space dimensions plus time, representing the conservation of mass, momentum and energy in a continuum system described in terms of mass density, the scalar field $\varrho(x, t)$, velocity, the vector field $V(x, t)$, and energy, the scalar $E(x, t)$,

$$\begin{aligned}\frac{\partial \varrho}{\partial t} + \nabla \cdot (\varrho V) &= 0 \\ \frac{\partial \varrho V}{\partial t} + \nabla \cdot (\varrho V \otimes V + pI) &= \frac{1}{\text{Re}} \nabla \cdot \Sigma \\ \frac{\partial \varrho E}{\partial t} + \nabla \cdot (\varrho V E) &= \frac{1}{\text{Re}} \nabla \cdot (\Sigma \cdot V) - \frac{1}{\text{Pr}} \nabla \cdot q.\end{aligned}$$

Certain constitutive equations must be included to close the system and to introduce certain phenomenological laws observed in fluids at a macroscopic level. Typically, the tension exchanged by neighboring particles is described in terms of the stress tensor $T = -pI + 1/\text{Re} \Sigma$, where $p(x, t)$ is the pressure field and $\Sigma(x, t) = \mu/\mu_0(\nabla V + \nabla V^T) + \lambda/\mu_0(\nabla \cdot V)I$ is the viscous stress in a simple Newtonian, compressible fluid – here I is the identity, and μ and λ are the so-called first and second viscosity coefficient with μ_0 their common characteristic value. Viscosity coefficients are generally expressed as functions strongly dependent on temperature – $\vartheta(x, t)$ – and weakly on pressure. A constitutive equation for the heat flux, the vector $q(x, t)$, should also be provided, e.g. in its simplest form as the Fourier law expressing the proportionality between heat flux and temperature gradient, $q = -k/k_0 \nabla \vartheta$.

Equipped with these constitutive laws, the above set of partial differential equations forms the Navier-Stokes system which is taken to describe the motion of a compressible fluid. For example, the fluid could be a gas obeying the perfect gas law relating pressure, density and temperature, $p = \varrho R/R_0 \vartheta$, where R is the gas constant. An additional equation of state must be specified to link the internal energy – the total energy density E is the sum of kinetic $1/2 V^2$ and internal energy e – to density and temperature, e.g., for a perfect gas, $e = c_v \vartheta$, where c_v is the specific heat at constant volume, here assumed to be constant, for the sake of simplicity. Written in dimensionless form, the Navier-Stokes equations for a non-reacting gas are characterized by two dimensionless parameters, the Reynolds number




Figure 1: Two kinds of layers in a flow domain. On the left, a visualization of the velocity profile in the boundary layer of an incompressible fluid over a flat plate. On the right, interior layers corresponding to shock waves generated by a supersonic blunt body.

$Re = \rho_0 U_0 L_0 / \mu_0$ and the Prandtl number $Pr = \mu_0 c_p / k_0$, where a subscript 0 denotes the proper characteristic quantity and $c_p - c_v = R$.

Even in this simple setting, solving these equations is a huge computational challenge for several reasons. At large Reynolds number the flow field is characterized by thin layers [1]. They can occur at the boundary of the domain as boundary layers (see figure 1.a) or in the interior, as in the case of shock waves or mixing layers (see figure 1.b). After Prandtl, the reason for the presence of such layers is clear. Increasing the Reynolds number, viscous effects will become less and less significant and most of the flow domain will experience an almost inviscid dynamics. However, locally viscous effects must keep their dynamical significance, e.g. near a solid wall they should force the relative velocity between fluid and wall to zero. This implies that, in these critical regions, the gradients should increase in proportion to keep the diffusive terms (in their dimensionless form)

$$\frac{1}{Re} \nabla \cdot \Sigma \approx \frac{1}{Re} \frac{\partial^2 V}{\partial n^2} \quad (1)$$

order one, where n denotes the normal to the layer. The fundamental point here is that the entire flow field is controlled by the dynamics of these thin layers. If a numerical simulation fails to capture these regions of strong gradient, the entire flow field is irremediably inaccurate and typically entirely wrong. To get an intuitive feeling, one can just think of the drag of a stream-lined body: it depends on the velocity gradient right at the wall, which is determined by the corresponding boundary layer. It is then clear that the momentum exchange between external forcing – the moving body – and fluid is entirely controlled by the boundary layer. The force the fluid will experience in the simulation depends on the ability to reproduce this thin sheet close to the wall. Even in laminar flows, this need of spatial resolution may translate in a huge number of degrees of freedom, and, propor-

tionally, in a significant demand of computational power, i.e. memory allocation and operation count.

Things can easily get worse in presence of a micro-structure. For example we can have flows with polymers, with solid particles, with micro-bubbles where we can be interested in drag reduction. Similarly, a micro-structure characterizes reacting flows and combustion problems. In this case we have to consider additional equations to take into account the different chemical species [2]

$$\frac{\partial \rho Y_\alpha}{\partial t} + \nabla \cdot (\rho V Y_\alpha) = \frac{1}{\text{ReSc}_\alpha} \nabla \cdot J_\alpha + \Omega_\alpha \quad (2)$$

where Y_α is the mass fraction, J_α the diffusive flux and Ω_α the source term of α^{th} specie. Clearly we have to introduce constitutive laws for the mass flux and the source term. When combustion occurs in presence of an abundant neutral specie, such as N_2 for combustion in air, the mass fluxes can be specified by the Fick law $J_\alpha = \chi_\alpha / \chi_{\alpha 0} \nabla Y_\alpha$, where χ_α is the binary diffusion coefficient of specie α into the abundant one. The source term depends on the reaction kinetics scheme we choose to simulate. In particular, considering a reaction mechanism characterized by N_R reactions between N_S species, a possible form is

$$\Omega_\alpha = P_\alpha \sum_{k=1}^{N_R} \nu_{k,\alpha} Da_k \omega_k \quad (3)$$

being P_α the molecular weight of the α^{th} specie and $\nu_{k,\alpha}$ the stoichiometric coefficient of the specie in the k^{th} reaction. Single reaction source terms ω_k can be modeled in the classical Arrhenius form, with an exponential term $\exp(-\vartheta_A/\vartheta)$, where ϑ_A is the activation temperature, and a prefactor depending on density and concentrations whose specific form depends on the particular system of reactions one is considering. The new parameters appearing in the system are the Schmidt number of specie α , $\text{Sc}_\alpha = \nu_0 / \chi_{\alpha 0}$ and the Damkohler numbers Da_k which represent the ratio between fluid dynamics characteristic time and the chemical time scale of the considered reaction. Necessarily, the constitutive equation for the internal energy must include the composition of the mixture. A simple one is given by $e = c_v(\vartheta - \vartheta_R) + \sum_{\alpha=1}^{N_S} \Delta h_\alpha Y_\alpha$, where Δh_α is the amount of heat released (either positive or negative) in the production of the unit mass of specie α at the reference temperature ϑ_R .

Simulations for turbulent combustion

To understand the main requirements of a numerical simulation able to catch the relevant physical aspects, let us consider the evolution of a flame in the com-




Figure 2: Visualizations of the flame front in a spark ignition internal combustion engine. The front is seen from the transparent ceiling of the piston. The different phases are characterized by the angle reported in each panel, which is proportional to the time elapsed from the ignition.

bustion chamber of a spark-ignition internal-combustion engine, figure 2. The reactive structure and its shape at four different instants is clearly visible. Reactions implies a very wide range of characteristic time scales. Moreover, the flame is usually a really thin sheet where all the processes related to the intense heat release take place. As a consequence, numerical simulations are forced to use strong constraints for temporal and spatial resolution. Quantitatively, a significant relationship between flame thickness, δ_F , and combustion speed, U_F , i.e. the average speed of the fresh mixture entering the combustion region is $U_F \delta_F/\nu_0 \approx 4$. This is interpreted by saying that the Reynolds number of the flame is order one. Considered that in a typical combustor $Re = U_F L/\nu_0 \approx 10^5$, we understand that the flame thickness is 10^5 smaller than the characteristic dimension of the apparatus. This can give a rough idea of the spatial resolution needed to follow the evolution of a flame sheet. This is not the only difficulty of this kind of flows, however.

In fact, most applications of interest deal with turbulent flows. The technological reason is that the combustion rate is largely increased by turbulence. Turbulent flows are characterized by a continuous range of spatial and temporal scales. As an example, figure 3 gives the distribution of turbulent kinetic energy among the different scales of the flow. Technically, the quantity reported in the figure for different experimental conditions is the one-dimensional energy spectrum, directly related to the Fourier transform of the correlation function of the fluctuating velocity field. Defining the fluctuating field as the difference between its actual value and its average, $v = V - \langle V \rangle$, the turbulent kinetic energy is $K = 1/2 \langle v^2 \rangle$. The energy spectrum $E(k)$ is such that

$$K = 1/2 \int_0^\infty E(k) dk,$$




Figure 3: Energy spectrum for a nearly homogeneous turbulent shear flow. Increasing the Reynolds number, the minus five-thirds slope predicted by Kolmogorov emerges clearly.

and describes how the energy associated to the fluctuations is distributed among the different scales $\ell = 2\pi/k$. According to Kolmogorov theory [3], the extension of the range of scales involved in a turbulent motion is given by the ratio between the integral scale L (of the order of the domain size) and the Kolmogorov scale $\eta = (\nu^3/\varepsilon)^{1/4}$, where ε is the average dissipation rate per unit mass associated to the fluctuating field. It is easy to understand that such ratio may become exceedingly large also in ordinary conditions. For instance, in a 1.5-liter household food mixer filled with water – $\nu \approx 1. \times 10^{-6} \text{ m}^2\text{s}^{-1}$ – and driven by a 300w electric motor η is so small as $.8 \times 10^{-2} \text{ mm}$, corresponding to $L/\eta \approx 10^4$. The same reasoning leads to the conclusion that the number of degrees of freedom active in a three-dimensional turbulent system, $N \approx (L/\eta)^3$, is of the order of $\text{Re}^{9/4}$. Actually, most applications are characterized by very large Reynolds numbers of the order $10^8 \div 10^{10}$ for an airplane or a ship, 10^5 for a combustion chamber.

A gas turbine combustion chamber is sketched in figure 4, which shows the complex geometry characterized by main and secondary air inlets and fuel injectors. The structure of the large scale flow field is crucial to control the position of the flame. Small scale turbulence instead controls the mixing, the detailed dynamics of the flame, and, overall, the combustion rate. The general conclusion is that the solution of combustion problems requires to follow a wide range of spatial and temporal scale in complex geometries, and systematically leads to large scale computations. Different approaches may be taken to limit such computational demand. These techniques manage to work in a broad range of parameters through the introduction of different kinds of approximations.




Figure 4: Sketch of the geometry of a gas turbine combustion chamber. Red and blue arrows denote main and secondary (dilution) airflow, respectively.

Direct Numerical Simulation (DNS) consists in the numerical solution of the exact Navier-Stokes equations. The present generation of supercomputers can nowadays simulate turbulent flows at $Re = 10^3 \div 10^4$. This approach, when coupled with high order of accuracy numerical methods in simple geometries, leads to a virtually exact solution of the Navier-Stokes equations.

Large Eddy Simulation (LES) can work at Reynolds number up to $10^4 \div 10^6$. In this case, the evolution of the large scales of the flow – the so-called large eddies which are expected to be strongly dependent on the specific flow geometry – are solved directly, accounting for the effect of the smaller, unresolved scales by means of suitable models which, in principle, are based on universal features of the fine scales of turbulence. LES is a challenging task for present industrial applications, particularly in the context of combustion problems.

For larger Reynolds numbers, one must resort to the Reynolds Averaged Navier-Stokes equations (RANS) to reach values of Re of the order of $10^6 \div 10^{11}$. This is the most common approach for industrial computations. Only mean quantities can be evaluated and problems in complex geometry can be handled routinely. As a major drawback, this technique has very small or no predictive power, since all the effect of turbulence is demanded to closure models which need to be validated for the specific problem at hand.

If the interest is on the physics of turbulent combustion it is mandatory to consider a substantial simplification of the geometry. Nonetheless, toplevel computational facilities are needed and specific numerical algorithms must be selected in view of the particular computer architecture. Under this respect, present-day supercomputers are based on a huge number of parallel processors, collected in nodes of typically eight-sixteen units. The available memory is distributed between the nodes, i.e. it is shared between processors within the same node, while suitable buses take care of inter-node communications.

Typical algorithms working on this kind of architectures are based on a topological decomposition of the flow domain, with the solution in each subdomain assigned to a specific node. Information should be exchanged between nodes devoted to neighboring sub-domains along the common boundary. Given the fast communication offered by the large-band inter-node buses, the exchange of boundary data, most frequently realized using the MPI (message passing instructions) protocol, does not penalize the efficiency too much. This makes these architectures suitable for the solution of large elliptic problems where each sub-domain is strongly influenced by all the others. Typically, in fluid-dynamics, computationally dominant elliptic features are found in the incompressible Navier-Stokes equations.

Recently an alternative approach is emerging, the so called Grid Computing where a certain number of different machines are connected via an external net. This approach forces to deal with a reduced communication efficiency, and with a certain degree of unpredictability on the data rate of the available transmission line. In principle, the advantage is that different super-computers may be called to cooperate into a single huge computation. The basic strategy could still be based on a domain decomposition, however the data exchange should be considerably reduced to cope with the limited data rate. This rules out elliptic systems, but leaves open the possibility to deal with problems dominated by hyperbolic features. In this case, the finite propagation speed reduces the coupling between sub-domains hopefully allowing for efficient implementations on geographically distributed machines.

Propagation of a flame-front

In order to study the physics of the flame evolution let us consider the geometrically simple problem of the propagation of a nominally plane turbulent flame. The corresponding experimental configuration may consist of a tube filled with fresh mixture of reactants where chemical reactions can be activated by a spark at a certain location. A flame front is thus generated and propagates along the axis of the tube. Depending on the terminal configuration of the tube we can have two different regimes. They are known in literature as *deflagration* and *detonation* regime, respectively. In the deflagration regime the combustion is sustained by molecular heat conduction that gives rise to the heating of the fresh mixture ahead of the flame front. In these conditions the propagation velocity and the Mach number are small. In the detonation regime heating of mixture is due to shock compression and the Mach number of the flame propagation is equal to one. We address here the deflagration regime which is found in the most common applications. In this case the low combustion Mach numbers imply a substantially incompressible dynamics of the near-flame region. Nonetheless, acoustic waves are present, and they influence combustion rate and flame dynamics to a great extent. In these conditions stringent limitations are intrinsically introduced in numerical methods by the relatively large sound speed, which is much faster than any other propagation velocity in the system.

The laminar deflagration regime is characterized by a propagation velocity s_L and a front thickness δ_L that can be evaluated numerically with good accuracy [4]. In turbulent conditions, a strong interaction between reactions and turbulence occurs. As long as the turbulence intensity, defined as $\sqrt{\langle v^2 \rangle}/U_F$, is low, the flame front retains, on average, its one dimensional propagation. The corresponding geometry is described in figure 5. From a numerical point of view, periodicity along the two homogeneous directions in the cross-flow plane normal to the average flame propagation velocity can be assumed. Concerning the mean flow direction, at the inflow boundary the fresh mixture must enter the domain with an average velocity corresponding to the average combustion rate. At the outflow, the burnt mixture must leave the domain with no spurious reflections from the boundary.

A fully three-dimensional simulation of turbulent flame propagation needs a certain amount of preliminary calculations. A good procedure is to start with initial conditions given in terms of a one-dimensional laminar flame with fresh-mixture conditions – e.g. density, temperature, composition – identical to the average values to be enforced to the turbulent simulation. This is a computational task per se, and implies the solution of the non-linear eigenvalue problem defining the flame – the eigenfunction – and the propagation speed – the eigenvalue. An additional preliminary three-dimensional simulation is then required to drive the fresh gas to an equilibrium turbulent state. This is achieved by stirring the fluid in a triply-periodic domain, by means of a direct numerical simulation evolving the cold, compressible Navier-Stokes equations. The actual initial conditions for the flame are then constructed as the superposition of the one-dimensional laminar flame and the turbulent fluctuations so prepared.

Concerning inlet and outlet, the propagative nature of the solution is exploited to generate appropriate inflow/outflow boundary conditions [5]. In particular, the algorithm we describe below is based on the conservative form of the Navier-Stokes equations

$$\frac{\partial U}{\partial t} = - \left(\frac{\partial F_z}{\partial z} + \frac{\partial F_x}{\partial x} + \frac{\partial F_y}{\partial y} \right) + \Omega , \quad (4)$$

where the conserved variables are

$$U = (\varrho, \varrho w, \varrho E, \varrho u, \varrho v, \varrho Y_a) , \quad (5)$$

with the corresponding fluxes

$$\begin{aligned} F_x &= (\varrho u, \varrho w u, \varrho E u, \varrho u^2 + p, \varrho v u, \varrho Y_a u) \\ F_y &= (\varrho v, \varrho w v, \varrho E v, \varrho u v, \varrho v^2 + p, \varrho Y_a v) \\ F_z &= (\varrho w, \varrho w^2 + p, \varrho E w, \varrho u w, \varrho v w, \varrho Y_a w) . \end{aligned} \quad (6)$$




Figure 5: Sketch of the geometry for the propagation of a turbulent flame.

The source term includes the chemical sources Ω_α and the divergence of the diffusive components of the fluxes,

$$\Omega = (0, \nabla \cdot \sigma_z, \nabla \cdot \sigma_E, \nabla \cdot \sigma_x, \nabla \cdot \sigma_y, \Omega_\alpha + \nabla \cdot \sigma_\alpha), \quad (7)$$

where $\sigma_{x/y/z}$ are the diffusive fluxes of the three components of momentum, σ_E is related to the energy (heat flux) and σ_α describes the diffusion of specie α .

To discuss the boundary conditions, it is however easier to consider the quasi-linear form of the equations,

$$\frac{\partial U}{\partial t} + A_z \frac{\partial U}{\partial z} = T, \quad (8)$$

where the attention is concentrated on the propagation in the z -direction and $A_z = dF_z/dU$ is the Jacobian of the corresponding component of the flux. All the other terms are grouped together in T . Decomposing the matrix A_z in terms of right/left eigenvectors – R and L – and eigenvalues – Λ – i.e. $A_z = R\Lambda L$, the increments of the Riemann-like variables are $\delta W = L\delta U$.

This set of intrinsic variables highlights the propagative behavior of the system

$$\frac{\delta W}{\delta t} + \Lambda \frac{\delta W}{\delta z} = LT, \quad (9)$$

since the direction of propagation is determined by the sign of the corresponding eigenvalue. Suitable projectors, $\Pi^\pm = R\Lambda^\pm/|\Lambda|L$, can then be used to select waves propagating in the different directions, $\delta U^\pm = \Pi^\pm \delta U$.

In this framework the boundary conditions are easily enforced. At the inlet plane we can assign a number of independent data corresponding to the number of waves entering the domain, e.g. the number of positive eigenvalues. The remaining information must be provided by the evolution equation itself. All this can be combined into a single evolution equation,

$$\frac{\partial U}{\partial t} = \Pi^- \left(T - \frac{\partial F_z}{\partial z} \right) + \Pi^+ \left(\dot{Q}_{\text{in}} \right), \quad (10)$$

where \dot{Q}_{in} stands for the time derivative of a complete collection of boundary conditions, written in terms of the conservative variables U , among which the system will be able to select by itself the necessary data via the projector Π^+ .

At the outlet, by enforcing no reflections from the end of domain, the same equation reduces to

$$\frac{\partial U}{\partial t} = \Pi^+ \left(T - \frac{\partial F_z}{\partial z} \right). \quad (11)$$

It is interesting to observe that in the interior of the domain the original equation may be formally rewritten as

$$\frac{\partial U}{\partial t} = T - \frac{\partial F_z}{\partial z} \quad (12)$$

to show how the inclusion of the boundary conditions can be achieved with a minor impact on the overall structure of the solution algorithm.

It should be emphasized that, despite the fact that only the z -component of the convective flux F_z appears explicitly, the three equations above actually coincide with the original Navier-Stokes for a compressible reacting mixture. The only modification concern the inclusion of the boundary conditions.

Given the periodicity in the cross-flow plane, the discretization in these directions is done with a classical truncated Fourier expansion – spectral discretization –

$$q(x_n, y_m, z_i) = \sum_{r=-N_x/2}^{N_x/2-1} \sum_{s=-N_y/2}^{N_y/2-1} \hat{q}(r, s, z_i) e^{j \left(n \frac{2\pi}{N_x} r + m \frac{2\pi}{N_y} s \right)}, \quad (13)$$

where (x_n, y_m, z_i) is a generic node on the three-dimensional, uniform grid used to discretized the flow domain. More precisely, $\hat{q}(r, s, z_i)$ is the double discrete Fourier transform of the sequence of nodal values $q(x_n, y_m, z_i)$ of a generic quantity q for given i . As main advantages of this approach, differentiation can be carried out exactly in Fourier space, and direct and inverse transforms can be efficiently evaluated by means of fast Fourier algorithms [6].

Concerning the streamwise direction, compact finite difference schemes are used for their favorable dispersion properties, which allow propagation to occur without substantial phase alterations [7]. The resulting algebraic system

$$N \left\{ \frac{\partial f}{\partial z} \right\} = M\{f\} , \quad (14)$$

has to be inverted to obtain the derivatives from the nodal values. Its banded nature allows for its efficient inversion, resulting in high accuracy algorithms of the same efficiency of pure pseudo-spectral methods.

For time advancement, a fourth-order Runge-Kutta method is used. Beyond its accuracy, its peculiar stability domain allows for the integration of strongly convection-dominated systems with no artificial viscosity.

As an example, figure 6 shows a snapshot of the field obtained from a relatively small simulation. The fresh mixture enters from the bottom of the computational box. The sharp front separating fresh mixture, below, from hot combustion products, above, is clearly apparent. The flame is wrinkled through the interaction with the incoming turbulence and its extension increases with respect to the one-dimensional laminar case. The resulting increase in surface area leads to an increased combustion rate. The quantities reported in the figure are the density (left) and the axial velocity (right). The incoming turbulent fluctuations are apparent below the flame front. Above the flame, fluctuations tend to be reduced, due to the thermal expansion of the gas undergoing combustion. The interest for this kind of simulations is the evaluation of the combustion rate for given inflow thermo-chemical conditions as a function of the inflow level of turbulence intensity. In fact the detailed dynamics of the flame may help understanding elusive phenomena such as local quenching of the flame as well as combustion instabilities related to acoustic interactions.

The algorithm, just outlined in its main aspects, can be efficiently implemented on distributed memory architectures. The best approach is to decompose the flow domain in sub-domains along the z -axis – see figure 7. In this way, when each sub-domain is assigned to a node, the Fourier transforms in the cross-flow directions are evaluated locally and no exchange of information between different nodes is needed. In the mean-flow direction the overall discretization maintains its general structure, with suitably defined boundary equations for each sub-domain. More specifically, the interface between neighboring domains is dealt with using a

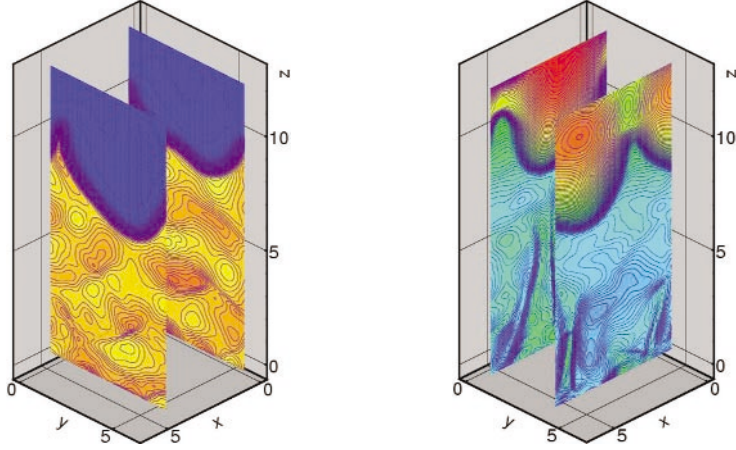


Figure 6: Results from simulations: slices with isosurfaces of density (left) and axial velocity (right).

slight generalization of the boundary treatment described before. The nodal points of the interface are duplicated, as part of either the sub-domain to the left and of that to the right, with the same duplicated evolution equation used for both right and left sub-domains. After denoting with a subscript L (R) quantities pertaining to the left (right) domain, the equation at the interface reads

$$\left[\frac{\partial U}{\partial U} \right]_{\text{IN}} = \left[\Pi^- \left(T - \frac{\partial F_z}{\partial z} \right) \right]_{\text{R}} + \left[\Pi^+ \left(T - \frac{\partial F_z}{\partial z} \right) \right]_{\text{L}}, \quad (15)$$

where the two terms in the right hand side are evaluated locally in the right and left sub-domain, respectively. In this way the forcing term is given by a combination of left and right propagating waves originating in the right and left sub-domain, respectively, to take advantage of the nearly hyperbolic nature of the solution. This is easily done in the context of the MPI protocol, for instance.

Perspectives for grid computing

The direct numerical simulation of turbulent flame propagation is particularly well suited to cooperative computing, whereby different computers – the nodes of the grid – connected through an external net are combined to perform a single task. Actually, the algorithm presented in the previous section can be straightforwardly generalized to this context. Since the flow domain is decomposed in a set of sub-domains, each machine can easily take care of a subset. The algorithm needs boundary conditions at the very leftmost and rightmost sections of the partial

domain handled locally. The corresponding data will be received through the external net from the two nodes of the grid which are providing the solution in neighboring portions of the flow domain. Everything else will proceed exactly as already specified, except for the necessity of embedding the code within an environment able to deal with grid applications.

One of the issues is that of load balancing. Actually, the different computers will have variable performances, depending either on the different hardware or on the different local load. In any case, at crucial points, typically when an event of data exchange takes place, they must be more or less strictly synchronized, with the fastest in completing its task waiting for the slowest. Hence, in order to preserve the efficiency of the simulation, each computer should reach the critical point almost simultaneously with the others. Achieving this result may be difficult, in general. However, it is not hard in the present context, since the operation count of the algorithm is strictly linear in the length of the flow domain. Consequently, the static balance can be straightforwardly designed before hand to achieve a fine tuning of the whole system. More complex would be the dynamical re-balancing of the load in response to variations of the status of the grid due to external disturbances, say new users logging in the system. Dynamic balancing could be achieved, in principle, but it would require the redistribution of the sub-domains across the nodes and would imply a substantial amount of data exchange through the net. This would lead, presumably, to a significant deterioration of the overall performances.

If load balancing can be managed, on the contrary the latency of the external net is a real difficulty also for the present application. A possibility is the adoption of the so-called deferred step technique. This approach consists in using interface data with a certain delay – of one time step, say. This will alleviate substantially the dramatic effect of fluctuations of the transmission rate. However, the naive version of the deferred step technique will spoil the precision of the solution. Improvements can however be devised to preserve the accuracy of the algorithm. As shown, updating the interface requires two contributions to the right-hand size, from the




Figure 7: Domain decomposition for the nearly one-dimensional combustion problem.

right and from the left, respectively. Assuming for the sake of definiteness to deal with the domain to the left of the interface, where the left contribution is available while that from the right is missing, a sufficiently accurate guess could follow by high order extrapolation from known values at previous time steps. The guess could be used instead of the missing true contribution to advance the solution in the left domain. A similar procedure could be used in the right domain, allowing the system to march without waiting for completing the communications. In this way data transfer can take place asynchronously, with only minor additional work needed to correct the extrapolation with computed values once the data become available. With this and other possible devices the issue of communication latency can in principle be overcome, opening the possibility of exploiting huge simulations on a system of geographically distributed high end machines to address basic issues in combustion theory.

REFERENCES

- [1] Batchelor G.K., *An introduction to fluid dynamics*. Cambridge University Press, 1967.
- [2] Williams F.A., *Combustion Theory*. Addison-Wesley, Massachusetts, 1965.
- [3] Kolmogorov A.N., *The local structure of turbulence in incompressible viscous fluid for very large Reynolds number*. *Dokl. Akad. Nauk. SSSR* **30**, 4 (1941). (Translation by Levin V.)
- [4] Zeldovich Ya.B., Barenblatt G.I., Librovich V.B., Makhviladze G.M., *The mathematical theory of combustion and explosions*. Consultants Bureau, New York and London, 1985.
- [5] Poinot T.J., Lele S.K., *Boundary Conditions for Direct Simulations of Compressible Viscous Flows*. *Journal of Computational Physics* **101**, 104 (1992).
- [6] Gottlieb D., Orzag S., *Numerical analysis of spectral methods: theory and applications*. Society for Industrial and Applied Mathematics, Philadelphia, Pa., **26**, CBMS-NSF Regional Conference Series in Applied Mathematics (1977).
- [7] Lele S.K., *Compact Finite Difference Schemes with Spectral-like Resolution*. *Journal of Computational Physics* **103**, 16 (1992).