High performance parallel computing of flows in complex geometries: II. Applications

COMPUTATIONAL SCIENCE&DISCOVERY

High performance parallel computing of flows in complex geometries: II. Applications

N Gourdain[1], L Gicquel[1], G Staffelbach[1], O Vermorel[1], F Duchaine[1], J-F Boussuge[1] and T Poinsot[2]

[1] Computational Fluid Dynamics Team, CERFACS, Toulouse, 31057, France
[2] Institut de Mécanique des Fluides de Toulouse, Toulouse, 31400, France
E-mail: Nicolas.gourdain@cerfacs.fr

Received 2 July 2009, in final form 21 October 2009
Published 12 November 2009
Computational Science & Discovery **2** (2009) 015004 (28pp)
doi:10.1088/1749-4699/2/1/015004

Abstract. Present regulations in terms of pollutant emissions, noise and economical constraints, require new approaches and designs in the fields of energy supply and transportation. It is now well established that the next breakthrough will come from a better understanding of unsteady flow effects and by considering the entire system and not only isolated components. However, these aspects are still not well taken into account by the numerical approaches or understood whatever the design stage considered. The main challenge is essentially due to the computational requirements inferred by such complex systems if it is to be simulated by use of supercomputers. This paper shows how new challenges can be addressed by using parallel computing platforms for distinct elements of a more complex systems as encountered in aeronautical applications. Based on numerical simulations performed with modern aerodynamic and reactive flow solvers, this work underlines the interest of high-performance computing for solving flow in complex industrial configurations such as aircrafts, combustion chambers and turbomachines. Performance indicators related to parallel computing efficiency are presented, showing that establishing fair criterions is a difficult task for complex industrial applications. Examples of numerical simulations performed in industrial systems are also described with a particular interest for the computational time and the potential design improvements obtained with high-fidelity and multi-physics computing methods. These simulations use either unsteady Reynolds-averaged Navier–Stokes methods or large eddy simulation and deal with turbulent unsteady flows, such as coupled flow phenomena (thermo-acoustic instabilities, buffet, etc). Some examples of the difficulties with grid generation and data analysis are also presented when dealing with these complex industrial applications.

1749-4699/09/015004+28$30.00

Contents

Nomenclature

CFD	computational fluid dynamics
CPU	computational power unit
DES	detached eddy simulations
DNS	direct numerical simulation
DTS	dual time stepping
FLOPS	FLoating-point Operation Per Second
HPC	high-performance computing
IGV	inlet guide vane
LES	large eddy simulation
LU-SSOR	lower-upper symmetric successive over-relaxation
Mb/Gb	megabytes/Gigabytes
MUSCL	monotone upstream-centered scheme for the conservation laws
PU	processing unit (i.e. a computing core)
(U)RANS	(unsteady) Reynolds-averaged Navier–Stokes
Re	Reynolds number
SMP	steady mixing plane
TSM	time spectral method
UWC	unsteady whole compressor

1. Introduction

Today, intensive numerical simulations requiring very large computing power can be performed by means of massively parallel platforms, leading to new challenges for computational fluid dynamics (CFD) flow solvers. However and despite the steady increase of the computing power accessible to CFD, its main difficulty comes

from the unsteady and/or turbulent nature of the flow to be simulated, which can be measured through the Reynolds number. Flows with large Reynolds numbers imply a wide range of time and length flow scales (i.e. many vortices of very different sizes) that either needs to be represented through a model or fully simulated on a mesh. For moderate Reynolds number applications, high-performance computing (HPC) allows using the direct numerical simulation (DNS) approach, where all vortices are explicitly resolved in time and space on a given mesh. While still limited to simple geometries, this method allows us to go into the smallest details of the flow, leading to a very fine understanding of the unsteady flow phenomena related to turbulence. For example, wake vortex instabilities that develop behind aircrafts have been studied by Nybelen and Deniau (2009) using the DNS method and a 110 M cells grid ($Re = 10^4$). The simulation uses 1024 computing cores (Blue Gene /L platform) and only 350 h are needed to obtain the solution (i.e. 350 000 computational power unit (CPU) hours). This work based on HPC strategy was one of the first spatial simulations showing the development of elliptic instabilities and the merging process of co-rotating vortices, which are of critical interest for the aircraft security. DNS is also being developed for simulating the flow around a rotating golf ball using an immersed boundary method (Smith *et al* 2008). Computations are performed using up to 500 computing cores on a range of mesh resolutions from 61 M points to 1200 M points (at $Re = 10^5$). Results show the development of instabilities in the near-surface flow as well as the complete local separation mechanisms, providing very interesting data to minimize the ball drag. This technique has also produced major breakthroughs in the field of turbulent reacting flows (Chen *et al* 2009) by addressing more and more complex reacting schemes needed for the understanding of pollutant formations. However, most industrial applications consider very high value of the Reynolds number ($>10^6$) and the DNS approach can no longer be used. In these cases, a part of the problem (such as turbulence) has to be modeled to reduce the simulation cost. This can be done through different modeling strategies that are Reynolds-averaged Navier–Stokes (RANS) or large eddy simulation (LES). For example, Weber *et al* (2000) present a RANS numerical simulation of the flow in a high-power pressurized water reactor (3800 MW). The flow simulation has been performed on a 240 M cells grid composed of several millions of fuel rods, by using 200 computing cores (IBM RS/6000 platform). This work demonstrates the capabilities of modern CFD codes in solving complex thermal-hydraulic problems, including cross-flow, mixing and heat transfer effects. Other complex physics has recently been addressed by Arakawa *et al* (2005) who investigated the tip vortex noise generated by wind turbines. The flow and acoustic field around the wind turbine are simulated using compressible LES (the far field is modeled with thanks to an acoustic analogy). The blade simulation uses a structured grid composed of 320 M points and the flow solution is obtained in 300 h by using 112 vector cores (earth-simulator platform). This work led to the design of an optimized shape, emitting a lower sound pressure level in the high-frequency domain. Other works are related to internal rotating flow simulations. Van der Weide *et al* (2006) investigate unsteady flow phenomena in compressors and turbines by means of an unsteady RANS approach. The flow in a whole turbine stage (88 M cells) and a five-stage compressor (220 M cells) have been computed each with 600 computing cores (IBM × Series335), requiring 450 h to simulate one revolution of the turbine row (3500 h for the compressor). A total of 2.65 million CPU hours is necessary to perform the computation on the whole configuration. A larger mesh grid has also been tested for the five-stage compressor configuration (van der Weide *et al* 2008), showing that the solution on a 750 M cells grid is obtained in only three weeks when considering 10 000 computing cores (Cray XT3). Large multi-component systems are also accessible to CFD flow solvers. The simulation of a complete aircraft gas-turbine performed by Schluter *et al* (2005) and Medic *et al* (2006) using a 350 M cells grid (figure 1) open new perspectives and proves the feasibility of multi-CFD solver coupling as a potential route to address multi-physics flows. This specific simulation of a Pratt and Whitney aircraft engine is based on an inventive RANS (rotating parts)—LES (combustion chamber) coupling strategy, including compressor, combustor and turbine.

This state-of-the-art shows that (massively) parallel computing impacts academic and fundamental research as well as in the industrial context either by increasing the size and the complexity of a problem, or by reducing the calculation return time. Tackling flow problems in increasingly complex geometries or physics can be only possible thanks to massively parallel platforms, usually due to memory requirements. In the second case, the objective is to reduce the computational time of a given problem to potentially offer fast assessment of technological solutions to a given industrial problem. For example and in today's context, unsteady flow

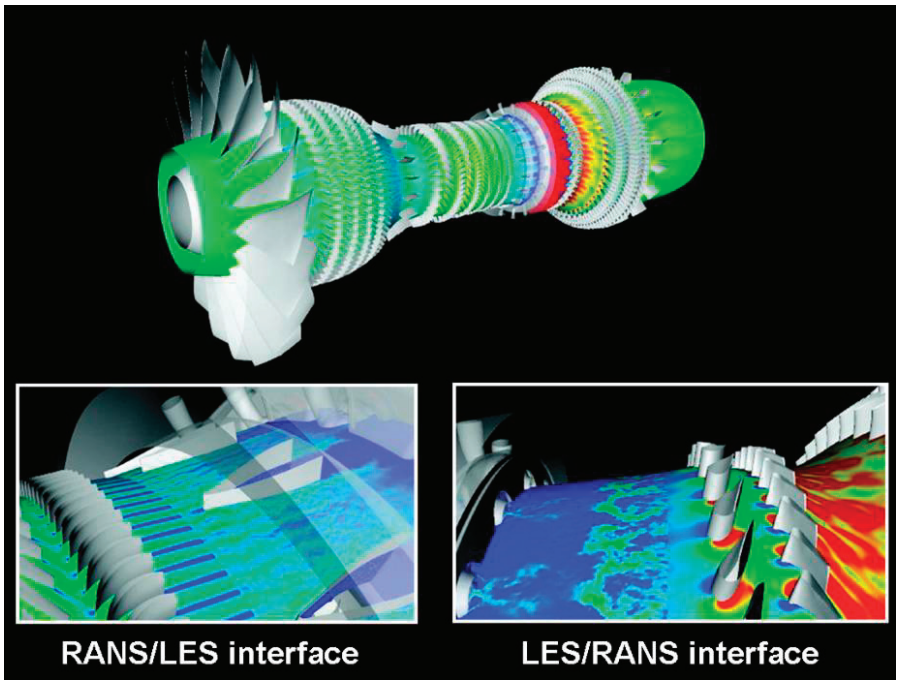

Figure 1. Simulation of a whole gas turbine using RANS/LES coupling (van der Weide 2008).

simulations are out of range at the design stage because the time to obtain the solution is too long (few days to few years, depending on the problem size) and incompatible with current industrial time constraints. Such simulations are, however, not impossible and recent super-computing facilities that usually offer a view of the next generation of computer mainframe to be used by industry within the next 5–10 years, definitely provide the necessary computing environment. As discussed in part I of the paper, code development and validations are still needed to obtain the best benefits from these high-end platforms. Moreover, all the examples of HPC applications described in the introduction also indicate that pre-processing and post-processing are two major issues, due to the size and the complexity of the problem (technological effects, multi-components systems, etc). Indeed, the capacity to design mesh-grids and to exploit/analysis results in a reasonable amount of time is necessary to get the best benefits from HPC (at least this time must be less than the computation itself).

This paper focuses on the application of HPC solutions for solving flows in complex industrial systems. Numerical simulations using thousands of computing cores are presented for advanced industrial applications dealing with flows in gas turbines and around aircrafts. Information about parallel computing methods and computing platforms are first provided. More precisely, the estimation of the parallel computing efficiency is discussed for complex industrial problems. Performance indicators are used such as normalized speed-up, computational time and electric consumption of computing platforms. The interest is to show the capacity of different computing platforms to deal with large and complex configurations (>100 M points, unsteady flows, etc). A few examples of industrial applications are then presented for aircraft, combustion chambers and turbomachines. These examples show a brief overview of the numerical challenges and the related difficulties such as computational time, grid generation and post-processing. A major difficulty with CFD for industrial applications is that geometry is often very complex, including a lot of components and different physics (combustion, aerodynamics, thermal effects, etc). Applications are chosen to give an example of this complexity. Unsteady flows influence the system performance and instabilities are always critical for the design of aeronautic systems. Indeed, the simulation of these phenomena would largely benefit from HPC methods. Specific challenges will be presented for aircraft, combustion chambers and turbomachines applications. For example, the simulation of aero-elastic effects in aircraft at high Reynolds number and the related instability (the buffet phenomenon) are described, based on the use of unsteady flow simulations such as time spectral method (TSM) or zonal detached eddy simulation (ZDES). The ignition of the flow in a whole combustion chamber and the development of thermo-acoustic instabilities are also studied, by means of (LES). Finally, the

use of HPC to simulate unsteady flows (rotor-stator interactions, vortex shedding, etc) and devices to control the development of aerodynamic instabilities is described for turbine and compressor test cases. Comparisons of steady RANS and unsteady RANS results are compared with LES results for a turbine inlet guide vane (IGV), and a coupling strategy between an aerodynamic flow solver and a thermal solver is presented. These results, obtained by means of a HPC strategy, show that considerable improvements of the flow physics understanding can be achieved thanks to these multi-physics and high-fidelity simulations.

2. Application of high-performance computing to complex industrial configurations

This section provides basic information about the flow solvers and the computing platform performance when dealing with complex configurations. The interest is to complete part I of the paper by giving tools to compare the computational resources required by the applications presented in the following sections. For example, computing strategies are different when the computing platform is a vector supercomputer or a massively parallel platform.

2.1. Overview of flow solvers

Two flow solvers are considered in this paper: the structured multi-block flow solver *elsA* (Cambier and Veuillot 2008) dedicated to aerodynamics problems and the unstructured flow solver AVBP (Garcia 2009) dedicated to reactive flows. The *elsA* software is a multi-application CFD simulation platform that deals with internal and external aerodynamics. The compressible three-dimensional (3D) Navier–Stokes equations for arbitrary moving bodies are considered in several formulations according to the use of absolute or relative velocities. In order to deal with unsteady and/or separated flows, high-fidelity methods such as unsteady RANS (URANS), detached eddy simulation (DES) and LES are available. High-flexibility techniques of multi-block structured meshes have been developed to deal with industrial configurations, such as non-coincident interfaces and sliding mesh methods. The semi-discrete equations are integrated either with explicit schemes (Runge–Kutta, etc) or with implicit schemes (backward Euler, lower-upper symmetric successive over-relaxation (LU-SSOR) method, etc). The flow solver AVBP is an unstructured flow solver that solves the laminar and turbulent compressible Navier–Stokes equations by means of LES. AVBP focuses on unsteady turbulent flows (with and without chemical reactions) for internal flow configurations. A reduced Arrhenius law chemistry model allows investigating combustion for complex configurations. Specific conditions have been developed for complex configurations such as moving meshes, allowing the simulation of moving bodies (piston in a cylinder, etc). More information about parallel computing strategies used by *elsA* and AVBP can be found in part I of this paper (Gourdain *et al* 2009c).

2.2. Speed-up and computational time

As discussed in part I of the paper, the normalized speed-up is often use to measure the scalability of a parallel task, especially for very large industrial configurations. However, results given by this criterion must be interpreted carefully. The main problem is that a different mesh partitioning is often considered for each number of processing units (PUs), leading to a different problem. Such a shortcoming is true for multi-block structured and unstructured grid solvers. The consequence is that the computational times obtained for solving different problems are compared to determine the initial problem scalability. To illustrate this remark, a study is provided for two mesh-partitioning strategies used to compute the flow in an industrial compressor. The configuration is composed of a structured multi-block grid with coincident and non-coincident interfaces. In the first case, the mesh partitioning is optimized for each number of PUs: i.e. the number of blocks and ghost cells can vary with the number of considered PUs. In the second case, the original problem is split only one time for a large number of PUs ($N = 4096$). This new problem is then load balanced for lower numbers of PUs. Finally, calculations are performed with *elsA* on 512, 1024 and 2048 PUs. The normalized speed-up values are compared with the theoretical speed-up estimated with the Amdahl's law (1967).

Results are presented on figure 2. The optimized mesh partitioning seems to give poor quality results in terms of normalized speed-up. When comparing with the Amdahl's law, the calculation efficiency is less

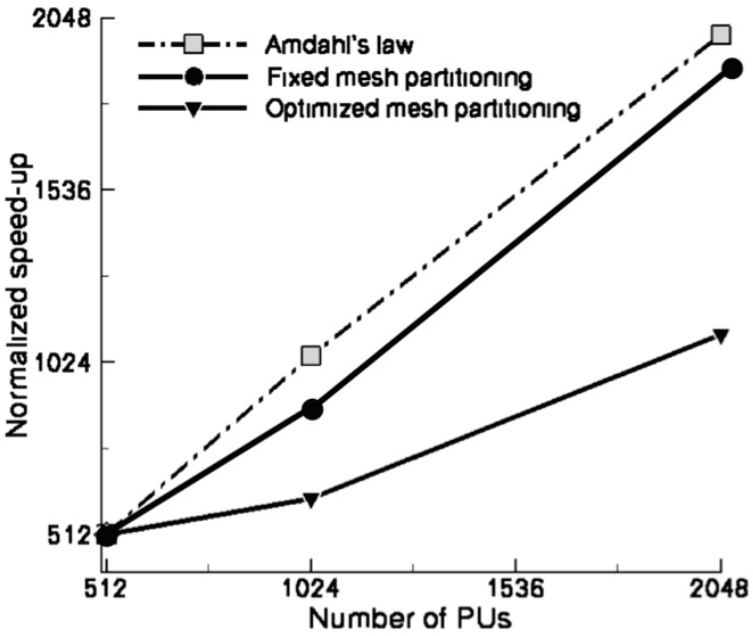

Figure 2. Effect of the mesh partitioning strategy on the speed-up (optimized mesh-partitioning and fixed mesh partitioning with the *elsA* flow solver).

than 50% with 2048 PUs. With the fixed mesh partitioning, the calculation scales well and the efficiency with respect to the Amdahl's law is around 90%. At this stage, a naive conclusion would be that the use of a fixed mesh partitioning with a large number of blocks is a better strategy than an optimized mesh partitioning for each number of PUs. However, the observed computational time is usually better with an optimized mesh partitioning since this strategy corresponds to the solution with the minimum number of blocks (and thus the minimum number of ghost cells). For the present application, no general rule can be found, mainly because the parallel efficiency depends on the repartition of blocks with non-coincident interfaces (that are not taken into account during the mesh-partitioning). On the one hand the simulation with 512 PUs runs 13% faster with the optimized mesh partitioning than with the fixed mesh partitioning. On the other hand, with 2048 PUs, the simulation runs 55% faster with the fixed mesh-partitioning method. In this case, the correct parameter to measure the parallel computation efficiency is the time required to obtain the solution, no matter which strategy is used (this is especially true for industrial applications).

2.3. Computing platforms performance

In the industrial context, the time to obtain the solution with CFD flow solvers ranges from a few hours to few days. Most industrial flow solvers such as *elsA* are running on both scalar and vector supercomputers. All these platforms provide sufficient performance for most the current applications but at different costs in terms of acquisition, maintenance and electric consumption. In this context, it is interesting to compare the costs induced to obtain the solution with both kind of platforms. Usually scalar platforms are well adapted to massively parallel computing by using thousands of computing cores concurrently (but with a limited memory by core) while the vector approach uses a few powerful computing cores with a large memory. Several benchmarks have been performed in order to evaluate the benefits of these approaches and the scalability of an industrial test case in terms of elapsed time (figure 3(a)) and power consumption (figure 3(b)). The test case is a whole multi-stage compressor (that uses a 134 M cell grid), which is a representative industrial application that does not scale perfectly well with a large number of PUs (see figure 2). Performance is evaluated by computing ten physical time steps for all benchmarks. Three platforms have been selected to show an overview of current available supercomputers. The first one is a vector supercomputer (NEC SX8++) with 16 nodes of 8 computing cores; each node is equipped with 128 Gb of RAM and each computing core develops 32 GFLOPS

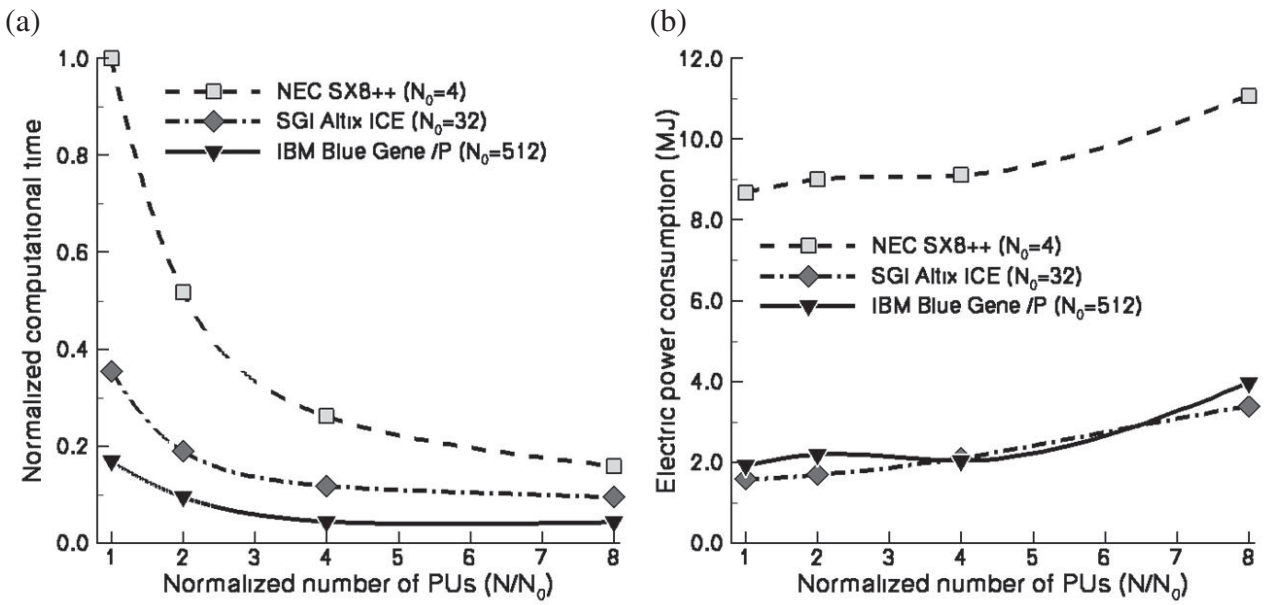

Figure 3. Comparisons of the time needed to compute one iteration (a) and electric consumption (b) with the *elsA* flow solver: application to an industrial multistage compressor.

(Giga FLoating-point Per Second, peak power). The second platform is a scalar platform (SGI Altix ICE) composed of 1536 bi-quad cores (i.e. the total is 12288 computing cores); each bi-quad core is equipped with 32 Gb of RAM memory and each core develops 12 GFLOPS. The third platform is a massively scalar computer (IBM Blue Gene/P) that is composed of 8192 nodes each with 4 computing cores; each node is equipped with 2 Gb of RAM and each core develops 3.3 GFlops.

The first comparison is done to estimate the number of computing cores needed to achieve a given elapsed time. Results shown in figure 3(a) are normalized by the time required to compute one physical time step with four vector cores (i.e. 2783 s). An identical elapsed time (around 450 s by iteration) is obtained with 512 scalar cores (IBM), 30 vector cores (NEC) and 72 scalar cores (SGI). These results are useful to underline the efficiency and the interest of parallel computations. On the one hand, the use of a small number of computing cores (<128) does not significantly affect the computational time that is still largely related to the computing core peak power: to achieve a given elapsed time, the ratio between the required number of vector (NEC) and scalar (SGI) computing cores ($= 2.7$) is very similar to their core peak power ratio ($= 2.4$). On the other hand, computations that consider a large number of computing cores (>128) are more affected by communication cost and load balancing errors, even with an optimized communication network: while the peak power ratio between scalar (IBM) and vector (NEC) computing cores is close to 9.7, the ratio between the number of computing cores required to achieve an identical elapsed time is increased to 17.1. Similar observations are done when comparing IBM and SGI platforms and results also indicate that a plateau is reached for the elapsed time when the number of computing cores is increased, showing the limit in terms of scalability.

Based on the manufacturer specifications (i.e. 8 W for an IBM core, 780 W for a NEC core and 49 W for a SGI core), it is possible to obtain the electric consumption to compute one physical time step of the flow solver (figure 3(b)). The massively scalar platform (IBM) requires 1.9 MJ with 512 computing cores, the vector supercomputer (NEC) needs 11.0 MJ with 30 cores and the scalar platform (SGI) requires 1.7 MJ with 72 cores. At identical performance in terms of elapsed time, the vector supercomputer is always more energy consuming than scalar supercomputers. When comparing IBM and SGI platforms, the result depends on the number of considered computing cores: for small numbers of computing cores, the SGI platform appears to be more interesting, while the IBM platform is better adapted to the use of a large number of cores. The same electric power is required for 768 cores (IBM) and 128 cores (SGI). The simulation of one compressor rotation is usually sufficient to obtain an unsteady flow solution at design operating conditions. Considering an identical elapsed time, the required power is 5440 MJ with the SGI scalar computer, 6080 MJ with the IBM

platform and 35 200 MJ with the NEC vector computer. To give an idea of the order of magnitude inferred by these consumptions, such powers are equivalent to run the experimental facility during 45 min (simulation with the SGI platform) or 4 h 54 min (simulation with the NEC supercomputer). This consumption is also equivalent to the consumption of the energy produced by a 'standard' nuclear plant (1000 MW) during 5 s (SGI platform) or 35 s (NEC platform). A second configuration dedicated to an aircraft application has been tested. The multi-block mesh contains more than 1000 blocks (28 M cells) and only coincident interfaces are used for block connectivity. In this case, the same computational time is observed with 1024 PUs of an IBM Blue Gene /L system (scalar) or 30 PUs of an NEC SX8+supercomputer (vector). The actual financial cost of these simulations is interesting but difficult to evaluate since it is different for each company (and most of the time confidential), depending on negotiations with manufacturers. The conclusion is that (massively) scalar platforms are usually more interesting for the presented applications regarding computational time and power consumption.

3. Application to civil aircrafts

Designers face complex challenges for the next generation of civil aircrafts; they will have to operate under strict environmental rules resulting from an ever increasing market and diminishing fuel availability. Performance estimation plays an extensive role at all stages of a commercial aircraft development program and design methodologies, together with computational methods, largely support the creation of these aircrafts. Innovative design requires interdisciplinary, multi-physics and interactive design methods that need more and more computational power. In this context, a rapid and detailed design of a greener and quieter aircraft could be largely helped by use of HPC platforms at various phases. Firstly, HPC can be useful at the design stage by reducing the time required to obtain the solution (at constant cost). Secondly, based on the engineering experience, the number of numerical simulations necessary to evaluate the complete flight domain (cruise conditions, high lift, take-off, etc) is of the order of one to a few thousands CFD runs. HPC is a very interesting tool to better manage the scheduling of such tasks. Finally a direct application of HPC is to achieve 'high-fidelity' simulations in a reasonable amount of time, taking into account unsteady flows (URANS/LES) and technological effects to better understand an apprehend potential difficulties. An overview of the flow phenomena encountered during flight conditions is proposed in figure 4 (Rossow and Kroll 2008). The region above the dark line represents the forbidden flight domain, due to undesirable phenomena. As shown, the design of new aircrafts has to tackle complex physics such as shock/boundary layer interaction (buffet phenomenon), massively separated flows or aero-elastic instabilities.

3.1. Grid design and post-processing

Post-processing is often quite straightforward in the context of industrial applications, since it mainly relies on steady coefficients such as drag and lift. Indeed, it requires only steady (or time-averaged) solutions directly provided by the flow solver and it does not need in itself costly storage operations. However, an aircraft can operate under a very large flight domain and a lot of simulations is usually necessary to fully describe all the flight conditions. Indeed aircraft manufacturers need to access a very large database, costly in terms of storage and computing capacities. A popular solution is thus to use a method based on a proper orthogonal decomposition (Delville *et al* 1999) that is very efficient to reduce both the size of the database and the number of required simulations. Another difficulty is related to the sensibility of aerodynamic parameters to the mesh-grid design, which is thus a critical step for the flow simulation around aircraft. Today, a grid including most technological effects (wing, tail, flap track fairing, nacelle, etc) requires 100 to 150 millions points. Engineers can use two strategies to design such a grid: either a structured multi-block grid (as required by *elsA*) or an unstructured grid (as required by AVBP). In the first case, the geometrical complexity of the configuration imposes very strong constraints on the mesh. Many people work in the same time to design the grid (the so-called task parallelism). Each person designs a part of the mesh (such as fuselage, wing, tail, etc) and non-matching conditions are then used to join the different parts. Even by means of this strategy, around one month is usually necessary to obtain the final grid (that contains more than 1000 blocks). When using

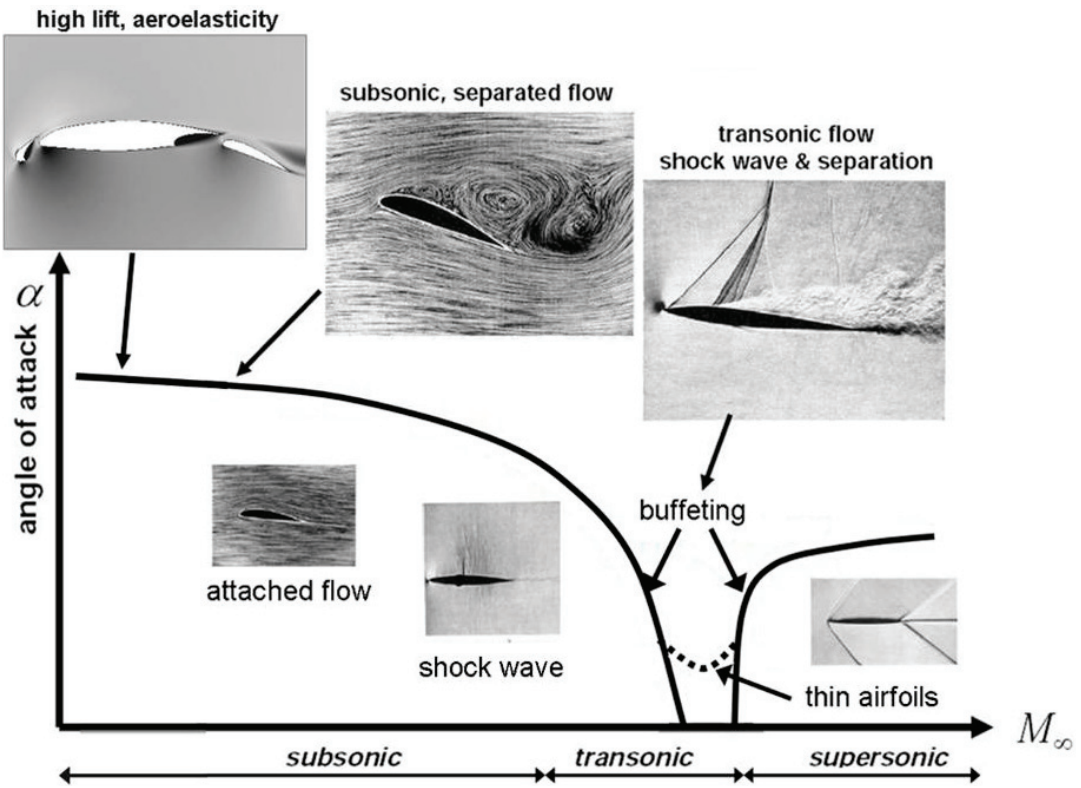

Figure 4. Overview of the flow phenomena, depending on the flight conditions (Rossow and Kroll 2008).

unstructured grids, the main difficulty comes from the description of the multiple boundary layers behind the aircraft (generated by wings, tails, fuselage, turbofan facilities, etc). In this case, about one month is also necessary to validate the final grid. This task is thus of primary importance since the flow simulation is performed on a daily basis (at least for industrial applications), meaning today more time is needed to design the mesh grid than for the flow simulation itself. However, the huge number of parameters required to design an aircraft imposes to run a lot of simulations and thus very efficient computing algorithms are needed to reduce the cost associated with this task. A typical example of that is the simulation of aero-elastic effects.

3.2. Aeroelasticity simulations

Compared to flight conditions, take-off and landing conditions necessitate that the aircraft lift be increased, which is usually obtained by setting higher angle of attack of the wing. However, this approach requires a flexible wing design that interacts with other parts (such as spoilers, flaps, ailerons, etc), inducing unsteady effects. The prediction of aero-elastic efforts and time-dependent cycles (flutter, etc) is necessary but requires large computational resources. For example, a simulation of a whole generic aircraft configuration (including propulsion system) has been performed with the flow solver *elsA* to underline the effect of spoilers (figure 5). Unfortunately, a direct fluid/structure coupling is still beyond industrial resources. A common practice is thus to perform unsteady forced motion simulations to obtain the unsteady loads on the wings of a complete aircraft configuration (Delbove 2006). Due to high-Reynolds numbers encountered during flight and the need to extend the computational domain far from the aircraft, a typical grid ranges from 50 to 150 millions cells, depending on the complexity of the configuration. A second point of concern is the unsteady nature of the studied phenomena. The long turnaround time associated with classical methods has prompted the development of alternative techniques to reduce the computational time (Hall *et al* 2003).

The classical unsteady approach to treat such configurations is the URANS method. Although it can potentially capture most physics nonlinearities, this method suffers from a high-computational cost: the

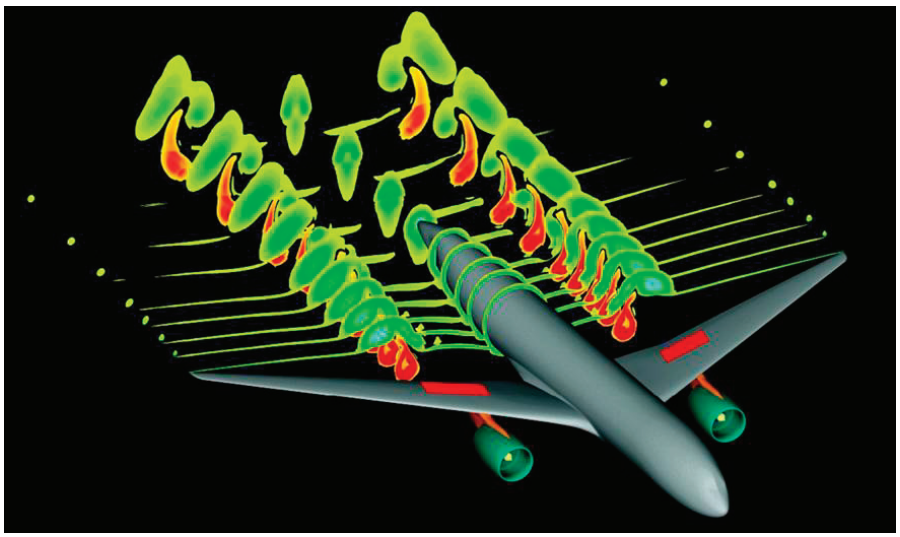

Figure 5. Simulation of aero-elastic effects in a whole generic long-range aircraft performed with *elsA*—instantaneous solution of total pressure.

Table 1. Comparisons of TSM and URANS costs for aero-elastic simulations (*elsA*).

Computing method	Computational time	Memory consumption
Standard aero-elastic solver	1	1
TSM-aero-elasticity, 1 harmonic	0.09	3
TSM-aero-elasticity, 2 harmonic	0.16	5
TSM-aero-elasticity, 3 harmonic	0.24	7

simulation must cover several periods of the unsteady phenomena, leading to a long transient phase (from 2 to more than 20 periods). For typical aircraft configurations, the computational need ranges from 50 to 150 Gb for the memory and 500–1500 CPU hours on a vector computer. A recently developed method for time-periodic flows is the TSM, also referred as the harmonic balance technique (Gopinath *et al* 2007, Sicot *et al* 2008). Based on a Fourier decomposition of the flow variables, it transforms an unsteady problem into several coupled steady calculations. The coupling is done by a source term that can be viewed as a spectral derivative operator, expressed in the time domain.

The TSM strategy is applied to both the CFD flow solver and the structural dynamics solver. Table 1 presents a comparison of the normalized computational time and the memory usage obtained with standard (URANS) and TSM aero-elasticity solvers: the use of TSM significantly reduces the total computational time. For an aircraft configuration, the CPU timesaving ranges from a factor 10 for a single harmonic representation to a factor 4 when three harmonics are considered. However, the counterpart for this reduction in CPU time is an increased memory requirement, due to the fact that several instants in the period are solved and stored simultaneously. The memory needed for the aircraft configuration is increased by a factor 3 (for a single harmonic) and up to 7 (for 3 harmonics). In practice, about 150–350 Gb are required to obtain the solution on coarse grids (30 M cells), imposing serious constraints on the hardware requirements. TSM, requiring high memory (depending on number of harmonics), is ideally suited to parallel computing. Such applications demonstrate the need for efficient computer architectures with high-memory capacity. In this context, HPC is an interesting solution to simulate aero-elasticity effects and design the next aircraft generation.

3.3. Application to buffeting

Aircraft buffeting is a vibratory phenomenon that may appear during manoeuvres or while cruising. Buffeting is eventually driven by aerodynamic mechanisms, involving separations of the wing boundary layer, but the

Table 2. Typical computing effort related to RANS and ZDES simulations of transonic buffet on a civil aircraft configuration—NEC-SX8+computations with *elsA*.

	RANS (steady)	ZDES (unsteady)
Number of points	10 000 000	10 000 000
Total CPU (hours) effort for converged state (RANS) or 30 ms of simulation (ZDES)	10	2000

performance limitation depends on the structure dynamic response. The main consequence is a reduction of the life duration of the various components subject to the high-vibration levels. As a consequence, manufacturers have to design aircraft with a large lift margin that tends to limit minimum flying speed and maximum cruising altitude of the aircraft. The understanding of this phenomenon is thus crucial for designers. However, the simulation of buffet with CFD tools is challenging, due to massive flow separations and unsteady behaviors. Considering realistic aircraft shapes, it is now well established that buffeting exhibits a broadband frequency distribution that presents high-instability levels. Buffeting on 3D configurations can be considered as an interaction between the separation unsteadiness and the shock, so that (U)RANS methods are not adapted to this kind of simulation since the basis of these approaches is a full modeling of the turbulent spectrum (only decoupled phenomena are correctly simulated such as Von Karman streets). A promising way is thus to perform a simulation that solves (at least a part of) the turbulent eddies. On the one hand, the LES approach is a potential solution but it is still very costly in terms of computational resources for aircraft configurations at realistic Reynolds numbers. On the other hand, RANS methods are accurate enough in boundary layers for a low CPU cost. Hence, RANS/LES hybrid methods, such as the DES approach (Spalart *et al* 1997) that proposes to use RANS for attached boundary layers and LES for separated regions, are very promising. The ZDES method is an adaptation of the original DES approach suggested by Deck (2005) to overcome some specific difficulties related to DES (the so-called 'gray region'). Such costly unsteady flow simulations take advantage of parallel computing platforms to reduce the computational time (and cost) and obtain a correct description of complex flow phenomena. The feasibility of a 3D buffet simulation with the ZDES method implemented in *elsA* was studied by Brunet and Deck (2008) on a realistic aircraft configuration ($Re = 2.8 \times 10^6$, $M_\infty = 0.82$). The computing efforts are compared in table 2 for RANS (using the Spalart–Allmaras turbulence model, 1992) and ZDES. The cost ratio between the two approaches is around 200. However, a pure RANS approach is not able to reproduce the unsteady behavior of buffet contrarily to ZDES, which correctly predicts most flow features related to buffet: i.e. the unsteady separated flow region that triggers the shock motion. Brunet and Deck (2008) have presented detailed comparisons with experiments, showing a major increase of the predictive capacity with ZDES compared to RANS.

4. Application to combustion chambers

Turbulence is a recurrent and recognized challenge for which the scientific community has not yet been able to provide reliable methodologies necessary for its prediction in complex industrial applications. In fact and throughout the past century that challenge has been clearly identified because of the tremendous impact such a contribution would have on existing industrial processes. Among all industries, the gas turbine and automotive companies are probably the most receptive to any new contributions in the field of turbulent reacting flows because of the upcoming new regulations and the existing pressure linked to petroleum consumption and pollutant emissions. From a purely scientific point of view, the treatment of combustion and turbulence yields additional difficulties.

The first difficulty stems from the length scale of a reacting front, i.e. the distance over which reactants are chemically transformed into products. For simple fuels, this distance is usually of the order of 1 mm at atmospheric pressure and greatly reduces when pressure increases. Turbulence as underlined previously covers a wide range of scales and dimensions, which will depend on the burner Reynolds number. Relatively to the reacting scale, interactions between chemistry and turbulence have different consequences. The major effect is the increase in flame surface and overall rate of combustion: turbulent burners are usually more compact

than laminar burners. The main drawback of these interactions is the potential quenching, flame blow out and burner extinction. Indeed a reacting front can withstand given strain and stretch levels beyond which chemistry is inhibited. Below these values, combustion is stimulated compared to the un-stretch configuration. Reproducing all the possible turbulent reacting interactions encountered in a complex burner is a very difficult task although RANS did provide significant contributions. Among the scientific interests involved in turbulent reacting flows, the following (non-exhaustive) list gives a brief overview of the subjects related to aeronautical applications:

- combustion including fuel consumption and pollutants emissions (Baum *et al* 1994, Chen *et al* 2006, Im *et al* 1999, Lignell *et al* 2007, Mizobuchi *et al* 2005, Seiser *et al* 2005)
- multi-phase flows (Fevrier *et al* 2005, Miller and Bellan 1999, Reveillon and Vervisch 2000, Vermorel *et al* 2003),
- radiation (Soufiani and Djavdan 1994, Wu *et al* 2005),
- primary and secondary break-up of liquid sheets (Gokalp *et al* 2000, Varga *et al* 2003),
- thermo-acoustic instabilities (Poinsot and Veynante 2005, Truffin and Poinsot 2005) and
- aero-acoustics (Bogey and Bailly 2006, Freund 2001).

The numerical approach routinely used in industry is the RANS approach, mainly developed in the 1970s (Lumley 1978, Tennekes and Lumley 1972). These methods only provide the mean solution (in the sense of time or ensemble average) and all effects of turbulence are modeled. As a consequence, this approach heavily relies on the turbulence models. Their range of applicability is theoretically restricted to statistically stationary flows although industry sometimes uses RANS beyond this limit to design new industrial systems. Here again, the increase of computing power offers an efficient turnover time that allows us to gauge new design points before industrial validation of the final design on expensive experimental test facilities. In order to extend the domain of validity of the numerical tools, especially in the industrial context, a new numerical approach has been developed. This numerical method called LES (Smagorinsky 1963), introduces the notion of scale separation, which filters out the high-frequency flow structures and retains only the large-scale motions, which usually depend on the geometry considered. With this separation of scales, the computing effort is increased when compared to RANS while remaining largely below the huge computer costs of DNS where all turbulent scales are simulated. The models, still needed with LES, are nonetheless more universal and more physics is naturally imbedded in the simulations. Indeed, LES offers access to a spatially and temporally dependent flow description, and is particularly valuable for unsteady flow phenomena. As such and despite the fact that LES is still under development, it opens new perspectives, especially in the industrial context where transient physical phenomena are known to be critical which is definitely the case for reacting flows.

As of today, the development of LES for an extensive use by industry at the design stages needs to meet several requirements. Among the important issues, scalability is of the foremost importance since these applications aim at running on HPC architectures to bring response times within the industrial requirements or to tackle ever more complex geometries. Systematic model and code validations are necessary to ensure reliability of the LES tool. In the following, a brief state-of-the-art for LES of turbulent reacting flows performed mostly with AVBP in complex configurations is proposed. Issues about modeling, code implementation and validation procedures are also discussed. Potential perspectives and closing remarks are given prior to the illustration of other aeronautical applications based on HPC codes developed at CERFACS.

4.1. Grid design and post-processing

For industrial applications with complex combustion chamber geometries, the use of structured meshes can be either very complex or impossible. Hexahedral unstructured meshes are sometimes used, for example in piston engine configurations. For aeronautical combustion chambers, however, the most common solution is a full tetrahedral grid mesh. An alternative solution to obtain the best flexibility is to use hybrid meshes, typically a combination of prismatic layers near the walls and tetrahedral elements in the rest of the domain. For a clean geometry (i.e. with well-defined curves and surfaces, and without unclosed surfaces), the generation of such a mesh is today quite simple and fast compared to the cost of the simulation itself (usually not more than 1

day). The post-processing of a typical unsteady turbulent and reacting flow computation requires many efforts both in terms of memory and storage capacities. It is mainly based on three kinds of diagnostics: instantaneous solutions, time-averaged solutions and local probes. Concerning the storage issue, instantaneous solution files are clearly the bottleneck. Typical unstructured meshes contain several million nodes and generate files going from a few tens of Mb up to several Gb, depending on the number of variables stored. To give a measure of the difficulty, one snapshot for the instantaneous field of a whole annular combustor weights about 1 Gb and a temporal evolution of the main flow structures usually necessitates four million iterations (i.e. the 4000 Tb is required to store the full solution). Of course, only primitive variables are strictly necessary and should be stored but in reality, many other variables are usually also stored (gradients, sensors, wall quantities, etc) since they are directly available in the code during the computation and thus avoid an additional computational cost *a posteriori*. Time-averaged solution files are less problematic in terms of storage since they are stored only once a run. Their analysis is of primary importance since the comparison with experimental data and the validation of any LES computation is only meaningful in a statistical sense. They can be calculated either from the instantaneous solution files with a post-processing tool or 'on the fly' during the calculation, the last solution being generally favored as it allows a better frequency resolution. Indeed, the storage capacities often restrict the number of instantaneous solutions. Unsteady flow post-processing is therefore often ensured thanks to local probes rather than full instantaneous solutions. Probes are specified grid points where the temporal evolution of all desired quantities are recorded during the computation. This diagnostics is little costly in terms of storage and allows a detailed analysis of the flow unsteadiness via frequency analysis for example. However, its use implies a strong knowledge of the flow *a priori* since these probes have to be determined before the computation, blindly. Although incomplete, part of these questions are presented and discussed by Gicquel *et al* (2008). The inferred difficulty when being able to perform LES of such problems and having at hand such an amount of data is to extract valuable information. New grid handling and post-processing tools will be needed to fully understand the physics at hand. In particular, analysis of the temporal means and fluctuations will remain of interest but new advanced diagnostics addressing the full spatial and temporal nature of the solution are still to be designed to identify the key coupling elements resulting in a given observation.

4.2. Simulation of the ignition stage

Current HPC facilities provide enough computing power to demonstrate the capacity of advanced LES codes in addressing complex industrial configurations (figure 6). Although still at the state of demonstration, numerical simulations clearly demonstrate that LES can provide interesting data for industrial configurations (Boileau *et al* 2008, Boudier *et al* 2007, Moin and Apte 2006, Staffelbach *et al* 2009). The final requirement that may be necessary and that usually dimensions the size and the type of the problem targeted is the access to a large number of computing cores. As indicated on figure 6(a), most simulations usually focus on a single sector of the full combustor. With the increasing number of ever-larger HPC centers (INCITE), full annular chamber computations will rapidly exit the realm of 'opportunity' or the need for specific 'challenges'. For example, the fully transient LES of the ignition phase for a complete helicopter gas-turbine combustion chamber shown on figure 6(b) was achieved thanks to a dedicated access to the most powerful supercomputers currently available in the world: the Bull machine (Tera-10 Nova-Scale 5160) from the CEA computing facility located in France, the IBM (blue-Gene/P) architecture located in Rochester, USA, the Barcelona Supercomputing Centre (Blade Centre JS21) IBM, Spain, and the Oak Ridge National Centre (Cray XT3/XT4) Cray, USA. It is also interesting to underline that although all the reported advanced LES results aimed at being purely demonstrative, the acquired knowledge yields important scientific contributions with several journal publications in the field of turbulence and combustion. This interesting consequence further emphasizes the importance of such pioneering works for industry but also for the scientific community.

4.3. Simulation of combustion instabilities

A typical knowledge acquired thanks to HPC LES of turbulent reacting flows in aeronautical chambers deals with the thermo-acoustic stability of such devices. Indeed, with the advent of fuel lean operating burners to control fuel consumption and pollutant emissions, aeronautical engines are more prone to combustion

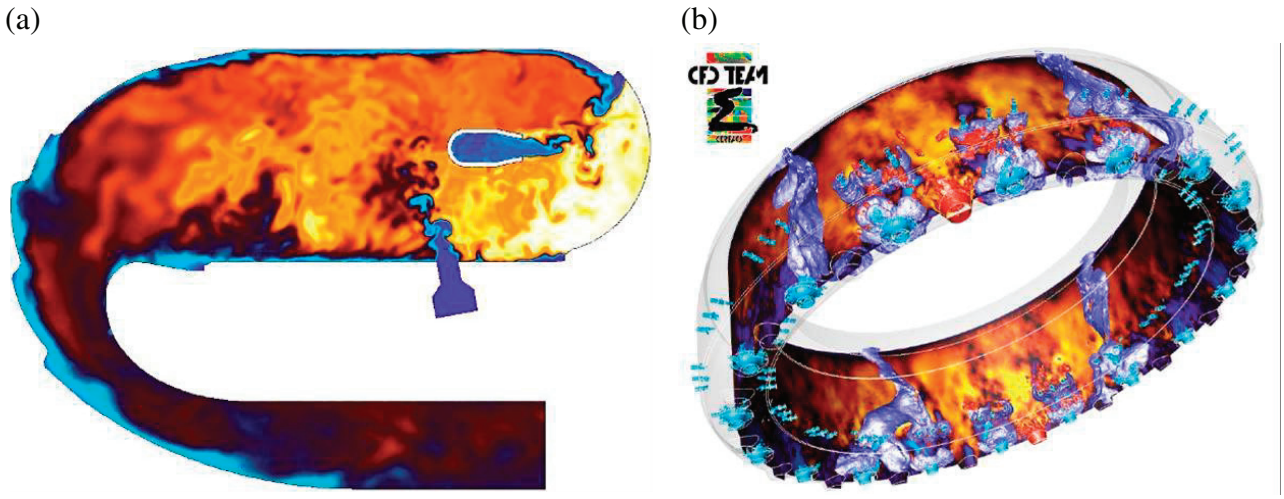

Figure 6. Instantaneous temperature flow fields obtained with AVBP (LES) in aero-engine combustion chambers: (a) Turbomeca single sector chamber (Boudier *et al* 2007) and (b) ignition sequence in a whole Turbomeca annular chamber (Boileau *et al* 2008).

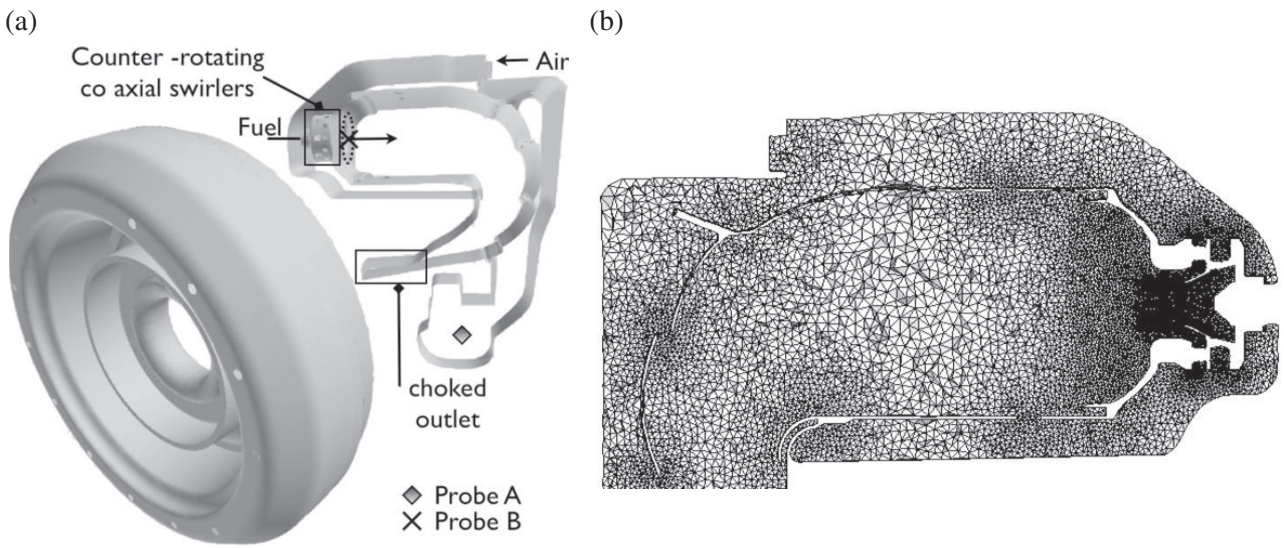

Figure 7. Full annular aeronautical burner computed with AVBP (LES) for a thermo-acoustic analysis of the burner: (a) computational domain and (b) typical mesh resolution in the injector region.

instabilities. Such a running condition of the engine is of course undesired since it can lead to its destruction. Thermo-acoustic instabilities are nevertheless unpredictable in the design phase and engineers often face the problem at the end of the project when the engine is ready for qualification. Advanced numerical tools such as LES are for this highly sensitive issue of valuable interest. Although the computer effort is important, this fully unsteady numerical method inherently contains the physics resulting from the coupling between the acoustic eigenmodes of the geometry and combustion that are at the root of thermo-acoustic instabilities. To study this phenomenon, the LES of the entire module was performed (Boudier *et al* 2009, Staffelbach *et al* 2009), i.e. the combustion chamber, the fuel injection system (also called 'swirler') and the casing (figure 7). All components were included in the computation in order to reduce the boundary condition effects, which are known to be critical for such problems. The unstructured grid is composed of 42 287 640 tetrahedra with local refinement as shown in figure 7(b), to guarantee proper operation of the different LES sub-models such as the turbulent combustion, turbulent mixing and turbulence generation models.

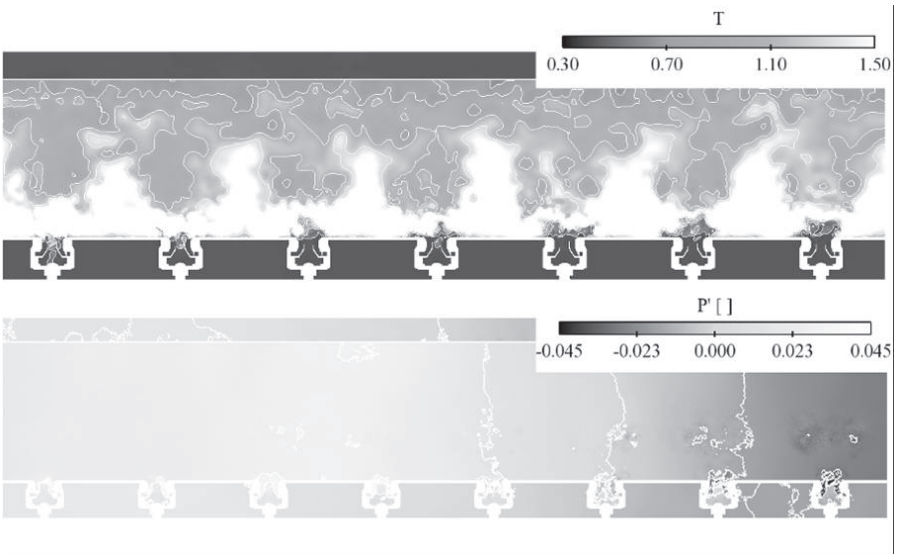

Figure 8. Typical snapshot computed with AVBP (LES) for the configuration pictured on figure 7. The view is produced for a given radial plane going through all swirlers at a given instant (only few burners are pictured for clarity).

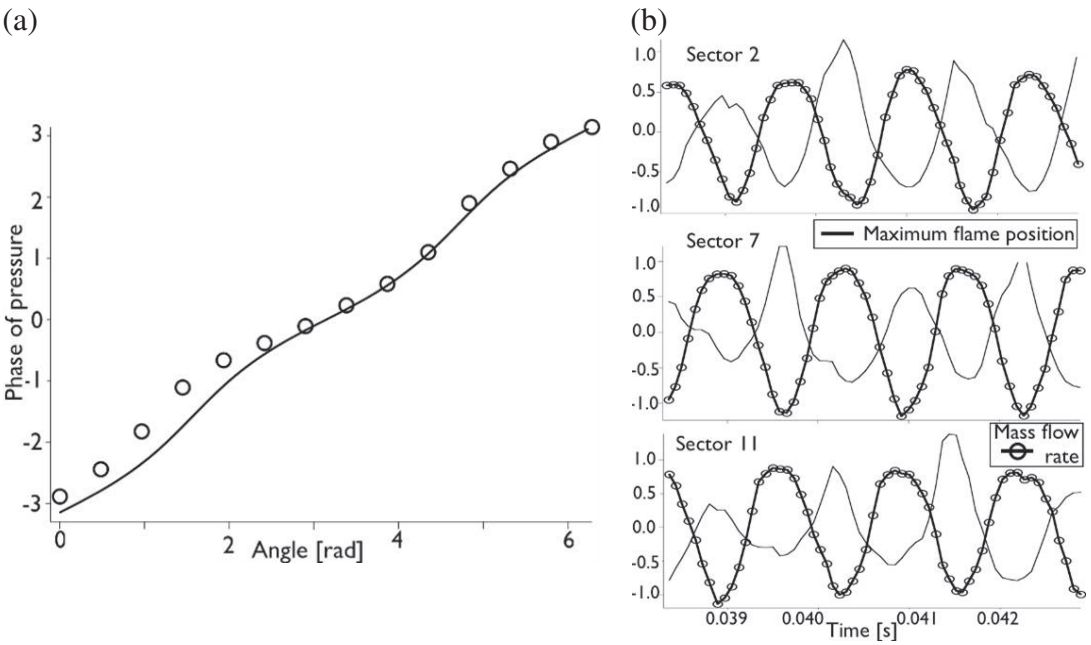

Figure 9. Thermo-acoustic analysis of the LES results obtained with AVBP for the full annular aeronautical gas-turbine combustion chamber: (a) phase-angle relationship of the azimuthal pressure signal (symbols: analytical model, curve: LES data) and (b) swirled mass flow rate and axial flame position.

Although the actual computation is based on a simplified version of the real engine (purely gaseous LES, fully axisymmetric computational domain, etc), the simulation captures the self-excited nature of the azimuthal thermo-acoustic oscillation (figure 8). Frequencies and shape of the acoustic eigenmode were captured as well as pressure variations observed on the engine. Finally, the thermo-acoustic nature of the oscillation, which was undetermined initially, is confirmed here thanks to the LES approach (figure 9).

Such computations, although of scientific interest, should be viewed as a first step toward the full industrial use of LES. Most results are obtained for fully gaseous LES when real engines operate on liquid fuels and chemistry models are simple. In fact, many phenomena controlling the real engine combustion chamber are not understood yet or modeled in such computations. Technological aspects introduced by the manufacturers to specifically address cooling of the chamber walls for example are open issues, which only start to be dealt and modeled by LES experts (James *et al* 2006, Lefebvre 1999, Mendez and Nicoud 2008). Note also that the validation of these LES simulations in a rigorous scientific way remains a challenge since experimental measurements are scarce for such real burners. Finally, issues specific to the LES approach (Boudier *et al* 2008), its analysis and the handling of the huge amount of data generated by such codes are basic questions that need to be addressed to fully capitalize on the information generated by LES in the industrial context.

5. Application to turbomachinery

Rotating machines, such as encountered in gas or wind turbines, are involved in most of the energy conversion processes. Increasing their efficiency is a necessary step to improve industrial processes and reduce pollutant emissions. Unsteady flows that take place in these systems are still not well understood, especially in multistage turbomachines. The consequence of this lack of knowledge is a difficulty to design highly efficient turbomachines. It is now well established that the overall performance of rotating machines are strongly dependent on complex unsteady flows such as end-all flows (Domercq and Escuret 2007) and rotor–stator interactions (He 1997). While experimental campaigns are the historic way to provide reliable data for studying these unsteady flow effects, this approach leads to complex measurement techniques and high costs and delays. Again, CFD offers a very attractive method to reduce cost and design time. However, flow solvers usually give poor predictive results regarding the complex flows that develop in these systems. The main reason for this lack of precision is that computational cost is directly related to the size and complexity of the computational domain. Most industrial simulations thus focus on a limited part of the system (such as an isolated blade) that is solved with a steady RANS approach.

5.1. Grid design and post-processing

Post-processing is an easy task in the case of purely steady flow simulation (most industrial applications), since the main interest is to extract global parameters from the database (such as pressure ratio, efficiency, etc). The problem becomes more complex when dealing with unsteady flows. In this case, the storage of (a part of) the unsteady solution is necessary and the storage capacity can be the limiting factor. For example, a grid with 130 millions cells grid (a 'coarse' grid to represent a whole three-stage compressor) and a moderate number of time steps (about 6400 iterations to discretize one rotation) generates about 42 Tb of data if all the information is conserved. Obviously, the whole solution cannot be stored and only the regions of interest are extracted. First, a time-averaged flow field can be computed. Since one run is usually not enough to perform the simulation, this work must be done at the end of each run by an external post-processing, combining both the current run solution and the previous run solution. Then to deal with the analysis of deterministic unsteady flow (related to the relative motion between blade rows), Fourier series decomposition can be used both for azimuthal and time signals. Following a well-known method used for phase-lagged simulations (He 1990), the interest is to store only the first Fourier modes that are assumed to contain the greatest part of information. Based on the previous example of a whole three-stage compressor, it is possible to reduce the whole database to about 50 Gb, by keeping only the three first harmonics of the blade passing frequency at 100 axial position. This database can still be reduced if the information is only stored at some radial position or less axial positions. When dealing with the LES method, the problem of storage is virtually unsolvable since very high frequencies can be of interest (vortex shedding, small flow scales in the boundary layers, etc) and only the physical sense of researchers and engineers can help to reduce the database (signal extraction in limited regions, frequency filtering, etc).

The mesh generation is also a critical step for turbomachine applications, since a grid including most technological effects (rotor tip clearance, hub leakage flow, cooling holes in the case of turbines, etc) requires

1–2 millions points for one blade passage (i.e. around 500 millions for a whole high-pressure compressor). Most current applications consider identical blade profiles (except for the next generation of turbofan), meaning the grid can be generated only for one given blade and is then simply duplicated around the annulus. Moreover, very efficient automatic meshing tools have been developed in the past few years, leading to major saving of time (no more than one week is necessary to design the grid, even for a complex multistage configuration). However, the difficulty comes from the technological effects. On the one hand unstructured grids are efficient to take into account such complex geometries but classical schemes with this kind of grids are limited to low orders (usually <3) that can be insufficient, especially near the walls where very strong flow gradients occur. On the other hand, classical numerical schemes associated with structured grids are accurate in the near wall region, but some technological effects such as cooling holes located on the turbine blade wall cannot be directly simulated. Indeed, there is no consensus in the literature about the best way to tackle with the complex geometry of internal components, but hybrid grids are probably a way to consider (Yang *et al* 2005).

5.2. Simulation of unsteady flows in a multistage compressor

A direct application of HPC is the simulation of deterministic unsteady flows in a multistage compressor (all blade passages are represented). The major interest from the scientific point of view is to understand the development of aerodynamic instabilities and the impact of rotor–stator interactions on the system performance. Some authors have recently shown that a correct physical description of these complex unsteady flows can be obtained by performing unsteady flow simulations with high-end computing platforms. For example, aerodynamic instabilities have been successfully simulated in a whole subsonic compressor stage (Gourdain *et al* 2006) and in a full helicopter multistage compressor (Hathaway *et al* 2004). The stability limit (the so-called 'surge' line) has been successfully estimated by these unsteady flow simulations, while current industrial CFD methods such as steady RANS simulations coupled with a mixing plane model (Denton and Singh 1979) usually fail to predict this critical design parameter.

The test case considered to highlight benefits of HPC is a research multistage compressor dedicated to aero-thermal and aerodynamic studies. This $3^1/_2$-stage axial compressor, named CREATE (Compresseur de Recherche pour l'Etude des effets Aérodynamiques et Technologiques), is representative of high-pressure compressor median-rear blocks of modern turbojet engines. The compressor along with measurement sections and a partial view of the numerical domain are presented in figure 10. The rotor shaft is driven at the design speed of 11 543 rpm by a 2 MW dc-drive coupled with a gearbox. At this rotational speed, the first rotor tip speed is 313 m s^{-1} and the mean Reynolds number based on the chord is 10^6. More information about the experimental facility is available in Ottavy *et al* (2003). Simulations are performed with *elsA* using steady RANS or URANS methods. The mesh represents either the whole compressor with 134 M points (unsteady flow is simulated in all the blade passages) or only one passage with 1.5 M points (steady flow is simulated in one passage of each blade row). Spatial convective fluxes are computed with Roe's scheme and the van Leer's monotone upstream-centered scheme for the conservation law (MUSCL) approach is used to reach third-order accuracy in space (van Leer 1979). A second order dual time stepping (DTS) method is applied for the time integration (Jameson 1991). The time marching for the inner loop is performed using an efficient implicit time integration scheme, based on the backward Euler scheme and a scalar LU-SSOR method (Yoon and Jameson 1987). The turbulent viscosity is computed with the two equations model of Wilcox (1988) based on a k-omega formulation. More information about the numerical method can be found in Gourdain *et al* (2009a). The flow is simulated at design condition using 512 PUs and at near stall conditions using 4096 PUs. Simulations are performed thanks to IBM Blue Gene/P (EDF) and IBM Blue Gene/L (CERFACS) platforms.

An overview of the required computational effort is shown in table 3 for RANS (industrial) and URANS methods. The fine mesh corresponds to the industrial standard in terms of grid density, while the coarse mesh corresponds to a quite low-quality mesh. The computational effort is based on performance obtained with a IBM Blue Gene/P supercomputer and at design operating conditions. This example shows that the computational effort should be increased by a factor of 5000 with respect to steady simulations to obtain a correct description of rotor–stator interactions. At near stall conditions (off-design), the computational effort can be increased by an additional factor equal to 10. This kind of unsteady flow simulation is thus a demonstrative step and can be performed only with a HPC strategy. In the future, such advanced simulations

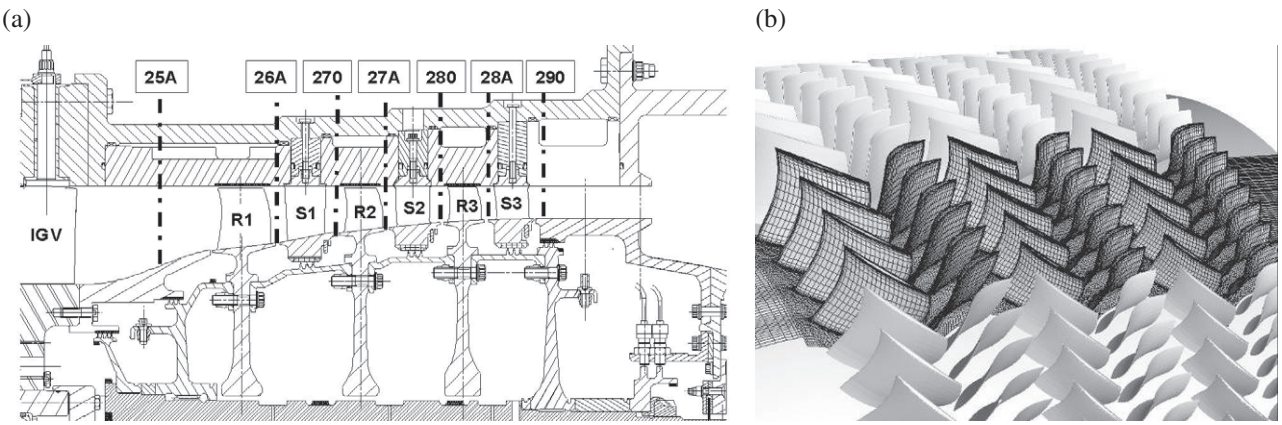

Figure 10. View of the experimental compressor (a) and the numerical domain (b).

Table 3. Typical computing effort related to steady RANS and unsteady RANS simulations in a multistage compressor (IBM Blue Gene/P and design operating conditions).

	Steady RANS		Unsteady RANS	
	Coarse mesh	Fine mesh	Coarse mesh	Fine mesh
Number of points	1 500 000	6 700 000	134 000 000	600 000 000
Total CPU (hours) effort for converged state (RANS) or 5 ms of simulation (URANS)	55	220	280 000	1 250 000

will clearly help industries to improve the efficiency of their products by better predicting multistage interactions and related unsteady effects.

The steady RANS approach combined with the steady mixing plane (SMP) assumption is compared to the URANS method applied to the unsteady whole compressor (UWC). The SMP method considers only a single passage for each blade row, so only 1.5 M of points are required for the whole grid. On IBM Blue Gene architectures with 16 PUs, the computational time to obtain the solution with the SMP method is 55 CPU hours at design conditions and 80 CPU hours at off-design conditions. On the same architecture, the UWC approach requires 280 000 CPU hours to obtain a periodic solution at design conditions (with 512 PUs) and 4 000 000 CPU hours at off-design conditions to simulate a rotating stall phenomenon (with 4096 PUs). The UWC method provides a result at design conditions in 22 days, while 40 days are needed at off-design conditions. Despite this large CPU time, this work demonstrates that simulations of complex unsteady flows in multistage turbomachines are now feasible by using HPC.

A comparison between SMP and UWC time-averaged results is done to estimate the benefit of such expensive simulations. The relative difference between the two solutions ψ is computed at sections 28A (see figure 10), based on the axial velocity V_x flow field:

$$\psi = \left| \frac{V_x(\mathrm{SMP}) - \bar{V}_x(\mathrm{UWC})}{V_x(\mathrm{SMP})} \right|. \tag{1}$$

Results are presented in figure 11. For visibility reasons, only a part of the circumference is shown (a 22.5° sector). In the case of purely steady flows, the SMP method should provide results close to the time-averaged results obtained with the UWC method. Observed differences are thus related directly to mean unsteady flow effects. The most visible feature is that the time-averaged flow does not remain identical in all the compressor passages, showing the impact of rotor-stator interactions on the time-averaged solution. At design conditions, the SMP approach overestimates the axial velocity inside the rotor wakes (white regions in figure 11(a)) and the sequence of over-predicted and under-predicted regions of axial velocity tends to compensate near the

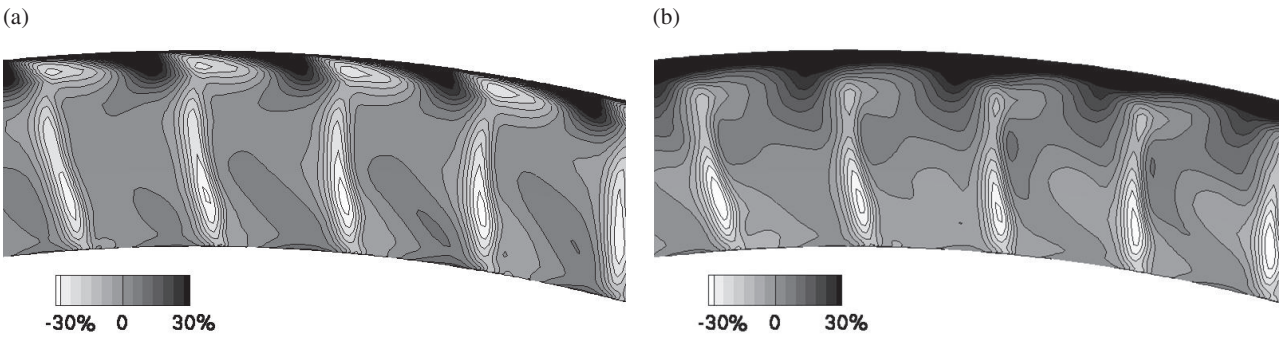

Figure 11. Relative differences measured on mean values of axial velocity at design (a) and off-design (b) conditions between the unsteady and the steady RANS methods (third stage, *elsA*).

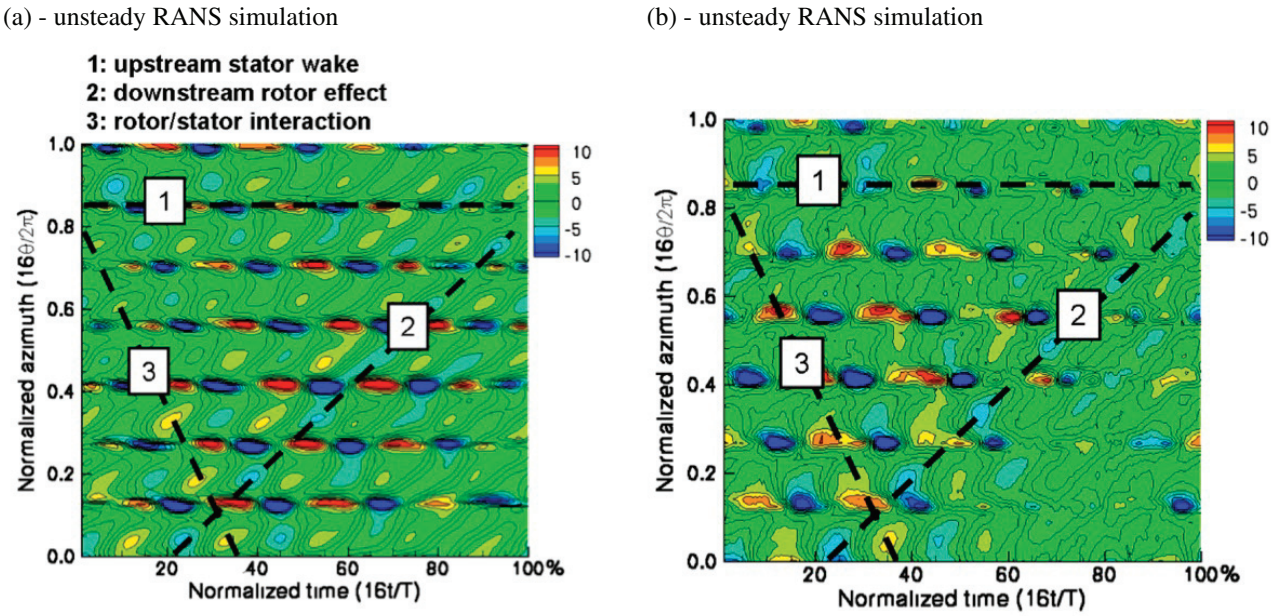

Figure 12. Fluctuation of velocity angle α' (in $^\circ$) at section 280 and at design operating point (near the casing, $h/H = 83\%$): (a) URANS simulation*(elsA)* and (b) measurements.

casing. Impact of unsteady flow effects is found to be reasonable at design conditions and the SMP method predicts global values reasonably well. However, as expected, this conclusion is no longer true at off-design conditions since the mixing plane method is far from its application domain. Development of reversed flow regions near the casing are responsible for the development of aerodynamic instabilities (the so-called 'rotating stall' phenomenon) that cannot be predicted by a steady flow method. In these conditions (figure 11(b)), large differences are induced by the SMP approach, especially near the casing (more than 30% of the time-averaged axial velocity). A simulation at design operating conditions is compared with experiments in figure 12 to point out the capacity of the UWC method to correctly estimate the flow fluctuations in this complex system. As indicated in figure 12, the simulation correctly predicts most flow features, including the unsteadiness level and the flow features generated by rotor–stator interactions.

It clearly shows that very promising improvements of the system design can be expected from these simulations since the prediction of such unsteady effects is of wide interest both for the scientific and the industrial communities. Instantaneous velocity flow fields are presented in figure 13 at design and off-design conditions. White regions represent the low momentum flows induced by the tip leakage flows. Solutions to

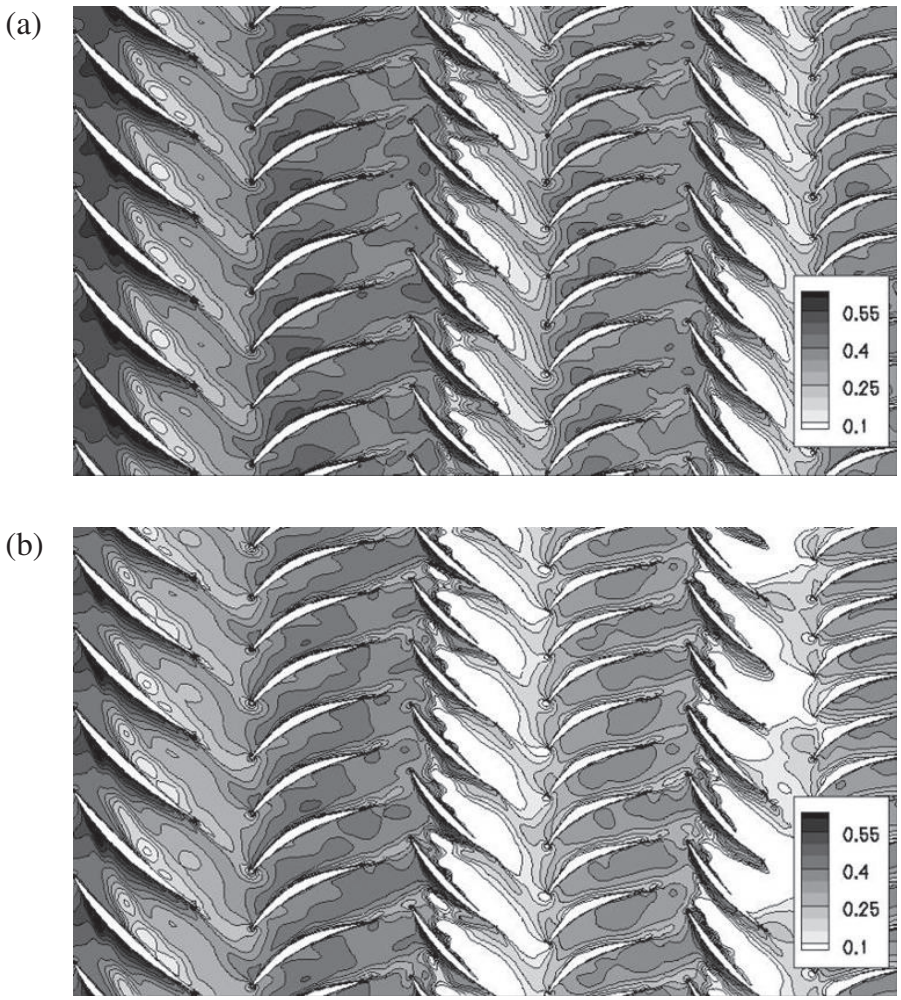

Figure 13. Instantaneous flow fields of axial velocity computed with *elsA* at design (a) and off-design (b) operating conditions in a multistage compressor (near the casing, $h/H = 83\%$).

control these flows have already been proposed in the literature (Hathaway 2006, Gourdain *et al* 2009b) but are usually applied on isolated blade row configurations. Based on the results obtained with the UWC approach, the flow in the last stage has been identified as the most penalizing for the compressor stability. The reason is that wakes generated by the stator of the second stage disturb the flow near the casing in the rotor of the third stage, leading to the development of a rotating instability. These observations are used to design an original control solution based on honeycomb. Simulations are then performed to assess the capacity of this control to improve the compressor performance. Results show that rotating stall is suppressed by the control and the compressor stability in terms of mass flow is increased by +50%. This work clearly underlines the capacity of flow solvers, coupled with HPC, to simulate unsteady flows in a whole multistage compressor. Even if the cost ratios between industrial standards and this advanced work ranges from 5000 to 50 000 (depending on the operating conditions), interesting information have been obtained about unsteady flow effects, validity of RANS methods and aerodynamic instabilities.

5.3. High-fidelity simulations of unsteady flows in turbines

Around the design conditions, industrial simulations based on RANS methods give satisfactory results for mean performance such as pressure ratio and efficiency. However, local flows (secondary flows, boundary layer separations, etc) have a major influence on the flow unsteadiness and cannot be correctly predicted with standard steady numerical simulations. An accurate estimation of wall temperature and heat coefficient, which

Table 4. Computing efforts related to different methods to simulate the flow in an IGV of a high pressure turbine (flow solver *elsA*, SGI Altix platform).

	RANS	URANS	LES
Number of cells	795 000	795 000	6 370 000
Total CPU (hours) effort for converged state (RANS) or 10 ms of simulation (URANS/LES)	21	230	2000

are design parameters for turbines, remains out of range with these methods. In this context, unsteady flow simulations that solve a part of the turbulent spectrum (LES) emerge as a promising way to increase the reliability of flow solvers.

An IGV of a high-pressure turbine has been chosen to assess the capacity of LES to correctly reproduce the main flow features. This configuration has been experimentally studied at VKI by Sieverding *et al* (2003). The Reynolds number based on chord is $Re = 2.8 \times 10^6$ (that corresponds to a high value for LES) and the outlet Mach number is 0.79. Measurements highlight the presence of large coherent structures in the turbine blade wakes (von Karman vortices) that largely affect the pressure distribution around the trailing edge. The objective is to investigate the capacity of (U)RANS and LES methods to reproduce the flow in the turbine configuration, especially near the trailing edge. (U)RANS calculations are performed using the two equations turbulence model of Wilcox (1988) and LES uses the subgrid model of Smagorinsky (1963). The convective fluxes are computed with the third-order AUSM + scheme proposed by Liou (1996) and the time integration is performed with an implicit DTS method (simulations are performed with the *elsA* software). The computational effort is compared for the three approaches (table 4). Based on the assumption that the flow is mainly 2D, the computational domain represents only 1% of the blade span. The grid used for (U)RANS simulations is composed of 0.795 M cells and corresponds to a very fine grid, insuring that mesh convergence is obtained. The grid considered for LES is more refined and 6.37 M cells are used. The cost ratio (in terms of computational time) between RANS and URANS is 11 and the cost ratio between URANS and LES is 9. The cost of LES is moderate because the implicit scheme for time integration allows using large time steps. A comparison of instantaneous flow field of the density gradient is displayed in figure 14. As expected, the steady RANS approach is not able to reproduce the vortex shedding (which is a fully unsteady phenomenon). But the major discrepancy with other methods is that RANS predicts the development of a non-physical shock-wave on the suction side of the blade, indicating a difficulty to correctly estimate the boundary layer thickness. Other approaches predict the development of a vortex shedding behind the trailing edge and a complex system of acoustic waves is observed in the blade passage. Periodic waves are emitted from the trailing edge in coherence with the von Karman vortices. These waves are reflected on the blade suction side and interact with the vortex shedding. On the one hand, the URANS method predicts the vortex shedding but other flow features (such as acoustic waves) are rapidly damped by artificial viscosity. The wake vortices are also largely damped and stretched due to anisotropic artificial viscosity (related to the mesh quality). On the other hand, LES demonstrates its capacity to transport the flow vortices and acoustic waves. Two LES are performed: one with the structured flow solver *elsA* (figure 14(c)) and the second one with the unstructured code AVBP (figure 14(d)). Due to a higher grid resolution, the LES done with *elsA* better propagates the waves and shows a more complex flow than the LES obtained with AVBP. However, far from the trailing edge, the structured mesh exhibits large disparities of cell dimensions and the vortex shedding is partially dissipated illustrating the potential difficulties and advantages of structured and unstructured flow solvers for such applications. Note that this specific phenomenon is not observed with AVBP since it uses a third-order low dissipative scheme (Colin and Rudgyard 2000) and a constant grid density in the vortex shedding direction.

A comparison of the results obtained with *elsA* are then compared with experimental data. First, the isentropic Mach number is shown in figure 15 for RANS and LES. As previously mentioned, RANS predicts a shock-wave on the suction side that is not reported by the experimental work (at an axial position equal to 0.6). RANS also under-predicts the value of the isentropic Mach number near the blade leading edge on the pressure side. When comparing with measurements, LES results agree well everywhere around the blade. URANS gives very similar results to LES regarding the isentropic Mach number (results not shown).

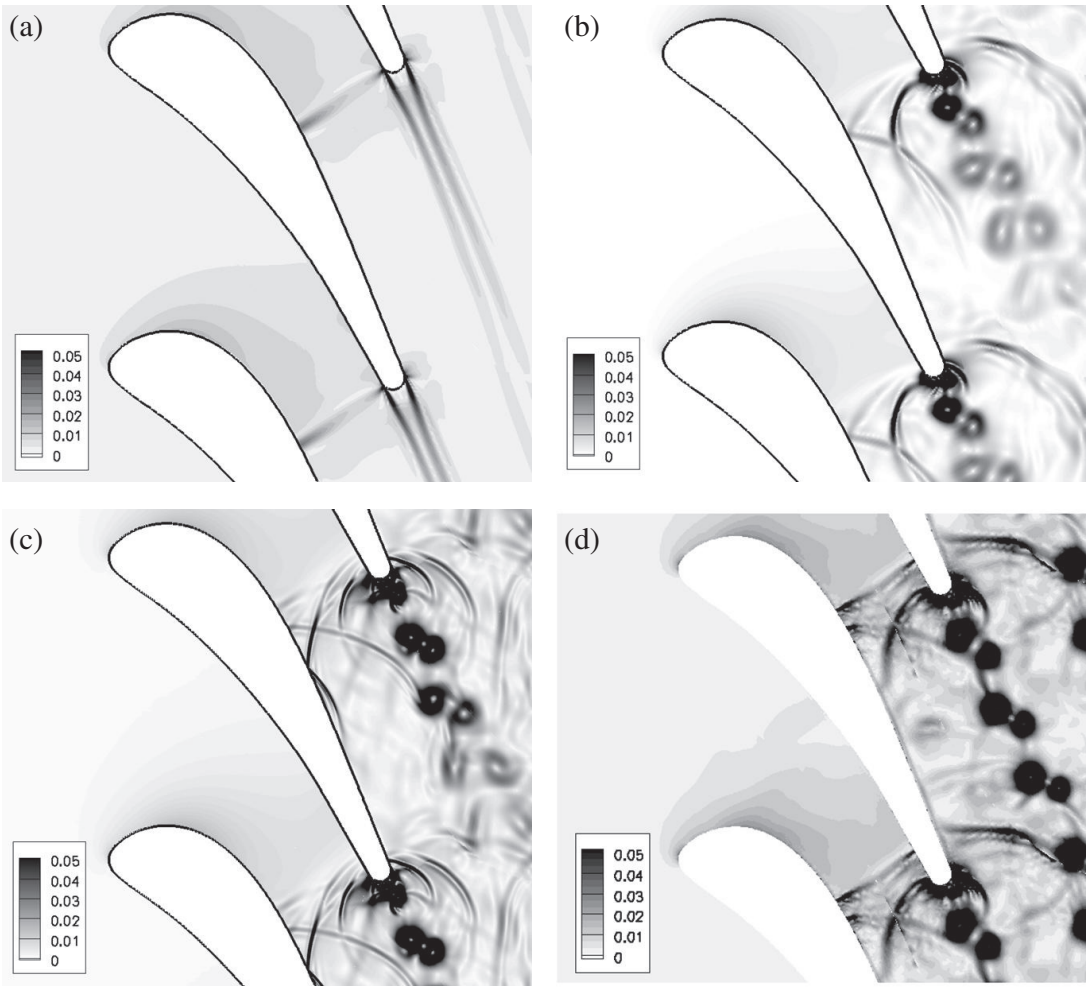

Figure 14. Comparisons of instantaneous flow field of grad ρ in the IGV of a high-pressure turbine: (a) RANS with *elsA*, (b) URANS with *elsA,* (c) LES with *elsA* and (d) LES with AVBP.

Another comparison is done between URANS and LES results, regarding the vortex shedding frequency. This parameter is expressed with the Strouhal number $St = f\ D/U$, where f is the frequency, D the trailing edge diameter and U the external velocity. The unsteady signals of axial velocity are registered at 20% of the blade chord behind the trailing edge. As shown in figure 15, URANS and LES do not predict identical unsteady flows. URANS produces only one frequency, while LES predicts a much more complex flow behavior with a large frequency spectrum, and with coherent flow features that also influence the trailing-edge vortex shedding. Actually, the Strouhal number estimated with the URANS method is far from the experimental value (+26%) while LES gives a correct estimation of this parameter (+4%). This section shows that high-fidelity simulations of turbomachine flows are now feasible in a reasonable amount of time (few days). In the future, more complex configurations such as rotor–stator configurations or combustion chamber–turbine blade rows should be investigated with LES.

5.4. Aero-thermal simulations of cooling blades

As already explained, one of the most important challenges today for turbine designers is the prediction of heat transfer on engine blade walls. The life duration of the turbine components directly depends on the wall heat transfer levels and therefore designers imperatively need an accurate estimation tool: a slight difference on the temperature prediction ($<1\%$) at mid-span of a blade corresponds to a reduction of its life duration by a factor 2. The difficulty is that complex flows observed in the turbine environment cannot be efficiently

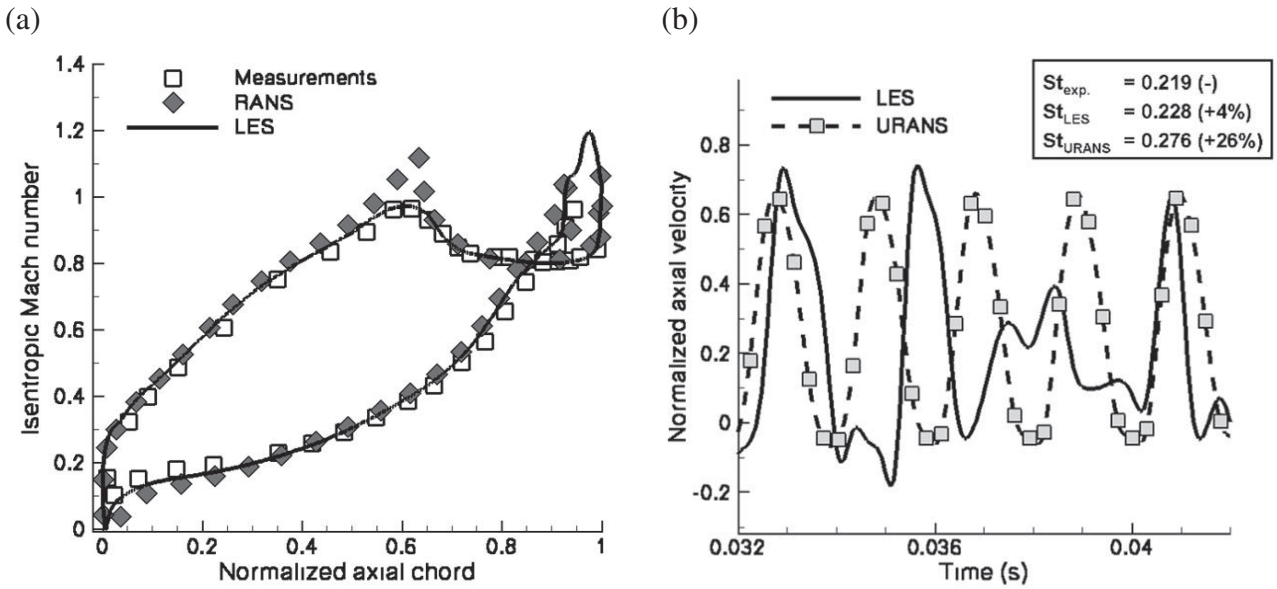

Figure 15. Vortex shedding simulations performed with *elsA* behind the turbine IGV—comparisons of the isentropic Mach number (a) and unsteady axial velocity signals (b).

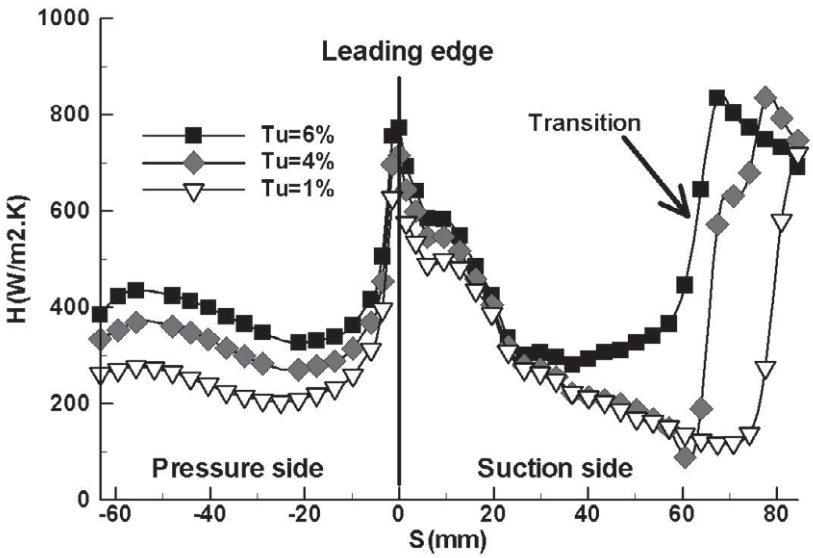

Figure 16. Effect of the laminar to turbulent transition on the heat transfer coefficient H (IGV configuration, Arts 1990).

computed with (U)RANS methods, especially when regarding thermal effects. For example, laminar to turbulent transition, hot spot incoming from the combustion chamber and temperature gradient at walls are among the difficulties that CFD solvers have to address. The turbulence effect on the heat transfer coefficient H is shown in figure 16 for an IGV of a highly loaded transonic turbine, experimentally studied by Arts *et al* (1990). Today, RANS simulations (even with transition models) performed on this configuration exhibit a very poor predictive capacity and lead to errors higher than 50% on the value of the heat transfer coefficient. As shown in figure 16, one of the difficulties comes from the strong impact of the inlet turbulence level on the transition region (mainly driven by the shock-wave position). A way to overcome this difficulty could be the use of LES for solving the flow around the blade (and possibly inside the blade in the case of internal cooling devices), coupled with a thermal solver for the solid inside the blade.

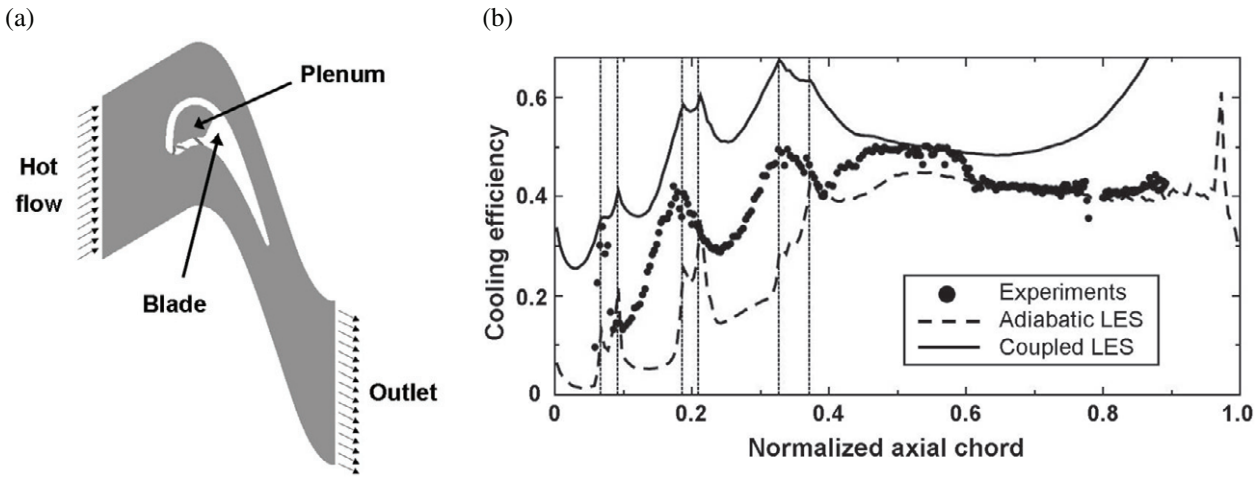

Figure 17. Aero-thermal simulation of a cooling turbine blade (a) and time-averaged solution of cooling efficiency (b) obtained with LES (AVBP simulations).

Such a computing strategy is applied to simulate the flow in the T120 blade cascade and assess the interest of aero-thermal simulations with respect to adiabatic simulations (figure 17(a)). Information about measurements and experimental configuration can be found in Homeier and Haselbach (2005). The outlet Mach number is $M = 0.87$ and the Reynolds number is $Re = 3.9 \times 10^5$. This highly loaded high-pressure turbine airfoil is designed to obtain a large separation of the boundary layer on the pressure side. The cooling device is composed of three cylindrical holes located on the pressure side. The 3D flow simulation is performed with the LES code AVBP, by using a low dissipative third-order explicit scheme (Colin and Rudgyard 2000). One blade passage is considered with spatial periodicity conditions. Wall effects induced by hub and casing are neglected by simulating only a part of the blade span (8 mm). The number of tetrahedral cells used to represent the flow domain is 6.5 M cells and the solid domain contains 600 000 cells.

Two simulations are performed: the first one considers adiabatic walls and the second one uses a coupling method to exchange data between aerodynamic and a thermal flow solvers (Duchaine *et al* 2009). For the coupled simulation, the quasi-periodic state is obtained in 4800 CPU hours, by using 32 PUs of a scalar IBM JS21 computing platform (30 PUs are allocated to the flow solver and 2 for the thermal code). Time-averaged results are shown in figure 17(b) for the cooling efficiency θ, defined as

$$\theta = \frac{T_\infty - T_{\text{wall}}}{T_\infty - T_{\text{cj}}}, \tag{2}$$

where T_∞ is the upstream temperature ($= 333.15$ K) and T_{cj} is the cold jet temperature (i.e. the plenum temperature $= 303.15$ K).

The cooling efficiency predicted with the adiabatic LES is under-estimated in the region of the cooling holes with respect to experimental measurements, indicating that the wall temperature is over-predicted. Some flow features such as the shock wave position are also not perfectly computed with adiabatic LES. However, the overall performance of the cooling device is quite well estimated, except near the trailing edge. Results obtained with the coupled LES strategy indicate that the wall temperature on the blade pressure side is reduced compared to the adiabatic LES. The consequence is that the cooling efficiency is largely increased (by a factor of 2 near the cooling holes). The global shape of the cooling efficiency computed with the coupled LES correctly matches experimental data, but the wall temperature is still not correctly predicted (probably due to insufficient wall resolution). While an important work is still required, this study indicates that an aero-thermal simulation based on a code coupling strategy is now feasible by using modern parallel computing platforms, showing promising results to accurately estimate the wall temperature in turbine blades.

6. Conclusion

Examples of CFD applications that benefit from HPC have been presented for aeronautic and propulsion domains. The most common architecture used for these applications is the massively parallel platform that allows using thousands of concurrently computing cores to solve a problem. A first difficulty in industrial context is related to the determination of the parallel efficiency. A good speed-up indicates either a very good parallel efficiency or a poor sequential efficiency. The normalized speed-up is commonly used instead of the strong speed-up that cannot be evaluated in most cases (due to memory limitations). However, this definition can lead to misunderstanding if the speed-up is normalized with a low-efficiency computation. Finally, one of the most relevant indicators of performance is the real time needed to obtain the solution. A guideline has been proposed to assess the performance of the computing platforms (vector, scalar, etc), mainly based on the electric consumption. It indicates that massively parallel platforms are very efficient to deliver a high-computing power at a low electric consumption (i.e. exploitation cost).

Thanks to the computing power provided by high-end computing platforms (both scalar and vector computers), very interesting numerical results have been obtained for industrial systems with a particular interest for the simulation of unsteady flow effects. Such high-fidelity simulations will help to improve the reliability of numerical solutions and better understand complex flow phenomena such as aerodynamic and thermo-acoustic instability. The development of greener, quieter and more economic products is thus clearly dependent on the capacity to obtain the best performance of flow solvers on computing platforms. For aeronautic industry, CFD is already a key technology for design and a reduction of the development time and costs. Based on several examples of industrial applications presented in this paper, it is clear that numerical simulation is also a very effective tool for physical investigations and design validation. Simulations of flight, off-design operating conditions or certification are among the challenges that CFD proposes to address in the next few years by using HPC. However, important work is still necessary to improve the computing algorithms and the code performance on current (and future) computing platforms. Future challenges for the scientific community will be the use of HPC for the simulation of multi-component configurations (such as combustor and turbine) and multi-physics applications (aerodynamic, thermic, combustion, elasticity, etc). The key point will be the development of efficient coupling strategies and the capacity to deal with different kinds of simulation (coupling between (U)RANS and LES, DES, etc). Another critical issue is the adaptation of the pre/post processing tools to these very efficient simulation methods. In fact, as discussed in this paper, the limiting factor for the application of CFD to industrial problems could become the grid generation and the analysis of a large database. If no care is brought to these requirements, the use of very costly and powerful computers can become useless in the next few years.

Acknowledgments

We thank the CERFACS' teams involved in the management and maintenance of computing platforms (in particular the CSG group). Furthermore, we acknowledge all people who contributed to this paper by means of discussions, support, direct help or corrections (including the anonymous reviewers). We also acknowledge industrial and research partners for supporting code developments and permission for publishing results. In particular, we thank Airbus, Onera, Snecma, Turbomeca, ECL-LMFA and VKI for collaborative work around aircraft, combustion and turbomachine projects. Finally, we acknowledge EDF, Météo-France and GENCI-CINES for providing computational resources that have been used for the applications presented in this paper.

References

Amdahl G 1967 Validity of the single processor approach to achieving large-scale computing capabilities *AFIPS Conf. Proc.* **30** 483–5

Arakawa C, Fleig O, Iida M and Shimooka M 2005 Numerical approach for noise reduction of wind turbine blade tip with earth simulator *J. Earth Simulator* **2** 11–33

Arts T, Lambert de Rouvroit M and Rutherford A W 1990 Aero-thermal investigation of a highly loaded transonic linear turbine guide vane cascade—a test case for inviscid and viscous flow computations *VKI Technical Note* 174

Baum M, Haworth D, Poinsot T and Darabiha N 1994 Direct numerical simulation $H_2/O_2/N_2$ flames with complex chemistry in two-dimensional turbulent flows *J. Fluid Mech.* **281** 1–32

Boileau M, Staffelbach G, Cuenot B, Poinsot T and Bérat C 2008 LES of an ignition sequence in a gas turbine engine *J. Combust. Flame* **154** 2–22

Bogey C and Bailly C 2006 Investigation of downstream and sideline subsonic jet noise using large eddy simulations *Theor. Comput. Fluid Dyn.* **20** 23–40

Boudier G, Gicquel L, Poinsot T, Bissières D and Berat C 2007 Comparison of LES, RANS and experiments in an aeronautical gas turbine combustion chamber *Proc. Comb. Inst.* **31** 3075–82

Boudier G, Gicquel L Y M and Poinsot T 2008 Effects of mesh resolution on large eddy simulation of reacting flows in complex geometry combustors *J. Combust. Flame* **155** 196–214

Boudier G, Lamarque N, Staffelbach G, Gicquel L Y M and Poinsot T 2009 Thermo-acoustic stability of a helicopter gas turbine combustor using large eddy simulations *Int. J. Aeroacoust.* **8** 69–94

Brunet V and Deck S 2008 Zonal-detached eddy simulation of transonic buffet on a civil aircraft type configuration *46th AIAA Aerospace Science Meeting and Exhibit (Reno, USA)*

Cambier L and Veuillot J-P 2008 Status of the elsA CFD software for flow simulation and multidisciplinary applications *46th AIAA Aerospace Science Meeting and Exhibit (Reno, USA)*

Chen J H *et al* 2009 Terascale direct numerical simulations of turbulent combustion using S3D *Comput. Sci. Discov.* **2** 015001

Chen J H, Hawkes E R, Sankaran R, Mason S D and Im H G 2006 Direct numerical simulation of ignition front propagation in a constant volume with temperature inhomogeneities, part I: fundamental analysis and diagnostics *J. Combust. Flame* **145** 128–44

Colin O and Rudgyard M 2000 Development of high-order Taylor–Galerkin schemes for unsteady calculations *J. Comput. Phys.* **162** 338–71

Deck S 2005 Numerical simulation of transonic buffet over a supercritical airfoil 2005 *AIAA J.* **43** 1556–66

Delbove J 2006 Unsteady simulations for flutter prediction *Proc. Third Int. Conf. on CFD—ICCFD3* pp 205–10

Delville J, Ukeiley L, Cordier L, Bonnet J and Glauser M 1999 Examination of large-scale structures in a turbulent plane mixing layer. Part 1. Proper orthogonal decomposition *J. Fluid Mech.* **391** 92–122

Denton J D and Singh U K 1979 Time marching methods for turbomachinery flow calculations *(VKI Lecture Series)* (Brussels: Von Karman Institute)

Domercq O and Escuret J-F 2007 Tip clearance effect on high-pressure compressor stage matching *J. Power Energy* **221** 759–67

Duchaine F, Corpron A, Pons L, Moureau V, Nicoud F and Poinsot T 2009 Development and assessment of a coupled strategy for conjugate heat transfer with large eddy simulation. Application to a cooled turbine blade *Int. J. Heat Fluid Flow* at press (doi: 10.1016/j.ijheatfluidflow.2009.07.004)

Février P, Simonin O and Squires K D 2005 Partitioning of particle velocities in gas-solid turbulent flows into a continuous field and a spatially uncorrelated random distribution: theoretical formalism and numerical study *J. Fluid Mech.* **533** 1–46

Freund J B 2001 Noise sources in a low-Reynolds number turbulent jet at mach 0.9 *J. Fluid Mech.* **438** 277–305

Garcia M 2009 Développement et validation du formalisme Euler–Lagrange dans un solveur parallèle et non-structuré pour la simulation aux grandes échelles *PhD thesis* University of Toulouse

Gicquel L Y M, Staffelbach G, Cuenot B and Poinsot T 2008 Large eddy simulations of turbulent reacting flows in real burners: the status and challenges, SciDAC 2008 *J. Phys.: Conf. Ser.* **125** 17

Gökalp I, Chauveau C, Morin C, Vieille B and Birouk M 2000 Improving droplet breakup and vaporisation models, including high pressure and turbulence effects *Atomization Sprays* **10** 475–510

Gopinath A, van der Weide E, Alonso J J, Jameson A, Ekici K and Hall K C 2007 Three-dimensional unsteady multi-stage turbomachinery simulations using the harmonic balance technique *45th AIAA Aerospace Sciences Meeting and Exhibit (Reno, USA)* paper 2007-0892

Gourdain N, Burguburu S, Michon G-J, Ouayahya N, Leboeuf F and Plot S 2006 About the numerical simulation of rotating stall mechanisms in axial compressors *ASME Turbo Expo* paper GT2006-90223 (Barcelona, Spain)

Gourdain N, Ottavy X and Vouillarmet A 2009a Experimental and numerical investigation of unsteady flows in a high-speed three stages compressor *8th European Turbomachinery Conference* paper 2009-107 (Graz, Austria)

Gourdain N and Leboeuf F 2009b Unsteady simulation of an axial compressor stage with casing and blade passive treatments *J. Turbomachinery* **130**

Gourdain N, Gicquel L, Montagnac M, Vermorel O, Gazaix M, Staffelbach G, Garcia M, Boussuge J-F and Poinsot T 2009c High performance computing of complex geometries: I. Methods *Comput. Sci. Discov.* **2** 015003

Hall K C, Dowell E H and Thomas J P 2003 Nonlinear reduced order modeling of limit cycle oscillations of aircraft wings and wing/store *Technical Report* Duke University

Hathaway M D, Herrick G, Chen J and Webster R 2004 Time accurate unsteady simulation of the stall inception process in the compression system of a US army helicopter gas turbine engine *Proc. DoD Users Group Conf. (Washington, DC, USA)* pp 182–93

Hathaway M D 2006 Passive endwall treatments for enhancing stability *Advances in Axial Compressor Aerodynamics (VKI Lecture Series)* (Brussels: Von Karman Institute)

He L 1990 An Euler solution for unsteady flows around oscillating blades *J. Turbomach.* **112** 714–22

He L 1997 Computational study of rotating-stall inception in axial compressors *J. Propulsion Power* **13** 31–8

Homeier L and Haselbach F 2005 Film cooling of highly loaded blades *17th ISABE Conf. paper AIAA-2005-1114 (Munich, Germany)*

Im H G, Chen J H and Chen J-Y 1999 Chemical response of methane/air diffusion flames to unsteady strain rate *J. Combust. Flame* **118** 204–12

James S, Zhu J and Anand M 2006 Large eddy simulations as a design tool for gas turbine combustion systems *AIAA J.* **44** 674–86

Jameson A 1991 Time dependent calculations using multigrid, with applications to unsteady flows past airfoils and wings *10th AIAA Computational Fluid Dynamics Conf.* paper 91–1596

Lefebvre A H 1999 *Gas Turbine Combustion* 2nd edn (London: Taylor and Francis)

Lignell D O, Chen J H, Smith P J, Lu T and Law C K 2007 The effect of flame structure on soot formation and transport in turbulent nonpremixed flames using direct numerical simulation *J. Combust. Flame* **151** 2–28

Liou M S 1996 A sequel to AUSM: AUSM+ *J. Comput. Phys.* **129** 364–82

Lumley J L 1978 Computational modelling of turbulent flows *Adv. Appl. Mech.* **18** 123–76

Medic G, Kalitzin G, You D, Herrmann M, Ham F, van der Weide E, Pitsch H and Alonso J J 2006 Integrated RANS/LES computations of turbulent flow through a turbofan jet engine *Annual Research Briefs* Stanford University

Mendez S and Nicoud F 2008 Large eddy simulation of a bi-periodic turbulent flow with effusion *J. Fluid Mech.* **598** 27–65

Miller R S and Bellan J 1999 Direct numerical simulation of a confined three-dimensional gas mixing layer with one evaporating hydrocarbon-droplet laden stream *J. Fluid Mech.* **384** 293–338

Mizobuchi Y, Shinjo J, Ogawa S and Takeno T 2005 A numerical study on the formation of diffusion flame islands in a turbulent hydrogen jet lifted flame *Proc. Combust. Inst.* **30** 611–9

Moin P and Apte S 2006 Large eddy simulation of realistic turbine combustors *AIAA J.* **44** 698–708

Nybelen L and Deniau H 2009 Spatial simulation of a co-rotating vortex merging process in unstable conditions (FAR-Wake workshop, Marseille, France, 2008) *Int. J. Methods in Fluids* **61** 23–56

Ottavy X, Trébinjac I, Vouillarmet A and Arnaud D 2003 Laser measurements in high speed compressors for rotor–stator interaction analysis *Proc. ISAIF 6th, Shanghai, April 2003*

Poinsot T and Veynante D 2005 *Theoretical and Numerical Combustion* 2nd edn (R T Edwards)

Reveillon J and Vervisch L 2000 Accounting for spray vaporization in non-premixed turbulent combustion modeling: a single droplet model *J. Combust. Flame* **121** 75–90

Rossow C and Kroll N 2008 High performance computing serves aerospace engineering: opportunities for next generation product development *46th AIAA Aerospace Science Meeting and Exhibit (Reno, USA)*

Schluter J, Apte S, Kalitzin G and van der Weide E 2005 Large-scale integrated LES-RANS simulations of a gas turbine engine *Annual Research Briefs* Stanford University

Seiser R, Frank J H, Liu S, Chen J H, Sigurdsson R J and Seshadri K 2005 Ignition of Hydrogen in unsteady nonpremixed flows *Proc. Combust. Inst.* **30** 423–30

Sicot F, Puigt G and Montagnac M 2008 Block–Jacobi implicit algorithms for the time spectral method *AIAA J.* **46** 3080–89

Sieverding C H, Richard H and Desse J-M 2003 Turbine blade trailing edge flow characteristics at high subsonic outlet mach number *J. Turbomachin.* **125** 298–309

Smagorinsky J S 1963 General circulation experiments with the primitive equations: I. the basic experiment *Mon. Weather Rev.* **91** 99–163

Smith C, Beratlis N, Squires K, Balaras E and Tsunoda M 2008 Direct numerical simulations of the flow around a golf ball: effect of rotation, *61st Annual Meeting of the APS Division of Fluid Dynamics (San Antonio, USA)*

Soufiani A and Djavdan E 1994 A comparison between weighted sum of gray gases and statistical narrow-band radiation models for combustion applications *J. Combust. Flame* **97** 240–50

Spalart P R and Allmaras S R 1992 A one equation turbulence model for aerodynamic flows *AIAA paper* 92-0439

Spalart P R, Jou W-H, Stretlets M and Allmaras S R 1997 Comments on the feasibility of LES for wings and on the hybrid RANS/LES approach, advances in DNS/LES *Proc. First AFOSR Int. Conf. on DNS/LES*

Staffelbach G, Gicquel L Y M, Boudier G and Poinsot T 2009 Large eddy simulation of self excited azimuthal modes in annular combustors *Proc. Combustion Institute (Pittsburgh, USA)* vol 32

Tennekes H and Lumley J L 1972 *A First Course in Turbulence* (Cambridge, MA: MIT Press)

Truffin K and Poinsot T 2005 Comparison and extension of methods for acoustic identification of burners *J. Comb. Flame* **142** 388–400

van der Weide E, Kalitzin G, Schluter J and Alonso J J 2006 Unsteady turbomachinery computations using massively parallel platforms *44th AIAA Aerospace Sciences Meeting and Exhibit (Reno, USA)*

van der Weide E, Mattson K, Svärd M, Nordström J and Gong J 2008 Recent developments of the SBP/SAT discretization technique on multi-block structured grids *Thermal and Fluid Sciences Affiliates Program, Stanford*

van Leer B 1979 Towards the ultimate conservative difference scheme, V: a second order sequel to Godunov's method *J. Comput. Phys.* **32** 101–36

Varga C, Lasheras J and Hopfinger E 2003 Initial breakup of a small-diameter liquid jet by a high-speed gas stream *J. Fluid Mech.* **497** 405–34

Vermorel O, Bédat B, Simonin O and Poinsot T 2003 Numerical study and modelling of turbulence modulation in a particle laden slab flow *J. Turbul.* **4** 025

Weber D P, Wei T Y C, Brewster R A and Rock D T 2000 High fidelity thermal-hydraulic analysis using CFD and massively parallel computers *4th Int. Conf. Supercomputing in Nuclear Applications (Tokyo, Japan)*

Wilcox D C 1988 Reassessment of the scale-determining equation for advanced turbulence models *AIAA J.* **26** 1299–310

Wu Y, Haworth D C, Modest M F and Cuenot B 2005 Direct numerical simulation of turbulence/radiation interaction in premixed combustion systems *Proc. Combust. Inst.* **30** 639–46

Yang H, Nuernberger D and Kersken H-P 2005 Towards excellence in turbomachinery CFD: a hybrid structured-unstructured RANS solver *J. Turbomachinery* **128** 390–402

Yoon S and Jameson A 1987 An LU-SSOR scheme for the euler and Navier–Stokes equations *25th AIAA Aerospace Sciences Meeting (Reno, USA)* paper 87-0600