On large scale turbomachinery computations

By E. van der Weide, G. Kalitzin, J. Schluter, G. Medic AND J.J. Alonso

1. Motivation and objectives

The three main components of a modern turbofan gas-turbine aircraft engine are the compressor (which includes the fan), the combustor, and the turbine (Fig. 1). The compressor and the turbine consist of a number of individual stages, each of which consists of a rotor that rotates on a shaft, and a stator, which is attached to the outer casing and is therefore fixed in space. Modern engines typically have two concentric shafts that allow for a lower rotation speed of the rotors of the low pressure part of the engine and for a higher speed of the high pressure portion. The functionality of the engine is limited by the stability of the high pressure compressor, which has a higher load, partly because of the higher rotation speed. In addition, stability issues can also arise because of the fan.

The present industry standard in the computational modeling of jet engines is driven by the need for fast simulation techniques. Different levels of simulation fidelity are used. The mixing plane approximation by Denton & Singh (1979) is used most commonly for the turbomachinery components. Alternative approaches include the use of phase-lagged boundary conditions by Chen & Barter (1998) and by Wang & Chen (2004), which account for first order unsteady effects, and the deterministic stress approach by Adamczyk (1999). The latter captures the effect of unsteadiness in steady simulations with moderate success.

The focus of this work is on unsteady simulation of turbomachinery components. A large variety of problems of interest to the turbomachinery industry have their origin in the unsteadiness of the flow. This unsteadiness is created either by the relative motion in the geometry, by instabilities in the flow, or by multi-component interactions. Although these large scale unsteady simulations currently consume very large amounts of CPU time, this time will be available in the future at moderate cost. The goal of the Center for Integrated Turbulence Simulations (CITS) is to identify and solve problems with modeling of physics, efficiency of numerical algorithms, parallel scalability and data management present for this type of simulations.

Future fully coupled simulations will, for example, allow the investigation of the effect of heat loads in the turbine with implications for durability, as well as of the effect of transient shocks in the turbine transition duct with implications for engine stability. Both of these are important questions for aircraft engines which cannot be addressed with single component simulations.

2. SUmb flow solver

The flow in turbomachinery components is computed with the compressible, multi-block, structured, cell-centered, parallel RANS flow solver SUmb. The code includes a turbine cooling model, sliding interfaces, moving grids, automatic block splitting and load balancing during run time. The inviscid fluxes are discretized using either central differences and a scalar dissipation scheme by Jameson *et al.* (1981) or an upwind scheme in combination with the approximate Riemann solver of Roe (1982). In the latter scheme,

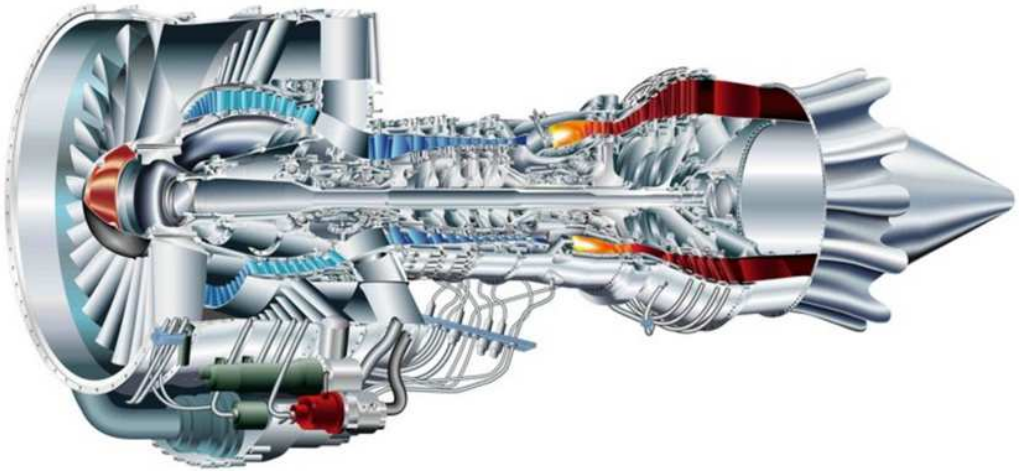

FIGURE 1. Schematic of a Pratt & Whitney aircraft engine.

overshoots in the MUSCL reconstruction are avoided by applying the Van Albeda limiter. A standard central discretization is used for the viscous fluxes. A second order implicit time integration scheme is used to approximate the time derivative. The resulting non-linear system of equations is solved using dual time stepping by Jameson (1998). The convergence is accelerated using a 3W multigrid cycle in combination with local time stepping and a 5-stage Runge Kutta time integration scheme. Usually, 25 3W multi-grid cycles are used per physical time step for the unsteady turbomachinery simulations. Turbulence is modeled with the two-equation k-ω model by Wilcox (1988). The model by Spalart & Allmaras (1994) and the v^2-f model by Durbin (1995) are also available. Adaptive wall functions by Kalitzin *et al.* (2005) and Medic *et al.* (2006) are used for the turbine simulations, and wall integration is used for the compressor.

The computational requirements for the turbine and compressor simulations are as follows. The CPU time required for a typical SUmb simulation on QSC is 0.7 processor hours per one million computational cells per physical time step, assuming 25 multi-grid cycles per physical time step for adequate convergence. The number of time steps per rotor revolution is estimated by assuming that one blade passing should be resolved with approximately 50 time steps. The maximum blade count in the engine is 144 blades per wheel, which results in approximately 7000 time steps per rotor revolution. Given that 1.55 revolutions are needed for one flow through time, it is estimated that the simulation will require 11,000 time steps per flow through time. Because of the smaller number of blades per wheel, the time step for the turbine simulation can be approximately twice as large as that for the compressor, and only 5,500 time steps are required for that part of the engine. However, for the following estimates, we will assume the same time step in compressor and turbine. The expected size of the grids for the compressor and turbine are 233 million (776 blade passages) and 74 million (246 blade passages) computational cells, respectively. This yields a total cost of the SUmb simulation of about 2 million total CPU hours on QSC for the compressor and 651,000 total CPU hours on QSC for the turbine. This leads to a total number of 2.65 million CPU hours for a computation on a total of 307 million cells. Experience has shown that SUmb scales almost linearly for at least 40,000 cells/processor (and at least 5,000 cells/processor on BlueGene). Based on this experience, the simulation could be expected to use about 7,500 processors efficiently

(for the worst-case scenario in scalability). The total wall clock time per computed flow through time is therefore approximately two weeks. The memory required for SUmb is about 1 GB per million grid cells. The total memory required is 307 GB or 41 MB per processor.

A final word to the inflow boundary conditions: at the inflow, radial distributions of total pressure, total temperature and the velocity direction are prescribed. At the outlet, a radial distribution of static pressure is given. One-dimensional Riemann invariants are used to determine the values in the ghost cells. At the turbine inlet, due to the large variation in the prescribed total temperature in the near wall region, a total enthalpy scaling had to be applied to the acoustic Riemann invariant $u \cdot n - 2a/(\gamma - 1)$. Without this scaling, the one-dimensional boundary condition treatment was not stable.

3. Large-scale turbomachinery computations

The compressor of a gas turbine engine is used to compress the ambient air to a very high pressure. The compressed air is passed to the combustor, where it is mixed with fuel and burned. High-performance jet engines usually employ multiple compressors consecutively; the Pratt & Whitney engine considered in this work has two: the low pressure compressor (LPC) and the high-pressure compressor (HPC). The LPC compresses the air from atmospheric pressure to a higher level. The rotational speed of this compressor is relatively low at about 5,000 rpm. The HPC compresses the air further and rotates usually at about 15,000 rpm. The focus here will be on the high-pressure compressor (HPC), which is computationally more demanding. It consists of eleven blade rows with six stators and five rotors. Due to the adverse pressure gradient, the compressor flow is prone to flow separation, and a proper modeling of the turbulence in the near wall region is mandatory.

The turbine is located at the outlet of the combustor and extracts energy to propel the compressor and the fan. Like the compressor, the turbine consists of a high-pressure section (high-pressure turbine, HPT), which powers the high-pressure compressor, and the counter rotating low-pressure spool (low-pressure turbine, LPT), which powers the low-pressure compressor. Since these spools are connected by a shaft, the rotational speeds in the two sections of the turbine correspond to their counterparts in the compressor. In the computations presented here, the HPT consists of one stator and one rotor, and the first stage of the LPT consists of an additional stator and a rotor. Due to the high temperatures at the inlet of the turbine, cooling flow is normally injected through holes in the blades to avoid burning of material. Although SUmb has the capability to model these cooling flows, all the results shown in sections 4 and 5 are obtained without the cooling model.

The operating conditions for these large-scale simulations have been determined in discussions with our industrial collaborators at Pratt & Whitney. Two conditions will be considered. These correspond to the regular cruise condition and the so called min cruise condition. The cruise condition will be used as a baseline simulation. At this condition, the performance of the engine, measured by the pressure ratio and mass flow rate at a given work input, is certainly crucial. This condition is also very important for questions related to fatigue, durability, and engine lifetime. The min cruise operating point corresponds to cruise at the minimum possible altitude. At min cruise, compressor stability becomes an important factor, since the compressor operating line approaches the stability line (surge line). The compressor becomes unstable when the operating line crosses the surge line. This is the so-called pinch point in the stability map.

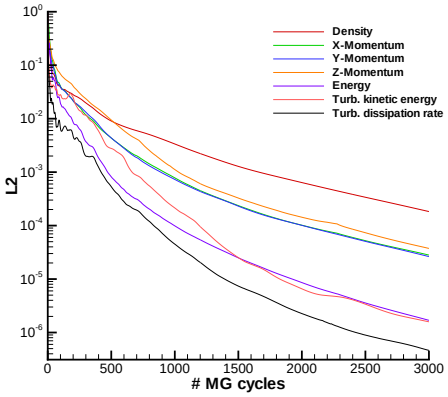

FIGURE 2. Convergence histories for the PW6000 turbine, mixing plane assumption. 3W cycle in combination with a 5 stage Runge Kutta smoother.

We start by running simple steady-state simulations using the mixing plane concept for both the turbine and the compressor. The solutions of those mixing plane simulations are used to initialize the 20^o sector and the full wheel unsteady simulations for both turbomachinery components.

3.1. *Geometry and grids*

Original full-wheel and scaled sector grids have been provided by Pratt & Whitney for the compressor and turbine. The tip gap regions have been neglected and shrouded casings are assumed for the rotating parts. Both the compressor and the turbine grids have the same topology consisting of an O-grid around the blade and an H-type grid in the passage. The compressor and turbine grids for one passage consist of 3,080,092 and 1,425,408 cells, respectively. The full wheel grid for the compressor consists of 1552 blocks (219,299,840 cells), while for the turbine it consists of 492 blocks (87,865,344 cells). The disk space needed to store these grids in double precision is 2.1 GB for the turbine and 5.7 GB for the compressor.

In these grids, the full wheel geometries of the compressor and of the turbine contain blade rows of different blade counts that do not have a common denominator. Circumferentially scaled grids that adjust the number of blades in each blade row have been provided to allow the creation of a 20^o periodic sector of the geometry. The scaling does not change the number of cells in one passage. The 20^o sections of compressor and turbine consist of 90 blocks (12,713,984 cells) and 28 blocks (4,982,784 cells), respectively.

For the flow conditions considered, the compressor grid is quite coarse. It is has an average y^+ value of about 3. It is, however, used as a wall integration grid. The average y^+ value of the turbine grid is about 60, thereby requiring the use of wall functions. Both grids are multigrid friendly to at least three levels with a proper one-to-one cell matching on the coarser levels between the O and H grids.

4. Steady computations

The flow is first computed for a single multistage passage of turbine and compressor for steady state conditions. The flow in the passage is assumed to be periodic in the circumferential direction, and a mixing plane approximation is used for the boundary

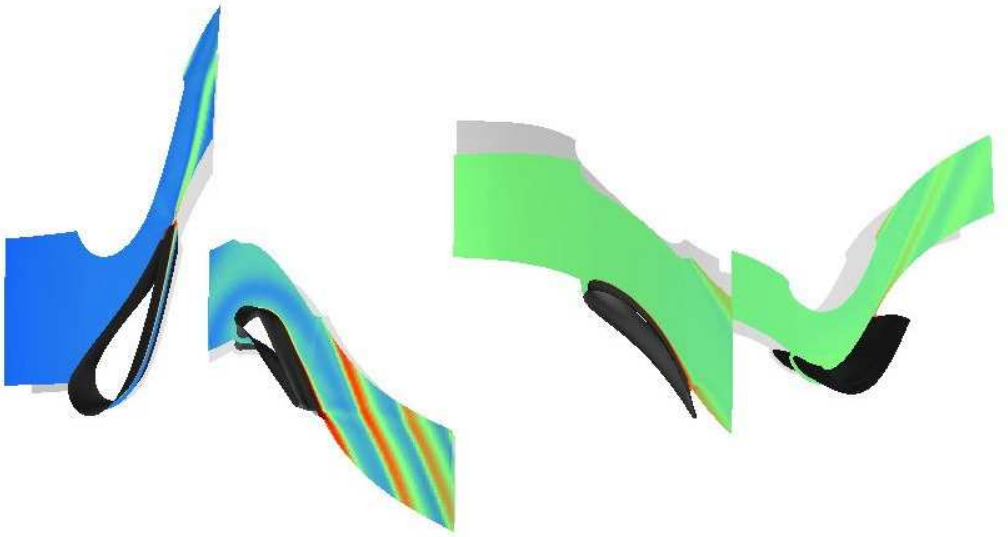

FIGURE 3. Entropy for mixing plane computation of one flow passage of the turbine.

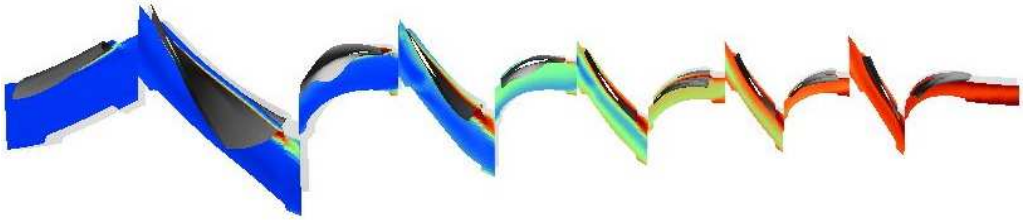

FIGURE 4. Entropy for mixing plane computation of one flow passage of the compressor.

condition between the blade rows. In the mixing plane approximation, the computed flow at the end of each blade row is circumferentially averaged and used as inflow condition for the downstream blade row. Similarly, the circumferentially averaged solution at the beginning of the flow passage is used as an outflow boundary condition for the upstream blade row. The averaging for the coupling of individual blade rows is necessary, since stator and rotor move relative to each other.

4.1. *Steady computations for a single turbine passage*

The convergence history for this test case is presented in Fig. 2. Although 3,000 multigrid cycles were used, the solution reached engineering accuracy after approximately 1,000 cycles. Note that the initial solution on the fine grid was obtained by applying grid sequencing on the coarser meshes. Consequently, the relative initial position for the convergence histories shown in Fig. 2 is already two orders of magnitude lower than when the solution is initialized from free stream conditions.

The turbine computation can be started relatively easily from an arbitrary initial guess. The pressure difference between the inflow and the exit drives the flow through the turbine. The entropy field of a steady state mixing plane solution in a turbine passage is shown in Fig. 3. The abrupt disappearance of the wakes between the blade rows is a direct consequence of the one-dimensional interpolation.

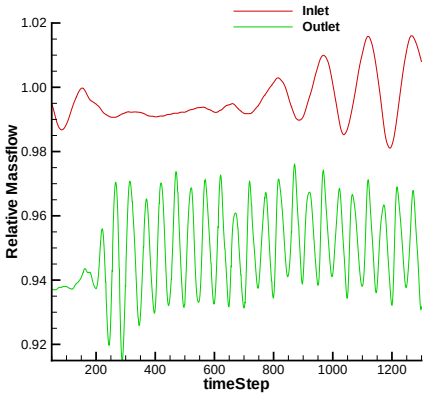

FIGURE 5. Evolution of the relative inlet and outlet mass flow for the scaled turbine geometry.

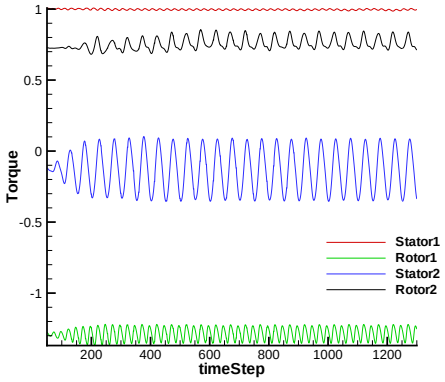

FIGURE 6. Evolution of the relative torque on the four blade rows for the scaled turbine geometry.

4.2. *Steady computations for a single compressor passage*

As for the turbine, the standard turbomachinery boundary conditions are applied. Radial profiles of stagnation pressure and temperature, as well as the velocity unit vector prescribing the flow direction, are specified at the inlet. For the compressor, the inlet velocity has - in addition to the component in the axial direction - an additional circumferential velocity component specified to mimic the flow direction from the upstream rotor. The static pressure is specified at the exit.

The compressor computation is far more challenging to start due to the high compression ratio that results from the work transmitted from the rotating blades to the fluid. Since the exit pressure is higher than the pressure at the inlet, the flow tends to reverse its direction as soon as large pockets of separation occur in the passage. As a consequence, the computations must be carried out in several steps, first by slowly raising the rotational speed of the wheels while keeping the pressure low at the exit, and then by increasing the pressure once the full rotational speed was achieved. The final pressure at the exit matches the pressure that is specified at the inlet of the combustor. The entropy field of a steady state mixing plane solution in a compressor passage is shown in Fig. 4.

5. Unsteady computations

A more accurate description of the flow field can be achieved by performing unsteady simulations. In contrast to the mixing plane approximation, the values in the halo cells on the sliding mesh interfaces are determined by interpolation from the neighboring blade row without any averaging. Interpolated values from the downstream blade row are imposed in the halo cells of the upstream blade row and vice versa; the flow field is continuous over the sliding mesh interface.

Unsteady simulations can be performed on a periodic sector of the geometry. This is often impossible, since each blade row can have a different number of blades that do not have a common denominator. This is done to avoid instabilities caused by resonance between two components. For a sector simulation, the number of blade counts is rescaled such that sector periodicity is obtained. The pitch and chord of the blades are adjusted

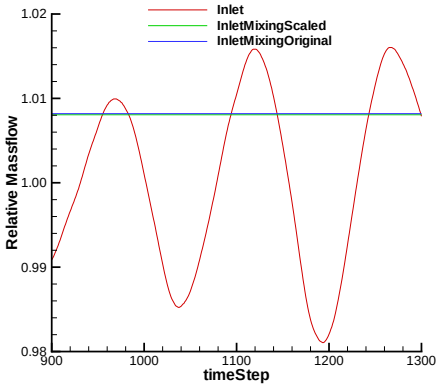

FIGURE 7. Comparison of the relative inlet mass flow of the unsteady solution and the two mixing plane solutions.

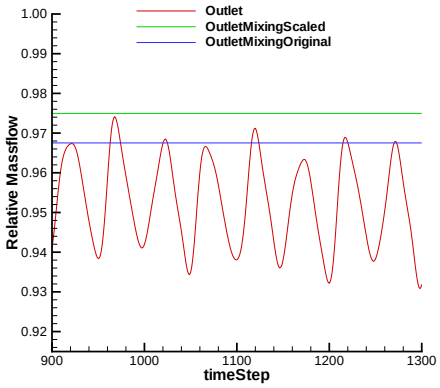

FIGURE 8. Comparison of the relative outlet mass flow of the unsteady solution and the two mixing plane solutions.

to preserve the same flow blockage. Unsteady simulation for the exact geometry can only be done for the entire wheel.

5.1. *Unsteady computations of* 20^o *turbine sector*

The 20^o turbine sector computations have been advanced for 1,300 timesteps. 50 time steps are used for one blade-pass of the first rotor. This corresponds to 2,700 time steps for a full wheel revolution. Transient effects are still present in the flow, as illustrated by time evolutions of the relative inlet and outlet mass flow. The torque on the blades seems less sensitive and has reached a seemingly periodic state.

Figures 7 and 8 show the mass flow evolution during the last 400 time steps compared to mixing plane solutions computed using both scaled and unscaled geometry. First, it can be seen that all solutions have a lower mass flow rate at the outlet than at the inlet, despite the use of a conservative scheme. In a conservative scheme only the sum of central and dissipative fluxes is conserved. This means that on coarse grids, the mass flow computed from central fluxes is not conserved, as shown in Figures 5, 7 and 8. For the mixing plane solution, this leads to a loss of 3 percent in the mass flow rate due to dissipation. The unsteady solution predicts an average loss of approximately 4 percent. This is a clear indication that the grid is too coarse, and the solution is not grid converged. This result was anticipated, but the decision was made to proceed with such coarse grids in order to learn some initial lessons about the unsteady nature of these large scale computations at somewhat lower computational cost.

Secondly, the average mass flow for the unsteady computation is lower than for the mixing plane solutions. Although this could still be a transient effect, Fig. 5 indicates that this is not the case. Finally, the unsteady mass flow rate at the outlet, Fig. 8, shows a much higher frequency than the mass flow rate at the inlet, Fig. 7. The flow at the outlet is influenced by all blade rows, while the flow at the inlet seems to be influenced only by the first blade row.

Figures 9 to 12 show the relative torque on the blade rows for the unsteady solution of the last 400 time steps, for the mixing plane solutions, and for the time averaged values. The difference between the value obtained with the mixing plane assumption and the time averaged value is between 1 and 3 %. However, due to the steady nature of

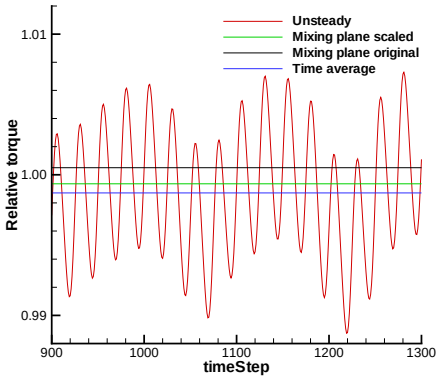

FIGURE 9. Comparison of the relative torque on the first stator of the unsteady solution and the two mixing plane solutions.

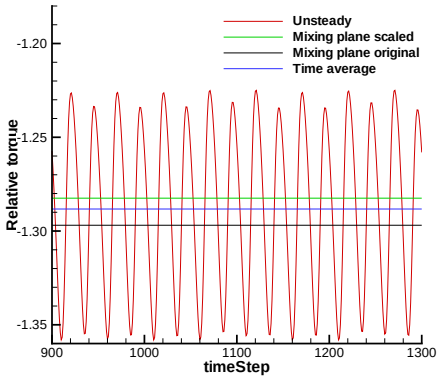

FIGURE 10. Comparison of the relative torque on the first rotor of the unsteady solution and the two mixing plane solutions.

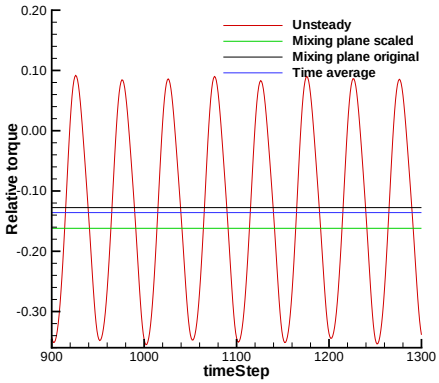

FIGURE 11. Comparison of the relative torque on the second stator of the unsteady solution and the two mixing plane solutions.

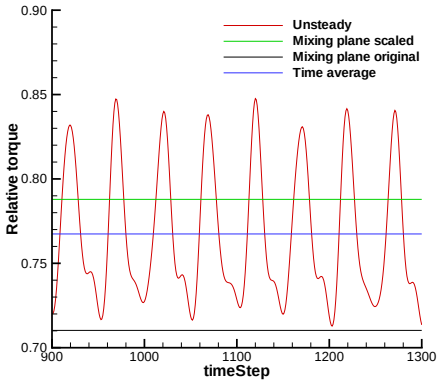

FIGURE 12. Comparison of the relative torque on the second rotor of the unsteady solution and the two mixing plane solutions.

the mixing plane assumption, it does not give any information about the amplitude and frequencies of the oscillations, which are important design properties.

Figure 13 shows the instantaneous entropy distribution in the same plane as Fig. 2. It is clear that the solution is now continuous over the sliding mesh interface, and the effect of the wakes on the downstream blades is evident.

5.2. *Unsteady computations of* 20^o *compressor sector*

For a full wheel revolution of the compressor, 6,300 time steps are needed. This value again corresponds to 50 time steps for a blade passing of the blade row with the highest blade count. Figure 14 presents the instantaneous entropy distribution after 1500 time steps, which corresponds to approximately 85^o of rotation.

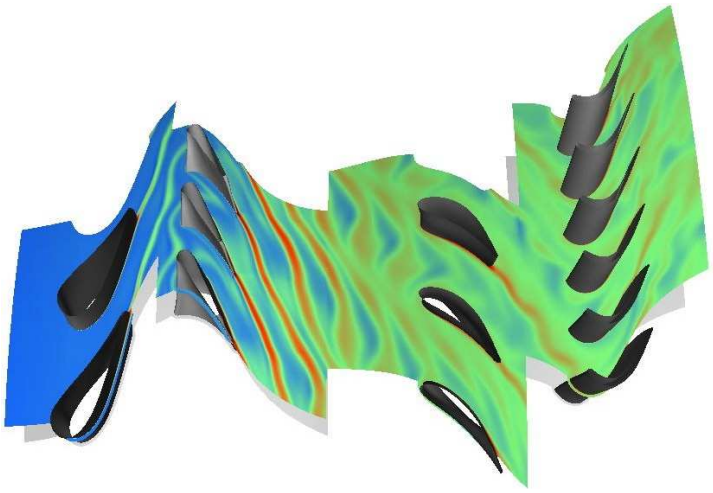

FIGURE 13. Turbine, 20^o sector; Instantaneous entropy distribution for the unsteady solution for the scaled geometry.

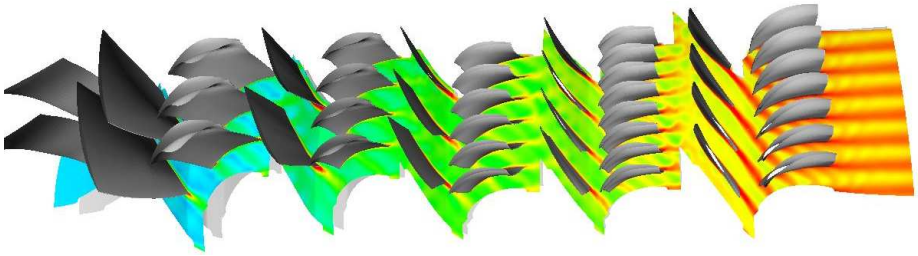

FIGURE 14. Compressor, 20^o sector; Instantaneous entropy distribution for the unsteady solution for the scaled geometry.

5.3. *Unsteady computations of full-wheel turbine*

Due to the smaller number of blade rows in the turbine, the computational costs for the unsteady simulation of the full wheel turbine are smaller than for the compressor. Again, a mixing plane solution is used as initial solution for the unsteady simulation. This case has been run on the LLNL ALC machine on either 300, 600 or 1,200 processors, depending on the available resources; SUmb writes the solution in a single file and a restart can be made on a different number of processors due to the fully integrated parallel preprocessor.

The time step is chosen such that a full revolution of the turbine takes 2,700 time steps. On 600 processors, the computation of one time step takes approximately 10 minutes. This includes 2 minutes for writing the grid and the solution files. A solution computed with 593 time steps, i.e. approximately 1/6 of a revolution, is compared to the steady solution using the mixing plane approximation and to the unsteady solution for the 20^o sector in Fig. 15. Transient effects are still present, and more time steps are needed. It is estimated that approximately half a revolution is needed to remove these transients, and another full revolution is needed to obtain a periodic solution. This corresponds to approximately 4,000 time steps and thus to 680 hours wall clock time on 600 processors, i.e. 500,000 CPU hours.

For these full wheel computations, the disk storage space also becomes an issue. If a solution file is written in double precision, this leads to a file size of 6.7 Gb. Together

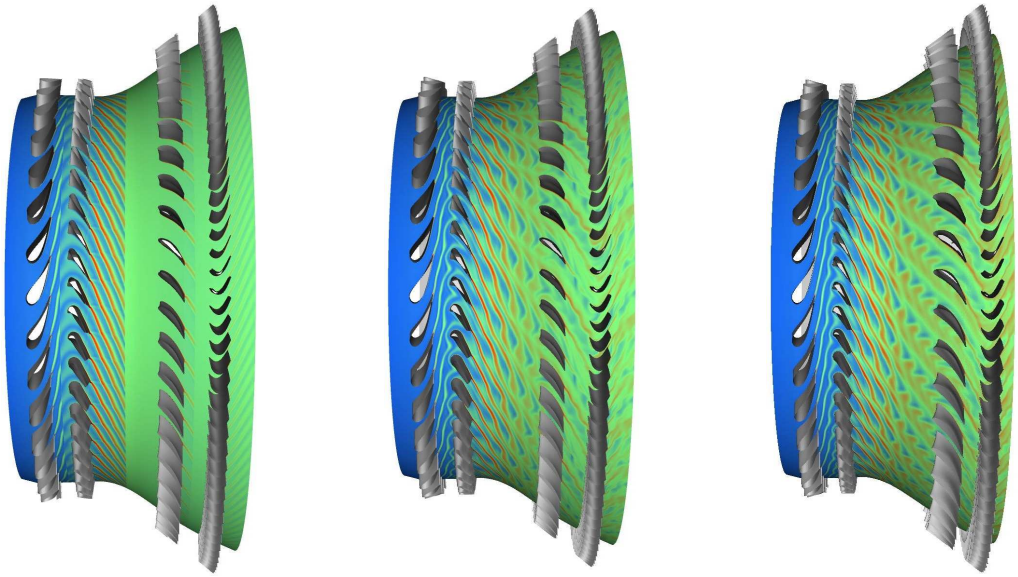

FIGURE 15. Instantaneous entropy distribution for turbine simulations after 593 time steps, i.e. 59.3^o rotation of the HPT rotor. Steady solution with the mixing plane approximation (left), unsteady solution for 20^o sector (middle), and unsteady solution for the full wheel (right).

with the 2.1 Gb for the grid file, this leads to a total of approximately 9 Gb of data per time step. When the solution is to be saved at every time step, this will lead to a disk storage requirement of approximately 35 Tb for the entire simulation; for single precision this is equal to 17.5 Tb. The question is whether to store the whole solution every time step or only to dump *a priori* chosen planes. The advantage of the latter approach is that the disk space requirements are reduced by at least one order of magnitude. The disadvantage is that some *a priori* knowledge of the solution is needed.

The numbers above are based on a grid of 88 million cells, i.e. approximately 300,000 per blade passage. It has been shown for the 20^o sector that this grid resolution is not enough to obtain a grid converged solution. It is estimated that such a grid converged solution is obtained with 1,500,000 cells per blade passage if wall-integrations is used, i.e. using a grid five times finer. The corresponding CPU and disk space requirements for such a grid then also increase by a factor of 5, under the assumption that 2,700 time steps for a full revolution are still enough to obtain a time converged solution.

5.4. *Unsteady computations of full-wheel compressor*

The full-wheel approach does not require rescaling of the turbomachinery blades. The current state-of-the-art turbomachinery computations rarely include more than one stage when computing the entire wheel. Here, however, we compute the full-wheel of the entire high pressure compressor for an extended period of time, which, to our knowledge, has not been previously attempted.

The time step chosen for this simulation is such that a full revolution of the HPC takes 7,000 time steps. This corresponds to 50 time steps per blade passing of the blade row with the most blades. We were able to run this computation on 600 processors for 20 hours obtaining 40 time steps (see Fig. 16), which corresponds to approximately 2^o of rotation.

The size of the solution file for current simulation is approximately 20 GB when stored

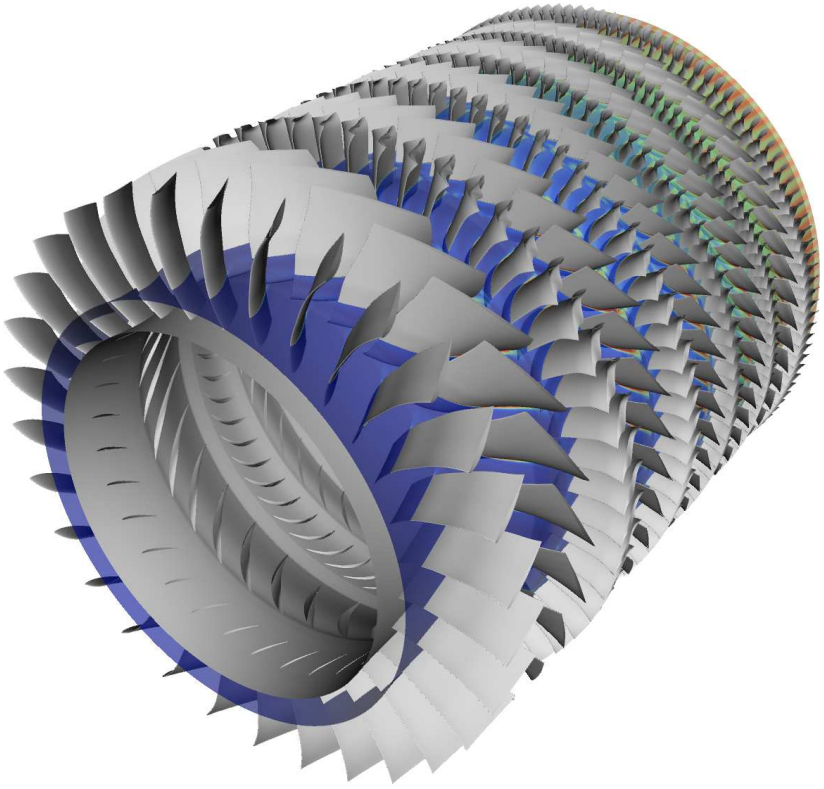

FIGURE 16. Instantaneous entropy distribution for the full wheel compressor simulation after 40 time steps, i.e. 2^o rotation of the HPC rotor.

as double precision. Combined with 5.7 GB for the grid, this leads to approximately 26 GB per time step. If the solution is stored at every time step, the total amount of disk space required for one full revolution is 180 TB. With single precision, this number is reduced to 90 TB.

The mesh for the full wheel high pressure compressor consists of 220 million control volumes, which corresponds to approximately 300,000 cells per blade passage. Note that this grid is a wall-integration grid. From the turbine simulations, we have already concluded that 300,000 cells per passage seems to provide insufficient resolution, even when wall-functions are used. A better resolution with wall-integration would require about 1.5 million cells per passage. This would lead to full-wheel grids consisting of approximately 1 billion cells.

The large number of control volumes of this case allows for the scalable execution on a high number of processors. Previous experience on the ALC platform suggests that the linear speed-up can be expected to be obtained by up to 10,000 processors, provided that good load balance can be assured.

6. Conclusions

This paper describes the large scale unsteady turbomachinery computations currently performed at the Center for Integrated Turbulence Simulations (CITS) at Stanford. The steady simulations of a single multi-stage passage for both turbine and compressor using the mixing plane approximation represent a challenging test case for compressible RANS

flow solvers. The impact of boundary conditions, numerical scheme, turbulence model and grid resolution can be investigated in detail. The unsteady simulations for a 20^o sector go beyond current industry standards and provide an insight into the unsteady effects that may significantly influence the important design properties. The computational cost, although high, is still acceptable and allows for a comprehensive investigation of unsteady features. However, the geometry usually needs to be modified by rescaling the blade counts. This change alters the flow physics and might have a significant impact on specific time dependent features. The full wheel computations circumvent this problem, albeit at an enormous cost. The grids used here consist of 88 million cells for the turbine and 220 million cells for the compressor. Unfortunately, these grids are still too coarse, and a higher resolution is needed; the estimate is that grids with at least 5 times as many cells are needed. However, the described pioneering large scale computations have identified the computational challenges and have provided an initial understanding of the complex three dimensional turbomachinery flow features.

7. Acknowledgment

The authors would like to acknowledge support from Saadat Syed, Dillip Prasad and William Sprout of Pratt & Whitney, who provided grids and boundary conditions.

REFERENCES

Adamczyk, J.J. 1999 Aerodynamic analysis of multistage turbomachinery flows in support of aerodynamic design. *ASME-99-GT-80*.

Chen, J.P. & Barter, J.W. 1998 Comparison of time-accurate calculations for the unsteady interaction in turbomachinery stage. *AIAA-98-3292*.

Denton, J.D. & Singh, U.K. 1979 Time marching methods for turbomachinery flows. *VKI LS 1979-07*.

Durbin, P.A. 1995 Separated flow computations with the k-ε-$\overline{v^2}$ model. *AIAA J.* **33**, 659-664.

Jameson, A., Schmidt, W. & Turkel, E. 1981 Numerical solution of the Euler equations by finite volume methods using Runge Kutta time stepping schemes. *AIAA-81-1259*.

Jameson, A. 1998 Time dependent calculations using multigrid, with applications to unsteady flows past airfoils and wings. *AIAA-91-1596*.

Medic, G., Kalitzin, G., Iaccarino, G. & van der Weide, E. 2006 Adaptive wall functions with applications. *36th AIAA Fluid Dynamics Conference and Exhibit*.

Kalitzin, G., Medic, G., Iaccarino, G. & Durbin, P. A. 2005 Near-wall behavior of RANS turbulence models and implications for wall function. *J. Comp. Phys.* **204**, 265-291.

Roe, P. 1982 Fluctuations and signals - a framework for numerical evolution problems. In *Numerical Methods for Fluid Dynamics*, Academic Press.

Spalart, P. R. & Allmaras, S. R. 1994 A one-equation turbulence model for aerodynamic flows. *La Recherche Aerospatiale* **1**, 1-23.

Wang, X. & Chen J. 2004 A post-processor to render turbomachinery flows using phase-lag simulations. *AIAA-04-615*.

Wilcox., D.C. 1988 Reassesment of the scale-determining equation for advanced turbulence models. *AIAA J.* **26**, 1299–1310.