C²A²S²E Newsletter

Issue 03 - October 2010

Center for Computer Applications in AeroSpace Science and Engineering

DLR Deutsches Zentrum für Luft- und Raumfahrt e.V. in der Helmholtz-Gemeinschaft

Niedersachsen

AIRBUS

C²A²S²E Center for Computer Applications in AeroSpace Science and Engineering

Europa fördert Niedersachsen

EUROPÄISCHE UNION

C²A²S²E - running into application

C²A²S²E project is widely focused on requirements from the aircraft development process. Flight Physics engineers appreciate to get hand on more powerful simulation tools with a wider scope, validated on their specific test cases. A number of examples demonstrate the extension of the current capability with respect to accuracy, efficiency and robustness. For sure there remain open questions and additional work to be done to provide a more comprehensive simulation of the final product.

Validated simulation of open rotor engines

Validation is the essential bit prior to any production use of new physical and numerical models. The two counter-rotating discs of the new open rotor engines marked a tough task for physical modeling as well as for the development of a best practice to resolve major physical effects in proper way. Investigations on an isolated engine test have been conducted in order to find out the necessary granularity of the actuator disc model. The solution implemented via C²A²S²E is able to predict physical phenomena with good accuracy.

New type of propulsion system

Rear mounted counter rotating propeller engines now put a next challenge on numerical simulation. Specifically the aerodynamic shape design in presence of a double disc of rotor blades requires a simplified physical model and a robust numerical approach. The actuator disc model previously developed through C²A²S²E provided a solution that could be used to simulate propulsion as well as reverse flow conditions. The numerical model now allows to optimize aerodynamic shapes of the rear end fuselage and tails for the design of a next generation aircraft.

Assessment of high lift flow prediction

Take-off and landing configurations pose a tough challenge on numerical flow simulation. Flow physics as well as the aircraft geometry are very complex, thus requiring a sophisticated shape and flow modeling. At the border of the flight envelope flow separation is becoming the dominant effect the modeling of which is unfortunately very difficult. Since some time C²A²S²E is working to provide better prediction of aerodynamic coefficients for mildly separated up to fully separated flow cases. First improvements over conventional simulations could be obtained from high level turbulence models.

Contact:

Dr. Klaus Becker
(klaus.becker@airbus.com)

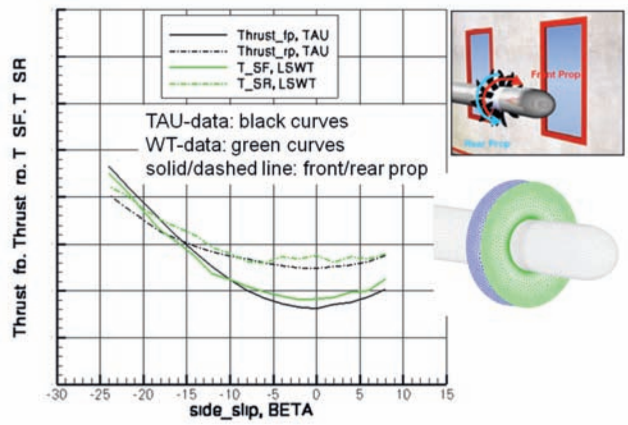

Double actuator disc model validation results for isolated engine

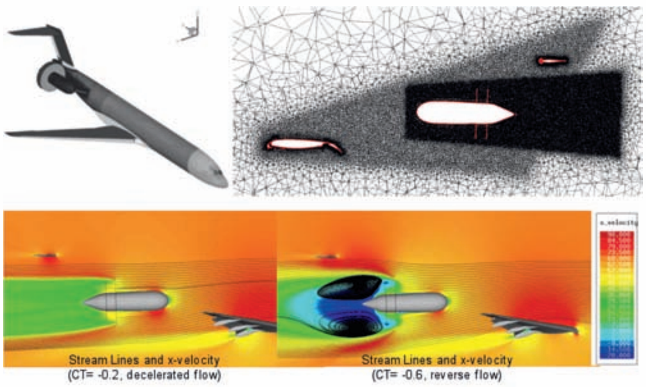

Installed engine simulation now applicable to even compute reverse flows

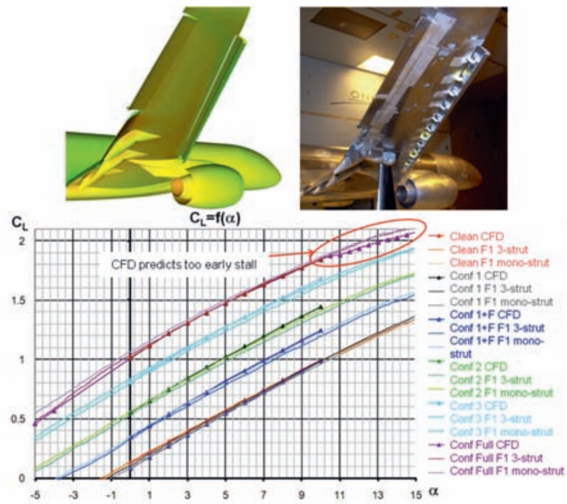

TAU RANS investigation on lift versus angle of attack for different high lift configurations compared to wind tunnel results

Enhanced simulation accuracy by high-level physical modelling

Pursuing the C²A²S²E vision of completely simulating a new aircraft before its first flight, the application of CFD is driven continuously towards the borders of the flight envelope. Consequently, on this way increasingly complex flow phenomena are encountered that have to be simulated accurately. In particular the prediction of separated flows is a demanding challenge for the underlying physical models, specifically of turbulence.

Standard turbulence models, like the Spalart-Allmaras or Menter SST model, rely on the assumption of an increased viscosity due to the turbulent motion. While being suitable for simple attached shear flows, this assumption is overly simple for predicting complex separating flows accurately. Therefore, simulating aircraft operating at the border of the flight envelope requires a higher level of physical modelling.

For this reason C²A²S²E has been following the approach of Differential Reynolds Stress Modelling (DRSM) of turbulence, exploiting the transport equations of the individual turbulent stress components instead of a simplified turbulent viscosity assumption. A model has been developed, transferring the ideas of Menter's famous SST model into the world of DRSM. This model has been implemented into the DLR TAU code in a numerically efficient and fairly robust manner, making it suitable to industrial applications.

The figure demonstrates the superiority of the DRSM technology for the transonic flow around the ONERA M6 wing. Up to the highest measured incidence of $\alpha=6.06°$ the Reynolds stress model's predictions are in good agreement with the experiment. In contrast the Spalart-Allmaras model and the Menter SST model both fail to represent the experiment for incidences of $\alpha \geq 4.08°$because they significantly over predict the shock induced separation.

Contact:

Dr. Bernhard Eisfeld
(bernhard.eisfeld@dlr.de)

Pressure distributions and skin friction lines on the ONERA M6 wing at $\alpha=4.08°$. Superior predictions by Reynolds stress model compared to standard eddy viscosity models due to improved prediction of shock induced separation.

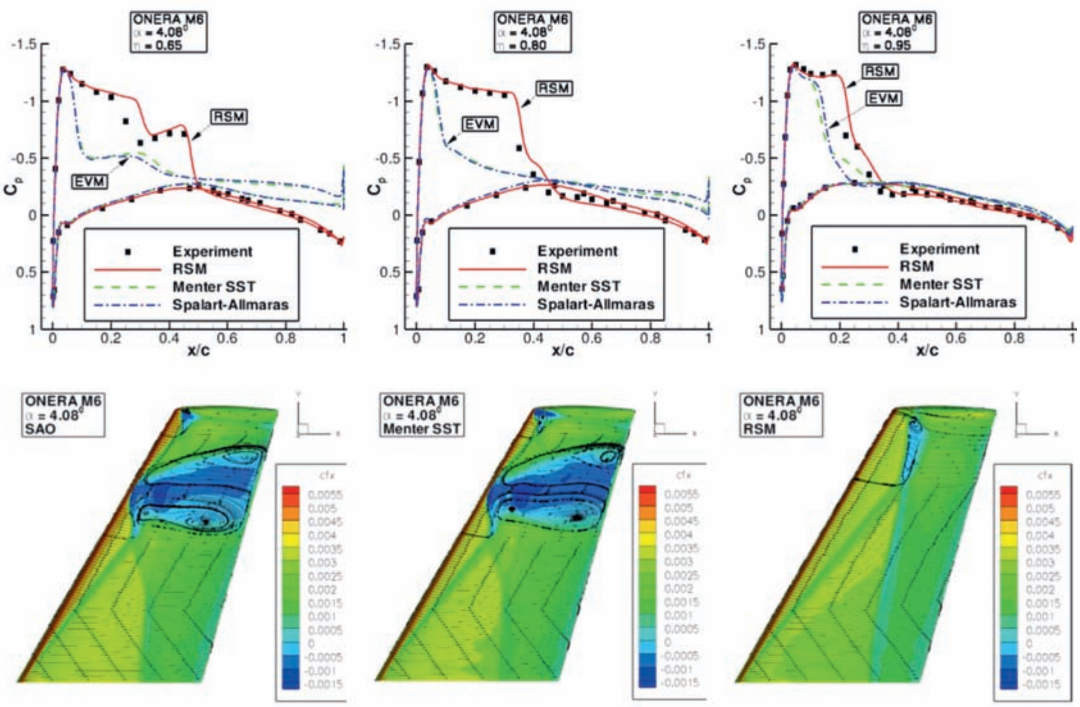

Deutsches Zentrum
DLR für Luft- und Raumfahrt e.V.
in der Helmholtz-Gemeinschaft

Fusing experimental and computational data for harmonized aerodynamic data

One of the objectives of C²A²S²E is to develop and demonstrate capabilities to provide harmonized aerodynamic data based on a set of inputs from different sources (low fidelity CFD, high fidelity CFD, wind tunnel, flight test, legacy Aircraft Data Bases). This capability is linked to the requirement to ensure consistent data for loads preparation, handling qualities analyses and performance prediction.

Typically, the integration of the (local) aerodynamic pressure distributions or section loads used in the aero-data for loads process results in (global) force and moment coefficients that do not match those considered for evaluating the handling qualities of an aircraft, i.e. the data is not harmonized. This is due to the fact that the underlying wind-tunnel data is determined using different measurement techniques – the total aerodynamic coefficients are determined from force and moment measurements using a wind-tunnel balance, while the surface pressures are measured using pressure taps at specific positions on the wind-tunnel model (typically wing, fuselage and tails sections). When integrating the experimental surface pressure values over the surface of the aircraft, assumptions have to be made, since the pressure measurements are available at only a few pressure ports and at a very limited number of sections (butt lines), thus resulting in a different integral value than that measured with the balance.

The Gappy POD method, which is based on Proper Orthogonal Decomposition (POD), provides a way to reconstruct full flow field information, using a combination of sparse experimental data and supplemental computational data. It was first developed in the context of reconstructing human face images from incomplete data sets – hence the name "Gappy" POD – and later extended to fluid dynamic applications, such as inverse airfoil design. The basic idea behind Gappy POD for data fusion is that the POD analysis of steady CFD solutions or snapshots at different flow conditions yields a set of empirical modes, which describes the dominant behavior or dynamics of a given flow problem in an optimal sense, i.e. for any given basis size, the error between the original and reconstructed data is minimized. Given a suitable set of POD modes derived from computational data, the Gappy POD solves a small and thus inexpensive least squares problem to determine a set of coefficients such that the reconstruction (linear combination of modes) optimally matches the experimental data. In this way, it is possible to predict a complete surface pressure distribution that will best match the few pressure measurements in a least-squares regression sense and, when integrated, will allow substituting the balance measurements if sufficient accuracy is obtained.

Besides its capability to harmonize data, the Gappy POD approach offers additional benefits, such as the prediction of quantities that were not measured (e.g. the skin friction distribution). This can be achieved by using global POD modes that account for all computed flow variables of interest simultaneously (e.g. surface pressure and skin friction) and a set of gappy experimental data containing only the measured quantities (e.g. pressure). Remarkably, the

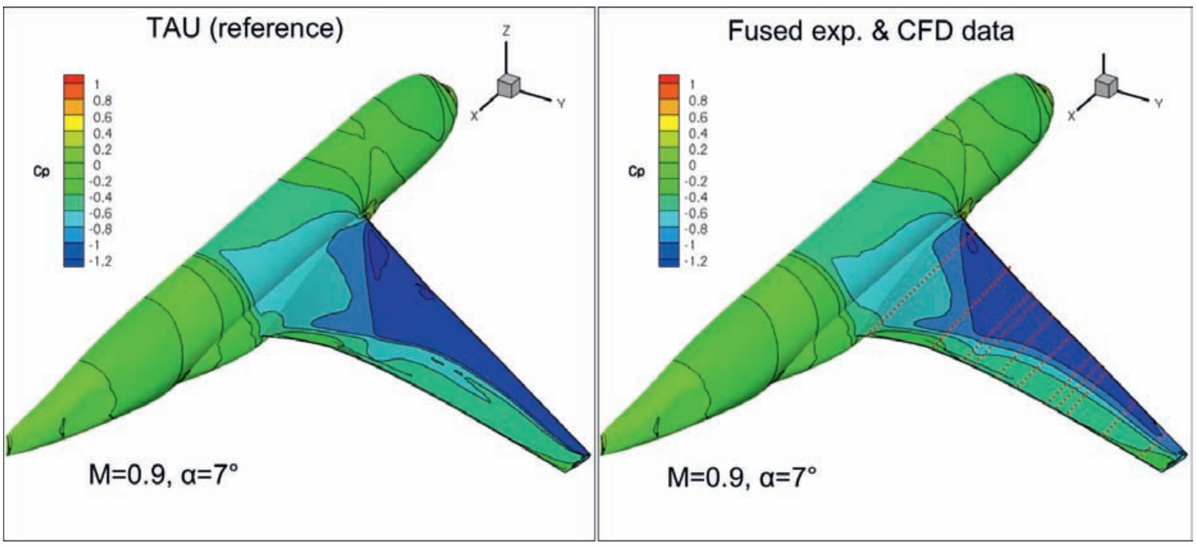

Figure 1

surface pressure distributions and integrated aerodynamic coefficients predicted by the Gappy POD are typically in better agreement with experimental data than the results of a CFD computation at the same nominal (wind-tunnel) flow condition, because the Gappy POD takes real measured data into account. It can thus also be seen as a way to correct complete CFD solutions (even off surface!) if limited experimental data is available.

A further application is the prediction of aerodynamic data on aircraft components not included in the experiment (e.g. fairing loads). Here the method makes use of the fact that data can be fused even if the geometry is not identical in all areas. If used with caution, this allows merging complete surface solution from CFD for an aircraft model with all relevant aircraft components with gappy experimental data. The result is a corrected CFD solution for all relevant aircraft components. This is especially useful when aiming to improve the accuracy of loads prediction on components that are difficult to integrate into a scale model in a wind-tunnel campaign.

At C²A²S²E, the Gappy POD method was developed in cooperation with researchers from the Institute of Computational Mathematics at TU Braunschweig, keeping a real engineering environment in mind. For the first time, the method was demonstrated for fusing real experimental and computational data of an industrial aircraft configuration. For this purpose, the wing-body configuration Megaflug was employed. This test case consists of wind tunnel data measured for transonic flow conditions at ARA (U.K.) and a numerical grid of about six million volume grid points and 163,204 surface nodes. Steady flow solutions at a Mach number of Ma = 0.9 and a Reynolds number of 3.3 million were computed for angles of attack of 4°, 6° and 8° with the DLR RANS flow solver TAU using the Spalart-Allmaras turbulence model. Gappy POD was then used to obtain the complete surface pressure distribution and skin friction distribution at an intermediate angle of attack of 7° and the same Mach number by considering sparse (“gappy”) experimental data at 11 sections (376 pres. ports) for this flow condition.

Figure 1 compares the surface pressure distribution obtained by fusing the experimental and CFD data using the Gappy POD (right) with that obtained by a reference CFD solution at the same nominal flow conditions (left). Although the overall agreement is good, this is not necessarily expected, as CFD has its own deficiencies. In other words, the Gappy POD prediction may be in better agreement with the measurements than the reference CFD. This can only be evaluated by integrating both pressure distributions over the surface of the aircraft. The results are compared with the measured aerodynamic coefficients in Table 1.

As expected the lift and drag coefficients predicted by Gappy POD are in better agreement with the experimental values than the reference CFD solution, demonstrating that the Gappy POD is a powerful method for data fusion and offers efficient means to harmonize data. It can thus be utilized for improving the consistency between the various sets of data used in aero-data for loads, handling qualities and performance.

Contact:

Dr. Stefan Görtz
(stefan.goertz@dlr.de)
Dr. Michael Mifsud
(michael.mifsud@dlr.de)
Petra Aumann
(petra.aumann@airbus.com)
Dr. Lars U. Hansen
(lars-uwe.hansen@airbus.com)

α=7°	C_L	C_D
Exp. (WTT)	0.815	0.114
GPOD	0.811 (-0.49%)	0.117 (+2.6%)
TAU (ref.)	0.843 (+3.4%)	0.121 (+6.1%)

Table 1

Next step in Petascale Computing for CFD in aerospace science

When the construction work in the computer center at the DLR started in Spring 2007 and the design for a C²A²S²E HPC cluster for the coming five years begun, it was already obvious that a hardware upgrade would be needed after about three years to ensure the cluster system remained virtually state-of-the-art. Therefore a replacement of the cluster blade servers using the most cutting-edge technology available at the time with at least three times increased performance was agreed with our partner T-Systems Solutions for Research.

The three years have now passed, so in 2010 the hardware upgrade has to be installed.

In the IT world three years is a long time and predictions of future commercial and technological developments are often difficult for example Sun Microsystems, the hardware supplier for our HPC cluster, had surely a turbulent time with its acquisition by Oracle over the last few months. But also the users of the C²A²S²E cluster discovered at latest during January 2010 that a hardware upgrade was inevitable. The monthly average utilization of the C²A²S²E cluster was constantly over 85 % this year.

The core design parameters for the hardware upgrade were already fixed in summer 2007:

- A desired performance increase by a factor of three on the basis of the TAU-Code relative to 650 nodes from the existing hardware must be reached.
- The electrical power consumption may not rise over the power budget from the existing hardware.

In the beginning of the year T-Systems carried out benchmark tests on the two eligible processors AMD Magny-Cour and Intel Xeon 5600 series (Westmere-EP architecture) running the Tau-Code. Looking at the processing performance, the Magny Cours at 2.1 GHz was just above the Intel Xeon X5680 at 3.3 GHz. The slightly lower clocked processors Magny-Cour at 1.7 GHz and Intel Xeon X5670 at 2.9 GHz were approximately equal. The important result from these benchmarks was, that the desired performance increase by a factor of three is reachable. However, the benchmarks also showed that the higher clocked processors would clearly have exceeded the specified electrical power consumption.

The Intel Xeon X5670, 2.9 GHz has finally been chosen due to the fact that with the Intel Xeon 5600 series significantly less processing cores were needed to reach the desired performance (8200 cores compared to 16800 cores

from the AMD Magny-Cour at 1.7 GHz). Thus, the complexity of the entire cluster system due to lower node numbers is reduced. Another important criterion for the selection of the Intel Xeon processors was not only the availability of the processors, but also the required components such as chipsets and mainboards. Mainboards for the Intel Xeon were announced for the third quarter 2010. AMD mainboards were announced not to ship before the fourth quarter 2010. Delays similar to the construction of the cluster system with the TLB bug in the Barcelona processor should be avoided.

In addition to the processors the complete InfiniBand (IB) will also be replaced. Sun Microsystems was interested in replacing the Magnum switch with the current M9 switch. The Magnum switch was only installed three times world wide. With this exchange the cluster would loose its non blocking IB communication network. The comparison of the existing IB network (DDR non-blocking) with the option QDR and 50% blocking shows that no significant time differences are expected in the operation of the cluster.

On the basis of this result the recommendation for the hardware upgrade was established jointly by Airbus, T-Systems SfR, and the Institute AS in March 2010:

- processor: Intel Xeon 5600 series (Westmere-EP architecture at 2.9 GHz)
- main memory: 24 GB, 1333 MHz DDR3 per node
- communication network: InfiniBand QDR

Publisher:

C²A²S²E - Center for Computer Applications in AeroSpace Science and Engineering

The Executive Board, represented by

Prof. Dr.-Ing.Norbert Kroll

Lilienthalplatz 7

38108 Braunschweig

Germany

Responsible in the sense of § 55, 2nd paragraph of Rundfunkstaatsvertrag (German Interstate Treaty on Broadcasting):

Prof. Dr.-Ing. Norbert Kroll

Deutsches Zentrum für Luft- und Raumfahrt e.V.
in der Helmholtz-Gemeinschaft

Niedersachsen

AIRBUS

- new central switch: Sun M9 (or comparable) with as many ports as possible to minimize the blocking factor
- cluster upgrade time: third quarter 2010
- cluster down time: as short as possible, first estimation minimum 14 days

Due to the acquisition of Sun Microsystems by Oracle and the specific legal conditions, the negotiations between T-Systems and Oracle delayed considerably in the summer months, so the favoured date for the upgrade could not be kept. The upgrade is now planed for November 2010. It will start with a proof of concept. One rack with 48 nodes will be upgraded first to verify the principal upgrade process. Afterwards the complete upgrade will start in two steeps. By end of December 2010 the cluster is expected to go online for production.

During the upgrade the blade servers will be completely substituted by new ones. Furthermore, the existing IB wiring will be completely removed out of the raised floor and replaced by a new installation. The Magnum switch as well as the leaf switches in the computing racks will also be replaced. The capacity of the storage nodes will be slightly increased to 660 TB. However the appearance of the cluster will not change. The computing and cooling racks will remain in operation.

Contact:

Andreas Ernst
(andreas.ernst@dlr.de)

Technical Details

- **13 and a half compute rack**
  - 48 blades per rack
  - two Intel Xeon X5670 at 2,9 GHz processors per blade
  - 6 cores per processor
- **648 nodes**
  - 12 cores per node
  - 24 GB main memory per node
- **total 15.552 GB main memory**
- **7776 cores with 156 TFlop/s peak performace**
- **Fileservice**
  - parallele filesystem (GPFS)
  - 660 TB RAID 6 (+ 25 %)
- **communication network**
  - InfiniBand QDR
  - 40 GB/s
  - Sun M9 Switch 648 Ports
- **12 APC cooling racks (unchanged)**
  - chilled water cooling racks
  - cooling in a closed water circulation
- **electrical power consumption**
  - 255 kW
- **weight 18 t**