ECERTA

Enabling Certification by Analysis

—

Marie Curie Excellence Team

Start: 01 January 2007
Duration: 48 months

www.cfd4aircraft.com

Searching for Transonic Aeroelastic Instabilities Using an Aerodynamic Model Hierarchy

Prepared by: **Sebastian Timme**

Document control data

Deliverable No.:	D **2.3**	Due date:	**01 October 2009**
Version:	**Version 1**	Team Leader:	**Prof. Ken Badcock**
Date delivered:	**18 November 2009**	**Host Organisation : University of Liverpool**	

Project co-funded by the European Commission within the Sixth Framework Programme (2002-2006)		
Dissemination Level		
PU	Public	**x**
PP	Restricted to other programme participants (including the Commission Services)	
RE	Restricted to a group specified by the consortium (including the Commission Services)	
CO	Confidential, only for members of the consortium (including the Commission Services)	

Searching for Transonic Aeroelastic Instabilities Using an Aerodynamic Model Hierarchy

S. Timme* and K. J. Badcock†

CFD Laboratory, University of Liverpool, Liverpool L63 3GH, United Kingdom

*Research Engineer, CFD Laboratory, Department of Engineering. email: sebastian.timme@liverpool.ac.uk

†Professor, Member AIAA, CFD Laboratory, Department of Engineering. email: K.J.Badcock@liverpool.ac.uk

Table of contents

1 Introduction

The physics and driving mechanisms in aeroelastic simulations have to be understood and modelled correctly. This is not likely to be achieved without the ability to consider the impact of uncertainties in the modelling. In structural dynamics, methods to assess uncertainties in modelling parameters are well established [1]. For instance a probability distribution in one (or many) input parameters is propagated through the simulation and the effects on the results are investigated. Propagation tools such as interval analysis, perturbation and polynomial chaos methods, or even brute force Monte Carlo simulations are routinely used. In recent years these tools have started to be transferred to computational fluid dynamics. As an early example, Ref. [1] investigates several propagation methods to address parametric uncertainty in the nonlinear Burgers equation. Approaches considering uncertainties in the physical modelling assumptions rather than the parameters are rare. Examples of dealing with uncertainties in the boundary treatment of a physical model include a geometrically uncertain domain boundary for the two–dimensional Laplace equation [2], or both deterministic and random perturbations on the boundary condition for the viscous Burgers equation subject to sensitivity to the boundary data [3].

As pointed out in Ref. [4], uncertainty from the aerodynamic model is often considerably larger than from the structural model. Also, transonic aerodynamics can have a major influence on nonlinear aeroelastic responses. Two aeroelastic phenomena are particularly associated with the nonlinear flowfield. One is the transonic dip where the presence of shock waves reduces the stability of the aeroelastic system. The second is limit–cycle oscillation (LCO) where the limiting mechanisms of the amplitude of the dynamic response are shock motions and separation. In Ref. [5] the influence of aerodynamic modelling assumptions on the amplitude of a store–induced LCO was investigated for the Goland wing. The dependence of LCO amplitudes on the modelling level, considering both inviscid and viscous flow, was investigated. It was argued that shock/boundary layer interaction in this case causes trailing edge separation and retards the shock movement (substantial in the inviscid case) thus limiting the LCO amplitude. Modelling both the inviscid/viscous interaction as well as the extent of the shock–induced separated regions is important in this case.

Another important contribution to the understanding of transonic LCO is the experimental studies of the supercritical NLR 7301 aerofoil [6–8]. The impact of shock/boundary layer interaction on LCO amplitude, and the importance of the correct prediction of the steady state solution, are discussed and conclusions drawn for the required model fidelity. Nonlinear inviscid aerodynamic modelling should predict LCO dominated by the shock dynamics, whereas aerodynamic modelling including viscous effects is required for LCOs associated with shock/boundary layer interaction [8, 9]. A comprehensive review of nonlinear aeroelastic phenomena including a detailed discussion of possible physical sources for these nonlinearities is given in Ref. [10].

One point missing from the aerodynamic modelling tools available for aeroelasticity is the ability to update lower order models with better available information, and to assess the impact of uncertainties in the prediction of aerodynamic phenomena like shock waves and regions of separation. In structural dynamics these approaches are routinely used when modelling difficult parts of the structure, such as joints. Simple parameterised models are used for these features [11] and experiments are used to tune the

parameters to match observations. The sensitivity to parameters can also be assessed. This approach is felt better than a highly detailed finite element model since it incorporates both the analysts insight into the form of the model and the available measurements. In the aerodynamic case, at least when nonlinear flow models are used, there are no well established methods for exploiting measured data or higher order predictions. There is also no way to assess the impact of uncertainties in the prediction of, for example, the shock wave location on the aeroelastic behaviour. The current paper describes an investigation into developing a method that will address these short–comings for the limited case of transonic aeroelastic predictions using an aerofoil sections moving in pitch and plunge.

The stability of an aeroelastic system can be inferred from time–accurate simulations following an initial excitation of the system. Calculations of complete aircraft configurations have been made [12,13]. The time–accurate approach is very capable due to its generality. However, the significant computational costs, in particular to solve for the unsteady, nonlinear transonic aerodynamics, is a major drawback in this approach. Emphasizing this point is the requirement to search a space of system parameters and flight conditions for critical conditions. The issue of cost generally limits the analysis to a few carefully chosen cases. Alternative approaches have been investigated to obviate the immense computational costs and to permit routine calculations over larger parameter ranges [14]. One popular method is reduced order modelling (ROM) based on proper orthogonal decomposition (POD). For a robust and reliable ROM to exist, the parameter space and flow phenomena of interest have to be covered by the set of system responses used to establish the ROM, and thus, creating the large number of system responses is the main cost in the approach. In addition, the reliability of the POD/ROM approach under parameter changes is a topic of investigation [15].

An alternative approach discussed in the present work uses the theory of dynamical systems to predict aeroelastic instabilities of the Hopf type commonly leading to flutter or LCO. Here, a stability problem for a steady state solution of the aeroelastic system is examined instead of performing unsteady simulations. Stability is lost by a Hopf bifurcation when a pair of complex conjugate eigenvalues of the system Jacobian matrix crosses the imaginary axis for some value of a critical system parameter. Following an approach first published in Ref. [16–18], the bifurcation method was successfully tested on a two–dimensional aerofoil configuration free to move in pitch and plunge. Convergence problems associated with applying a direct solver to a large linear system were resolved by using an iterative sparse linear solver [19]. The method was extended to a larger problem investigating a flexible AGARD 445.6 wing using a modal structural model [20]. Later, the shifted inverse power method was adapted to allow tracing of the critical eigenvalues with changing values of the system parameter. This provides information about the damping and frequency of the aeroelastic modes [21]. Also, a model reduction technique based on the centre manifold theory was investigated to simulated an LCO response in the vicinity of the linear instability point [21].

An improved version of the basic method used a Schur complement formulation to enhance computational performance [22] and was applied to several wing structures and also complete aircraft configurations to study uncertainty in the predicted instability due to structural variability [23]. The current paper takes a next step in the development of the Schur complement framework by forming a reduced order model based on an interpolation technique for the representation of the fluid contribution in the

aeroelastic system and by using a hierarchy of flow models to complement to previous Euler–based results. Therefore, two aerodynamic codes are discussed, namely an established RANS/Euler code and a newly developed full potential/integral boundary layer code. Then, the reduced model will be exploited for aeroelastic prediction and related issues will be discussed.

2 Governing equations and numerical schemes

In Ref. [24] four main aerodynamic modelling levels are discussed; linear and nonlinear potential, Euler and Navier–Stokes. The physics in the simulation can be built up from linear potential flow. The nonlinear potential model adds nonlinear compressible flow effects. The Euler equations add entropy and vorticity effects, while the Navier–Stokes equations include viscous and heat–conducting effects. For almost all aerodynamic flows of practical interest, the Reynolds–averaged form of these equations must be used, introducing the requirement for turbulence models. The impact of the flow modelling (and discretisation) on aeroelastic stability is addressed in the framework by isolating different physical effects, in particular viscous and entropy effects, in separate submodels. In the current study, the RANS and the Euler equations define the high–fidelity models and provide the reference solutions for the update of the lower–level models which consists of the unsteady full potential equations (with their ability to represent shock waves), a simple wake model for unsteady convection of vorticity behind a lifting surface, and an integral boundary layer model to add viscous effects. Additionally, it is planned to include a shock–correction model to deal with strong shock waves violating the assumptions of the full potential formulation.

An eigenvalue–based stability module is used to compute aeroelastic stability details of which are given in Section 4.1. The critical eigenvalues of the Jacobian matrix of the coupled aeroelastic system are traced to find the crossing with the imaginary axis indicating the loss of stability through a Hopf bifurcation [22]. This method proves significantly more efficient than the unsteady time domain approach.

2.1 Higher–order models — RANS and Euler equation

For the higher–level modelling, an established research code is used [25]. The governing equations are discretised using a block–structured, cell–centred, finite–volume scheme. Here, a finite number of non–overlapping control volumes (grid cells) constitutes the computational domain with the governing equations applied to each control volume in turn. The governing equations are formulated in a time–varying curvilinear coordinate system to facilitate the numerical solution. This is convenient since engineering applications have rather complex geometries which require body–conforming grids of arbitrary local orientation and density. Basic features of the code include; implicit time marching for steady state solves, dual time stepping for unsteady simulations [26] with second–order temporal accuracy, the approximate Riemann solver of Osher and Chakravarthy [27] gives the convective fluxes, MUSCL variable extrapolation achieves essentially second–order spatial accuracy, while van Albada's limiter preserves the monotone behaviour of a first–order scheme, viscous fluxes are evaluated using central differences, resulting linear systems are solved by a preconditioned Krylov subspace iterative method. Boundary conditions are set using two layers of halo cells on the outside of the computational domain. For the interested reader, details are given in the cited literature.

2.2 Lower–order models — FP and FPv equations

The lower–level modelling of the flowfield contains several submodels to include different physical effects. The baseline flow model uses the unsteady full potential formulation (FP) [28]. The equations are written

in terms of the time derivative for density from conservation of mass and of the velocity potential from the unsteady Bernoulli's equation. The unknowns are stored at the vertices of an unstructured triangular grid. Non–overlapping dual cells are formed around the vertices and the equation for conservation of mass is discretised by a second–order finite volume scheme with the fluxes evaluated at the centres of dual cell edges using a linear least squares reconstruction. The velocity is reconstructed at the edge centres whereas a slope limited gradient based upwind formulation is applied for the density. The usual jump in the velocity potential based on the Kutta condition is applied to the potential function along the wake cut, and the unsteady convection–based version of this is used. Boundary conditions are set using a layer of halo vertices. A transpiration boundary condition is applied to solid walls. The Jacobian matrix for the implicit scheme as well as the eigenvalue solver is evaluated analytically.

Viscous effects (FPv) are added to the FP model by a two–equation, dissipation–type closure, integral boundary layer formulation for both laminar and turbulent compressible boundary layers as well as free wakes [29–33]. Transition is fixed and user–defined. The basic integral model is augmented by a third equation to account for upstream history effects on the turbulent shear stresses. In its original formulation the model uses the steady integral momentum and mean kinetic energy equations, whereas the present formulation is meant to solve dynamic aeroelastic problems. Therefore, the unsteady governing integral equations are derived from the unsteady Prandtl boundary layer equations [34]. The system of equations contains more unknowns than equations and the additional unknowns have to be expressed in terms of the primary unknowns. The primary effect of the boundary layer and wake on the outer inviscid flow is to displace the inviscid flow by a distance equal to the displacement thickness [35]. The approach used to model this displacing effect is the 'blowing velocity' concept which calculates a wall–normal velocity based on the information from the viscous solution and imposes it on the solid wall of the original geometry and the wake centre line. Derivatives in the integral boundary layer equations are evaluated by a first–order upwind scheme similar to the wake convection equation, whereas the boundary layer correlations are evaluated at the current streamwise station. The boundary layer residual as well as the expressions of the closure correlations are expanded and differentiated with respect to the primary unknowns to calculate the terms of all blocks in the Jacobian matrix analytically.

Newton's method is used to converge the solution to a steady state. Jameson's dual time stepping [26] with a second–order backward temporal discretisation is applied for unsteady simulations with a steady state being calculated in pseudo time at each real time step.

2.3 Structural model

The structural part in the aeroelastic problem is described by the dynamics of a two degree–of–freedom aerofoil [36]. The 'typical section' aerofoil with oscillating pitching and plunging motion represents the torsional and bending behaviour of a wing structure. The linear model is idealized as a point mass defining the centre of gravity as well as a torsional and a translational spring attached to the elastic centre located some distance away from the centre of gravity.

3 Benchmark validation results

3.1 Steady state and unsteady results

Standard test cases are presented for building confidence in the developed full potential solver coupled with an integral boundary layer model. Steady measurements of pressure distributions and boundary layer quantities in sub– and transonic flow regimes are given in the experimental data base of Ref. [37] for the supercritical RAE 2822 aerofoil. Results of pressure distributions are shown in Fig. 1 for cases 2 (subsonic) and 9 (transonic). The simulations were done assuming free flight conditions. Hence, the numerical flow conditions were adjusted to match the experimental data subject to strong wall interference effects. Table 1 lists numerical as well experimental flow conditions. The RANS simulations were done on a 3–block C–type structured grid with 524×78 control volumes whereas an unstructured grid with 6k control volumes was used for the FPv simulations. The FPv simulations give excellent agreement to the experiments for both cases just as the RANS results. No differences can be found between the results of the two turbulence models. For the RANS simlations both the Spalart–Allmaras (SA) and the $k-\omega$ (SST) turbulence models were used. The integrated aerodynamic coefficients of normal force C_n, pitching moment C_m and drag force C_d are included in the figures and correspond in their agreement to the experimental data with the pressure distributions.

	M / M_{exp}	α / α_{exp}	Re	x_{tr}
Case 2	0.685 / 0.676	$-2.35°$ / $-2.18°$	$5.7E{+}6$	0.10
Case 9	0.735 / 0.730	$2.80°$ / $3.19°$	$6.5E{+}6$	0.03

Table 1. Flow conditions for RAE 2822 aerofoil test cases.

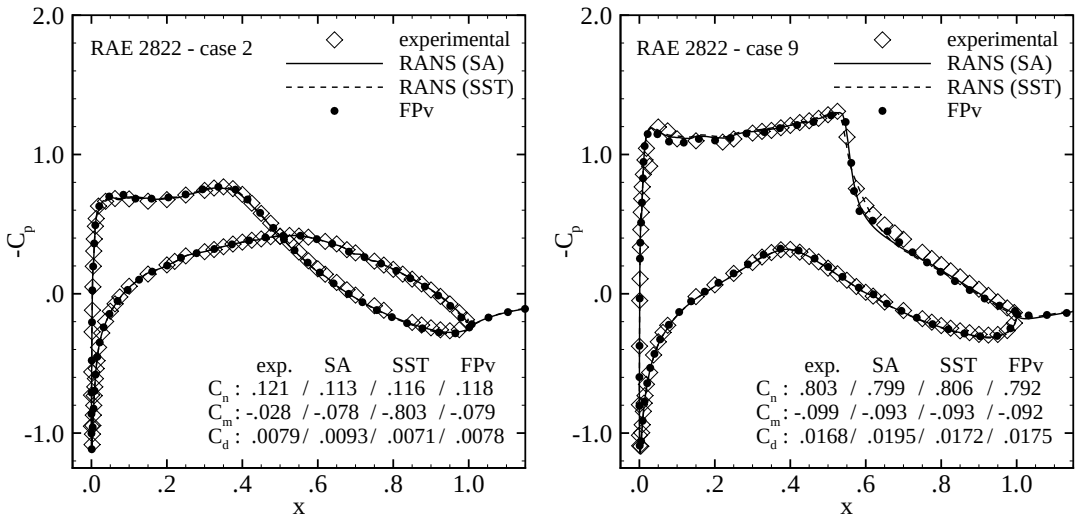

(a) Case 2 (b) Case 9

Figure 1. Pressure distributions of RAE 2822 aerofoil showing comparisons of measurements as presented in Ref. [37] and simulations using RANS and FPv.

Figure 2 presents results of forced pitching motion simulations of the AGARD CT1 test case and compares with experimental data [38]. The configuration CT1 is defined for the NACA 0012 aerofoil at a freestream Mach number of 0.6 and a Reynolds number based on chord length of 4.8 million. The forcing about the quarter chord is prescribed by a sinusoidal motion with reduced frequency of $k = 0.0808$, a mean incidence of $\alpha_m = 2.89°$ and a pitch amplitude of $\alpha_a = 2.41°$. The RANS simulation as well the FPv simulations are assumed to be fully turbulent with a forced transition location in the latter case at about 3 percent chord length. Five motion cycles with 128 steps each were simulated. The RANS (Euler) simulations were done on a 3–block C–type structured grid with 30k (15k) control volumes, and the FP simulations used an unstructured grid with 12k control volumes. The overall agreement of the FP model with or without BL coupling should be considered as excellent, also in view of results of coupled inviscid/viscous simulations as presented in Ref. [39]. Also, a quasi–steady (qs) simulation is shown in Fig. 2(b) where all unsteady terms in the BL model are omitted and a steady boundary layer is calculated at each real time step of the unsteady simulation. The results suggest indeed that a quasi–steady assumption of the BL model is a sufficient simplification as it is commonly done for integral boundary layer simulations. The inviscid results show a consistent trend compared to viscous simulation results and experiments. To explain the relatively large difference between RANS results and experiments in Fig. 2 compare to Ref. [40] where the centre for the moment calculation was shifted by a few percent of the chord length to obtain better agreement.

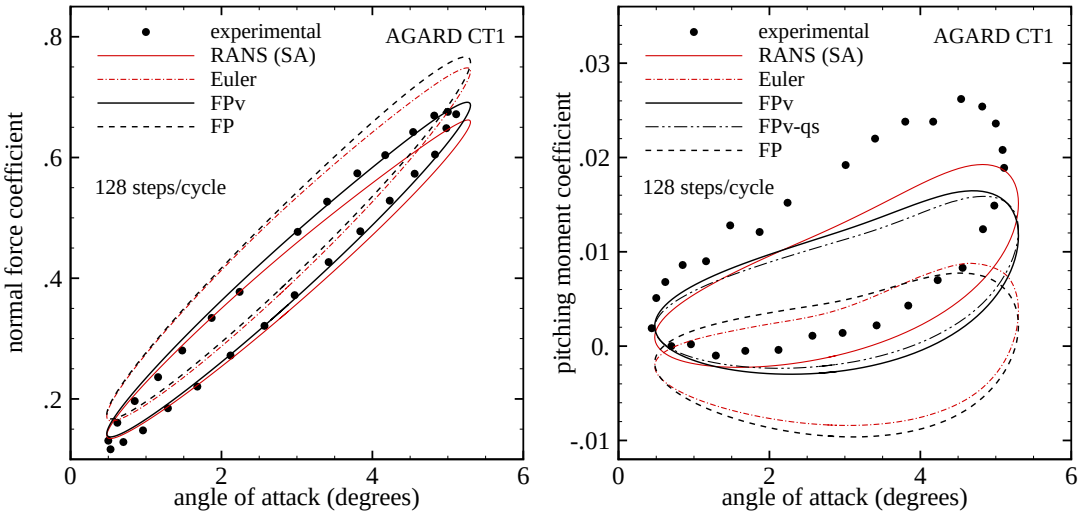

(a) Normal force coefficient

(b) Pitching moment coefficient

Figure 2. Normal force and pitching moment coefficient for AGARD CT1 [38] test case.

The performance of the different flow solver is discussed in terms of the AGARD CT1 case and summarized in Tabs. 2 and 3. Simulating the steady state flow solution at mean conditions is analysed first. For the RANS simulation with 30k control volumes (CV) reducing the residual by 6 order of magnitude takes about 2 minutes of CPU time. Here, 100 explicit steps to smooth out the initial flow field are followed by 270 implicit steps. The Euler simulation for a grid with 15k cv required about 20 seconds of CPU time to reduce the residual by 6 order of magnitude taking 100 explicit and 160 implicit updates. The FP code with and without boundary layer model was tested for two grid sizes. Once

the current approximation is close enough to the converged solution the FP solver achieves quadratic convergence. An inviscid solution without a strong shock wave is converged to machine accuracy within 6 full Newton updates for both the finer grid (12k cv) and the coarser grid (5k cv) requiring 7 and 2 seconds of CPU time, respectively. A viscous simulation takes more attention before quadratic convergence is found. First, an inviscid solution is needed as the initial guess adding to the overall cost of a viscous simulation. Secondly, the viscous updates have to be underrelaxed for stability reasons until the residual has dropped a few orders of magnitude (typically 1 to 2). Thus, it takes another 15 and 5 seconds for the fine and the coarse grids to achieve machine accuracy for a viscous simulation, respectively. An advantage of the FPv formulation is that the same computational grids can be used as for the FP formulation. The costs multiply when a strong shock wave is present in the flow field. Then, the updates have to be underrelaxed to converge to a good approximation required for full Newton updates in both inviscid and viscous simulations.

Model	CV	CPUTIME	ITIME	LSOLV	LTOL	CFL	ITER	TOL
Euler	15k	20	0.12	0.08	10^{-2}	50	260	10^{-6}
RANS	30k	120	0.38	0.16	10^{-2}	50	370	10^{-6}
FP	12k (5k)	7 (2)	1.17 (0.34)	1.09 (0.30)	/	1	6 (6)	10^{-13}
FPv	12k (5k)	15 (5)	1.28 (0.47)	1.18 (0.42)	/	0.5	11 (10)	10^{-13}

Table 2. Performance of different steady flow solver for AGARD CT1 [38] test case.

Model	CV	CPUTIME	ITIME	LSOLV	LTOL	CFL	ITER	TOL
Euler	15k	660	0.1	0.04	10^{-2}	50	10	10^{-2}
RANS	30k	4800	0.3	0.09	10^{-2}	50	20	10^{-2}
FP	12k (5k)	2755 (675)	1.00 (0.31)	0.90 (0.27)	/	1	3–10	10^{-4}
FPv	12k (5k)	6135 (1850)	1.20 (0.43)	1.10 (0.36)	/	0.5	6–12	10^{-4}

Table 3. Performance of different unsteady flow solver for AGARD CT1 [38] test case.

Starting from the steady state solutions the unsteady simulations run for 5 cycles with 128 steps each. The RANS equations require about 20 pseudo iterations at each real time step to achieve the demanded convergence criteria summing up to about 1.5 hours of CPU time. The Euler model takes about 10 pseudo iterations and runs for 11 minutes of CPU time. The unsteady routine in the low–order model solver runs about 45 (11) minutes using the finer (coarser) grid for an inviscid simulation taking a few Newton updates (3–10) depending on the angle of attack to converge to the specified tolerance at each real time step. The viscous simulations require about 2 to 3 times the inviscid runtime. Although the given numbers can only be considered as estimates, they indicate a trend. It can be said that the FP (FPv) solver compares excellent to the established high–level multi–block solver (having been optimized for more than 10 years) for the subsonic and low transonic regime. However, for the higher transonic regime the performance of the FP (FPv) reduces more dramatically compared to the multi–block solver

due to initial underrelaxation of the updates to maintain stable iterations. One of the main issues in the chosen implicit formulations lies in solving linear systems efficiently. Currently a direct solver is employed in the potential flow solver whereas the multi–block solver uses a tailored linear solver with a preconditioned Krylov subspace algorithm outreaching the direct solver dramatically with increasing problem size. Also, the FP solver uses the exact Jacobian of the spatial scheme to achieve quadratic convergence while the high–level solver uses approximate Jacobians for efficiency. For the FP solver, neither an exact Jacobian matrix nor an exact solution to the linear system are useful when the updates are underrelaxed.

Details are given in the tables; overall CPU time (CPUTIME) for steady state solve and unsteady simulation for 640 real time steps, CPU time per iteration to evaluate residuals, Jacobian matrices, to solve the linear system and to update the solution (ITIME), CPU time per one solve of the linear system (LSOLV), CFL number for high–level solver or underrelaxation factor for low–level solver (CFL), and number of iterations (ITER) to achieve the specified tolerance (TOL) either for the steady state or for the current time step in unsteady simulation. The number of iterations includes the number of explicit or underrelaxed updates. All times are given in seconds and represent a fair average.

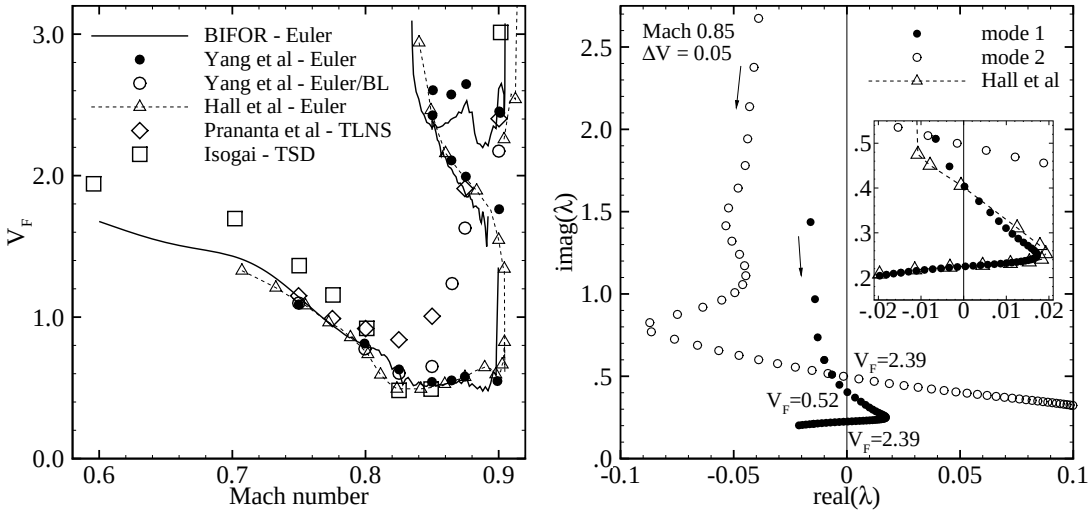

(a) Instability boundary

(b) Mode–tracing

Figure 3. Isogai [41] benchmark case for NACA 64A010 aerofoil configuration compared to numerical results.

3.2 Aeroelastic stability results

The 'typical section' aerofoil of Isogai [41] using the NACA 64A010 aerofoil at zero mean angle of attack is a numerical benchmark case for methods predicting aeroelastic instabilities. The structural parameters of the configuration (presented in Table 5) were chosen to represent the dynamics of an outer section of a swept–back wing. Figure 3(a) shows a comparison between results from different numerical methods (a detailed discussion of which can be found in Refs. [41–44]) illustrating the instability boundary as critical flutter speed index $V_F = \bar{U}/\sqrt{\mu_s}$ versus freestream Mach number. An overall good agreement of the current Euler–based results can be found compared to the different numerical solutions with varying

modelling fidelity. The complex s–shape of the curve in the deep transonic region, giving a second stable branch for higher values of the flutter speed index, is distinct for the inviscid aerodynamic modelling approaches. This behaviour with a double–valued critical flutter speed index disappears for flow models considering viscous flow effects such as thin–layer Navier–Stokes (TLNS) [42] or integral boundary layer (BL) [44] modelling.

The mode–tracing is visualized in Fig. 3(b) showing the migration of the eigenvalues corresponding to the two dominant aeroelastic modes originating in the wind–off structural modes. At fixed Mach number of $M = 0.85$ the s–shaped appearance of the instability boundary is described. Three bifurcations can be found; the first (lower frequency) aeroelastic mode experiences its first bifurcation at a low value of the flutter speed index ($V_F = 0.52$) when entering the unstable region with a positive real part in the critical eigenvalue. Then, at a higher value of the flutter speed index ($V_F = 2.39$) the first mode goes back to a stable behaviour having its second bifurcation. The second aeroelastic mode, experiencing only one bifurcation, is only critical for high values of the flutter speed index and high values of the freestream Mach number between approximately 0.84 and 0.9. The inlay of the figure presents the migration of the eigenvalue of the first mode compared to results of the mode–tracing done in Ref. [43] using a reduced–order model based on a proper orthogonal decomposition technique.

4 Generating Schur correction matrix

4.1 Schur complement formulation

Write the aeroelastic system in semidiscrete form as

$$\frac{d\boldsymbol{w}}{dt} = \boldsymbol{R}(\boldsymbol{w}, \mu) \tag{1}$$

where the vector of unknowns $\boldsymbol{w} = [\boldsymbol{w}_f, \boldsymbol{w}_s]^T$ contains fluid and structural contributions, and $\boldsymbol{R}$ is the corresponding residual vector. The system depends on an independent parameter μ representing for instance dynamic pressure or altitude. An equilibrium solution $\bar{\boldsymbol{w}}$ of the nonlinear system satisfies $\boldsymbol{R}(\bar{\boldsymbol{w}}, \mu) = 0$. In transonic flow the importance of the equilibrium manifests itself in the observation that a shock nonlinearity (strength and location) is defined in the steady flow, while unsteady perturbations about this steady state can be considered to be linear [45]. The theory of dynamic systems gives criteria for an equilibrium to be stable. In particular, stability is determined by eigenvalues $\lambda = \sigma \pm i\omega$ of the system Jacobian matrix $A(\bar{\boldsymbol{w}}, \mu)$ evaluated at the steady state and arbitrary values of μ. A stable system has all its eigenvalues with a negative real part. In many aeroelastic applications a pair of complex conjugate eigenvalues with vanishing real part defines the onset of an instability of the Hopf type. Linear aeroelastic stability is predicted by solving the general eigenvalue problem,

$$(A - \lambda I)\,\boldsymbol{p} = 0 \tag{2}$$

where the Jacobian matrix is conveniently partitioned in blocks expressing the different dependencies

$$A = \frac{\partial \boldsymbol{R}}{\partial \boldsymbol{w}} = \begin{pmatrix} A_{ff} & A_{fs} \\ A_{sf} & A_{ss} \end{pmatrix}. \tag{3}$$

For convenience, the eigenvector $\boldsymbol{p}$ is written just as the vector of unknowns in fluid and structural contributions [22]. Then, the Schur complement formulation can be given,

$$S(\lambda)\,\boldsymbol{p}_s = 0, \tag{4}$$

which defines a small nonlinear eigenvalue problem for the stability analysis. The method to solve this system is outlined in detail in Ref. [22]. The Schur complement is explicitly written as

$$S(\lambda) = (A_{ss} - \lambda I) - A_{sf}\,(A_{ff} - \lambda I)^{-1} A_{fs}, \tag{5}$$

where λ is an eigenvalue of the structural part. This formulation has two major advantages compared to solving the eigenvalue problem as given in Eq. (2); first the ill–conditioning of the system is avoided since λ is not an eigenvalue of A_{ff}, secondly a parallel implementation required for realistic problems is simplified. The Schur complement is formed as the sum of a structural part $S^s = (A_{ss} - \lambda I)$ and a correction matrix $S^c = -A_{sf}\,(A_{ff} - \lambda I)^{-1} A_{fs}$ arising from the coupling with the fluid model. As the coupling is reduced to zero the structural Jacobian is recovered and a structural eigenvalue of the uncoupled system is then obtained. The correction matrix depending on the eigenvalue produces a frequency matching with the aerodynamics. There is another convenient aspect of this formulation; the aerodynamics are stripped free in the Schur complement.

To solve the small complex–valued nonlinear eigenvalue problem given in Eq. (4), the system is augmented by

$$\boldsymbol{c}_s^T \boldsymbol{p}_s - i = 0 \tag{6}$$

to scale the structural eigenvector $\boldsymbol{p}_s$ against a real–valued constant vector $\boldsymbol{c}_s$. Then, the nonlinear system of the Schur residual

$$\boldsymbol{R}_S(\boldsymbol{w}_S) = \begin{pmatrix} S(\lambda)\,\boldsymbol{p}_s \\ \boldsymbol{c}_s^T \boldsymbol{p}_s - i \end{pmatrix} = 0 \tag{7}$$

is solved for $\boldsymbol{w}_S = \left[\boldsymbol{p}_s, \lambda\right]^T$ using Newton's method. The required Schur Jacobian matrix $\partial \boldsymbol{R}_S / \partial \boldsymbol{w}_S$ is given by

$$\frac{\partial \boldsymbol{R}_S}{\partial \boldsymbol{w}_S} = \begin{pmatrix} S & S_\lambda\,\boldsymbol{p}_s \\ \boldsymbol{c}_s^T & 0 \end{pmatrix} \tag{8}$$

with subscript λ expressing a differentiation with respect to the eigenvalue. The improvement is obvious. While the full formulation in Eq. (2) solves a problem with $N_f + N_s + 1$ unknowns, the Schur complement formulation only has $N_s + 1$ where the number of structural unknowns is generally small. While the correction term of the Schur complement is solved directly for the evaluation of the Schur residual, a series approximation is used for calculating the Schur Jacobian. Generally, the series approximation reads

$$\left(A_{ff} - \lambda I\right)^{-1} = A_{ff}^{-1} + \lambda A_{ff}^{-1} A_{ff}^{-1} + \lambda^2 A_{ff}^{-1} A_{ff}^{-1} A_{ff}^{-1} + \ldots \tag{9}$$

where λ must be small for the series to converge [46]. To overcome this restriction, write the eigenvalue as $\lambda = \lambda_0 + \lambda_\varepsilon$ where λ_0 is the initial guess, e.g. a natural frequency or a previous converged solution, and λ_ε is a small variation to the initial guess, and solve the system for the variation instead of the eigenvalue itself. Then, the series can be written as

$$\left(A_{ff} - \lambda I\right)^{-1} \approx \left(A_{ff} - \lambda_0 I\right)^{-1} + \lambda_\varepsilon \left(A_{ff} - \lambda_0 I\right)^{-1} \left(A_{ff} - \lambda_0 I\right)^{-1} \tag{10}$$

to form the term S_λ in the Schur Jacobian. Details about approximations to decrease the computational costs as well as the parallel implementation are given in Ref. [22].

Evaluating the correction term S^c accounts for the highest cost of the stability analysis since $2n$ linear system have to be solved on the fluid Jacobian matrix for a system with n degree–of–freedom whereas the costs to form the term S^s are negligible. For the 2 degree–of–freedom aerofoil structural model the size of the Schur complement matrix is 4×4. Out of these 16 elements there are 6 nonzero complex–valued elements in the correction term,

$$S^c = -A_{sf} \left(A_{ff} - \lambda I\right)^{-1} A_{fs} = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & S_{22}^c & S_{23}^c & S_{24}^c \\ 0 & 0 & 0 & 0 \\ 0 & S_{42}^c & S_{43}^c & S_{44}^c \end{pmatrix}. \tag{11}$$

Here, the structural unknowns are ordered as $\boldsymbol{w}_s = [h, \dot{h}, \alpha, \dot{\alpha}]^T$. The first column of the correction matrix is found to be zero due to the independence of the fluid response on the plunge state h. The first and the third row are zero due to the composition of matrix A_{sf} projecting the fluid response onto the structural states.

4.2 Extracting elements based on Fourier series

Solving the linear systems for the steady state fluid Jacobian matrix A_{ff} to form the Schur correction matrix directly is referred to as linear frequency domain approach. Alternatively, it is also possible to evaluate the correction matrix from a Fourier analysis of unsteady time–domain responses forced in the structural states. Therefore, write the unknowns as the sum of a steady state solution $\bar{\boldsymbol{w}}$ and an unsteady perturbation $\mathfrak{w}$ about this mean;

$$\boldsymbol{w}_f(t) = \bar{\boldsymbol{w}}_f + \mathfrak{w}_f(t) \qquad \text{and} \qquad \boldsymbol{w}_s(t) = \bar{\boldsymbol{w}}_s + \mathfrak{w}_s(t). \tag{12}$$

Then, the fluid contribution of the aeroelastic system in Eq. (1) can be given in its time–linearized form as

$$\frac{d\mathfrak{w}_f}{dt} = A_{ff}(\bar{\boldsymbol{w}})\,\mathfrak{w}_f + A_{fs}(\bar{\boldsymbol{w}})\,\mathfrak{w}_s, \tag{13}$$

where the fixed Jacobian matrices A_{ff} and A_{fs} are evaluated at the steady state. The latter equation essentially constitutes a standard state–space representation, where $\mathfrak{w}_s$ acts as the input vector of external forcing. This system is solved for forced periodic motions in the structural unknowns driving the solution of the unsteady fluid perturbation. Write the unsteady perturbation in a truncated Fourier series [47]

$$\mathfrak{w}_f = \sum_{n=-N}^{N} \boldsymbol{\alpha}_n\, e^{i\,n\,\omega\,t} \qquad \text{and} \qquad \mathfrak{w}_s = \sum_{n=-N}^{N} \hat{\boldsymbol{\alpha}}_n\, e^{i\,n\,\omega\,t}, \tag{14}$$

where ω is the fundamental frequency of the problem, and substitute latter expressions into Eq. (13). Rearranging gives

$$\sum_{n=-N}^{N} \boldsymbol{\alpha}_n = -\sum_{n=-N}^{N} \left(A_{ff} - i\,n\,\omega I\right)^{-1} A_{fs}\,\hat{\boldsymbol{\alpha}}_n. \tag{15}$$

The vector of complex–valued Fourier coefficients $\boldsymbol{\alpha}_n$ (and $\hat{\boldsymbol{\alpha}}_n$) is evaluated from the time signal over a period $T = 2\pi/\omega$ as

$$\boldsymbol{\alpha}_n = \frac{1}{T} \int_{-T/2}^{T/2} \mathfrak{w}_f(t)\, e^{-i\,n\,\omega\,t}\, dt, \qquad n = 0, \pm 1, \pm 2, \ldots \tag{16}$$

and accordingly for the structural contribution. There are $2N + 1$ equations for the Fourier coefficients with $n = 0$ corresponding to mean flow which is automatically satisfied since the unsteady perturbation about the steady state is considered.

Comparing to Eq. (5) it can be seen that, after multiplying with the Jacobian matrix A_{sf}, the Fourier coefficients of the fluid responses in each control volume of the computational domain constitute the correction term in the Schur matrix with a purely imaginary eigenvalue. More precisely, the fluid responses in the control volumes contributing to the aerodynamic forces acting on the structure have to be analysed. Evaluating the Fourier coefficients at integer multiples of the fundamental frequency (provided the system was excited accordingly) gives the correction matrix at these discrete frequencies. The matrix A_{sf} used to project the fluid response onto the structural states can be evaluated easily by finite differences. The Fourier coefficients of the forced excitation $\hat{\boldsymbol{\alpha}}_n$ set the column of the correction matrix and scale for the amplitude of the excitation; for instance forcing the system of the 2 degree–of–freedom aerofoil with $\mathfrak{w}_s = [0, 0, \hat{a}\,\sin(\omega\,t), 0]^T$ gives the third column of the 4×4 matrix at the frequency ω.

Solutions of a fully nonlinear system approach the time–linearized results if the amplitude of the forced motion is sufficiently small. This is that the assumptions of a linearized solution hold where the unsteadiness in the flow is linearly dependent on the structural motion. The step of using the nonlinear aeroelastic system or more precisely the nonlinear aerodynamic system with a forcing applied in the structural unknowns $\boldsymbol{w}_s$,

$$\frac{d\boldsymbol{w}_f}{dt} = \boldsymbol{R}_f(\boldsymbol{w}_f, \boldsymbol{w}_s), \tag{17}$$

is required if the Jacobian matrices for the fluid contribution, as given in the linearised system in Eq. (13), are not available explicitly. Also, it is possible to apply this method to arbitrary flow solver giving access to an arbitrary fidelity in the flow modelling.

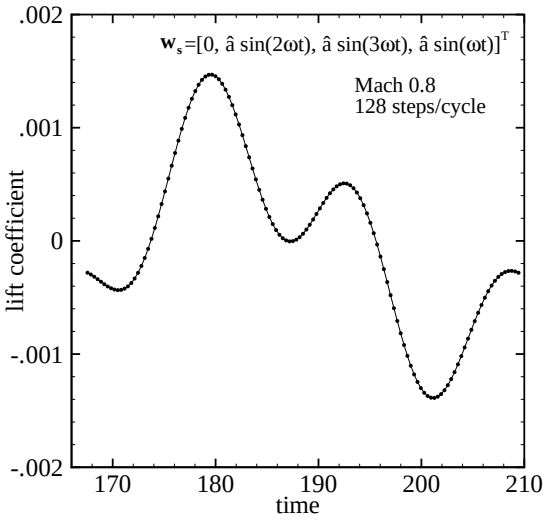

Figure 4. One period of motion with simultaneous forcing in three relevant structural states showing integrated aerodynamic coefficient of lift to represent fluid response in all control volumes.

An example simulation to demonstrate the three approaches of linear frequency domain, linearized time domain and nonlinear time domain is presented in Figs. 4 and 5. Here, a NACA 0012 aerofoil configuration as defined in Ref. [19] was excited in all three structural states of interest simultaneously at a fundamental frequency of $\omega = 0.15$ and an amplitude of $\hat{a} = 1.0 \times 10^{-4}$. Three simulations were required to obtain the complete correction matrices at three frequencies while swapping around the factors multiplying the fundamental frequency, for instance

$$\begin{aligned} \mathfrak{w}_s &= [h, \dot{h}, \alpha, \dot{\alpha}]^T \\ &= [0, \hat{a}\sin(2\,\omega\,t), \hat{a}\sin(3\,\omega t), \hat{a}\sin(\omega t)]^T. \end{aligned} \tag{18}$$

Exciting the plunge coordinate h is irrelevant. Following the transition to stable periodic cycles, one motion cycle of period $T = 2\pi/\omega$ is simulated with 128 time steps and used for the extraction of the Fourier coefficients. Two freestream Mach numbers at 0.7 and 0.8 are considered and represent a subsonic and a transonic case with strong shock wave. Figure 5 describe the development of two elements of the Schur correction matrix showing the real and the imaginary parts individually. The agreement is excellent between the three approaches.

Figure 5 also includes results for a nonlinear time domain simulation with a higher excitation amplitude of $\hat{a} = 1.0 \times 10^{-2}$. Intuitively, looking at the results for higher excitation amplitudes suggests that at

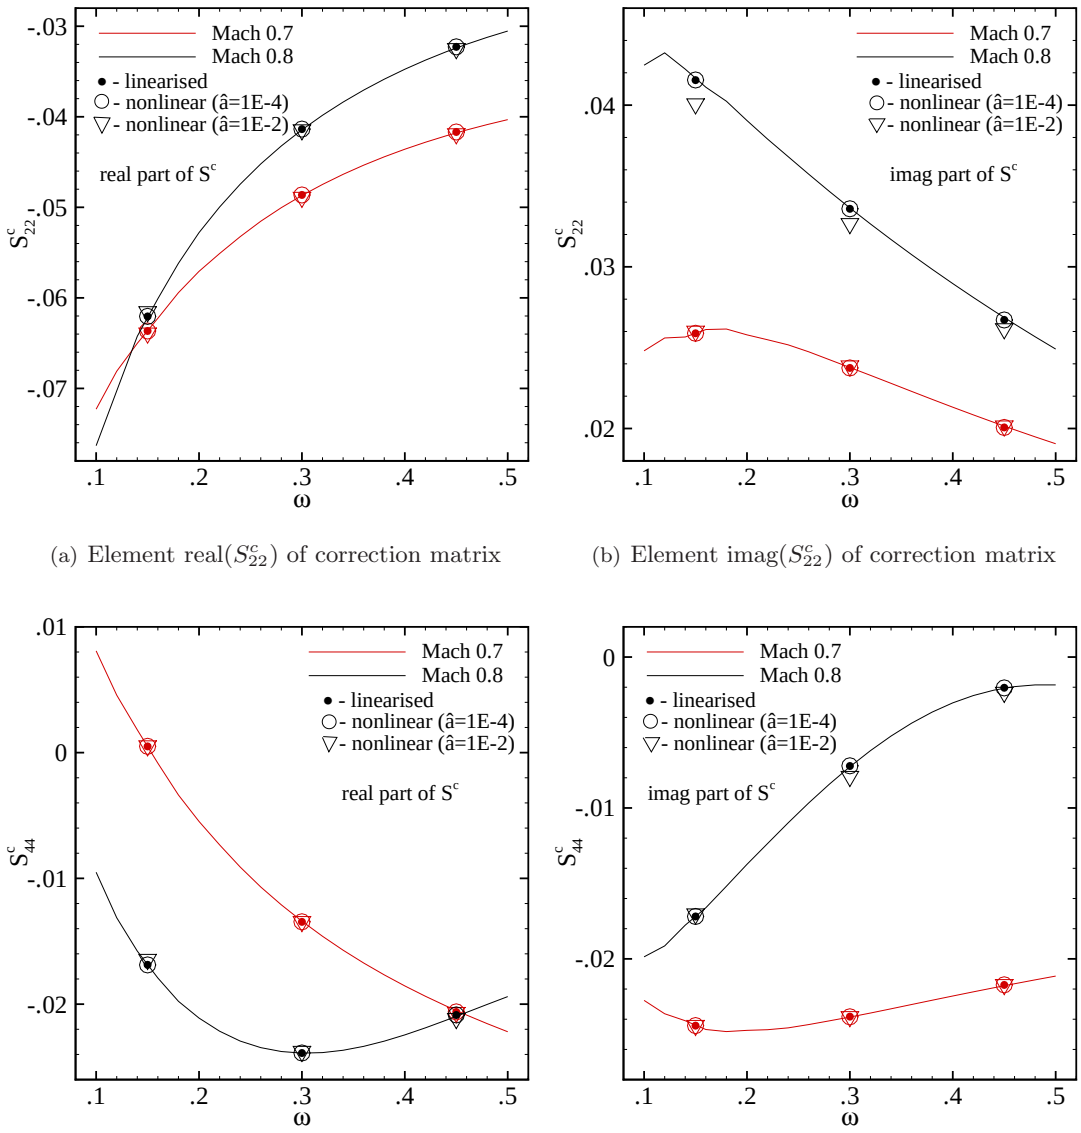

(a) Element real(S^c_{22}) of correction matrix

(b) Element imag(S^c_{22}) of correction matrix

(c) Element real(S^c_{44}) of correction matrix

(d) Element imag(S^c_{44}) of correction matrix

Figure 5. Extracted elements of Schur correction matrix for NACA 0012 configuration using the Euler equations; linearised and nonlinear time–domain simulations compared to linear frequency domain results.

Mach 0.8 the nonlinear time domain results lose accuracy compared to the linear frequency or linearized time domain results. However, at Mach 0.8 the appearance of a transonic shock wave introduces an additional issue in the analysis. In Ref. [48, 49] an oscillatory behaviour in transonic flow was discussed. Therein, the instability boundary of the NACA 0012 aerofoil configuration showed an oscillatory trend with changes in Mach number caused by a discrete movement of the shock wave. Due to the discrete numerical representation the movement of the shock wave is restricted to the resolution of the computational grid. The oscillatory instability boundary was related to an oscillation in the elements of the Jacobian matrix. Figure 6 presents such a nonsmooth behaviour for one element with real and imaginary parts of the Schur correction matrix using the Euler and the RANS flow models where the two turbulene model are shown for the RANS results. Nonlinear time domain results at a range of Mach numbers and four different excitation amplitudes $\hat{a}$ with increasing strength are presented. The time domain results

are compared to linear frequency domain predictions for the Euler flow model. As discussed in Ref. [49] for the influence of the initial disturbance on the unsteady simulations, a dependence on the amplitude can be found. Small excitation amplitudes give results very close to the linear frequency analysis with the distinct oscillatory phenomenon, whereas higher values eliminate the oscillations. This was related to the location of the shock wave depending on the deflection of the structure. A weak variation of the pressure distribution is found for small motion amplitudes with a strong influence of the discrete steady–state shock resolution throughout the unsteady forcing. The dynamic effects due to larger amplitudes on the other hand dominate the influence of the steady state. It seems to be more meaningful physically to use a higher excitation amplitude since the nonsmooth behaviour, which cannot be explained with arguments of a continuous change of a system parameter, disappears as seen in Fig. 6. However, if such a nonsmooth trend is also found in viscous simulations then more consideration is necessary due to the presence of possible boundary layer separation.

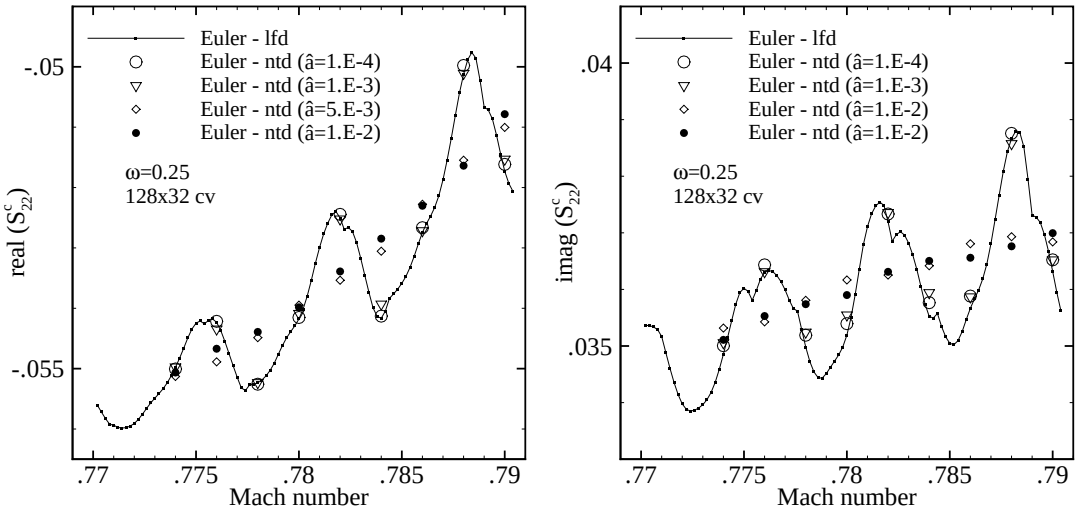

(a) Element real(S^c_{22}) for Euler flow model

(b) Element imag(S^c_{22}) for Euler flow model

Figure 6. Effect of excitation amplitude on oscillatory behaviour in transonic regime showing results of Euler flow model.

The integrated aerodynamic loads showed the same oscillatory behaviour in the transonic regime not only for inviscid [49] but also for viscous simulations (not presented herein). Thus, this unwelcome oscillatory phenomenon should also be expected for RANS simulations as presented in Fig. 7. The figure shows such a nonsmooth behaviour for one element of the Schur correction matrix using the RANS flow model with two different turbulence models. Three distinct regions can be distinguished; a shock free region (subsonic and very low transonic), a region with a distinct shock wave, and a region with shock–induced boundary layer separation. For shock free flow, in the figure shown at Mach 0.75, the excitation amplitude (chosen within reason) is irrelevant. If at all, slight differences can be found between different turbulence models. Having a distinct shock wave between Mach 0.78 and 0.81, the amplitude of the forcing becomes an important factor and an oscillatory behaviour as for the Euler simulations is found. The two smaller amplitudes clearly scatter around a "mean value". Interestingly, the assumed peaks and troughs for the Schur correction element are different for the two considered

turbulence models. This should be attributed to the differences in how the boundary layer is predicted, and consequently to the slight differences in shock wave location at a given Mach number. The results for the two higher amplitudes, leaving the constraints of the discrete grid resolution, show a converging trend in the prediction. As for the Euler results this smooth trend for the elements should be considered as physically more meaningful. A trend for the last region with shock–induced boundary layer separation at Mach 0.83 is less easy to establish. While the real part of the considered element shows relatively little (SA model) to essentially no spread (SST model) between the different amplitudes, the imaginary part behaves more or less like in the second region of interest. Also, the real part gives a distinct difference between the two turbulence model which should be related to the behaviour of the individual turbulence model in separated flow. It must also be said that more than one Mach number would be required to establish a clear picture of whether or not there is large spread since it is not known where on the nonsmooth curve the simulation is done.

When dealing with uncertainty in the prediction of aeroelastic stability, the impact of this nonsmooth trend should be included in the discussion in particluar when using an analysis algorithm based on steady state flow solutions. For the extraction of elements for the Schur correction matrix based on forced time domain simulations, an excitation amplitude with an order of about 10^{-2} seems to be sufficient and will be used.

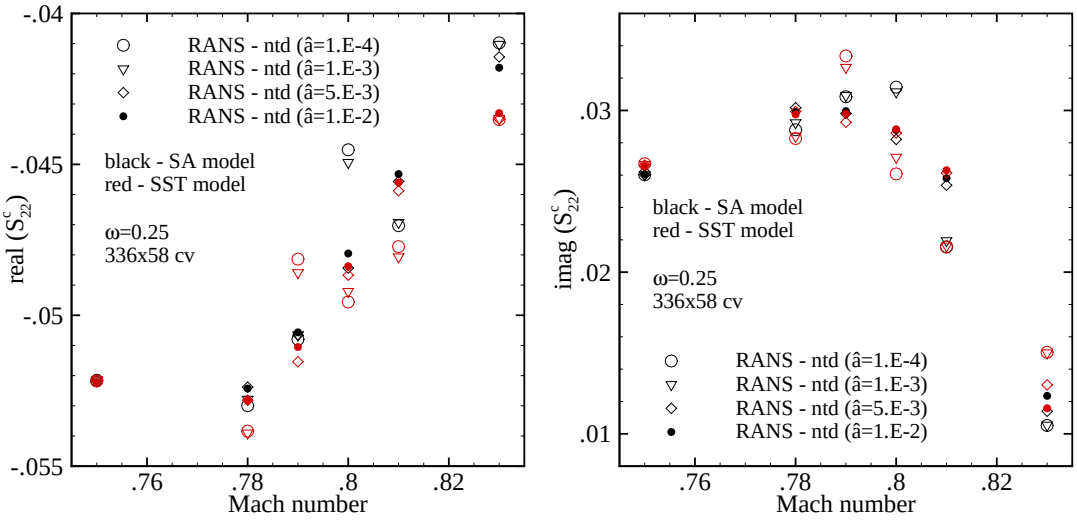

(a) Element real(S^c_{22}) for RANS flow models

(b) Element imag(S^c_{22}) for RANS flow models

Figure 7. Effect of excitation amplitude on oscillatory behaviour in transonic regime showing results of RANS flow models using Spalart–Allmaras (SA) and Menter's shear–stress transport (SST) turbulence modelling.

Alternatively, instead of exciting all structural states simultaneously, the structural state of interest could be forced at several frequencies at once which would give the corresponding column of the correction matrix at these distinct frequencies, or each structural state could be excited at each frequency individually. Of course, the time–step Δt for the forced motion simulation has to be chosen to resolve the dynamic content of the highest frequency accurately. Therefore, a number of simulations has been made with the target to extract the Schur correction matrix at four discrete frequencies using different

	Mach 0.6			Mach 0.8		
	32 steps	64 steps	128 steps	32 steps	64 steps	128 steps
1f (ltd)	1.26	0.32	0.10	1.01	0.33	0.25
4f (ltd)	7.12	1.87	0.46	14.6	2.78	0.61
4f (ntd)	/	/	0.49	/	/	0.66
fp (lfd)		17.5			108	

Table 4. Root mean square error of different approaches to extract Schur correction elements compared to linear frequency domain Euler results; lfd – linear frequency domain, ltd – linear time domain, ntd – nonlinear time domain.

approaches. In Table 4 the root mean square error (RMSE) of the different approaches is summed for all Schur elements and for all four discrete frequencies and compared to a linear frequency domain Euler solution as reference. Mach numbers at 0.6 and 0.8 are considered. The discussed approaches include linear time domain simulations exciting one frequency per structural state at a time (1f) and exciting four frequencies per structural state simultaneously (4f). Obviously, exciting only one frequency at a time requires four unsteady simulations per structural state. Also, a nonlinear time domain solution with a simultaneous excitation in four frequencies and a linear frequency domain solution for a FP simulation are presented. The RMSE is calculated as

$$RMSE = \sqrt{\frac{1}{N}\sum_{n}^{N}\left(1 - \frac{S_n^c}{S_n^{c,ref}}\right)^2} \times 100\,[\%] \tag{19}$$

where the index n runs over all elements of the correction matrix at all frequencies. The large error of the lower–level FP simulations compared to the higher–level Euler formulation is distinct building confidence in the quality of the time domain based extraction. Also, the table gives an good impression that a finer time step is required to have a comparable error when exciting several frequencies at the same time. The method of choice would depend on the information that is sought, for instance if there is a dominant term in the matrix then it is useful to excite this term at several frequency to get many data points.

Another interesting point was observed. It appears that the correction elements in the second column (sensitivity to plunge rate) have the same absolute value with opposite sign compared to the elements of the third column (sensitivity to pitch angle). An explanation to this can be found for instance by looking at the classical aerodynamic theory of Theodorsen [36]. Here, an increment in the pitch angle produces a contribution to the overall lift equal to an increment in the plunge rate but with opposite sign. Thus, a positive deflection in the pitch angle produces a positive lift whereas a positive deflection in the plunge rate (positive upward) produces a negative lift. This is illustrated in Fig. 8 for two relevant elements showing real and imaginary part for the Euler and the FP models at a Mach number of 0.75. It can be seen that the real part of the chosen elements have almost identical absolute values for the Euler formulation while the difference in the FP results is more distinct. The imaginary part also shows a bigger difference for the Euler model. It is interesting to remark that the response to the pitch deflection (coloured black) of the FP model agrees better to the Euler values than the response to the plunge rate (coloured red). This behaviour has been found to be consistent for all considered Mach numbers. The

question is whether this is a feature of the flow model, the discretisation or an indicator of an error in the code.

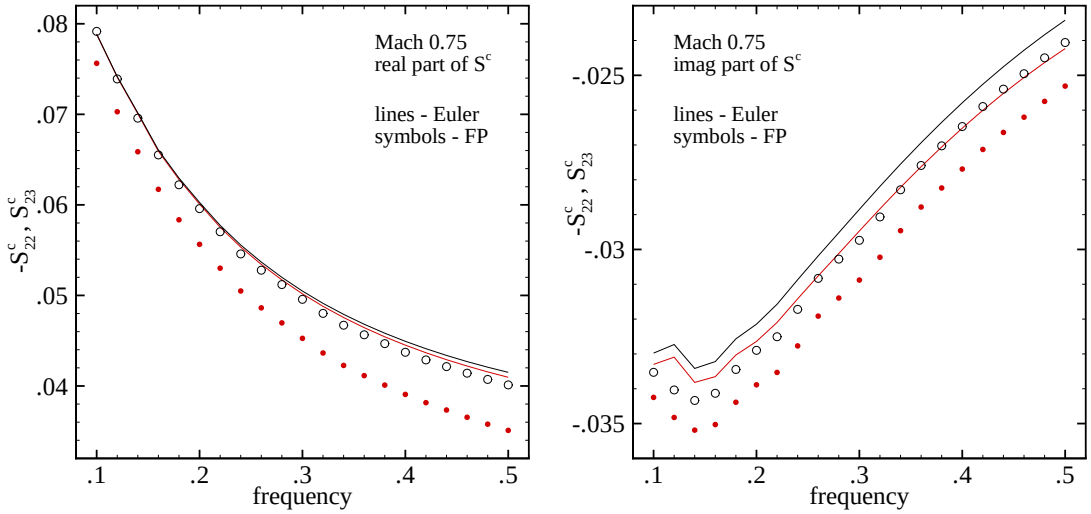

(a) Element real(S^c_{22}) of correction matrix

(b) Element real(S^c_{22}) of correction matrix

Figure 8. Illustrating the effect of deflections in plunge rate and in pitch angle for the NACA 0012 aerofoil configuration.

Apparently, the forced time domain simulations contain the main cost of the approach considering that the eigenvalue–based formulation was developed in the first place to avoid excessive unsteady simulations in predicting aeroelastic instabilities. For how much it accounts is considered next. As discussed in Section 3.1, an Euler steady state solution in subsonic and low transonic flow on a grid with 15k control volumes is found in less than half a minute of CPU time. In the higher transonic regime this should be multiplied by a factor of about three. Performing the forced motion simulations was found to be more or less independent of the considered Mach number. Creating the Schur correction matrix at one distinct combination of Mach number and frequency requires one unsteady simulation for each relevant column of matrix A_{fs}, i.e. states of the structural model with plunge rate, pitch and pitch rate. The plunge state is irrelevant. With the time–linearized formulation it involves about 3 to 5 times the cost of a subsonic (or low transonic) steady simulation for one structural state. Here, the simulation is run for four forced motion cycles with 32 time steps each while the Fourier coefficients are evaluated only from the response during the last cycle. The nonlinear time domain approach has about twice the cost of the linearized version. The linear formulation is faster since the unsteady residual is evaluated by a simple matrix vector multiplication. Evaluating the Jacobian matrices and one Schur correction matrix at an individual frequency with the linear frequency domain solver is about the same cost as simulating a steady state. It involves solving a linear system on the second–order fluid Jacobian matrix against each relevant column of matrix A_{fs}.

Summarizing, the linearized time domain approach involves the costs of about 9 to 15 steady state simulations to extract one Schur correction matrix at an individual frequency, while the nonlinear version is about 18 to 30 steady state simulations. These results however should only be considered as estimates. Also, the costs of the time–domain approach can be decreased by reducing the number of forced motion

cycles (while a converged motion cycle is still required). Using the observation that the responses to plunge rate and to pitch angle produce results of the same absolute value could be useful as well, however this needs to be confirmed for cases with very strong shock waves and flow separation.

5 Constructing a reduced order model

In the previous section different approaches to create the Schur correction matrix have been discussed. Even though the Schur complement eigenvalue–based approach to predict aeroelastic stability is already by many order of magnitude more efficient than common brute force time marching simulations, especially when sensitivity to parameters requires a large number of evaluations, an alternative reduced order model of the Schur complement formulation is introduced. Solving the full order Schur problem requires solving $2n$ linear systems on the exact (second–order) fluid Jacobian matrix where n is the number of DOF of the structure. Generating the Schur Jacobian matrix to apply Newton's method takes some considerable computing time extra; in the current formulation another $2n$ linear solves are needed to build the series as given in Eq. (10). Ways to approximate both the residual and the Jacobian matrix for the Schur complement formulation are discussed in Ref. [23]. The reduced order model (ROM) of the Schur correction matrix uses an interpolation technique which is based on samples of the full order model. The kriging interpolation technique will be discussed first. Then, the ROM applied to the aeroelastic stability analysis is presented.

5.1 Approximating the Schur correction matrix (the kriging predictor)

For computationally expensive simulations (numerical experiments) it is useful to generate a cheap approximation (surrogate) based on relatively few runs of the expensive model to provide information about its response at untried locations (parameter combinations of the model). Numerical experiments are different from physical experiments in a sense that a repeated (deterministic) simulation gives exactly the same result whereas replicated physical trials at the same conditions scatter around a mean answer due to inherent variability. Computational experiments lack the random error. Thus, an approximation model should both predict the calculated responses exactly and adapt to the functional behaviour of the responses. The kriging (nonlinear least squares) interpolation technique (first developed in the field of geostatistics) was introduced in Ref. [50] for deterministic simulations in computer–based disciplines. It is also referred to as DACE (design and analysis of computer experiments) model. Different kriging methods can be described which are distinguished by the way a functional trend is modelled; for instance simple, ordinary and universal kriging. Simple kriging is applied for a constant and known mean value whereas ordinary kriging estimates an unknown constant mean in the vicinity of interest. Universal kriging (kriging with a trend) uses a smoothly varying low–order polynomial regression model which is identical to ordinary kriging for a zeroth–order polynomial.

In the kriging approach a multidimensional deterministic response $\boldsymbol{y}(\boldsymbol{x})$ of a simulation is treated as a realisation of a stochastic process composed of a low–order regression model and a random normally distributed signal with zero mean and covariance $\sigma^2 R$. Here, σ^2 denotes the process variance, i.e. the variance of the input samples, and $R = R(\boldsymbol{x}, \boldsymbol{w})$ is the correlation between locations $\boldsymbol{x}$ and $\boldsymbol{w}$. Thus, the second term (the error term) is not independent at different locations but is related to the distance between points in the parameter space. Focusing on the correlation structure rather than on a more sophisticated regression model is an important feature of the kriging approach to reconstruct response surfaces. The parameters of the computationally cheap kriging model are fitted to a provided set of

numerical samples of the full order formulation by an optimisation process a detailed discussion of which can be found for instance in Ref. [50, 51].

Consider a given set of n numerical experiments (samples) as $S = \big[\boldsymbol{s}_1, \ldots, \boldsymbol{s}_n\big]^T$ and the corresponding system response $\boldsymbol{y}_s = \big[y(\boldsymbol{s}_1), \ldots, y(\boldsymbol{s}_n)\big]^T$. For convenience a single scalar response y is assumed to be a function of the m–dimensional input vector $\boldsymbol{s}$. The expressions generalize for multidimensional responses. The best linear unbiased predictor, herein referred to as the kriging predictor, minimizes the error of the interpolation and was given in Ref. [50] as

$$\hat{y}(\boldsymbol{x}) = \boldsymbol{f}(\boldsymbol{x}) \cdot \boldsymbol{\beta} + \boldsymbol{r}(\boldsymbol{x}) \cdot \big(R^{-1}(\boldsymbol{y}_s - F\boldsymbol{\beta})\big), \tag{20}$$

where the first term is the regression model and the second term adjusts the prediction based on the correlation. The kriging predictor approximates the system response at an unsampled location $\boldsymbol{x}$ in the parameter space at the expense of only two scalar products on $\boldsymbol{f}(\boldsymbol{x})$ and $\boldsymbol{r}(\boldsymbol{x})$ once the model is formed. It can be shown that the exact system response with zero error is interpolated at a sampled location. The mean squared error φ^2 of the predictor is evaluated by

$$\varphi^2(\boldsymbol{x}) = \sigma^2\Big(1 - \boldsymbol{r}(\boldsymbol{x}) \cdot R^{-1}\boldsymbol{r}(\boldsymbol{x}) + \boldsymbol{u} \cdot \big(F^T R^{-1} F\big)^{-1}\boldsymbol{u}\Big) \tag{21}$$

with the vector $\boldsymbol{u}(\boldsymbol{x}) = F^T R^{-1}\boldsymbol{r}(\boldsymbol{x}) - \boldsymbol{f}(\boldsymbol{x})$ and the process variance

$$\sigma^2 = \frac{1}{n}(\boldsymbol{y}_s - F\boldsymbol{\beta})^T R^{-1}(\boldsymbol{y}_s - F\boldsymbol{\beta}). \tag{22}$$

The root mean squared error φ is referred to as standard error of the model and is a measure of uncertainty in the prediction. The second term in Eq. (21) reduces the prediction error since the unsampled location $\boldsymbol{x}$ is correlated with the design samples S, whereas the third terms adjust for errors in estimating the regression model [51]. The vector of regression parameters $\boldsymbol{\beta}$ is the generalized least squares estimator of the overdetermined regression problem $F\boldsymbol{\beta} = \boldsymbol{y}_s$ and is given by the expression

$$\boldsymbol{\beta} = \big(F^T R^{-1} F\big)^{-1} F^T R^{-1}\boldsymbol{y}_s, \tag{23}$$

as can be found in literature of statistics. Here, the $n \times p$ regression matrix F is written as

$$F = \big[\boldsymbol{f}(\boldsymbol{s}_1), \ldots, \boldsymbol{f}(\boldsymbol{s}_n)\big]^T \tag{24}$$

with $\boldsymbol{f}$ as the basis vector of the regression model. For a zeroth–order regression model with $p = 0$, matrix F is a $n \times 1$ column vector filled with ones, whereas in the case of a linear regression model with $p = m + 1$ the columns of the matrix are given by

$$\boldsymbol{f}(\boldsymbol{s}) = \big[1, s_1, \ldots, s_m\big]^T. \tag{25}$$

For many situations a regression model with zeroth–order was found to be sufficient while a higher order did not offer advantages [52]. The $n \times n$ correlation matrix R of the samples is built from the elements

$$\mathcal{R}_{ij}(\boldsymbol{\theta}, \boldsymbol{p}, \boldsymbol{s}_i, \boldsymbol{s}_j) = \prod_{k=1}^{m} scf\big(\theta_k, p_k, s_k^{(i)} - s_k^{(j)}\big) \qquad 1 \le i, j \le n \tag{26}$$

where scf is a spatial correlation function of the arguments $\boldsymbol{\theta}$, $\boldsymbol{p}$ and the difference between samples $\boldsymbol{s}_i$ and $\boldsymbol{s}_j$. The correlation parameter θ_k indicates the activity of the independent variable k, while p_k is

Aerofoil	Ref.	x_{cg}	x_α	r_α	ω_r	ζ_h	ζ_α	α_0	μ_s
NACA 64A010	[41]	0.4	−1.8	1.865	1	0	0	0	60
NACA 0012	[19]	0.5	−0.2	0.539	0.343	0	0	0	100

Table 5. Parameters of aeroelastic aerofoil configurations.

a measure of the smoothness of the predictions in coordinate direction k. Several correlation functions have been given in the literature reflecting characteristics of the system output. In this study exponential and spline functions are used [53]. The exponential correlation function is written as

$$scf = \exp\left(-\theta_k \left|s_k^{(i)} - s_k^{(j)}\right|^{p_k}\right) \qquad p_k \in [1,2] \tag{27}$$

whereas the spline function is given by

$$scf = \begin{cases} 1 - 15\,\xi_k^2 + 30\,\xi_k^3 & \text{for } 0 \le \xi_k \le 0.2 \\ 1.25\,(1-\xi_k)^3 & \text{for } 0.2 < \xi_k < 1 \\ 0 & \text{for } \xi_k \ge 1 \end{cases} \tag{28}$$

with argument $\xi_k = \theta_k \left|s_k^{(i)} - s_k^{(j)}\right|$. The optimal correlation coefficients $\boldsymbol{\theta}$ and $\boldsymbol{p}$ are given by the maximum likelihood estimate [50, 52] and minimize the expression $\det\left(R^{1/n}\sigma^2\right)$ having the vector of regression coefficients $\boldsymbol{\beta}$ and process variance σ^2. Finally, the vector $\boldsymbol{r}$ is written as

$$\boldsymbol{r}(\boldsymbol{x}) = \left[\mathcal{R}_{1,1}(\boldsymbol{\theta}, \boldsymbol{p}, \boldsymbol{s}_1, \boldsymbol{x}), \ldots, \mathcal{R}_{n,1}(\boldsymbol{\theta}, \boldsymbol{p}, \boldsymbol{s}_n, \boldsymbol{x})\right]^T, \tag{29}$$

and contains the correlations between the samples S and an unsampled location $\boldsymbol{x}$.

The gradient of the predicted system response is

$$\nabla\,\hat{y}(\boldsymbol{x}) = \left[\frac{\partial \hat{y}}{\partial x_1}, \ldots, \frac{\partial \hat{y}}{\partial x_m}\right]^T \tag{30}$$

where the subscripts denote differentiation of the scalar system response with respect to the input parameters. It is given as

$$\nabla\,\hat{y}(\boldsymbol{x}) = A_{\boldsymbol{f}}^T\,\boldsymbol{\beta} + A_{\boldsymbol{r}}^T\left(R^{-1}(\boldsymbol{y}_s - F\boldsymbol{\beta})\right), \tag{31}$$

with $A_{\boldsymbol{f}}$ and $A_{\boldsymbol{r}}$ expressing the analytically evaluated Jacobian matrices of the vector of basis functions $\boldsymbol{f}$ and the vector of correlations $\boldsymbol{r}$ with respect to the unsampled location $\boldsymbol{x}$.

5.2 Model validation

In Ref. [51] the authors expressed the need to validate the kriging prediction, and therefore introduced a number of diagnostic tests. Apparently, the best way to validate the model is to compare the prediction with true response surfaces. This however is not feasible in general since if it was possible to evaluate the true response surface then there is no need in evaluating the approximation. Instead the true response could be evaluated at a few additional locations and compared to the kriging prediction. Since the latter test requires additional runs of the expensive full order model, the technique of leaving–one–out cross validation was discussed. Here, one design sample at a time is left out for evaluating the approximation

model denoted by $\hat{\boldsymbol{y}}_{-i}(\boldsymbol{x})$ and the prediction of the model based on $n-1$ samples is compared to the true response of the left–out sample. This is done n times. Additionally, the so–called standardized cross–validated residual was given as a second test and is written as,

$$cvr(\boldsymbol{s}_i) = \frac{y(\boldsymbol{s}_i) - \hat{y}_{-i}(\boldsymbol{s}_i)}{\varphi_{-i}(\boldsymbol{s}_i)} \qquad 1 \leq i \leq n. \tag{32}$$

A valid model should have the cross–validated residual in the interval from about -3 to 3 expressing a confidence interval of 99.7%. These two diagnostic test are visualised in Fig. 9 for the kriging evaluations of the Schur correction elements for the FPv and the RANS flow models as shown in Fig. 10.

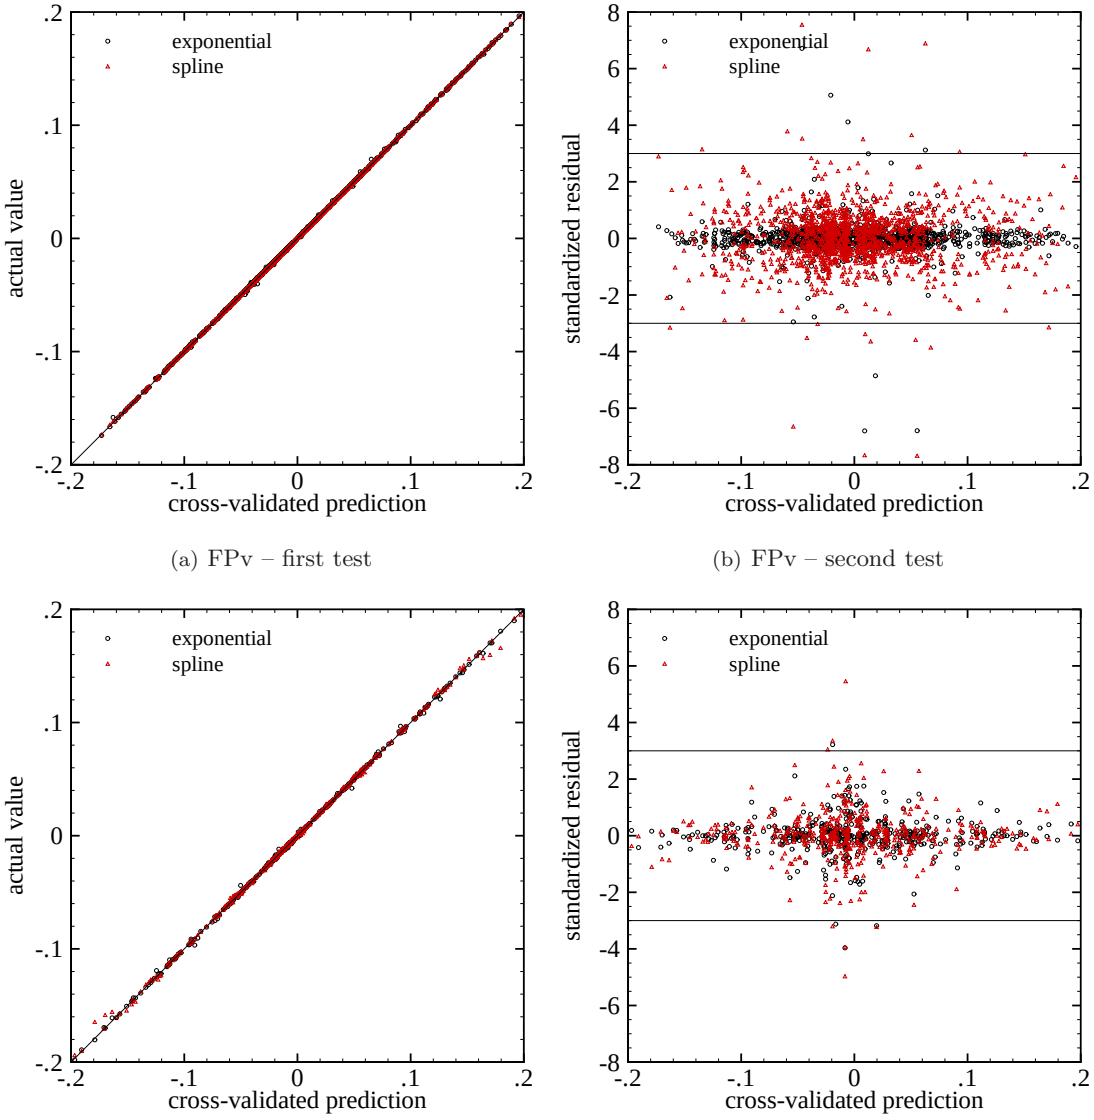

(a) FPv – first test

(b) FPv – second test

(c) RANS – first test

(d) RANS – second test

Figure 9. Two diagnostic tests for prediction model of Schur correction elements showing FPv and RANS flow models; first test showing actual values vs cross–validated prediction and second test showing standardized cross validated residual vs cross–validated prediction.

Figure 10 shows the samples extracted from the full order solver and the corresponding kriging evaluation of the element S^c_{22} of the Schur correction matrix for a range of Mach numbers and frequencies including real and imaginary parts. The considered structural model belongs to the NACA 0012 aerofoil

aeroelastic configuration as shown in Table 5. Four aerodynamic modelling levels are presented including FP, FPv, Euler and RANS. The samples were extracted using either one of the three discussed approaches of linear frequency domain, linearized time domain, or nonlinear time domain. For the RANS nonlinear time domain extraction the excitation amplitude was chosen to 7.5×10^{-3} following the results presented in Section 4.2. The Euler results compare all three ways and include samples created in the Section 4.2. Also, the Euler interpolation results show an exponential correlation function whereas the responses of the other three flow models are correlated by the spline function. This is simply done to illustrate the two correlation functions. For the relatively large number of samples presented in the figure either correlation model gives good results as can be seen in the figure with the diagnostic tests. It was found however that the exponential model proved to be more robust for the majority of cases considered at the expense of losing smoothness compared to the spline model.

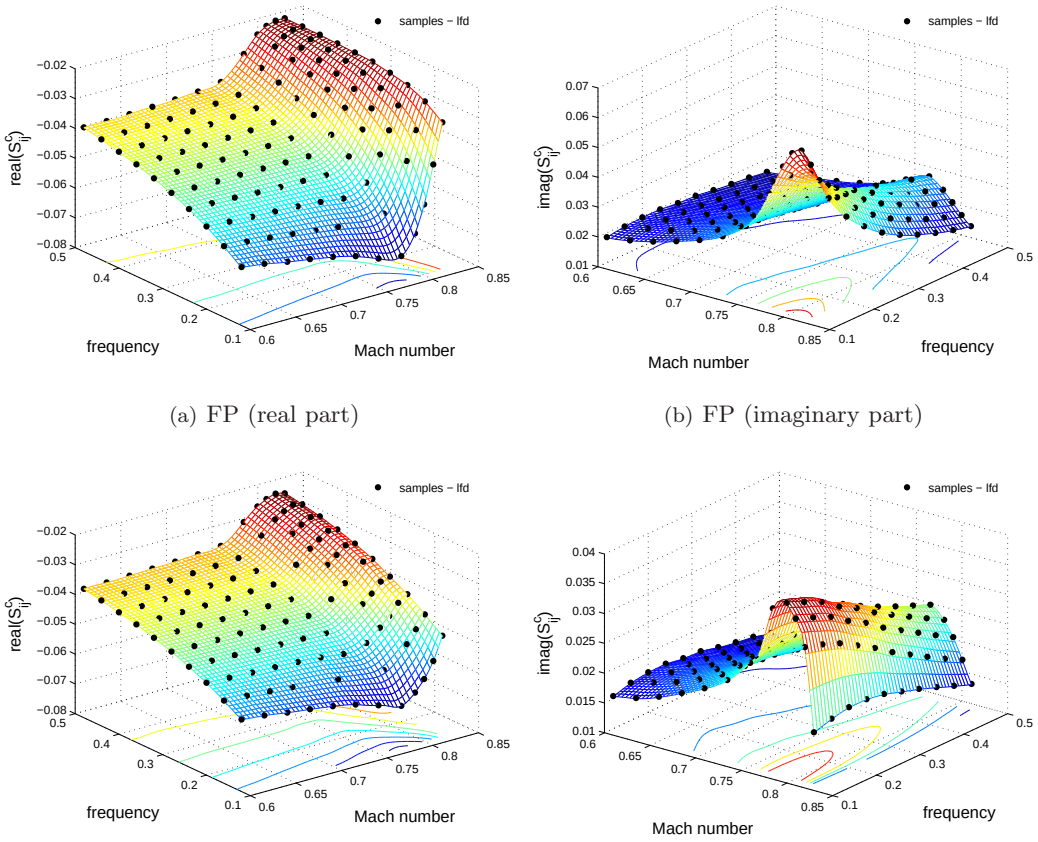

(a) FP (real part)

(b) FP (imaginary part)

(c) FPv (real part)

(d) FPv (imaginary part)

Figure 10. Extracted and interpolated element S_{22}^c of Schur correction matrix including real and imaginary parts and using four levels of aerodynamic modelling; lfd – linear frequency domain, ltd – linear time domain, ntd – nonlinear time domain.

The diagnostic tests shown in Fig. 9 include all 12 nonzero complex–valued elements of the matrix using a zeroth–order regression model with exponential and spline correlation functions. The first test shows very good agreement between the prediction and the true value as can be seen by the close scatter around the 45 degree line. Especially for the FPv results this should not be too surprising since the prediction is based on a very dense sampling. The less densely sampled RANS results build confidence

in the approximation model. The second test also shows good agreement for both the spline and the exponential correlation model. A few values are scattered outside the 99.7% interval for the second diagnostic test. In Ref. [51] the target function is transformed by a simple logarithmic or inverse scaling to improve the kriging model. However, this should not be attempted at this point since the vast majority of test points lie well within the 99.7% confidence interval and in the following the constructed approximation models applied to the aeroelastic stability analysis give excellent agreement to the full order results.

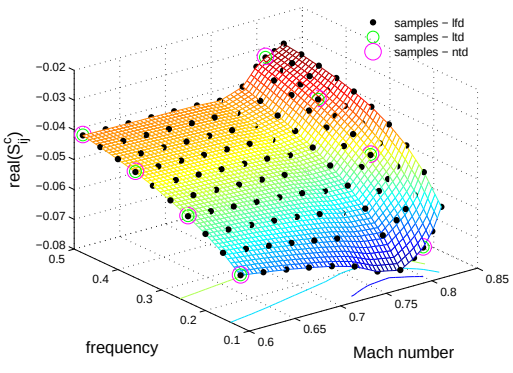

(e) Euler (real part)

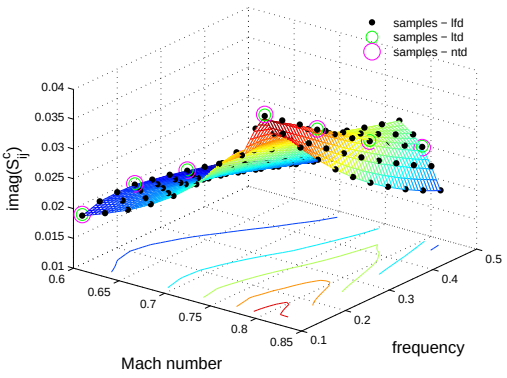

(f) Euler (imaginary part)

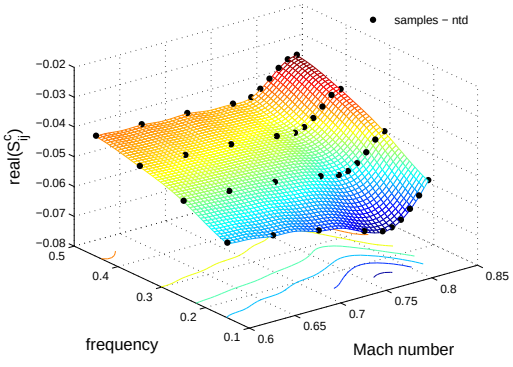

(g) RANS (real part)

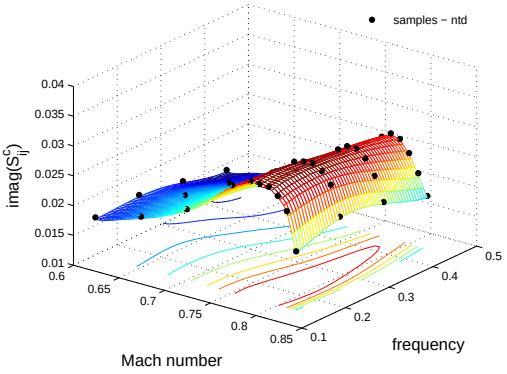

(h) RANS (imaginary part)

Figure 10. (con't)

5.3 Application to aeroelastic stability analysis

As discussed in Section 4.1, evaluating the Schur correction matrix S^c based on the full order model (FOM) of the flow solver accounts for the highest costs in the stability analysis. The kriging interpolation technique makes it possible to form a reduced order model (ROM) of the Schur correction matrix based on full order samples. Once the kriging predictor is evaluated the stability problem can be solved without relying on the FOM solver. To solve the ROM for the prediction of instabilities the same approach (using Newton's method) as outlined in Section 4.1 is applied with the Schur correction matrix readily available through the kriging predictor. The critical eigenvalue with zero damping is detected using a bisection method applied to the bifurcation parameter at fixed Mach number. There is another convenient aspect of the kriging method as discussed in Section 5.1. The derivative of the correction

term in the Schur complement S^c_λ required in forming the Schur Jacobian matrix for Newton's method is evaluated analytically from the chosen regression model and the correlation function as part of the kriging prediction as shown in Eq. (31). This avoids using finite differences to obtain the Schur Jacobian matrix.

In the current formulation the ROM of the correction term is based on a purely imaginary eigenvalue with zero damping $\lambda = i\,\omega$ (which allows extracting the Schur correction terms from forced time marching simulations at the specified frequency ω), whereas the structural part uses the complete eigenvalue including nonzero real part. Thus, the approximate Schur complement used for the stability analysis is written as

$$S = S^s(\lambda, \bar{U}) + \hat{S}^c(\omega, M) \tag{33}$$

where $\hat{S}^c$ is the kriging predicted correction term of S^c. The dependence of the approximation model on the Mach number is given by M. Interestingly, for the two degree–of–freedom aerofoil discussed in the present study the correction term S^c is independent of the bifurcation parameter which is the reduced velocity $\bar{U}$. For more general three–dimensional cases the kriging model of the correction term would then be built for instance at varying values of altitude as additional model dimension. At the critical eigenvalue $\lambda_F = i\,\omega_F$ the ROM is exact within the limits of the interpolation algorithm. The implications of the approximation $\hat{S}^c(\omega, M)$ are presented in Fig. 11.

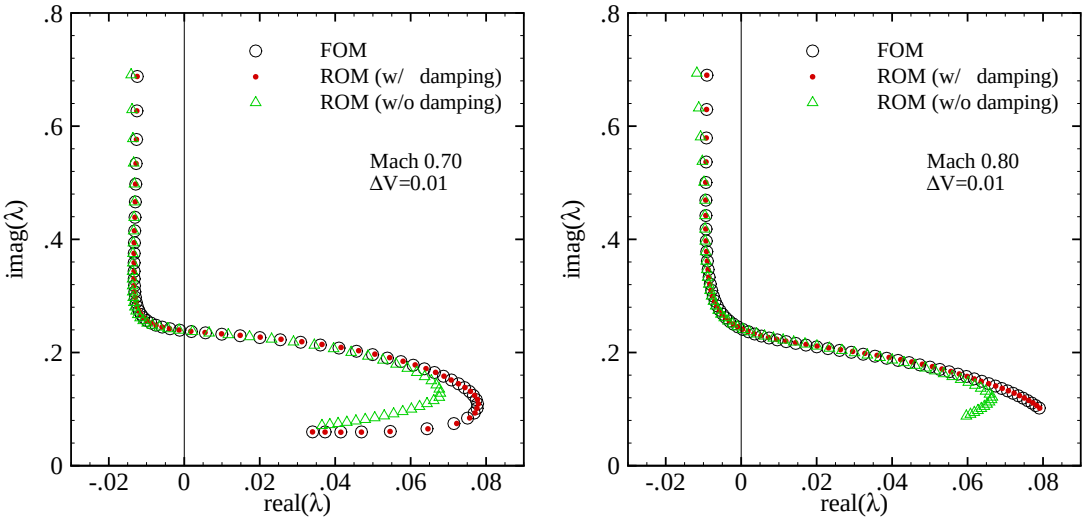

(a) Mode tracing at Mach 0.7

(b) Mode tracing at Mach 0.8

Figure 11. Mode tracing of least stable mode for NACA 0012 aerofoil configuration using Euler flow model and showing a comparison between full order model (FOM) and the approximate reduced order models (ROM) for the Schur correction term with damping terms $\hat{S}^c(\lambda)$ and without damping terms $\hat{S}^c(\omega)$ included in the design samples.

The figure shows the tracing of the least stable aeroelastic mode at two freestream Mach numbers for the NACA 0012 aerofoil configuration using the Euler flow model. The structural parameters for this configuration are described in Table 5. The increment in the reduced velocity $\bar{U}$ is 0.1 (giving $\Delta V = 0.01$) starting from 1.0. The calculation of about 60 points on the root locus with increasing bifurcation parameter takes less than a second of CPU time with the ROM, whereas the same number of points with the FOM using a grid with 15k control volumes takes more than an hour of CPU time.

Two simulations using different ROMs are presented in the figure. One uses full order samples with zero damping and varying frequency to build the kriging predictor and is denoted as ROM (w/o damping), whereas the second ROM was constructed with samples at fixed Mach numbers for both varying frequency and damping. All samples were extracted using the linear frequency domain approach.

It can be seen that as a first approximation the relevant eigenvalue with the smallest value of damping can be traced quite precisely even away from the imaginary axis without including damping in the design matrix of the predictor. In the discussed case the error for the tracing introduced by the approximate Schur correction term $\hat{S}^c(\omega)$ is very small in the relevant region close to the imaginary axis. This basically suggests that the variation of the Schur correction elements with the real part of the eigenvalue is small or at least that the influence of this variation on the eigenvalue problem is small. This observation is confirmed in Fig. 12 showing the element S^c_{22} as a function of damping and frequency. In the figure it was necessary to remove samples at low damping and low frequency since these adversely affected the construction of the interpolation model. The discussed samples interferred with fluid eigenvalues and caused an ill–conditioning of the problem in Eq. (5). The tracing of the ROM based on the cheap model illustrated in Fig. 12 follows precisely the results of the FOM. Also, it is found for the ROM based on zero damping samples that outside the coverage with full order samples which is between frequencies of 0.12 and 0.5 the error of the tracing gets worse compared to the FOM. The predictor based on damped samples covers a larger frequency region.

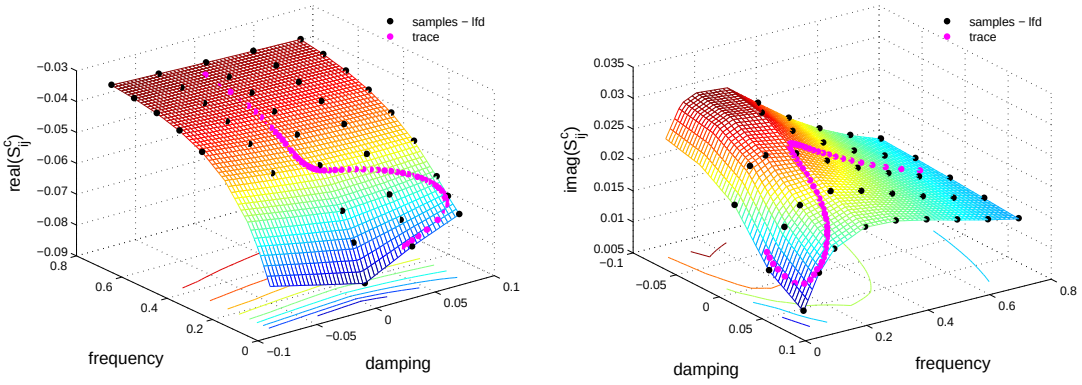

(a) real part

(b) imaginary part

Figure 12. Extracted and interpolated element $S^c_{22}(\lambda)$ of Schur correction matrix including trace of eigenvalue for Euler flow model at Mach 0.7.

Figures 13 and 14 present more results of the NACA 0012 aerofoil aeroelastic configuration. Results for a range of sub- and transonic Mach numbers are given for all four considered aerodynamic modelling levels. Figure 14 shows one element of the Schur correction matrix obtained by kriging interpolation from the full order samples together with the trace of the instability points in terms of combinations of Mach number and critical frequency. In Fig. 13 a comparison between the FOMs and the ROMs is shown for the instability boundary presented as critical flutter speed index $V_F = \bar{U}_F/\sqrt{\mu_s}$ and the instability frequency ω_F. More than 100 instability points were calculated with the ROM for each flow modelling level. The agreement between ROMs and FOMs is excellent as should be expected since the resolution of the full order samples is high as can be seen in Figs. 10 and 14. Also, time–accurate simulations to

confirm the RANS predictions are included for two Mach numbers with the plus sign (tilde) indicating a stable (unstable) response due to an initial disturbance. The somewhat constant offset between Euler and RANS results suggest that the boundary layer as predicted by latter modelling level has a stabilizing effect on the configuration. Furthermore, it seems that the shock dynamics, which are correctly predicted by the Euler model, act as the dominant mechanism for the aeroelastic instability compared to viscous effects (in this case and at the shown Mach number range). Indeed, comparing to the flow solutions it can be seen that shallow separation due to shock wave impingement is first encountered at about Mach 0.81 to Mach 0.82. The effect of separation on the stability prediction will be discussed in a later section.

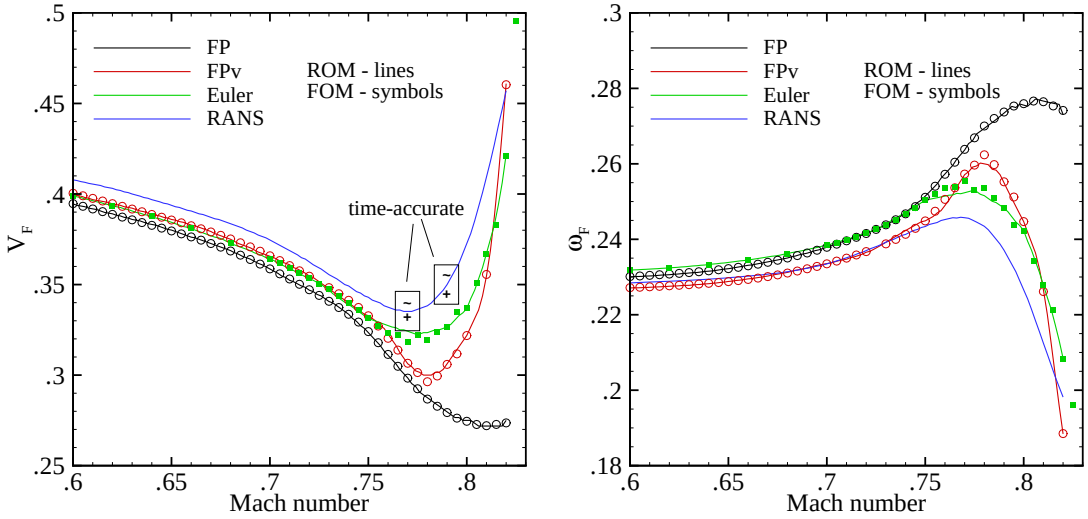

(a) Instability boundary

(b) Instability frequency

Figure 13. Instability boundaries and frequencies for four levels of aerodynamic modelling showing comparison between full order model (FOM) and reduced order models (ROM).

The differences between high–level and low–level results require further attention. First, even for lower (purely subsonic) Mach numbers there is an offset between both the inviscid (Euler and FP) and viscous (RANS and FPv) critical flutter speed indices, while the instability frequencies agree better. The reason for the disagreement in the flutter speed index is not found to be the grid resolution since the results present grid–converged solutions meaning that inspected finer grids (results of which are not presented) did not change the results notably. However, the influence of the grid topology (structured vs unstructured) cannot be estimated. Another possible factor could be in the distinct applied spatial discretisation schemes including different treatment of boundary conditions for the high–level and the low–level solver. For instance, the low–level solver uses a transpiration boundary condition with fixed geometry on solid surface whereas the high–level solver explicitly deflects the geometry. On the other hand, the discrepancy is about one percent and should not be too much of a concern. Secondly, in the transonic dip the violation of the isentropic and irrotational assumptions of the FP solver becomes evident. The formed shock wave is overpredicted by the FP method compared to the Euler prediction resulting in an increased difference in the instability boundary starting from about Mach 0.75. Interestingly, despite estimating the shock location quite accurately, the FPv instability prediction also deviates considerably in the transonic dip region. However, the other side of the dip with the strong rise in the flutter speed

index is simulated correctly compared to Euler and RANS. These points will need more efforts to be understood, and including a shock–correction model as an additional submodel in the lower–level solver may give answers.

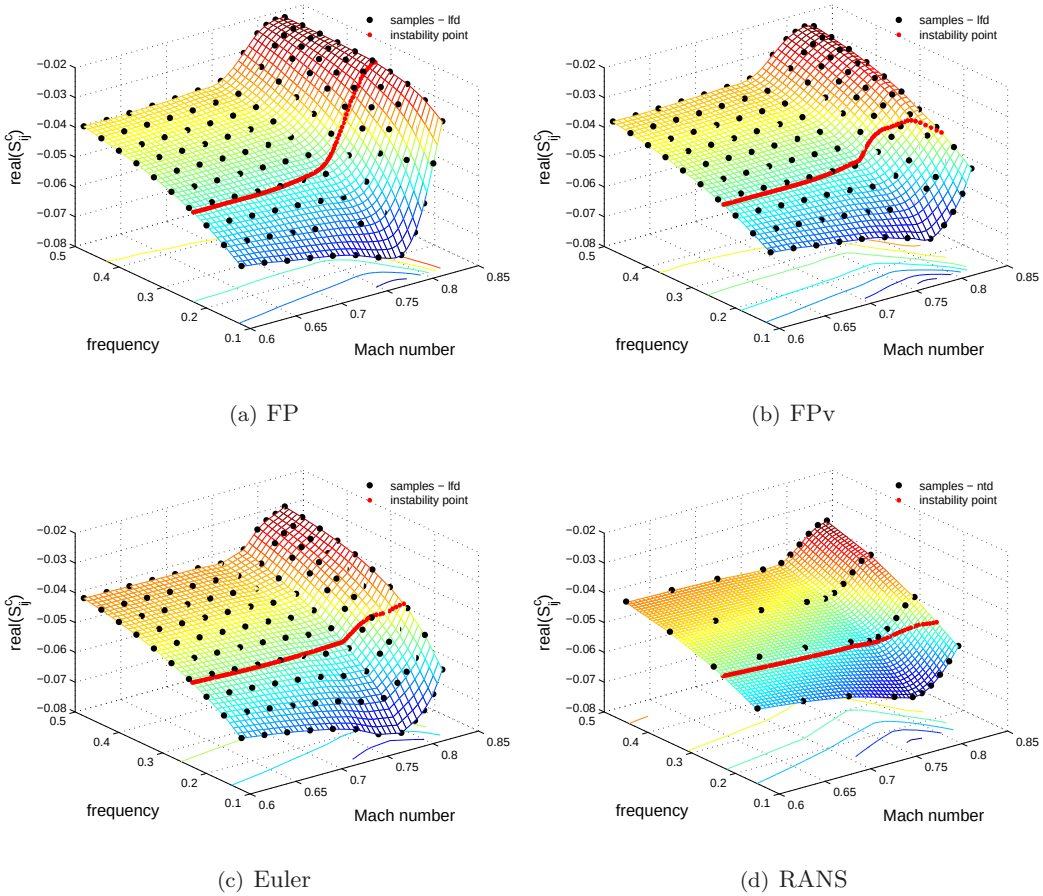

(a) FP (b) FPv

(c) Euler (d) RANS

Figure 14. Extracted and interpolated element real(S^c_{22}) of Schur correction matrix including trace of instability boundary; lfd – linear frequency domain, ntd – nonlinear time domain.

As a first attempt to explain the differences, the different boundary conditions are considered. Since a FP formulation with moving grids is currently not available, a transpiration boundary condition was included in the Euler formulation. Interestingly, for the stability analysis as presented herein the transpiration boundary condition only affects the Jacobian matrix block A_{fs} in Eq. (3) and implementing the boundary condition is easily done. Therefore, the computational mesh is rigidly deflected to set the values in the halo cells used to enforce a tangential flow condition at the solid surface. Then, the original grid is restored while keeping the updated halo values to evaluate the fluxes. The results for the stability analysis with the transpiration boundary condition are shown in Fig. 15. Indeed, it can be found that the modified boundary condition has a slight influence on the aeroelastic stability. Compared to the results of the original boundary condition the critical flutter speed index is reduced and approaches the FP results. The differences due to the boundary conditions can be found in the structure of matrix A_{fs}. Moving the entire computational domain as done for the original boundary condition creates a dense matrix A_{fs}, i.e. the structural response is felt in the entire flow domain, whereas the transpiration boundary conditions fills the matrix sparsely for control volumes close to the surface and affected by the

halo information. The critical flutter speed index of the FP formulation is closely followed until about Mach 0.75. As discussed before at this point the underlying modelling assumptions are violated by the forming strong shock wave. Still, the difference in the solid wall boundary condition does not explain the behaviour in the transonic dip since the FPv formulation predicts the shock location accurately.

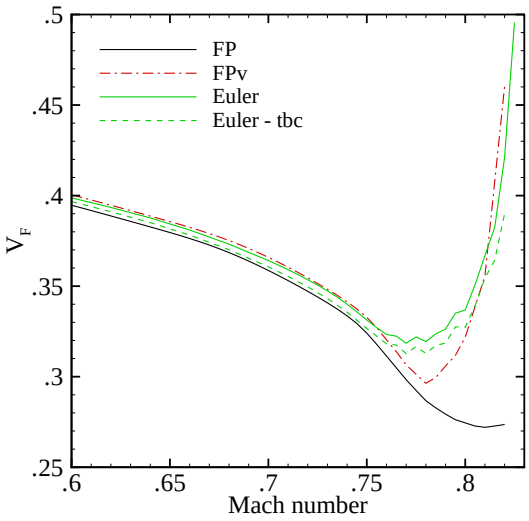

Figure 15. Influence of transpiration boundary condition on aeroelastic stability prediction for Euler flow modelling; tbc – transpiration boundary condition.

Using the FOM, the prediction of one instability point takes more than 10 minutes of CPU time for the Euler model with 15k control volumes not including the time to obtain the steady state solution, while the FP model on a grid with 5k control volumes requires about 15 seconds of CPU time depending on the number of iteration required to achieve the convergence criterium on the real part of the eigenvalue. On the other hand, the evaluation of 100 instability points using the ROM takes about 10 seconds of CPU time independent of the size or fidelity of the original full order model. The difference in CPU time between the Euler and the FP solver, despite having such an efficient linear solver for the high–level model as shown in Tables 2 and 3, comes from the way of solving the involved linear system for the fluid Jacobian matrix of the form $\tilde{A}\boldsymbol{y} = A_{fs}$ for each column of the right–hand side matrix to find the Schur residual and the Schur Jacobian matrix. Here, $\tilde{A}$ represents the matrix $(A_{ff} - \lambda I)$. The lower–level code uses a direct solver which as a first step factorizes the coefficient matrix as $\tilde{A} = LU$. The factorization done for each new value of the eigenvalue covers most of the costs. Then, the second step is to solve the linear system by performing forward and back substitution on each of the right–hand sides using the factors L and U which is cheap compared to the factorization. The higher–level solver uses a tailored linear solver which does not offer such an initial factorization but solves the complete linear system for each right–hand side. Thus for the considered cases with about 10k control volumes, the direct solver performs better especially with increasing number of modes (right–hand sides). However, the costs to evaluate the factorization directly scale badly with larger dimensional systems. Consider Table 6 showing the CPU time for one linear solve (LSOLV) on the fluid system (FS) and the eigenvalue problem (EV) to achieve the required tolerance (LTOL). The CPU time for the eigenvalue part in FP model includes the LU–factorization and solving for one right–hand side, whereas for the Euler model the CPU times

Model	CV	FS–LSOLV	FS–LTOL	EV–LSOLV	EV–LTOL
FP	5k	0.3	/	$0.5 + 0.03$	/
	15k	1.6	/	$2.8 + 0.13$	/
	33k	4.5	/	$8.0 + 0.30$	/
Euler	15k	0.08	10^{-2}	$0.6 + 5.00$	10^{-9}

Table 6. Performance of direct and linear solver for one linear solve.

for the preconditioner and one linear solve are given. The FP grid with 33k control volumes creates a linear system of about the size of the Euler grid with 15k control volumes since there are 4 unknowns in the two–dimensional Euler equations compared to 2 unknowns in the FP model.

Remarks

Strictly speaking, one contribution is missing in the reduced order models compared to the full order versions. The structural part of the Schur matrix S^s contains terms of the form $\partial C_l/\partial\alpha$ and $\partial C_m/\partial\alpha$, i.e. the dependence of the structural residual on the structural unknowns through the integrated aerodynamic forces. However, it was found that these contributions are very small for the considered cases and hardly affect the Schur complement elements, and thus not justifying the efforts of inclusion at the present point. This can be seen in the agreement between ROM and FOM results as presented herein. The uncertainty due to the kriging interpolation is considered to be far more significant.

As with any ROM, the costs to construct the ROM have to pay off. The model has to be used in repeated simulations to establish a picture about aeroelastic stability for the space of parameters that are covered by the ROM. There are however already two positive aspect justifying the costs in building the ROM as it stands. One is the access to the RANS modelling level (and flow modelling of any higher fidelity). Full blown RANS aeroelastic stability analysis over a large range of Mach numbers (i.e. avoiding time–accurate simulations) has not be done within the Schur framework before. The second aspect is the observation that evaluating the Schur correction matrix at discrete Mach numbers, frequencies (and damping values) to build a ROM used in the stability analysis is computationally more efficient than directly using the FOM since iterating on the full order system to converge the solution is avoided. On the other hand forming the Schur correction matrix does not require iterations at a given parameter combination while iterating on the ROM is very cheap. Another aspect will be in addressing aerodynamic uncertainty in aeroelastic prediction.

6 Reducing costs to construct reduced order model

The costs to create the reduced order model, i.e. the number of design samples to adequately represent physics, is a limiting factor in the analysis especially when an expensive high–fidelity flow model is used. There are two main aspects to this which are discussed in this context. First, the large number of full order samples as used for the kriging predictor in Fig. 10 does not seem to be necessary, and the number actually required to accurately build the response surface can be reduced by an appropriate sampling technique. Secondly, a cheap low–fidelity model can be exploited to get valuable information about the aeroelastic stability prediction using an expensive high–fidelity model.

6.1 Sampling techniques

Sampling methods to generate valuable information used for the construction of an interpolation (surrogate) model can be divided into a priori and a posteriori methods. A priori methods initially generate samples without information about the functional behaviour of the considered problem hereby simply relying on the dimensions of the input parameter space. Using intermediate information of the approximation model such as a measure of the maximum error allows placing samples in these critical regions. Thus, the samples are generated sequentially based on the current realisation of the approximation model. This is referred to as a posteriori.

A first sampling technique has already been presented in Fig. 10. Here, the so–called (rectangular) grid sampling has been applied. It should be referred to as brute–force sampling due to the (unnecessary) large number of design sites. Latin hypercube (LH) sampling can be considered as an improved version of random (Monte Carlo) sampling [54]. While random sampling creates parameter combinations independently (and possibly without providing additional information), LH sampling ensures that all parts of the parameter space are represented. Therefore, each dimension of the parameter space is divided into a specified number of non–overlapping bins of equal probability. One sample per dimension is randomly chosen from each bin and then randomly combined with the samples from the other parameter dimensions. LH sampling is presented in Fig. 16 for the FP flow model showing the element real(S^c_{22}) of the Schur correction matrix for combinations of Mach number and frequency. Kriging predictions based on 8 and 32 design points are presented.

Initially, four samples have been placed at the corners of the parameter space to avoid extrapolation. The dimension of the parameter space is defined to provide a good range for initial blind search with the Mach number limits covering the region of interest and the frequency limits based on typical flutter frequencies. The remaining design points have been generated by LH sampling. Looking at the figures it can said that even a few samples can approximate the target reasonably precisely. However, more LH samples are required to guarantee that the response surface is constructed correctly. Despite having some local extrema around the domain boundaries of the parameter space shown in Fig. 16(c) the region of immediate interest for the aeroelastic stability analysis is located away from the boundaries, and then the LH samples provide a good basis for the kriging predictor. The latter point is illustrated in Fig. 17(a) showing the critical flutter speed index (as the true measure for the quality of the approximation) for all four kriging models based on LH sampling as given in Fig. 16 and compares to the full order reference

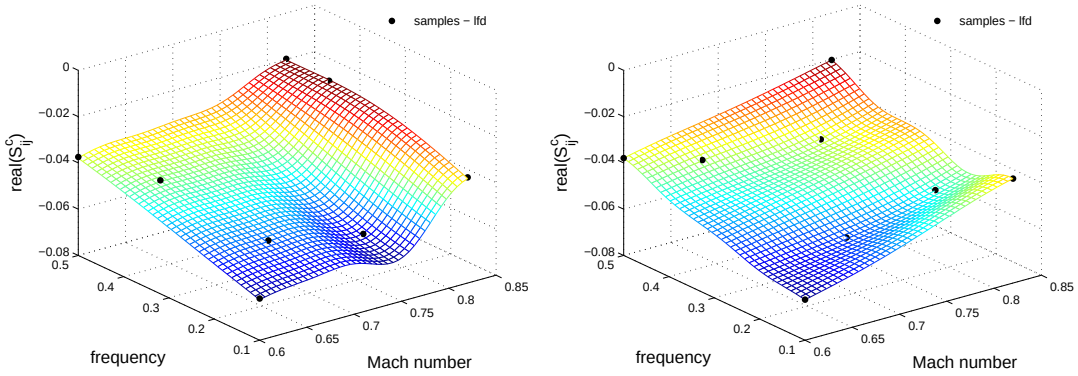

(a) 8 LH samples

(b) 8 LH samples

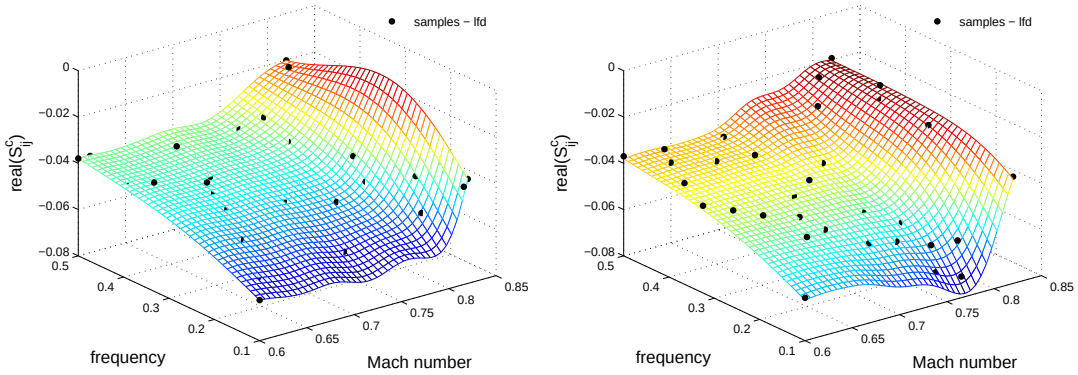

(c) 32 LH samples

(d) 32 LH samples

Figure 16. Element real(S_{22}^c) of Schur correction matrix for NACA 0012 aerofoil configuration using latin hypercube (LH) sampling and FP flow modelling.

solution of the FP flow model. The figure shows that the approximation models based on only 8 samples are not good enough compared to the reference solution, particularly the model denoted by (a) and corresponding to Fig. 16(a). The kriging predictions based on 32 samples show good agreement until about Mach 0.8 where model (c) deviates considerably.

Grid sampling and LH sampling are examples of a priori techniques. A first a posteriori sampling approach would naturally use a standard measure of the kriging predictor error provided in the framework. The location in the parameter space where this measure of error is maximized would then be the logical choice for the next sample. The technique herein referred to as MSE sampling (for mean squared error) uses the expression in Eq. (21). It is shown in Fig. 18. Initially a number of LH samples is created to allow evaluating the kriging model and to provide a somewhat filled parameter space. In the figure it can be seen that the response surface is well predicted showing less irregularities compared to the LH sampling with the same number of samples. It should be remarked that MSE sampling is a space–filling approach just as LH sampling except that the parameter space is filled a posteriori according to the error of the kriging model. Since the kriging error depends on the chosen correlation model with the correlation being weighted by the distance between samples, the new sample location is likely to be found near the point maximizing the distance to all surrounding samples while also including the degree of correlation between the samples. Compare the expression in Eq. (21). Thus, MSE sampling is an

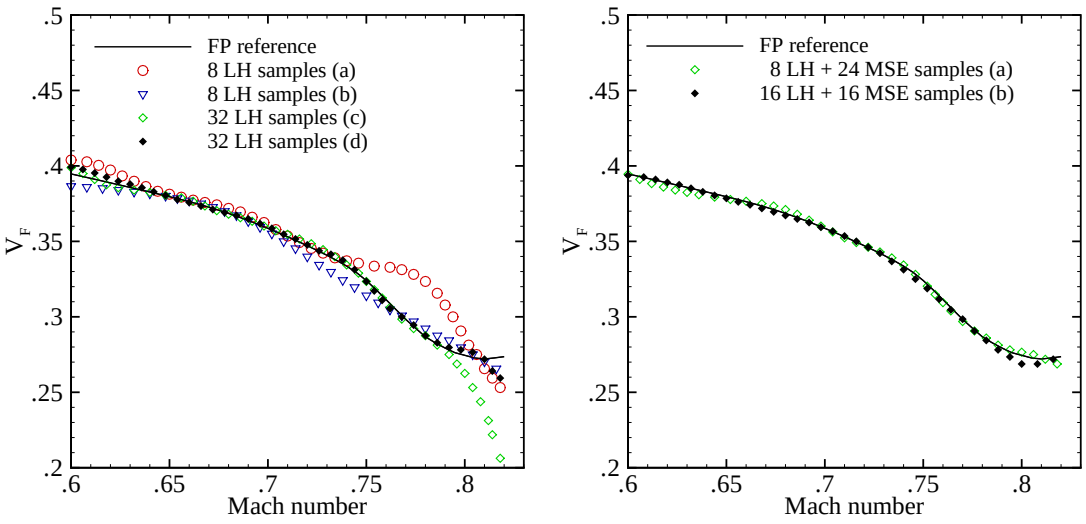

(a) LH sampling (b) MSE sampling

Figure 17. Critical flutter speed index shown for different sampling techniques using the FP flow model and compared to the full order reference solution.

improved space–filling algorithm. The predicted critical flutter speed index important for judging the quality of the kriging model is presented in Fig. 17(b). The two shown sample resolutions give a very good agreement to the reference solution. The difference between the two approximations lies in the number of initial LH samples.

The approach of cross–validation as discussed in Section 5.2 is another useful measure of error to locate samples for the kriging model. Leaving–one–out cross validation evaluates the kriging predictor as many times as there are samples while one sample is left out each time. The differences between either the full and the cross–validated predictions or the associated errors are analysed to give a measure relating the sample sensitivity and the standard kriging error [55]. However, it was found that sampling based on cross–validation did not prove to be superior to standard MSE sampling in the current application. Therefore, it should not be considered further at this point.

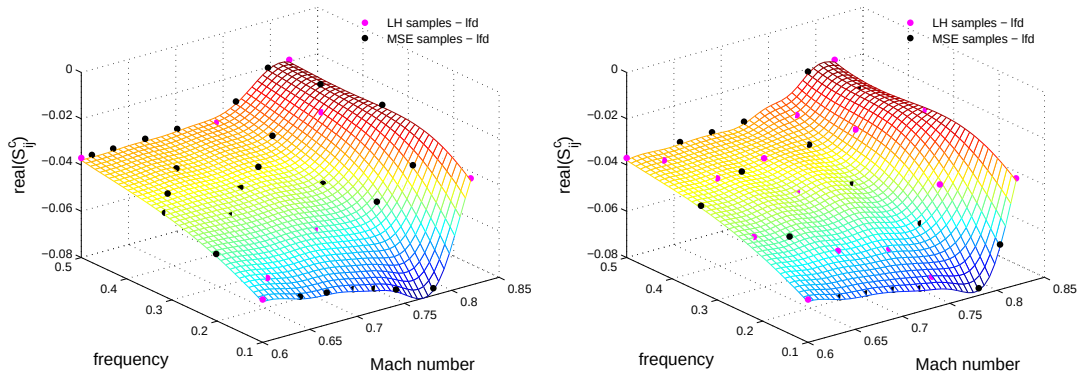

(a) 32 samples (b) 32 samples

Figure 18. Element real(S^c_{22}) of Schur correction matrix for NACA 0012 aerofoil configuration using sampling based on maximum of mean squared error (MSE) in the kriging formulation with different numbers of initial latin hypercube (LH) samples.

Inspecting the figures with the elements of the Schur correction matrix, it can be said that changes in the freestream Mach number have a more significant influence on the shape of the elements than changes in the eigenvalue (frequency). The dependence of the Schur elements on the damping part of the eigenvalue was found to be almost linear in the vicinity of the imaginary axis, compare to Figs. 11 and 12. Therefore, more samples should be placed in the dimension of the Mach number, less in the dimension of the eigenvalue. Another observation is that in the subsonic regime the response of the Schur correction elements can be represented by simple polynomials, and therefore only a few points are required in both dimensions. Once the steady state flow solution exhibits shock waves sampling in Mach number should be refined. Lets refer to sampling based on these observations as directed grid sampling; directed in a sense that general information of the dependencies on the steady state flow solution is included in choosing sample locations. The relatively little number of samples required for the eigenvalue dimension is convenient for exciting several frequencies simultaneously in the time domain approach as discussed in Section 4.2.

The approach of directed grid sampling is presented in Fig. 19 for the FPv flow modelling using the NACA 0012 aerofoil configuration, and in Fig. 20 for a refined search space since the initial search space showed that a third of the samples are basically redundant. The corresponding instability boundaries showing critical flutter speed index and critical frequency are given in Fig. 20. It is evident that the refined search space provides an even better prediction compared to the results for the initial parameter space. Both sample resolutions do not precisely predict the bottom of the transonic dip compared to the full order reference solution. Here, the changes with Mach number would require more samples. Nevertheless it is a good conservative prediction of either reduced model.

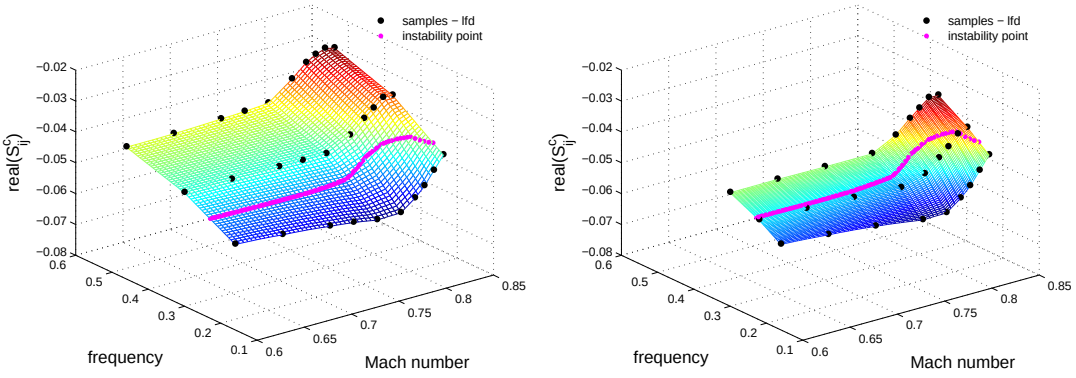

(a) Directed grid sampling

(b) Refined directed grid sampling

Figure 19. Element real(S_{22}^c) of Schur correction matrix for NACA 0012 aerofoil configuration using directed grid sampling for the FPv flow model with 30 samples.

6.2 Using lower fidelity information to establish higher fidelity model

The construction of a high fidelity reduced order model could be assisted by information provided by a lower fidelity model. A lower fidelity model could be established by either using a lower aerodynamic modelling level or by using the higher fidelity model on a coarser computational mesh. In both cases the lower fidelity model should capture the relevant physics qualitatively correct (to provide the trend/limits

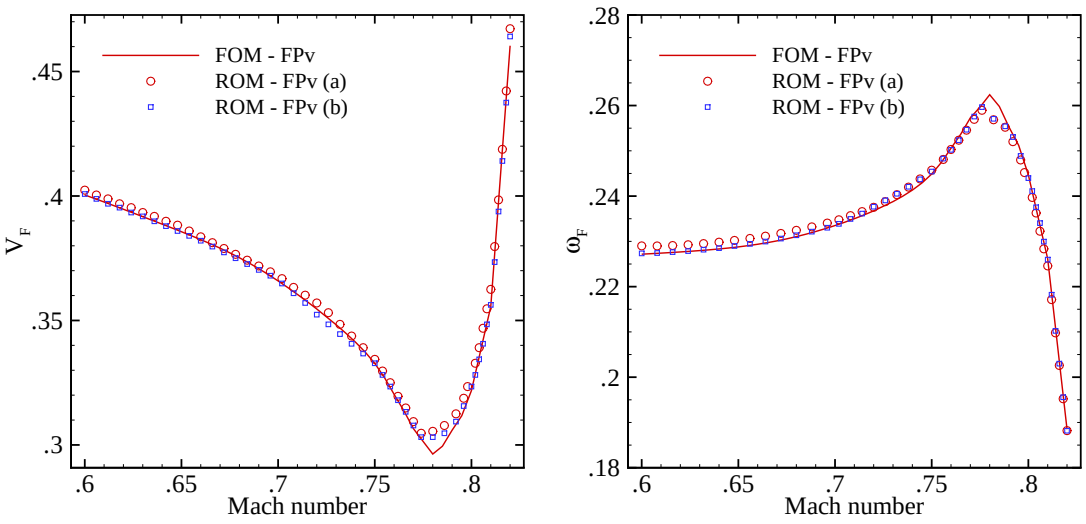

(a) Initial search space

(b) Refined search space

Figure 20. Instability boundary of NACA 0012 aerofoil configuration using reduced order model based on directed grid sampling and showing critical flutter speed index and frequency for FPv flow modelling.

of the prediction) while the higher fidelity model would then refine the prediction at reasonably little additional costs. For instance the FPv formulation could be used to assist in predicting the effects of separation on aeroelastic stability. Of course, the fundamental limitations of both the FP formulation and the integral boundary layer formulation would hemper an accurate prediction compared to RANS. FPv would provide a general trend though and helps in choosing correctly a Mach number and frequency interval (or more generally a critical parameter range) for the RANS search which then simulates the relevant physics correctly.

A first example is shown in Fig. 19 with the directed grid sampling. Imagine not much is known about the aeroelastic response of a configuration. Then, the response surface is constructed using the cheaper FPv formulation with a large number of samples to cover a larger parameter range. Doing the stability analysis on the kriging–based ROM for the FPv flow modelling reduces the parameter domain of relevance for the sample locations of the higher–fidelity RANS model saving costs instead of searching a large domain with an expensive model. (Note that Fig. 19(b) shows a refined search using FPv.)

Taking a step further, the shape of the response surfaces of the lower–level flow models clearly shows similar features compared to the higher–fidelity model as can be seen in Fig. 10. This is not surprising since the same physics are considered. To speak in terms of the statistical approach used for the kriging prediction; the low–fidelity response is spatially correlated with the high–fidelity response. This can be exploited using the co–kriging modelling technique. There is some confusion in the literature across the different disciplines and subjects about what techniques should be referred to as co–kriging. In this work all kriging–like techniques that use additional information on the functional behaviour of the response rather than just the response itself to establish the cheap kriging predictor shall be referred to as co–kriging. Three distinct approaches have been found in the literature related to aeronautical sciences and applied to the prediction of aerodynamic loads and in generating aerodynamic look–up tables for conceptual aircraft design. For convenience, only a single scalar response for a multidimensional input

is considered in the following, compare Section 5.1.

First, the gradient–enhanced kriging (GEK) approach includes both the function and its gradient in the model construction,

$$\boldsymbol{y}_s = \left[y(\boldsymbol{s}_1), \ldots, y(\boldsymbol{s}_n), \frac{\partial y(\boldsymbol{s}_1)}{\partial \boldsymbol{s}_1}, \ldots, \frac{\partial y(\boldsymbol{s}_n)}{\partial \boldsymbol{s}_n} \right]^T, \tag{34}$$

where the scalar response y is differentiated with respect to all input dimensions of $\boldsymbol{s}$. Since more information is provided in terms of the gradient this approach is meant to be more efficient in higher dimensional problems with fewer sample compared to standard kriging.

A second approach uses a spatially correlated co–variable (or auxiliary variable) to augment the input design parameter space of the primary variable. Generally, the primary variable represents the system response of a more expensive higher–fidelity (HF) model which, consequently, is only sampled at relatively few points in the parameter space, whereas the co–variable is found from a cheaper lower–fidelity (LF) model which can be sampled densely. The cheaper model could be established in this context either by a lower–level aerodynamic modelling approach or by a higher–level model run on a coarse grid as mentioned earlier. Apparently, since the HF and the LF models describe the same phenomenon the system response is generally correlated. Then, the LF model would provide a trend of the system response with the HF data updating the prediction. Thus, the kriging predictor for the primary variable is constructed using

$$\boldsymbol{y}_s = [y(\tilde{\boldsymbol{s}}_1), \ldots, y(\tilde{\boldsymbol{s}}_n)]^T, \tag{35}$$

where $\tilde{\boldsymbol{s}} = \left[\boldsymbol{s}, \hat{y}_{\mathrm{lf}}(\boldsymbol{s})\right]^T$. The co–variable response $\hat{y}_{\mathrm{lf}}$ is required at the location of the HF samples and obtained either by an intentionally placed sample or by its kriging–based prediction.

The last approach uses so–called bridge functions (scaling functions) which construct an unknown function to assist in updating the lower–fidelity model. In this approach essentially two kriging predictors are formed; one for the lower–fidelity model and one for the additional function. Considering an additive bridging formulation, the unknown function η is formed as the difference (error or offset) between the models of variable fidelity,

$$\boldsymbol{\eta}_s = \left[y_{\mathrm{hf}}(\boldsymbol{s}_1) - \hat{y}_{\mathrm{lf}}(\boldsymbol{s}_1), \ldots, y_{\mathrm{hf}}(\boldsymbol{s}_n) - \hat{y}_{\mathrm{lf}}(\boldsymbol{s}_n) \right]^T \tag{36}$$

where the LF data are interpolated at the HF sample locations. Then, the HF response at an unsampled location $\boldsymbol{x}$ is predicted using

$$\hat{y}_{\mathrm{hf}}(\boldsymbol{x}) = \hat{y}_{\mathrm{lf}}(\boldsymbol{x}) + \hat{\eta}(\boldsymbol{x}), \tag{37}$$

where $\hat{\eta}$ is kriging–predicted bridge function. Various variants have been presented including multiplicative and hybrid formulations [56] details of which shall not be given herein.

Theoretically, all flow models give the same response in the sub- and low transonic regime with the exception of viscous effects, i.e. the Euler model should agree to the FP formulation while RANS and FPv prediction are accurate compared to each other. The differences as found in the preceeding sections should be attributed to the applied discretisation schemes as well as the treatment of boundary conditions, compare for instance Section 5.3. Then, one higher–fidelity modelling sample should be sufficient to give the correct level of the response while a densely sampled lower–fidelity approximation model provides the trend in the dimensions of the response surface. Using the co–variable approach with augmenting

the high–fidelity data with the low–fidelity prediction would require a minimum number of high–fidelity design samples depending on the order of the regression model, for instance a linear regression model then requires at least four expensive samples for the three input dimensions (two high–fidelity dimensions plus the augmented dimension). Instead, a standard low–fidelity kriging model can be used with the prediction shifted according to the offset between high– and low–fidelity information. This shift however would be small for most elements of the Schur correction matrix except for the elements that are strongly dependent on the flow model discretisation. Also, the low–fidelity data would indicate where to put the important high–fidelity sample. Then, the samples saved in the subsonic and low transonic regime can be invested in the nonlinear transonic regime. These points however will need further discussion.

7 Isogai case — strong shock/boundary layer interaction

The NACA 0012 aerofoil aeroelastic configuration used for developing the methods presented a case with dominant shock dynamics. Now, a case with strong shock/boundary layer interaction is considered. The Isogai case as introduced in Section 3.2 shows a distinct s–shaped appearance of the instability boundary in the deep transonic regime for inviscid modelling levels. Several authors using either thin–layer Navier–Stokes (TLNS) [42] or coupled Euler/integral boundary layer [44] modelling have demonstrated that viscous effects fill the transonic dip and remove the double–valued critical flutter speed indices.

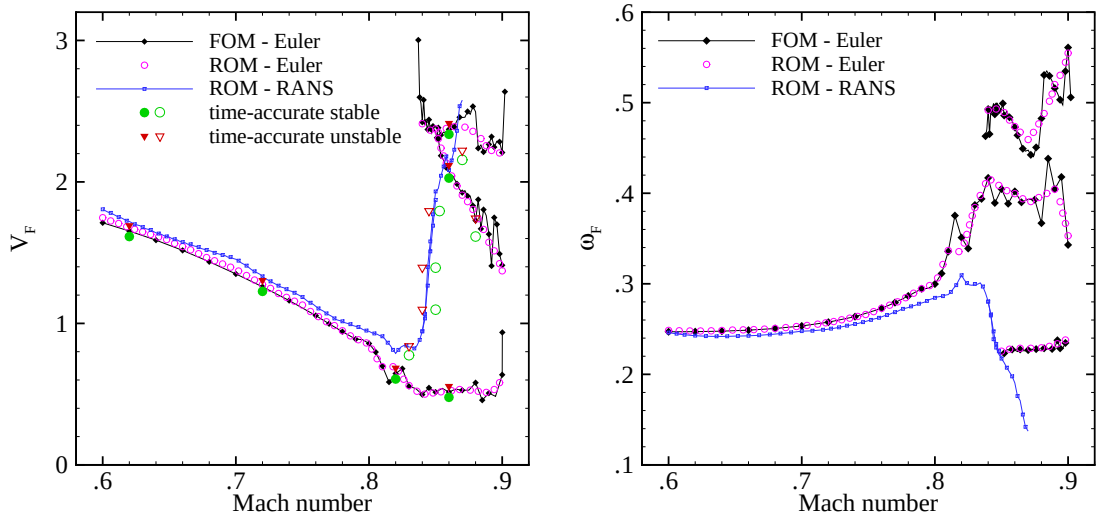

(a) Critical flutter speed index

(b) Instability frequency

Figure 21. Instability boundary of Isogai test case showing critical flutter speed index and critical frequency; flow modelling with Euler and RANS equations.

Figure 21 presents the instability boundary showing critical flutter speed index V_F and critical frequency ω_F. A comparison between Euler predictions from the full order and the reduced order model is discussed where the full order samples for the construction of the ROM were obtained by the linear frequency domain approach. The computational grid used for the Euler simulations with 16k control volumes differs from the grid applied for the discussion in Section 3.2. The RANS predictions are based on a ROM constructed from forced nonlinear time domain simulations and illustrate the effect of strong shock/boundary layer interaction on aeroelastic stability. Time–accurate results are included to build further confidence in the eigenvalue based results and to confirm the RANS prediction.

Figure 22 describes element S^c_{22} of the Schur correction matrix for the Euler–based ROM including the design samples and the trace of the critical eigenvalue. Using this ROM the instability boundary is correctly predicted for both the critical flutter speed index and the critical flutter frequency compared to the full order results as can be seen in Fig. 21. The oscillations in the instability boundary found for the full order predictions have been discussed in Ref. [49] and related to the numerically discrete shock movement with increments in Mach number. For the used resolution of full order samples it can not be expected that these oscillations are predicted by the interpolating ROM. The second unstable aeroelastic mode for high values of the flutter speed index is not completely resolved by the ROM though is still a very good prediction. The agreement to the time–accurate results is excellent throughout.

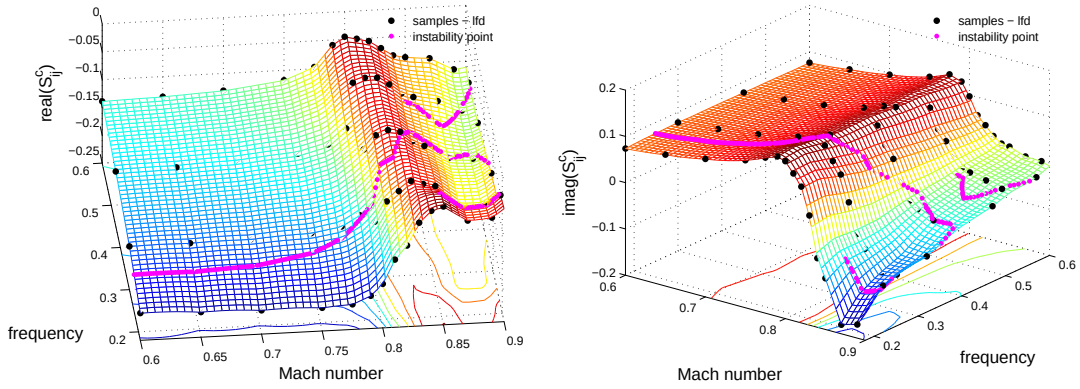

(a) Real part (b) Imaginary part

Figure 22. Real and imaginary part of element S_{22}^c of Schur correction matrix for Euler flow model reconstructed by globally defined reduced order models for the Isogai test case including traces of critical eigenvalue.

The initial guess of the parameter range of interest (blind search) to evaluate the ROM for the RANS prediction is based on the Euler results. However, the Euler predictions are misleading for the RANS simulations once strong boundary layer separation due to shock impingement is encountered. This can be seen in Fig. 23 describing element S_{22}^c for the RANS–based ROM including the design samples and the trace of the critical eigenvalue. Here, the variation in frequency ranges from 0.2 to 0.6 which is reasonable for the Euler results to include the second branch of the first unstable mode and to find the second unstable mode. In the figure the trace of the critical eigenvalue leaves the coverage of the full order RANS samples and one should refrain from using the kriging predictor for extrapolation as can be seen in Fig. 21 where the time–accurate results clearly deviate from the eigenvalue—based prediction. Therefore, it was not found to be useful to extend the instability boundary based on the current samples for the RANS flow modelling beyond about Mach 0.87. The ideas as presented in Section 6.2 about using a lower–fidelity model that contains the important phenomena, i.e. flow separation, to build a higher–fidelity model become evident. Instead of basing the critical parameter range on the Euler results the FPv results should be used since the latter formulation would be able to predict the separated flow effect.

While the Euler–based ROM can be considered as a global ROM for the whole domain of interest, it was found to be more convenient to split the RANS–based ROM into two parts; one for lower and one for higher Mach numbers with an overlap in between. The lower Mach numbers reach from the subsonic to the transonic regime with slight shock–induced boundary layer separation, whereas the second part covers strong shock/boundary layer interaction (SBLI). The definition of slight and strong SBLI is somewhat ambiguous; slight SBLI should include limited separated regions in the vicinity of the shock impingement, while strong SBLI includes massively separated flow extending to the trailing edge. A globally defined RANS–based ROM covering the complete Mach number range was difficult to obtain, i.e. an optimal correlation coefficient $\boldsymbol{\theta}$ for the kriging predictor was not found, due to the complexity of the shape in the response for strongly separated flow and the prediction in the sub- and lower transonic regime were adversely affected (results of which are not presented herein).

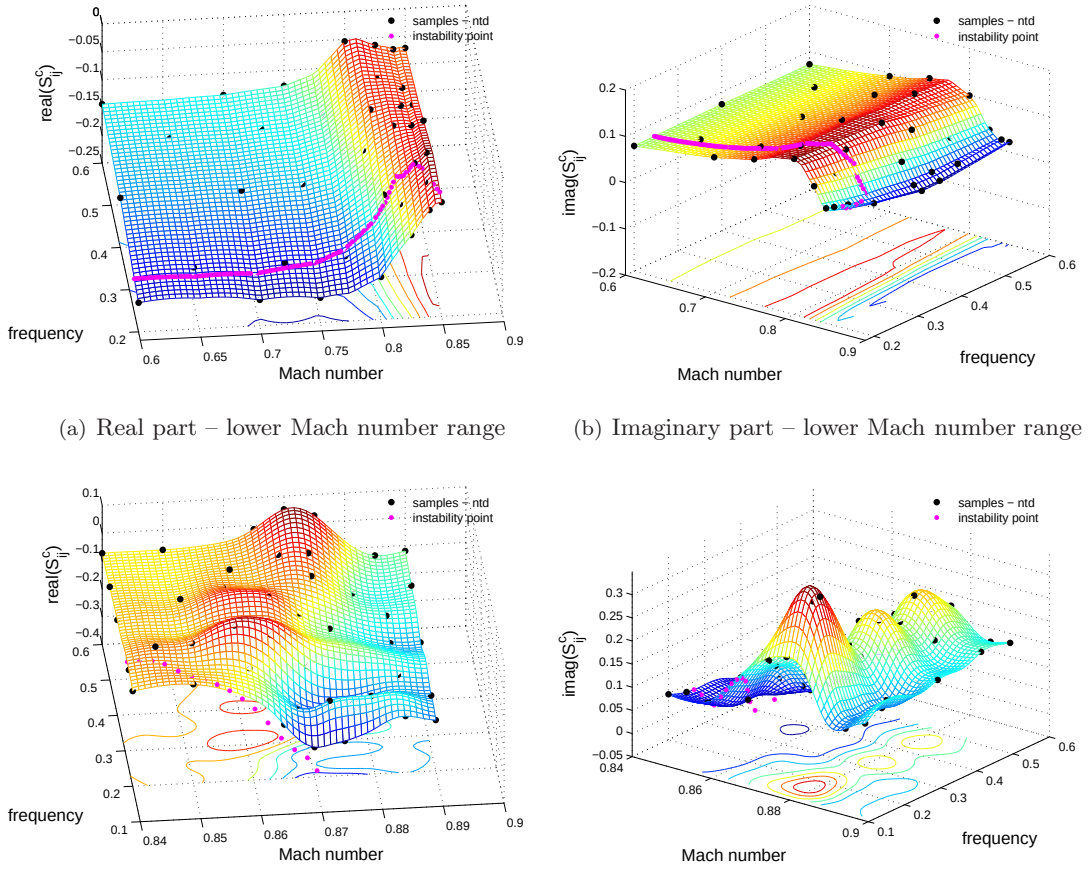

(a) Real part – lower Mach number range

(b) Imaginary part – lower Mach number range

(c) Real part – higher Mach number range

(d) Imaginary part – higher Mach number range

Figure 23. Real and imaginary part of element S_{22}^c of Schur correction matrix for RANS flow model reconstructed by two locally defined reduced order models for the Isogai test case including traces of critical eigenvalue.

The RANS–based prediction of aeroelastic stability shows differences to the Euler–based prediction. As for the NACA 0012 aerofoil configuration for sub- and low transonic Mach numbers the viscous effects seem to have a slightly stabilizing effect on the configuration. In the Mach number region where the impinging shock wave causes distinct boundary layer separation the sudden (almost vertical) increase in the critical flutter speed index can be seen removing the multiple bifurcations, and this is consistent with the observations given in Refs. [42, 44]. A comparison on pressure distributions C_p and skin friction coefficient C_f is presented in Fig. 24 for three distinct Mach numbers of interest. In the current study time–marching full blown RANS simulations confirm this trend accurately.

It shall be attempted to interpret the differences in the shape of the predicted Schur correction element S_{22}^c between the Euler and the RANS flow models. For subsonic Mach numbers (critical Mach number of NACA 64A010 is about 0.765) there are hardly changes to the shape of the predicted element for both real and imaginary parts. A first order polynomial seems to be appropriate to describe the changes in the dimensions of Mach number and frequency analytically. Then, with the development of the transonic shock wave there is a rapid change of the shape more distinct in the dimension of the Mach number than the frequency. With the development of strong shock–induced flow separation (at about Mach 0.85) there are clear differences between the RANS and Euler results (especially for the imaginary

part). The RANS–based elements develop a second plateau stopping the rapid change as found for the Euler elements which now also describe a stronger variation in the dimension of the frequency. The latter observation could be related to the multiple bifurcations found in the Euler prediction. A similar trend is found for all non–zero elements of the Schur correction matrix.

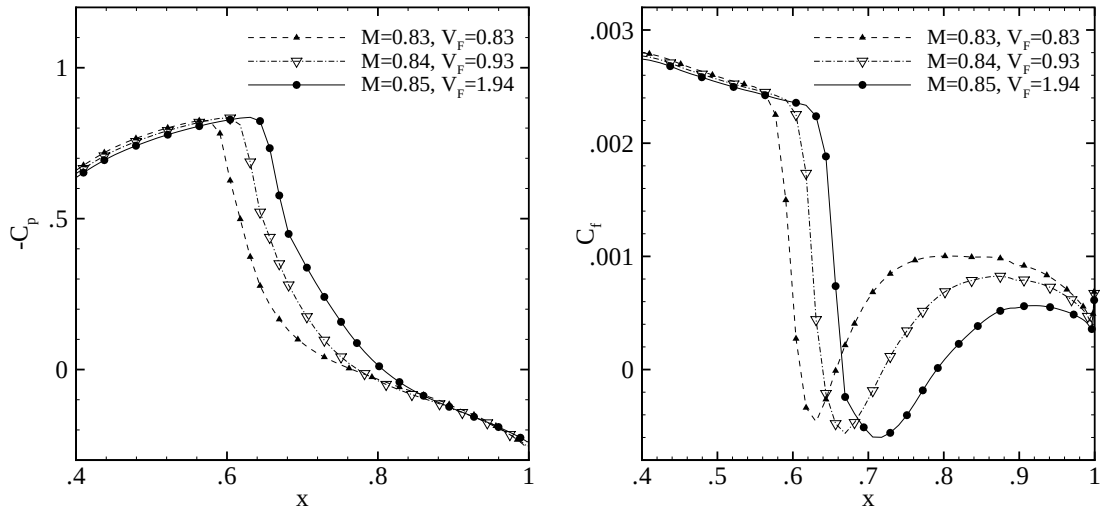

(a) Pressure distribution

(b) Skin friction coefficient

Figure 24. Pressure distribution and skin friction coefficient at three distinct Mach numbers for Isogai test case.

The strongly varying shape of the RANS elements for high Mach numbers requires further attention as it can not be said definitely that the extraction of the Schur elements based on nonlinear forced motion is precise for strongly separated flow. Repeated simulations at the same Mach number/frequency combinations but at different motion amplitudes seem to be useful to find an estimate of the involved error. Also, comparisons between individual vs simultaneous (either in frequency or in structural states) excitations should be attempted.

8 Conclusions and Outlook

The approach presented in this paper exploits the formulation of the Schur complement framework and builds a reduced order model using the kriging interpolation algorithm for the fluid dominated part of the Schur complement matrix based on samples of the full order model. The flow models are presented along with benchmark results. Ways to extract the samples are discussed and related issues such as shock–induced oscillations in the system response are analysed. While the basic Schur complement formulation is already significantly faster than common time–accurate approaches, the reduced order formulation proves to be computationally more efficient than the basic formulation despite the costs spent in the construction of the reduced order model itself. Another convenient aspect of the proposed method is the access to higher fidelity flow models in the aeroelastic stability analysis such as Reynolds–averaged Navier–Stokes modelling which has not been done within the framework before.

First ideas to exploit the proposed reduced order formulation have been presented and this work will continue. For instance, the dimension of independent variables in the kriging prediction can be extended to include altitude effects for more realistic three–dimensional cases or even structural parameters. Obviously, at this point more efforts have to be put on the sampling techniques and the kriging prediction than is done now. Also, the important issue of addressing aerodynamic modelling uncertainty will be considered in more detail and the ideas of model updating using co–kriging will be extended.

Acknowledgments

This research has been supported by the European Union for the Marie Curie Centre of Excellence 'ECERTA' under contract MEXT–CT–2006–042383.

References

[1]Walters, R. W. and Huyse, L., "Uncertainty Analysis for Fluid Mechanics with Applications," Tech. Rep. NASA/CR–2002–211449, NASA Langley Research Center, Hampton, VA, 2002.

[2]Walters, R. W., "Towards Stochastic Fluid Mechanics via Polynomial Chaos," *AIAA Paper 2003–0413*, 2003.

[3]Xiu, D. and Karniadakis, G. E., "Supersensitivity due to uncertain boundary conditions," *Int. J. Numer. Meth. Engng*, Vol. 61, 2004, pp. 2114–2138.

[4]Pettit, C. L., "Uncertainty quantification in aeroelasticity: recent results and research challenges," *Journal of Aircraft*, Vol. 41, No. 5, 2004, pp. 1217–1229.

[5]Beran, P. S. and Lucia, D. J., "Numerical analysis of store–induced limit–cycle oscillation," *Journal of Aircraft*, Vol. 41, No. 6, 2004, pp. 1315–1326.

[6]Schewe, G., Mai, H., and Dietz, G., "Nonlinear effects in transonic flutter with emphasis on manifestations of limit cycle oscillations," *Journal of Fluids and Structures*, Vol. 18, 2003, pp. 3–22.

[7]Dietz, G., Schewe, G., and Mai, H., "Experiments on heave/pitch limit–cycle oscillations of a supercritical airfoil close to the transonic dip," *Journal of Fluids and Structures*, Vol. 19, 2004, pp. 1–16.

[8]Dietz, G., Schewe, G., and Mai, H., "Amplification and amplitude limitations of heave/pitch limit–cycle oscillations close to the transonic dip," *Journal of Fluids and Structures*, Vol. 22, 2006, pp. 505–527.

[9]Bendiksen, O. O., "Nonlinear mode interactions and period–tripling flutter in transonic flow," *Journal of Fluid and Structures*, Vol. 19, 2004, pp. 591–606.

[10]Dowell, E. H., Edwards, J. W., and Strganac, T. W., "Nonlinear aeroelasticity," *Journal of Aircraft*, Vol. 40, No. 5, 2003, pp. 857–874.

[11]Friswell, M. I. and Mottershead, J. E., *finite element model updating in structural dynamics*, Kluwer Academic Press, Dorndrecht, The Netherlands, 1995.

[12]Farhat, C., Geuzaine, P., and Brown, G., "Application of a three–field nonlinear fluid–structure formulation to the prediction of the aeroelastic parameters of an F–16 fighter," *Computers & Fluids*, Vol. 32, No. 1, 2003, pp. 3–29.

[13]Woodgate, M. A., Badcock, K. J., Rampurawala, A. M., Richards, B. R., Nardini, D., and deC Henshaw, M. J., "Aeroelastic calculations for the Hawk aircraft using the Euler equations," *AIAA Journal*, Vol. 42, No. 4, 2005, pp. 1005–1012.

[14]Lucia, D. J., Beran, P. S., and Silva, W. A., "Reduced–order modeling: new approaches for computational physics," *Progress in Aerospace Sciences*, Vol. 40, 2004, pp. 51–117.

[15]Amsallem, D. and Farhat, C., "Interpolation method for adapting reduced–order models and application to aeroelasticity," *AIAA Journal*, Vol. 46, No. 7, 2008, pp. 1803–1813.

[16]Morton, S. A. and Beran, P. S., "Hopf–bifurcation analysis of airfoil flutter at transonic speeds," *AIAA Paper 96–0060*, 1996.

[17]Morton, S. A. and Beran, P. S., "Hopf bifurcation analysis applied to deforming airfoils at transonic speeds," *AIAA Paper 97–1772*, 1997.

[18]Morton, S. A. and Beran, P. S., "Hopf–bifurcation analysis of airfoil flutter at transonic speeds," *Journal of Aircraft*, Vol. 36, No. 2, 1999, pp. 421–429.

[19]Badcock, K. J., Woodgate, M. A., and Richards, B. E., "Hopf bifurcation calculations for a symmetric airfoil in transonic flow," *AIAA Journal*, Vol. 42, No. 5, 2004, pp. 883–892.

[20]Badcock, K. J., Woodgate, M. A., and Richards, B. E., "Direct aeroelastic bifurcation analysis of a symmetric wing based on Euler equations," *Journal of Aircraft*, Vol. 42, No. 3, 2005, pp. 731–737.

[21]Woodgate, M. A. and Badcock, K. J., "Fast prediction of transonic aeroelastic stability and limit cycles," *AIAA Journal*, Vol. 45, No. 6, 2007, pp. 1370–1381.

[22]Badcock, K. J. and Woodgate, M. A., "Prediction of bifurcation onset of large order aeroelastic models," *AIAA Paper 2008–1820*, 2008.

[23]Marques, S., Badcock, K. J., Khodaparast, H. H., and Mottershead, J. E., "CFD based aeroelastic stability predictions under the influence of structural variability," *AIAA Paper 2009–2324*, 2009.

[24]Raj, P., "Computational uncertainty: Achilles' heel of simulation based aircraft design," AVT–147 Symposium on 'Computational Uncertainty in Military Vehicle Design', 2007.

[25]Badcock, K. J., Richards, B. E., and Woodgate, M. A., "Elements of computational fluid dynamics on block structured grids using implicit solvers," *Progress in Aerospace Sciences*, Vol. 36, 2000, pp. 351–392.

[26]Jameson, A., "Time dependent calculations using multigrid with applications to unsteady flows past airfoils and wings," *AIAA Paper 91–1596*, 1991.

[27]Osher, S. and Chakravarthy, S., "Upwind schemes and boundary conditions with applications to Euler equations in general geometries," *Journal of Computational Physics*, Vol. 50, 1983, pp. 447–481.

[28]Holst, T. L., "Transonic flow computations using nonlinear potential methods," *Progress in Aerospace Sciences*, Vol. 36, 2000, pp. 1–61.

[29]Whitfield, D. L., "Analytical description of the complete turbulent boundary layer velocity profile," *AIAA Paper 78–1158*, 1978.

[30]Whitfield, D. L., Swafford, T. W., and Jacocks, J. L., "Calculation of turbulent boundary layers with separation and viscous–inviscid interaction," *AIAA Journal*, Vol. 19, No. 10, 1981.

[31]Swafford, T. W., "Analytical approximation of two–dimensional separated turbulent boundary layers," *AIAA Journal*, Vol. 21, No. 6, 1983, pp. 923–926.

[32]Drela, M., *Two–dimensional transonic aerodynamic design and analysis using the Euler equations*, Ph.D. thesis, Department of Aeronautics and Astronautics, Massachusetts Institute of Technology, 1985.

[33]Drela, M. and Giles, M. B., "Viscous–inviscid analysis of transonic and low Reynolds number airfoils," *AIAA Journal*, Vol. 25, No. 10, 1987, pp. 1347–1355.

[34]Schlichting, H. and Gersten, K., *Boundary–layer theory*, Springer–Verlag, Inc., Berlin, Germany, 8th rev. and enl. ed., 2000.

[35]Lighthill, M. J., "On displacement thickness," *Journal of Fluid Mechanics*, Vol. 4, 1958, pp. 383–392.

[36]Fung, Y. C., *An introduction to the theory of aeroelasticity*, John Wiley & Sons, Inc., New York, NY, 1955.

[37]Cook, P. H., McDonald, M. A., and Firmin, M. C. P., "Aerofoil RAE 2822 – Pressure distributions, and boundary layer and wake measurements," Tech. Rep. AGARD AR 138, 1979.

[38]AGARD, "Compendium of unsteady aerodynamic measurements," Tech. Rep. 702, 1982.

[39]Howlett, J. T., "Calculation of Unsteady transonic flows with mild separation by viscous–inviscid interaction," Tech. Rep. 3197, NASA Langley Research Center, Hampton, VA, 1992.

[40]Badcock, K. J. and Gaitonde, A. L., "An unfactored implicit moving mesh method for two–dimensional unsteady N–S equations," *Int. J. Numer. Meth. Fluids*, Vol. 23, 1996, pp. 607–631.

[41]Isogai, K., "Transonic–dip mechanism of flutter of a sweptback wing: part II," *AIAA Journal*, Vol. 19, No. 9, 1981, pp. 1240–1242.

[42]Prananta, B. B., Hounjet, M. H. L., and Zwaan, R. J., "Two–dimensional transonic aeroelastic analysis using thin–layer Navier–Stokes method," *Journal of Fluids and Structures*, Vol. 12, 1998, pp. 655–676.

[43]Hall, K. C., Thomas, J. P., and Dowell, E. H., "Proper orthogonal decomposition technique for transonic unsteady aerodynamic flows," *AIAA Journal*, Vol. 38, No. 10, 2000, pp. 1853–1862.

[44]Yang, S., Zhang, Z., Liu, F., Luo, S., Tsai, H.-M., and Schuster, D. M., "Time–domain aeroelastic simulation by a coupled Euler and integral boundary–layer method," *AIAA Paper 2004–5377*, 2004.

[45]Försching, H. and Hönlinger, H., "The expanding domain of aeroelastic simulation," *Euromech–Colloquium 349*, 1996, pp. 1–14.

[46]Bekas, C. and Saad, Y., "Computation of smallest eigenvalues using spectral Schur complements," *SIAM J. Sci. Comput.*, Vol. 27, No. 2, 2005, pp. 458–481.

[47]Oran Brigham, E., *The fast Fourier transform and its applications*, Prentice–Hall, Englewood Cliffs, NJ, 1988.

[48]Timme, S., "CFD method formulation for CFD sensitivity: Oscillation of instability boundary due to steady state resolution," Tech. Rep. Deliverable D 3.1, CFD Laboratory, University of Liverpool, 2008.

[49]Timme, S. and Badcock, K. J., "Oscillatory behavior of transonic aeroelastic instability boundaries," *AIAA Journal*, Vol. 47, No. 6, 2009, pp. 1590–1592.

[50]Sacks, J., Welch, W. J., Mitchell, T. J., and Wynn, H. P., "Design and analysis of computer experiments," *Statistical Science*, Vol. 4, No. 4, 1989, pp. 409–435.

[51]Jones, D. R., Schonlau, M., and Welch, W. J., "Efficient global optimization of expensive black–box functions," *Journal of Global Optimization*, Vol. 13, No. 4, 1998, pp. 455–492.

[52]Welch, W. J., Buck, R. J., Sacks, J., Wynn, H. P., Mitchell, T. J., and Morris, M. D., "Screening, predicting, and computer experiments," *Technometrics*, Vol. 34, No. 1, 1992, pp. 15–25.

[53]Lophaven, S. N., Nielsen, H. B., and Sœndergaard, J., "DACE – A matlab kriging toolbox," Tech. Rep. IMM-TR-2002-12, Technical University of Denmark, Denmark, 2002.

[54]McKay, M. D., Conover, W. J., and Beckman, R., "A comparison of three methods for selecting values of input variables in the analysis of output from a computer code," *Technometrics*, Vol. 21, No. 2, 1979, pp. 239–245.

[55]Laurenceau, J. and Sagaut, P., "Building efficient response surfaces of aerodynamic functions with kriging and cokriging," *AIAA Journal*, Vol. 46, No. 2, 2008, pp. 498–507.

[56]Han, Z. H., Görtz, S., and Zimmermann, R., "On improving efficiency and accuracy of variable–fidelity surrogate modeling in aero–data for loads context," *CEAS 2009 European Air and Space Conference, Proceedings of the Conference, Manchester, UK, 26–29 October 2009*, 2009.