ECCOMAS Congress 2016
VII European Congress on Computational Methods in Applied Sciences and Engineering
M. Papadrakakis, V. Papadopoulos, G. Stefanou, V. Plevris (eds.)
Crete Island, Greece, 5–10 June 2016

LINKING PARAMETRIC CAD WITH ADJOINT SURFACE SENSITIVITIES

Ilias Vasilopoulos[1,3], Dheeraj Agarwal[2], Marcus Meyer[1], Trevor T. Robinson[2] and Cecil G. Armstrong[2]

[1]Rolls-Royce Deutschland (RRD)
Eschenweg 11, Dahlewitz, D-15827 Blankenfelde-Mahlow, Germany
{ilias.vasilopoulos, marcus.meyer}@rolls-royce.com

[2]School of Mechanical and Aerospace Engineering, Queen's University Belfast (QUB)
Ashby Building, Belfast BT9 5AH, Northern Ireland, UK
{d.agarwal, t.robinson, c.armstrong}@qub.ac.uk

[3]School of Mechanical Engineering, National Technical University of Athens (NTUA)
Iroon Polytecniou 9, 15780 Zografou, Athens, Greece
vasilopoulos@central.ntua.gr

Keywords: CAD, CFD, Adjoint, Design Velocities, Sensitivity Map, Parametric Sensitivities.

Abstract. *The goal of this work is to present an efficient CAD-based adjoint process chain for calculating parametric sensitivities (derivatives of the objective function with respect to the CAD parameters) in timescales acceptable for industrial design processes. The idea is based on linking parametric design velocities (geometric sensitivities computed from the CAD model) with adjoint surface sensitivities. A CAD-based design velocity computation method has been implemented based on distances between discrete representations of perturbed geometries. This approach differs from other methods due to the fact that it works with existing commercial CAD packages (unlike most analytical approaches) and it can cope with the changes in CAD model topology and face labeling. Use of the proposed method allows computation of parametric sensitivities using adjoint data at a computational cost which scales with the number of objective functions being considered, while it is essentially independent of the number of design variables. The gradient computation is demonstrated on test cases for a Nozzle Guide Vane (NGV) model and a Turbine Rotor Blade model. The results are validated against finite difference values and good agreement is shown. This gradient information can be passed to an optimization algorithm, which will use it to update the CAD model parameters.*

1 INTRODUCTION

Throughout industry, engineers use Computer Aided Design (CAD) software packages for design and Computer Aided Engineering (CAE) tools for analysis. The CAD provides the environment for the geometry generation of a model and its parametrisation, whereas the CAE, such as Finite Element Analysis (FEA) and Computational Fluid Dynamics (CFD) solvers, performs an analysis on a computational mesh which approximates the geometry.

The engineer is interested in obtaining the optimum geometry (in terms of some objective function of interest), while satisfying predetermined (e.g. manufacturing) constraints. The use of efficient gradient-based optimization algorithms requires the derivatives of the objective function w.r.t. the design variables. In case of CAD-based optimization, the CAD parameters which define the geometric features can be considered as design variables. The derivative of the model's performance w.r.t. each design parameter can be obtained in a finite difference sense, where the effect of a parameter change is computed by analysing both (baseline and updated) geometries and comparing the results. For each parameter, a perturbed geometry is created in CAD and then fed to CAE for analysis (including geometry healing and mesh generation processes), where the resulting difference in performance enables the derivative calculation. This finite difference approach is computationally expensive, since the cost scales with the number of design parameters, and can thus be prohibitive for industrial applications.

As an alternative and much more efficient way of computing gradient information, adjoint methods have been the subject of considerable research in recent years [1, 2, 3, 4, 5]. The advantage of an adjoint method arises from the fact that the gradient computation becomes independent of the number of design variables. This is achieved by solving a linear problem of computational cost comparable to that of the primal analysis. As a result, most adjoint-based optimization strategies use the model's surface mesh node coordinates as design variables, which represents the richest design space the CFD can evaluate. On the other hand, since all grid surface nodes can move independently, the implementation of a smoothing algorithm is required to prevent the appearance of oscillatory shapes during the optimization process. Moreover, such CAD-free methods produce as their output an optimized mesh, which then has to be converted back to CAD in order to obtain the optimized geometry for further analysis or manufacturing. This mesh-to-CAD step is a non-trivial task and it may require extensive user interaction. For these reasons, CAD-based optimization is considered more desirable and thus a link between adjoint and CAD is of utmost importance.

Very limited work on CAD-based optimization is currently available [6, 7]. Three promising approaches to include the CAD system into the design loop are explored within the EC ITN IODA project: a) NURBS-based parametrisation, b) Parametric CAD sensitivities through finite-differencing and c) CAD derivatives through automatic differentiation (AD).

Non-uniform rational B-splines (NURBS) patches are a generic surface description and are used by all CAD packages (STEP format). Queen Mary University of London (QMUL) has developed a NURBS-based approach [8], where the NURBS control points are considered as design variables and the required derivatives are obtained by AD of the NURBS analytical description. The approach is not limited to a single patch, but allows to enforce continuity constraints across patch interfaces. Recent work has extended this CAD-based parametrisation method to include geometric constraints (such as thickness and trailing edge radius) [9].

Geometric constraints are best incorporated by a restriction of the design space, which is straightforwardly achieved through the parametrisation of the CAD model. However, if the native CAD model parameters are to be used as design variables, then the derivatives of the

model's surface w.r.t. these parameters are required, since they are not provided by today's commercial CAD packages. One approach to obtain these geometric sensitivities (also found in the literature as design velocities [10]) is to compute finite differences between the CAD model before and after a parameter perturbation. This approach has been identified by some researchers [11, 12, 13] to lack robustness against topology and labelling changes, which can occur even under small parameter perturbations. Nevertheless, Queen's University of Belfast (QUB) has developed an approach based on projections between discrete representations of the CAD geometries [14], which is efficient, robust and has been successfully applied to CFD [15].

To compute the exact geometric derivatives, one can differentiate the CAD kernel, either analytically or via AD tools [16]. Differentiating the CAD kernel is possible only when you have access to the CAD software source code, as in OpenCASCADE [27]. Although this approach is promising, the automatic differentiation of a complete design chain incorporating a CAD tool still remains a challenging task to be tackled.

In this work, the link to CAD is achieved by applying an improved QUB design velocity approach, in which the CAD modeling system is treated as a black-box. The approach uses the generic STEP surface description and thus it can be easily coupled to any feature-based CAD tool. It is immune to topology changes caused by parametric variations and it requires no access to the CAD source code. This makes it suitable for implementation within an industrial context, where commercial CAD packages (such as Siemens NX or CATIA V5) are widely used.

The idea is based on linking two concepts: adjoint surface sensitivities and design velocities. Adjoint surface sensitivities (which form a sensitivity map) can be obtained from an adjoint CFD solution and represent the derivative of the objective function w.r.t. normal perturbations of the boundary mesh at each node. Fig. 1(a) shows the sensitivity map provided by Rolls-Royce Deutschland (RRD) for a Nozzle Guide Vane (NGV) case. Areas of high sensitivity are shown in red or blue, representing areas where the body should be displaced outwards or inwards (respectively) to achieve an increase in objective function, whilst areas of low sensitivity (plotted in green) show that the boundary in this region has no large influence on the objective function.

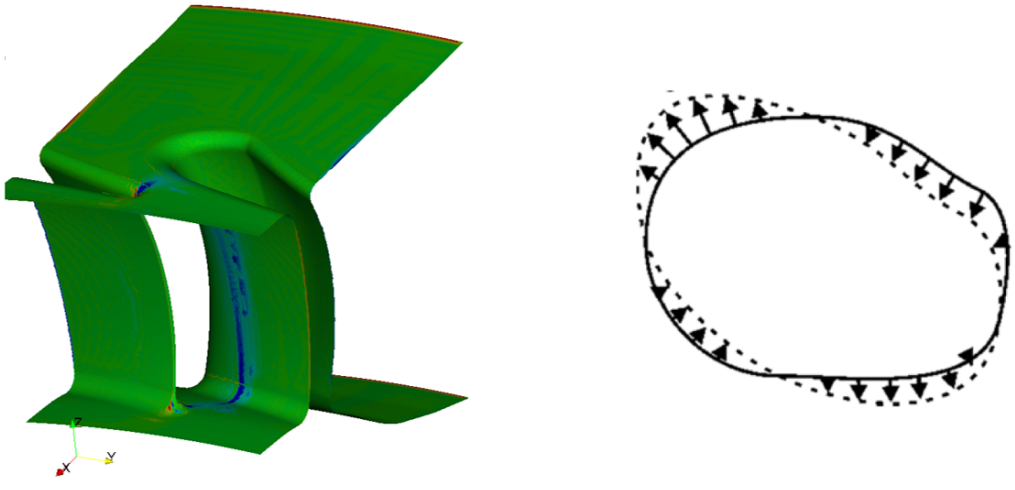

a) b)

Figure 1: a) Adjoint sensitivity map for NGV test case and b) Simple example of 2D design velocity field.

On the other hand, design velocities are geometric sensitivities computed from the CAD model, representing the displacement of a point on the model boundary due to the perturbation of a CAD model parameter. This work makes use of the normal component of the design velocity, V_n, which is computed as

$$V_n = dp \cdot \hat{n} \tag{1}$$

where dp is the displacement vector of a point on the boundary due to a parameter perturbation and $\hat{n}$ is the outward unit normal of the surface at that point. This measure was initially developed in the context of adjoint-based structural optimization problems [17]. In Fig. 1(b), which is taken from [13], the arrows represent the design velocities as the boundary changes from the solid to the dashed line. Due to the convention adopted throughout this paper that the boundary normals are pointing outwards, a positive design velocity represents an outward movement of the boundary, and negative is inward.

The CAD-based adjoint process chain for calculating parametric sensitivities (derivatives of the objective function w.r.t. the CAD parameters), followed in this work, is outlined in Fig. 2.

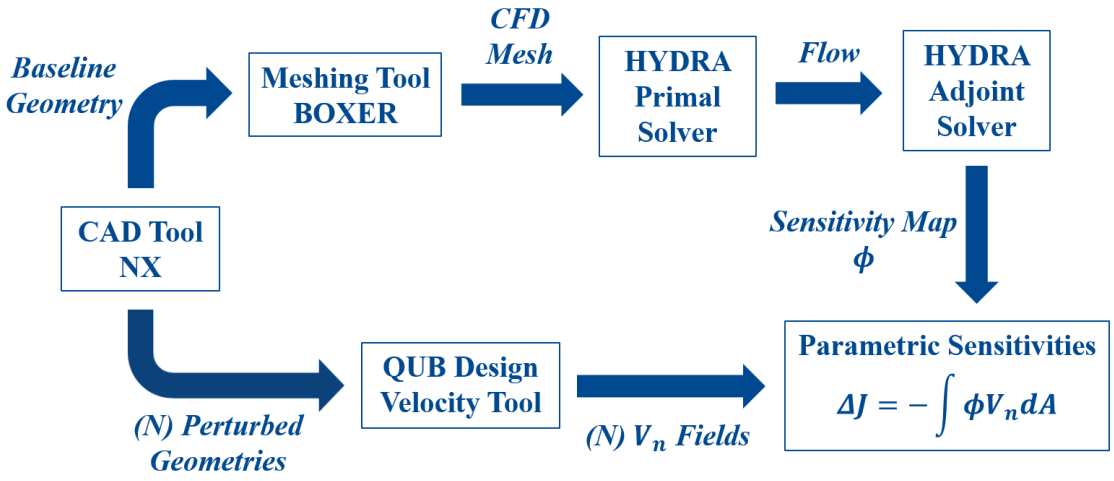

Figure 2: Workflow for parametric sensitivities computation.

Use of the above approach allows the computation of parametric sensitivities at a cost of one adjoint analysis per objective function and one design velocity field evaluation per parameter. Since design velocities can be computed for all of the parameters in a much shorter time than that needed to converge the primal and adjoint solutions, gradient computation is achieved in timescales acceptable for industrial design processes. As shown in the last step of Fig. 2, parametric sensitivities are calculated by applying the chain rule between the sensitivity map ϕ and the design velocities V_n.

The remainder of this paper is structured as follows: Sec. 2 explains the mathematical background of the discrete adjoint solver and how the sensitivity map is derived. Sec. 3 describes the method used for the computation of design velocities. The gradient results obtained for two turbomachinery test cases are presented in Sec. 4, followed by the conclusions in Sec. 5.

2 DISCRETE ADJOINT METHOD

2.1 Mathematical background

In this work, a RANS compressible flow solver and its corresponding discrete adjoint solver are used for computing the flow and sensitivity map. These are part of the Rolls-Royce in-house HYDRA suite of codes, which have been extensively validated and applied to various industrial cases (see for example [18, 19]). More details regarding the underlying theory and implemented algorithm can be found in [3, 20, 21]. The nonlinear flow solver uses a node-based finite-volume discretisation method and the pseudo-time-marching to steady state is accelerated by a block-Jacobi preconditioner and a geometric multigrid technique.

During the convergence of the primal solution, the nonlinear residual for each control volume is driven to zero

$$R(U, X) = 0 \,, \tag{2}$$

where U is the primal solution and X stands for the mesh coordinates. The objective function considered is a function of both the flow solution and the mesh, while the flow solution implicitly depends on the mesh. In addition, the mesh coordinates are defined by the CAD design variables α, which leads to the relation

$$J = J(U(X(\alpha)), X(\alpha)) \,. \tag{3}$$

The parametric sensitivities can be computed using the chain rule as

$$\frac{dJ}{d\alpha} = \frac{dJ}{dX}\frac{dX}{dX_S}\frac{dX_S}{d\alpha} \,, \tag{4}$$

where X_S stands for the surface mesh coordinates. The adjoint solver is responsible for the computation of the volume mesh sensitivities dJ/dX, which is achieved by solving the equations shown below.

The objective function is augmented as

$$J = J + \lambda^T R \,, \tag{5}$$

where λ is the adjoint solution. Differentiating Eq. (5) w.r.t. X leads to

$$\begin{aligned} \frac{dJ}{dX} &= \frac{dJ}{dX} + \lambda^T \frac{dR}{dX} \\ &= \frac{\partial J}{\partial X} + \frac{\partial J}{\partial U}\frac{dU}{dX} + \lambda^T \Big(\frac{\partial R}{\partial X} + \frac{\partial R}{\partial U}\frac{dU}{dX} \Big) \\ &= \frac{\partial J}{\partial X} + \lambda^T \frac{\partial R}{\partial X} + \Big(\frac{\partial J}{\partial U} + \lambda^T \frac{\partial R}{\partial U} \Big) \frac{dU}{dX} \,. \end{aligned} \tag{6}$$

The idea is to eliminate the computationally expensive term dU/dX. To do so, its multiplier in Eq. (6) is set to zero, giving rise to the adjoint equation

$$\frac{\partial R}{\partial U}^T \lambda = -\frac{\partial J}{\partial U}^T \,. \tag{7}$$

Obtaining the adjoint solution λ, by solving the linear Eq. (7), allows for the efficient computation of dJ/dX, given by the remaining terms of Eq. (6) as

$$\frac{dJ}{dX} = \frac{\partial J}{\partial X} + \lambda^T \frac{\partial R}{\partial X} \,. \tag{8}$$

As shown in Eq. (4), the computation of parametric sensitivities also requires the terms dX/dX_S and $dX_S/d\alpha$.

2.2 Sensitivity map

The term dX/dX_S represents the inverse operation of a mesh moving algorithm. It allows the computation of the surface mesh sensitivities dJ/dX_S as

$$\frac{dJ}{dX_S} = \frac{dJ}{dX}\frac{dX}{dX_S} \, . \tag{9}$$

For this projection of sensitivities from the volume to the surface, a simple spring-based mesh deformation algorithm is used. Eq. (9) results in

$$\frac{dJ}{dX_S}^T = \mathbf{K}_{s2v}^T \mathbf{K}^{-T} \frac{dJ}{dX}^T \, , \tag{10}$$

where $\mathbf{K}$ is the stiffness matrix and $\mathbf{K}_{s2v}^T$ is a permutation matrix mapping each volume node to its corresponding surface node if any, and to null if none. The only computationally expensive part is to approximately invert $\mathbf{K}^T$. A Jacobi preconditioned conjugate gradient solver is used in this work, resulting in negligible CPU cost compared to the primal and adjoint solution.

After this step, Eq. (4) is rewritten as

$$\frac{dJ}{d\alpha} = \frac{dJ}{dX_S}\frac{dX_S}{d\alpha} \, . \tag{11}$$

The adjoint sensitivity map is derived by projecting each surface node's sensitivity vector dJ/dX_S onto its corresponding boundary normal. The same is done for the missing term $dX_S/d\alpha$, which represents the design velocities (explained in Sec. 3), and the parametric sensitivities $dJ/d\alpha$ are finally obtained from the inner product of those two components.

3 DESIGN VELOCITY APPROACH

The approach presented in this paper for computing parametric design velocities is based on a finite difference between the CAD model's shape before and after changing the value of a parameter defining the geometry. It is a development of the methodology presented in [14] and is characterized by an increased robustness. It works with existing commercial CAD packages and it can cope with the changes in CAD model topology and face labeling, as discussed in the introduction.

The CAD geometry is represented using a surface mesh of triangular elements. In this work, the mesh representations (for the unperturbed model and for each perturbed model) are created using either Gmsh [28] or CADfix [29] grid generator. The displacement of the model boundary due to a perturbation is approximated by projecting the centroid of each of the elements in the unperturbed model onto the elements in the perturbed model, as shown in Fig. 3. The projection location Pp is defined by the intersection of the surface normal direction $\hat{n}$ at the centroid C of the original model element with the plane the perturbed element exists on. When the correct projection is identified, the displacement vector dp is calculated as the difference between the centroid and its projection. The design velocity is then given by Eq. (1).

The question that arises is which perturbed element should each element centroid in the unperturbed model be projected onto. In order to increase the efficiency of this search, a multidimensional binary search tree known as KD-tree [22] is employed to find the perturbed element with the closest centroid, which is tested first. If the result is an unsuccessful projection, Barycentric coordinates [23] are used to determine which perturbed element should

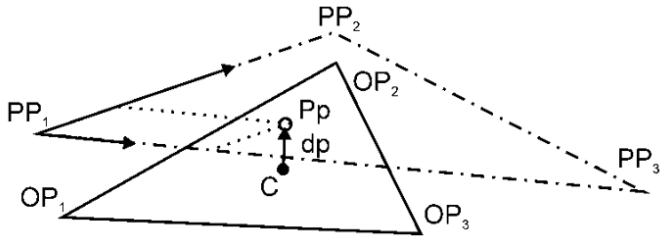

Figure 3: The centroid C of the original model element $OP_1OP_2OP_3$ is projected onto the perturbed element $PP_1PP_2PP_3$.

be tested next. The process continues until a successful projection onto an element face, edge or vertex is found, allowing the displacement to be calculated.

One additional check to ensure that the appropriate projection is identified is to demand the unperturbed boundary normal at the centroid that is projected to be similar to the perturbed boundary normal at the identified projection point. This is important in thin regions where otherwise the projection could occur onto the wrong face of the model. In this work, the surface normals are considered to be similar when

$$\hat{n}_{unperturbed} \cdot \hat{n}_{perturbed} > \frac{\sqrt{3}}{2} \, . \tag{12}$$

The design velocity computation process is outlined in Fig. 4 and it is fully automated using a Python script.

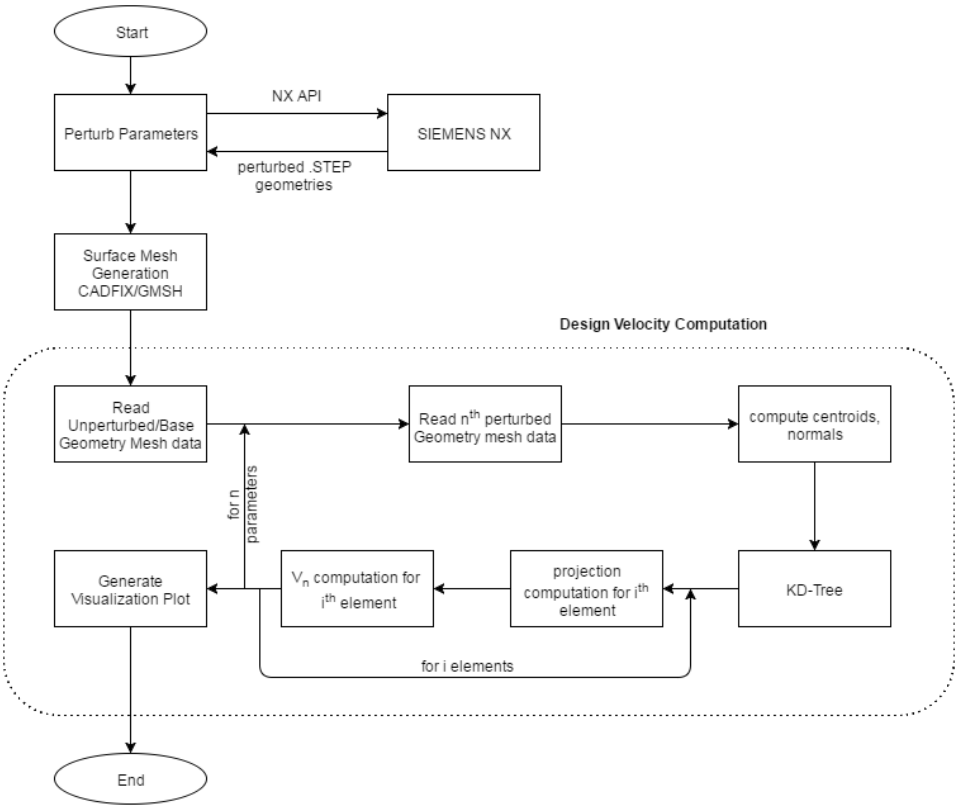

Figure 4: Design velocity computation process.

4 APPLICATIONS

4.1 Nozzle Guide Vane

In order to test the proposed approach, the first case used in this work is a state-of-the-art nozzle guide vane (NGV) of a high pressure turbine (HPT) provided by Rolls-Royce. The NGV geometry is shown in Fig. 5, where a 3D parametric CAD model is built in Siemens NX. The NGV has fillets at both ends and a cooling slot feature at the trailing edge (TE). This geometry is also investigated in [24], which illustrates different methods in order to quantify the impact of geometric variations on the NGV performance with focus on the capacity. Typically, the engine mass flow (and by association the capacity) is governed by the NGV design, which defines the narrowest cross section of the turbine. As a result, in this test case the capacity is considered as the objective function.

The CFD simulation requires only one periodic section of the engine's annulus. Consequently, the sector domain (Fig. 5) was used for both CFD and design velocities computation.

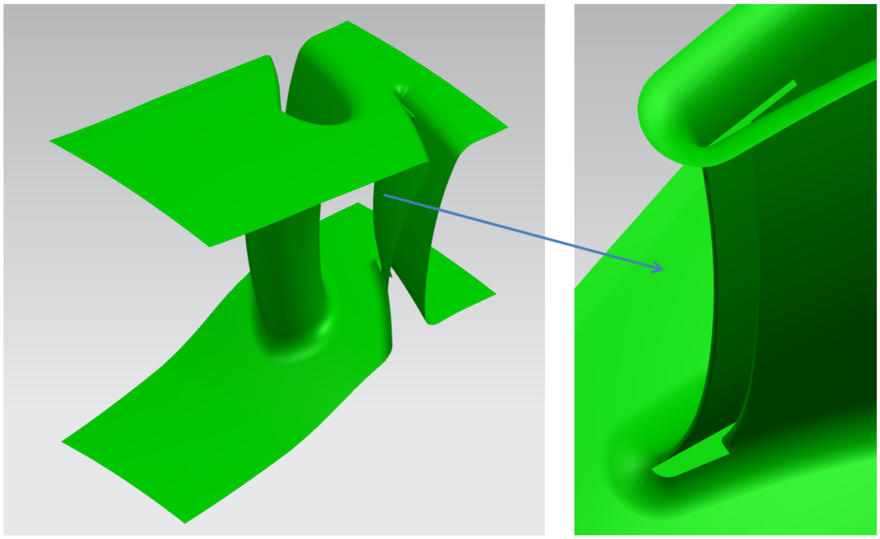

(a) (b)

Figure 5: 3D parametric CAD model in Siemens NX: a) NGV geometry and b) TE cooling slot.

The CAD model is parametrised using NX expressions and perturbed geometries are created using an iSIGHT workflow. Fig. 6 shows some of the 12 CAD parameters considered as exemplary design variables in this test case, together with the corresponding geometric variation that they cause.

For each perturbation of these parameters, a design velocity field was obtained using the QUB design velocity tool. In this case, the QUB tool was coupled with the open source grid generator Gmsh, which provided the mesh representations used for the projections. Design velocity fields for some of the parameters are shown in Fig. 7.

In order to be able to validate the results produced by linking these design velocity fields with the adjoint sensitivity map, a finite difference (FD) study has been conducted. Thus, 24 perturbed geometries (for + and - perturbations of the parameters) were generated using the iSIGHT workflow and an automated meshing procedure was executed using the BOXER

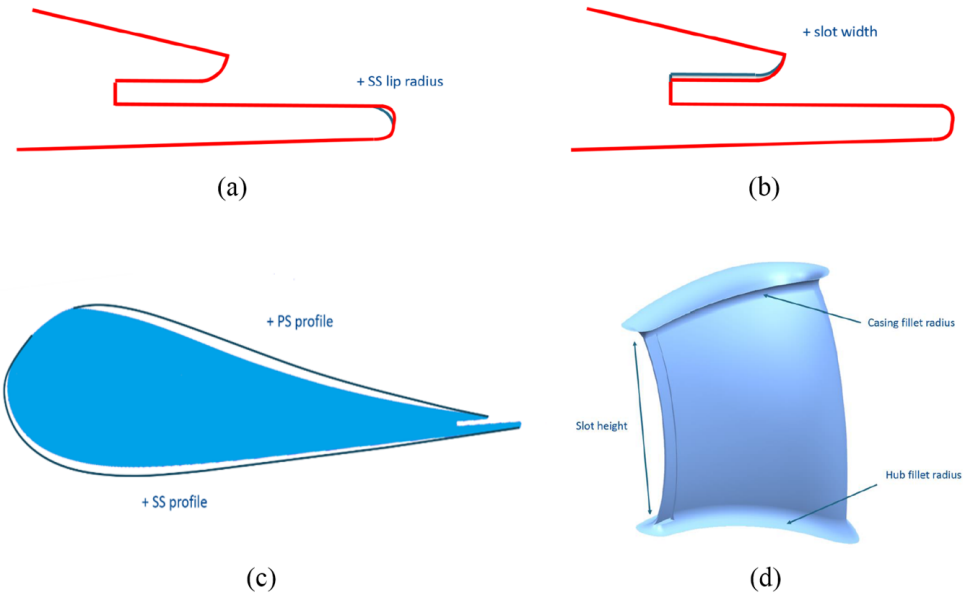

Figure 6: CAD feature parameters considered as design variables (not to scale): a) suction side lip radius, b) slot width, c) pressure side and suction side profile and d) slot height and fillet radius at hub and casing.

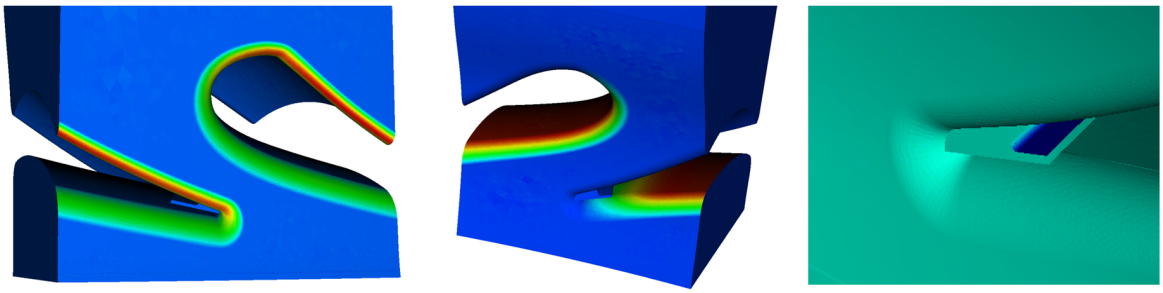

Figure 7: Design velocity fields due to negative parametric perturbations of: a) Casing fillet radius, b) SS profile and c) slot width.

meshing software. The produced meshes have approximately 9 million nodes and 13 million cells. Prior to conducting the CFD simulations, a mesh refinement study was performed to achieve a mesh-independent prediction of capacity. Fig. 8, illustrates the CFD domain, including a detailed view of mesh properties in the TE area.

The capacity results were obtained using the Rolls-Royce in-house CFD solver HYDRA, solving the steady RANS equations with Spalart-Allmaras turbulence model and wall functions. After convergence, the capacity q at the inlet was calculated as

$$q = \dot{m}\frac{\sqrt{T_t}}{p_t}\ , \tag{13}$$

where $\dot{m}$ denotes the inlet mass flow, T_t the total temperature and p_t the total pressure at the inlet using mass averaged values. The resulting capacity for the baseline geometry was found to match the expected capacity value very well and is taken as the reference value for all

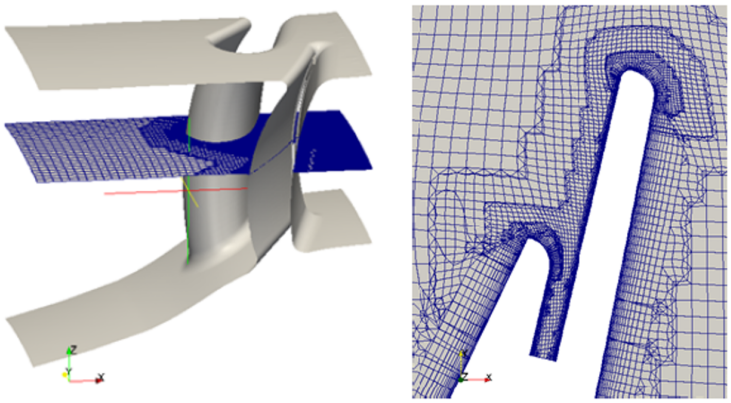

(a) (b)

Figure 8: a) NGV CFD domain and b) CFD mesh around TE.

subsequent comparisons with perturbed geometries.

The CFD simulations have shown that a linear relationship between parameter variation and capacity change applies for all of the parameters considered. This means that the adjoint approach should be able to capture the relation and predict accurate results.

Apart from the primal CFD results, the adjoint sensitivity map for the baseline geometry (needed for the proposed approach) was obtained using the HYDRA adjoint solver. The result is illustrated in Fig. 1(a), according to which, the most sensitive area is the suction side.

Finally, the change in capacity caused by each parametric perturbation is predicted by taking the inner product of the sensitivity map with the corresponding design velocity field. The parametric sensitivities can then be computed by dividing with the perturbation step. The obtained derivatives are shown in Fig. 9, where they are compared with central FD values.

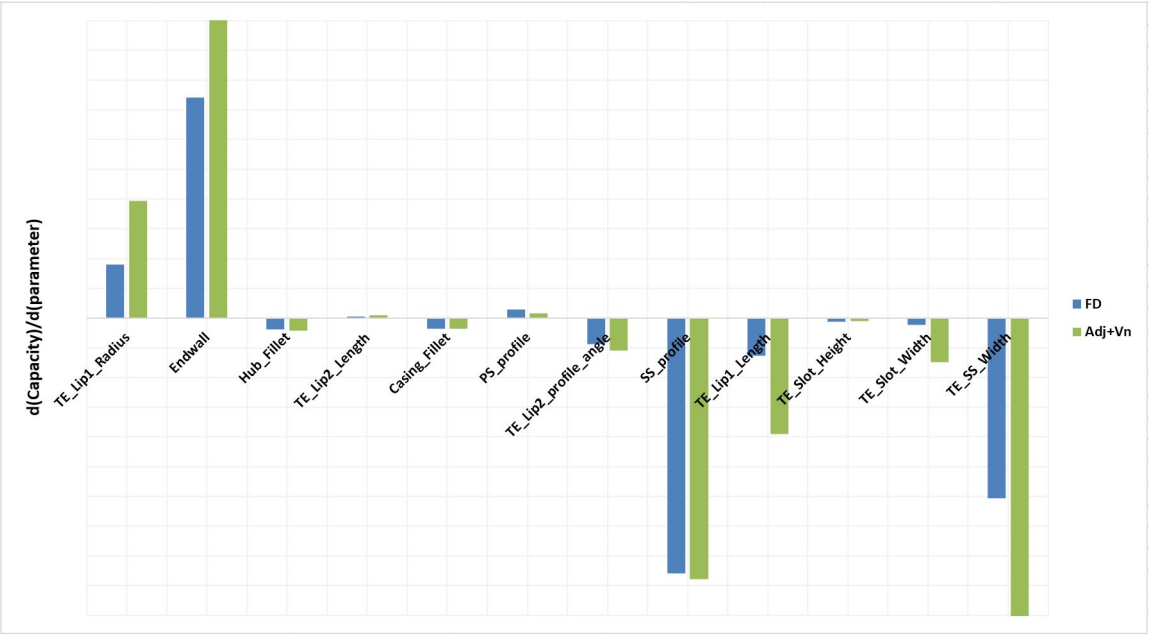

Figure 9: Capacity derivatives predicted by adjoint results (in green) and validation against FD (in blue).

As demonstrated in Fig. 9, for most parameters there is a good agreement between FD and adjoint derivatives. The CPU cost of the design velocity tool was around 2 hours (5 min per parametric perturbation), which is acceptable compared to the flow and adjoint solution. The main advantage of this approach is its efficiency, since it only requires 1 flow solution, 1 adjoint and 2 hours of post-processing, instead of 25 flow solutions (needed for the central FD).

4.2 Rotor blade with winglet

The second case tested in this work is a turbine rotor blade, also provided by Rolls-Royce. The rotor blade geometry is shown in Fig. 10(a), where a 3D parametric CAD model is built again in Siemens NX. The key feature of this blade is its winglet, a detailed view of which is included in Fig. 10(b). This feature is used to reduce the leakage mass flow, which leads to an increased rotor efficiency, by decreasing the driving pressure difference across the rotor's tip. The fluid stays inside the winglet and recirculations occur, as shown in Fig. 13(a), which may cause an increase in total pressure losses within the tip gap. A detailed study on winglets and squealers has been conducted in [25]. In what follows, the study is focused on the winglet geometry and total pressure loss between inlet and outlet is considered as the objective function.

A similar case has been used in [26], where the same approach for calculating parametric sensitivities is implemented and promising results are obtained. However, in this work the further developed QUB tool is combined with the latest HYDRA adjoint version, leading to results of increased accuracy.

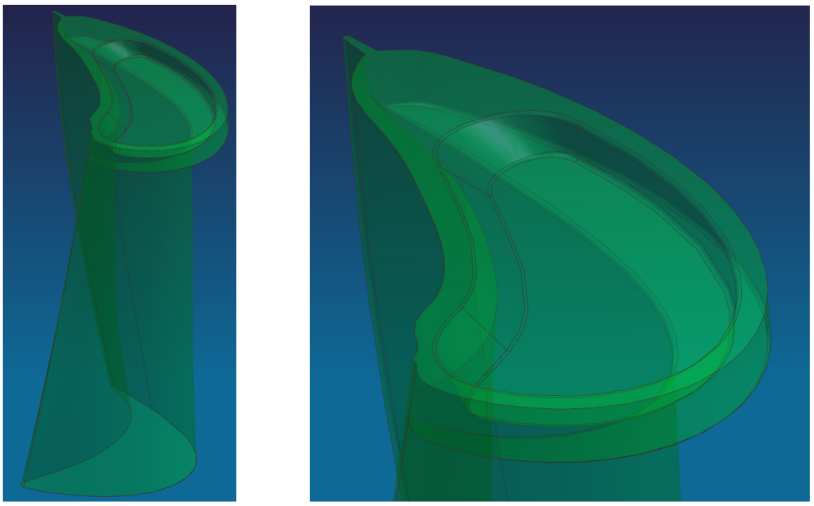

(a) (b)

Figure 10: 3D parametric CAD model in Siemens NX: a) Rotor blade geometry and b) winglet.

7 CAD parameters controlling the winglet geometry are considered as design variables. For each perturbation of these parameters, a design velocity field was calculated using the QUB tool, which was this time coupled with the commercial mesh generator CADfix, instead of Gmsh. Moreover, bigger perturbation steps were used in this case, in order to test the robustness of the approach. Some of the resulting design velocity fields are shown in Fig. 11.

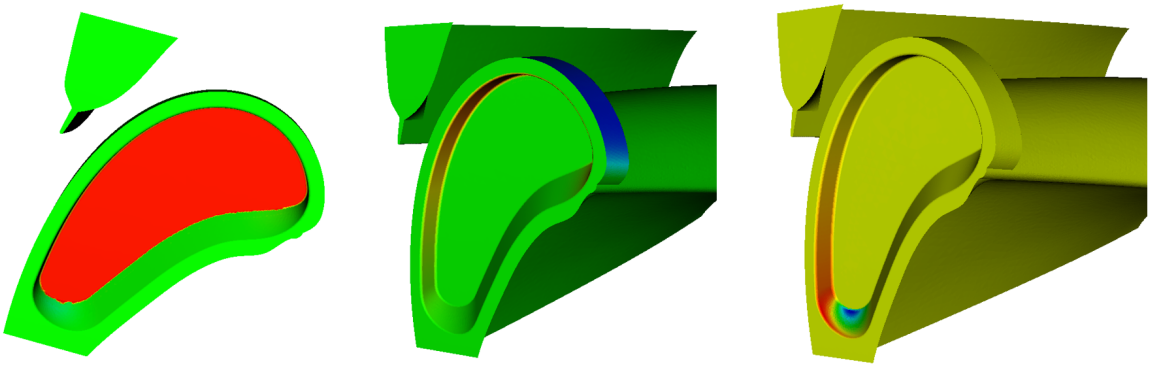

(a) (b) (c)

Figure 11: Design velocity fields due to parametric perturbations: a) -CAVITY_HEIGHT, b) -SS_FRONT_OFFSET and c) -SS_REAR_ANGLE_PRIMARY.

The CFD mesh was generated again in BOXER, where a periodic sector was created from the initial geometry, as illustrated in Fig. 12(a). The produced CFD domain of the baseline geometry, shown in Fig. 12(b), has around 9 million nodes.

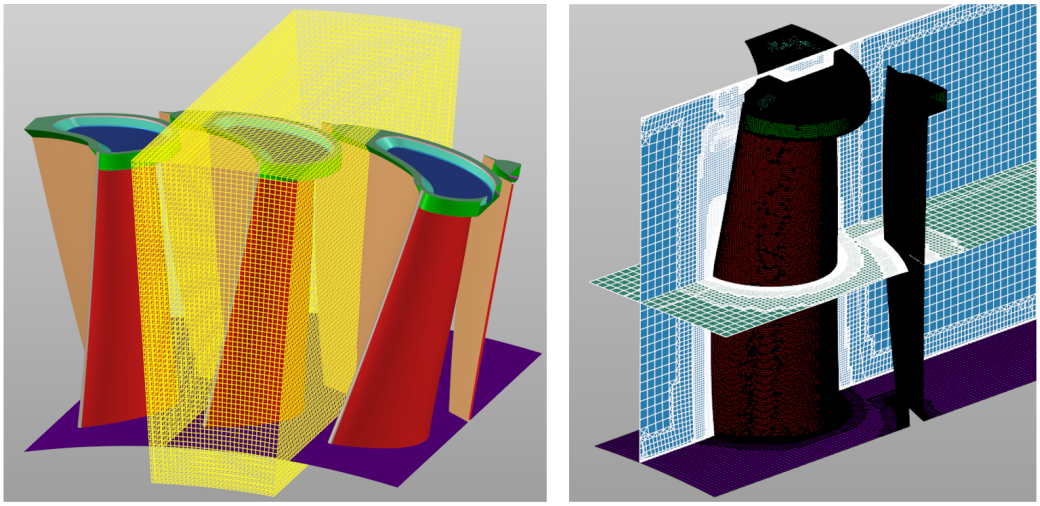

(a) (b)

Figure 12: BOXER mesh generation: a) Periodic section of initial geometry and b) Rotor CFD domain.

The primal and adjoint solutions were obtained from the HYDRA solvers, by using the steady RANS solver with a one equation SA turbulence model and wall functions. The converged flow for the baseline geometry is shown in Fig. 13(a) and the corresponding adjoint sensitivity map is illustrated in Fig. 13(b). The main sources of total pressure loss are the rotor tip vortices.

As in the previous test case, a finite difference (FD) study was also conducted. More specifically, 14 perturbed geometries were exported from Siemens NX and a Python script was used to execute an automated meshing (BOXER) and solving (HYDRA) procedure.

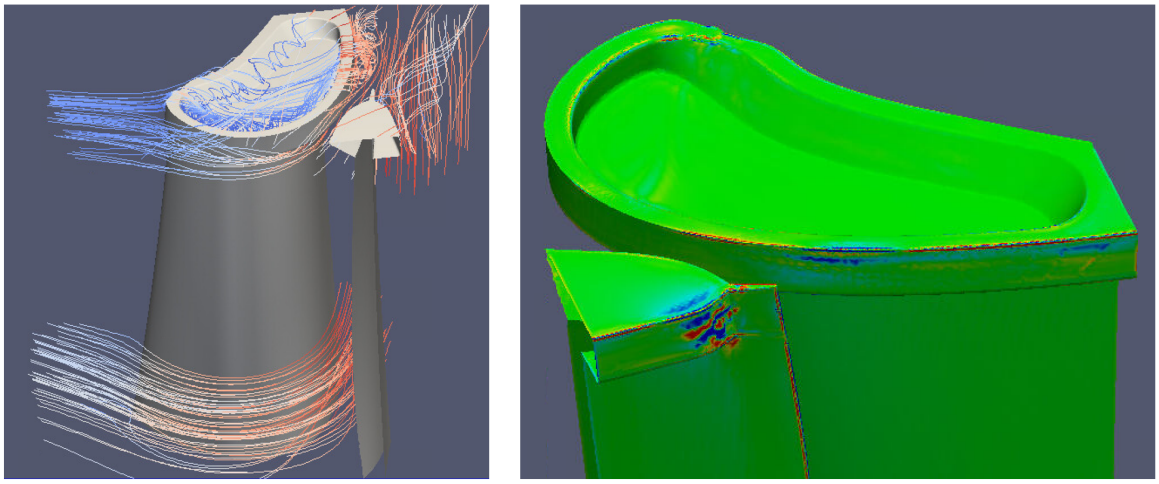

(a) (b)

Figure 13: Flow and adjoint solution: a) Streamlines coloured by relative velocity magnitude and b) Sensitivity map focused on winglet.

The derivatives of the pressure losses w.r.t. each parameter, predicted by linking the design velocities with the sensitivity map, are compared with the corresponding central FD values in Fig. 14. Since the results computed by the adjoint approach are again similar to the FD, further confidence is gained that the proposed method works as expected. The CPU cost of the design velocity tool was around 48 min, which is again negligible compared to the primal and adjoint solution.

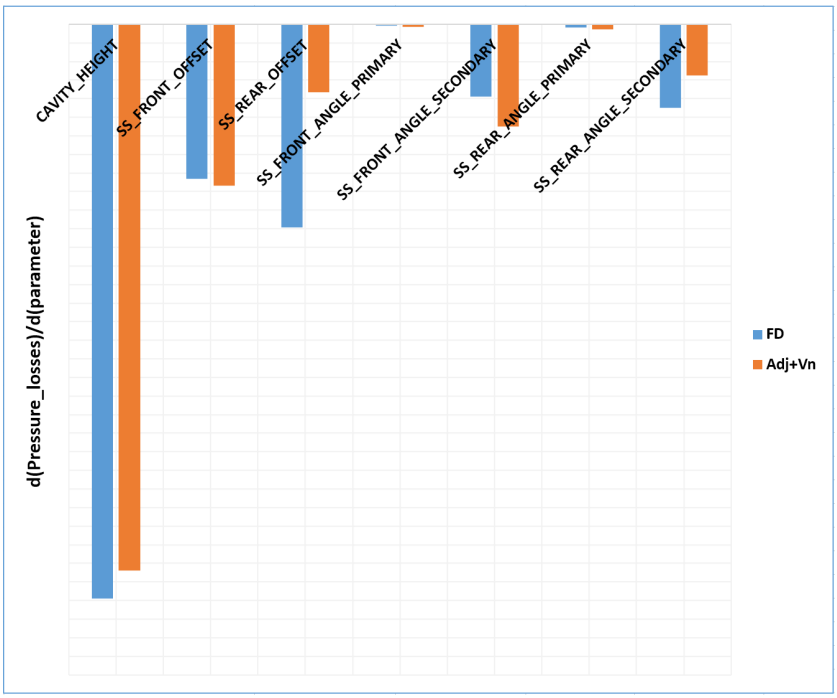

Figure 14: Pressure losses derivatives predicted by adjoint results (with orange) and validation against FD (with blue).

5 CONCLUSIONS

This paper has presented an efficient and robust method for calculating derivatives w.r.t. CAD parameters, based on link of adjoint surface sensitivities with geometric sensitivities, called design velocities, computed from a parametric CAD model. Design velocities are obtained from finite differences between the CAD model before and after a parameter perturbation. The method has been applied to two turbomachinery test cases and the gradient results have been shown to match those computed using finite differences.

Possible sources of error causing the remaining differences between the adjoint derivatives and the FD could be: a) the design velocities, b) the finite differences and c) the adjoint sensitivity map. The design velocity fields have been extensively checked for both cases and they were found to be smooth and without spurious design velocities caused by wrong projections. The FD cannot be considered to be exact, especially for the Rotor case where there was no refinement study carried out. The adjoint sensitivity map is also contributing to the differences, due to the noise that characterizes the surface sensitivities. However, since the sign of the derivatives is correctly calculated, the approach can already be used for steepest descent optimization.

In terms of future work, the possibility of improving the sensitivity map for BOXER meshes will be examined. Then, the work will be focused on passing the gradient information to an optimization algorithm, in order to create a fully automated CAD-based aerodynamic optimization system.

ACKNOWLEDGMENTS

This work has been conducted within the IODA project (http://ioda.sems.qmul.ac.uk), funded by the European Union HORIZON 2020 Framework Programme for Research and Innovation under Grant Agreement No. 642959. The authors would like to thank Rolls-Royce Deutschland for the permission to publish the work.

REFERENCES

[1] M.B. Giles, N.A. Pierce, *An introduction to the adjoint approach to design*. Flow, turbulence and combustion **65:** 393–415, 2000.

[2] A. Jameson, *Aerodynamic shape optimization using the adjoint method*. Lectures at the Von Karman Institute, Brussels, Belgium, 2003.

[3] M.B. Giles, M.C. Duta, J.D. Muller, N.A. Pierce, *Algorithm developments for discrete adjoint methods*. AIAA Journal **41:** 198–205, 2003.

[4] D.I. Papadimitriou, K.C. Giannakoglou, *Compressor blade optimization using a continuous adjoint formulation*. ASME Turbo Expo GT2006/90466, Barcelona, Spain, May 8-11, 2006.

[5] C. Othmer, *A continuous adjoint formulation for the computation of topological and surface sensitivities of ducted flows*. International Journal for Numerical Methods in Fluids **58:** 861–877, 2008.

[6] D.M. Fudge, D.W. Zingg, R. Haimes, *A CAD-free and a CAD-based geometry control system for aerodynamic shape optimization*. AIAA Paper 2005-0451, 2005.

[7] J.F. Dannenhoffer, R. Haimes, *Design sensitivity calculations directly on CAD-based geometry*. AIAA Paper 2015-1370, 2015.

[8] S. Xu, W. Jahn, J.D. Muller, *CAD-based shape optimisation with CFD using a discrete adjoint*. International Journal for Numerical Methods in Fluids **74:** 153–168, 2014.

[9] S. Xu, D. Radford, M. Meyer, J.D. Muller, *CAD-based adjoint shape optimisation of a one-stage turbine with geometric constraints*. ASME Turbo Expo GT2015-42237, Montreal, Canada, June 15-19, 2015.

[10] K.K. Choi, Kuang-Hua Chang, *A study of design velocity field computation for shape optimal design*. Finite Elements in Analysis and Design **15:** 317–341, 1994.

[11] E. Hardee, Kuang-Hua Chang, J. Tu, K.K. Choi, I. Grindeanu, X. Yu, *A CAD-based design parameterization for shape optimization of elastic solids*. Advances in Engineering Software **30:** 185–199, 1999.

[12] J. Kripac, *A mechanism for persistently naming topological entities in history-based parametric solid models*. Computer-Aided Design **29:** 113–122, 1997.

[13] J. Chen, M. Freytag, V. Shapiro, *Shape sensitivity of constructively represented geometric models*. Computer Aided Geometric Design **25:** 470–488, 2008.

[14] T.T. Robinson, C.G. Armstrong, H.S. Chua, C. Othmer, T. Grahs, *Optimizing parameterized CAD geometries using sensitivities based on adjoint functions*. Computer-Aided Design and Applications **9:** 253–268, 2012.

[15] P.M. Thompson, T.T. Robinson, C.G. Armstrong, *Efficient CAD-based aerodynamic design optimization with adjoint CFD data*. 21st AIAA Computational Fluid Dynamics Conference, San Diego, California, United States, 2013.

[16] A. Griewank, G.F. Corliss, *Automatic differentiation of algorithms: theory, implementation, and application*. Defense Technical Information Center, 1992.

[17] K.K. Choi, N.H. Kim, *Structural sensitivity analysis and optimization 1: Linear systems*. Springer, 2005.

[18] S. Shahpar, L. Lapworth, *Padram: parametric design and rapid meshing system for turbomachinery optimisation*. ASME Turbo Expo GT2003-38698, collocated with the 2003 International Joint Power Generation Conference, Atlanta, Georgia, USA, June 16-19, 2003.

[19] S. Shahpar, S. Caloni, *Aerodynamic optimization of high-pressure turbines for lean-burn combustion system*. Journal of Engineering for Gas Turbines and Power **135:** 055001, 2013.

[20] M.B. Giles, *On the iterative solution of adjoint equations*. Automatic Differentiation of Algorithms: 145–151, 2002.

[21] P. Moinier, J.D. Muller, M.B. Giles, *Edge-based multigrid and preconditioning for hybrid grids*. AIAA Journal **40:** 1954–1960, 2002.

[22] J.H. Friedman, J.L. Bentley, R.A. Finkel, *An algorithm for finding best matches in logarithmic expected time*. ACM Transactions on Mathematical Software **3:** 209–226, 1977.

[23] C. Ericson, *Real-time collision detection*. Morgan Kaufmann Publishers, 2005.

[24] L. Hogner, M. Meyer, A. Nasuf, P. Voigt, M. Voigt, K. Vogeler, C. Berridge, F. Goenaga, *Analysis of high pressure turbine nozzle guide vanes considering geometric variations*. In Proceedings of ASME Turbo Expo GT2016-57502, Seoul, South Korea, June 13-17, 2016.

[25] Z. Schabowski, H. Hodson, *The reduction of over tip leakage loss in unshrouded axial turbines using winglets and squealers*. ASME Turbo Expo GT2007-27623, Montreal, Canada, May 14-17, 2007.

[26] M.W. Finnegan, *Shape sensitivity of a turbine blade*. Master's thesis, School of Mechanical and Aerospace Engineering, Queen's University Belfast, 2012.

[27] OpenCASCADE: Open CASCADE Technology, 3D modeling & numerical simulation, 2012. [http://www.opencascade.com]

[28] C. Geuzaine, J.F. Remacle, *Gmsh: a three-dimensional finite element mesh generator with built-in pre- and post-processing facilities*. International Journal for Numerical Methods in Engineering **79:** 1309–1331, 2009.

[29] CADfix: [http://www.iti-global.com/cadfix]