AUTOMATIC DESIGN OF ORC TURBINE PROFILES USING EVOLUTIONARY ALGORITHMS

Pablo Rodriguez-Fernandez, Giacomo Persico

Laboratorio di Fluidodinamica delle Macchine,
Politecnico di Milano,
Via Lambruschini 4, I-20156 Milano, Italy
e-mail: giacomo.persico@polimi.it
Web page: http://www.lfm.polimi.it

ABSTRACT

In this paper, an automated design tool for Organic Rankine Cycle (ORC) turbines is presented. Supersonic flows and real-gas effects featuring ORC turbines complicate significantly their aerodynamic design, which may benefit significantly from the application of systematic optimization methods. This study proposes a complete method to perform shape optimization of ORC turbine blades, constructed as a combination of a generalized geometrical parametrization technique, a high-fidelity Computational Fluid Dynamic (CFD) solver (including real gas and turbulence models) and an evolutionary algorithm. As a result, a non-intrusive tool, with no need for gradients definition, is developed. The high computational burden typical of evolutionary methods is here tackled by the use of a surrogate-based optimization strategy, for which a Gaussian model is applied. Application to ORC turbines has been proved to be successful, resulting in a comprehensive method for a very wide range of applications. In particular, the present optimization scheme has been applied to the re-design of the supersonic nozzle of an axial-flow turbine. In this design exercise very strong shocks are generated in the rear blade suction side and shock-boundary layer interaction mechanisms occur. Optimization aiming at a more uniform flow at the blade outlet section is shown to minimize the shock losses, resulting in a significant improvement in the nozzle efficiency. The optimal configuration determined with the present design tool is also successfully validated against the outcome of a previous optimization performed with a gradient-based method, demonstrating the reliability and the potential of the design methodology here proposed.

1. INTRODUCTION

Thanks to the progressive increase of computational capability, optimization techniques based on high-fidelity flow models play a key role in the present-day design process of turbomachinery. The turbomachinery design process offers optimization challenges at different fidelity levels, from the zero-dimensional stage-by-stage definition (Pini et al. (2013)), to the axisymmetric design (Larocca (2008); Pasquale et al. (2014)), up to the detailed blade shape definition (Verstraete et al. (2010); Pini et al. (2014)).

In the last decades several CFD-based shape optimization procedures were specifically developed in Aerodynamics, such as inverse design methods (Demeulenaere et al. (1997)), adjoint-based gradient methods (Peter and Dwight (2010)), or evolutionary algorithms (Coello (2000)). These latter techniques allow to explore a wider range of feasible solutions, identifying the best individual, and also allow to handle multi-objective optimization problems (Pierret et al. (2006)).

In this paper, a novel optimization package is presented, based on evolutionary algorithms, specifically oriented to the design of turbomachinery blades. Thanks to the high-fidelity flow model, which includes turbulence models and a generalized thermodynamic treatment of the working fluid, the method is particularly attractive for ORC turbines, that feature severe supersonic flows and strong real-gas effects. Several optimization algorithms are tested and the application of the design tool to a highly complex

ORC turbine indicates that dramatic performance improvement is achievable by systematic application of the proposed optimization method.

The paper is structured as follows: in Section 2, the methodology behind the shape optimization tool is described in detail; in Section 3, the different optimization strategies applied and the environment in which they are implemented are discussed. Section 4 finally reports the results of the application to a supersonic ORC turbine nozzle, considering the effectiveness of optimization as well as its implications on the cascade aerodynamics.

2. METHODOLOGY

The optimization strategy here presented is constructed by combining four main blocks, namely a geometry parameterization code, a high-fidelity CFD solver, a library of evolutionary algorithms, and a meta-model interpolation tool. All these items are discussed in detail in the present Section. In particular, Subsection 2.1 describes how B-Spline curves are used to generate the blade geometry; in Subsection 2.2, the CFD model employed for the present high-fidelity calculations is summarized; in Subsection 2.3, genetic algorithms and evolutionary strategies are presented; finally, in Subsection 2.4 the surrogate model used in this study is defined.

2.1 Geometry Parameterization

A key feature of shape-optimization problems in aerodynamics is the technique used to reconstruct the shape of profiles by employing a minimum number of variables. In this work, B-Spline curves are used to parameterize the blade geometry, as the shape can be easily described by a certain number of so-called control points. Thanks to this and other features, B-Spline curves are presently recognized as a powerful tool in both application and theory for aerodynamic designs (Farin (2002)).

A B-Spline can be defined as a piecewise curve with components of degree n that provide local support and whose smoothness and continuity can be adjusted. Thus, a B-Spline curve can be described as a weighted sum of basis functions as follows:

$$x(u) = \sum_{j=0}^{L} d_j N_j^n(u) \tag{1}$$

where $\{d_j\}$, with $j = 0, \ldots, L$, are the control points, and $N_j^n(u)$ are the corresponding n-degree B-Spline bases. These can be defined recursively in the form:

$$\begin{aligned} N_j^k(u) &= \frac{u - u_{j-1}}{u_{j+k-1} - u_{j-1}} N_j^{k-1}(u) + \frac{u_{j+k} - u}{u_{j+k} - u_j} N_{j+1}^{k-1}(u) \\ N_j^0(u) &= \begin{cases} 0 & if \quad u_{j-1} \leq u < u_i, \\ 1 & otherwise \end{cases} \end{aligned} \tag{2}$$

where $\{u_j\}$, with $j = 0, \ldots, K$, is the knot sequence and u is a parameter. Notice that $K = L + n - 1$.

At this point, a first algorithm can be developed to generate a B-Spline curve from a given set of control points, as it is depicted in Figure 1.

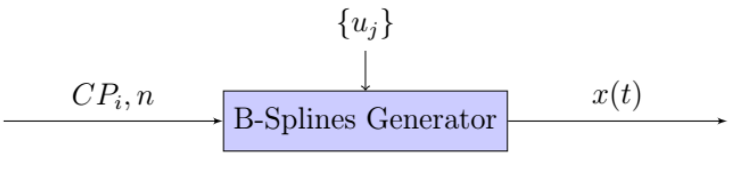

Figure 1: Geometry generation algorithm. CP_i **and** n **are the control points and the degree of the curve, respectively;** $\{u_j\}$ **is the knot sequence. As output, the B-Spline curve** $x(t)$ **is generated.**

The optimization tool described in this study is able to build an approximate representation of the shape of the blade using a B-Spline curve, defined and manipulated by the position of the control points. In order to generate this curve, the baseline geometry needs to be interpolated. In this work, a least squares interpolation method is used.

It is assumed that $P + 1$ data points p_i are given, with $i = 0, \dots, P$, and we seek to find the approximated B-Spline curve $x(u)$ of degree n and $K+1$ knots u_k, with $k = 0, \dots, K$. This B-Spline curve will be defined by $L + 1$ control points d_j, with $j = 0, \dots, L$, such that $L = K - n + 1$. The error of the approximation for a given point can be expressed as $\|p_i - p(w_i)\|$, where w_i, with $i = 0, \dots, P$, are the data parameters of the problem. Therefore, the objective is to minimize the total approximation error:

$$f(x) = \sum_{i=0}^{P} \|p_i - x(w_i)\| \tag{3}$$

If we rewrite Equation (3) using Equation (1) and perform a least squares minimization process, a final expression for the $L + 1$ normal equations is derived:

$$\sum_{j=0}^{L} d_j \sum_{i=0}^{P} N_j^n(w_i) N_k^n(w_i) = \sum_{i=0}^{P} p_i N_k^n(w_i); \tag{4}$$

This equation leads to a linear system $A \cdot x = B$ that is solved, in this work, by using Cholesky Decomposition. At this point, a new algorithm can be defined, to find the control points that best represent a given geometry, as it is depicted in Figure 2.

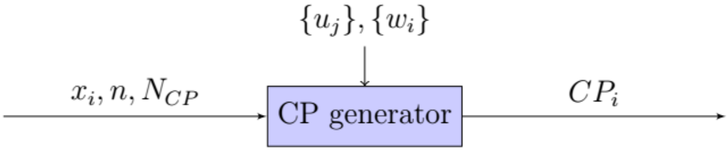

Figure 2: Interpolation algorithm. x_i**,** n **and** N_{CP} **are the data points, the degree of the B-Spline curve and the desired number of control points, respectively.** $\{u_j\}$ **and** $\{w_j\}$ **are the knot sequence and the data parameters. As output, the control points** CP_i **of the B-Spline are generated.**

The sequence of data parameters $\{w_i\}$ is built using a centripetal parameterization. With regard to the knot sequence, it can control the shape of the final interpolated B-Spline. It divides the curve into segments, defined on a knot span. In this work, the knot sequence is used to regulate the position of the control points that serve as design variables of the optimization problem. In this way, one can identify areas prone to optimization and thus assign a different number of control points for each region during the interpolation process.

In this study, both pressure and suction sides of the blade are generated with a unique B-Spline curve. However, the trailing edge will be considered separately, as it is best represented by a circular arc. As a result, the B-Spline curve is constrained to pass through the first and last data points to create a closed curve along with the circular-shaped trailing edge. Continuity and regularity are ensured by imposing same derivative. Additionally, no weights for the control points are used during the interpolation.

2.2 CFD Codes

The present optimization strategy makes use of high-fidelity numerical simulations of the selected blade configurations, performed applying a fully-turbulent and real-gas CFD model based on the ANSYS-CFX solver. As only blade-to-blade effects are of interest, quasi-3D simulations are carried out, using 2D profiles generated with the geometry parameterization algorithm and considering a straight stream-tube around midspan. The effects of turbulence are introduced by resorting to the $\kappa - \omega$ SST model, providing a proper clustering of the cells in the near-wall region so to ensure y^+ below unity along the blade. The real-gas thermodynamic behavior of the working fluid is treated by means of a Look-up-Table (LuT) interpolation method; the LuT was constructed in primitive variables (P, T) by resorting to the Span-Wagner equations of state implemented in the FluidProp database (Colonna et al. (2008)). Total conditions, flow angles, and turbulence quantities are assigned at the inlet, and static pressure is given at the outlet. The onset of spurious pressure wave reflections from the downstream boundary is avoided by placing the outflow boundary three axial chords away from the trailing edge. High resolution numerical schemes are used for advective and diffusive fluxes.

2.3 Genetic Algorithms

The interest in single and multi-objective optimization has grown dramatically in the last decades, thanks to the progressive increase of computational power. In engineering, this has led to the development of several methods which apply the concepts of optimization to support the design of components. Thus, design-oriented optimization can be pursued by applying inverse design methods, gradient-based methods, or heuristic methods. For aerodynamic design purposes, the application of heuristic methods is particularly attractive, as it can be performed by using direct calculation models (such as CFD).

Within the class of heuristic methods, Evolutionary Algorithms (EAs) have become very interesting in a wide range of applications, due to many advantages that make them outperform other optimization methods. Among them, Genetic Algorithms (GAs) have the possibility of dealing with oscillating or smooth-less objective functions; they also allow to introduce constraints in a relatively easy way, and to treat multi-objective optimization as well. Furthermore, genetic algorithms are global optimization methods and, hence, are best suited for optimization problems with multiple local optima, for which gradient methods are too computationally expensive or not readily available. In many optimization problems, GAs quickly identify promising regions of the design space where the global optimum might be located. The interested reader is invited to consult Reeves and Rowe (2002) for a complete description of different GA approaches.

However, the flexibility and simplification provided by GAs are achieved through a massive application of the direct computational model of interest, and thousands of evaluations are usually needed to identify the optimum. Shape-optimization problems in aerodynamics require the application of CFD models, which are expensive tools requiring at least some minutes of calculation to achieve convergence (a reliable optimization requires, in general, fully turbulent flow models and an appropriate grid resolution). As a consequence the direct application of CFD-based genetic optimization is usually not feasible for aerodynamic design purposes. To tackle this unacceptable computational cost, surrogate models can be used. Surrogate models, also known as meta-models or response surfaces, are analytical functions that relate the design variables with performance (i.e., the objective function) in an approximate way. A mathematical representation is selected for the objective function with no relation with the physical phenomena of the real problem (namely, with the CFD model). As it will be discussed later, the mathematical model is trained along the process, resulting in a dramatic reduction of computational burden, as the genetic algorithm is applied directly to the meta-model.

2.4 Surrogate Models

An extensive theory about surrogate models has been developed and many schemes are currently available (the interested reader is referred to Simpson et al. (2001) for a review of available techniques). In this study, the Kriging model is used as the mathematical approximate objective function or non-linear constraint. The Kriging technique is based on a set of interpolation methods, sometimes called Gaussian Processes, originally developed for geostatistic problems and nowadays widely used in many engineering fields. The mathematical model of Kriging can be understood as the linear combination of a trend function and the implementation of a stochastic process. The most common form of a Kriging model is as follows:

$$\tilde{f}(x) = g(x)^T\beta + r(x)^T R^{-1}(f - G\beta) \tag{5}$$

where x is the point to be evaluated; $g(x)^T\beta$ is the trend function evaluated at x, whose coefficients are estimated using a least squares approach; $r(x)$ is the correlation vector with the data points; R is the correlation matrix for all of the data points; f is the response values vector; and $G\beta$ is a vector that contains the trend function evaluated at all data points. The terms in the correlation vector ($r(x)$) and correlation matrix (R) are computed using a Maximum Likelihood Estimation (MLE) procedure.

The implementation of this meta-model in the evolutionary optimization strategy is described in Section 3.

3. IMPLEMENTATION

In this study the optimization problem has been assembled in the object-oriented framework Dakota (Adams et al. (2013)). Dakota provides optimization algorithms, i.e. single-objective and multi-objective genetic algorithms, as well as surrogate models and optimization strategies.

To perform the optimization using genetic algorithms, the JEGA library is used. JEGA (Java Engine for Genetic Algorithms) is a framework that provides a flexible and extensible optimization environment for computational models. Different optimization approaches have been considered and tested to investigate the performance and eventually improve the automatic design tool developed in this research. The initial database is built using Latin Hypercube sampling technique. Usually the size of the population is taken as 2 to 4 times the number of design variables. For example in the present case, a population of 50 individuals is chosen, in line with the 16 design variables used. The probability of crossover and mutation are 80% and 20%, respectively. Elitism is used as selection technique.

To determine the Kriging parameters, the Surfpack library is employed. In this case, it uses a global search method called *DiRect* algorithm (dividing rectangles), a derivative-free global optimization method that balances search in promising and unexplored regions. The trend function is built using a reduced quadratic expression. When working with non-linear constraints, a Kriging model is also built as a mathematical expression for each of them.

Regarding optimization strategies, local and global schemes are tested. In **surrogate-based local optimization** (SBLO), also called Trust Region technique, the optimization algorithm operates directly on a surrogate model that is built using an initial database composed by a certain number of individuals. However, the surrogate model has a limited reliability (especially at the beginning of the process), and hence the fidelity of the approximation is assessed by comparing with the high-fidelity expensive tool (by running the CFD solver). The main feature of the local optimization is the use of a trust region approach, which defines the extent of the approximation. SBLO method needs to generate and update the data fit in each trust region, performing high-fidelity evaluations over a design of experiments. Although a local approach, each sampling in the trust region can be performed globally, which allows to extract relevant global design trends. The comparison between the surrogate and the high-fidelity evaluations, formulated as a trust region ratio, is used to define the step acceptance and trust region size and position of the next iteration.

On the other hand, in a **surrogate-based global optimization** (SBGO), the algorithm is not supported by a trust-region approach. It starts from an initial sample of points and the optimizer operates on that surrogate, by updating its parameters after a new optimum is found and added to the sample. This approach should be used carefully, as there is no guarantee of convergence. It should be used either when there exists the need of using an initial database or when the surrogate needs, somehow, to be updated globally. The surrogate becomes more accurate as the iterations progress. In the present study, both global and local schemes are compared and tested.

4. RESULTS

An exercise of shape-optimization is discussed in the following to demonstrate the capabilities of the design tool presented in this paper. To show the flexibility and the wide range of application of the technique, a particularly severe test-case is considered, namely a converging-diverging supersonic nozzle for an axial-flow ORC turbine operating with MDM. The original blade geometry, shown in Figure 3 and called Baseline in the following, was designed by means of the method of characteristics (MOC) for the diverging part and features a highly smooth leading edge to reduce the sensitivity to incidence variations. The optimization process aims at maximizing the performance of the cascade operating with an expansion ratio of 8 starting from a superheated condition ($P_{T,in} = 8bar$, $T_{T,in} = 272C$) close to the saturation line. As a result, supersonic flows are induced (in fact, the cascade-exit Mach number exceeds 2) and strong real-gas effects occur in the expansion, justifying the use of a LuT approach for the thermodynamic modeling of the fluid.

The blade geometry has been first interpolated and parametrized using the method described in Subsection 2.1. Once the right number of control points is established, their relative position can be adjusted by modifying the knot sequence. In this work, 30 control points have been found to be sufficient to provide an accurate interpolation. From the complete set, the vertical positions (y-coordinate) of 16 control points define the set of design variables for the optimization problem (see Figure 3). The leading edge and the front part (roughly up to the throat) are kept fixed during the optimization process and a larger number of control points is chosen in these regions. The trailing edge width is kept constant so to guarantee the structural resistance of the blade, and therefore only one control point is needed to determine the location of the trailing edge.

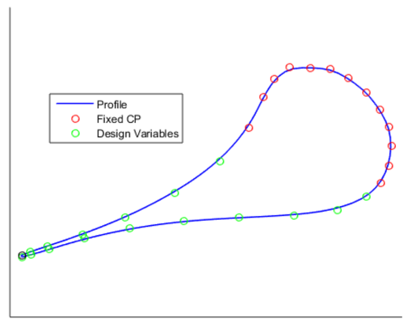

Figure 3: Baseline profile shape and control points distribution. Green circles indicate design variables, while red ones are kept fixed during optimization. The black circle indicates de control point that moves accordingly to keep the trailing edge width constant.

High-fidelity calculations were performed on structured grids composed by 400,000 hexahedral elements. The reliability of the numerical model used in this context was previously assessed against experiments performed by the authors themselves on a research turbine stage installed at Politecnico di Milano (Persico et al. (2012)). The CFD model was shown to accurately predict the fully three-dimensional and unsteady flow physics of the whole turbine stage, and provided estimates of stage efficiency within 1% of the experimental datum, i.e. comparable to the uncertainty of the measurement technique.

In the following, several optimization tests are discussed, with the aim of investigating the impact of different approaches on both the computational cost and the fitness of the design outcome. Different surrogate strategies are considered, also in comparison to gradient-based optimization techniques; it is shown how the proposed automatic design tool provides, in addition to the specific optimized geometry, some intuitive design guidelines for supersonic ORC turbines.

4.1 Impact of the optimization strategy

At first, the comparison between the Global (training) and Local (trust regions) optimization strategies is performed. For both cases the same objective function is used, defined as the standard deviation of the azimuthal pressure distribution evaluated half axial chord downstream the blade trailing edge. The minimization of this quantity in a supersonic cascade is expected to reduce the shock strength, thus increasing the cascade performance.

For the construction of the Kriging surrogate model, a minimum sample size of 5 times the number of design variables is commonly considered. In this study, 16 design variables are set to optimize the geometry, and hence a database of 80 individuals is, in all cases, considered. In the global approach,

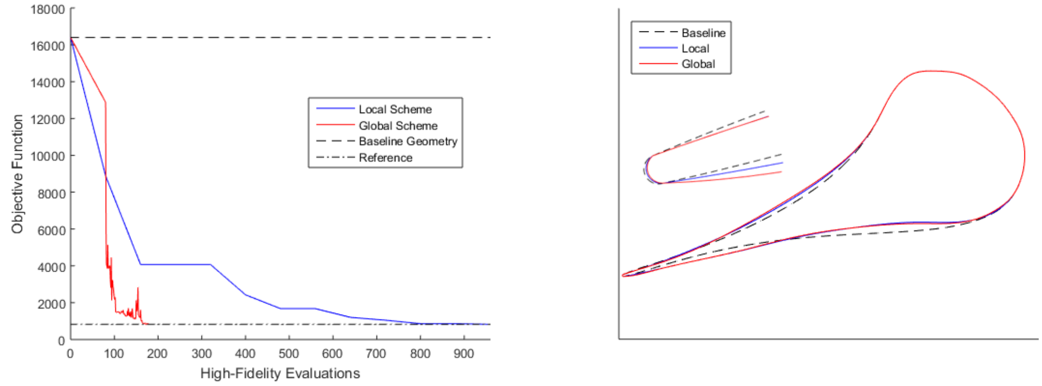

Figure 4: Left: convergence history of global and local strategies for surrogate-based optimization. Right: optimized geometry using global and local strategies for surrogate-based optimization.

this means that the code builds the initial Kriging model over this initial sample. After a complete GA-optimization applied only to the mathematical model, the result is assessed via the high-fidelity tool and added to the population (updating the Kriging parameters in each iteration). On the other hand, the local approach builds a new Kriging model after each GA-optimization. These different features are clearly visible in the left frame of Figure 4, where the paths towards optimization are compared for the two methods; in particular, it is observed how the local scheme advances only each 80 high-fidelity evaluations, while the global method advances continuously after the initial 80 iterations. As a result the local approach shows a much slower trend, even though the local and global schemes obtain a very similar minimization of the objective function, significantly reduced with respect to the baseline configuration. For this reason, in the following only the global scheme will be used.

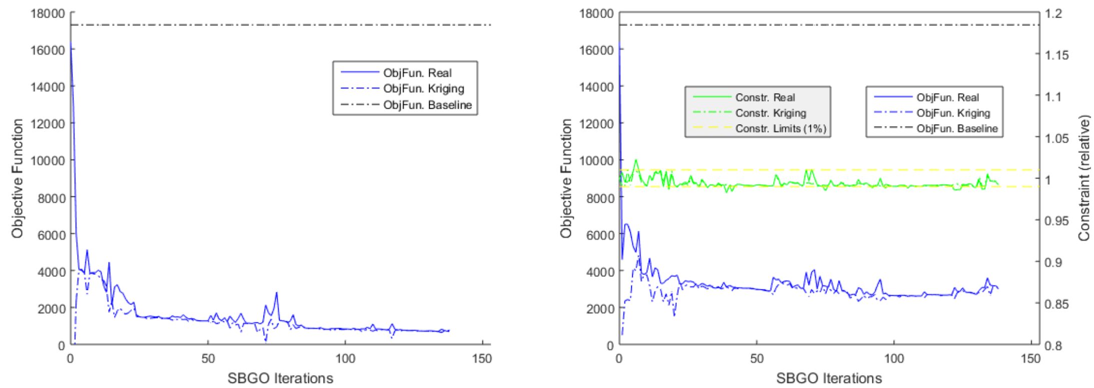

Figure 5: Convergence history of Kriging model for (left) non-constrained, and (right) constrained optimization, using a global strategy for surrogate-based optimization.

When performing evolutionary optimization using surrogate strategies, the convergence of the approximate model to the actual response surface of the problem needs to be verified. The left frame of Figure 5 shows that the surrogate model quickly matches the high-fidelity tendency, except for some spikes of progressively reduced amplitude as the algorithm converges to the optimum. This allows to conclude that the training procedure chosen for the Kriging model is appropriate for the present design problem. The smooth convergence trends observed so far have been achieved without the application of explicit constraints to the optimization. However from the engineering perspective it is interesting to investigate how the optimization proceeds when a relevant quantity is constrained; for example, when looking to the turbine cascade performance, the flow rate is usually a fixed parameter whose variation needs to be limited in a narrow range. The right-hand side of Figure 5 shows the results a non-linear constrained optimization, by forcing the flow rate to stay within ±1% with respect to the baseline value. A specific Kriging model is built for the constraint. It is observed how, after a significant reduction of the objective function in the initial phase, the constraint prevents the optimization from progressing further, resulting in a lower fitness of the outcome. It is also very interesting to note that the flow rate of the optimized cascade achieved via non-constrained global optimization is, in fact, within the limits. This is probably because the front part of the blade up to the throat is kept fixed and hence the flow rate is somehow imposed indirectly in the present chocked flow configuration. This seems to be a more effective procedure when dealing with supersonic cascades.

4.2 Aerodynamic analysis

The flow configuration established in the cascade optimized via the non-constrained global strategy is now discussed, in comparison to the baseline configuration and to another optimized case. This latter configuration was obtained by applying an adjoint-based gradient method developed by Pini et al. (2014); it was performed using the same objective function used here and was based on an inviscid flow solver; the result here reported is, however, the high-fidelity calculation of that optimized configuration.

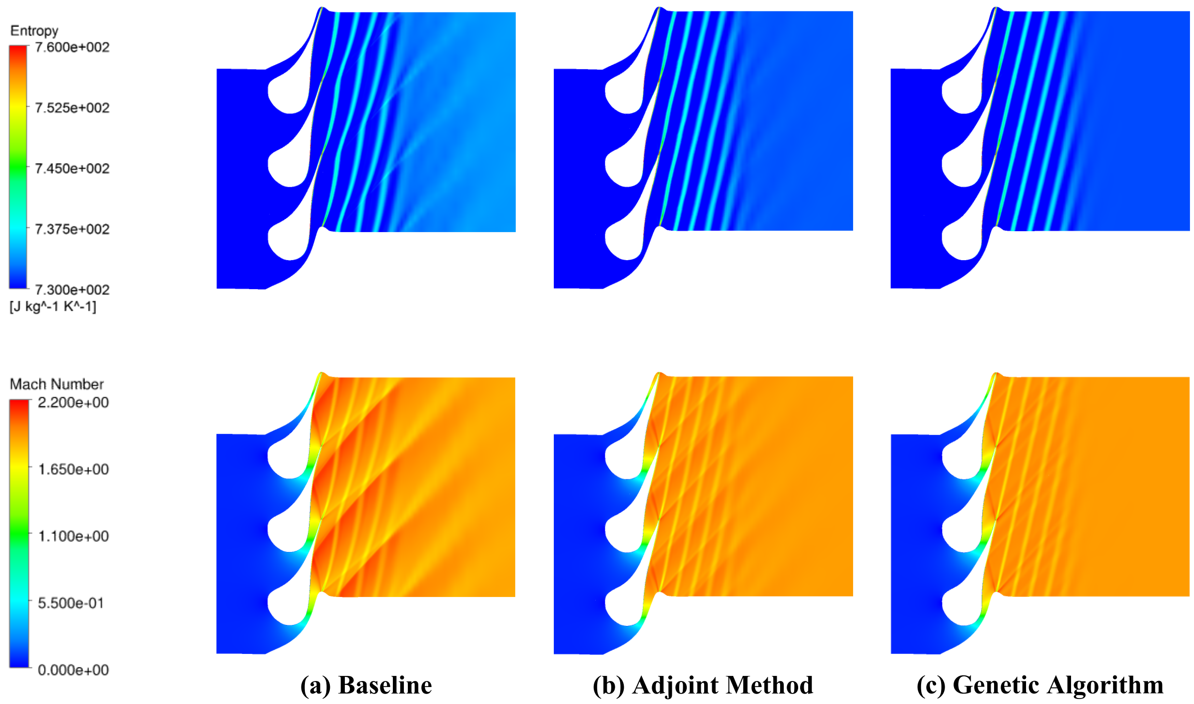

Figure 6: Comparison of the Mach number (bottom) and Entropy (top) distributions for (a) Baseline, (b) Adjoint method (Pini et al. (2014)), and (c) present study using the global strategy.

The entropy and Mach number field on the blade-to-blade surface are reported for the three cases in Figure 6. The effect of the optimization is clearly visible by comparing the optimized and baseline configurations; in particular the minimization of the pressure oscillations downstream of the cascade leads to a dramatic reduction of the main shock strength. Hence, the severe pressure gradient observed in the baseline case is highly weakened and the weaving path of the blade wakes is almost eliminated. As a result, the overall entropy generation is significantly reduced in the optimized cases. The two optimized blades exhibit very similar flow configurations, even though the two approaches make use of very different methods based on different flow models; this suggests that both the methods are converging towards the same optimum, which is probably the global optimum for the present problem. The evolutionary optimization leads, in fact, to a slightly more uniform flow, due to the slightly weaker fishtail shock system at the blade trailing edge; such a shock system is not accurately captured by the inviscid flow model used in the adjoint-based optimization (due to the missing displacement thickness of the boundary layers), and this may explain the slight differences observed. However the outcome of the adjoint optimization still remains an excellent result considering its extremely limited computational cost (less than an hour on a standard PC) with respect to that of the present GA strategy (20 hours on a 15-processor cluster).

Figure 7 reports the pressure distribution on the three blades, in the form of isentropic Mach number, and explains the reason for the improved performance of the optimized configurations. Both optimal blades move ahead the acceleration of the flow on the suction side, just downstream of the sonic throat and still within the bladed channel; as a result the over-speed on the rear suction side is limited and hence the subsequent diffusion is eliminated. In this way the onset of the strong shock observed in the baseline configuration is prevented. Once again the two optimized configurations show very similar trends, with local differences especially close to the trailing edge.

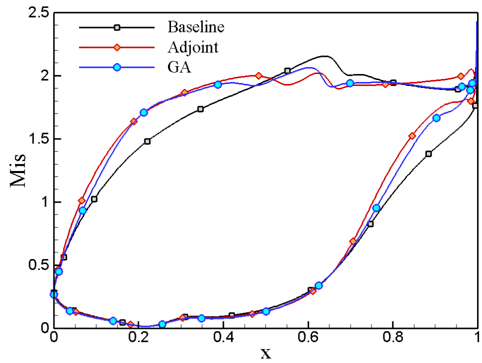

Figure 7: Isentropic Mach Number distribution over the blade for the Baseline configuration, the optimized blade using the Adjoint Method (Pini et al. (2014)) and using genetic algorithms.

To quantify the impact of the optimization on the cascade performance the Total Pressure Loss Coefficient is used, defined as the total pressure loss referred to the exit dynamic pressure ($Y = \frac{P_{T,in} - P_T}{P_{T,in} - P}$). The results, considering an outlet placed at 2 chords downstream the blade (where the flow can be considered mixed-out) are summarized in Table 1. It can be concluded that the more uniform flow achieved in the downstream region leads to a relevant decrease of the total pressure loss coefficient, which reduces from 15.0 to 9.3%.

	Baseline	Adjoint Method	GA
Y	0.15	0.11	0.093

Table 1: Total Pressure Loss Coefficient at two axial chords downstream the blade.

5. CONCLUSIONS

This paper has presented a novel package for the automatic design of ORC turbines based on an evolutionary strategy. Detailed descriptions of all the steps of the optimization scheme have been provided, namely the geometry parametrization, the high-fidelity flow solver and the genetic algorithm.

The blade shape is parametrized via B-Splines, whose local control capability allowed a detailed shape reconstruction while preserving surface smoothness. The implementation of advanced high-fidelity flow models, of paramount importance for ORC turbines, is easily attained thanks to the non-intrusive character of the evolutionary optimization strategy here used. To tackle the computational burden typical of CFD-driven evolutionary strategies, the genetic algorithm is coupled to a surrogate model that reflects the influence of the design variables on the objective function. Several optimization strategies have been discussed to evaluate the convergence process and the associated computational cost.

Application to a supersonic ORC turbine nozzle has demonstrated that relevant performance improvements can be achieved by maximizing flow uniformity downstream the blade. A further comparison, for the same test case, against an alternative optimization approach has assessed the validity of the present design methodology. Results have also allowed to quantify the impact of the application of high-fidelity flow models within the optimization process. Future research will be addressed towards the application of the present shape-optimization tool to the design of novel turbine blade configurations for ORC applications.

REFERENCES

Adams, B., Bauman, L., Bohnhoff, W., Dalbey, K., Ebeida, M., Eddy, J., Eldred, M., Hough, P., Hu, K., Jakeman, J., Swiler, L., and Vigil, D. (2013). Dakota, a multilevel parallel object-oriented framework for design optimization, parameter estimation, uncertainty quantification, and sensitivity analysis: Version 5.4 user's manual. *Sandia Technical Report SAND2010-2183*.

Coello, C. (2000). An updated survey of GA-based multiobjective optimization techniques. *ACM Computing Surveys*, 32(2):109–143.

Colonna, P., Nannan, N., and Guardone, A. (2008). Multiparameter equations of state for siloxanes: [(ch3)3-si-o1/2]2-[o-si-(ch3)2]i=1,...,3, and [o-si-(ch3)2]6. *Fluid Phase Equilibria*, 263(2):115 – 130.

Demeulenaere, A., Leonard, O., and Van Den Braembussche, R. (1997). A two-dimensional Navier-Stokes inverse solver for compressor and turbine blade design. *Journal of Power and Energy*, 211(4):299–307.

Farin, G. (2002). *Curves and Surfaces for CAGD: A Practical Guide*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 5th edition.

Larocca, F. (2008). Multiple objective optimization and inverse design of axial turbomachinery blade. *Journal of Propulsion and Power*, 24(5):1093–1099.

Pasquale, D., Persico, G., and Rebay, S. (2014). Optimization of Turbomachinery Flow Surfaces Applying a CFD-Based Troughflow Method. *Journal of Turbomachinery*, 136(3):031013 1–11.

Persico, G., Mora, A., Gaetani, P., and Savini, M. (2012). Unsteady aerodynamics of a low aspect ratio turbine stage: Modeling issues and flow physics. *Journal of Turbomachinery*, 134(6).

Peter, J. and Dwight, R. (2010). Numerical sensitivity analysis for aerodynamic optimization: A survey of approaches. *Computers & Fluids*, (39):373–391.

Pierret, S., Coelho, R., and Kato, H. (2006). Multidisciplinary and multiple operating points shape optimization of three-dimensional compressor blades. *Structural and Multidisciplinary Optimization*, 33(1):61–70.

Pini, M., Pasquale, D., Persico, G., and Rebay, S. (2014). Adjoint method for shape optimization in real-gas flow applications. *Journal of Engineering for Gas Turbines and Power*, 137:032604.

Pini, M., Persico, G., Casati, E., and Dossena, V. (2013). Preliminary design of a centrifugal turbine for ORC applications. *ASME Journal of Engineering for Gas Turbines and Power*, 135(042312).

Reeves, C. R. and Rowe, J. E. (2002). *Genetic Algorithms: Principles and Perspectives: A Guide to GA Theory*. Kluwer Academic Publishers, Norwell, MA, USA.

Simpson, T., Poplinski, J., Koch, P. N., and Allen, J. (2001). Metamodels for computer-based engineering design: Survey and recommendations. *Engineering with Computers*, 17(2):129–150.

Verstraete, T., Alsalihi, Z., and Van Den Braembussche, R. (2010). Multidisciplinary Optimization of a Radial Compressor for Microgas Turbine Applications. *Journal of Turbomachinery*, 132:031004.