High-Fidelity Shape-Optimization of Non-Conventional Turbomachinery by Surrogate Evolutionary Strategies

Giacomo Persico, Pablo Rodriguez-Fernandez[1], Alessandro Romei

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

ABSTRACT

This paper presents a novel tool for the shape-optimization of turbomachinery blade profiles operating with non-ideal working fluids in complex flow configurations. The non-conventional turbomachinery layouts as well as the non-ideal thermodynamics of the working fluids featuring, for example, Organic Rankine Cycle and supercritical CO_2 power systems complicate significantly the blade aerodynamic design. This class of turbomachinery may considerably benefit from the application of systematic optimization methods, especially if high-fidelity Computational Fluid Dynamic (CFD) models are implemented within the optimization process. The proposed technique is implemented in the package FORMA (Fluid-dynamic OptimizeR for turbo-Machinery Aerofoils) developed in-house at the Politecnico di Milano to perform shape optimization of blade profiles. FORMA is constructed as a combination of a generalized geometrical parametrization technique, a high-fidelity CFD solver and multiple surrogate-based evolutionary strategies. The application to the re-design of a supersonic turbine nozzle shows the capabilities of applying a high-fidelity optimization, consisting in a 50% reduction in cascade loss coefficient and in a much-increased flow uniformity at the inlet of the subsequent rotor. Two alternative surrogate-based evolutionary strategies and different fitness functions are tested and discussed, including non-linear constraints within the

[1] Presently at Plasma Science and Fusion Center, MIT, Cambridge, MA 02139, USA

design process. The optimization study reveals relevant insights on the design of supersonic turbine nozzles as well on the performance, reliability and potential of the proposed design technique.

Keywords: high-fidelity CFD simulations; evolutionary optimization; surrogate models; constrained optimization; non-ideal turbomachinery flows; supersonic turbines

1. INTRODUCTION

The design of turbomachinery for non-conventional or novel applications has often to face new challenges that prevent classical empirical methods from obtaining high performance. At the same time, the combined increase of computational performance, of flow-model reliability, and of algorithm effectiveness has triggered the introduction of high-fidelity optimization techniques within aerodynamic design. Three main design steps compose the path that brings to the geometric definition of a turbomachine, each of them corresponding to a specific level of approximation of the flow and worth for a dedicated optimization: (i) the preliminary design, based on the mean-line model; (ii) the spanwise design, based on the throughflow model; (iii) the aerodynamic design of the blades, based on Computational Fluid Dynamics (CFD). A presentation of the methods and of their combination in a comprehensive design procedure is reported in Persico and Pini, 2016. In an effort of constructing a set of design tools suitable for application to non-conventional turbomachinery, a mean-line code (Pini et al., 2013) and a throughflow code (Persico and Rebay, 2012) have been recently developed at Politecnico di Milano. The work presented in this paper is the logical continuation of those ones, focusing on the blade aerodynamic design, here proposed via a shape-optimization technique.

Shape-optimization in Aerodynamics can be pursued by applying deterministic procedures such as inverse methods (Demeulenaere et al., 1997) or adjoint-based optimization (Peter and Dwight, 2010, Pini et al., 2015, Vitale et al., 2017). These methods are attractive for their rigorousness and their computational effectiveness; nevertheless, they require to work-out the CFD source code (thus preventing the use of well-established and validated commercial codes). Alternative to deterministic techniques, heuristic shape-optimization procedures based on evolutionary methods have been widely

applied to the design of turbomachinery blades (Pierret et al., 2006, Verstraete et al., 2010, Pasquale et al., 2013). These methods exhibit several advantages with respect to the deterministic ones, as they are non-intrusive, in the sense that they only require the availability of the CFD code, thus inheriting its robustness and reliability. Moreover, they allow to scan the entire design space in searching for the best individual, thus resulting less prone to be stuck in local optima with respect to deterministic methods. As a further consideration, evolutionary methods allow to handle constrained and multi-objective optimization problems in a straightforward way (Coello, 2000). Evolutionary methods, however, are affected by scalability issues, due to the huge computational cost resulting from the high number of individuals (i.e., geometric configurations) that must be tested to identify the fittest individual. To tackle this cost, neural networks (Pierret and Van den Braembussche, 1999) and analytical surrogates (Simpson et al., 2001) were proposed as meta-models of the system response to be used within the main optimization algorithms.

This paper presents a surrogate-assisted evolutionary shape-optimization methodology for the high-fidelity optimization of non-conventional turbomachinery blades, implemented in the form of a design package named FORMA (Fluid-dynamic OptimizeR for turbo-Machinery Aerofoils). The optimization technique is constructed as a combination of: (i) a flexible geometric parametrization method based on B-splines; (ii) a high-fidelity and experimentally validated flow model based on CFD, including turbulence effects and a generalized thermodynamic treatment of the working fluid; (iii) multiple robust and effective surrogate-based evolutionary strategies.

By virtue of the aforementioned features, FORMA allows to perform a novel class of high-fidelity optimization for non-conventional turbomachinery adopted in supercritical CO_2 and ORC power systems. This paper presents the application of FORMA to a particularly severe test-case, namely the high-fidelity optimization of a turbine nozzle working with an organic fluid, which combines high supersonic flows with non-ideal gas effects. Thanks to the flexibility of the optimization tool, the paper reports comparative tests between several optimization strategies, shedding light on both the cascade aerodynamics and on the optimization method.

2. The FORMA Package

FORMA is constructed by combining several blocks in multiple optimization strategies. This Section presents in detail the single items as well as their combination.

2.1. Geometry parametrization

As the determination of the optimal geometric shape of turbomachinery blades is the ultimate outcome of FORMA, a key element of the procedure is the way blade profiles are parametrized. In this work, the profile geometry is parametrized by using B-Spline curves. A B-Spline is defined as a piecewise curve with adjustable smoothness and continuity, allowing for a local control of the shape by moving a limited number of so-called Control Points (CPs). More details on the mathematical foundations of B-Splines can be found in Farin, 2002. In this study, the pressure and suction sides of the blade are generated with a single B-Spline curve of order 3, ensuring a second-order continuity in the inner part. The trailing edge (TE) is considered separately and represented by a circular arc.

To set-up the optimization, at first the baseline blade shape is approximated by the so-defined B-Spline curve, whose CPs coordinates are found via a least squares interpolation method. The spatial coordinates of a subset of CPs (hereinafter referred to as movable CPs) become the design variables of the optimization problem. The movable CPs can be displaced either in a specific coordinate direction or along the local orthogonal direction with respect to the blade surface. Three-dimensional blade representation is then achieved by combining several profiles in spanwise direction.

2.2. Computational flow model

The present optimization strategy makes use of high-fidelity CFD simulations of the selected blade configurations, performed applying a flow model featuring turbulence and non-ideal gas models. The flow model is based on the ANSYS-CFX solver. The effects of turbulence are introduced by resorting to the k-ω SST model; to properly calculate the near-wall region of the boundary layer, a proper clustering of the cells is assigned at the blade wall, so to keep the $y^+ < 1$ on the blade. The

non-ideal thermodynamic behavior of the working fluid is treated by resorting to a Look-up-Table (LuT) interpolation method; the LuT is constructed using primitive variables (P, T), by calling an external thermodynamic library, and including tabulated transport properties.

The mathematical discretization of the equations is obtained as follows. The advection term of both the flow and turbulence equations is discretized by using a high-resolution numerical scheme, which exploits a non-linear recipe based on the boundedness principles (Barth and Jesperson, 1989). The numerical algorithm was proved to be total variation diminishing when applied to 1D problems, which is crucial for the proper treatment of supersonic flows and shock waves. Diffusive fluxes are discretized with second-order central differences.

All the calculations performed during and after the optimization process are run on structured grids composed by hexahedral elements. The mesh is re-generated for each CFD run performed throughout the optimization process.

The reliability of the flow model used here was assessed against experiments performed on a cold-flow gas turbine stage installed at Politecnico di Milano (Persico et al., 2012).

2.3. Genetic Algorithm and Surrogate model

To pursue the shape-optimization of turbomachinery blades, FORMA makes use of evolutionary algorithms that drive the geometric design of the blade from a baseline configuration to another one featuring the highest performance, considering the assigned design space, the definition of performance, and the fulfillment of constraints. Within the class of evolutionary methods Genetic Algorithms (GAs) were selected as they are global optimization techniques which allow to deal with oscillating or smooth-less objective functions, introduce non-linear constraints, and treat multi-objective optimization in a relatively easy way (Reeves and Rowe, 2002).

The flexibility and simplification achieved using GAs are counterbalanced by the very high computational cost that they require. As a matter of fact, the construction of several generations and the very high number of genome combinations lead to a massive application of the computational

model of interest. The computational cost of CFD simulations makes the direct application of CFD-based genetic optimization usually unfeasible for design purposes. However, the introduction of surrogates (or meta-models), that map the design variables to the objective function and the constraints in an approximate way, allow drastically reducing the computational cost of the optimization, as the evolutionary algorithm is only applied to the surrogate model.

An extensive theory about surrogate models was developed and many formulations are currently available (Simpson et al., 2001). After preliminary trials using several types of meta-models (also including artificial neural networks, radial basis functions, multi-variate adaptive regression splines), a universal Kriging formulation was selected, with non-zero covariance based on a Gaussian correlation function.

2.4. Surrogate evolutionary strategies

The application of GAs to surrogates introduces reliability issues. The optimization outcome becomes crucially dependent on the reliability of the meta-model, which demands initialization and an active training during the optimization process. This offers the opportunity to set-up surrogate evolutionary strategies, two of them are considered in this work and are now discussed in detail.

The surrogate model has first to be built by interpolation using an initial database of blade shapes tested with the CFD model; this initial database is constructed by resorting to a Design of Experiments (DoE) technique to sample the design space. Then, the GA finds the optimum of the surrogate. Since the surrogate model is constructed with a limited database, its reliability must be improved following either a Local (Surrogate-Based Local Optimization, SBLO) or a Global (Surrogate-Based Global Optimization, SBGO) strategy, whose operation is schematized in Figure 1.

The SBLO method stems from the idea that a local reliability of the surrogate is sufficient to find the optimum (Toropov et al., 1993, Giunta and Eldred, 2000). In practice, SBLO progressively selects different subsets of the original design space, and builds the surrogate functions only over these ones. The subsets are called trust regions, and their extension is modified throughout the optimization

process in order to both find the optimal region and maximize locally the reliability of the surrogate function. To this end, SBLO needs a metric to evaluate this reliability and a criterion to evolve the trust region, starting from the design space. This reliability index, called trust region ratio r^k, is defined as:

$$r^k = \frac{f(x_c^k) - f(x_{opt}^k)}{\hat{f}(x_c^k) - \hat{f}(x_{opt}^k)} \tag{1}$$

where f is the real (CFD-evaluated) objective function, $\hat{f}$ is the surrogate model, x_c^k is the center of the trust region (optimum of iteration k-1, or $x_c^k = x_{opt}^{k-1}$), and x_{opt}^k is the optimal configuration identified by the GA applied to the surrogate interpolated at the iteration k. By virtue of the above definition, r^k measures how reliable is the surrogate model in predicting a fitness improvement.

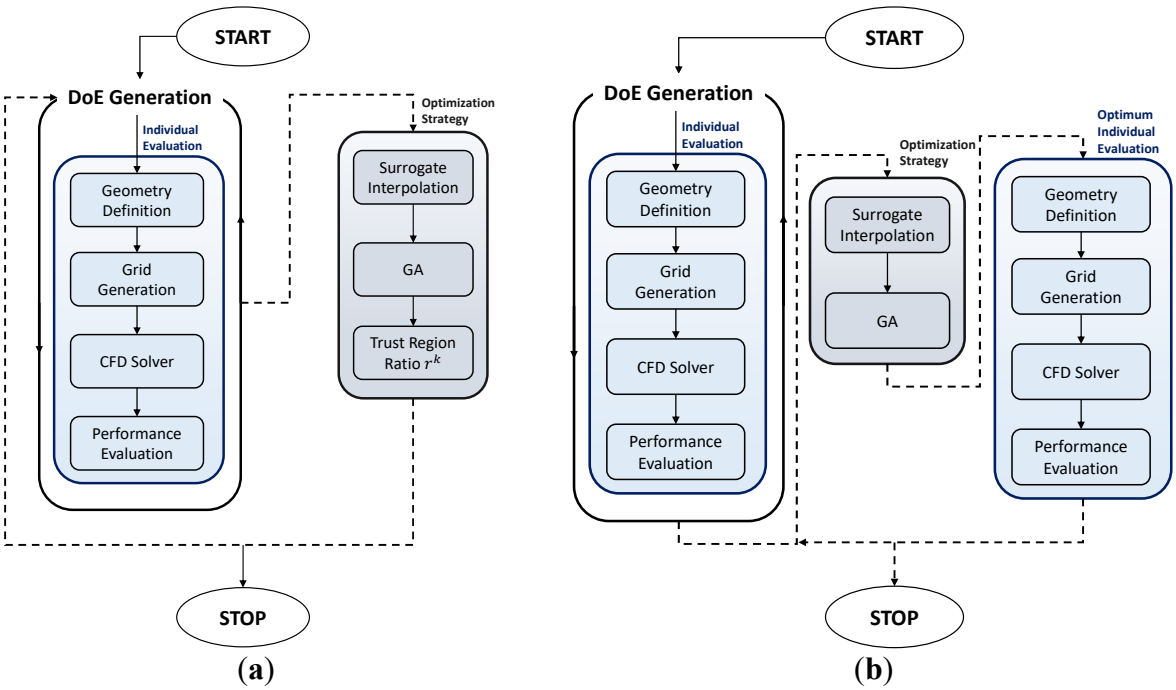

Figure 1. Flow Charts for surrogate-based optimization: (a) SBLO; (b) SBGO

The trust region ratio r^k is used to properly re-define the trust region for the next iteration. Values of r^k very different from the unity (either too low or too high) indicate that the gap between the physical and surrogate models is large, and that the surrogate model needs improvement. This is done by 'centering' the new trust region on the present optimal solution x_{opt}^k, and by shrinking the present trust region by a certain factor (which means to reduce the maximum displacements available for the

movable CPs). Then, as schematically shown in the left frame of Figure 1, a new DoE of CFD runs is performed over the new trust region. The surrogate is re-interpolated, a new optimum is found, a new CFD run is performed on the optimal configuration, and a new trust region ratio is evaluated. By iterating this procedure, the center of the trust region is guided towards the optimum of the real problem, somehow mimicking the deterministic action of the gradient. If the trust region ratio approaches unity, indicating that the surrogate predicts a reliable change in optimal values, the trust region is enlarged, with the aim of extending as much as possible the validity of the surrogate and avoiding overfitting. In the formulation used in FORMA, the trust region is shrunk for $r^k \leq 0.25$ and enlarged for $0.75 \leq r^k \leq 1.25$. To ensure convergence, an elitist filter is applied, so that the optimum point can be accepted only if it is better than any previous evaluations. The convergence criterion is specified as a maximum number of iterations with improvement lower than a certain tolerance.

Alternative to the local technique, a SBGO strategy can be constructed by resorting to the concept of training. The procedure is much simpler and intuitive than the local one. The first steps are identical (initialization, first interpolation, optimization of the surrogate, CFD assessment of the optimal configuration). The CFD result of the optimal configuration is then introduced into the database and used to update the interpolation of the surrogate. The new optimum can either substitute an element present in the sample (so that the size of the database remains constant), or be added to the database. The second choice is attractive as it should lead to concentrate the samples in the region close to the optimum, thus enhancing the reliability of the surrogate in the region of interest.

Both surrogate strategies do not guarantee to find the actual global optimum: even though the GA is theoretically able to identify the global optimum of the surrogate model, this might not correspond to the actual one due to the limited reliability of the surrogate function at the beginning of the process. From this perspective, a proper initialization of the surrogate is crucial for the success of the procedure. Mathematically SBLO has some advantages, as it is more robust and guarantees convergence (by virtue of the elitist filter). However, it requires significant computational efforts, as a DoE is required at each iteration, while SBGO requires the construction of one database only and it

adds one single point (i.e., one CFD evaluation) at each iteration. The implications of these features on the computational cost of the optimization are non-trivial; as a matter of fact, the 'training' iterations of the SBGO need to be serial, while DoE sampling can take great benefit by parallelization, resulting in a very significant scalability of the SBLO if high computational resources are available.

2.5. Constrained formulation

In most practical cases, the aerodynamic design of turbomachinery blades has to cope with limitations coming from the previous steps in the design process (such as the flow rate, the stage reaction degree, etc…). Some of these limitations are automatically imposed by acting on the geometry (i.e., by limiting the design space), or by assigning proper boundary conditions; however, other quantities cannot be assigned directly and have to be constrained in an indirect way.

The introduction of constraints within an evolutionary strategy can be obtained, for example, simply discarding the individuals that do not satisfy the condition. However, such a procedure might severely alter the convergence process and ultimately might prevent the algorithm from exploring the region where the actual optimum is placed. As an alternative, a penalty formulation of the constraints can be used. The degree of unfulfillment of the constraints, expressed as a difference or a ratio between the value of the constrained quantities and the specified thresholds, is added to the objective function after multiplication by a 'penalty' coefficient, which both ensures quantitative homogeneity with the objective function and defines the weight of the constraint. A proper selection of the penalty coefficient makes the optimization converge to a configuration that minimizes both the actual objective and the degree of unfulfillment of the constraints. This penalty formulation of the constraint is used in FORMA; coherently with the surrogate optimization approach, a specific Kriging model is built also for the constraint.

2.6. Implementation

FORMA implements constrained global and local evolutionary strategies, formulated in both single and robust multi-point (Pini et al., 2014) fashion, making use of the Dakota framework (Adams

et al., 2017). In the following, the techniques implemented in FORMA are tested and compared considering a technically relevant and scientifically challenging turbomachinery application.

3. Shape-optimization of a supersonic turbine nozzle working with an organic fluid

With the double aim of investigating the capabilities of the optimization technique and the challenges of optimizing a non-conventional turbomachine, a particularly severe test-case is considered, namely a converging-diverging supersonic nozzle cascade for an ORC axial-flow turbine working with the siloxane MDM as working fluid.

The original blade geometry, shown in Figure 2 and called Baseline in the following, features a large leading edge region that also acts as a converging section upstream of the throat and a diverging section downstream of the throat, this latter designed by applying the Theory of Characteristics. This cascade was studied by Colonna et al., 2008b. The optimization process aims at maximizing the performance of the cascade operating with an expansion ratio of about 7.5 starting from a superheated condition ($P_{T,in}$ = 8 bar, $T_{T,in}$ = 272 °C) close to the saturation line.

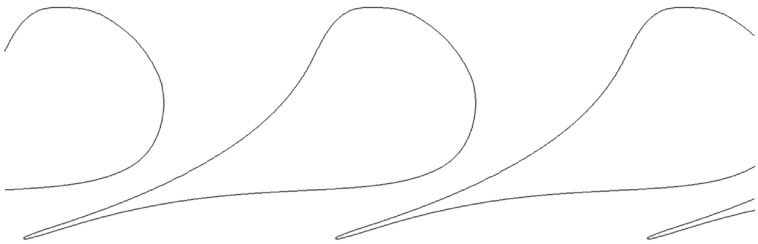

Figure 2. Baseline converging-diverging configuration.

In the conditions of interest, supersonic flows occur in the rear section of the blades, triggering the onset of shock waves and related viscous phenomena such as shock-boundary layer interaction and shock-induced flow separation at the training edge. In such a context, the use of a high-fidelity flow model is crucial to obtain a realistic picture of the cascade aerodynamics.

Moreover, in the prescribed operating conditions the MDM exhibits relevant non-ideal effects; the use of an appropriate thermodynamic model is crucial to obtain an accurate flow representation

and, hence, to exploit the benefits of optimization (as shown by Pini et al., 2015). The fluid thermodynamic behavior was introduced via a LuT constructed by interrogating the RefProp library which implements, for the MDM, the Span-Wagner Equation of State as formulated by Colonna et al., 2008a, and provides dedicated correlations for transport properties.

Total pressure, total temperature and axial flow is prescribed at the cascade inlet, where 5% turbulence intensity is also imposed. At the outflow boundary, the static pressure is imposed; to avoid the onset of spurious pressure wave reflections from the downstream, the outflow boundary is treated in a 'weak' form, by imposing the mean pressure value and accepting an oscillation of 5% around the mean; moreover, the outflow boundary was placed four axial chords away from the trailing edge.

As only blade-to-blade effects are of interest in this design exercise, quasi-3D simulations are carried out, stacking spanwise the same profile generated with the geometry parameterization algorithm, and considering a straight stream-tube around midspan.

Preliminary to the optimization, high-fidelity CFD simulations of the baseline cascade were performed using several structured grids (composed by hexahedral elements), with cell number varying in the range 25k—400k. Figure 3 reports the grid dependence analysis in terms of entropy production across the cascade (frame 3a) and standard deviation of pressure at the cascade exit (frame 3b) for the baseline blade shape. The results clearly show an asymptotic trend towards grid independence, but also indicate that grids composed by at least 400k cells are required for a fully reliable CFD simulation. This result is motivated by the onset of multiple shock waves and shock-induced viscous effects, which require appropriate mesh resolution in order to capture their related gradients. Such high-resolution calculations require more than 1 hour of computational time on a 16-processor cluster. However, the use of lower mesh resolutions leads to a significant abatement in computational cost. Intermediate-resolution grids composed by 150k cells require about 15 minutes for each CFD run and guarantee an error of 3% with respect to the grid independent mesh. Low-resolution grids composed by 50k cells require 5 minutes of each CFD run and produce estimates within 5% of the grid independent one.

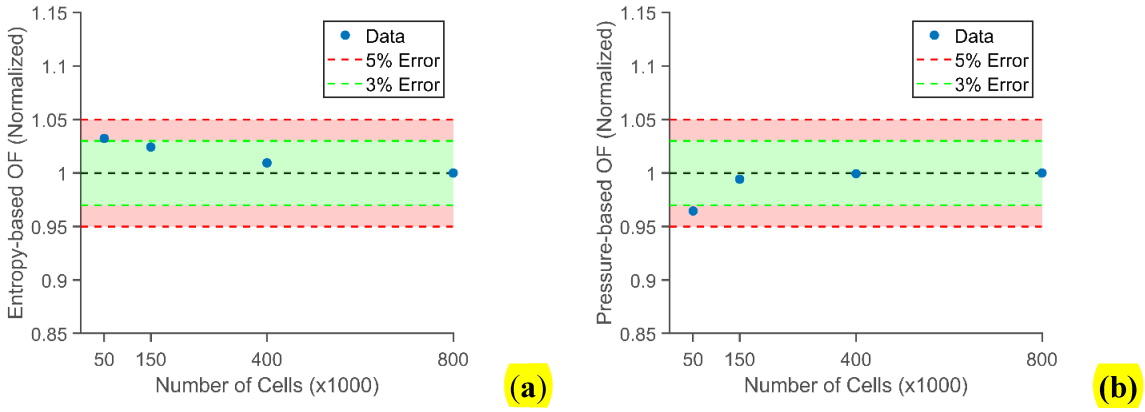

Figure 3. Grid dependence analysis for the calculation of the flow in the supersonic turbine nozzle in the baseline configuration, for two quantities: (a) entropy generation across the cascade; (b) standard deviation of the pressure downstream of the cascade.

As discussed in the following, most of the optimization tests were performed using a low-resolution grid, in order to limit the computational cost of the whole optimization process. However, further optimization with higher mesh refinement were also performed for comparison; the impact of using a low-resolution meshes demands a specific discussion, which is reported in Section 3.2. The cascade performance and the computed flow fields reported in this paper were evaluated by using the grid-independent mesh composed by 400k elements.

3.1. Baseline blade parametrization

The baseline blade geometry was parametrized using the method described in Section 2.1. The primary rule used to identify the number and distribution of CP is to keep the geometrical error (i.e., the distance between the points of the original blade and those of the interpolated one) below a certain threshold. 30 CPs resulted sufficient to provide an accurate interpolation in this case. By acting on the knot sequence, a non-uniform distribution of CPs was prescribed, reducing the distance between the CPs in regions of higher curvature. The CP distribution used in this work is reported in Figure 4.

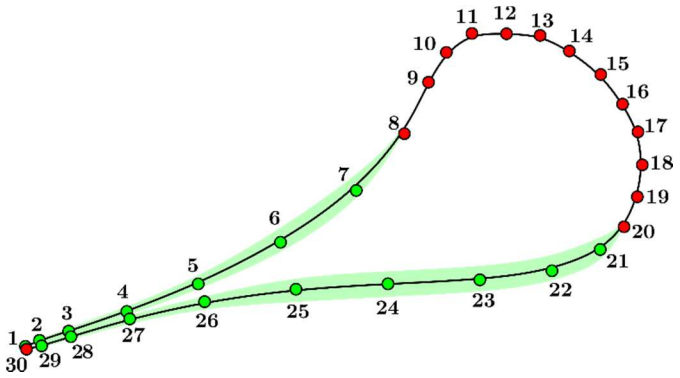

Figure 4. Baseline blade parametrization with 16 movable Cps (marked as green circles); fixed Cps are the red circles; the design space is highlighted by the green shadow.

Considering the aerodynamics of the cascade (Colonna et al., 2008) the supersonic region downstream of the throat is the most critical section of the blade, due to the onset of shock waves on the rear suction side and at the trailing edge of the blade. This was further confirmed by both an adjoint-based sensitivity analysis (Pini et al., 2015) and an ANOVA analysis (Romei and Persico, 2018) devoted to evaluate the impact of the control point position. Considering these indications, and with reference to the CP distribution reported in Figure 4, the objective function is considered insensitive to CPs 8-20 in the front part of the blade, so they were kept fixed for all the optimization tests presented in this paper. The trailing edge (TE) thickness was also geometrically constrained (mainly due to structural reasons as the TE is the zone of minimum blade thickness). In FORMA, this is done by keeping movable only the CP on one side, while the other one moves rigidly.

With reference to Figure 4, a subset of 16 movable CPs is prescribed to the present optimization: 6 on the pressure side, 9 on the suction side, and 1 in the TE region. Further subsets of movable CPs are considered in Section 3.3, in an effort of reducing the computational cost. The design space considered for the movable CPs on the two blade sides was set as composed by symmetric vertical displacements varying from a minimum of ±0.2 mm at the TE edge up to a maximum of ±1.5 mm in the central part of the blade.

Before considering the optimization of the supersonic turbine nozzle with a technically-relevant setup, an assessment study is reported in Section 3.2 to discuss the convergence process of local and

global strategies, their outcome and cost.

3.2. Assessment of optimization strategies

For the local-global comparison the objective function was defined as the standard deviation of the azimuthal pressure distribution at the cascade exit, evaluated half an axial chord downstream of the blade TE, i.e.:

$$\Delta P = \sum_{i=1}^{np} \sqrt{\frac{(P_i - P_{mix})^2}{np}} \tag{2}$$

where np is the number of computational points along the azimuthal direction on the downstream traverse. This quantity is a good marker of the presence of shock waves; in the present supersonic cascade, the minimization of this specific objective function is expected to reduce the shock strength, with positive implications on the cascade performance and on the forcing acting on the subsequent rotor. The high-fidelity flow model used here allows to select other objective functions, as discussed in section 3.5. For this assessment study, no constraints were introduced into the optimization process.

The initialization of the surrogate, required by both strategies, was performed by resorting to the Latin Hypercube Sampling technique. As a rule of thumb, the number of samples was chosen as 5 times the number of design variables, i.e. 80 individuals were generated and tested with CFD to initialize the surrogate with 16 movable CPs. In the case of SBLO, each iteration requires further 80 CFD runs. In case of SBGO, after the initial DoE, only one CFD run is required for each iteration.

The features of the two surrogate strategies are clearly visible by comparing the two frames of Figure 5, which report the convergence process of two methods. Despite the differences in the convergence processes, both the strategies exhibit a smooth convergence trend and both lead to a dramatic minimization of the objective function, reduced by more than one order of magnitude with respect to that evaluated in the baseline configuration. Interestingly, the two methods converge to a very similar optimum, which actually corresponds to a nearly identical blade shape, as shown by Figure 6. The optimal blades are, moreover, very different from the baseline one, especially in the

diverging section of the suction side. In particular, the optimal blades exhibit much higher divergence rate in the bladed duct downstream of the sonic throat, followed by an almost straight blade shape on the rear suction side, in the region of unguided turning. The action of the evolutionary strategy finds justification with proper aerodynamic considerations, as discussed in Section 3.4.

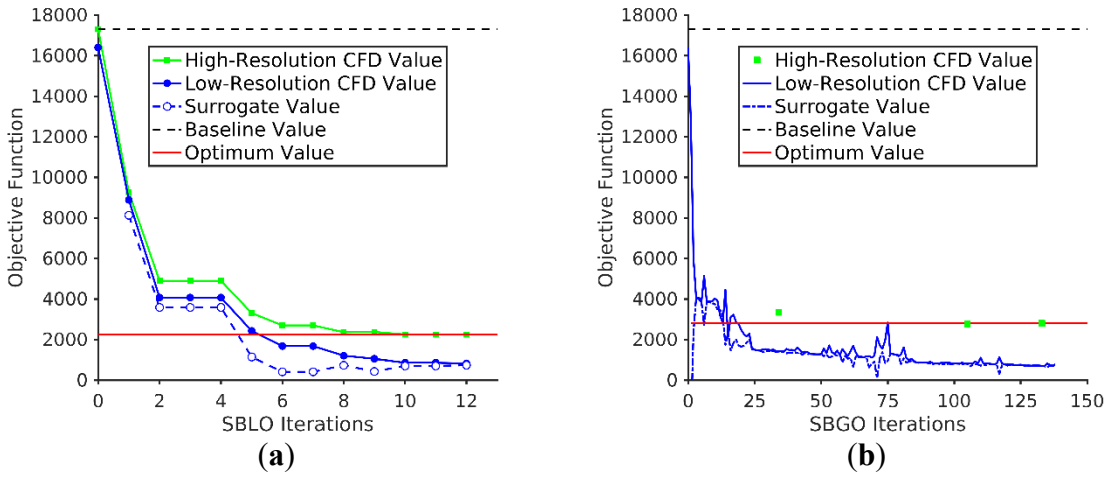

Figure 5. Convergence history for SBLO (a) and SBGO (b). The objective function (OF) is defined according to Eq. 2, expressed in Pa.

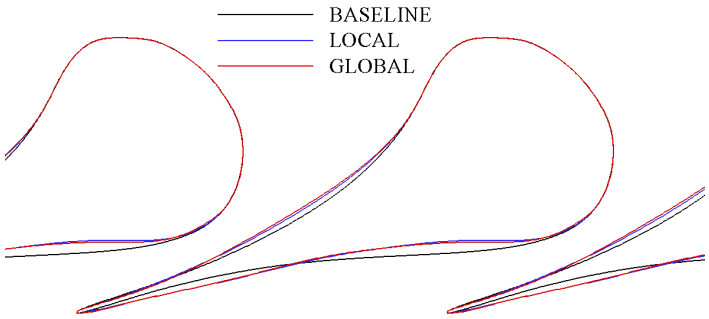

Figure 6. Comparison between baseline, optimal-SBLO, and optimal-SBGO blades.

In order to understand the motivation for the equivalent outcome of the two methods, a more refined analysis of the convergence trends is now reported, with focus on the surrogate convergence. Figure 5 shows explicitly the evolution of the surrogates alongside the one of the real flow model. These trends indicate that the surrogate models provide a relatively reliable representation since the very early stages of the process. In case of SBLO the surrogate and the flow model are aligned in the

first two iterations, which provide most of the minimization. In these two initial iterations, the trust region ratio is close to unity. This is a first indication that the DoE technique used to initialize the Kriging model is appropriate for the present design problem. Then iterations 3 and 4 do not improve the fitness and their minima are discarded; correspondingly, the trust region ratios are negative and, hence, the trust region shrinks. In the following iterations 5 and 6 the minimization restarts, but the difference between the surrogate and the flow model amplifies with respect to the first two iterations; the correspondent trust region ratio rises but it does not exceed 0.5. The 7th iteration fails to produce a further minimization, and a negative trust region ratio is also found. After that, the procedure acquires a continuous minimization trend, accompanied by a progressive convergence of the surrogate model to the CFD result. This is also testified by the corresponding values of the trust region ratio, which are greater than 0.5 in all the last iterations. Ultimately, 12 iterations are required to achieve convergence, which means about 1000 CFD runs to achieve convergence (equivalent to 100 hours of calculation on a 16-processor cluster in which no parallelization of the DoE is exploited).

The situation changes significantly when the SBGO is used. The experience with SBLO indicates that the initialization of the surrogate with 80 individuals is appropriate for the present problem, as testified by the excellent reliability of the surrogate model built from the initialization only ($r^0 \approx 1$). This is the ideal context for a training-based method. Thanks to the reliability of the initial surrogate, the first optimum is already close to the actual one and the progressive increase of sampling in this region allows to refine very effectively the surrogate just in the area of the design space close to the optimum. As a result, the surrogate quickly matches the high-fidelity tendency, except for some spikes of progressively reduced amplitude as the algorithm converges to the optimum. This results in a much 'cheaper' convergence process, indicating that the training procedure chosen for the Kriging model is appropriate for the present design problem. As a result, about 220 CFD runs are required for finding the optimum, i.e. about 20% of the ones required by the SBLO.

The trends in Figure 5 also show a very relevant feature. The aforementioned computational costs were achieved running CFD simulations with low resolution grids (50k cells), so to limit the

computational cost of the optimization. To investigate the impact of the mesh resolution on the optimization, the optimal configuration as well as some intermediate ones were subsequently simulated using high-resolution meshes (400k cell) for assessment. Regarding the SBLO, the optimal configurations for all the iterations were assessed and are shown in Figure 5a. As clearly visible, the high-resolution values are always higher than those obtained with the low-resolution model (as expected, since the pressure gradients are better captured by the high-resolution CFD model). However, the trend of the high-resolution values perfectly matches the one achieved during the optimization, and the high-resolution calculations indicate that the optimum of the final iteration is actually the best individual among the tested ones. The few assessments performed for the SBGO provide the same indication.

As conclusive test, a global optimization was run also using meshes composed by 150 kcells. The result of the convergence process is reported in Figure 7. The convergence process resembles in any feature the one observed for the SBGO with lower resolution. When evaluated with the grid independent mesh, the objective function value of the optimal configuration recovers the one achieved by running the optimization with the low-resolution mesh.

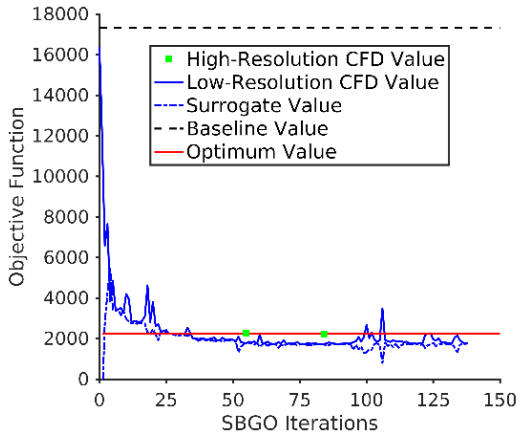

Figure 7. Convergence process for SBGO performed with finer CFD mesh resolution (150k cells vs 50k cells). The objective function (OF) is defined according to Eq. 2, expressed in Pa.

This analysis shows that both the surrogate strategies successfully identify the optimal configuration even though prescribing a low-resolution mesh to the computational flow model,

despite the complexity of the flow configuration, which features shocks and related viscous interactions. Moreover, the SBGO allows performing reliable optimizations at the lowest cost. Such set-up was then used for the technically-relevant optimization studies discussed in the following

3.3. Constrained Optimization

Turbomachinery design has usually to cope with constraints. In subsonic cascades the flow angle is typically constrained, as it has a double implication on the flow rate and the velocity triangles. In case of supersonic turbines featuring a sonic throat in the bladed channel, the flow rate becomes a crucial quantity to constraint; as a matter of fact, an erroneous design might prevent from matching the design flow rate, limiting the power release of the system. In case of rotor blades, some geometrical features (such as the area or the moments of inertia of the profile section) might also be constrained to ensure structural integrity.

In the present ORC turbine nozzle configuration, the supersonic flow conditions lead to consider the flow rate as the most relevant quantity to constraint. More specifically, the flow rate discharged by the cascade was constrained to remain within ±1% of the baseline value. A constrained optimization of the cascade was then performed, using the design space as reported in Figure 4, the objective function as given in Eq. 2, and using a SBGO strategy.

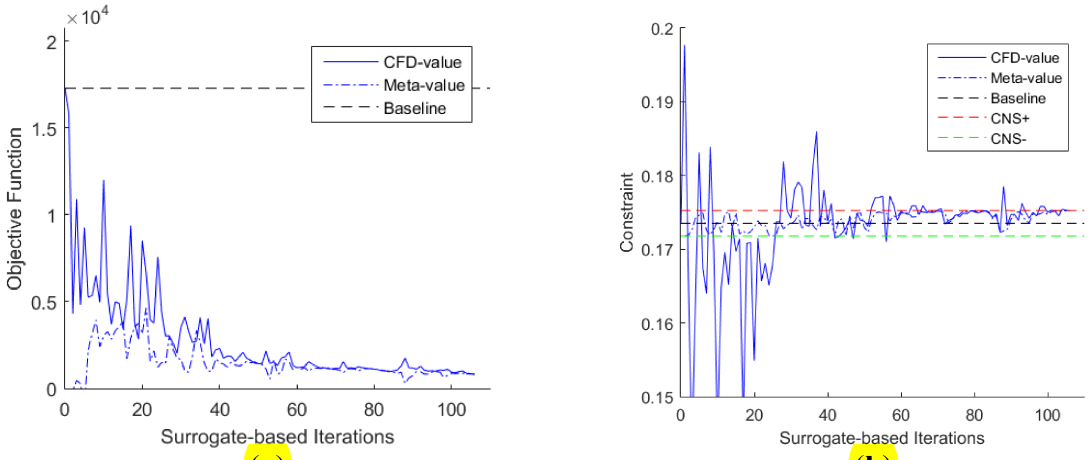

Figure 8. Convergence history of the constrained SBGO: (a) objective function; (b) constraint.

The objective function (OF) is defined according to Eq. 2, expressed in Pa on mesh; the constraint is defined as the mass flow rate, expressed in kg/s.

Figure 8(a) reports the SBGO iterations of the constrained optimization, while Figure 8(b) reports the convergence history of the constrained quantity. The two trends indicate that the process requires about 110 iterations to reach convergence, similarly to what found for the non-constrained optimization. The convergence trends of the objective function and of the constraint are relatively smooth, with few spikes in the initial phases of the process; both the surrogates converge to the high-fidelity value in about 60 iterations. It is interesting to note that the optimization drives the blade design towards the maximum threshold of the acceptability range in flow rate. From this perspective, it is important to remark that the flow rate is evaluated with very high accuracy already by the low-resolution mesh (0.013% of error with respect to the value obtained with the grid-independent calculation). This is because, in supersonic cascades, the flow rate is imposed by the upstream thermodynamic conditions, the fluid thermodynamic model, and the throat size, corrected by the displacement thickness of the boundary layers. Since high resolution numerical schemes are used, the mesh may only affect this latter term; however, the blade boundary layer is properly simulated even by the coarse mesh, as this latter features high cell density in the boundary layer and a wall y^+ lower than one. The reliable estimate of the flow rate achieved during optimization guarantees that the constrained is fulfilled also in the high-resolution assessment.

The high-resolution analysis of the optimal blade shows that the optimization allows reducing the objective function by 85%, resulting 2500 Pa (high-resolution value) from a baseline value of about 17000 Pa. Further interesting findings emerge when the high-resolution CFD simulations are processed to evaluate the cascade performance. In this work, the total pressure loss coefficient Y was used as indicator of the cascade performance, defined as:

$$Y = \frac{P_{T,IN} - P_{T,OUT}}{P_{T,IN} - P_{OUT}} \tag{3}$$

in which the total and static pressure at the outlet ($P_{T,OUT}$, P_{OUT} respectively) are evaluated four axial chords downstream of the blade, where the flow is completely mixed-out. The baseline blade exhibits a total pressure loss coefficient of 18.3%. The optimal blade obtained as a result of the constrained

optimization features a loss coefficient of 9.7%; the optimization process has, hence, reduced by almost 50% the loss coefficient; as it will be clarified in Section 3.5, this have been achieved since, by minimizing the pressure azimuthal variability, the shocks have been drastically weakened.

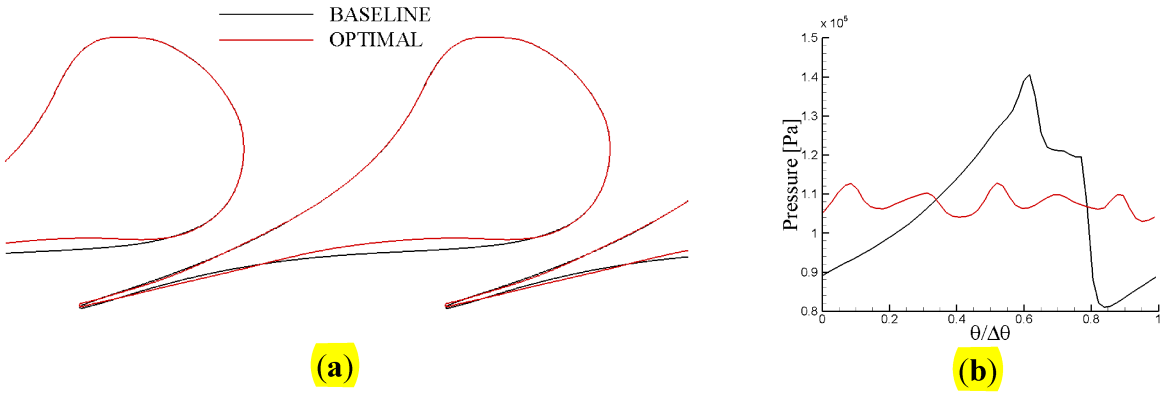

Figure 9. Outcome of constrained optimization for SBGO with 16 CPs: (a) baseline and optimal blade shapes; (b) pressure distribution half an axial chord downstream of the blade.

Figure 9 reports the optimal blade, compared to the baseline one, alongside the pressure distribution at the cascade exit. It is evident that the optimal blade manages to strongly reduce the pressure variability at the cascade exit, eliminating the sharp pressure gradients induced by the shocks that affect the baseline cascade aerodynamics.

An accurate analysis of the blade geometry reveals that the optimization changes the shape of the blade in the rear part but it tends to retain the shape of the baseline blade in throat area, acting as a geometric constraint assigned to the throat size. This is fully consistent on the application of the constraint on the flow rate, which can only depend on the throat size in this cascade. This behavior is illuminating on the capability of the evolutionary strategy in 'learning' through surrogate interpolation, finally producing physically-sound optimal configurations.

3.4. Impact of the design space

The optimization tests discussed so far were performed considering the design space reported in Figure 4, which was selected on the basis of previous studies and aerodynamic considerations. However, it is interesting to evaluate the penalty (if any) in the optimization outcome by using a

reduced design space, which provides an evident reduction of computational cost. Analysis of the optimal blades shown in Figure 6 and 9 indicates that the pressure side of the blade (CPs 3—7) might have a limited impact on the constrained optimization, as well as the first movable CP on the suction side (CP 21), which is slightly upstream of the cascade throat, especially considering the constraint imposed on the flow rate. The trailing edge region seems also to have a limited impact.

Therefore, several constrained optimizations were run, reducing progressively the number of movable CPs, while maintaining the same maximum displacement. Three configurations were considered: (i) one test fixing the most upstream CPs (CP6, CP7, and CP21), resulting in 13 movable CPs; (ii) one test fixing the upstream CPs and the four CPs closest to the trailing edge (CP1, CP2, CP29, CP30), resulting in 10 movable CPs; (iii) one test fixing the entire pressure side, the trailing edge, and CP21, resulting in 7 movable CPs. Table 1 summarizes the most relevant results of these optimization trials, alongside the original one, in terms on computational cost, objective function minimization, and cascade performance (all these cases, of course, satisfy the flow rate constraint).

Table 1. Setting, cost, and outcome of multiple optimizations with different design spaces

CPs	DoE runs	SBGO runs	Tot- CFD runs	ΔP [Pa]	$\Delta P / \Delta P_{BASE}$	Y [%]
16	80	106	186	2.5×10^3	0.15	9.7
13	65	96	161	2.5×10^3	0.15	9.7
10	50	63	113	4.0×10^3	0.23	10.2
7	35	33	68	4.0×10^3	0.23	9.9

The quantitative results reported in Table 1 show that a 65% reduction of computational cost can be achieved by progressively reducing the number of movable CPs, still obtaining a very significant minimization (almost 75%) of the objective function and a practically identical cascade performance. However, an almost double value of the minimum objective function has to be accepted if both the pressure side and the trailing edge of the blade are fixed (7 movable Cps). These results allow to derive three relevant conclusions on the shape-optimization of this supersonic cascade: (i) the design of the suction side of the blade downstream of the throat in both the bladed and in the un-guided turning regions is crucial for the cascade optimization; (ii) the position of the trailing edge has a

limited but measurable impact on the design, and it should be assigned movable to properly exploit the blade evolution in the rear sections; (iii) the rear pressure side has a minor impact of the optimization. Conversely, the upstream CPs have no influence on the optimization; the optimization with 13 movable CPs leads to the same result of the original set-up, hence providing the ideal combination of the optimization outcome and computational cost.

3.5. Aerodynamic analsys

The flow field established in the cascade optimized using 13 movable CPs is now discussed, in comparison to the one of the baseline configuration. To this end, the entropy and Mach number distributions on the blade-to-blade surface are reported for the two cases in Figure 10. First considering the baseline cascade, reported in the frames of Figure 10a, a smooth acceleration occurs in the front part of the blade and on both the blade sides up to the sonic throat, followed by a supersonic acceleration within the diverging channel. As a consequence, the supersonic flows approaching the trailing edge region trigger the onset of a typical fishtail shock system.

The higher Mach number on the rear suction side leads to the generation of a shock oriented downstream; on the pressure side, despite the lower Mach number, the supersonic rotation around the circular trailing edge generates a further shock, oriented towards the adjacent blade. This latter shock reflects on the suction side of the adjacent blade, and it is re-oriented downstream. Contrarily on what could be expected, the strength of the reflected shock grows as it develops downstream; moreover, it has a different inclination with respect to the one originated on the suction side of the trailing edge. So the two shocks eventually coalesce in a single strong shock at about one and a half axial chord downstream of the trailing edge. The unexpected behavior of the reflected shock is caused by the shape of the blade on the rear suction side. In particular, the curved shape of the rear suction side imposes to the flow a significant unguided turning, and hence the expansion proceeds leading to a large over-speed. This process is terminated by a compression wave that eventually originates a shock, close to the position where the trailing edge shock generated by the adjacent blade is reflected.

(**a**): baseline configuration (**b**) optimal configuration

Figure 10. High-fidelity assessment of the optimization outcome – flow field: (a) entropy (top) and Mach number (bottom) distributions for the baseline blade configuration; (b) entropy (top) and Mach number (bottom) distributions for the optimal blade configuration.

So, it can be concluded that the leading shock appearing downstream is mostly due to the peculiar shape of the blade. This is further demonstrated by inviscid simulations of this cascade performed using sharp trailing edge (see Colonna et al., 2008b and Pini et al., 2015), in which the fishtail shock system was eliminated while the leading shock remained almost unaltered with respect to the one observed in the present high-fidelity simulations. The onset of severe azimuthal pressure gradients is also responsible for the weaving character of the wake appearing from the entropy distribution. This latter also marks the presence of the main shock and the loss generation due to its mixing downstream.

It should be noted that no classic semi-analytical method can guarantee to avoid these effects when designing a highly supersonic cascade operating with a non-ideal fluid. This further motivates the application of a systematic optimization tool. The action of the optimization can be clearly understood when comparing the flow configuration obtained for the optimal cascade, reported in Figure 10b, with the baseline one. With the aim of minimizing the pressure oscillations downstream

of the cascade, the optimizer has generated an optimal blade with higher curvature in the diverging channel and an almost straight shape downstream of the cascade opening. As a result, the pressure gradients in the region of unguided turning are drastically minimized, thus eliminating the main shock observed in the baseline case. Consequently, the flow is discharged by the cascade with a much more uniform direction, as marked by the straight wake avenues leaving the optimal cascade.

The computed flow fields were further processed to extract the pressure distribution on the blade surface and the streamwise evolution of total pressure loss, evaluated as expressed in Eq. 3. The pressure distribution on the blades, reported in Figure 10a as isentropic Mach number (M_{IS}) along the non-dimensional curvilinear abscissa (S/S_{MAX}) of the blade, supports the aforementioned interpretation of the optimization effect. The shape of the optimal blade moves ahead the flow acceleration on the suction side, just downstream of the sonic throat and still within the bladed channel; consequently, the over-speed on the rear suction side is limited and thus the subsequent deceleration is mostly reduced, as testified by the almost constant pressure for $0.8 \leq S/S_{MAX} \leq 1.0$.

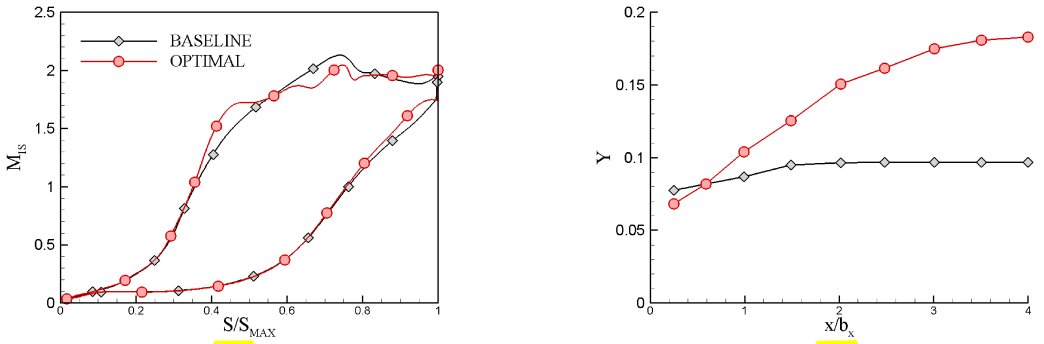

Figure 11. High-fidelity assessment of the optimization outcome – processing: (a) isentropic Mach number distribution on the baseline and optimal blade; (b) axial evolution of total pressure loss coefficient downstream of the cascade.

The indirect outcome of these flow features on the cascade performance is best synthesized by Figure 11b, which reports the evolution of total pressure loss coefficient (Y) for the two cascades along the axial direction, normalized by the blade axial chord (x/b_x). The two configurations exhibit similar loss levels at the TE, with even a small of advantage of the baseline blade. In this

region, the loss is mainly caused by the boundary layers on the blade sides. Moving downstream, the onset of the shock has a dramatic impact on the loss coefficient of the baseline blade; in particular, the entropy production connected to the mixing of the shock gradients makes rise the loss coefficient from 7% at the TE to 18.3% four axial chords downstream, where the flow is completely mixed-out. In the case of the optimal blade, the elimination of the main shock greatly minimizes the mixing loss, which now amounts to 2%. This leads to a mixed-out loss coefficient equal to 9.7%, almost halved with respect to the one of the baseline blade, and especially in a much faster mixing, with evident advantages for the aerodynamics of the subsequent rotor.

As a further indication, the outlet flow angle of the optimal cascade (evaluated four axial chords away from the trailing edge, where the flow is mixed-out) results 76.8°, namely less than 0.2° different from the corresponding one of the baseline cascade. This indicates that the optimal blade is aerodynamically equivalent to the baseline one in terms of the both flow rate and flow angle, while providing very relevant advantages in terms of performance and flow uniformity.

3.6. Impact of using an entropy-based objective function

The analysis on the loss coefficient in Section 3.5 has shown that minimizing the pressure oscillations at the exit of a supersonic blade row also allows increasing its aerodynamic performance. In fact, most often the reduction of losses is the ultimate goal of turbomachinery design. From this perspective, it is interesting to perform optimizations using an objective function more directly connected with the aerodynamic losses, such as the entropy production across the cascade (which is a direct measure of aerodynamic loss in adiabatic flows). It is to be remarked that the present optimization method is inherently suitable for such 'entropy-based' optimization, as the high-fidelity flow model used for the CFD runs performed during optimization includes all the sources of cascade aerodynamic loss (turbulent boundary layers, turbulent mixing, and shock mixing).

In order to investigate the outcome of an entropy-based optimization, the baseline blade was re-optimized using SBGO and 13 movable Cps, and still constraining the flow rate to stay within ±1%

of the design vale. The objective function was defined as the entropy production from the inlet to the outlet of the cascade, whose value amounts to about 4.4 J/(kg K) for the baseline cascade. Frames 12(a) and 12(b) report the convergence process, showing a quick convergence of the surrogate model to the high-fidelity CFD result for both the objective function and the constraint. The resulting optimal blade is reported in Figure 12(c), alongside the pressure-based optimal one. The two optimal blades appear extremely similar, with only minor differences in the diverging regions of the bladed channel downstream of the throat, on both the pressure and suction sides.

When analyzed with high-resolution CFD, the entropy-based optimal blade exhibits a mixed-out loss coefficient of 9.7%, nearly identical to the one obtained with the pressure-based optimization. However, the entropy based optimal blade features a standard deviation of the pressure at the cascade exit of 4000 Pa, i.e. almost twice the one obtained with the pressure-based optimal blade. These differences indicate that, despite the heuristic character of the optimization method, the design outcomes are reliable and fully consistent with the objective function definition.

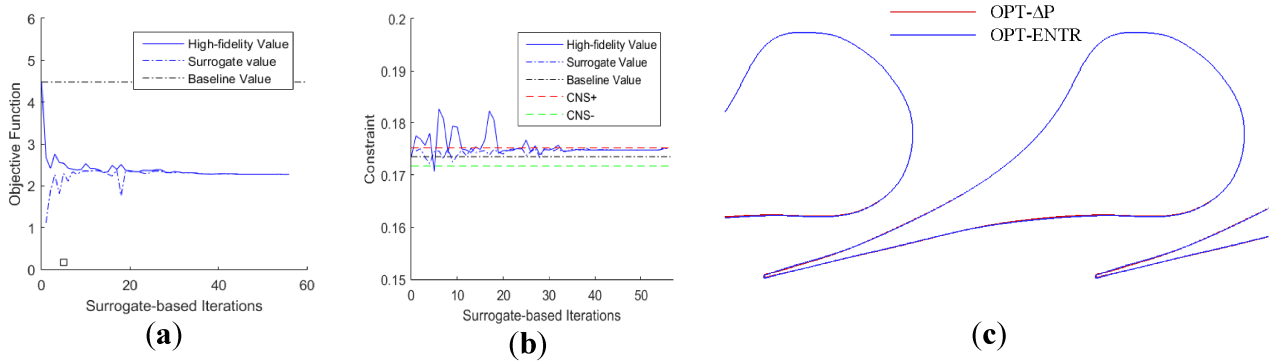

Figure 12. Results of SBGO performed with 13 movable CPs and the entropy production across the cascade as objective function: (a) convergence process – objective function [J/kgK]; (b) convergence process – constraint [kg/s]; (c) baseline and optimal blades

4. CONCLUSIONS

This paper has presented a novel high-fidelity optimization approach for the design of blades for non-conventional turbomachinery applications. Detailed descriptions of all steps of the optimization scheme have been provided, namely the geometry parametrization, the high-fidelity flow solver, the

genetic algorithm, the surrogate model, and their combination in two evolutionary strategies. The blade shape is parametrized via B-Splines, whose local control capability allows a detailed shape reconstruction while preserving surface smoothness. The implementation of advanced high-fidelity flow models, of paramount importance for non-conventional turbomachinery, is easily achieved thanks to the non-intrusive character of the evolutionary optimization strategy here used. To tackle the computational burden typical of CFD-based evolutionary strategies, the genetic algorithm is coupled to a surrogate model that reflects the influence of the design variables on the objective function and on the constraints. Two different optimization strategies have been presented and discussed to evaluate the convergence process and the associated computational cost. Application to a supersonic Organic Rankine Cycle turbine nozzle has led to an almost 50% reduction of loss coefficient, achieved by minimizing both the azimuthal pressure variation downstream of the cascade (reduced by 85%) and the entropy generation across the cascade (reduced by 50%), while maintaining aerodynamic equivalence in terms of flow rate and outlet flow angle. The application of a constraint in the flow rate, in particular, has shown that the evolutionary optimization is able to produce not only efficient individuals, but also physically-sound optimal configurations. This suggests that in complex cases the method can be used not just as a technical design tool, but also as machine-learning device in view of constructing novel design rules. This makes the method particularly suitable for application to non-conventional machines, for that well-established design guidelines are still unavailable.

ACKNOWLEDGEMENTS

The work of Pablo Rodriguez-Fernandez was supported by the EAGLES program under a European Union - United States Atlantis grant.

NOMENCLATURE

Roman Symbol

r	Trust Region Ratio (Equation 1)
M	Mach Number
P	Pressure [Pa]
S/SMAX	Non-Dimensional Curvilinear Abscissa
T	Temperature [K]

x/bX	Axial direction normalized by the blade axial chord
Y	Total Pressure Loss Coefficient (Equation 2)

Acronym

CFD	Computational Fluid-Dynamics
CP	B-Spline Control Point
CNS	Constraint
DoE	Design of Experiments
DEVST	Pressure Standard Deviation
EoS	Equation of State
ENTR	Entropy-based
FORMA	POLIMI Shape Optimizer
GA	Genetic Algorithm
LuT	Look-up-Table
OF	Objective Function
ORC	Organic Rankine Cycle
OPT	Optimum Blade
SBGO	Surrogate-Based Global Optimization
SBLO	Surrogate-Based Local Optimization
TE	Trailing Edge
ΔP	Standard deviation of the azimuthal variation of static pressure at the cascade exit

Subscript

IS	Value obtained through isentropic transformation
IN	Inlet of Cascade
OUT	Outlet of Cascade
T	Total Quantities

REFERENCES

Persico, G., Pini, M., "Fluid dynamic design of Organic Rankine Cycle Turbines", 2016, book chapter for "Organic Rankine Cycle (ORC) Power Systems: Technologies and Applications", Woodhead Publishing, ISBN 978-0-08-100510-1.

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

Persico, G., Rebay, S., "A penalty formulation for the throughflow modeling of turbomachinery", 2012, Computers & Fluids, Vol. 60, pp. 86–98.

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

Peter, J., Dwight, R., "Numerical sensitivity analysis for aerodynamic optimization: A survey of approaches", 2010, Computers & Fluids, Vol. 39, pp. 373–391.

Pini, M., Persico, G., Pasquale, D., and Rebay, S., "Adjoint method for shape optimization in real-gas flow applications", 2015, Journal of Engineering for Gas Turbines and Power, Vol 137 (032604).

Vitale, S., Albring, T., Pini, M., Gauger, N., Colonna, P., "Fully turbulent discrete adjoint solver for non-ideal compressible flow applications". Journal of the Global Power and Propulsion Society. Vol. 1, pp. 252–270

Pierret, S., Coelho, R., Kato, H., "Multidisciplinary and multiple operating points shape optimization of three-dimensional compressor blades", 2006, Structural and Multidisciplinary Optimization, Vol. 33, Issue 1, pp. 61–70.

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

Pasquale, D, Ghidoni, A., Rebay S., "Shape optimization of an Organic Rankine Cycle radial turbine nozzle", 2013, Journal of Engineering for Gas Turbines and Power, Vol. 135 (042308)

Coello, C., "An updated survey of GA-based multiobjective optimization techniques", 2000, ACM Computing Surveys, Vol. 32, Issue 2, pp. 109–143.

Pierret S., Van den Braembussche, R.A., "Turbomachinery Blade Design Using a Navier–Stokes Solver and Artificial Neural Network", 1999, Journal of Turbomachinery, Vol. 121, Issue 2, pp. 326–332.

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

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

Barth, T.J., Jesperson, D.C, "The Design and Application of Upwind Schemes on Unstructured Meshes", AIAA Paper 89-0366, 1989.

Persico, G., Mora, A., Gaetani, P., Savini, M. (2012). Unsteady aerodynamics of a low aspect ratio turbine stage: Modeling issues and flow physics. Journal of Turbomachinery, Vol. 134 (061030).

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

Toropov, V. V., Filatov, A. A., Polynkin, A. A., "Multiparameter Structural Optimization Using FEM and Multipoint Explicit Approximations", Structural Optimization, vol. 6, pp. 7-14, 1993

Giunta, A. A., Eldred, M. S., "Implementation of a Trust Region Model Management Strategy in the Dakota Optimization Toolkit" Proceedings of the 8th Symposium on Multidisciplinary Analysis and Optimization, 2000, Long Beach, CA; United States; 6-8 September 2000

Pini, M., Persico, G., Dossena, V., "Robust Adjoint-based Shape Optimization of Supersonic Turbomachinery Cascades", ASME Turbo Expo 2014: Turbine Technical Conference and Exposition, GT 2014; Dusseldorf; Germany; 16-.20 June 2014

Adams, B.M., Bauman, L.E., Bohnhoff, W.J., Dalbey, K.R., Ebeida, M.S., Eddy, J.P., Eldred, M.S., Hough, P.D., Hu, K.T., Jakeman, J.D., Stephens, J.A., Swiler, L.P., Vigil, D.M., and Wildey,

T.M., "Dakota, A Multilevel Parallel Object-Oriented Framework for Design Optimization, Parameter Estimation, Uncertainty Quantification, and Sensitivity Analysis: Version 6.6 User's Manual," Sandia Technical Report SAND2014-4633, Updated May 2017

Colonna, P., Nannan, N., and Guardone, A., "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", 2008a, Fluid Phase Equilibria, 263(2):115 – 130.

Colonna, P., Harinck, J., Rebay, S., Guardone, A., "Real-Gas Effects in Organic Rankine Cycle Turbine Nozzles", 2008b, Journal of Propulsion and Power, Vol 24, Issue 2, pp. 282-294

Romei, A., Persico, G., "Novel Shape Parametrization Technique Applied to the Optimization of a Supersonic ORC Turbine Cascade", 2018, ASME Turbo Expo 2018, ASME paper GT2018-76732