Optimal Design of ORC Turbine Blades Under Geometric and Operational Uncertainties

Nassim Razaaly, Giacomo Persico, Pietro Marco Congedo

OPTIMAL DESIGN OF ORC TURBINE BLADES UNDER GEOMETRIC AND OPERATIONAL UNCERTAINTIES

Nassim Razaaly[1*], Giacomo Persico[2], Pietro Marco Congedo[1]

[1] DeFI Team (INRIA Saclay IDF, Ecole Polytechnique), CMAP Laboratory,
1 Rue Honoré d'Estienne d'Orves, 91120 Palaiseau, France
* email: nassim.razaaly@inria.fr

[2] Laboratorio di Fluidodinamica delle Macchine, Politecnico di Milano,
Via Lambruschini 4, I-20156 Milano, Italy

ABSTRACT

Typical energy sources for Organic Rankine Cycle (ORC) power systems feature variable heat load, hence turbine inlet/outlet thermodynamic conditions. The use of organic compounds with heavy molecular weight introduces uncertainties in the fluid thermodynamic modeling and complexity in the turbomachinery aerodynamics, with supersonic flows and strong shocks, which grow in relevance in the aforementioned off-design conditions. These features also depend strongly on the local blade shape, which can be influenced by the geometric tolerances of the blade manufacturing. This study presents a Robust Optimization (RO) analysis on a typical supersonic nozzle cascade for ORC applications under the combined effect of uncertainties associated to operating conditions and geometric tolerances: a classical formulation consisting in minimizing the mean of a well-suited performance function, constraining the average mass flow rate to be within a prescribed range is addressed, by means of a bi-level Gaussian Process (GP) surrogate-based approach. Influence of the operating conditions range and geometric variability are investigated considering several scenarios, in which the different effects act in combination or separated; results indicate that the combination of different classes of uncertainites has an impact on the robust-optimal blade shape and, in turn, in their response in the frame of uncertain scenarios.

1. INTRODUCTION

Small-medium scale ORC power plants have received great interest in both the technical and academic community, in particular due to its ability to recover mechanical energy from low-grade heat sources such as, solar, geothermy, biomass or waste heat. As well known (Macchi (2013); Meroni et al. (2018)), the performance of the ORC power system is strongly linked to the efficiency of the turbine. The turbine aerodynamics is complicated by the use of organic fluids, which combine low enthalpy drops, high-expansion ratio per stage and low speed of sound, leading to transonic or supersonic ORC turbines, which demand the use of converging-diverging cascades and are commonly prone to strong shock waves and chocked flow conditions.

Recent advances of fluid dynamic simulation tools accounting for so-called *dense gas* effects, induced by the use of organic fluid described by complex Equation of State (EOS), has permitted the development of Fluid-dynamic Shape Optimization (FSO) approaches for automated design of ORC blade cascades. Nevertheless, in ORC applications, the operational variability of the hot and cold sources often results in significant variations of inflow and outflow conditions at the turbine. Moreover, geometric imperfections of turbine cascades, stemming from manufacturing processes, may have a detrimental influence of the machine performance, which grow in relevance in the aforementioned off-design conditions. The fluid-dynamic design of ORC turbines could benefit for automated design methodologies, possibly integrating uncertainties.

This work presents the results of the application of a *robust* shape optimization method to the design of a typical converging-diverging turbine nozzle for ORC applications, first introduced by Colonna et al. (2008), and extensively studied in the context of deterministic optimization (Pini et al. (2015a); Vitale et al. (2017); Persico et al. (2019)) and multi-point optimization (Pini et al. (2014)). Some recent works investigated design of the aforementioned blade under epistemic uncertainties (Razaaly et al. (2019a)) due to turbulence modelling, and investigated the 'robustness' of the blade design under aleatoric uncertainties due to variability in the operating conditions, uncertainty in the thermodynamic model parameters, and geometric tolerances (Razaaly et al. (2019b)).This study presents a RO framework under bother operational and geometrical (aleatory) uncertainties. The mean of the performance function is minimized (classic RO formulation), while the mean mass-flow rate is constrained to lie within a prescribed range centered to the nominal value. Different scenarios involving operational and geometrical uncertainties are considered, the latter being modeled using a non-stationary gaussian random field, discretized with a Karhunen-Loeve expansion.

The paper is organized as follows. Section 2 describes the turbine cascade, the parametrization adopted to control the blade geometry and provides details on the physical solver and the mesh. Section 3 and 4 presents respectively the uncertainty modelling framework and the optimization methods. Section 5 analyses the resulting optimal blades discussing in particular, their performances under uncertainties.

2. TURBINE BLADE MODEL

2.1 Cascade Geometry and Parametrization

The Biere represents a refence two-dimensional benchmark geometry to test the design of devices operating with the siloxane fluid MDM (Octamethyltrisiloxane, $C_8H_{24}O_2Si_3$). This blade profile is meant to obtain a convergent-divergent cascade passage which serves to accelerate the fluid up to a supersonic speed. Across the cascade, the fluid is expanded from superheated conditions. As the flow past the cascade is highly supersonic ($M \approx 2$ at the blade trailing edge), compressibility effects play a key role. Indeed, because of the high Mach number achieved at the nozzle exit, a typical fish-tail shock pattern is generated downstream the trailing edge. The presence of strong shocks past stator vanes may result in large losses and thus the design of the trailing edge region is critical to the turbine efficiency. Moreover, shock-waves propagate through the vane and usually interact with the boundary layer developing over the suction side of the neighboring blade, thus further compromising the efficiency of the cascade.

The Biere pressure and suction sides are parametrized using a unique B-spline curve of degree 3 (Farin (2002)), defined over a total number of 30 Control Points (CP). The design vector $\boldsymbol{x} \in \Omega \subset \mathbb{R}^9$ parametrizing the 2D cascade is constituted by a subset of 9 CPs allowed to be displaced in the direction normal to the baseline geometry, Ω referring to the design space.

2.2 Numerical Simulations and mesh

Since the study aims at the aerodynamic optimization of the blade profile, the flow model focuses on the two-dimensional flow at the midspan section of the cascade. The numerical domain is periodic with a pitch spacing of 45-mm. The flow is simulated up to a distance of 0.5 and 2 chord-lengths ahead and past the blade, respectively. The SU2 open-source suite (Economon et al. (2016)) was used for the CFD simulations, based on inviscid models. We used a generalized Approximate Riemann solver (ARS), of Roe type, with the SU2's library of thermodynamic models for complex fluid flows in the non-ideal regime (Vitale et al. (2015)). using the improved Peng-Robinson-Stryjek–Vera (iPRSV) Equation Of State (EOS). An implicit Monotone Upstream-centered Schemes for Conservation Laws (MUSCL) scheme, with van Albada slope limiter, is used to ensure second-order accuracy and prevent spurious oscil-

lations in the steady-state solution. Non-Reflecting Boundary Conditions (Giles (1990)) are also implemented to suppress the non-physical reflection of acoustic pressure perturbations at outflow boundaries. Detailed convergence analyses (not shown) have been performed, leading to a computational mesh of the flow domain with $36,000$ triangular elements, which represents a trade-off between accuracy and computational cost.

During the optimization process, several blade profiles are progressively generated. A dedicated mesh deformation tool based on Radial Basis Functions (RBF), originally presented in De Boer et al. (2007) and successfully applied in Pini et al. (2015b); Razaaly et al. (2019b) allows high flexibility and robustness while maintaining the grid connectivity.

The objective function (ΔP) is defined as the standard deviation of the azimuthal distribution of static pressure half an axial chord downstream of the blade Trailing Edge (TE). Indeed, minimizing ΔP within the optimization is convenient for such highly supersonic cascade since it allows achieving a severe reduction of the shock strength, and hence of the shock loss, thus improving the cascade performance and, at the same time, reducing the perturbations entering the downstream rotor (Persico et al. (2019)). The mass-flow rate per unit span $\dot{m}$, normalized with respect to the nominal value is the second Quantity of Interest (QoI), as used to formulate the constraints.

Total Pressure P_{in}^t, total Temperature T_{in}^t, and axial flow direction are assigned at the inlet, while static pressure P_{out}^s is given at the outlet.

3. UNCERTAINTY TREATMENT

The geometric variations due to blade manufacturing are assumed to be represented by a non-stationary gaussian random field of null mean Dow and Wang (2014), fully described by its auto-covariance function (see Razaaly et al. (2019b) for details), using a constant standard deviation $\sigma_0 = 3 \times 10^{-5}$m, the latter quantifying the level of manufacturing variability. A Karhunen-Loeve (KL) expansion, based on a spectral decomposition of the autocovariance function (Betz et al. (2014)) permits to parametrize the random field with a standard gaussian vector of dimension 7, $\boldsymbol{u} \sim \mathcal{N}(0, \boldsymbol{I}_7)$.

Following Colonna et al. (2008), we consider first a nominal (or full-load) operating condition for this nozzle cascade, characterized by an inlet thermodynamic state close to the saturation curve, and a high expansion ratio equal to ≈ 6. As well known, ORC power systems are requested to operate at part-load for long periods during their technical life, due to changes in the thermal power made available by the heat source and in the condenser temperature; this variation implies a large change in the turbine pressure ratio, resulting in a variation of aerodynamic loading on each cascade. In Colonna et al. (2008), the implication of part-load operation for this cascades was estimated so to reduce the pressure ratio to ≈ 4, by an increase of cascade outlet pressure. The latter comment lead us to consider the operating conditions variability as reported in Table 1, featuring a large variability w.r.t. the outlet pressure, assuming independence and uniform distributions for each component of the resulting random vector $\boldsymbol{\xi}$ modelling the input variability. A scenario accounting only for operating conditions leading to a random vector $\boldsymbol{\xi}$ of dimension 3 is considered, while another ore accounts also for both geometric and operational conditions variability, permitting to investigate the impact of geometric uncertainties on the RO profiles.

Within the robust optimization process, the first order statistics of the QoIs, namely ΔP, and $\dot{m}$, are evaluated by means of 24 CFD samples $\{\boldsymbol{\xi}_i, f(\boldsymbol{\xi}_i)\}_i$, f denoting a QoI, where the $\boldsymbol{\xi}_i$ are selected following Latin-Hypercube-Sampling (LHS). A combinaison of Gaussian-Process (GP) surrogate and Monte-Carlo (MC) is then carried out to estimate the statistics.

Table 1: Operating Conditions (Nominal and Uncertain). Random Vector $\boldsymbol{\xi} = (P^s_{out}, P^t_{in}, T^t_{in}, \boldsymbol{u}) \in \mathbb{R}^{10}$ or $\boldsymbol{\xi} = (P^s_{out}, P^t_{in}, T^t_{in}) \in \mathbb{R}^3$, $\boldsymbol{u}$ denoting the components of the KL expansion describing the random field modeling geometric variations due to manufacturing processes.

Condition	P^t_{in} [bars]	T^t_{in} [K]	P^s_{out} [bar]	$\boldsymbol{u}$
Nominal (DO)	8.0	543.65	1.333	$0_{\mathbb{R}^7}$
Scenario 1 (RO-1)	$\mathcal{U}[7.95, 8.05]$	$\mathcal{U}[543.3, 544.0]$	$\mathcal{U}[1.333, 2.015]$	$\mathcal{N}(0, \boldsymbol{I}_7)$
Scenario 2 (RO-2)	$\mathcal{U}[7.95, 8.05]$	$\mathcal{U}[543.3, 544.0]$	$\mathcal{U}[1.333, 2.015]$	-

A comprehensive Uncertainty Quantification (UQ) assessment is performed *a posteriori* on the resulting optimized profiles based on 100 CFD samples, chosen using LHS in the stochastic space (see Razaaly et al. (2019b) for full details). It permits in particular to reconstruct the full Probability Density Function (PDF) of scalar QoIs and statistics of high dimensional QoIs. A dedicated convergence analysis performed on the baseline profile has shown that using 24 CFD samples permit to recover first order statistics with a relative error lower than 1%, the reference value obtained with 200 CFD simulations. Additionally, PDFs of different QoIs reconstructed using 100, 200 or 400 CFD simulations almost merge, while higher order statistics corresponds, with negligible relative errors.

4. OPTIMIZATION FRAMEWORK

The optimization framework is described in the following, based on bayesian optimization (E. Brochu and de Freitas (2010)).

4.1 Deterministic Optimization

The following constrained Deterministic Optimization (DO) is performed, at nominal conditions $\boldsymbol{\xi}_0$ (Table 1):

$$\begin{aligned} \text{Minimize} \quad & \Delta P(\boldsymbol{x}, \boldsymbol{\xi}_0) \\ \text{s.t.} \quad & \dot{m}(\boldsymbol{x}, \boldsymbol{\xi}_0) \in \Delta \\ & \boldsymbol{x} \in \Omega \end{aligned} \tag{1}$$

where $\Delta = [0.98, 1.02]$, $\boldsymbol{x}$ the vector of control points belonging to the design space $\Omega \subset \mathbb{R}^d$, $d = 9$. We rely on a classical Bayesian framework for Surrogate-Based Optimization (SBO), the popular Efficient Global Optimization (EGO, E. Brochu and de Freitas (2010)). Specifically, we consider a sequential approach where Gaussian Process (GP)-based surrogate models C. E. Rasmussen (2006) are built to approximate the objective and constraint functions. At each step of the sequence, a new design is obtained maximizing the so-called Expected Improvement (EI), whose expression analytically depends on the current GP surrogate. This approach permits both to identify promising regions of Ω and to explore portions of the design space characterized by high uncertainty on the GP surrogate. The latter one is initialized considering $5d$ initial samples in the design space, using LHS. The GP is iteratively refined as the proposed sample is evaluated by means of CFD, the procedure being repeated until either a stopping criterion is met, or the budget of CFD evaluations is exhausted.

4.2 Robust Optimization

The following constrained mean-based RO is solved, for both the uncertainties reported as scenarios 1 and 2 in Table 1:

$$\begin{aligned} \text{Minimize} \quad & \mathbb{E}[\Delta P(\boldsymbol{x}, \boldsymbol{\xi})] \\ \text{s.t.} \quad & \mathbb{E}[\dot{m}(\boldsymbol{x}, \boldsymbol{\xi})] \in \Delta \\ & \boldsymbol{x} \in \Omega \end{aligned} \tag{2}$$

where $\mathbb{E}$ refers to the expectation operator induced by the random vector $\boldsymbol{\xi}$ (Table 1). A simple nested (a.k.a. bi-level) approach is considered here, following a strategy very similar to DO. Indeed, instead of performing a single CFD evaluation at nominal conditions for each design vector $\boldsymbol{x}$ representing a blade profile, the following procedure is conducted. An UQ analysis is performed based on 24 CFD simulations, yielding an estimation of the mean of both the objective and constraint functions.

5. RESULTS

In this section, we present the results of the optimization problems described in Section 4. DO, RO considering scenario 1 and RO (Table 1) considering scenario 2 resulting profiles will be referred respectively as DO, RO-1 and RO-2 profiles. The UQ assessment proposed are based on the variability induced by scenario 1, since including both operational conditions and geometric variability.

5.1 Optimization Procedure

Figure 1 reports optimization convergence curves featuring the best normalized objective function as a function of the number of designs considered. Note that during the optimization process, some CFD computations lead to an objective function value lower than the current best, but not retained as such since violating one constraint. The DO requires approximately 80 designs (80 CFD simulations) to reach convergence, while ROs feature slower convergence trends; in both cases, the optimization algorithm stops after around 200 designs ($\sim 24 \times 200$ CFD simulations). This phenomemon is largely due to the increased numerical noise associated to first order statistics of the objective function ΔP, w.r.t. its deterministic evaluations, leading to a global surrogate used directly within the bayesian deterministic optimization framework of lower quality, hence delaying the convergence process. This observation is indeed emphasized for RO-1, which considers more uncertainties than RO-2. Inviscid CFD simulations were performed on a cluster equipped with Intel(R) Xeon(R) CPU X5650 at 2.67 GHz, using 24 CPUs in parallel. An adaptive CFD between 10 and 100 is used. Simulations are assumed to be converged when density residuals are decreased by ten orders of magnitude (which always holds here), with a computation cost of around 500s for 1 CFD simulation on 1CPU.

5.2 Probability Density Functions

To gain further understanding of the ROs over DO results, the Probability Density Functions (PDF) of ΔP, of the normalized mass flow rate $\dot{m}$ and of the mass-average outlet angle w.r.t. the axial direction α_{out} are presented in Figure 2. W.r.t. the DO profile, the *robust* optimal profiles RO-1 and RO-2, are characterized by PDF distributions of ΔP significantly displaced towards the lowest values and featuring tighter support. This is an evident marker of a reduced variability, namely of an improved robustness of the cascade performance w.r.t. the variability induced by operational and geometric variability. The two RO blades are substantially equivalent from a stochastic perspective, even though RO-2 seems to feature a slightly reduced variability over RO-1 in terms of ΔP. The $\dot{m}$ PDFs have a similar gaussian-like shapes, the mean of the two RO's profiles being approximately at the upper limit of the constraint (102%),

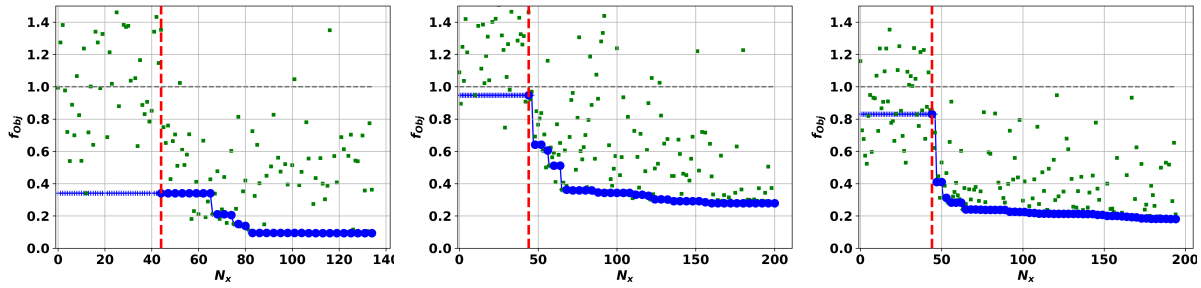

(a) DO (b) RO-1 (c) RO-2

Figure 1: Convergence curves during the optimization. The red vertical line indicates the number of designs to consider necessary for the GP-based DO/RO initialization. Green (resp. blue) dots indicate the (resp. best) objective function. For ROs (b-c), the corresponding number of CFD evaluations is $24 \times N_x$, N_x being the number of designs.

while the DO profile features a mass-flow rate at nominal conditions close the lower limit: its PDF considering uncertainties tends to shift its mean value towards the reference mass-flow rate. The PDFs associated to the outlet angle α_{out} features noticeable observations. Indeed, the DO profile distribution exhibits a slightly narrower support w.r.t. RO' ones, sign of a lower variability, even if one might expect the opposite behaviour due to their respective deterministic and robust design aspects. However, from a pure mathematical perspective, this finding is totally compatible with the adopted formulation, since the flow angle is not accounted in the optimization process.

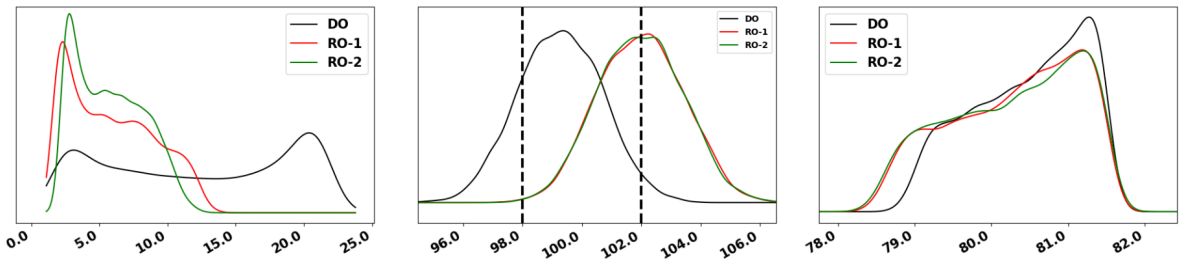

(a) ΔP (kPa) (b) $\dot{m}$ (%) (c) α_{out} (°)

Figure 2: PDF of the objective function ΔP, the normalized mass flow rate (constraint function) $\dot{m}$ and outlet angle α_{out}, for the three resulting designs (DO, RO-1, RO-2). (b) Vertical lines correspond to the upper/lower bounds of the mass-flow constraint, formulated in the mean value for ROs, and nominal one for DO.

5.3 Blade Profiles

The optimal blade profiles obtained with the three optimization processes discussed above are shown in Figure 3.

The DO profile features a slightly more accentuated curvature in the diverging section the bladed channel, i.e. between the (sonic) throat and the cascade opening, with a almost straight suction side in the region of unguided turning. As discussed in previous design exercises performed on this cascade (with different nominal conditions though), such shape eliminates the strong shock originated in the baseline cascade flow[1] with beneficial effects on both the uniformity of the

[1]The baseline profile aerodynamics is not documented here for sake of brevity, since largely discussed in previously cited works.

pressure field downstream of the cascade and its related loss.

The RO' designs, very similar to each other, differ significantly from the DO layout, especially at the Trailing Edge (TE), in the throat region and on the suction side. Nevertheless, the general action of the optimization is analogous to the one discussed for DO. The throat opening is now slightly larger w.r.t. the one of the DO case, consistently with the 2% larger mass-flow rate already commented in Figure 2(b). Crucial differences emerge on the pressure and suction side of the blade in correspondance to the blade opening. As a matter of fact, the cross-section of the cascade passage is much lower to the one of the DO blade. This is consistent with the fact that, in the uncertain scenario, the mean outlet pressure is higher than the deterministic value. As a result, the RO procedure tends to reduce the the throat-to-opening area ratio w.r.t. the DO. The two RO blades also feature minor, though visible differences, mostly concentrated in the diverging section of the suction side, i.e. the one most affecting the ΔP QoI, indicating that the manufacturing tolerances may have an impact in the definition of the optimal blade shape in the most sensitive regions.

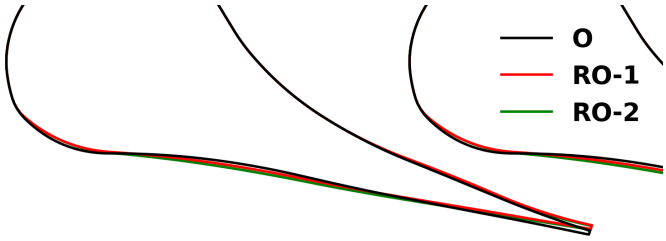

(a)

Figure 3: Blade profiles comparison.

5.4 UQ Physical Flow

On the basis of physical considerations, the aerodynamics of the optimized blades are investigated under uncertain flow conditions. The mean and standard deviation of Mach contours[2] are reported respectively in Figures 4 and 5. It's worth mentioning that the analysis of contour statistics is not straightforward, as the Mach number value at each node of the mesh results from a statistical procedure equivalent to a locally *independent* UQ treatment, and hence, the resulting field is not referring to a specific condition. As a consequence, the analysis of detailed features might be misleading and only general trends will be discussed. Moreover, the analysis of the flow field at nominal conditions and perfect blade geometry relevance is questionable within this uncertain treatment, as the discussion of the PDFs of QoIs suggest. The mean distribution computed for the DO exhibits the detrimental effects due to the rear shock generated at the TE for the whole variability range. The throat-to-opening area ratio of the DO blade, larger than the RO ones, is indeed too large for most of the realizations occurring within the prescribed variability, so the cascade mostly operates in post-compression conditions. RO cascades significantly reduce the aforementioned effect, featuring an almost uniform mean Mach number distribution at the cascade exit. This consideration is, to some extent, less established for the RO-2 layout, whose optimization formulation does not incorporate geometric variability. In order to highlight the local variability over the flow field, the Mach Standard Deviation (Std) is reviewed in the following. A large Std is established in the region influenced by the compression wave/shock and in the wake. A visible reduction in Std is observed when investigating the RO profiles, which however exhibit a Std flow pattern similar to that of DO.

[2]under uncertain conditions, evaluated using a Kriging-Principal Component Analysis(PCA)-MC framework as detailed in Razaaly et al. (2019b), considering around 70 modes targeting to a cumulative energy conservation beyond 99.99%.

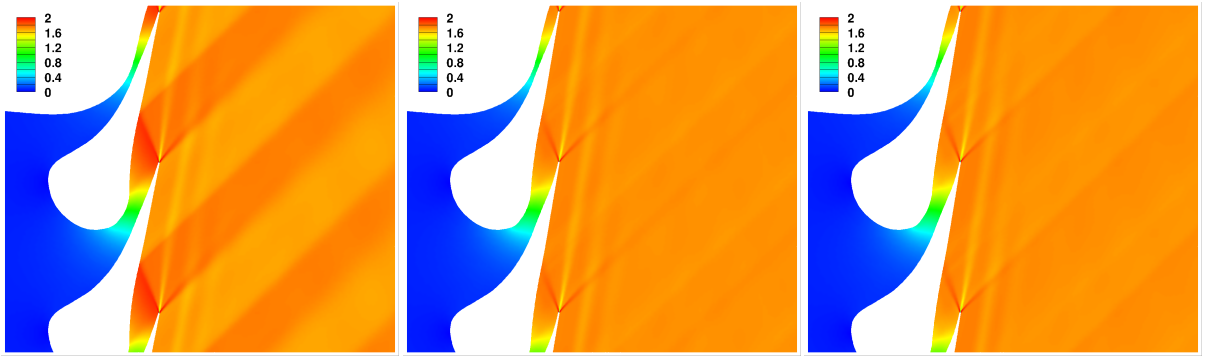

(a) DO (b) RO-1 (c) RO-2

Figure 4: Mean Mach contours (first order statistics).

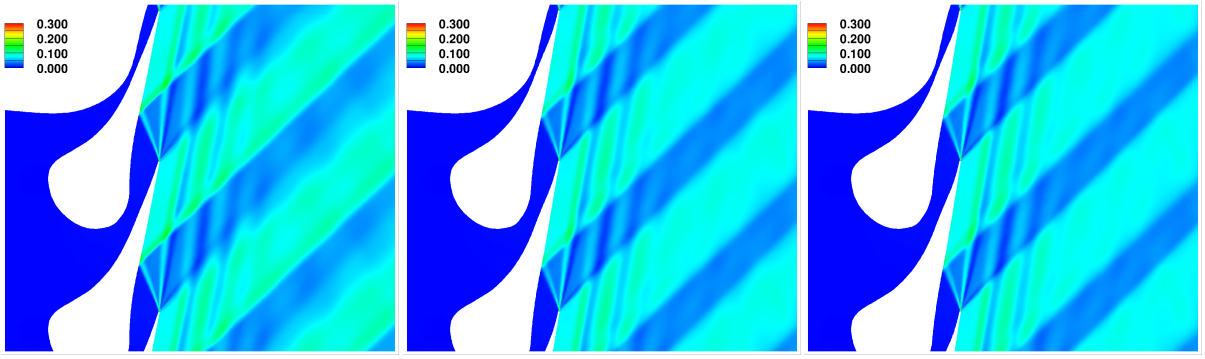

(a) DO (b) RO-1 (c) RO-2

Figure 5: Standard Deviation Mach contours (second order statistics).

Additionally, the first order statistics (mean) of the Mach number distribution over the blade surface is reported in Figure 6, the variability being represented by the two standard deviation bounds centered in the mean. The mean distributions appear relatively similar among the DO and RO blades; however a further inspection reveals a higher peak Mach number for the DO blade, followed by a post-compression effect already recalled, penalizing the cascade performane in off-design conditions. Moreover, the variability appears also larger for the DO than for the two ROs. Interestingly, but not surprisingly due to the chocked-flow condition of all the cascades, the variability in Mach number is negligible everywhere but on the suction side, downstream of the blade opening. In this region, high variability characterizes all the optimal blades, but the DO one exhibit the largest std. This evidence, combined with the corresponding one visible in terms of Mach number contours, suggests that the *robust* design procedure allows reducing the variability in the entire flow field.

6. CONCLUSION

This paper has illustrated a robust optimization framework for turbine design considering uncertainties due to large variations of operating conditions and geometric variability due to manufacturing. Two robust designs, one of them accounting for operational variability only, are compared to a deterministic optimal cascade obtained at nominal conditions. In particular, the influence on the PDF of QoIs as well as on the flow field is examined. The technique makes use of a bi-level surrogate-based algorithm, combining a classic bayesian optimization loop and classic Kriging-based UQ tools. The blade profile is parametrized via B-splines, and

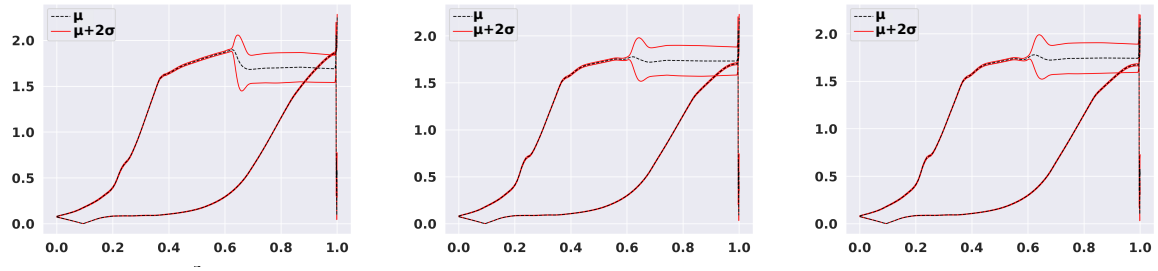

(a) DO (b) RO-1 (c) RO-2

Figure 6: Mach number distribution over the optimized blades: mean and two standard deviation bounds.

a mesh-deformation tool based on RBF permits to handle modified geometries within the optimization process and UQ analysis accounting for geometric variability modeling. The results have demonstrated that, for a benchmark supersonic ORC turbine, the robust optimization provides relevant benefits over classic deterministic optimization at nominal point, in terms of the variability of the performance QoI. The proposed methodology is, therefore, able to reduce the sensitivity to operating condition of the performance of a typical ORC first-stage nozzle. Since such component provides a very relevant contribution to the whole turbine aerodynamic loss, the obtained design has the potential for improving significantly the performance of the whole ORC power system, by reducing the efficiency penalty at part-load. Short-term future studies will investigate further test-cases featuring different variability and design conditions to better explore the potential of RO for ORC turbines.

REFERENCES

Betz, W., Papaioannou, I., and Straub, D. (2014). Numerical methods for the discretization of random fields by means of the karhunen–loève expansion. *Computer Methods in Applied Mechanics and Engineering*, 271:109–129.

C. E. Rasmussen, C. W. (2006). *Gaussian Processes for Machine Learning*. the MIT Press.

Colonna, P., Harinck, J., Rebay, S., and Guardone, A. (2008). Real-gas effects in organic rankine cycle turbine nozzles. *Journal of Propulsion and Power*, 24(2):282–294.

De Boer, A., Van der Schoot, M., and Bijl, H. (2007). Mesh deformation based on radial basis function interpolation. *Computers & structures*, 85(11-14):784–795.

Dow, E. A. and Wang, Q. (2014). Optimal design and tolerancing of compressor blades subject to manufacturing variability. In *16th AIAA Non-Deterministic Approaches Conference*, page 1008.

E. Brochu, V. M. C. and de Freitas, N. (2010). A tutorial on bayesian optimization of expensive cost functions, with application to active user modeling and hierarchical reinforcement learning. *arXiv:1012.2599*.

Economon, T. D., Mudigere, D., Bansal, G., Heinecke, A., Palacios, F., Park, J., Smelyanskiy, M., Alonso, J. J., and Dubey, P. (2016). Performance optimizations for scalable implicit {RANS} calculations with {SU2}. *Computers & Fluids*, 129:146 – 158.

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

Giles, M. B. (1990). Nonreflecting boundary conditions for euler equation calculations. *AIAA journal*, 28(12):2050–2058.

Macchi, E. (2013). The choice of working fluid: the most important step for a successful organic rankine cycle (and an efficient turbine). In *Second International Seminar on ORC Power Systems, Rotterdam, The Netherlands, Oct*, pages 7–8.

Meroni, A., Andreasen, J., Persico, G., and Haglind, F. (2018). Optimization of organic rankine cycle power systems considering multistage axial turbine design. *Applied Energy*, 209:339–354.

Persico, G., Rodriguez-Fernandez, P., and Romei, A. (2019). High-fidelity shape-optimization of non-conventional turbomachinery by surrogate evolutionary strategies. *Journal of Turbomachinery*, 141(8):081010.

Pini, M., Persico, G., and Dossena, V. (2014). Robust adjoint-based shape optimization of supersonic turbomachinery cascades. In *ASME Turbo Expo 2014: Turbine Technical Conference and Exposition*, pages V02BT39A043–V02BT39A043. American Society of Mechanical Engineers.

Pini, M., Persico, G., Pasquale, D., and Rebay, S. (2015a). Adjoint method for shape optimization in real-gas flow applications. *ASME Journal of Engineering for Gas Turbines and Power*, 137(3).

Pini, M., Persico, G., Pasquale, D., and Rebay, S. (2015b). Adjoint method for shape optimization in real-gas flow applications. *ASME Journal of Engineering for Gas Turbines and Power*, 137(3).

Razaaly, N., Gori, G., Iaccarino, G., and Congedo, P. (2019a). Optimization of an orc supersonic nozzle under epistemic uncertainties due to turbulence models. In *GPPS 2019*.

Razaaly, N., Persico, G., and Congedo, P. M. (2019b). Impact of geometric, operational, and model uncertainties on the non-ideal flow through a supersonic orc turbine cascade. *Energy*, 169:213–227.

Vitale, S., Albring, T. A., Pini, M., Gauger, N. R., and Colonna, P. (2017). Fully turbulent discrete adjoint solver for non-ideal compressible flow applications. *Journal of the Global Power and Propulsion Society*, 1:Z1FVOI.

Vitale, S., Gori, G., Pini, M., Guardone, A., Economon, T. D., Palacios, F., Alonso, J. J., and Colonna, P. (2015). Extension of the SU2 open source CFD code to the simulation of turbulent flows of fluids modelled with complex thermophysical laws. Number AIAA Paper 2015-2760.