

ADJOINT-BASED SHAPE OPTIMIZATION FOR INDUSTRIAL HEAT EXCHANGERS

Tobias Kattmann¹, Ole Burghardt², Nicolas R. Gauger³, Nijso Beishuizen^{4,5}

¹ TU Kaiserslautern
Paul-Ehrlich-Straße 34, 67663, Kaiserslautern, Germany
tobiaskattmann@gmail.com

² TU Kaiserslautern
Paul-Ehrlich-Straße 34, 67663, Kaiserslautern, Germany
ole.burghardt@scicomp.uni-kl.de

³ TU Kaiserslautern
Paul-Ehrlich-Straße 34, 67663, Kaiserslautern, Germany
nicolas.gauger@scicomp.uni-kl.de

⁴ Bosch Thermotechniek B.V.
Postbus 3, 7400 AA Deventer, The Netherlands
nijso.beishuizen@nl.bosch.com

⁵ Eindhoven University of Technology
Postbus 513, 5600 MB Eindhoven, The Netherlands
n.a.beishuizen@tue.nl

Key words: Multiphysics Problems, CHT, Discrete Adjoint, Industrial Applications

Abstract. *This paper showcases the multiphysics discrete adjoint solver of the open Source solver SU2 at the example of a pin-array heat exchanger, simulated via a simplified 2D unit cell to evaluate the pin-shapes performance. The shape derivatives of the utilized conjugate heat transfer case are validated against finite differences and show good agreement. In a constrained shape optimization, the presented workflow proves its robustness over a full optimization. The utilized mesh deformation setup for the investigated downstream periodic and internal flow is presented in detail, as retaining a valid mesh under deformation was found to be challenging.*

1 INTRODUCTION

Heat exchangers are frequently used in applications like power electronics cooling in e.g. electric vehicles or residential gas/hydrogen boilers. A commonly used design choice is the pin-fin array consisting of tens to thousands of individual pins. Numerous factors determine the pin-array's performance: in- and outlet geometries of the heat exchanger, total number of pins, arrangement of the pins (e.g. inline or staggered) to name a few and last but not least the pin-shape. While a reasonable overall design can be achieved by e.g. design of experiments (DOE), gradient based shape optimization of the pin-shape promises additional performance gains. Although cost-competitive manufacturability, especially for large quantities is a concern, the rise of 3D metal printing processes (e.g. Binder-Jetting, SLM, etc.)

could be a key enabler for the usage of free-shape optimized pins in pin-fin heat exchangers.

In the present work, a conjugate heat transfer case is considered with one fluid and one solid zone. In the fluid zone the incompressible Navier-Stokes equations, RANS to be precise, are solved together with an SST turbulence model. The used density based-based incompressible solver is explained in-depth in [3]. Heat conduction is solved in the solid, and the respective fluid-solid coupling mechanisms used in SU2 are outlined in [1].

The discrete adjoint method proved to be efficient for the computation of the required shape gradients for a large number of design variables. While the discrete adjoint method is widely used for single physics, e.g. most prominently aerodynamic wing optimization, in the present heat exchanger application, conjugate heat transfer (CHT) is simulated: solid heat conduction is coupled to fluid flow via an interface boundary. This paper makes use of recently developed methods for multiphysics discrete adjoint [2], where the testcase presented in this paper was already introduced with a preliminary gradient validation.

This contribution is organized as follows. Section (2) briefly explains the derivation of the multiphysics discrete adjoint solver for two, following the approach presented in [1], and some notes on the utilization of automatic differentiation in that application. Section (3) gives a detail introduction of the examined testcase with an emphasis on the mesh deformation setup. Section (4) presents the results of the gradient validation and optimization. This paper concludes with section (5) and a remark on reproducibility (6).

2 DISCRETE ADJOINT METHOD

The PDE-constrained optimization problem of the present two zone CHT problem is formulated as follows:

$$\begin{aligned} \min_{\alpha} \quad & J(U_1(\alpha), U_2(\alpha), X(\alpha)), \\ \text{subject to} \quad & U_1(\alpha) = G_1(U_1(\alpha), U_2(\alpha), X(\alpha)), \\ & U_2(\alpha) = G_2(U_1(\alpha), U_2(\alpha), X(\alpha)), \\ & X(\alpha) = M(\alpha) \end{aligned} \tag{1}$$

where J , the objective function, depends on the mesh state $X(\alpha)$ and the PDE state variables U_i of the discretized primal problem in each zone, i.e. pressure, velocities and temperature for the fluid zone and just temperature for the solid zone. Shape modifications will result in an altered mesh and PDE state, therefore both are implicit functions of the design variables α . The mesh results from a surface and volume mesh routine, where the former is executed by free form deformation (FFD) and the latter by linear elasticity equations. This grid deformation process is denoted by the operator M . Lastly, the fixed-point operator for the governing equations is defined as:

$$U_i^{n+1} = G_i(U_1^n, U_2^n), \tag{2}$$

with the subscript i representing either of the zones and the superscript n being the iteration number. The fixed-point operators G_i depend on the PDE solutions of both zones due to the necessary interface coupling in CHT. Note, that in (2) the fixed G_i depends on the solution states at time step n , where in practice, sequentially solved zones might already use the updated solution states at $n + 1$.

The constrained optimization problem (1) is transformed into an unconstrained optimization problem

by setting up the Lagrangian:

$$L(\alpha, U_1, U_2, X) = J(U_1, U_2, X) + \lambda_1^\top (G_1(U_1, U_2, X) - U_1) + \lambda_2^\top (G_2(U_1, U_2, X) - U_2) + \phi^\top (M(\alpha) - X), \quad (3)$$

with λ_i and ϕ being the adjoint variables for the PDE state and mesh. These adjoint variables are defined by the following disappearing partial derivatives:

$$\frac{\partial L}{\partial U_1} = 0, \quad \frac{\partial L}{\partial U_2} = 0, \quad \frac{\partial L}{\partial X} = 0. \quad (4)$$

Evaluation of these partial derivatives results in a coupled fixed-point iteration system for the adjoint PDE states in equations (5) and (6). The adjoint mesh state can be evaluated by equation (7) once the adjoint PDE states are computed:

$$\lambda_1^\top = \frac{\partial J}{\partial U_1} + \lambda_1^\top \frac{\partial G_1}{\partial U_1} + \lambda_2^\top \frac{\partial G_2}{\partial U_1}, \quad (5)$$

$$\lambda_2^\top = \frac{\partial J}{\partial U_2} + \lambda_1^\top \frac{\partial G_1}{\partial U_2} + \lambda_2^\top \frac{\partial G_2}{\partial U_2}, \quad (6)$$

$$\phi^\top = \frac{\partial J}{\partial X} + \lambda_1^\top \frac{\partial G_1}{\partial X} + \lambda_2^\top \frac{\partial G_2}{\partial X}, \quad (7)$$

$$(8)$$

Finally, once equations (5), (6) and (7) are satisfied, the formula for evaluating the total derivative of the objective function becomes:

$$\frac{dJ}{d\alpha} = \frac{\partial L}{\partial \alpha}(\alpha, U_1(\alpha), U_2(\alpha), X(\alpha)) = \phi^\top \frac{\partial M}{\partial \alpha} \quad (9)$$

The terms in the adjoint equations (5), (6) and in (7) are evaluated using automatic differentiation (AD). SU2 makes use of the operator overloading AD-tool CoDiPack [6]. Consider the second term used in (5):

$$\lambda_1^\top \frac{\partial G_1}{\partial U_1}. \quad (10)$$

The evaluation of that expression can be formulated in a four stage process. First, the computation to evaluate G_1 starting from U_1 , i.e. one iteration of the zonal solver, is recorded and stored onto a so-called tape. In the second step, the current adjoint solution λ_1 is set as the adjoint of G_1 , often referred to as 'seeding'. Note that 'adjoint' is an AD-specific term as well and should not be confused with the adjoint solution λ_i . The derivatives are evaluated in the third step by reversing the previously recorded tape. Finally, as the fourth step, the solution to the above considered expression is extracted from the adjoint of U_1 , here the AD-specific adjoint is meant again.

3 TESTCASE DESCRIPTION

A typical pin array for e.g. power electronics cooling can consist of tens up to a couple of thousand individual pins. Instead of a simulation on the entire pin array, a common simplification is to only use a characteristic unit cell, see Fig. (1a). Through mirroring at symmetry boundaries and translating at periodic boundaries, one recovers an approximation to the initial/full array. It is however within reason to presume that general design rules can be derived from such a reduced, fast to evaluate, unit cell approach.

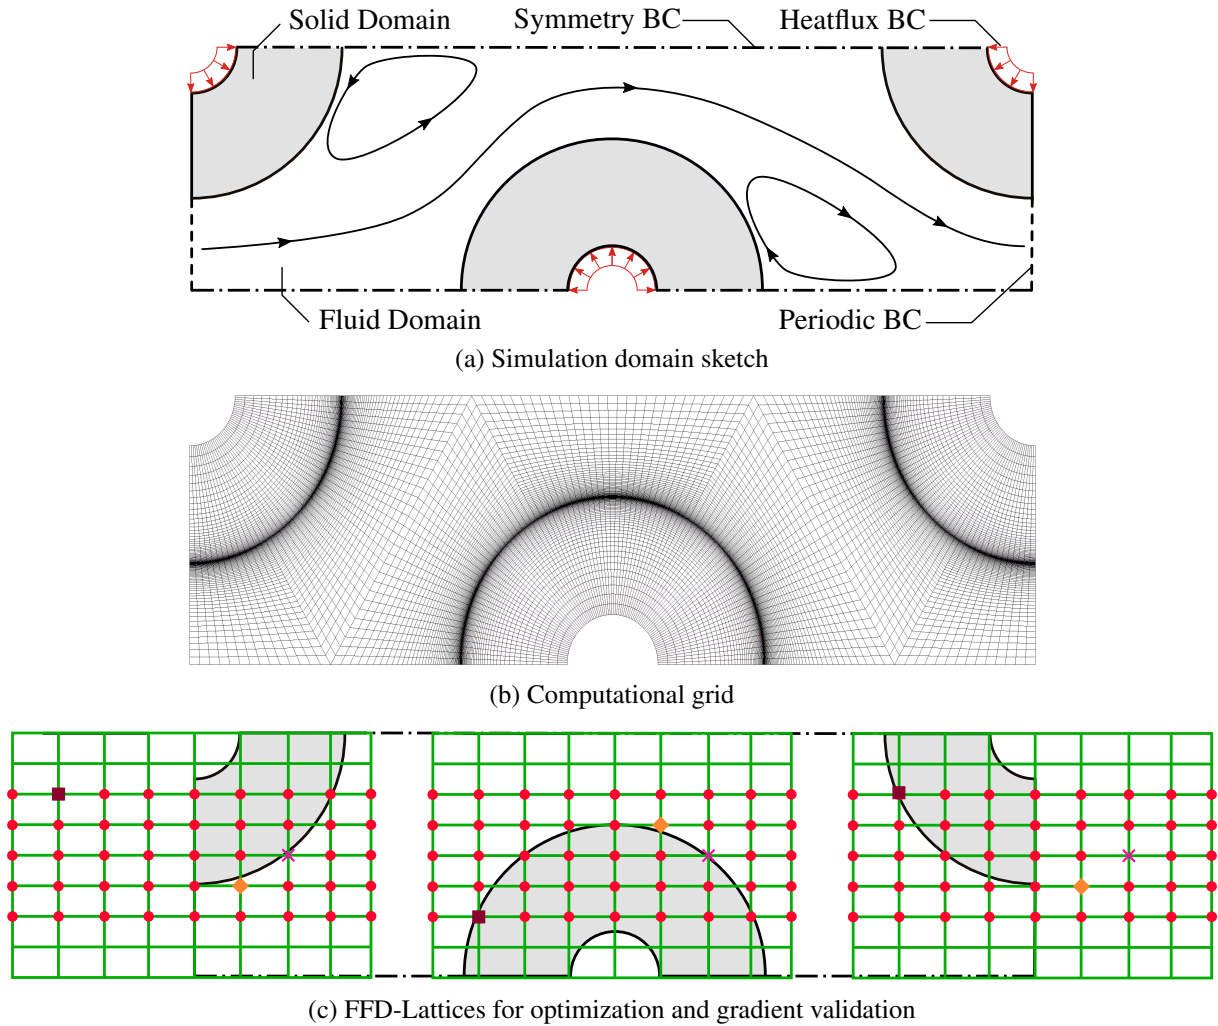


Figure 1: (a) Domain sketch pointing out the boundary conditions. (b) The fully structured grid consists of a total of 15.224 quadrilateral elements. (c) Pin enclosing FFD-boxes. The front and back box are mirrored compared to the middle one, ensuring equal deformation. Three special symbols (brown square, orange diamond, pink cross) are used to indicate equal points. Only points with a symbol (e.g. red dot) are allowed to deform and only in y -direction as a further restriction.

3.1 Geometry, boundary conditions and mesh description

The geometry of the simulation domain is fully characterized by three quantities: distance between pin midpoints 6.44 mm (three neighboring pins form an equilateral triangle which is oriented such that the downstream pin rows are staggered), the inner pin radius 0.6 mm and outer pin radius 2 mm.

Inlet and outlet in the fluid domain are prescribed as periodic boundaries, and streamwise periodicity (or fully developed flow) is used [5]. The massflow over the domain is set to 0.85 kg/s and the inlet bulk temperature to 338 K. The massflow is achieved by iterating over a pressure drop that is ultimately used in the associated source term, and the inlet bulk temperature is enforced by and outlet heat sink.

Symmetry boundaries on top and bottom are indicated in Fig. (1a) by dash-dotted lines, and handling of the interfaces between fluid and solid is described in literature [1]. On the inner pin arcs of the solids a heatflux of $5e5 \text{ W/m}^2$ is prescribed, and all remaining walls are adiabatic. Note, that during optimization, the inner pin surface remains clamped such that the integrated heatflux, i.e. the power, into the domain does not change.

The computational mesh is fully structured, e.g. only quadrilaterals are used, and consists of 8477 elements in the fluid and 6747 elements in the solid, see (1b). A $y^+ < 1$ is obtained everywhere on the walls. Whether a mesh independent solution is achieved with the grid was not tested. The purpose of the grid is to verify the gradient accuracy of the discrete adjoint method and to showcase the optimization abilities.

3.2 Geometry and mesh deformation setup

Ensuring a suitable and robust geometry and mesh deformation setup for the optimization was one of the major challenges for the present testcase. This is mostly due to the two quarter pins on the periodic interface that have to undergo the same deformation as the half pin in the middle of the domain. This is to guarantee the required matching surface boundary and mesh on the periodic interface. That approach also enforces, that the simulated pin-array consists of a single pin-shape only. The main focus here will be on the surface deformation procedure, as for the volume mesh deformation all here used surfaces are simply considered clamped for the linear-elasticity based volume deformation algorithm.

The surface mesh and volume mesh deformation are separated in SU2. For the surface mesh deformation, free form deformation (FFD) boxes are used and their setup for the presented testcase is shown in figure (1c). The green lattices represent separate FFD-boxes and the symbols (e.g. red dots) indicate actively used points, the remaining points stay clamped. Two rows are clamped as B-Splines of 3rd order in y-direction are used for the FFD-boxes, which ensures no point movement of the top and bottom symmetry boundaries. A movement of FFD-points in x-direction is not considered in this work, as this would create further challenges when it comes to retaining a valid shape for the connected quarter pins. The front and back box are mirrored, and of course translated, along the y-axis compared to the middle FFD-box. Some symbols (brown square, orange diamond, pink cross) are used for easy identification of the same point in each FFD-box.

In order to apply an equal movement to each FFD-box, one can either just compute sensitivities for one FFD-box (in this case the middle one) and map these appropriately onto the other FFD-boxes, or compute sensitivities for all FFD-boxes individually, average the values and then apply these averaged values to all FFD-boxes. The former is used for the optimization presented section (4.2).

Each surface patch is either clamped, i.e. remains unaltered, or is affected by the FFD-Box movement. The inner-pin circle segments and the symmetry boundaries are clamped, while consequently all other surface segments (fluid-solid interfaces and periodic front- and backface) are moved by the FFD-box. An alternative, only deforming the fluid-solid interface and allowing in-plane movement of surface mesh points in the volume mesh deformation phase, was not viable due to a resulting non-matching mesh on the periodic interface.

3.3 Material properties

Constant fluid properties are density $\rho = 1045 \text{ kg/m}^3$, specific heat $c_p = 3540 \text{ J/(kg K)}$, laminar viscosity $\mu_{dyn} = 1.385e-3 \text{ Pa s}$, laminar and turbulent Prandtl numbers $Pr_{lam} = 11.7$ and $Pr_{turb} = 0.9$. The fluid's

material properties represent Glysantin in a 50/50 mixture with water, a commonly used liquid for automotive cooling circuits. The constant material properties in the solid represent a non further specified aluminum-alloy with density $\rho = 2719 \text{ kg/m}^3$, specific heat $c_p = 871 \text{ J/(kg K)}$ and thermal conductivity $\lambda = 200 \text{ W/(m K)}$.

3.4 Numerical methods

Convective fluxes in the incompressible solver are discretized using a flux-difference-splitting (FDS) with second order accuracy achieved by a MUSCL approach [3]. No limiters are used for variable reconstruction, and gradients are computed via the Green-Gauss theorem. The pseudo-time derivative is solved using an implicit Euler scheme. The Menter SST turbulence model without wall function is used. Gradients in the heat conduction solver are computed via the Green-Gauss theorem, and the pseudo-time derivative is evaluated with the implicit Euler scheme. In both solvers, the resulting linear system is solved using FGMRES with an ILU preconditioner.

4 RESULTS

4.1 Gradient Validation

The sensitivity computed via the discrete adjoint method is validated against 1st order forward differences for a specific design variable. The chosen design variable is the y -coordinate deflection of the FFD-box point that is indicated with an orange diamond on the middle pin only, see figure (1c). Table (1) shows the gradient validation results for the average temperature (columns 2 and 3) and pressure drop objective functions over multiple finite difference stepsizes. Average temperature is evaluated on the surface of the middle inner pin and pressure drop between the two periodic interfaces which represent the fluid 'inlet' and 'outlet'. The first row lists the sensitivity reported by the discrete adjoint (DA) method. In the remaining rows the sensitivity of the finite difference scheme with the specified stepsize is shown, matching digits are colored black and non-matching in red, as well the respective relative difference to the discrete adjoint solution:

$$\text{rel. diff.} = \frac{|DA_{Sens} - FD_{Sens}|}{|DA_{Sens}|}. \quad (11)$$

With an outer pin radius of $1e-3\text{m}$, the largest stepsize of $1e-4\text{m}$ is not expected to provide an accurate comparison but serves as an indication how aggressive the first geometry modification in the optimization run for this specific case can be chosen. The optimal stepsize for the average temperature objective is $1e-6\text{m}$ with a relative difference between the discrete adjoint and finite difference sensitivity of just 0.0012% . For pressure drop, this optimum is found at $1e-7\text{m}$ with 0.0754% relative difference. For other design variables not shown here, the relative differences have a similar magnitude. Note the opposite sign of the gradient between the two objectives. Moving the FFD-box point up (enlarging the pin) leads to a decreased average temperature but increased pressure drop. This conflict will be handled in more detail in the following optimization section (4.2).

Overall, the agreement between DA and FD sensitivities is very good, and the gradients are certainly suitable for accurately finding local optima. A clear justification for the remaining deviations, especially for pressure drop, cannot be given.

Table 1: Gradient validation for average temperature and pressure drop objective functions over multiple finite difference stepsizes. The first row shows the discrete adjoint (DA) results. Relative difference is computed after equation (11).

FD step [m]	avg. temp. sens. [K/m]	rel. diff. [%]	pressure drop sens. [Pa/m]	rel. diff. [%]
DA	-2.340265e3	-	1.5445492e4	-
1e -4	-3.111159e3	32.940	2.9906661e4	93.627
1e -5	-2.433546e3	3.9859	1.6744531e4	8.4105
1e -6	-2.340294e3	0.0012	1.5551885e4	0.6888
1e -7	-2.330330e3	0.4245	1.5433860e4	0.0753
1e -8	-2.329331e3	0.4672	1.5422077e4	0.1516

4.2 Optimization

After successfully validating the gradients computed via the discrete adjoint method, a constrained optimization is performed. The objective function being average temperature on the inner surface of the middle pin, as introduced in the gradient validation, and the inequality constraint of the pressure drop not exceeding the value of the initial geometry (208.023Pa). The design variables are FFD-box points, as presented in section (3.2).

The optimization framework FADO [4] was used in conjunction with the SciPy SLSQP optimizer. Figure (2) shows the optimization history of the monitored primal values as well as the gradient norms of both, objective function and constraint. Points are drawn for feasible values of converged simulations and are only connected with a line segment, if subsequent iterations are feasible. Infeasible values for the primal occur, if e.g. the optimizer proposes a self-intersecting grid or even just a bad quality mesh, resulting in divergence. Interruptions in the gradient plot occur as well if the optimizer decides to not compute an additional gradient, either because the previous design is not accepted and the second to last is recycled or simply no considerable improvement can be achieved without violating the constraint, see the final design iterations.

In the first ~ 30 design iterations the objective function drops in a zigzag course while the constraint oscillates around the upper constraint limit. The optimizer first follows an OF dominated gradient that leads to a constraint violation, which is consolidated in a following design iteration that predominantly follows the constraint gradient. This effect shows the conflicting nature of the chosen objective and constraint for this testcase, i.e. the gradients, by and large, are showing in opposite directions. In the second ~ 30 design iterations the modifications are less significant and only minor additional gains are made until the optimization converges. The average temperature is decreased by ~ 1.1 K while the pressure drop remained constant, i.e. fulfills the imposed inequality constraint.

The plotted gradient norms are normalized by their initial norm to allow for an easy comparison. The norm for the objective function drops by just 35% as the pressure drop constraint limits the feasible design space dramatically. As an interesting, but otherwise not important, sidenote the constraints' gradient norm drops by 20% as well.

The final pin design of the constrained optimization is shown in the top right corner of figure (3) together with the initial design and unconstrained optimization results for both involved objectives. Additionally, the FFD-box lattice is given for each design.

As expected, the pressure drop (dp) optimization results in the pin collapsing and is only limited by

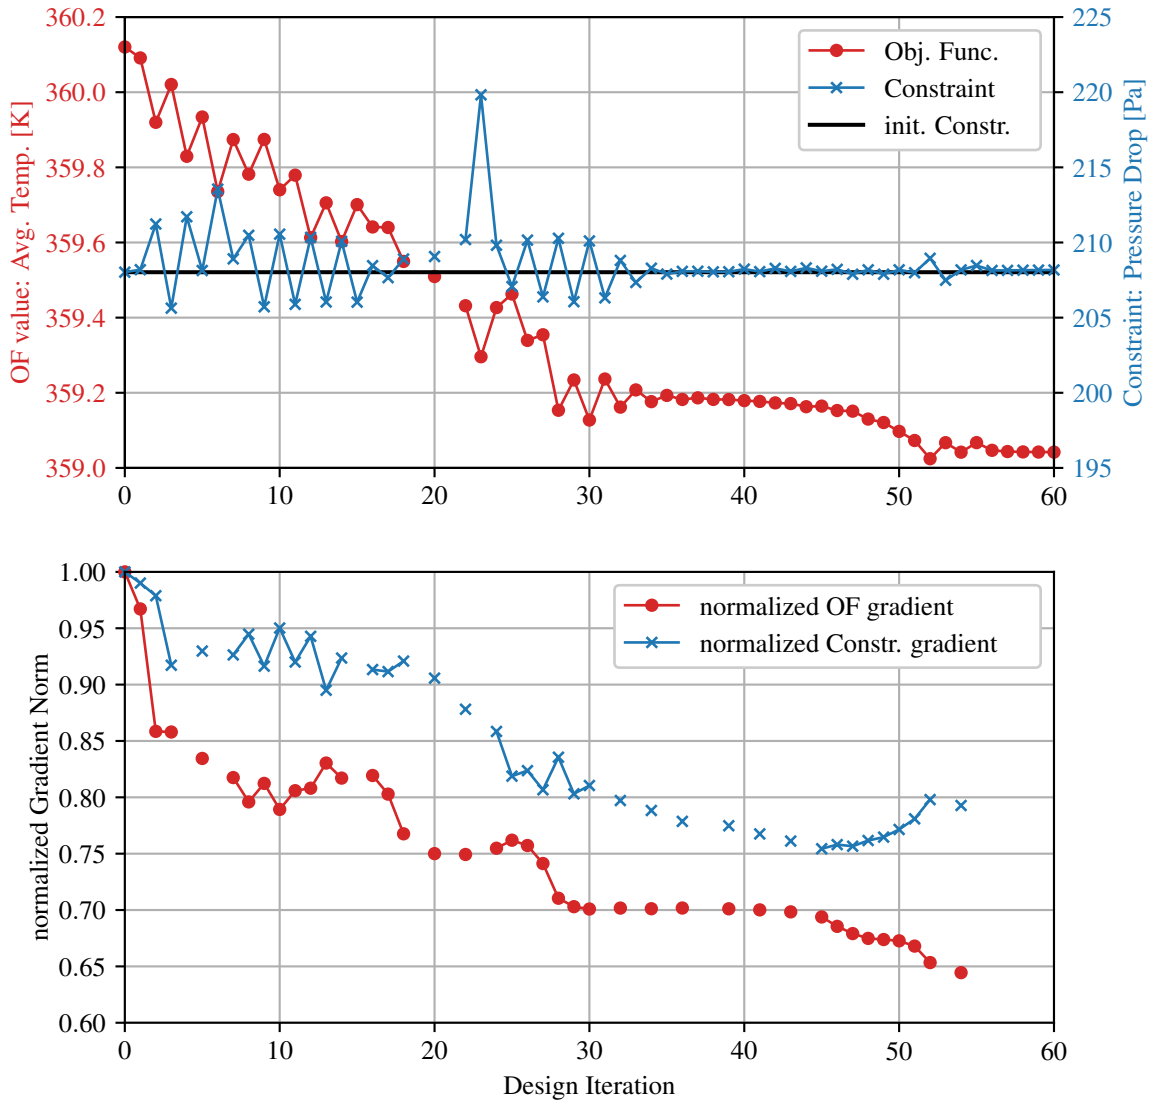


Figure 2: Optimization history: objective function (OF) value and Constraint in the upper image and their respective Gradients in the lower.

the mesh deformation algorithm that produces self-intersecting meshes beyond that shown state. A re-meshing and optimization restart would certainly result in a further flattened pin. The greatly thickened average temperature optimized pin is limited by the prescribed FFD-box bounds. A greater surface area between fluid and solid leads to a higher cooling, as the high heat conduction in the solid body spreads the temperature rather uniformly over the whole surface. As a secondary effect, the thickened pin creates a choke point that increases the flow velocity through the domain.

The defining feature of the constrained design is the downstream shifted hilltop and thus downstream shifted separation point. This more wedge- or v-shaped design has an increased surface area with attached flow, which has a greater contribution to the overall cooling, compared to the recirculation area

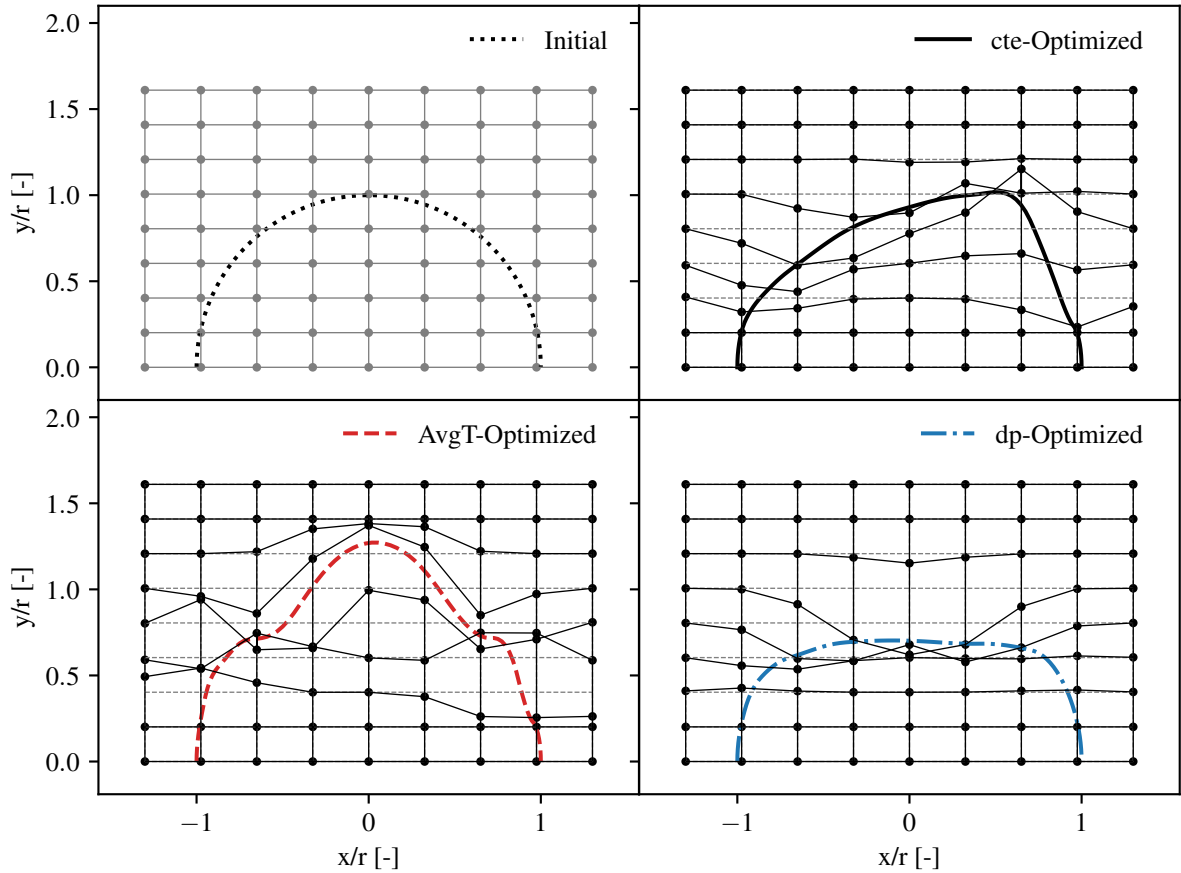


Figure 3: Optimized half-pin geometries and their respective FFD boxes for the initial, the constrained and unconstrained optimized setups.

behind the pin.

5 SUMMARY AND CONCLUSION

The derivation of the discrete adjoint equations for the considered CHT solver in order to efficiently and accurately compute derivatives with respect to objective functions depending on the PDE solution was outlined. The studied 2D pin-array testcase was presented in depth with an emphasis on the challenges in the mesh deformation routine using FFD-boxes, that is required for shape optimization. The gradients computed by the discrete adjoint solver are compared against finite differences and show relative differences of below 0.1% for both considered objective functions. In the subsequent gradient-based and constrained optimization, the average temperature on the inner pin arc is lowered by $\sim 1.1\text{K}$ while keeping the necessary pressure drop constant, compared to the initially round pin.

6 REPRODUCIBILITY

All utilized methods are implemented in the open Source multiphysics and design solver SU2 [7] and the results shown in this paper are created using the git tag [v7.3.1](#). The presented case is part of the SU2 testcase suite and is available with some deviations as a tutorial on the projects' website [8]. In order to recreate the results shown in this paper, please contact the first author for the exact configuration files.

7 ACKNOWLEDGMENTS

The authors would like to gratefully acknowledge the support of the Bayerische Forschungsstiftung (Grant No. BFS AZ-1232-16), as well as the computational resources provided by the RHRK high-performance computing center via the 'Elwetritsch' high-performance cluster at the TU Kaiserslautern.

References

- [1] Ole Burghardt, Nicolas R. Gauger, and Thomas D. Economon. "Coupled Adjoint for Conjugate Heat Transfer in Variable Density Incompressible Flows". In: *AIAA Aviation 2019 Forum*. 2019. DOI: [10.2514/6.2019-3668](#).
- [2] Ole Burghardt et al. "Discrete adjoint methodology for general multiphysics problems". In: *Structural and Multidisciplinary Optimization* 65.28 (2022). DOI: [10.1007/s00158-021-03117-5](#).
- [3] Thomas D. Economon. "Simulation and Adjoint-Based Design for Variable Density Incompressible Flows with Heat Transfer". In: *AIAA Journal* 58.2 (2020), pp. 757–769. DOI: [10.2514/1.J058222](#).
- [4] *Framework for Aerostructural Design Optimization (FADO)*. <https://github.com/pcarruscag/FADO>. Accessed: 2022-05-19.
- [5] S. V. Patankar, C. H. Liu, and E. M. Sparrow. "Fully developed flow and heat transfer in ducts having streamwise-periodic variations of cross-sectional area". In: *Journal of Heat Transfer* 99.2 (May 1977), pp. 180–186. ISSN: 0022-1481. DOI: <https://doi.org/10.1115/1.3450666>.
- [6] Max Sagebaum, Tim Albring, and Nicolas R. Gauger. "High-Performance Derivative Computations Using CoDiPack". In: *ACM Trans. Math. Softw.* 45.4 (Dec. 2019). ISSN: 0098-3500. DOI: [10.1145/3356900](#). URL: <https://doi.org/10.1145/3356900>.
- [7] *SU2 Codebase*. <https://github.com/su2code/SU2>. Accessed: 2022-05-19.
- [8] *SU2 Website*. <https://su2code.github.io>. Accessed: 2022-05-19.