

Evolving an Aircraft Using a Parametric Design System

Jonathan Byrne, Philip Cardiff, Anthony Brabazon and Michael O'Neill

Natural Computing Research & Applications Group
University College Dublin, Ireland
jonathanbyrne@gmail.com, philip.cardiff@ucd.ie,
anthony.brabazon@ucd.ie, m.oneill@ucd.ie

Abstract. Traditional CAD tools generate a static solution to a design problem. Parametric systems allow the user to explore many variations on that design theme. Such systems make the computer a generative design tool and are already used extensively as a rapid prototyping technique in architecture and aeronautics. Combining a design generation tool with an evolutionary algorithm provides a methodology for optimising designs. This work uses NASA's parametric aircraft design tool (OpenVSP) and an evolutionary algorithm to evolve a range of aircraft that maximise lift and reduce drag while remaining within the framework of the original design. Our approach allows the designer to automatically optimise their chosen design and to generate models with improved aerodynamic efficiency.

1 Introduction

Parametric systems are changing the conceptual design process in the same way spreadsheets changed finance. Both operate on the same principle. The user defines the relationships in a system and then changes variables in that system to rapidly explore alternative possibilities. Instead of manually creating a CAD model by dragging and dropping components, the parametric design is specified using variables and functions. Just as changing the value in a cell causes the spreadsheet to recalculate all related values, changing a variable that defines part of a model will adapt all the connected components so as to maintain a coherent design. Although there is a longer lead time to implement the initial model, once it is encoded the user can easily create endless variations on the original.

Evolutionary algorithms (EA) have shown their ability to optimise the shape and form of designs [11, 1]. One of the primary considerations when applying an evolutionary algorithm to a design problem is the representation used. The representation limits the search space by defining all the designs the algorithm could possibly generate. Poor representations generate designs that are invalid (internal faces, unconnected parts), infeasible (wrong scale) or missing the desired functionality. Creating a suitable representation is a difficult task that requires knowledge of both programming and of the specific domain.

Parametric systems provide a novel solution to the representation problem. A well implemented parametric system will only generate valid designs and incorporates domain knowledge. It also allows a designer with no formal programming experience to define the representation for the evolutionary algorithm. The designer provides the initial model and specifies the range limits so as to generate appropriate variations of their design. Parametric models make evolutionary optimisation directly accessible to the designer and allows them to use their domain knowledge to create a representation that generates feasible designs.

This work combines NASA’s parametric aircraft system (OpenVSP) and a computational fluid dynamics solver (OpenFOAM) with an evolutionary algorithm to generate a variety of optimised and novel designs. Sect. 2 gives an overview of parametric design systems and their application in industry. Sect. 3 describes the fluid dynamics solver used to generate the fitness values for the model. Sect. 4 discusses previous aircraft optimisation examples that used evolutionary approaches. Sect. 5 describes the parametric blended wing body model that was used and the two experiments that were carried out. Finally sections 6 and 7 examine the results of the experiments and the conclusions that can be drawn from them.

2 Parametric Design

Parametric design defines the relationships between components in a design. Generating a model consisting of hierarchical and geometric relations allows a exploration of possible variations on the initial design while still limiting the search space. Instead of manually placing and connecting components as is done in traditional CAD, component generating algorithms are linked with user definable variables. Defining the relationship between the components prevents invalid design generation. A change to one component will automatically effect a change on any connected component.

Parametric systems traditionally consist of basic components tailored for a particular design problem. An example of this would be the wing, fuselage and engine components in OpenVSP. Pre-defined components allow for domain knowledge to be embedded in the software and simplifies the design process. Although the user can explicitly define design components by programming them, normally model creation is done by combining existing components using a graphical interface. Many parametric design systems, such as grasshopper [5], are implemented using a drag and drop interface, shown in Fig. 1. The user can then manipulate the input and evaluate the benefit of the component to the overall design. An important aspect of parametric design is that the user observes the effects caused by manipulating a variable in real time, allowing the user to treat the underlying algorithm as a black box. Showing the effect of changing input to the system means that the user does not require an understanding of the underlying mechanics of the system, but instead gives them an intuitive understanding of how the components in a system are related to each other.

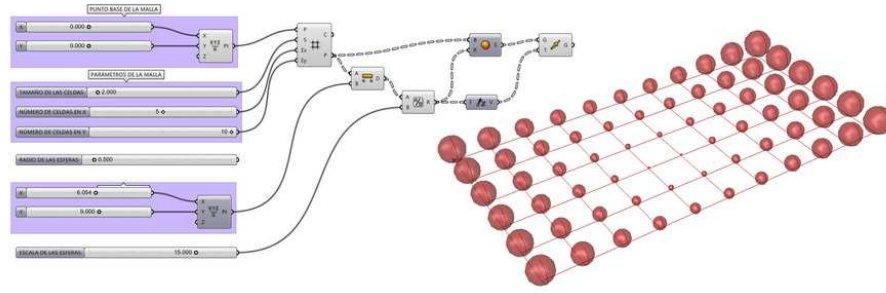


Fig. 1. The GUI for the Grasshopper parametric system. The variables are shown in the purple boxes on the left and are connected to the shape generating functions. The output design is on the right.

Parametric design tools have now been introduced into mainstream design software. There is the Grasshopper parametric design tool plug-in for the Rhino modelling system [5], Bentley Systems have implemented a program called Generative Components [23] based on the parametric design paradigm and Dassault Systems have developed CATIA, a CAD system combined with a parametric design tool. Parametric functionality was introduced to AutoCAD 2010 to allow for algorithmic manipulation of a design.

Combining parametric systems with structural analysis allows the user to make informed decisions about the geometric alterations during the conceptual design stage [9]. EIFForm is a parametric design system that optimises lattice structures by using structural analysis and a simulated annealing algorithm. The results have been used to design a structure in the inner courtyard of Schindler house [20]. Bollinger et al. [3] have developed parametric design systems that incorporate structural considerations and have used it to generate roofing structures for the BMW Welt Museum, Munich and the Rolex learning centre, EPFL, Lausanne. CATIA was combined with GSA structural analysis software [22] to evolve roofing structures for a football stadium [9].

The software used in this work is open vehicle sketch pad (OpenVSP). It was originally developed by NASA and Sterling Software as a rapid geometry modeler for conceptual aircraft [8] and has since developed into a stand-alone aircraft modelling tool. It was released as open-source software in 2012 under the NASA open source agreement. This work combines aerodynamic analysis with OpenVSP to analyse the lift and drag of the models. The next section discusses how the aerodynamic analysis was performed and the solver that was used.

3 Computational Fluid Dynamics

Computational Fluid Dynamics (CFD) uses numerical methods to solve how liquids and gases interact with surfaces. Although the calculations are computationally intensive, the dramatic increase in the power of standard hardware

enables basic CFD analysis to be carried out on standard desktop machines. OpenFOAM (open-source field operations and manipulation) [24] is used as the CFD solver in the experiments. Although primarily used for fluid dynamics simulations, it provides a toolbox of different solving techniques for applications such as combustion, electromagnetism, solid mechanics and heat transfer. It is designed for parallel execution due to the high processor demand of CFD modelling. It is highly extensible and has been adapted for calculating transonic aerodynamics [25], marine cavitation models [2] and orthotropic solid mechanics [4].

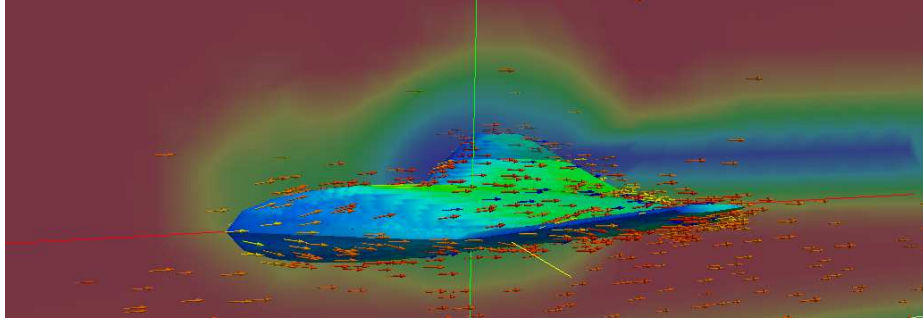


Fig. 2. The relative wind velocity and turbulence caused by the blended wing body model.

The solver used in the experiments is the semi-implicit method for pressure linked equations (SIMPLE) algorithm [17]. It is a steady state numerical solver for efficiently solving the Navier-Stokes equations that describe fluid motion. The algorithm forms the basis of CFD software and has been adopted to calculate the transfer of mass and momentum in a discretised three dimensional environment. The solver iteratively calculates the pressure and velocity within the system. Post processing then calculates the lift and drag forces generated by the model and these are used as the fitness value.

4 Evolutionary Aircraft Optimisation

“Since design problems defy comprehensive description and offer an inexhaustible number of solutions the design process cannot have a finite and identifiable end. The designer’s job is never really done and it is probably always possible to do better.” [13].

Design problems inevitably involve some trade off between desirable attributes [21]. In aircraft design there is a trade off between lift and drag which is known as aerodynamic efficiency. A design must have a minimal cross-sectional area to reduce drag but it must also have a large wing to maximise lift. Conflicting

objectives mean there is no one perfect solution, instead there is a pareto front of equally viable designs. Multi-objective problems are difficult to optimise but the population based approach of evolutionary algorithms has been shown to be a successful approach [26]. Multi-objective evolutionary algorithms (MOEA) have been shown to be a useful approach for finding the best compromise when tackling a multi-objective problem [6].

Accordingly there have been several MOEA approaches to evolving aerodynamically efficient aircraft. Due to the computational expense of CFD analysis most approaches focus on 2D optimisation of airfoils [18, 1, 14]. Different components have been optimised individually, such as the wing [15] or the turbine blade positions [19]. Although some large scale optimisation examples have been carried out [7, 16] the difficulty in defining such a complex representation has limited its application. The next section describes the aircraft model that is the basis for optimisation and the multi-objective algorithm used to optimise the aerodynamic efficiency.

5 Optimisation of Blended Wing Body Design

In traditional aircraft the fuselage provides little or no lift to the craft. Originally developed by NASA, the blended wing body (BWB) flattened the fuselage into the shape of an airfoil so that the entire craft generated lift. The BWB model has been used extensively as a test case for Multidisciplinary design optimisation (MDO) [12]. MDO uses optimisation techniques to solve design problems that span multiple disciplines. A parametric model of the BWB is provided with OpenVSP and was used as a test case. The model is shown in Fig. 3. One of the main advantages of parametric design optimisation is that it is easy to optimise specific features of a design. In order to highlight this two separate experiments were carried out. The first experiment solely optimised the airfoils while maintaining the predefined shape, so as to improve the design while remaining visually the same. The second experiment varied the sections and airfoils of wing structure, allowing the algorithm to explore different design possibilities.

The initial experiment only allows variation of the airfoil sections of the wing. The airfoil is defined by a National Advisory Committee for Aeronautics (NACA) profile system [10]. The NACA profile combines mean lines and thickness distribution to obtain the desired airfoil shapes. The NACA system allows the airfoil to be defined using only three parameters: thickness, camber and camber location. The wing on the BWB consists of 3 distinct wing sections. Only the camber and thickness were varied while the camber location remained fixed. Fixing the camber location of the airfoils means that the overall shape and configuration of the aircraft remain close to the original model.

The second experiment increases the number of variables in the representation to include the span, sweep, tip chord, root chord and dihedral angle of the wing. These features of the wing are illustrated in Fig. 6. Although changing this many features means that the model will vary greatly from the original design, it examines if the optimiser can be used as an explorative tool. Increas-

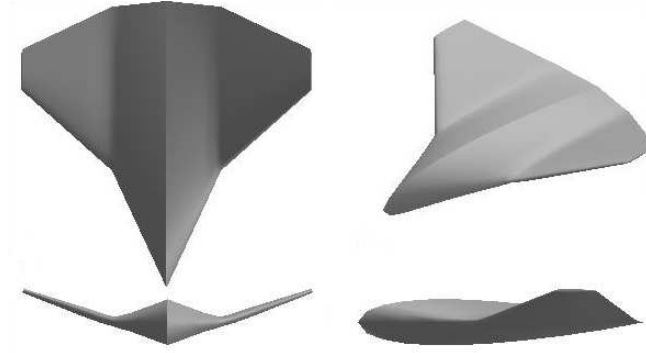


Fig. 3. The blended wing body model.

```

<aircraft> ::= <foil0><foil1><foil2>
<airfoil>  ::= {'Camber':<r>, 'Thickness':<r>}
<foil0>    ::= self.plane['foil0'] = <airfoil>
<foil1>    ::= self.plane['foil1'] = <airfoil>
<foil2>    ::= self.plane['foil2'] = <airfoil>
<r>        ::= 0.<digit><digit><digit><digit><digit>
<digit>    ::= 1|2|3|4|5|6|7|8|9|0

```

Fig. 4. The encoding used to describe the camber and thickness of each airfoil on the wing.

ing the amount of variability in the representation will generate more infeasible design but does open up the possibility of finding an improved yet unexpected configuration.

5.1 Experimental Settings

A standard genetic algorithm (GA) was used in the experiments. The settings used by the GA are shown in Table 1. Both lift and drag are being used as fitness values to evaluate the designs. In order to evolve designs that incorporated these features, the non sorting genetic algorithm II (NSGA2) multi-objective fitness function was used for selection and replacement [6]. Multi-objective search algorithms do not assume there is a globally optimal solution but that there are a set of non-dominated solutions. The non-dominated solutions are solutions that are better than the rest of the population for at least a single constraint and at least equivalent for all other constraints. The NSGA2 algorithm selects the least dominated solutions to create the child population.

In order to evaluate the performance of the evolutionary algorithm, the results were compared against randomly generated designs from the search space, essentially a brute force approach. This comparison examines if any useful genetic information is being transferred between individuals and whether the parametric representation is amenable to evolutionary search. Due to limited available com-

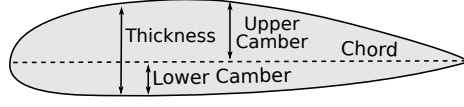


Fig. 5. NACA profile of an airfoil.

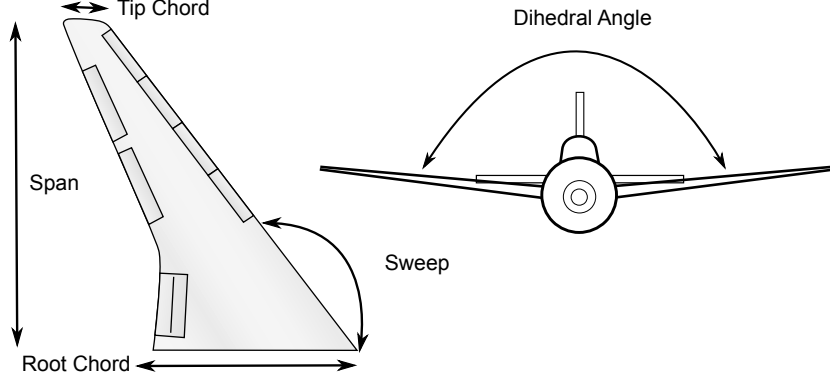


Fig. 6. The features of a wing section.

```

<aircraft> ::= <section0><section1><foil0><foil1><foil2>
<section> ::= {'Span':<r>, 'TC':<r>, 'RC':<r>, 'Sweep':<r>,'Dihedral':<r>}
<airfoil>  ::= {'Camber':<r>, 'Thickness':<r>}
<section0> ::= self.plane['section0'] = <section>
<section1> ::= self.plane['section1'] = <section>
<foil0>    ::= self.plane['foil0'] = <airfoil>
<foil1>    ::= self.plane['foil1'] = <airfoil>
<foil2>    ::= self.plane['foil2'] = <airfoil>
<r>        ::= 0.<digit><digit><digit><digit><digit>
<digit>    ::= 1|2|3|4|5|6|7|8|9|0

```

Fig. 7. The encoding used to vary each section and airfoil of the wing.

Table 1. Experimental Settings.

Property	Setting
Population Size	50
Generations	50
No. of Runs	2
Mutation Operator	Per Codon
Mutation Rate	1.5%
Crossover Operator	Single Point
Crossover Rate	70%
Selection & Replacement	NSGA2
Random Number Generator	Mersenne Twister

puting power only two runs were carried out for each experiment. Although this does not constitute a sufficient sample size to support the efficacy of stochastic methods such as an EA, the intention of these experiments is to examine if the aerodynamic efficiency of a pre-specified model could be optimised. As such the pareto-efficiency of the individuals in the final population will be used to judge the effectiveness of the algorithm as an active design tool.

6 Optimisation Results

A scatter plot of airfoil optimisation results is shown in Fig. 8(a). The graph shows how well the design maximised lift on the x-axis and how well it reduced drag on the y-axis. The original model is shown in black. The evolved solutions and brute force solutions are shown in red and green respectively with a line connecting individuals on the pareto front. Overall the pareto front of the evolved solutions is equivalent to the randomly generated solutions, indicating that no benefit was provided by the genetic information.

That an evolutionary approach did not outperform a brute force approach could be the result of the constrained nature of the representation. Each of the three airfoil sections had two variables. Although each individual was encoded by thirty integers, the range of each variable was limited to viable designs. Such a constrained representation could generate good solutions by random variation. This conclusion would be supported by the fact that both approaches generated pareto optimal designs that outperformed the original model. A sample of individuals from the pareto front are shown in Fig. 9. Limiting the evolvable representation to the airfoils produced optimised solutions that maintained the same overall design as the BWB aircraft.

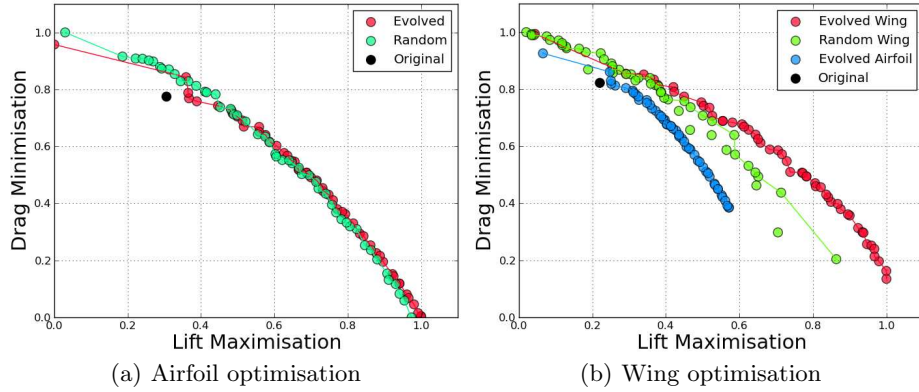


Fig. 8. The Pareto front for the final generation of aircraft. The results from the airfoil optimisation are shown in blue in the wing optimisation for comparison

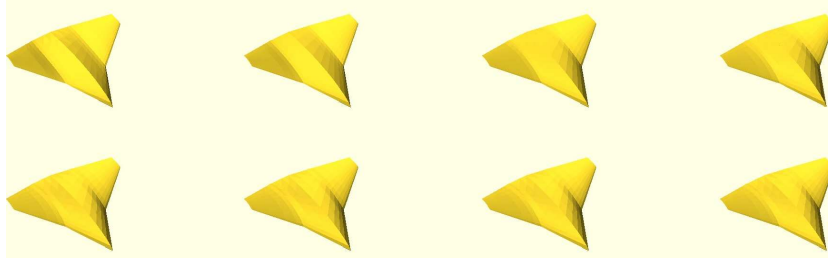


Fig. 9. Airfoil optimisation in order of increasing lift. The overall shape of the design remains the same

A scatter plot of wing and airfoil optimisation are shown in Fig. 8(b). Again the original model is shown in black and the evolved and brute force solutions are shown in red and green respectively. The graph shows how well the design maximised lift on the x-axis and how well it reduced drag on the y-axis. The increased variability of the representation greatly increased the range of the Pareto fronts when compared to the airfoil optimisation results, shown in blue.

The evolved Pareto front is distinct from the brute force approach. The randomly generated individuals tend to cluster around minimal drag designs as it is easy to find a design with a smaller wing, all the individual has to do is reduce the size of the aircraft. It is more difficult to find a design with an aerodynamically viable wing and this is where the evolutionary algorithm excels.

This result is highlighted by examining the average population fitness during the course of a run, as shown in Fig. 10. The NSGA2 selection operator compares child and adult populations so the graphs start at the second generation. The evolutionary algorithm is already populated with high fitness designs at this point while the selection pressure quickly improves average fitness of the brute force approach up to a point. In both drag and lift graphs the brute force approach plateaus after five generations. The evolutionary approach on the other hand continues to improve lift (while sacrificing drag efficiency) for the duration of the run.

A sample of the individuals on the pareto front are shown in Fig. 11. The relaxing of the evolvable representation resulted in many different wing configurations being generated. The amount of variation shows that such design problems are highly open-ended with no single optimal design configuration. It also suggests that allowing the algorithm to evolve more components of the representation could result in novel yet highly efficient designs.

7 Conclusions

A parametric system allows the designer, not the programmer, to specify the design to be evolved. The experiments showed that the level of design optimisation could be varied. Specific components of the model can be optimised or the model can be used as the basis for generating entirely different aircraft configurations.

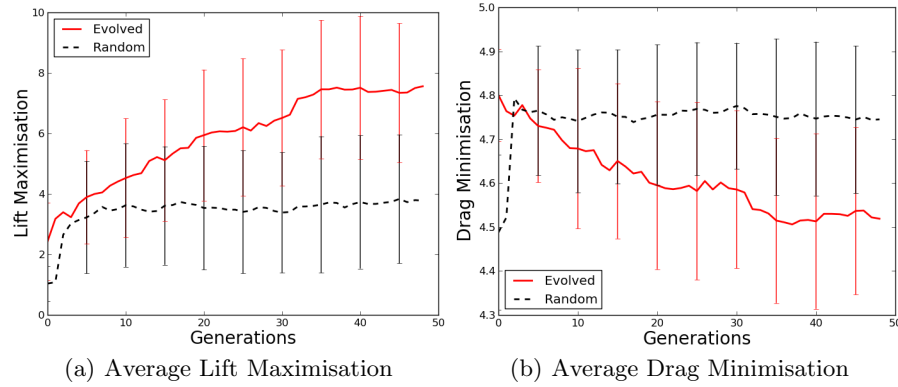


Fig. 10. The change in average lift/drag during the course of the run

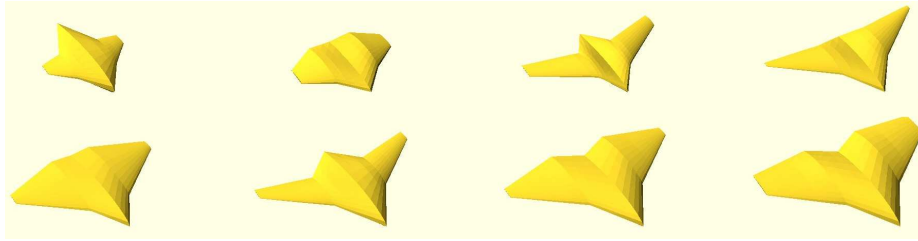


Fig. 11. Wing optimisation in order of increasing lift. The increased number of variables resulted in different wing configurations.

Although the sample size of the experiment is too small to draw any significant conclusions, initial results indicate that this representation is capable of being optimised. Even in experiment where brute force approaches performed comparably to evolutionary approaches, both generated designs that outperformed the original parametric model. This approach could be potentially applied to any existing parametric design to generate optimised solutions, turning the computer into an active design tool in the conceptual design process.

8 Acknowledgments

We would like to thank Science Foundation Ireland, the Financial Mathematics Computation Cluster and Andrea McMahon for her help during this project. We also wish to acknowledge the DJEI/DES/SFI/HEA Irish Centre for High-End Computing (ICHEC) for the provision of computational facilities and support. This work was funded by the SFI grants 08/RFP/CMS1115, 08/IN.1/I1868 and 08/SRC/FM1389.

References

- [1] Emre Alpman. Airfoil shape optimization using evolutionary algorithms. *Aerospace Engineering Department, Pennstate University*, 2004.
- [2] Rickard E Bensow and Goran Bark. Simulating cavitating flows with les in open-foam. In *V European Conference on Computational Fluid Dynamics*, pages 14–17, 2010.
- [3] K. Bollinger, M. Grohmann, and O. Tessman. Form, force, performance: Multi-parametric structural design. *Architectural Design*, 78(2):20–25, 2008.
- [4] P Cardiff, A Karač, and A Ivanković. A large strain finite volume method for orthotropic bodies with general material orientations. *Computer Methods in Applied Mechanics and Engineering*, 2013.
- [5] Martyn Day. Grasshopper, generative modelling. <http://www.grasshopper3d.com/>, 2010.
- [6] K. Deb, A. Pratap, S. Agarwal, and T. Meyarivan. A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE transactions on evolutionary computation*, 6(2):182–197, 2002. ISSN 1089-778X.
- [7] George S Dulikravich. Aerodynamic shape design and optimization-status and trends. *Journal of aircraft*, 29(6):1020–1026, 1992.
- [8] James R Gloudemans, Paul C Davis, and Paul A Gelhausen. A rapid geometry modeler for conceptual aircraft. In *34th Aerospace Sciences Meeting and Exhibit, Reno, NV Jan*, pages 15–18, 1996.
- [9] D. Holzer, R. Hough, and M. Burry. Parametric design and structural optimisation for early design exploration. *International Journal of Architectural Computing*, 5(4):625–643, 2007.
- [10] Eastman N Jacobs, Kenneth E Ward, and Robert M Pinkerton. The characteristics of 78 related airfoil sections from tests in the variable-density wind tunnel. Technical report, DTIC Document, 1933.
- [11] Rafal Kicinger, Tomasz Arciszewski, and Kenneth De Jong. Evolutionary computation and structural design: A survey of the state-of-the-art. *Computers and Structures*, 83(23-24):1943 – 1978, 2005. ISSN 0045-7949. doi: DOI: 10.1016/j.compstruc.2005.03.002.
- [12] Yan-Yee Andy Ko. *The multidisciplinary design optimization of a distributed propulsion blended-wing-body aircraft*. PhD thesis, Virginia Polytechnic Institute and State University, 2003.
- [13] B. Lawson. *How designers think: the design process demystified*. Elsevier/Architectural, 2006. ISBN 9780750660778. URL <http://books.google.ie/books?id=IPvqZJNAdG8C>.
- [14] Boris Naujoks, Lars Willmes, Werner Haase, Thomas Bäck, and Martin Schütz. Multi-point airfoil optimization using evolution strategies. In *Proc. European Congress on Computational Methods in Applied Sciences and Engineering (ECCOMAS00)(CD-Rom and Book of Abstracts)*, page 948, 2000.
- [15] Shigeru Obayashi. Multidisciplinary design optimization of aircraft wing planform based on evolutionary algorithms. In *Systems, Man, and Cybernetics, 1998. 1998 IEEE International Conference on*, volume 4, pages 3148–3153. IEEE, 1998.
- [16] Ian C Parmee and Andrew H Watson. Preliminary airframe design using co-evolutionary multiobjective genetic algorithms. In *Proceedings of the Genetic and Evolutionary Computation Conference*, volume 2, pages 1657–1665, 1999.
- [17] S.V Patankar and D.B Spalding. A calculation procedure for heat, mass and momentum transfer in three-dimensional parabolic flows. *International Journal of Heat and Mass Transfer*, 15(10):1787 – 1806, 1972.

ISSN 0017-9310. doi: [http://dx.doi.org/10.1016/0017-9310\(72\)90054-3](http://dx.doi.org/10.1016/0017-9310(72)90054-3). URL <http://www.sciencedirect.com/science/article/pii/0017931072900543>.

- [18] Domenico Quagliarella, Domenic D'ambrosio, and Angelo Iollo. Airfoil design using navier-stokes equations and hybrid evolutionary optimization techniques. Technical report, DTIC Document, 2003.
- [19] T Rogalsky, RW Derksen, and S Kocabiyik. An aerodynamic design technique for optimizing fan blade spacing. In *Proceedings of the 7th Annual Conference of the Computational Fluid Dynamics Society of Canada*, pages 2–29. Citeseer, 1999.
- [20] K. Shea, R. Aish, and M. Gourtovaia. Towards integrated performance-driven generative design tools. *Automation in Construction*, 14(2):253–264, 2005. ISSN 0926-5805.
- [21] H.A. Simon. *The sciences of the artificial*. The MIT Press, 1996.
- [22] Oasys Software. GSA, structural analysis version 8.5. <http://www.oasys-software.com/gsa-analysis.html>, 2011.
- [23] Bentley Sytems. Generative components, v8i. <http://www.bentley.com/getgc/>, 2011.
- [24] Henry G Weller, G Tabor, Hrvoje Jasak, and C Fureby. A tensorial approach to computational continuum mechanics using object-oriented techniques. *Computers in physics*, 12:620, 1998.
- [25] Benjamin Wüthrich and YJ Lee. *Simulation and validation of compressible flow in nozzle geometries and validation of OpenFOAM for this application*. PhD thesis, ETH, Swiss Federal Institute of Technology Zurich, Institute of Fluid Dynamics, 2007.
- [26] Eckart Zitzler and Lothar Thiele. Multiobjective evolutionary algorithms: a comparative case study and the strength Pareto approach. *IEEE Trans. Evolutionary Computation*, 3(4):257–271, 1999.