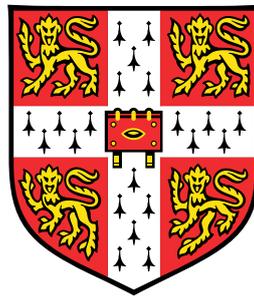


Facilitating the Use of Optimisation in the Aerodynamic Design of Axial Compressors



Samuel Tudy Phillips

Department of Engineering
University of Cambridge

This dissertation is submitted for the degree of
Doctor of Philosophy

Declaration

I hereby declare that except where specific reference is made to the work of others, the contents of this dissertation are original and have not been submitted in whole or in part for consideration for any other degree or qualification in this, or any other university. This dissertation is my own work and contains nothing which is the outcome of work done in collaboration with others, except as specified in the text and Acknowledgements. This dissertation contains approximately 59,000 words including footnotes, tables and equations and has 115 figures.

Samuel Tudy Phillips
December 2020

Facilitating the Use of Optimisation in the Aerodynamic Design of Axial Compressors

Samuel Tudy Phillips

There is commercial pressure to design axial compressors exhibiting high levels of performance more quickly. This is despite the performance of these machines approaching an asymptote in recent years, with further gains becoming increasingly difficult to achieve. One tool that can be used to help is optimisation, effectively harnessing the speed of computational analysis to accelerate the design process and unlock additional performance improvements. The greatest potential for optimisation exists at the preliminary design stage, however, current methodologies struggle when applied at this early point in the design process due to inadequate problem formulations, an inability to fulfil the role of enhancing designer understanding and a lack of high-fidelity analysis due to computational cost. The goal of this thesis is to facilitate the use of optimisation in the preliminary aerodynamic design of axial compressors by developing an improved methodology that overcomes these limitations.

The multiple dominance relations (MDR) formulation enables a larger number of performance parameters to be incorporated in a way that accurately reflects the desires of the designer. This is implemented within a Tabu Search (TS) that is capable of providing interpretable design development information to enhance designer understanding. The combined MDRTS algorithm, overcoming the limitations associated with formulation and understanding, outperforms existing methods when applied to analytic, aerofoil and six-stage axial compressor test cases, generating computational savings of up to 80%.

Multi-fidelity techniques are used to accelerate the search by conducting analysis on a “need-to-know” basis. Computational savings of over 70% are observed compared to the single-fidelity version of the algorithm across the analytic, aerofoil and six-stage axial compressor test cases, enabling high-fidelity analysis to be employed in a computationally efficient manner. The resultant methodology represents a novel and inherently flexible multi-level multi-fidelity optimisation technique.

Application to an N-stage axial compressor test case, in which the optimiser is given control over the number of stages in the machine, demonstrates the capabilities of the accelerated MDRTS approach. The complex design space is effectively navigated, generating computational savings of over 90% compared to existing methodologies and producing designs that are more likely to be of interest to the designer. Interpretable design development information is also provided for this problem to enhance designer understanding. These results show that the improved methodology successfully facilitates the use of optimisation in the preliminary aerodynamic design of axial compressors, overcoming the problems associated with formulation, understanding and speed that limit existing approaches.

“And whatever you do, whether in word or deed,
do it all in the name of the Lord Jesus,
giving thanks to God the Father through him.”

Colossians 3:17

Acknowledgements

Firstly, I would like to thank my supervisor, Dr. Jerome Jarrett, for his guidance, support and wise counsel over the past few years. Our discussions, both PhD related and otherwise, have been a highlight of my time studying in Cambridge and it goes without saying that this dissertation would not exist without him. I would also like to thank Dr. Laurence Cook for granting me early access to his manuscript introducing optimisation using multiple dominance relations (MDR) and for the many helpful discussions that enabled the successful application and development of MDR in this thesis. Finally, I would like to thank my family for their love and support, particularly my wife, Lizzie, whose company and remarkable ability to make me laugh have been invaluable whilst attempting to write-up during the ongoing COVID-19 pandemic.

This research was funded by the Engineering and Physical Sciences Research Council in the UK. Data generated from the research is not shared with this dissertation, but will be made available when published via the peer-review process. The results of Chapters 4 and 5 have already been published in conference proceedings with myself as lead author [191, 192], whilst the work from Chapters 7 and 8 has been prepared for publication in two articles yet to be submitted to a journal, for which I am also lead author.

Contents

List of Figures	xvii
List of Tables	xxv
Nomenclature	xxvii
1 Introduction	1
2 Axial Compressors and Design Optimisation	7
2.1 Axial Compressors	7
2.1.1 Basic Operation	8
2.1.2 Traditional Design Process	9
2.1.3 Computational Analysis	11
2.1.3.1 Meanline Methods	11
2.1.3.2 Throughflow Methods	12
2.1.3.3 Computational Fluid Dynamics	12
2.1.3.4 Accuracy and Academia	12
2.2 Design Optimisation	13
2.2.1 The Basics	13
2.2.2 Role in the Design Process	14
2.2.3 Single-Objective and Trade-Off Scenarios	14
2.2.4 Search Algorithms	15
2.2.5 Multiple Operating Points, Multiple Disciplines and Uncertainty	16
2.3 Review of Axial Compressor Design Optimisation	17
2.3.1 Preliminary Design Stage	18
2.3.2 Detailed Design Stage	20
2.3.3 Applications in the Wider Turbomachinery Field	22
2.3.4 Limitations of Existing Applications	23
2.4 Summary	23

3	Formulation, Understanding and Speed	25
3.1	Formulation	25
3.1.1	The Traditional Approach	25
3.1.2	Limitations of Objectives and Constraints	26
3.1.3	Techniques for Many-Objective Optimisation	28
3.1.4	Beyond Objectives and Constraints	30
3.1.5	Optimisation Using Multiple Dominance Relations	30
3.2	Understanding	32
3.2.1	The Need for Understanding	33
3.2.1.1	Justifying Optimisation Results	33
3.2.1.2	Promoting Creativity and Innovation	34
3.2.1.3	Parallels With Explainability and Interpretability	34
3.2.2	Search Algorithms and Their Ability to Enhance Understanding	35
3.2.2.1	Genetic Algorithms	35
3.2.2.2	Particle Swarm Optimisation	36
3.2.2.3	Bayesian Optimisation	36
3.2.2.4	Gradient Methods	36
3.2.2.5	Tabu Search	37
3.3	Speed	38
3.3.1	The Need for Speed	38
3.3.2	Methods for Accelerating Design Optimisation	39
3.3.2.1	Using Low-Fidelity Analysis Codes	39
3.3.2.2	Faster High-Fidelity Analysis Codes	40
3.3.2.3	Parallelisation	41
3.3.2.4	Reduced Order Modelling	41
3.3.2.5	Multi-Fidelity Methods	41
3.4	An Improved Optimisation Methodology	43
4	Tabu Search Using Multiple Dominance Relations	45
4.1	Tabu Search	45
4.1.1	Basic Search Mechanism	46
4.1.2	Heuristics to Enhance the Searching Capability	46
4.1.3	Previous Applications to Engineering Design Problems	48
4.1.4	Limitations of Tabu Search	48
4.2	Incorporating Multiple Dominance Relations	49
4.3	Additional Algorithmic Alterations	50
4.4	Implementation Validation	52
4.4.1	Analytic Test Case	52
4.4.2	Comparative Implementations	52

4.4.3	Results	54
4.5	Summary	56
5	A Simplifying Application Framework	57
5.1	An Additional Performance Parameter Classification	57
5.2	Arrangement Within a Nested Structure	58
5.2.1	Aerofoil Test Case	59
5.2.2	Experimental Set-Up	60
5.2.3	Single-Objective Results	62
5.2.4	Trade-Off Results	65
5.3	Similarities With Other Constraint Handling Approaches	69
5.4	Summary	72
6	Comparison to Existing Methods Using an Aerofoil Test Case	73
6.1	Experimental Set-Up	73
6.2	Single-Objective Scenario	75
6.2.1	Comparison to Traditional Formulations	75
6.2.2	Comparison to Alternative MDR Implementations	78
6.2.3	Emergent Constraints	81
6.3	Trade-Off Scenario	87
6.3.1	Comparison to Traditional Formulations	87
6.3.2	Comparison to Alternative MDR Implementations	90
6.3.3	Emergent Constraints	93
6.4	Information to Enhance Designer Understanding	96
6.5	Summary	99
7	Initial Application to a Six-Stage Axial Compressor	101
7.1	Six-Stage Axial Compressor Test Case	101
7.1.1	Performance Parameters	101
7.1.2	Parameterisation	102
7.1.3	Open Source Analysis System	102
7.1.4	Initial Geometry	103
7.2	Suitability of Different Algorithms	104
7.3	Single-Objective Scenario	106
7.3.1	Problem Formulations	107
7.3.2	Results	108
7.3.3	Emergent Constraints	110
7.4	Trade-Off Scenario	114
7.4.1	Problem Formulations	114

Contents

7.4.2	Results	115
7.4.3	Emergent Constraints	118
7.5	Information to Enhance Designer Understanding	119
7.6	Summary	123
8	Multi-Fidelity Acceleration	127
8.1	Incorporating Multiple Fidelities Within Optimisation Using MDR	127
8.1.1	Analysis on a “Need-to-Know” Basis	128
8.1.2	Avoiding Low-Fidelity Masking	129
8.2	Similarities With Existing Multi-Level Approaches	133
8.2.1	Sequential Approaches	133
8.2.2	Inexact Pre-Evaluation	133
8.2.3	Simultaneous Search Methods	134
8.2.4	Different Performance Parameters for Different Fidelity Levels	134
8.2.5	Other Approaches Conducting Analysis on a “Need-to-Know” Basis	134
8.2.6	A Distinct and Flexible Multi-Fidelity Methodology	135
8.3	Application to Test Cases	135
8.3.1	Analytic Test Case	135
8.3.1.1	Problem Formulations	135
8.3.1.2	Response Surface Construction Techniques	137
8.3.1.3	Experimental Details	137
8.3.1.4	Results	138
8.3.1.5	Comparing Response Surface Methodologies	139
8.3.1.6	Summary	139
8.3.2	Aerofoil Test Case	140
8.3.2.1	Multi-Fidelity Analysis	140
8.3.2.2	Problem Formulations	140
8.3.2.3	Experimental Details	142
8.3.2.4	Single-Objective Results	143
8.3.2.5	Trade-Off Results	145
8.3.2.6	Summary	148
8.3.3	Six-Stage Axial Compressor Test Case	148
8.3.3.1	Multi-Fidelity Analysis	148
8.3.3.2	Problem Formulations	149
8.3.3.3	Experimental Details	152
8.3.3.4	Single-Objective Results	152
8.3.3.5	Trade-Off Results	154
8.3.3.6	Summary	156
8.4	Potential Applications of the Flexible MFMDRTS Algorithm	157

8.5	Summary	158
9	Application to an N-Stage Axial Compressor	159
9.1	N-Stage Axial Compressor Test Case	159
9.1.1	Parameterisation	160
9.1.2	Analysis	160
9.1.3	Performance Parameters	161
9.1.4	Initial Geometry	161
9.1.5	Experimental Details	162
9.2	Single-Objective Scenario	163
9.2.1	Problem Formulations	163
9.2.2	Results	164
9.2.3	Emergent Constraints	168
9.2.4	Summary	170
9.3	Trade-Off Scenario	171
9.3.1	Problem Formulations	171
9.3.2	Results	172
9.3.3	Emergent Constraints	177
9.3.4	Summary	180
9.4	Handling a Variable Number of Stages	180
9.5	Information to Enhance Designer Understanding	181
9.6	Summary	185
10	Conclusion	189
10.1	Tabu Search Using Multiple Dominance Relations	190
10.2	A Simplifying Application Framework	191
10.3	Comparison to Existing Methods Using an Aerofoil Test Case	192
10.4	Initial Six-Stage Axial Compressor Application	193
10.5	Multi-Fidelity Acceleration	193
10.6	Application to an N-Stage Axial Compressor	194
10.7	Suggestions for Further Research	195
	References	197

List of Figures

1.1	Cross-section of a Rolls-Royce Trent engine in aero and gas turbine configurations	1
1.2	The axial compressor design process	2
2.1	Schematic of a four-stage axial compressor	8
2.2	A basic optimisation loop	14
3.1	Market share curves showing the commercial optimum and the effect of accelerating the design process	39
4.1	Flowchart for MOTS and MDRTS	47
4.2	Progression of the best overall performance measure found using the multi-objective and MDR formulations applied to the analytic test case	54
4.3	Progression of the best overall performance measure found using the different MDR implementations applied to the analytic test case	55
5.1	Baseline mesh used in SU ² RANS analysis	61
5.2	Mach number around the baseline RAE2822 aerofoil	61
5.3	Performance of the best designs found using Formulations D and F applied in the single-objective scenario with a minimum C_L requirement of 0.6	63
5.4	Progression of the best overall performance measure found using Formulations D and F applied in the single-objective scenario with a minimum C_L requirement of 0.6	64
5.5	Number of designs produced using Formulations D and F applied in the single-objective scenario with a minimum C_L requirement of 0.6 that exhibit performance within a given Euclidean distance of the idealised reference point	64
5.6	Progression of searches using Formulations D and F applied in the single-objective scenario with a minimum C_L requirement of 0.6	66
5.7	Progression of searches using Formulations D and F applied in the single-objective scenario with a minimum C_L requirement of 0.7	67

List of Figures

5.8	Performance of designs on the C_L - C_D Pareto fronts found using the six formulations applied in the trade-off scenario	68
5.9	Progression of searches using the six formulations applied in the C_L - C_D trade-off scenario	70
5.10	Progression of searches using the six formulations applied in the A_c - C_D trade-off scenario	71
6.1	Performance of the best designs found during 10 runs of the MDRTS algorithm and methods employing the multi-objective formulation applied to the aerofoil test case in the single-objective scenario	77
6.2	Progression of the best overall performance measure found by the MDRTS algorithm and methods employing the multi-objective formulation applied to the aerofoil test case in the single-objective scenario	77
6.3	Number of design produced by the MDRTS algorithm and methods employing the multi-objective formulation applied to the aerofoil test case in the single-objective scenario that exhibit performance within a given Euclidean distance of the reference point	78
6.4	Performance of the best designs found during 10 runs of the different MDR implementations applied to the aerofoil test case in the single-objective scenario	79
6.5	Progression of the best overall performance measure found by the different MDR implementations applied to the aerofoil test case in the single-objective scenario	80
6.6	Number of designs produced by the different MDR implementations applied to the aerofoil test case in the single-objective scenario that exhibit performance within a given Euclidean distance of the reference point	81
6.7	Progression of the desirable features during one run of the MDRTS algorithm applied to the aerofoil test case in the single-objective scenario	82
6.8	Progression of the desirable features used during 10 runs of the MDRTS algorithm applied to the aerofoil test case in the single-objective scenario	82
6.9	Performance of the best designs found during 10 runs of the MDRTS algorithm and methods employing the constrained single-objective formulation applied to the aerofoil test case in the single-objective scenario	84
6.10	Progression of the best overall performance measure found by the MDRTS algorithm and methods employing the constrained single-objective formulation applied to the aerofoil test case in the single-objective scenario	85
6.11	Number of designs produced by the MDRTS algorithm and methods employing the constrained single-objective formulation applied to the aerofoil test case in the single-objective scenario that exhibit performance within a given Euclidean distance of the reference point	85

6.12	Number of sufficient points found by the MDRTS algorithm and methods employing the constrained single-objective formulation applied to the aerofoil test case in the single-objective scenario	86
6.13	Performance of the best designs found during 10 runs of the MDRTS algorithm and methods employing the four-objective formulation applied to the aerofoil test case in the trade-off scenario	88
6.14	Progression of the best overall performance measure found by the MDRTS algorithm and methods employing the four-objective formulation applied to the aerofoil test case in the trade-off scenario	89
6.15	Number of designs produced by the MDRTS algorithm and methods employing the four-objective formulation applied to the aerofoil test case in the trade-off scenario that exhibit performance within a given Euclidean distance of the reference point	90
6.16	Performance of the best designs found during 10 runs of the different MDR implementations applied to the aerofoil test case in the trade-off scenario . . .	91
6.17	Progression of the best overall performance measure found by the different MDR implementations applied to the aerofoil test case in the trade-off scenario . . .	92
6.18	Number of designs produced by the different MDR implementations applied to the aerofoil test case in the trade-off scenario that exhibit performance within a given Euclidean distance of the reference point	92
6.19	Progression of the desirable features during 10 runs of the MDRTS algorithm applied to the aerofoil test case in the trade-off scenario	93
6.20	Performance of the best designs found during 10 runs of the MDRTS algorithm and methods employing the constrained two-objective formulation applied to the aerofoil test case in the trade-off scenario	94
6.21	Progression of the best overall performance measure found by the MDRTS algorithm and methods employing the constrained two-objective formulation applied to the aerofoil test case in the trade-off scenario	95
6.22	Number of designs produced by the MDRTS algorithm and methods employing the constrained two-objective formulation applied to the aerofoil test case in the trade-off scenario that exhibit performance within a given Euclidean distance of the reference point	95
6.23	Geometries and C_P distributions of the aerofoils exhibiting the lowest C_D values found during 10 runs of the different MDR implementations applied to the aerofoil test case in the single-objective scenario	96
6.24	PARSEC aerofoil parameterisation	97

List of Figures

6.25	Variation of the PARSEC parameters along the development path to the aerofoil exhibiting the lowest C_D value produced by the MDRTS algorithm in the single-objective scenario	98
7.1	Annulus of the initial six-stage axial compressor	104
7.2	Number of successful geometries generated by the TS, GA and PSO applied to a meanline single-objective axial compressor optimisation	106
7.3	Performance of the best designs found by the formulations applied to the six-stage axial compressor test case in the single-objective scenario	109
7.4	Progression of the best overall performance measure found by the formulations applied to the six-stage axial compressor test case in the single-objective scenario	110
7.5	Number of designs produced by the formulations applied to the six-stage axial compressor test case in the single-objective scenario that exhibit performance within a given Euclidean distance of the reference point	111
7.6	Progression of the objective and desirable features during the search using the MDR formulation applied to the six-stage axial compressor test case in the single-objective scenario	112
7.7	Number of sufficient points found by the formulations applied to the six-stage axial compressor test case in the single-objective scenario	113
7.8	Progression of M_{exit} and α_{exit} during searches using the MDR formulation and the emergent constraint limits applied to the six-stage axial compressor test case in the single-objective scenario	114
7.9	Performance of the best designs found by the formulations applied to the six-stage axial compressor test case in the trade-off scenario	116
7.10	Progression of the best overall performance measure found by the formulations applied to the six-stage axial compressor test case in the trade-off scenario.	117
7.11	Number of designs produced by the formulations applied to the six-stage axial compressor test case in the trade-off scenario that exhibit performance within a given Euclidean distance of the reference point	117
7.12	Progression of the objectives and desirable features during the search using the MDR formulation applied to the six-stage axial compressor test case trade-off scenario	118
7.13	Annuli of the initial design and that exhibiting the highest value of η_p found by the MDRTS algorithm applied to the six-stage axial compressor test case in the single-objective scenario	120
7.14	Changes to overall performance parameters during development of the design exhibiting the highest value of η_p found by the MDRTS algorithm applied to the six-stage axial compressor test case in the single-objective scenario	121

7.15	Changes to stage parameters during development of the design exhibiting the highest value of η_p found by the MDRTS algorithm applied to the six-stage axial compressor test case in the single-objective scenario	124
7.16	Changes to blade parameters during development of the design exhibiting the highest value of η_p found by the MDRTS algorithm applied to the six-stage axial compressor test case in the single-objective scenario	125
8.1	Flowchart for the MFMDRTS algorithm	130
8.2	The procedure used by the MFMDRTS algorithm to find non-dominated designs	131
8.3	Progression of the best overall performance measure found using the formulations applied to the multi-fidelity analytic test case	138
8.4	Time taken to train the RBF and GP response surfaces during the optimisations applied to the multi-fidelity analytic test case	139
8.5	Performance of the best designs found during 10 runs of the formulations applied to the multi-fidelity aerofoil test case in the single-objective scenario .	143
8.6	Progression of the best overall performance measure found by the formulations applied to the multi-fidelity aerofoil test case in the single-objective scenario .	144
8.7	Number of designs produced by the formulations applied to the multi-fidelity aerofoil test case in the single-objective scenario that exhibit performance within a given Euclidean distance of the reference point	145
8.8	Performance of the best designs found during 10 runs of the formulations applied to the multi-fidelity aerofoil test case in the trade-off scenario	146
8.9	Progression of the best overall performance measure found by the formulations applied to the multi-fidelity aerofoil test case in the trade-off scenario	147
8.10	Number of designs produced by the formulations applied to the multi-fidelity aerofoil test case in the trade-off scenario that exhibit performance within a given Euclidean distance of the reference point	148
8.11	Performance of the best designs found by the formulations applied to the multi-fidelity six-stage axial compressor test case in the single-objective scenario	153
8.12	Progression of the best overall performance measure found by the formulations applied to the multi-fidelity six-stage axial compressor test case in the single-objective scenario.	153
8.13	Number of designs produced by the formulations applied to the multi-fidelity six-stage axial compressor test case in the single-objective scenario that exhibit performance within a given Euclidean distance of the reference point	154
8.14	Performance of the best designs found by the formulations applied to the multi-fidelity six-stage axial compressor test case in the trade-off scenario . .	155

List of Figures

8.15	Progression of the best overall performance measure found by the formulations applied to the multi-fidelity six-stage axial compressor test case in the trade-off scenario.	156
8.16	Number of designs produced by the formulations applied to the multi-fidelity six-stage axial compressor test case in the trade-off scenario that exhibit performance within a given Euclidean distance of the reference point	157
9.1	Annulus of the initial eight-stage axial compressor used for the N-stage test case	162
9.2	Performance of the best designs found by the formulations applied to the N-stage axial compressor test case in the single-objective scenario	165
9.3	Performance of the best designs found by the formulations applied to the N-stage axial compressor test case in the single-objective scenario with the number of stages indicated	166
9.4	Progression of the best overall performance measure found by the formulations applied to the N-stage axial compressor test case in the single-objective scenario	167
9.5	Number of designs produced by the formulations applied to the N-stage axial compressor test case in the single-objective scenario that exhibit performance within a given Euclidean distance of the reference point	168
9.6	Progression of the desirable features during searches using the single- and multi-fidelity MDR formulations applied to the N-stage axial compressor test case in the single-objective scenario	169
9.7	Number of sufficient points found by the formulations applied to the N-stage axial compressor test case in the single-objective scenario	170
9.8	Progression of M_{exit} and α_{exit} during searches using the single- and multi-fidelity MDR formulations and the emergent constraints method applied to the N-stage axial compressor test case in the single-objective scenario	171
9.9	Performance of the best designs found by the formulations applied to the N-stage axial compressor test case in the trade-off scenario	173
9.10	Performance of the best designs found by the formulations applied to the N-stage axial compressor test case in the trade-off scenario with the number of stages indicated	175
9.11	Performance of designs on the η_p - SM Pareto front found by the MFMDRTS approach applied to the N-stage axial compressor test case in the trade-off scenario and the number of iterations taken for the CFD analysis to converge	176
9.12	Progression of the best overall performance measure found by the formulations applied to the N-stage axial compressor test case in the trade-off scenario. . .	176
9.13	Number of designs produced by the formulations applied to the N-stage axial compressor test case in the trade-off scenario that exhibit performance within a given Euclidean distance of the reference point	177

9.14 Progression of the desirable features during searches using the single- and multi-fidelity MDR formulations applied to the N-stage axial compressor test case in the trade-off scenario 178

9.15 Number of sufficient points found by the formulations applied to the N-stage axial compressor test case in the trade-off scenario 179

9.16 Progression of M_{exit} and α_{exit} during searches using the single- and multi-fidelity formulations and the emergent constraints method applied to the N-stage axial compressor test case in the trade-off scenario 179

9.17 Annuli of the initial design and that exhibiting the highest value of η_p found by the MFMDRTS algorithm applied to the N-stage axial compressor test case in the single-objective scenario 182

9.18 Changes to overall parameters during development of the design exhibiting the highest value of η_p found by the MFMDRTS algorithm applied to the N-stage axial compressor test case in the single-objective scenario 184

9.19 Changes to stage parameters during development of the design exhibiting the highest value of η_p found by the MFMDRTS algorithm applied to the N-stage axial compressor test case in the single-objective scenario 186

9.20 Changes to blade parameters during development of the design exhibiting the highest value of η_p found by the MFMDRTS algorithm applied to the N-stage axial compressor test case in the single-objective scenario 187

10.1 Potential improvement to the design process offered by the MFMDRTS algorithm 190

List of Tables

4.1	Parameter settings for the TS, GA and PSO	53
5.1	The six permutations of the performance parameter classifications within the nested hierarchy of dominance relations	59
5.2	Hicks-Henne bump locations and bounds	60
6.1	Formulations applied to the aerofoil test case in the single-objective scenario .	76
6.2	Formulations applied to the aerofoil test case in the trade-off scenario	87
7.1	Performance of the initial six-stage axial compressor	105
7.2	Formulations applied to the six-stage axial compressor test case in the single-objective scenario	108
7.3	Formulations applied to the six-stage axial compressor test case in the trade-off scenario	115
7.4	Variable changes and their impact on the key performance parameters during development of the design exhibiting the highest value of η_p found by the MDRTS algorithm applied to the six-stage axial compressor test case in the single-objective scenario	122
8.1	Formulations applied to the multi-fidelity analytic test case	136
8.2	Formulations applied to the multi-fidelity aerofoil test case in the single-objective scenario	141
8.3	Formulations applied to the multi-fidelity aerofoil test case in the trade-off scenario	142
8.4	Performance parameter classifications for the multi-fidelity six-stage axial compressor test case	150
8.5	Formulations applied to the multi-fidelity six-stage axial compressor test case in the single-objective scenario	150
8.6	Formulations applied to the multi-fidelity six-stage axial compressor test case in the trade-off scenario	151

List of Tables

9.1	Performance of the initial eight-stage axial compressor used for the N-stage test case	162
9.2	Formulations applied to the N-stage axial compressor test case in the single-objective scenario	163
9.3	Formulations applied to the N-stage axial compressor test case in the trade-off scenario	172
9.4	Variable changes and their impact on the key performance parameters during development of the design exhibiting the highest value of η_p found by the MFMDRTS algorithm applied to the N-stage axial compressor test case in the single-objective scenario	183

Nomenclature

Roman Symbols

A_c	cross-sectional area
c	constraint function
C_D	drag coefficient
C_L	lift coefficient
C_M	pitching moment coefficient
C_P	surface pressure coefficient
DF	diffusion factor
DH	de Haller number
f	objective function
L_{ax}	axial length
M_{exit}	exit Mach number
PR	overall pressure ratio
SM	surge margin
\mathbf{x}	vector of design variables
\mathbb{X}	design space - set of values that design variables can take
$\mathbb{X}_{optimal}$	optimal elements of the set of all designs

Greek Symbols

α_{exit}	exit flow angle
η_p	polytropic efficiency

Nomenclature

η_s	isentropic efficiency
Λ	degree of reaction
ϕ	flow coefficient
ψ	stage loading coefficient

Superscripts

$^\circ$	degrees
----------	---------

Other Symbols

\preceq_n	n th dominance relation
-------------	---------------------------

Acronyms / Abbreviations

AI	Artificial Intelligence
AR	aspect ratio
BO	Bayesian Optimisation
CFD	Computational Fluid Dynamics
CP	Bézier curve control point
GA	Genetic Algorithm
GP	Gaussian Process
H&J	Hooke and Jeeves
HPC	high pressure compressor
IPC	intermediate pressure compressor
IPE	Inexact Pre-Evaluation
LTM	long term memory
MDO	Multidisciplinary Design Optimisation
MDR	Multiple Dominance Relations
MDRTS	Multiple Dominance Relations Tabu Search
MFMDRTS	Multi-Fidelity Multiple Dominance Relations Tabu Search

min	minimum elements of a set
ML	Machine Learning
MOTS	Multi-Objective Tabu Search
MTM	medium term memory
PSO	Particle Swarm Optimisation
RANS	Reynolds-Averaged Navier-Stokes
RBF	Radial Basis Function
SBO	Surrogate Based Optimisation
sep.	trailing edge separation
STM	short term memory
TS	Tabu Search

Chapter 1

Introduction

Axial compressors are one of the key components in gas turbines employed across the world to generate electricity and power commercial aircraft. Situated towards the front of the engine, their task is to raise the pressure of the working fluid in preparation for combustion. According to UK government data [54], industrial gas turbines similar to that depicted in the lower half of Figure 1.1 produced over 40% of the electricity consumed across the country in 2019. Meanwhile, there are approximately 27,500 commercial aircraft operating worldwide each of which is powered by two or more turbofans similar to the upper half of Figure 1.1 [40].

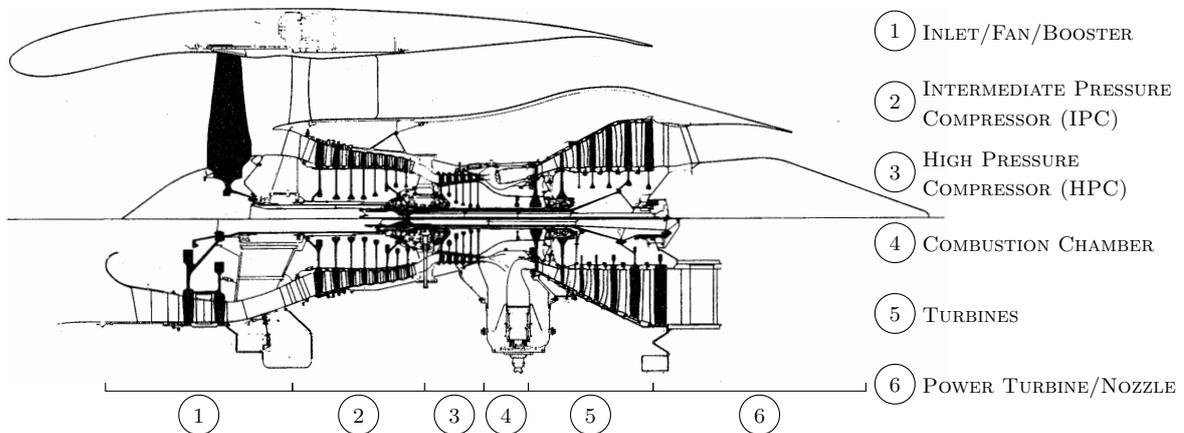


Fig. 1.1 Cross-section of a Rolls-Royce Trent engine in aero (upper) and gas turbine (lower) configurations [210].

In the absence of significant advances in renewable technology or investment in new nuclear power plants it is likely that gas turbines will continue to play an important role in electricity generation both in the UK and abroad for the foreseeable future. In aviation, the two major manufacturers, Airbus and Boeing, forecast a market for around 40,000 new aircraft over the next 20 years [2, 227]. The majority of these are likely to rely on gas turbines,

Introduction

and by extension axial compressors, for their propulsion, either directly or as part of hybrid configurations [3]. The design of good axial compressors therefore remains an important topic for engineering research, as it has been for over a century [106].

The performance of these turbomachinery components has approached an asymptote in recent years, with further improvements becoming increasingly difficult to achieve [237]. Despite this challenge there is continued commercial pressure for designs exhibiting high levels of performance to be delivered more quickly [53, 92]. According to Molinari and Dawes [176] the analysis techniques available to axial compressor designers seeking these improvements are sufficient; it is the design process itself that requires attention if the goal of increased performance in reduced design cycle time is to be achieved.

Gallimore [76] splits the axial compressor design process into four stages, shown in Figure 1.2. During preliminary design a wide-ranging search is conducted for machines that meet the design requirements, with key features of the compressor determined, such as mass flow rate, rotational speed and the number of stages. The most promising candidates advance to the throughflow stage where a more detailed view of the compressor is adopted, taking into account radial variations in the flow. With the configuration and annulus geometry set, the latter stages of the design process focus on the blading, with two-dimensional section design followed by further three-dimensional alterations.

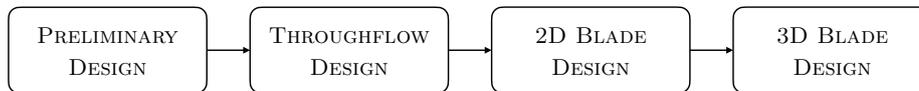


Fig. 1.2 The axial compressor design process [76].

One of the tools used in attempts to improve this design process over the past two decades has been optimisation. By removing designers from the tedious and time-consuming trial-and-error loop these automated routines enable many more designs to be considered than would be possible using traditional human-driven methods [186]. Efficient search algorithms are capable of unlocking previously untapped areas of the design space potentially leading to performance improvements [129].

Most applications of optimisation to the design of axial compressors have focussed on the detailed design stage [156], altering and improving specific aspects of a machine such as the shape of the blades (e.g. [118]) or endwalls (e.g. [134]). This is despite the fact that around 80% of the final performance of a machine is locked in during preliminary design [214]. In fact, according to Gallimore [76], no amount of optimisation later in the design process can overcome errors made at the preliminary stage. There is potential, therefore, for large performance improvements to be realised if the benefits of optimisation can be effectively harnessed at this early point in the design process.

Several authors have recognised this potential and applied optimisation to the preliminary design of axial compressors (e.g. [82, 115, 130, 193]). However, Li and Zheng [156], in their review of turbomachinery optimisation, suggest that there is room for improvement, outlining several factors limiting current methodologies when applied to this complex problem. These include:

- Inadequate problem formulations
- A lack of physical understanding
- Unacceptable computational expense

Firstly, Li and Zheng suggest that the problem formulations utilised by current methods are unable to adequately represent the true desires of the designer. When designing axial compressors numerous performance parameters exist that are relevant to the problem, particularly at the preliminary stage [238], and formulating these into a problem statement that accurately reflects what the designer actually wants from an optimisation can be challenging. Li and Zheng conclude that current approaches, relying on combinations of objectives and constraints, are unable to generate sufficient problem definitions, restricting the utility of the results they produce.

Secondly, there appear to be problems associated with a lack of physical understanding on the part of axial compressor designers. In mentioning this limitation Li and Zheng are primarily concerned with the impact this has on the effective choice of parameterisations and goals for optimisation schemes. However, Denton and Dawes [51, 53] note that physical understanding is also important when attempting to justify designs produced by an automated search. Without comprehensive explanations for the performance improvements achieved industrial experience suggests that engineers will be unlikely to trust designs generated by computerised algorithms [14, 214]. Optimisation is also used to explore the design space, helping the designer to gain physical insight into the problem under investigation [184, 240]. Discussing the history of turbomachinery aerodynamics, Cumpsty and Greitzer [43] state that “machines produce ideas just as surely as ideas produce machines”. This sentiment appears equally applicable to optimisation, with the automated routines also producing ideas just as surely as they produce machines. To facilitate the use of optimisation in the axial compressor design process algorithms need to be fully equipped to enhance designer understanding, assisting with the justification of the final designs produced and potentially generating ideas through improved physical insight.

Thirdly, there is the problem of computational cost. According to Shahpar [214] speed is often a deciding factor when choosing whether or not to employ optimisation in an industrial setting. The recursive nature of optimisation algorithms means that numerous calls are made to analysis codes, often leading to long runtimes that could potentially delay the design process rather than accelerate it [53]. One way around this problem might be to use

computationally cheap analysis techniques. However, this could result in fruitful areas of design space being missed due to a lack of modelling fidelity and may also promote designs with flaws that are only discovered when more accurate analysis is employed later in the design process [28]. By this stage remedial action is likely to be both expensive and time consuming, if it is possible at all [76]. Li and Zheng [156] suggest that the only way to produce meaningful optimisation results is through the use of high-fidelity analysis tools. The computational cost associated with employing these more accurate techniques is particularly high at the preliminary design stage where analysis of entire machines is required. Methods for accelerating high-fidelity optimisation, well established in the wider design optimisation literature, will therefore need to be applied to facilitate the use of any new axial compressor optimisation methodology in an industrial setting.

In addition to the three limitations discussed above, Li and Zheng [156] also highlight the need to consider manufacturing uncertainties in optimisation problems. These uncertainties can cause the actual performance of a final machine to be worse than that predicted by the optimisation routine used to develop it. Accounting for uncertainties can lead to robust designs that are more likely to satisfy the design requirements when constructed. However, incorporating stochastic elements adds to both the computational expense and complexity of an optimisation.

In contrast to the problems associated with formulation, understanding and speed discussed above, optimisation under uncertainty has received a lot of attention in the research community. Recent reviews by Chatterjee et al. [32] and Moustapha and Sudret [177] discuss several techniques that efficiently incorporate uncertainties to improve both the robustness and the reliability of designs produced by optimisation algorithms. This thesis therefore focuses on the alternative problem areas of formulation, understanding and speed, as it is believed that these offer the greatest potential for improvement compared to current methods.

In summary, to facilitate the use of optimisation in the preliminary design of axial compressors an improved methodology is required that addresses problems associated with formulation, understanding and speed. The primary goal of this thesis, therefore, is to develop and assess a new optimisation approach that fulfils the following criteria:

- Is capable of efficiently handling a large number of performance parameters, using a problem formulation that accurately reflects the desires of the designer.
- Is fully equipped for the role of enhancing designer understanding.
- Is computationally efficient, making minimal use of expensive high-fidelity analysis runs whilst retaining sufficient accuracy to ensure that the results produced are useful.

For the avoidance of doubt, the aim of this thesis is not to develop improved axial compressors or suggest alterations to the configuration or design of these complex machines.

Instead the goal is to improve the design process by introducing a new methodology that facilitates the use of optimisation, potentially enabling designers to generate performance improvements. The experiments undertaken therefore primarily assess the performance of the optimisation algorithms themselves rather than the machines produced.

To begin the development process it is first necessary to understand what makes axial compressor design such a challenging problem. Therefore, in Chapter 2 the traditional design process for these machines is introduced in more detail, along with the role played by optimisation routines. A review of previous applications of optimisation at the different stages of the design process is also conducted. Attention then turns, in Chapter 3, to the wider body of aerospace design optimisation literature in search of potential donor techniques that could be used to overcome the limitations of existing approaches associated with formulation, understanding and speed.

An initial methodology addressing the first two of these concerns is implemented and developed in Chapters 4 and 5, with performance compared to existing approaches in Chapter 6 using an aerofoil test case. The ability of the new technique to overcome the limitations associated with formulation and understanding when applied to the preliminary design optimisation of axial compressors is assessed using a six-stage test case in Chapter 7, again through comparisons to existing methodologies. Chapter 8 is dedicated to accelerating the resultant algorithm, addressing the final problem of speed. The new methodology facilitating the use of optimisation in the preliminary aerodynamic design of axial compressors is further assessed in Chapter 9 through application to an N-stage axial compressor problem in which the optimiser is given full authority over the compressor geometry and configuration, including the number of stages. Finally, Chapter 10 concludes the thesis by evaluating the approach against the criteria set out above and suggesting potential avenues for further research.

Chapter 2

Axial Compressors and Design Optimisation

To facilitate the use of optimisation in the aerodynamic design of axial compressors it is first necessary to understand what makes designing these machines so difficult and the assistive role that can be played by optimisation. This chapter begins by outlining how axial compressors work before describing the traditional design process and computational analysis tools used in their development. An introduction to design optimisation follows, covering the role it plays in the design process, some of the different search algorithms available and the types of problem they can be applied to. These strands are combined in the latter stages of the chapter with previous applications of optimisation to the design of axial compressors reviewed. This highlights the need for improvement in the areas of formulation, understanding and speed, with the main goal of the thesis being to develop a new methodology addressing these concerns.

2.1 Axial Compressors

As discussed in the previous chapter, axial compressors are one of the main components in gas turbines. Their role, as the name suggests, is to compress the working fluid, raising pressure and temperature in preparation for combustion and subsequent energy extraction. In axial compressors flow is parallel to the axis of rotation, in contrast to radial machines that exploit centrifugal effects to achieve larger pressure rises at the expense of efficiency. Whilst the first jet powered aircraft did utilise these centrifugal compressors, attention quickly turned to axial machines due to their potential for higher overall pressure ratios, higher efficiencies and a reduced cross-sectional area for a given mass flow rate [176, 210].

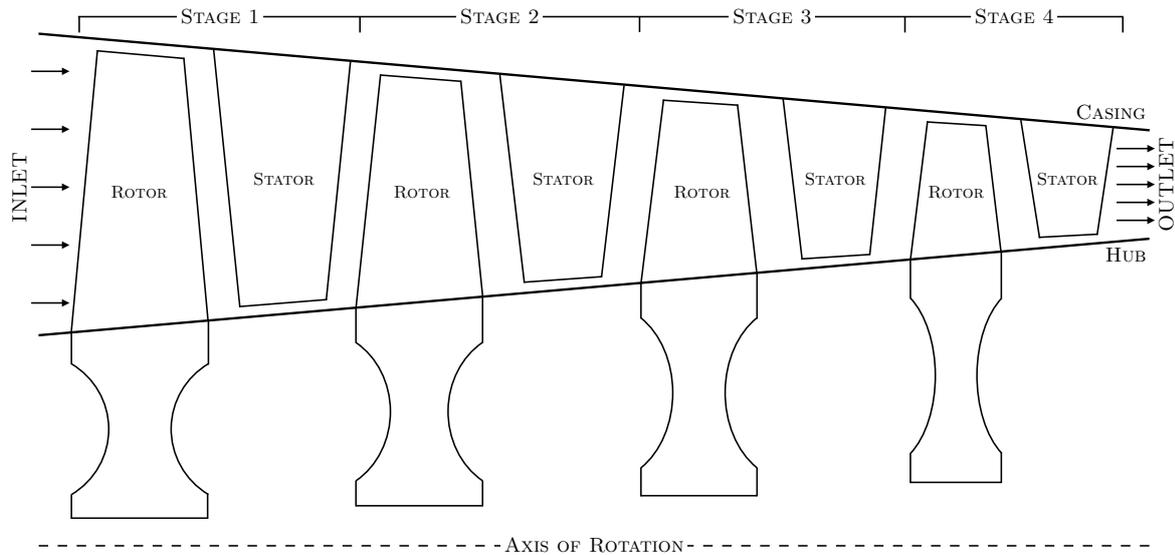


Fig. 2.1 Schematic of a four-stage axial compressor.

2.1.1 Basic Operation

Figure 2.1 shows a meridional view of a typical axial compressor. It is annular in shape, with the flow area decreasing towards the rear of the machine to ensure that the axial component of the fluid velocity remains roughly constant as the density increases.

Compression is achieved through interaction of the flow with a series of aerodynamically shaped blades arranged in rows. These rows work in pairs known as stages, with the first row in each stage rotating and the second stationary. The rotating row, referred to as the rotor, does work on the fluid, adding to the circumferential component of velocity. This additional swirl is removed by the stationary row, the stator, transferring the kinetic energy to static pressure. Both of these processes are diffusive relative to the blades, meaning the flow is at risk of separating if any single row is overworked. This limits the pressure ratio achievable using a single stage to being $O(1.2-1.5)$. If higher pressure ratios are desired several stages must be concatenated, with the flow leaving the stator of one stage becoming the input to the rotor of the next. The resultant multi-stage machines are complex with alterations made to one stage impacting the flow passing through the rest of the compressor.

One problem faced by multi-stage machines occurs when the rotational speed falls below the design value, as must happen during start-up when the compressor is accelerated from rest. The lower rotational speed reduces the amount of kinetic energy imparted to the flow by the rotors, lowering the amount of diffusion and leading to a reduction in the density of the flow through the machine. This results in higher velocities in the rear stages, with the flow approaching the blades in these rows at a lower incidence angle than they were designed for. These blade passages eventually become choked, restricting the mass flow rate through the front stages. The incidence angle for the front rows increases, ultimately leading to the

flow separating from the blades in a phenomenon known as stall. If the rotational speed of the compressor increases beyond the design value then the opposite occurs, with the front stages choking and the latter becoming more likely to experience flow instabilities.

There are three main approaches to ensure safe and reliable operation when the rotational speed departs from the design value. These may be used in isolation or, as is the case in most modern aircraft engines, in combination.

The first is to reduce the flow rate through the rear stages by bleeding off some of the working fluid part way along the compressor. The lowered flow rate delays the onset of choke in the rear stages and allows increased flow velocities in the front regions, reducing the likelihood of stall occurring when the rotational speed falls below the design value.

The second method addresses the incidence angles by introducing variable guide vanes at the compressor inlet or allowing the stators themselves to rotate about a radial axis. Angling these blades away from the axial direction decreases the flow rate for a given speed, again delaying stall in the front stages and choke in the latter. Walsh and Fletcher [238] suggest that one row of variable stator vanes is required for every additional stage beyond five to maintain satisfactory part-speed performance.

Another way to redress the incidence angles of the flow relative to the blades would be to increase the rotational speed of the last stage and decrease the rotational speed of the first. With a single compressor this is impossible, but when high pressure ratios are desired the compression can be split between multiple machines, each with a reduced number of stages, rotating at a speed that is more appropriate for the conditions being faced. This approach results in multi-spool architectures, with each compressor driven by a separate turbine. Multi-spool configurations are particularly important for the turbofan engines used in commercial aviation due to the low rotational speeds necessitated by the large diameter of the propulsive fan. An example of a three-spool layout can be seen in Figure 1.1, where compression is split between a fan, an intermediate pressure compressor (IPC) and a high pressure compressor (HPC).

Further detail of the physics that underpins axial compressor operation can be found in several textbooks [6, 42, 210, 238], whilst the historical development of these machines has been charted by Cumpsty and Greitzer [43] and Molinari and Dawes [176].

2.1.2 Traditional Design Process

The traditional design process for axial compressors, as outlined in the previous chapter, consists of four key stages: preliminary design, throughflow design, two-dimensional blade design and three-dimensional blade design [76, 176]. At the preliminary stage the basic outline of the compressor is determined. The number of stages and the number of blades in each row are selected, with the geometry of the hub and casing contours also defined along with the axial location of the different stages. Some of these attributes, in particular the hub

and casing contours, are refined during throughflow design, whilst the latter two phases are primarily concerned with the detailed design of the blade shapes.

Preliminary design has a significant impact on the final performance of a machine [57, 214], with several key attributes of the compressor being fixed at this early stage. Gallimore [76] emphasises the importance of getting these decisions right, warning that even large alterations to the blade shape in the latter stages of the design process cannot overcome fundamental problems associated with poor choices made at the preliminary stage. Unfortunately, making good choices when designing axial compressors is difficult, with Holt and Bassler [104] suggesting that they are more challenging to design than any other gas turbine component.

One reason for this is the large number of variables that affect the geometry. These range from machine-level parameters, such as the number of stages, to detailed aspects like the shape of individual blades. Complex interdependencies exist between different parameters and it can be difficult for human designers to explore the full extent of these within acceptable timeframes. As will be seen later in this chapter, automated optimisation routines can enable more thorough exploration of these relationships.

As well as the large number of variables, there are also several aspects of performance for a compressor designer to consider. The most obvious are efficiency, either isentropic, η_s , or polytropic, η_p , the latter of which is used throughout this thesis, and overall pressure ratio, PR , with one of the main goals being to produce a sufficient pressure rise as efficiently as possible.

However, additional performance attributes exist that need to be taken into account. As discussed in the previous section, the stability of axial compressors can be a significant problem. Flow through the machines is moving against an adverse pressure gradient, increasing the likelihood of separation and subsequent deterioration of the internal flow structures. The resultant instability primarily occurs in two forms, rotating stall and surge, both of which negatively impact compressor performance and potentially cause lasting damage to the machine¹. The designer therefore needs to ensure that the compressor is operating with a sufficient buffer before the onset of either rotating stall or surge. The most common measure of this buffer is the surge margin, SM , defined mathematically in Equation 2.1.

$$SM = \frac{PR_{stall} - PR_{design}}{PR_{design}} \times 100\% \quad (2.1)$$

As well as the SM of the machine as a whole, the designer also needs to consider the diffusion across individual blades, ensuring it is within acceptable limits for each row. Empirical measures are often used for this purpose, with two of the most common being the de Haller number [47], DH , and the diffusion factor [157], DF . These were developed in the early 1950s using experimental wind tunnel data. The former is calculated as the ratio of

¹For detailed descriptions of these phenomena see Cumpsty [42]

the flow velocities relative to the blade at the outlet and inlet of each row, with diffusion levels deemed acceptable when the value is greater than 0.72. The latter is a more complex relationship accounting for the effect of spacing between blades on the peak flow velocity at their surface, with values below 0.6 indicating suitable blade loading.

The designer must also acknowledge that the compressor forms only part of an overall engine. One area this consideration affects is the exit of the compressor, where the Mach number, M_{exit} , and flow angle, α_{exit} , need to be sufficiently small to avoid excessive losses being incurred in downstream components [238].

The numerous performance parameters discussed in the previous paragraphs arise from a purely aerodynamic consideration of the axial compressor design problem. There are also important mechanical and material factors that need to be taken into account if the final machine is to work successfully and efficiently. For the purposes of this thesis axial compressor design is considered from an aerodynamic standpoint alone in order to maintain tractability, with further disciplines taken into account through heuristic limits on certain design attributes.

2.1.3 Computational Analysis

The tools available to assess performance are important in any engineering design process. The analysis techniques most commonly employed by axial compressor designers fall into three categories varying in fidelity and computational cost. Fidelity in this sense refers to the accuracy of a model [69] or the degree to which it represents the real-world phenomena of interest [27].

2.1.3.1 Meanline Methods

The computationally cheapest tools are so-called meanline analysis codes. These employ analytic solution of the Euler work equation in combination with empirical loss models to generate performance predictions based on a limited amount of information about a design [6, 210]. Typical data requirements include a mean radius for each stage, the blade angles at this radius, and details relating to the inlet conditions and rotational speed.

Whilst seemingly lacking in physical realism, the use of empirical loss models allows these methods to implicitly incorporate a large amount of historical experience. Accurate predictions can be produced and, perhaps more importantly, the results are trusted by seasoned designers [29, 176]. Meanline approaches are also quick, with a single analysis taking less than a second to complete on a modern workstation. This speed enables a large number of potential designs to be assessed in a relatively short period of time, making meanline codes a popular choice at the preliminary design stage.

A limitation of meanline methods is their inability to produce reliable performance predictions when applied to novel designs. The use of empirical loss models means accuracy

can only be assured when the machines being analysed are similar to those used to generate the experimental data. In this thesis, where the focus is on facilitating the use of optimisation in the design process rather than improving axial compressors themselves, conventional machines are considered, reducing the impact of this limitation.

2.1.3.2 Throughflow Methods

At the next level up in terms of fidelity and computational cost are throughflow analysis codes. These extend meanline methods by introducing radial equilibrium, with the full extent of the compressor being considered from hub to casing. This allows important endwall effects to be taken into account, resulting in more accurate modelling of compressor performance compared to meanline methods. Throughflow approaches, like their meanline counterparts, often rely on empirical correlations for accurate performance predictions.

Due to their consideration of a greater extent of the compressor geometry, being two- rather than one-dimensional, throughflow methods are computationally more expensive than meanline approaches. They have therefore traditionally been employed at the second stage of the design process [76].

2.1.3.3 Computational Fluid Dynamics

The highest-fidelity analysis currently employed routinely in axial compressor design is solution of the Reynolds-Averaged Navier-Stokes (RANS) equations using purpose-built computational fluid dynamics (CFD) packages. These require a detailed geometric representation of the entire compressor to build a mesh on which to carry out the analysis. The accuracy and computational expense of CFD approaches depends on the quality and fineness of this mesh as well as the efficiency of the solver employed.

Computational cost initially limited these simulations to small regions of a machine, such as single blade passages [53], however advances in recent years have made RANS analysis of entire multi-stage compressors feasible [52]. Conducting three-dimensional CFD ensures that important secondary flow effects are accounted for and provides the most accurate and reliable measure of compressor performance before resorting to more complex and computationally expensive large-eddy or direct numerical simulations or physical experiments.

2.1.3.4 Accuracy and Academia

Whilst the analysis codes discussed in this section are established tools for use in the axial compressor design process, they are not infallible. Denton [51] warns that the exact numerical results of computational analysis should be treated with caution, as they rely heavily upon any loss or turbulence models employed within the solver. However, they can be used reliably on a comparative basis, assessing relative changes in performance between designs [51, 76]. It

is this comparative capability that is of most concern when seeking to utilise analysis codes within optimisation schemes, as is the case in this thesis.

Another limitation that is particularly applicable to compressor design is a lack of robustness. Baert et al. [10] highlight the regularity with which axial compressor analysis codes fail to converge, with a large number of design variable combinations also generating physically infeasible geometries. Any routine utilising computational analysis of axial compressors needs to account for this lack of robustness.

Finally, problems arise due to this research being conducted without an industrial partner. Day [46], discussing the prediction of surge behaviour, noted that the academic community lags several years behind industry when it comes to turbomachinery research due to the proprietary nature of a lot of information. This extends to analysis tools, with each of the major manufacturers using in-house solvers whose capabilities exceed those of the open source methods available to researchers working in academia. A move against this trend has recently been made by Denton [52] who issued the public release of a turbomachinery analysis system featuring meanline, throughflow and RANS analysis codes. These tools are employed throughout this thesis, representing the first use of exclusively open source analysis methods in the preliminary design optimisation of axial compressors known to the author. Whilst the exact numerical data may suffer due to the use of open source codes, the lack of industrial partnership does allow freedom in the presentation and discussion of results.

2.2 Design Optimisation

The complex nature of the axial compressor design process means there is potential for optimisation routines to have a beneficial impact. These methods are introduced in the following pages. Their role in the design process is discussed along with some of the different search algorithms that are available and the types of problem they can be applied to.

2.2.1 The Basics

Optimisation requires three components: a parameterisation, a method for predicting performance and a search algorithm. The parameterisation summarises the design in question, with alterations to variables within this parameterisation producing new candidate geometries. These are assessed using the method for performance prediction, with the results used by the search algorithm to inform further changes to the variables within the parameterisation. After several iterations this optimisation loop, shown schematically in Figure 2.2, should produce final designs that exhibit improved performance compared to the starting points.

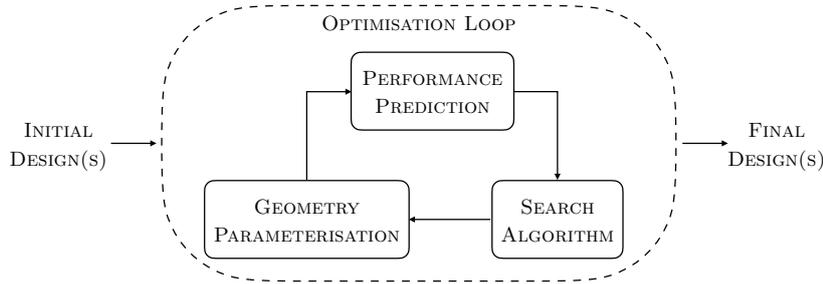


Fig. 2.2 A basic optimisation loop.

2.2.2 Role in the Design Process

Optimisation does not replace the human designer [112], and in practice is not used to seek designs that are optimal in any sense [38]. Instead it is a tool that enables systematic exploration of the design space, providing opportunities to learn about the interdependencies between different parameters [129, 171]. Whilst the designer is not replaced, they are removed from the trial-and-error loop, allowing them to focus on the more creative aspects of design such as specifying goals and interpreting results [112, 186, 233]. The speed of computational analysis allows many more designs to be considered than would be possible using manual procedures [186], resulting in regions of the design space being visited that may otherwise have been unreachable or overlooked [104]. As part of this exploratory role it is useful if optimisation algorithms can provide information about the design space to the designer, enhancing their understanding of the underlying problem. This is particularly important during the initial design stages [138], with improved knowledge enabling better decisions to be made [184]. The role of enhancing designer understanding becomes more important as optimisation is applied to increasingly complex problems [9].

With all of these potential benefits there is a risk of optimisation being viewed as a silver bullet. However, Denton and Dawes [53] advise caution when applying optimisation methodologies. The iterative nature of the loop shown in Figure 2.2 results in several analyses being required for progress to be made. Without a sufficient parameterisation or search algorithm many of these analyses may be unnecessary, ultimately resulting in a significant amount of time being wasted that could have been spent advancing the design process. Optimisation does have potential and has been applied effectively to a number of problems. However, effort is required to ensure that the most efficient use of computational resources is made at all times.

2.2.3 Single-Objective and Trade-Off Scenarios

Optimisation is commonly applied in two main scenarios. The first seeks to improve a single quantity of interest, resulting in a sole design exhibiting the best value of the selected

performance parameter. The second investigates a trade-off, with multiple quantities of interest considered and the optimiser tasked with finding a range of designs that perform well in terms of each of them. These trade-off studies are often employed during the early stages of the design process.

An important concept for the latter scenario is Pareto dominance. Under this criterion design A is only considered to be better than design B if it is at least as good in terms of all of the quantities of interest, and better in at least one. If this is the case then A is said to dominate B. If, however, some quantities are better in A but others are better in B then neither design dominates; they are said to be non-dominated. An optimiser conducting a trade-off study attempts to find the set of non-dominated designs using this rule of Pareto dominance. The non-dominated set is often referred to as a Pareto front or trade-off curve and these terms are used interchangeably throughout this thesis.

2.2.4 Search Algorithms

Several search algorithms have been developed to make progress on optimisation problems. These can be split into two families: gradient methods and global heuristics.

Gradient methods use sensitivity information to find the optimum that is closest to the starting design. They can do this efficiently if second order information is readily available, but otherwise must resort to potentially expensive finite differencing to calculate gradients. Popular algorithms in this category include SLSQP and SNOPT, both of which have seen renewed interest in recent years following the introduction of the adjoint approach that enables sensitivity information to be calculated at the cost of just one additional function evaluation irrespective of the number of design variables [111].

Despite their efficiency, gradient methods are limited to single-objective problems and must therefore employ weighted aggregate objective functions if the designer wishes to improve multiple quantities of interest in a single optimisation. There is also no guarantee that the design found will be the global optimum, as better designs might have been found by starting from a point elsewhere in the design space.

This is where global heuristics have a distinct advantage. These methods lack the rigorous convergence properties of gradient methods but have been shown to provide a wider exploration of the design space, increasing the likelihood of finding the global optimum rather than just the best design that is close to whatever happened to be the starting point.

Among the most popular global heuristics are Genetic Algorithms (GAs) and Particle Swarm Optimisations (PSOs). The former start with a population of designs and mimic evolutionary processes to generate new points. This involves breeding the best members of the population by combining their design attributes, as well as randomly mutating certain features in search of further performance enhancements. The GA that has been applied most widely is NSGA-II developed by Deb et al. [48].

PSO similarly starts with a population of designs, known as particles, with each simulating the behaviour of an insect in a swarm. Particles have a velocity and direction of travel that are updated based on information such as the best design found by that particle so far and the best design found overall. The velocity and direction are used to move to a new design which can be analysed. As with GAs, PSOs have no provable convergence properties but have been shown to be effective when applied to a number of problems [35].

The final heuristic mentioned here is Tabu Search (TS), originally developed by Glover and Laguna [88]. In contrast to GAs and PSO, TS is not a nature-inspired algorithm. Instead, new designs are generated using a pattern search, with one variable changed at a time by a small amount known as the step size. This allows the optimiser to move through the design space in an ordered manner, making small changes to the design vector between each analysed point. TS uses memory to enhance this basic searching technique. The short term memory (STM) prevents the optimiser from returning to previously visited points, rendering them “tabu”, enabling the search to climb out of local minima. A medium term memory (MTM) stores the best points visited during the search and a long term memory (LTM) contains all designs found by the optimiser. Both of these memories are used to reinvigorate the search after a set number of moves without improvement through procedures known as intensification and diversification. The optimisation can also be restarted with a reduced step size in order to refine good designs already found.

Due to handling just a single design, as opposed to a population, TS lacks some of the explorative capabilities of the other global heuristics. It is also devoid of the rigorous convergence guarantees that bolster gradient-based approaches. However, the ordered progression through design space enables this algorithm to navigate even the most complex design spaces [110], with particular effectiveness demonstrated when applied to the design of axial compressors [82, 115, 137].

Each of these algorithms have benefits and drawbacks, and the best method to employ is likely to be problem dependent [129]. The suitability of each algorithm for the preliminary design of axial compressors is discussed in the following chapter, paying particular attention to the ability of the optimiser to fulfil the role of enhancing designer understanding.

2.2.5 Multiple Operating Points, Multiple Disciplines and Uncertainty

In the simplest form of optimisation the performance of each design is considered at a single, deterministic operating point according to one discipline. However, it has been recognised that this approach fails to adequately represent the full range of factors that need to be taken into account during a real design process [129].

Firstly, any final design is likely to face many different conditions when operated in the real world. Optimising for a single operating point, even if it is the one that is expected to occur most commonly, may lead to a design that is over-specialised, resulting in poor performance

2.3 Review of Axial Compressor Design Optimisation

or even failure at other operating conditions [149, 228]. Secondly, the performance of a design is likely to be effected by more than one discipline, with interactions between different disciplines often being important [234]. Failing to account for these interactions during an optimisation may lead to improvement opportunities being missed. Finally, the real world is not deterministic, with any engineered product facing uncertainty both in how it is manufactured and the operating conditions it faces. Without taking these uncertain parameters into account the true performance of a final design may fall significantly below that predicted by the optimisation scheme used to produce it [129].

Optimisation methodologies have been developed for each of these scenarios. Multipoint methods consider the performance of designs at several operating points, leading to improved performance across the entire operational envelope [149, 228, 242]. Multidisciplinary Design Optimisation (MDO) incorporates analysis from several disciplines into the performance assessment of different designs, accounting for interactions between these disciplines and potentially unlocking synergistic performance improvements [168]. Optimisation under uncertainty, meanwhile, can lead to robust final designs that are more likely to produce the predicted levels of performance when assembled and operated in the non-deterministic real world [32, 177].

Whilst each of these approaches make the results of an optimisation more effective and useful, they also add to the complexity and computational cost. Later in this chapter, and more so in the following one, it is shown that current optimisation methodologies are limited when applied to the preliminary design of axial compressors due to problems associated with formulation, understanding and speed. It is only once these factors have been addressed that extensions accounting for multiple operating points, multiple disciplines and uncertainty can be considered. This thesis therefore focusses on overcoming the basic problems of formulation, understanding and speed using single-point, single-discipline, deterministic optimisation scenarios.

2.3 Review of Axial Compressor Design Optimisation

Having introduced the two topics separately, previous applications of optimisation to the design of axial compressors are reviewed in this section. The gaps in capability outlined in Chapter 1 are highlighted, with the discussion also providing context as the first steps are made towards developing an improved methodology to facilitate the use of optimisation in axial compressor design in later chapters. The review is organised by the design stage at which optimisation is applied, with a further section covering relevant work conducted in the wider turbomachinery field.

2.3.1 Preliminary Design Stage

Among the earliest applications of optimisation to the preliminary design of axial compressors appears in the work of Hearsey [100]. Using a throughflow analysis code, Hearsey improved the η_p of a 10-stage compressor by varying parameters including the stage pressure ratios, flow velocities, blade solidities and the radius of the hub. Similar early work was carried out by Holt and Bassler [104], optimising a smaller two-stage machine using a meanline code, and Tuccillo [231], who varied the hub and casing radii, blade inlet angles and deflections by hand to increase η_p .

Whilst most of these initial studies, along with work by Júnior [121], were limited to raising η_p and treating other parameters using constraints, Massardo and Satta [169] noted the importance of trade-off studies at the preliminary stage. Optimising a single-stage compressor using meanline analysis, these authors employed a weighted combination of η_p , SM and the mass of the machine as the objective function, taking more factors into account when comparing different designs. This work was continued at a later date by Maleki [164].

Jeshke et al. [117], meanwhile, developed a gas turbine design system known as MOPEDS (MODular Performance and Engine Design System) and used it to conduct meanline optimisation of a 10-stage HPC. The approach was limited to just three design variables (radius ratio at the compressor exit, rotational speed and number of stages) and used a weighted sum of the specific fuel consumption, the number of blades and the mass of the disks as the objective function. Moving away from weighted aggregate methods, Oyama and Liou [181] conducted a Pareto-based multi-objective optimisation of a four-stage axial compressor, maximising η_p and PR using a throughflow analysis code and an evolutionary search algorithm.

Ghisu et al. [81] compared a gradient method to a TS routine for the multi-objective optimisation of a seven-stage machine using a meanline code. With η_p and SM selected as objectives the TS was found to be superior due to the ability to extract a trade-off curve in a single run. The gradient method, in contrast, required several separate optimisations to be carried out to generate the Pareto front, resulting in greater computational expense overall. Taghavi-Zenouz and Afzali [225] also indicated that gradient-based approaches may not be the best choice for these preliminary studies due to the non-linearity inherent to the problem.

Further examples of optimisation being used at the preliminary design stage include work by Siller et al. [216], who altered the hub and casing contours and blade profiles of a single-stage compressor simultaneously, and Park et al. [185], who optimised a three-stage low pressure compressor using an enhanced meanline code. In more recent work, He et al. [99] conducted a meanline optimisation of a five-stage compressor using η_p and SM as objectives.

As well as realising the importance of taking multiple performance parameters into account, other authors, such as Panchenko et al. [183], Jarrett et al. [114] and Turner et al. [232], incorporated multiple disciplines into their optimisation studies, including stress testing and mechanical analysis alongside aerodynamic considerations. Dorca-Luque and Perrot [56]

2.3 Review of Axial Compressor Design Optimisation

coupled meanline aerodynamic analysis with structural and stability codes in the optimisation of a two-stage machine. In this work the authors were limited by an inability to handle all of the quantities of interest in a single problem formulation. A somewhat convoluted three-step procedure was undertaken consisting of a single-objective PR optimisation, followed by a two-objective SM and cost trade-off study, before a final three-objective optimisation maximising PR and η_p and minimising cost.

Similar problems were faced by Keskin and Bestle [130] when conducting preliminary design optimisation of a nine-stage compressor using a meanline analysis code. Having replaced human design procedures with automated methods the authors resorted to what was openly referred to as a “compromise method” due to the inefficiency of the optimisation approach when attempting to handle a large number of objectives. The compromise involved selecting one performance parameter as the objective and carrying out several optimisations with different limits for the other quantities of interest to observe the impact on the final designs produced.

This need for multi-step procedures and compromise methods demonstrates a limitation of the problem formulations being used. It is expected that improved final designs could be produced at a lower computational cost if all of the desires expressed in the different steps could be accommodated within a single optimisation run.

As well as being unable to handle the large number of potential objectives, other authors have found preliminary optimisations to be hampered by certain constraints imposed upon them. Bell et al. [15] noted the existence of so-called “process-intrinsic” constraints that are particularly limiting and exist purely as a result of the decomposition of the problem. This type of constraint often appears at the interface with other components, such as the the compressor exit where M_{exit} is limited to avoid excessive pressure losses downstream. Bell et al. demonstrated that performance improvements could be achieved by relaxing these process-intrinsic constraints.

Jarrett et al. [116] attempted an alternative approach, removing the interface constraints by incorporating larger swathes of the gas turbine into a single optimisation routine. The integrated method was extended in the work of Ghisu et al. [79, 82, 83] resulting in an optimisation system comprising the IPC and HPC from a three-spool gas turbine. Although limited to using meanline analysis codes due to computational expense, the authors demonstrated the benefits of this integrated approach compared to a more traditional segregated method. Non-deterministic elements were also incorporated into this framework [84], with robust optimisation conducted using polynomial chaos expansions to propagate the uncertainties.

Jarrett and Ghisu [113, 115] sought to increase the analysis fidelity of this integrated approach. High-fidelity RANS analysis was successfully applied through the use of a trust-region model management framework [37], with computationally cheaper performance predictions provided by a meanline code accelerating the optimisation. In the later work [115] the

authors discussed how the design process could be improved by exiting the refinement phase early, conducting further exploration of the design space to seek step-changes in performance rather than being satisfied with marginal gains. Eastwood et al. [60] employed the same integrated technique accelerated by meanline methods to investigate the importance of the starting geometry in preliminary optimisation studies. They found that the initial machine had a large impact on the final designs produced and that generating physically realisable starting points was extremely challenging for this complex problem. Hendler et al. [101] also considered integrated compressor optimisation, placing the machine in the context of whole-engine design by using interface parameters as design variables.

Whilst most other works have employed either meanline or throughflow analysis codes, Poehlmann and Bestle [193] attempted to utilise both in a sequential optimisation scheme. However, they struggled to get the results of the meanline optimisation of a seven-stage machine to converge in the throughflow code. The final technique required 20,000 design evaluations, too many for this method to be considered in conjunction with high-fidelity analysis techniques.

Overall, these applications of optimisation to the preliminary design of axial compressors have been characterised by a low level of analysis fidelity and problem formulations that struggle to effectively handle all of the relevant performance parameters. With the exception of the integrated approach developed by Jarrett and Ghisu [60, 113, 115], analysis has been limited to meanline and throughflow fidelities. Whilst these can provide useful initial performance approximations, without consulting higher-fidelity methods there is a risk that improvement opportunities could be missed [28]. Perhaps more importantly, defects that are only revealed by high-fidelity analysis may not be detected until late in the design process, at which point remedial action is likely to be expensive and time-consuming, if it is possible at all [76].

The persistent use of weighted aggregate objective functions [117, 164, 169], multi-step approaches [56] and methods openly referred to as a “compromise” [130] highlight that the formulations currently being employed struggle to effectively handle the large number of performance parameters that need to be considered during the early stages of the axial compressor design process. The work of Bell et al. [15], noting the process-intrinsic nature of many constraints, indicates that this is a wider problem and that a more capable formulation is required.

To facilitate the effective use of optimisation in the preliminary design of axial compressors these problems associated with fidelity and formulation need to be overcome.

2.3.2 Detailed Design Stage

Whilst there have been several applications of optimisation to the preliminary design of axial compressors, more attention has been paid to the detailed design phase [156]. The

2.3 Review of Axial Compressor Design Optimisation

majority of these works use high-fidelity RANS analysis to optimise the shape of the blades. The use of a computationally expensive high-fidelity code has limited most approaches to considering just single rows or stages. Benini [16] and Brooks et al. [25] optimised the blade shapes of an isolated compressor blade row, with the latter using coarsened grids to accelerate the procedure. Kampolis and Giannakoglou [123] also sought acceleration using semi-isolated searches conducted by viscous and inviscid methods. Li and Liu [154] included the effects of surface roughness in their single rotor optimisation, concluding that roughness had a negligible effect on the final results, whilst Li et al. [155] applied robust optimisation techniques to the shapes of rotor blades taking into account tip clearance uncertainties. In a break from the normal use of high-fidelity RANS analysis, the single-stage blade optimisation conducted by Pasquale et al. [187] employed a throughflow code

Noting that isolated compressor rotors rarely occur in practice, a number of authors sought to optimise the blade shapes of single rows or stages within the context of a larger compressor. Massardo et al. [170] was perhaps the first to do this, following up the preliminary optimisation mentioned in the previous section with a blade shape study using a throughflow code. Ellbrant et al. [64] carried out a two-step method to optimise the blade shapes for the first two rows of a three-stage compressor, using two-dimensional techniques followed by more intricate three-dimensional alterations. Goinis and Nicke [89], Kipouros et al. [137], Lejon et al. [144] and Ning et al. [179] all undertook similar studies, optimising the blade shapes of select rows within multi-stage machines. Whilst these methods analysed the entire compressor using the same code, Becker et al. [13] opted for a different approach, restricting high-fidelity RANS analysis to the row that was being optimised and using a computationally cheaper throughflow code to assess the effect of the changes made on the compressor as a whole.

As with work at the preliminary design stage, authors have also sought to extend optimisation methods to consider multiple disciplines [126, 161, 167] and multiple operating points [131, 132]. More recent research documents initial attempts to apply the adjoint method to compressor blade optimisation, progressing from two-dimensional studies [239], to isolated blades [162], to single stages [241] and rows within multi-stage machines [163, 240, 245]. Puente et al. [196] compared the results of an adjoint-based optimisation method to a design produced by a human. The final compressors exhibited similar performance, with the adjoint method requiring just two days to generate the machine compared to two weeks for the human designer.

Many of these detailed studies have noted the problem of non-converged solutions, with infeasible geometries also resulting in regular failure of the analysis codes. Baert et al. [10] and Joly et al. [118] developed methods to predict these failures using surrogate models, with the goal being to avoid wasting time on analysis that is unlikely to converge.

Whilst optimisation of the blade shapes has been most common at the detailed design stage, some authors have investigated alternative specific features of the compressor annulus. Dinh et al. [55] varied the shape and location of air injection holes in the stator shroud to improve performance, whilst Kim et al. [134] and Song et al. [222] included the shape optimisation of casing grooves in their formulations. Li et al. [146] also used optimisation to inform the best settings for variable stators within a five-stage machine, employing a meanline analysis code. Reitenbach et al. [201] conducted a similar study for a 10-stage compressor, using throughflow analysis coupled with a zero-dimensional full engine model.

These optimisation applications at the detailed design stage benefit from using higher-fidelity analysis than has been possible at the preliminary stage. However, they are limited to specific features of the compressor, in most cases altering the blades of a single row with little consideration given to the effect on the machine as a whole. Any improvements at the detailed stage are also likely to be marginal, as up to 80% of the performance of the machine will already have been determined during preliminary design [214].

2.3.3 Applications in the Wider Turbomachinery Field

In addition to the axial compressor applications discussed in the previous sections optimisation has been applied in the wider turbomachinery field. Dornberger et al. [57] discussed MDO in turbomachinery, highlighting the importance of trade-off studies at the preliminary stage and single-objective refinements during the detailed design phase. Gramatyka et al. [92] focussed on the role of speed in the design process, noting how even machines exhibiting the best technical performance are commercially worthless if they are not delivered in good time.

Several studies have also applied optimisation to the design of turbines, mainly focussing on the detailed design stage. Baert et al. [9] optimised the blades of a three-stage turbine, whilst authors including Kamenik et al. [122] and Keane [128] have conducted robust turbine blade optimisations. Li et al. [151] and Song et al. [221] applied MDO approaches to turbine blade design, and more recently Rodrigues and Marta [202] and Vitale et al. [237] have used the adjoint method for this problem. Turbine optimisation has not been limited to altering the shape of the blades, with Bergh et al. [17] varying the endwall geometry and Kim et al. [135] investigating the best placement of the air injection holes required for film cooling.

Optimisation has also been used to assess the potential of novel turbomachinery configurations, such as counterrotating compressors [96, 119, 190], as well as to repurpose existing machines, with Goryachkin et al. [91] and Popov et al. [194] used automatic blade alteration techniques to optimise aircraft engine components for new land-based operating conditions. Research has also extended beyond aerodynamics, with Toal et al. [229] conducting whole-engine transient thermomechanical design optimisation.

Whilst all of these applications are worthy of note, this thesis focuses solely on axial compressors. Given the similarities of different turbomachinery components it is expected that

the improved tool developed in later chapters to facilitate the use of optimisation in preliminary axial compressor design will also be applicable to and useful in the wider turbomachinery field, potentially enabling similar performance enhancements for these components.

2.3.4 Limitations of Existing Applications

This review of axial compressor design optimisation highlights a number of factors limiting current approaches. The majority of the optimisation applications have focussed on detailed design despite greater potential for improvement existing at the preliminary stage. Approaches that have been applied during preliminary design have often been restricted to using low-fidelity analysis, with the problem formulations employed also struggling to adequately handle the large number of relevant performance parameters. Moreover, none of the methodologies surveyed in this section appear to acknowledge their role as tools for enhancing designer understanding, with an optimisation algorithm more likely to be useful for a designer if it is equipped to fulfil this knowledge-based role.

2.4 Summary

In this chapter the basics of axial compressors and design optimisation have been introduced, demonstrating why the design of the former is such a difficult task and the role the latter can play in improving the design process. Reviewing previous work applying optimisation to the design of axial compressors demonstrates the limitations of existing approaches associated with formulation, understanding and speed outlined in Chapter 1. The next chapter discusses these shortcomings in more detail, emphasising their significance and surveying the wider literature in search of potential donor techniques that could be applied to overcome them. These solutions form the building blocks for an improved methodology facilitating the use of optimisation in the preliminary design of axial compressors.

Chapter 3

Formulation, Understanding and Speed

The review at the end of the previous chapter highlighted a number of factors limiting current optimisation methodologies when applied to the preliminary design of axial compressors. These can be summarised as inadequate formulations, an inability to fulfil the role of enhancing designer understanding, and use of low-fidelity analysis due to the computational cost associated with more accurate methods. In this chapter each of these problems is discussed in more detail, with the wider literature surveyed in search of potential donor techniques that could be used to overcome them. The chapter closes with a plan for an improved methodology to facilitate the use of optimisation in the preliminary design of axial compressors, making use of the donor solutions uncovered in the preceding discussion and review. This plan is then implemented in the remainder of the thesis.

3.1 Formulation

The first factor limiting current optimisation methodologies when applied to the preliminary design of axial compressors is their inability to effectively handle the large number of performance parameters that need to be considered early in the design process. In this section the limitations of the current state-of-the-art approach are outlined before alternative techniques that might offer improvement are discussed.

3.1.1 The Traditional Approach

The traditional formulation for optimisation problems is expressed in Equation 3.1. The optimiser is tasked with minimising¹ one or more objective functions, f , whilst satisfying

¹Maximisation can be accommodated by minimising the negative of a performance parameter.

a series of inequality constraints, c . To do this it searches the design space, \mathbb{X} , by altering elements within the vector of design variables, \mathbf{x} .

$$\begin{aligned} & \text{minimise} && f_1(\mathbf{x}), f_2(\mathbf{x}), \dots \\ & \text{subject to} && \mathbf{x} \in \mathbb{X} \\ & && c(\mathbf{x}) \leq 0 \end{aligned} \tag{3.1}$$

Formulating a problem using this objectives-and-constraints approach is relatively simple for the designer. All they are required to do is select the parameters they want to minimise as objectives and set upper or lower limits for any quantities they wish to constrain.

3.1.2 Limitations of Objectives and Constraints

The objectives-and-constraints formulation has been in use for several decades, enabling successful applications of optimisation to a plethora of problems. However, it does have limitations, particularly when dealing with a large number of performance parameters. This was evidenced in the work of Dorca-Luque and Perrot [56] and Keskin and Bestle [130] reviewed in the previous chapter in which the inadequacies of objectives and constraints led to the authors resorting to multi-step and compromise methods to incorporate all of the quantities of interest into their optimisation routines. Li and Zheng [156], reviewing turbomachinery optimisation, suggest that new methodologies are required that are capable of handling a larger number of performance parameters in an effective and efficient manner.

Including a large number of performance parameters in the traditional formulation shown in Equation 3.1 results in either a large number of constraints or a large number of objectives. Having lots of constraints can restrict the progress of the optimiser through the design space, potentially preventing it from generating significant performance improvements [8]. Applying constraints effectively is also challenging [34]. The most common approach is the penalty method, recommended by Coello [34], which involves adding values to the objective functions corresponding to the amount of constraint violation. This relies upon designer input in the form of factors for the penalties being applied, with an increasing number of constraints meaning more user-defined terms are needed.

The alternative to lots of constraints is lots of objectives. This too is problematic, with current optimisation algorithms rapidly losing search efficiency as the number of objectives increases. Schütze et al. [213] present a mathematical argument for why this is the case. If N_{par} is the number of solutions required to sufficiently span a two-dimensional Pareto front, then the number required to achieve the same level of convergence when there are k objectives would be N_{par}^{k-1} . As the number of objectives increases beyond three or four even requiring $N_{par} = 10$ results in a large number of solutions being needed to sufficiently span the trade-offs under investigation. Experience with current optimisation algorithms

confirms this theory, with most losing effectiveness when handling more than three or four objectives [30].

These problems associated with having too many constraints or objectives mean that the traditional objectives-and-constraints formulation is ill-equipped to handle the large number of performance parameters relevant to complex problems such as preliminary axial compressor design. A more fundamental concern is whether objectives and constraints are able to adequately represent the desires of the designer for some quantities of interest at all. Fonseca and Fleming [72] suggest that many constrained quantities are actually “soft objectives”, being quantities that the designer ideally wants to improve but that are not important enough to be considered as main objectives. A similar view was expressed in early axial compressor optimisation studies, with Hearsey [100] referring to all performance parameters other than η_p as “secondary objectives”.

Consider the Mach number and flow angle at the exit of an axial compressor (M_{exit} and α_{exit} respectively) as examples of the type of parameters being alluded to by Fonseca, Fleming and Hearsey that cannot be adequately handled using objectives or constraints. These two quantities are commonly treated using constraints in existing problem formulations for the preliminary design optimisation of axial compressors [99, 100, 130]. However, when discussing these parameters Walsh and Fletcher [238] state the following:

“These values must be minimised to prevent excessive downstream pressure loss.
 [...] Mach number should not be higher than 0.35 and ideally 0.25. Exit swirl
 should ideally be zero but certainly less than 10° .”

Applying constraints to these quantities focusses solely on the requirement for them to be less than the stated respective upper limits of 0.35 and 10° , failing to accommodate the need for M_{exit} and α_{exit} to be “minimised to prevent excessive downstream pressure loss”. Constraints therefore represent a poor approximation of what the designer actually wants for these parameters and as a result an optimisation formulation utilising constraints is unlikely to produce designs that are of real interest to the designer.

Bell et al. [15] noted that parameters such as M_{exit} and α_{exit} are primarily considered due to the modularised nature of the design process. While quantities like η_p and SM directly represent the performance of the compressor itself, M_{exit} and α_{exit} are instead mainly concerned with how the final design interacts with other components in the engine as a whole. Bell et al. investigated the η_p improvements that could be achieved by relaxing the resultant “process-intrinsic” constraints, allowing quantities such as M_{exit} to increase beyond the limits suggested by Walsh and Fletcher [238]. Despite the performance enhancements demonstrated, violation of the suggested limits is likely to prove problematic for engineers tasked with designing downstream components. One option might be to incorporate these components into the same optimisation procedure, as was done by Ghisu et al. [79]. However, the modularised commercial structure of modern engine manufacturers and the computational

cost associated with high-fidelity analysis of large regions of a complex machine represent significant obstacles to any integrated approach [14, 214].

An alternative would be to treat M_{exit} and α_{exit} as additional objectives. However, in addition to the problems associated with having lots of objectives, this too would probably be an inadequate representation of the desires of the designer. Holt and Bassler [104] stated the following:

“A lower compressor exit Mach number will reduce downstream losses, and is pursued as long as explicit optimisation goals are not affected”

The latter part of this quote suggests that these parameters should be minimised but with secondary importance compared to the main objectives. Treating M_{exit} and α_{exit} as additional objectives alongside a parameter such as η_p would therefore be a poor approximation of what the designer actually wants.

To facilitate the use of optimisation in the preliminary design of axial compressors a new methodology is required that incorporates problem formulations capable of adequately representing the desires of the designer for all of the relevant performance parameters. This need for good formulations only becomes more important as optimisation is applied to increasingly complex problems [9]. The limitations discussed above show that sufficient formulations cannot be achieved using the traditional objectives-and-constraints approach meaning an alternative, more sophisticated technique is required.

3.1.3 Techniques for Many-Objective Optimisation

A number of approaches have been developed to improve the efficiency of current search algorithms when handling a large number of objectives. Chand and Wagner [30], along with other authors [33, 147, 197, 235], review these techniques for so-called many-objective optimisation, mainly focussing on their application within evolutionary algorithms.

One of the main methods involves eliciting preference information from designers and incorporating this into an optimisation. Physical programming, developed by Messac [172, 173], was one of the first formal frameworks developed for this purpose, combining a series of functions to generate a single objective that more accurately reflects the desires of the designer for the relevant performance parameters. Friedrich et al. [74] and Nguyen et al. [178] used similar preference elicitation techniques to generate appropriate weightings for aggregate objective functions and the ideas of utility theory, applied by Du and Leifsson [58], are also related. An alternative method developed by Branke et al. [23] built on the notion that a designer would prefer a new design that was much better in terms of one objective even if it was slightly worse in one or more of the others. The proposed methodology required the user to specify maximum acceptable trade-offs between objectives, incorporating their preferences into the problem formulation. Schütze et al. [213] and Daskilewicz and German

[45] used preference information to explore Pareto fronts, first finding a single point on the trade-off curve before moving along the front using metrics specified by the designer. Further techniques for incorporating preferences into optimisation are found in works by Cvetkovic and Parmee [44], Fonseca and Fleming [72], He et al. [98] and Li et al. [150].

Instead of attempting to elicit preferences from designers, other authors sought to develop alternatives to the Pareto dominance criterion. A number of these alternatives were reviewed by Batista et al. [12], with one of the most promising being α -dominance developed by Ikeda et al. [109]. Under this criterion design A dominates design B if A exhibits slightly worse performance than B in one objective but much better in another, with the user-defined α parameter determining how much worsening is acceptable. The philosophy behind this dominance criterion is similar to that of Branke et al. [23] that led to the development of their preference-elicitation technique. Farina and Amato [67] used fuzzy logic to improve upon Pareto dominance in a similar way, allowing the optimiser to prioritise large improvements in objective values rather than penalise minuscule reductions and potentially miss out on overall gains. Corne and Knowles [41] found that there was no clear winner when comparing several of these Pareto alternatives, but that all were more effective than the basic approach when applied in a many-objective scenario.

As well as eliciting preferences and introducing alternative dominance criteria, user-defined reference points have been employed to improve the search efficiency of algorithms applied to many-objective problems. In these techniques the designer specifies the performance of an idealistic design and the optimiser is tasked with finding points whose performance is as close as possible to this reference design in a space containing the objectives. Deb et al. [49] developed an updated version of their popular NSGA-II algorithm that used a reference point approach to improve many-objective efficiency, with this extended further in the work of Mohammadi et al. [175]. Other notable reference point methods include that of Ruiz et al. [206] and Li et al. [153] in which the many-objective problem is converted into an optimisation with just two objectives: the proximity to the reference Pareto front and a measure of the crowding of solutions.

Whilst these different methods for many-objective optimisation have shown benefit in development studies, they have seen limited use within practical design optimisation, particularly in the aerospace field. The most likely reason for this is their complexity and reliance on designer input. With optimisation primarily being used in an exploratory role early in the design process it may be difficult for designers to express their preferences, particularly in quantitative terms, with little information about the design space available. It is more likely that these preferences will evolve as the designer learns about the problem and the potential performance of different design concepts emerge.

These methods also fail to address the problem discussed towards the end of the previous section that some parameters cannot be adequately treated using objectives or constraints.

Even when employing sophisticated many-objective optimisation techniques the user is limited to these two classifications of performance parameter and as a result may still be unable to adequately express their desires for some quantities of interest.

3.1.4 Beyond Objectives and Constraints

If one of the main limitations of current approaches is that certain parameters cannot be adequately represented using objectives or constraints then the improved methodology being developed in this thesis needs to employ an alternative problem formulation that goes beyond them. Fortunately, this type of parameter is not unique to axial compressor design, with the existence of quantities of interest that are similarly difficult to define recently highlighted by Cook et al. [39] in the design of two-dimensional aerofoils. Maximising lift and minimising drag are commonly used as objectives for this problem, with the designer also wanting the pitching moment and trailing edge separation to be as small as possible to alleviate trim and buffet concerns. One way to ensure the values of these latter quantities are acceptable would be to apply constraints, however it can be difficult for the designer to specify appropriate limit values before any optimisation has been carried out. Treating them as additional objectives alongside lift and drag would also be unsatisfactory as pitching moment and separation are unlikely to be considered to have the same level of importance as those main objectives. Pitching moment and separation are therefore difficult to adequately represent using either constraints or objectives, closely resembling the axial compressor performance parameters, M_{exit} and α_{exit} , discussed in Section 3.1.2.

To enable more accurate handling of this type of parameter Cook et al. developed a novel formulation for optimisation in which good designs according to pitching moment and separation could be sought within the set that perform well in terms of lift and drag. This efficiently guides the search towards designs that are of interest to the designer without requiring quantitative preference information that can be difficult to specify at the outset of an optimisation.

3.1.5 Optimisation Using Multiple Dominance Relations

Mathematically, the approach developed by Cook et al. [39] uses the set-based view of optimisation shown in Equation 3.2, where the optimiser seeks the minimal elements of the set of all designs, \mathbb{X} , subject to a dominance relation, \preceq , that is simply the criterion used to select between designs [18, 247].

$$\mathbb{X}_{optimal} := \min(\mathbb{X}, \preceq) \tag{3.2}$$

Cook et al. recognised that the expression in Equation 3.2 could be nested, as in Equation 3.3, to produce the more sophisticated problem of seeking designs that are optimal according

to a second dominance relation, \preceq_2 , within the set that is non-dominated according to the first, \preceq_1 .

$$\mathbb{X}_{optimal} := \min(\min(\mathbb{X}, \preceq_1), \preceq_2) \quad (3.3)$$

In theory there is no limit to the number of dominance relations that can be nested in this way, with the case for n dominance relations shown in Equation 3.4.

$$\mathbb{X}_{optimal} := \min(\dots(\min(\min(\mathbb{X}, \preceq_1), \preceq_2)\dots), \preceq_n) \quad (3.4)$$

This use of several criteria for selecting between designs leads to the formulation being referred to as optimisation using multiple dominance relations (MDR).

The MDR formulation allows designers to generate sophisticated problem definitions, more accurately handling parameters that were previously difficult to satisfactorily designate as either objectives or constraints. According to Cook et al. this approach ensures that the optimiser is searching for designs that the designer will actually be interested in.

Cook et al. developed a method to solve the problem in Equation 3.4 using ranking functions. Each design is given a separate ranking for each of the dominance relations present in the formulation, with this ranking usually just a count of the number of other designs dominating a given point according to a particular dominance relation [39]. When selecting between designs the optimiser compares these rankings using lexicographic ordering in place of directly comparing objectives. If two designs have the same ranking according to the first dominance relation, \preceq_1 , the optimiser compares the rankings of the designs according to the second dominance relation, \preceq_2 , and so on to the n th relation, \preceq_n . Two designs are only considered incomparable if they have identical rankings for all dominance relations.

It is in this hierarchical consultation of a series of dominance relations that the new formulation has key advantages over the objectives-and-constraints approach. The inclusion of several comparative criteria within a single optimisation run leads to fewer incomparable designs, increasing selection pressure towards points that are more likely to be of interest to the designer. The presence of several layers also allows a larger number of parameters to be taken into account by the optimiser in an effective and efficient manner.

Parallels may be drawn between this approach and the technique of bilevel optimisation, reviewed by Colson et al. [36] and Sinha et al. [219]. In bilevel methods two distinct optimisations are carried out, with each iteration of the higher level search requiring a sub-optimisation to be completed. This allows, for example, individual sections of a wing to be optimised whilst the planform is altered by an overarching upper level optimisation scheme [63]. In contrast to bilevel methods, the MDR formulation does not require separate optimisations. Instead, several dominance criteria are used to select between designs within a single search.

Another similar approach is the priorities method developed by Fonseca and Fleming [72]. This technique employs a hierarchical structure, with objectives assigned different priorities and designs compared using the objectives in priority order. If two designs exhibit exactly the same values of the highest priority objectives they are compared using the second most important objectives, and so on. This reliance on equality, rather than just non-dominance, to move to the next level of the hierarchy is a distinct limitation of the priorities method. In an engineering context it is rare for two designs to have exactly the same value of an objective, making it unlikely that performance parameters lower down the hierarchy will ever be considered during an optimisation. This could be avoided by assigning goals for higher-priority parameters, allowing less important quantities to be taken into account once the upper-level goals have been achieved. However, setting suitable values for these goals at the outset of an optimisation might be difficult, with the exact values chosen likely to have a large impact on the final designs produced.

One advantage of the priorities method, and indeed most other relations used to select between designs, is that they can be exactly evaluated given any two performance vectors. The MDR approach cannot as it relies upon previously analysed points to generate the rankings for each design. The restriction on being able to exactly evaluate a relation have been loosened to allow a hierarchical method to be employed that is expected to drive the search towards good designs more quickly when used in practice. In part this is due to the requirement to move to the next level in the hierarchy being just non-dominance rather than equality, increasing the likelihood of attributes lower down the hierarchy being taken into account during the decision-making process.

In their original work, Cook et al. demonstrated the effectiveness of the novel formulation using an optimisation under uncertainty problem, a car suspension test case and the aerofoil example discussed in the previous section. Given these promising results the MDR formulation appears to be a good choice for the new methodology being developed in this thesis to facilitate the use of optimisation in the preliminary design of axial compressors. The flexible and detailed formulation enabled by the use of MDR should overcome the problems faced in previous work where combinations of objectives and constraints were unable to efficiently handle the large number of relevant performance parameters. Moreover, application to a complex problem such as axial compressor design will enable a more rigorous assessment of this promising technique and allow for the maturing of the approach, ultimately facilitating future use on a wider range of problems.

3.2 Understanding

The second factor limiting current optimisation methodologies when applied to the preliminary design of axial compressors is that they are ill-equipped to fulfil the role of enhancing designer

understanding. According to empirical research by Bradner et al. [21] “professionals use design optimisation to gain understanding about the design space, not simply to generate the highest performing solution”. This view is also expressed by Gaier et al. [75] who referred to optimisation as “illumination” of the available design space. Kontagiannis et al. [138] suggested that “the most critical attribute of an optimisation tool [is] the ability to provide the maximum amount of information to the designer”, whilst Papageorgiou and Ölvander [184] argued that the results of an optimisation are useless if designers cannot understand them.

Given the importance of this knowledge-based role it would be useful if optimisation algorithms were equipped to provide relevant information to the designers that use them. Reasons for the importance of enhancing understanding are discussed in this section, along with an assessment of some of the most popular optimisation algorithms available in terms of their ability to fulfil this role of knowledge provision.

3.2.1 The Need for Understanding

In highlighting a lack of physical understanding as one of the key factors limiting current optimisation methodologies when applied to the design of axial compressors, Li and Zheng [156] were primarily concerned with the ability of designers to generate effective problem formulations and parameterisations. This is a relevant concern, particularly as the complexity of the problems being addressed increases [9]. However, designer understanding is also needed to justify the results of optimisation routines and to promote creativity and innovation.

3.2.1.1 Justifying Optimisation Results

According to Denton [51] any improvement predicted by computational analysis techniques is worthless if it cannot be justified. Determining this physical reasoning behind performance improvements requires understanding on the part of the designer.

While optimisers are good at producing designs exhibiting high levels of performance, many are ill-equipped to help designers answer this important question of why the final designs perform so well. According to Shahpar [214] the lack of justification is a significant barrier to the adoption of optimisation techniques within the turbomachinery industry. Being inquisitive people by nature, engineers often treat optimisation “with suspicion when a final good design is obtained going from A to B, without knowing all the steps in between that the optimiser has taken and why” [214]. Belie [14] suggests that many aerospace program managers view optimisation as a “mix of magic, hope, and hype all wrapped up in a tidy, yet unverifiable, black box”.

One way to facilitate the use of optimisation in an industrial setting and increase the likelihood of the designs produced being adopted would be for the algorithms themselves to provide useful information to assist designers in justifying the observed performance

improvements. This information will most likely be in the form of design rationale, detailing the decisions made by the optimiser during development of the final designs. The importance of capturing design rationale has been discussed by Bracewell et al. [20], Lee [142] and Szykman et al. [224], whilst Herschel et al. [102] review the related idea of provenance. This type of information can be particularly useful when remedial action is required at a late stage in the design process or if a fault is discovered when a product is in service. It can be difficult to determine the best course of action without information about why a particular design decision was made [20], again suggesting that it would be useful for optimisers to provide this data to users.

3.2.1.2 Promoting Creativity and Innovation

According to Cumpsty and Greitzer [43] and Jameson and Vassberg [112] it is the inventiveness and creativity of human designers, not automated optimisation routines, that lead to step-changes in performance during a design process. In order for the best possible final designs to be produced optimisation should therefore be a tool that facilitates this innovative activity.

One of the keys to encouraging innovation is information [203, 204], with deep insight required to guide the process of developing novel solutions [50]. Keane and Nair [129] and Rubbert [205] state that the best way to support innovation is to improve the processes by which this knowledge and understanding is acquired. Optimisation is one such process, and therefore has an important role to play in facilitating innovation through the provision of information. The need for understanding suggests that it is not good enough for an automated search to simply produce a series of designs exhibiting high levels of performance. To encourage and assist creative behaviour more information is required about the rationale used to generate those designs. If provided in a timely manner to the right people this data may promote innovation that leads to step-changes in product performance.

3.2.1.3 Parallels With Explainability and Interpretability

At this point it is informative to draw parallels between the present discussion and recent advances in the fields of Machine Learning (ML) and Artificial Intelligence (AI). Researchers in these areas noticed that despite the predictive capability of modern methods there was resistance to their adoption due to a lack of trust from potential users. In response, there has been a shift towards developing ML and AI methods that are inherently explainable and interpretable. The topic is reviewed by Biran and Cotton [19], with interpretable systems defined as being those whose operations “can be understood by a human, either through introspection or through a produced explanation”.

The reasoning behind the need for explainability and interpretability is similar to that suggested in the previous two sections discussing the need for optimisation algorithms to facilitate improved designer understanding. They include a requirement to justify decisions

made by algorithms to satisfy skeptical users [1, 87] and to promote learning and innovation [1, 94]. Vellido [236] also highlights how a lack of interpretability may prevent the uptake of computer-based systems in some scenarios, as has been the case with optimisation in the turbomachinery industry [214].

However, Lipton [159] sounds a note of caution, warning that to seek interpretability over predictive power may go against the wider goals of ML and AI to improve decision making. This must also be considered in optimisation; no matter how good a technique is at enhancing designer understanding it should not sacrifice the ability to produce designs exhibiting high levels of performance in a computationally efficient manner.

Developing an optimisation technique that fully aligns with the formal requirements of explainability and interpretability is beyond the scope of this thesis. However, there are similarities between the need for understanding in design optimisation and this related decision-making field. Investigating these links and determining whether any of the techniques for explainability and interpretability developed in the ML and AI literature are applicable to design optimisation could represent a fruitful avenue of further research.

3.2.2 Search Algorithms and Their Ability to Enhance Understanding

The discussion in the previous section highlighted the importance of understanding in the design process and that without it designs produced by optimisation are more difficult to justify and so are unlikely to be adopted in an industrial context. Optimisation algorithms therefore have an important role to play in enhancing designer understanding. In this section a number of the most popular optimisation methodologies are assessed in terms of their ability to provide useful information to designers, with the goal being to select the best for the new approach under development in this thesis to facilitate the use of optimisation in the preliminary design of axial compressors.

3.2.2.1 Genetic Algorithms

One of the most widely-used optimisation approaches is the Genetic Algorithm (GA). As discussed in the previous chapter, this mimics natural selection to generate new designs, breeding parents by combining aspects of the different design vectors. Whilst this method has been shown to be an effective and efficient way of finding good designs, following and interpreting the search process can be challenging for human users. When studying a design produced by a GA it may be difficult for the designer to determine how the optimiser has produced that design. There is no easily understandable or interpretable development information available from the optimiser due to the pseudo-random combinatorial techniques used to generate new geometries. Whilst the optimiser is adept at producing good final designs in a computationally efficient manner the search process itself is essentially a black-box, providing little useful information to the designer. This results in the user attempting

to determine the justification for a final design with minimal assistance available from the optimisation algorithm that generated it. GAs are therefore ill-equipped to fulfil the role of enhancing designer understanding.

3.2.2.2 Particle Swarm Optimisation

A similar problem is faced by another popular search algorithm mentioned in the previous chapter, particle swarm optimisation (PSO). In this case new designs are generated by making changes to the design vector using a velocity and a direction that are informed by the best points found so far. This approach does result in a certain amount of information about the development of a final design being available to the designer, as the positions taken by a particle can be traced through the design space. However, the path taken often involves large changes being made to the design vector with several variables altered simultaneously. This can make it difficult for users to determine the rationale used by the optimiser to produce the final designs. These PSO methods are therefore unable to provide the user with useful, interpretable information to enhance their understanding.

3.2.2.3 Bayesian Optimisation

A method that has seen increasing use in recent years is Bayesian Optimisation (BO). This approach, historically referred to as Efficient Global Optimisation (EGO) [120], uses an initial sample of designs to construct a response surface of the objective function. This response surface is used to suggest the best design, which is analysed and added to the set of training data. A new best design is then suggested, with the approach continuing until the accuracy of the response surface is considered to be sufficient.

As with GAs and PSO, this method has been shown to be effective at finding designs exhibiting high levels of performance in a computationally efficient manner. However, there is again little useful information provided to the user regarding the rationale behind the designs that have been developed. In selecting new designs the optimiser is essentially taking a series of educated guesses based on the performance of previously analysed designs. At the end of the optimisation the user is furnished with a final design and a response surface that models the design space, but is left unassisted when it comes to assessing the provenance of this final point.

3.2.2.4 Gradient Methods

Another important family of optimisation methods are those that make use of gradient information. Ranging from Quasi-Newton methods to more intricate approaches such as SLSQP and SNOPT, these utilise sensitivity information to inform the changes made to the design vector. In contrast to the optimisation algorithms discussed so far, the design vector

changes made by these methods are often small and can be tracked easily, making it possible to work back through the search and visualise the alterations made by the optimiser. The sensitivity information may also provide useful learning opportunities for designers.

However, one negative feature of these methods in terms of enhancing understanding is that they often change a large number of design variables simultaneously. As with PSO this may inhibit the interpretability of the design rationale information as it can be difficult for designers to determine the physical impacts of changing different variables independently. Gradient information is also not always available, and if finite differencing methods are required it can be computationally expensive to obtain. Moreover, these approaches have been demonstrated to be unsuitable for preliminary axial compressor design problems due to the highly non-linear design space and the existence of discrete variables such as the number of stages in the machine [225]. Even for problems where sensitivity information is available, gradient methods are still limited by the fact that they are inherently local, requiring multiple restarts from different initial designs to fully investigate the design space [81]. Eastwood et al. [60] demonstrated how challenging it can be to generate a single suitable starting geometry for axial compressor design problems, never mind the several required for multi-start gradient-based approaches.

Ultimately, despite the potential learning opportunities provided by sensitivity information, gradient methods do not represent a good choice of underlying search algorithm for the new technique under development in this thesis to facilitate the use of optimisation in the preliminary design of axial compressors. Altering several design variables simultaneously can make it difficult to determine the reasoning behind changes in performance, negatively impacting the interpretability of design development information. In addition, gradients may not be available for the axial compressor design problem that is of primary interest and generating starting geometries to ensure a more global search also represents a significant challenge.

3.2.2.5 Tabu Search

One method that has fallen out of favour in recent years is Tabu Search (TS). In this approach, originally developed by Glover and Laguna [88], new designs are generated using a pattern search. The technique of Hooke and Jeeves [105] is employed most commonly, where each design variable is changed in turn by a small amount. The optimiser moves to whichever new point exhibits the best performance, with that design used as the centre for the next iteration of the pattern search.

Generating new designs by making small changes of a known amount to one design variable at a time means that tracing the development of final designs is simple and results in rationale information that is interpretable for designers. The search process conducted by the optimiser to generate the final designs can be visualised and understood, making it

easier for designers to determine the physical reasoning behind the observed performance improvements. This rich development information, tracking improvements and the design changes that led to them throughout the search, is invaluable when considering the use of an optimiser in an industrial setting. It facilitates the enhancement of designer understanding, assisting them as they attempt to determine the justification for final designs and providing learning opportunities that may lead to creative innovation.

Despite lacking rigorous convergence guarantees, TS has been successfully applied to the optimisation of axial compressors in a number of works [82, 115, 137] and is therefore a suitable underlying search algorithm for the new approach under development in this thesis.

3.3 Speed

The final factor limiting current optimisation methodologies when applied to the preliminary design of axial compressors is associated with speed. The iterative nature of optimisation algorithms means that several hundred designs often need to be considered before significant performance improvements can be achieved. Analysing these designs is computationally expensive, particularly when high-fidelity analysis is desired. If care is not taken these analyses can delay the design process [53], going against the original goal of accelerating development of the final product. According to Shahpar [214] a reduction of a factor of 10 in the computational cost of an approach can be significant when deciding whether to employ optimisation in an industrial setting. In this section the need for speed is discussed in more detail alongside various techniques that could be used to achieve the required acceleration.

3.3.1 The Need for Speed

In engineering design time is money, with economic pressure for designs exhibiting high levels of performance to be produced in shorter design cycle times [53]. According to Keane and Nair [129] “[t]he simple pursuit of improved nominal performance at any cost is no longer commercially viable, even in military [and] space applications”.

Rubbert [205] discussed the idea of a commercial optimum, the point at which a product will achieve the greatest market share, and therefore revenue, for a company. Figure 3.1a shows how this commercial optimum relates to the development process for engineered products, demonstrating that even the most efficient machine is commercially worthless if it takes too long to design [92, 115]. However, if the design process can be accelerated, as seen in Figure 3.1b, then the commercial optimum not only achieves a greater market share for the company, but produces better technical performance as well [115]. The optimum point is also reached more quickly, potentially leaving time for further research and development work that could result in step-changes in technological capabilities.

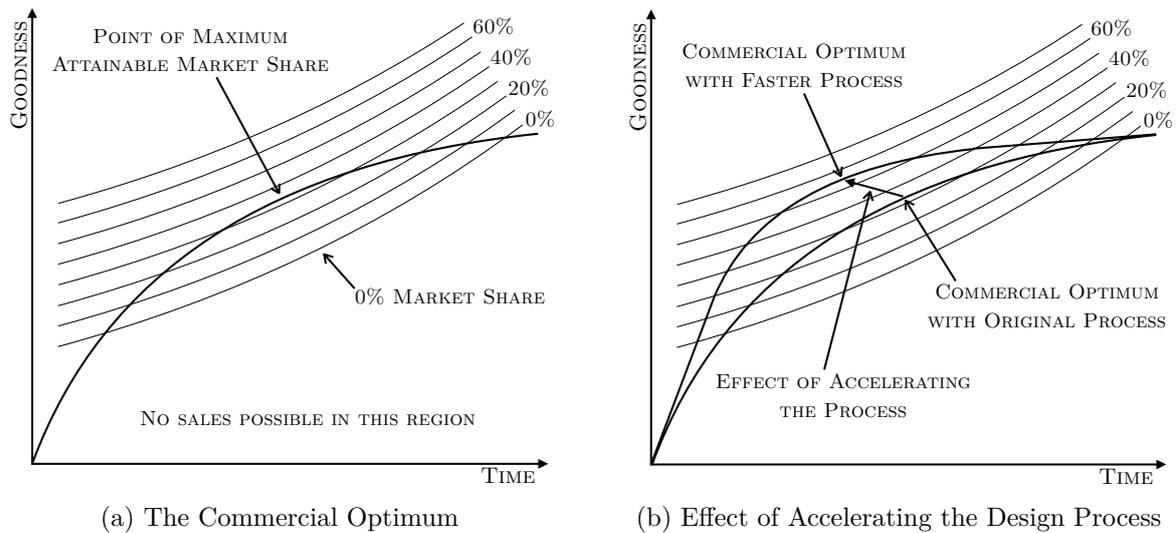


Fig. 3.1 Market share curves showing the commercial optimum and the effect of accelerating the design process [205].

3.3.2 Methods for Accelerating Design Optimisation

Optimisation can contribute to the acceleration of the design process by harnessing the rapid analysis capabilities of computers and employing methodologies for efficient search of the design space. However, as mentioned previously, if care is not taken then the iterative nature of optimisation can result in numerous unnecessary and expensive simulations being carried out, ultimately delaying the development of the final product [53]. It is therefore important to consider computational cost when developing optimisation methodologies and in the following sections various techniques for achieving acceptable runtimes are discussed alongside their suitability for use within the new approach developed in this thesis.

3.3.2.1 Using Low-Fidelity Analysis Codes

One way to ensure that the computational cost associated with an optimisation is acceptable would be to employ low-fidelity analysis codes for performance prediction. This is the approach adopted in the traditional axial compressor design process described in the previous chapter, with preliminary design carried out using computationally cheap meanline analysis codes, enabling a large number of candidates to be assessed in a short period of time. The majority of previous applications of optimisation to the preliminary design of axial compressors opt for a similar approach, employing meanline [56, 81, 99, 104, 117, 130, 169, 185] or throughflow [100, 181] analysis codes to provide performance predictions during the search.

Whilst these low-fidelity models are able to provide meaningful performance information that is trusted by compressor designers [29, 76], failing to employ high-fidelity analysis during the early stages of the design process may result in fruitful avenues of improvement being

missed, with these opportunities only becoming apparent when a more realistic representation of the underlying physics is considered [28]. The low-fidelity analysis codes available for axial compressor design also rely on empirical correlations, limiting their applicability when novel designs are sought. High-fidelity analysis is essential in this scenario to ensure that performance of the new concepts has been adequately assessed. Another risk of employing low-fidelity codes is that defects in a design may not be detected until a later stage in the design process [28]. By this point remedial action is likely to be both expensive and time-consuming, and in some cases may prove impossible [76]. Using high-fidelity analysis at the preliminary design stage gives greater assurance that the predicted performance will be achievable in the final machine.

Missed improvement opportunities, a lack of reliability when analysing novel designs and potentially allowing defects to go undetected represent significant drawbacks to achieving acceptable computational cost through the use of low-fidelity models. Each of these problems can be addressed using high-fidelity analysis, and other techniques are therefore needed to accelerate optimisation using these more computationally expensive codes.

3.3.2.2 Faster High-Fidelity Analysis Codes

One way to provide the necessary reduction in computational cost might be to accelerate the high-fidelity analysis codes themselves. Over the years significant speed-up has been achieved in CFD using techniques such as multi-grid approaches [24] and running on graphics processing units (GPUs) [22]. These methods require significant expertise in the solution procedures themselves, and this limits acceleration of the high-fidelity analysis codes as a technique for reducing the computational cost associated with optimisation.

It is also often the case that any acceleration of the analysis codes is used to increase the overall fidelity of the design process rather than reduce the time taken to generate a final product. This is evidenced by the increased use of RANS analysis in academic optimisation studies as computing power has risen over the past two decades. As RANS analyses become cheaper, already on the horizon large-eddy and direct numerical simulation techniques are appearing with computational costs that are not too far from being feasible for use within optimisation routines [223].

In addition, the acceleration of the analysis codes themselves can only go so far. Even with efficient solution schemes and multi-grid approaches, CFD analysis of a multi-stage turbomachine can still take upwards of an hour to complete on a single workstation [52]. If optimisation methodologies are to be considered computationally feasible within an industrial context there is a need for a greater degree of acceleration than can be provided by just speeding up the high-fidelity analysis codes.

3.3.2.3 Parallelisation

A simple and effective means of speed-up is through parallelisation. This relies upon suitable computational resources being available with the capability of running several analyses in tandem to reduce the overall time taken for the optimisation. Whilst the ultimate saving achieved through parallelisation depends on the computational environment being used it is a potentially potent means of reducing the computational cost associated with an optimisation method. Parallelisation, and the degree to which it is possible, is therefore an important consideration when developing the new methodology to facilitate the use of optimisation in the preliminary design of axial compressors. However, it cannot be relied upon as the sole source of acceleration due to the degree of speed-up being highly dependent on the computational architecture available to the user.

3.3.2.4 Reduced Order Modelling

Another popular method for accelerating optimisation methodologies is through the use of reduced order models. These cost-efficient representations of large scale systems can be generated using various techniques [5], with the most popular being proper orthogonal decomposition [31]. Whilst effective at reducing cost, these approaches are unlikely to be compatible with the requirements for enhancing designer understanding discussed earlier in this chapter. Reducing the order of the system detaches the changes made by the optimiser from the design variables themselves, potentially making it difficult for users to determine the exact physical changes that have led to the observed performance improvements. For this reason reduced order modelling is not considered to be a useful acceleration technique in this instance.

3.3.2.5 Multi-Fidelity Methods

Multi-fidelity methods for the acceleration of high-fidelity optimisation have received a lot of attention in the literature over the past two decades. Peherstorfer et al. [189] define these as being approaches that use low-fidelity models to provide computational speed-up whilst ensuring accuracy through a reduced number of calls to a high-fidelity analysis code. Suitable low-fidelity models include those generated by fitting response surfaces to sparse high-fidelity data, those based on coarsened grids and relaxed convergence criteria for iterative solvers, and those employing reduced physics models, for example through the removal of viscosity to produce the Euler equations [69, 189].

Utilising multi-fidelity techniques to accelerate the new methodology under development in this thesis is appealing for a number of reasons. Firstly, the presence of trusted physics-based low-fidelity codes, in the form of the meanline and throughflow methods discussed in Section 2.1.3, should allow for significant speed-up to be achieved. Including these low-fidelity

codes within the decision-making process could also help to overcome trust issues that have hampered the uptake of optimisation methods in the turbomachinery industry [214].

Secondly, multi-fidelity approaches have a good track record in the field of compressor optimisation. At the detailed design stage, where three-dimensional optimisation of the blade geometry is performed, the use of data-fit models in surrogate based optimisation (SBO) frameworks has become standard practice [156, 209]. Response surfaces are built from an initial dataset generated using design of experiments techniques, with several different methodologies employed to construct the low-fidelity models including Kriging [55, 96, 122, 128, 131, 154, 155], artificial neural networks [126, 146, 179] and radial basis functions (RBFs) [64, 144, 222]. The surrogates are used to predict performance within an optimisation loop, often being updated sequentially as in BO (see Section 3.2.2.3) to improve their accuracy as the optimisation progresses.

Some authors have also made use of physics-based low-fidelity codes to further reduce the number of high-fidelity analyses required. Brooks et al. [25] and Kim et al. [135] utilised coarsened grids in conjunction with co-Kriging and hierarchical Kriging respectively, two multi-fidelity response surfaces that incorporate data from different analysis codes to produce more accurate performance predictions. Goinis and Nicke [89] also employed a co-Kriging technique in conjunction with *SM* approximations of differing fidelities. Schemmann et al. [211], optimising the impeller of a centrifugal compressor, used one-dimensional analysis to filter the initial design of experiments, halving the number of high-fidelity points required to produce a response surface of a given accuracy.

These surrogate based methods have successfully been used to accelerate high-fidelity optimisation at the detailed design stage. However, the need for a large initial dataset limits their applicability to preliminary design. Eastwood et al. [60], showed that producing even a single viable starting geometry for such problems can be challenging, never mind several hundred that are suitably spread out around the design space to form a valid training dataset for a response surface. Despite this difficulty, Hendler et al. [101] did manage to construct a RBF response surface to accelerate optimisation of a 10-stage machine using a throughflow analysis code.

Other authors have employed alternative multi-fidelity techniques to accelerate optimisation at the preliminary design stage. Jarrett and Ghisu [113, 115] used a trust-region model management framework to accelerate the RANS-based optimisation of an integrated multi-stage compression system. At each iteration a corrected low-fidelity model was used to find promising points within a restricted area of the design space (the trust-region), with these points then analysed using the high-fidelity code. The analysis results informed an update to the size of the trust-region, with the low-fidelity stage given more or less freedom depending on how well it approximated the high-fidelity performance. This approach was

shown to provide significant speed-up compared to a single-fidelity version of the underlying search algorithm [113, 115].

In an alternative approach, Poehlmann and Bestle [193] carried out sequential optimisations using meanline and throughflow analysis codes. The meanline results were used as starting designs for the throughflow search in an attempt to accelerate convergence of the higher-fidelity optimisation. The authors reported some success, but found that the results of the low-fidelity meanline stage often failed to converge when analysed using the throughflow code. This highlights the importance of making regular accuracy checks using a high-fidelity model to ensure that errors in the low-fidelity predictions are not leading the search astray.

Multi-fidelity approaches have a proven track record in both axial compressor design specifically and in the wider aerospace field [69, 189] and therefore represent a suitable acceleration technique for the new optimisation methodology developed in this thesis.

3.4 An Improved Optimisation Methodology

In the preceding sections three donor techniques have emerged with the potential to address the factors limiting current optimisation methodologies when applied to the preliminary design of axial compressors. Use of the MDR formulation should enable sophisticated problem definitions to be employed that more accurately reflect the desires of the designer, overcoming the limitations of current approaches when attempting to deal with the large number of performance parameters present early in the design process. The TS algorithm, meanwhile, results in rich and interpretable design development information being available to enhance understanding, assisting designers as they attempt to determine the physical justification for the observed performance improvements and supporting creativity and innovation through improved physical insight into the underlying problem. Finally, multi-fidelity methods, making use of the available physics-based low-fidelity analysis codes, should enable adequate acceleration of the search to satisfy the strict computational cost requirements that exist when employing optimisation in an industrial setting.

In this thesis these techniques are combined to generate an improved methodology that facilitates the use of optimisation in the design of axial compressors. Applications focus on the preliminary design stage due to the potential for larger performance improvements to be achieved at this early point in the design process. To maintain tractability just the aerodynamic design of machines is considered, with other disciplines accounted for through the application of heuristic constraints. As mentioned in the introduction, the aim is to improve the axial compressor design process rather than suggest enhancements to the machines themselves. Assessment of the developed approach therefore focusses on the ability to generate good designs in a computationally efficient manner, rather than specifics of the axial compressors that are produced.

Whilst TS and multi-fidelity methods have already been successfully applied to axial compressor optimisation [113, 115], the MDR formulation is a new technique that is at an early stage of the development process. So far the formulation has only been implemented within a GA and a PSO, with applications limited to the demonstrative trial problems presented in the introductory paper [39]. The first step towards a successful methodology that facilitates the use of optimisation in the aerodynamic design of axial compressors by overcoming the limiting factors outlined in this chapter is therefore to implement a TS algorithm that is capable of using the MDR formulation. This is carried out in the following chapter. Due to the infancy of the MDR approach there is then a requirement to mature the technique through the generation of a simplifying framework before it can be applied to the more complex problem of axial compressor design. The performance of the new TS implementation of the MDR approach is assessed and compared to existing methods using an aerofoil test case, with the ability to overcome the limitations associated with formulation and understanding confirmed through application to an initial six-stage axial compressor design problem. With these capabilities demonstrated, the technique can be accelerated by facilitating the use of low-fidelity models in conjunction with their high-fidelity counterparts. The resultant algorithm is shown to address each of the three factors limiting current optimisation methodologies when applied to the preliminary design of axial compressors.

Chapter 4

Tabu Search Using Multiple Dominance Relations

In the previous chapter potential solutions to the problems limiting current optimisation methodologies when applied to the preliminary design of axial compressors were presented. These included the MDR formulation to enable more adequate handling of the large number of performance parameters that need to be considered early in the design process, the TS algorithm to ensure the optimisation process is fully equipped for the role of enhancing designer understanding, and use of multi-fidelity methods to accelerate the search for good designs and reduce computational cost.

During initial development of the MDR formulation it was only implemented within a GA and a PSO. Therefore the first step towards an improved methodology that facilitates the use of optimisation in the aerodynamic design of axial compressors is to implement the new formulation within a TS algorithm that is capable of providing interpretable design development information to enhance understanding. The new implementation is presented in this chapter, with the resultant algorithm validated and compared to existing methods using an analytic test case.

4.1 Tabu Search

As discussed in the previous two chapters, TS is an optimisation algorithm based on pattern search in which new designs are generated by altering one variable at a time by a known amount. This behaviour allows the path taken through design space by the optimiser from the initial to final designs to be reconstructed, resulting in interpretable design development information that can help the user determine the physical reasoning behind the observed performance improvements. The approach was originally developed in single-objective form by Glover and Laguna [88] before being extended to multi-objective scenarios by Jaeggi et

al. [110]. This multi-objective TS (MOTS) algorithm has seen extensive use in engineering design optimisation and forms the basis of the new implementation incorporating the MDR formulation.

4.1.1 Basic Search Mechanism

A flowchart for the MOTS algorithm described by Jaeggi et al. [110] is shown in Figure 4.1. At each iteration new candidate designs are generated in the vicinity of the current point using a pattern search method. The approach of Hooke and Jeeves (H&J) [105] is usually employed in which each variable is incremented and decremented by a set amount in turn to produce new design vectors. A random sample of these candidate designs is selected for analysis which can be conducted in parallel to reduce the time taken for the optimisation if the appropriate computational architecture is available. If any of the newly analysed points dominate the current design then one of these is selected as the pattern search centre for the next iteration. If none of the newly analysed candidates dominate the current point then a further random sample is analysed until no candidates remain. At this stage the best available candidate is selected even if it is dominated by the central point, providing a way for the optimiser to move through the design space and climb out of local minima.

4.1.2 Heuristics to Enhance the Searching Capability

TS uses a series of heuristics to enhance the searching capabilities of this basic method. The first involves a short term memory (STM) which stores all of the points already used as the centre of a H&J iteration. These points are considered “tabu”, with the optimiser unable to select them as the central point for the next pattern search. This improves the ability of the optimiser to climb out of local minima and encourages a more widespread search of the design space by preventing it from returning along the same path.

The remaining heuristics rely on the use of a counter. Following each H&J iteration a medium term memory (MTM) containing the current non-dominated solutions is updated. If at least one new point is added to the MTM during this update then the counter is reset to zero. Otherwise, if no new points are added to the MTM, the counter is incremented to give a measure of the time since the last new non-dominated design was found. When the counter reaches user-defined critical values a series of procedures are triggered that aim to further improve the searching capability of the algorithm.

The first of these, intensification, attempts to reinvigorate the search by moving to the best point found that is yet to be used as the centre of a H&J pattern search. The second, diversification, encourages global exploration by moving the search to the most remote point in design space analysed so far. Finding the most remote point relies upon the use of an archive known as the long term memory (LTM) that stores details of every design analysed throughout the optimisation. Finally, step size reduction aims to further improve upon the

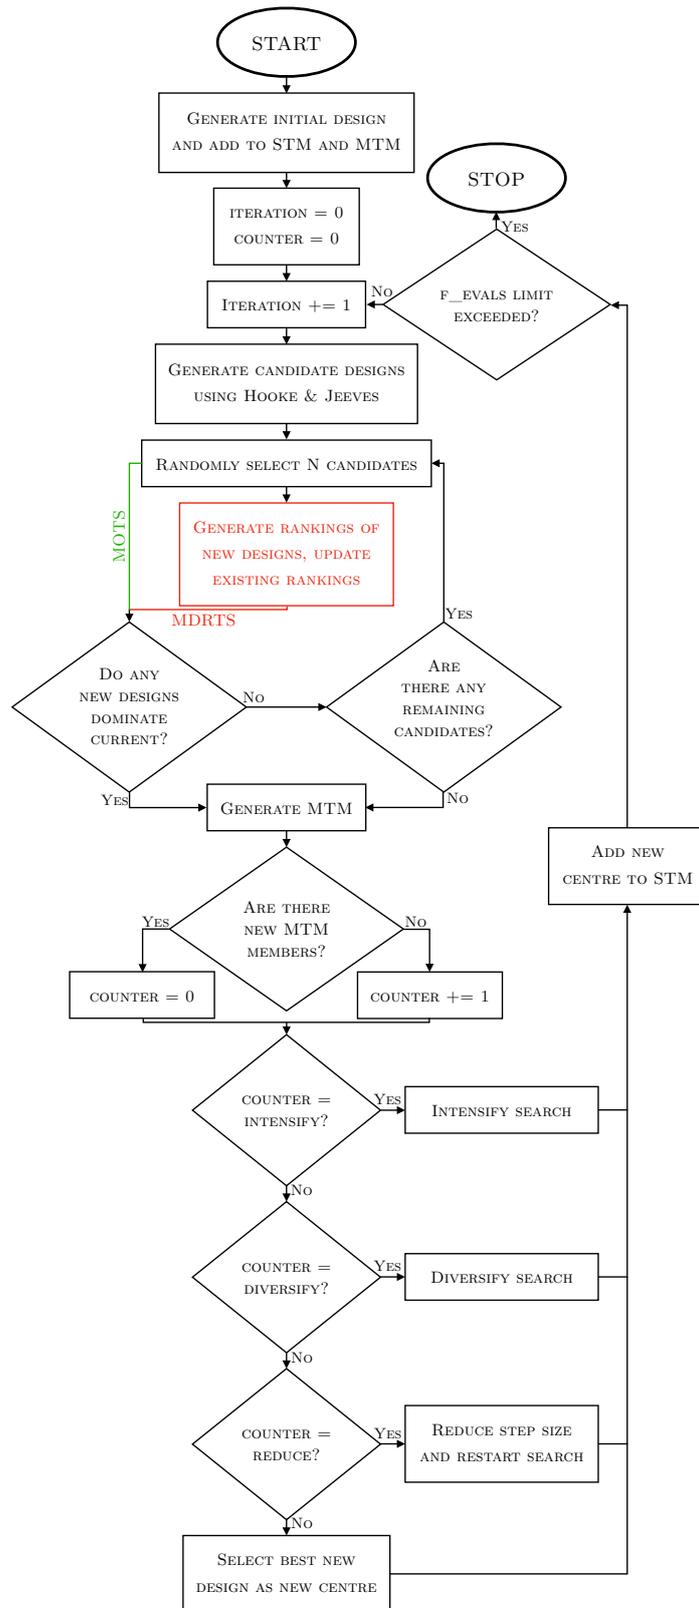


Fig. 4.1 Flowcharts for MOTS (green) and MDRTS (red).

best designs found by moving to a random member of the MTM and reducing the amount by which each variable is changed during a H&J iteration. The STM is emptied following step size reduction as new information can be obtained by conducting pattern searches around previously visited points using smaller design variable changes.

4.1.3 Previous Applications to Engineering Design Problems

The MOTS algorithm has been successfully applied to a variety of engineering design problems. Ghisu et al. [82], Jarrett and Ghisu [113, 115] and Kipouros et al. [137] all used MOTS in their axial compressor optimisation studies, with the latter employing a multi-fidelity version of the basic algorithm built around a trust-region methodology. Further work by Ghisu et al. [78] applied MOTS to robust aerofoil optimisation allowing for the effects of ice accretion, whilst the algorithm has also seen use in the renewables sector to optimise a horizontal axis wind turbine [70] and the blades of a Wells turbine [93]. Other areas of application include turbomachinery blade design [136], high-lift aerofoil optimisation [230] and development of a civilian unmanned aerial vehicle [7].

4.1.4 Limitations of Tabu Search

Despite these successful applications TS does have limitations, primarily in the lack of a rigorous convergence guarantee. When employing the algorithm in engineering optimisation studies this is not usually a problem as the designer is rarely interested in finding truly optimal designs. Instead the optimiser is used to explore the design space and produce the largest performance improvement possible given a limited computational budget [38].

The local nature of the algorithm, considering just a single design at each iteration also restricts the explorative capability of the search, especially when compared to population-based methods such as GAs and PSO. However, the likes of Ghisu et al. [81] have still successfully used MOTS to generate trade-off curves in a computationally efficient manner.

Finally, the use of a H&J pattern search, changing only one design variable at a time, means that any synergies involved with altering different parameters simultaneously will not be uncovered, potentially resulting in improvement opportunities being missed by the algorithm [104]. Evidence in the literature suggests that MOTS is still able to make progress on axial compressor design problems despite this limitation [82, 113, 115, 137]. Perhaps more importantly, changing one design variable at a time allows for the provision of interpretable development information to enhance designer understanding, outweighing any minor losses in exploratory performance.

4.2 Incorporating Multiple Dominance Relations

As discussed in the previous chapter, the MDR formulation can be incorporated into existing optimisation algorithms through the use of ranking functions [39]. Designs are given a separate ranking for each dominance relation, indicating how good that design is according to that relation. Cook et al. [39] suggest two possible ranking functions to fulfil this purpose. The first is the non-dominated sorting method [90] in which the rank of a design is equal to the number of non-dominated fronts that must be removed before it becomes non-dominated. The second is a count of the number of other designs that dominate a given point [71]. Results in the original work [39] show little difference between these two ranking functions, with the simpler second method preferred here.

When determining whether one design is better than another these rankings are compared in place of the usual direct comparison of objectives. This is carried out using lexicographic ordering, meaning the rankings corresponding to the first dominance relation (\preceq_1 in Equation 3.4) are compared first, with the rankings corresponding to \preceq_2 only consulted if those corresponding to \preceq_1 are equal. This continues to the n th relation, meaning two designs are only considered equivalent (i.e. neither dominates the other) if their rankings are identical for every dominance relation.

Using this rankings approach Cook et al. [39] implemented the MDR formulation within a GA and a PSO, making only minor changes to the underlying algorithms. Similarly, only limited modifications to the MOTS algorithm described in the previous section are required to facilitate the use of MDR. The sole notable addition, highlighted in Figure 4.1, is the generation and updating of the rankings when new designs are analysed. The rankings of each new design are calculated by comparing it to all previously analysed points and counting the number that dominate it according to each dominance relation. During this comparison the rankings of any existing designs that are dominated by the new point are updated accordingly to reflect the new information.

A method for carrying out this process is presented in Algorithm 1. The computational cost of the procedure is $O(nMN^2)$, where n is the number of dominance relations being used, M is the number of values defining each of these relations, and N is the number of previously analysed points [39], with all designs found during the search needing to be considered to avoid cyclic behaviour [18]. For optimisation using high-fidelity analysis N is usually restricted to $O(1000)$ due to time constraints, meaning the additional computational expense of calculating the rankings is negligible when compared to the cost of even a single high-fidelity analysis. However, the reliance on a library of previously assessed points may limit the MDR formulation if applied to problems where a much larger number of designs are considered.

Tabu Search Using Multiple Dominance Relations

Algorithm 1 Procedure for generating and updating the rankings following analysis of a new design.

```
function UPDATE_RANKINGS(new_design, existing_designs, dominance_relations)
  for design in existing_designs do
    for dom_rel in dominance_relations do
      if design dominates new_design then
        new_design.rankings[dom_rel] += 1
      else if new_design dominates design then
        design.rankings[dom_rel] += 1
      end if
    end for
  end for
  return new_design, existing_designs
end function
```

The only other alteration to the original MOTS algorithm required to incorporate the MDR formulation is to replace the direct comparison of objectives with the lexicographic comparison of the rankings whenever a choice is made between designs. This occurs when deciding whether to advance the optimisation to a newly analysed candidate and when filling the MTM with non-dominated designs.

With these modifications the MOTS algorithm is equipped to utilise the MDR formulation, producing a new multiple dominance relations tabu search (MDRTS) method. This implementation can handle more sophisticated problem definitions without adding significant computational overhead and retains the ability to provide interpretable design development information to enhance designer understanding.

4.3 Additional Algorithmic Alterations

In addition to the modifications facilitating the use of MDR a number of further changes are made to enhance the effectiveness of the MOTS algorithm developed by Jaeggi et al. [110].

Originally the STM was limited in size, storing only the most recently visited points to ensure memory allocation constraints were not breached [88, 110]. In the applications considered in this thesis the maximum number of iterations is $O(100)$ due to strict limits being imposed on the number of computationally expensive analyses that can be carried out. Given the relatively small number of iterations sufficient memory is not a problem, allowing the limit on the size of the STM to be removed. This ensures that each H&J pattern search is conducted around a new central point, increasing the extent of the design space that is explored.

Intensification in this implementation follows that of Jaeggi et al. [110], with the best design not yet present in the STM selected as the centre for the next H&J iteration. In the

4.3 Additional Algorithmic Alterations

diversification approach the range of each variable is split in half, resulting in $2^{n_{var}}$ regions of design space, where n_{var} is the number of design variables. All previously analysed designs are allocated to the appropriate region using their design vectors and a design from the region with the smallest number of existing points within it is selected as the new centre.

When a pattern search centre is suggested by either intensification or diversification in the original algorithm the step size used during the subsequent H&J iteration remains unchanged. However, initial experiments by the author found that this often resulted in small step sizes being used in the latter stages of the search to explore regions of design space that had yet to receive significant coverage. It would be more appropriate to use a larger step size within these under-explored regions, allowing a greater extent of the design space to be covered more efficiently, accelerating the search for good designs.

In the implementation used here, therefore, the step size is reset to the level at which the newly selected central point was discovered. This means that if a design is generated by a pattern search using a large step size, but is only selected as a central point through intensification or diversification at a later stage when the step size has been reduced, the algorithm returns to using the larger step size when continuing the search from that point. This requires separate STMs to be stored for each step size, as designs should only be considered “tabu” if they have been used as the central point for a H&J pattern search employing the current step size or smaller. Again due to the relatively small number of iterations in the optimisations conducted in this work the memory allocation required for these additional STMs is not significant. The ability to alter the step size so that it more adequately reflects the extent to which a region of design space has been explored should ensure more efficient use of the limited computational budget available to the optimiser.

A final alteration also enhances the explorative capabilities of the search algorithm. When intensification and diversification fail to produce new members of the MTM the TS algorithm resorts to reducing the step size in the hope of further improving good designs that have already been found. In the current algorithm this may result in a smaller step size being selected before all members of the MTM have been used as centres for pattern searches employing the existing amount of design variable change. This could lead to improvement opportunities being missed and to an unnecessary reduction of the step size, limiting the explorative capabilities of the algorithm.

To avoid this scenario when step size reduction is triggered the new implementation first checks whether all of the designs in the MTM are also present in the STM for the current smallest step size. If they are all in both memories then step size reduction is carried out as normal. If any are not then one of these is selected at random as the new pattern search centre and the value of the counter is reduced by one. The algorithm then conducts a H&J iteration around the chosen MTM design using the current smallest step size. If no new additions are made to the MTM then the counter is incremented, triggering step size reduction again. This

process is repeated until either a new addition is made to the MTM, resetting the counter to zero, or all MTM designs have been used as pattern search centres, resulting in the step size being reduced. This approach ensures that all potential improvement opportunities are adequately explored before the algorithm reduces the step size, improving the efficiency of the search as it moves through the design space.

4.4 Implementation Validation

To validate the TS implementation of the MDR approach it is applied to an analytic test case developed by Cook et al. [39]. The performance of the new MDRTS algorithm is compared to the two population-based implementations developed in the original work [39], as well as to multi-objective formulations of the problem solved using the MOTS algorithm. These comparisons validate the implementation of the MDR formulation and demonstrate whether incorporating the new approach through ranking functions is effective within an algorithm based on pattern search. They also provide an initial indication of the performance of the new MDRTS algorithm compared to existing methods.

4.4.1 Analytic Test Case

The analytic test case employed for this initial validation was developed by Cook et al. [39] specifically to demonstrate the benefits of the MDR formulation. It consists of four design variables, each in the range $[-5, 5]$, and three performance metrics defined in Equation 4.1.

$$\begin{aligned} f_1 &= \sqrt{(x_1 - 4)^2 + 25(x_2 - 0)^2 + 25(x_3 - 0)^2 + 25(x_4 - 0)^2} \\ f_2 &= \sqrt{(x_1 + 4)^2 + 25(x_2 - 0)^2 + 25(x_3 - 0)^2 + 25(x_4 - 0)^2} \\ f_3 &= \sqrt{25(x_1 - 3)^2 + 25(x_2 - 0)^2 + 25(x_3 - 0)^2 + (x_4 - 1)^2} \end{aligned} \tag{4.1}$$

The goal is to find the design that minimises f_3 within the Pareto front defined by minimising f_1 and f_2 , and a single optimum is located at $\mathbf{x} = [3, 0, 0, 0]$.

4.4.2 Comparative Implementations

Two comparisons are carried out to validate and assess the performance of the new MDRTS algorithm. Firstly, results generated by the new method are compared to those found using two more traditional multi-objective problem definitions and MOTS as the underlying algorithm. This reveals whether the MDR formulation is behaving as expected when combined with TS. The three problem formulations used are the same as those adopted by Cook et al. [39] when applying this test case for a similar purpose. They are a two-dimensional multi-objective optimisation using f_1 and f_2 as objectives; a three-dimensional multi-objective

Table 4.1 Parameter settings for the TS, GA and PSO.

Algorithm	Parameter	Value	Description
Tabu Search	N_SAMPLE	8	N ^o randomly selected for analysis
	INTENSIFY	5	Counter value for intensification
	DIVERSIFY	10	Counter value for diversification
	REDUCE	15	Counter value for step size reduction
	STEP_FAC	0.5	Step size reduction factor
	INIT_STEP	0.1	Initial step size 10% of variable range
Genetic Algorithm	TOURN_SIZE	2	Tournament size for parent selection
	CROSS_MAG	0.1	Blended crossover magnitude
	MUT_RATE	0.1	Mutation rate
	MUT_MEAN	0.0	Mean of Gaussian used for mutation
	MUT_STD	1.0	Std. dev. of Gaussian for mutation
Particle Swarm Optimisation	MUT_RATE	0.5	Mutation rate

optimisation with all three performance metrics used as objectives; and the MDR approach using the nested relation presented in Equation 3.3, with Pareto dominance of f_1 and f_2 as the first dominance relation, \preceq_1 , and real number ordering of f_3 as the second, \preceq_2 . The two multi-objective problem definitions could be seen as special cases of the MDR formulation with Pareto dominance of the selected objectives used as the sole dominance relation. The MDRTS approach therefore inherits all of the capabilities of the MOTS algorithm, but with the added facility to incorporate more sophisticated problem definitions.

The second comparison is to results generated using the previously developed GA and PSO implementations of the MDR approach [39]. These are based on NSGA-II [48] and a multi-objective PSO developed by Coello et al. [35]. The same MDR problem formulation is employed by both algorithms, allowing an initial performance assessment of the TS implementation compared to these population-based methods.

One problem faced when attempting to compare population-based algorithms with those utilising pattern search is that the former start from an initial set of designs whereas the latter require just a single starting point. This can make it challenging to determine whether differing results are due to the underlying operation of the algorithms or simply a consequence of one receiving a more favourable set of starting points. In order to provide a fairer comparison of the search mechanisms themselves both TS algorithms are modified to start from an initial population of designs, with the best selected as the central point for the first H&J pattern search. This allows direct comparison of the results generated by the different methods.

Each approach is run from the same 20 randomly generated sets of starting points using a population size of 50 and a computational budget of 1000 calls to the function calculating the objectives. Various parameter settings for the different algorithms are shown in Table 4.1.

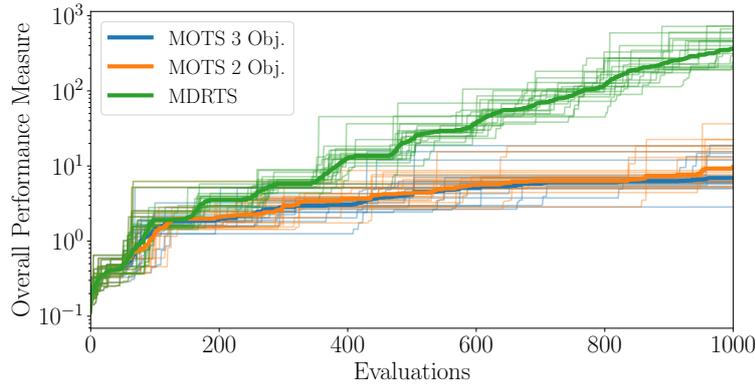


Fig. 4.2 Progression of the best overall performance measure found using the multi-objective and MDR formulations applied to the analytic test case. All 20 runs are shown as faint lines with the mean plotted in bold.

4.4.3 Results

Figure 4.2 compares the performance of the new MDRTS algorithm to the two multi-objective formulations using MOTS. The progression of the best value of an overall performance measure is tracked for each of the different runs, with the average shown in bold. This overall performance measure is defined as the reciprocal of the minimum Euclidean distance to the known optimum in a space containing the four design variables.

The results demonstrate similar behaviour to that observed by Cook et al. [39] using the population-based implementations. The multi-objective formulations make initial progress towards the optimum but stall after around 160 evaluations due to their inadequate expression of the true goal of the problem. The two-objective approach fails to account for the desire to minimise f_3 , whilst the three-objective method does take this into account but in an inefficient manner, requiring a three-dimensional trade-off that cannot be sufficiently resolved using the given computational budget. The formulation utilised by the new MDRTS algorithm, in contrast, represents the goal of the problem more accurately and as a result produces values of the overall performance measure that are on average around two orders of magnitude higher than those found using the existing methods. The speed of convergence does not diminish even after 1000 function evaluations, with values of the overall performance measure growing throughout the optimisation. The superiority of the MDR approach appears to increase as more points are analysed. This is perhaps a result of the comparative relation improving as more information is generated to calculate the rankings. With a greater extent of the design space represented these functions start to reflect the true performance of different designs in relation to one another more accurately. In some senses the fidelity of the MDR relation itself is being elevated as the search progresses and more designs are included in the rankings calculation. The results in Figure 4.2 validate the TS implementation of the MDR approach,

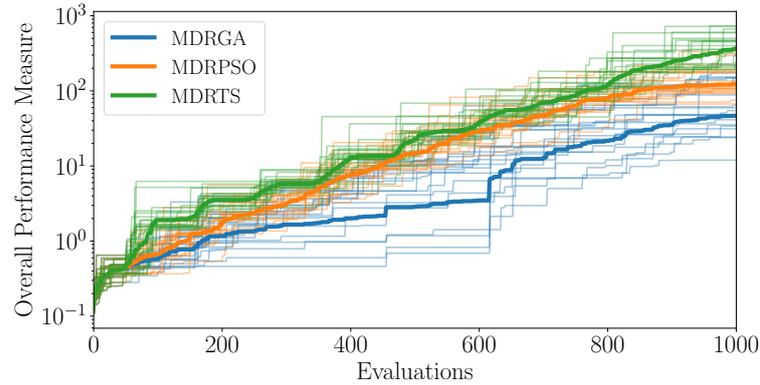


Fig. 4.3 Progression of the best overall performance measure found using the different MDR implementations applied to the analytic test case. All 20 runs are shown as faint lines with the mean plotted in bold.

demonstrating that the benefits of the more sophisticated problem formulation persist within this algorithm based on pattern search.

Figure 4.3 compares the performance of the three different implementations of the MDR formulation, with this again plotted in terms of the overall performance measure defined earlier. The new TS method outperforms both population-based approaches when applied to this analytic test case. The MDRTS algorithm produces good designs more efficiently, on average generating equivalent performance to that found using the PSO in 820 evaluations, and matching the results produced by the GA after 650 calls to the function calculating the objectives. These correspond to respective computational savings of 18% and 35%. The new implementation also continues to produce designs closer to the known optimum throughout the search, with the final values of the overall performance measure being on average three times greater than the PSO results and 10 times that produced using the GA.

One possible reason for this superior performance is the natural balance between exploration and exploitation in the two types of algorithm. By making large changes to the design vector, population-based methods inherently carry out a more thorough search of the entire design space than is possible using the smaller changes of an algorithm based on pattern search. A wide-ranging search reduces the likelihood of promising regions being missed, but can mean that the potential of good designs is not fully exploited. Pattern search approaches, in contrast, are good at exploiting the potential of designs due to their incremental movement through the design space. In this analytic test case comprising just four variables the design space is relatively small and easy to navigate, reducing the importance of exploration and favouring the TS implementation that is more adept at exploitation.

The results presented in this section validate the MDRTS algorithm as a successful optimisation methodology and motivate further development to enable application to the preliminary design of axial compressors.

4.5 Summary

The promising MDR approach has been successfully implemented within a TS routine. The resultant MDRTS algorithm, validated using an analytic test case, is well equipped to overcome problems associated with both formulation and understanding that limit current optimisation methodologies when applied to the preliminary design of axial compressors. These capabilities are assessed in later chapters. However, before that can be carried out further maturing of the MDR formulation is required through the development of a simplifying application framework that is the subject of the following chapter.

Chapter 5

A Simplifying Application Framework

Whilst the TS algorithm has been applied to a range of engineering design optimisation problems, the new MDR formulation has only recently been introduced and lacks maturity. This means there is an absence of guidance available for how to apply the new formulation, and by extension the MDRTS algorithm developed in the previous chapter, to more complex problems such as the preliminary design of axial compressors. Faced with a long list of performance parameters it may be difficult for a designer to determine the best way for them to be arranged within the nested hierarchy of dominance relations at the heart of the new formulation. To overcome this problem a simplifying framework is proposed in this chapter that harnesses the main benefits of the MDR formulation without adding complexity from the perspective of the designer tasked with setting up the optimisation problem.

5.1 An Additional Performance Parameter Classification

One positive feature of the traditional objectives-and-constraints formulation for optimisation is simplicity. When formulating a problem all the designer needs to do is designate the relevant performance parameters as being either objectives or constraints and specify values for any limits imposed. Ideally, formulating a problem using MDR should be no more difficult, and here it is suggested that the main benefits of the new approach can be realised by introducing one additional performance parameter classification alongside the objectives and constraints. This new classification incorporates parameters that the designer wants to improve but with secondary importance compared to those designated as objectives. Quantities of interest assigned to this classification are therefore referred to as “desirable features”.

This type of performance parameter formed part of the original motivation for developing the MDR formulation discussed in previous chapters. The desires of the designer for quantities

A Simplifying Application Framework

of interest such as pitching moment and trailing edge separation in aerofoil design and M_{exit} and α_{exit} in axial compressor problems are often difficult to adequately represent using objectives or constraints. Instead, using the new framework, these performance parameters can be assigned to the desirable features classification, aligning more closely with what the designer actually wants.

Performance in terms of the objectives, constraints and desirable features can be assessed using three separate dominance relations, with designs given distinct rankings based on how they perform in terms of the quantities of interest assigned to each classification. Arranging these dominance relations in the nested formulation shown in Equation 5.1 produces an optimisation problem definition that is able to reflect the desires of the designer more accurately than would be possible using objectives and constraints alone.

$$\mathbb{X}_{optimal} := \min(\min(\min(\mathbb{X}, \preceq_1), \preceq_2), \preceq_3) \quad (5.1)$$

This simplifying application framework makes formulating optimisation problems using MDR no more difficult than when employing the traditional objectives-and-constraints approach. The designer simply has one additional classification available when assigning performance parameters. Whilst complexity from the perspective of the designer has not increased, the new approach allows sophisticated problem formulations to be employed, facilitating more accurate handling of quantities of interest that were previously difficult to satisfactorily designate as either objectives or constraints.

5.2 Arrangement Within a Nested Structure

The dominance relations assessing performance in terms of the three performance parameter classifications need to be arranged within the nested hierarchy shown in Equation 5.1. There are six possible permutations for the arrangement of these three dominance relations, outlined in Table 5.1. The presence of multiple potential orderings means an experiment is required to determine whether the arrangement is significant and if a particular permutation of the dominance relations can be recommended as most likely to produce designs that are of interest to the designer.

Before conducting any experiments it can be shown that not all of the permutations in Table 5.1 are permissible for all problems. For example, if just one parameter is assigned to the objectives classification then this must appear as the final dominance relation, \preceq_3 , in any ordering. This is because only a single design will ever be considered non-dominated according to that relation as dominance is determined by ordering the parameter values and it is rare for two designs to exhibit exactly the same value of a given quantity of interest. With only one design being considered non-dominated the lexicographic comparison would not need to go past the objectives dominance relation to select between designs, resulting

5.2 Arrangement Within a Nested Structure

Table 5.1 The six permutations of the performance parameter classifications within the nested hierarchy of dominance relations shown in Equation 5.1.

	A	B	C	D	E	F
\preceq_1	Objectives	Objectives	Des. Feat.	Des. Feat.	Constraints	Constraints
\preceq_2	Des. Feat.	Constraints	Objectives	Constraints	Objectives	Des. Feat.
\preceq_3	Constraints	Des. Feat.	Constraints	Objectives	Des. Feat.	Objectives

in subsequent dominance relations being ignored. Therefore, to ensure the constraints and desirable features are taken into account by the optimiser, the objectives must appear as the final dominance relation, \preceq_3 , in any problem formulation where just one parameter is assigned to the objectives classification. This is the case in two of the six permutations in Table 5.1, D and F, with these being the only permissible orderings of the dominance relations for problems using a single objective. When a trade-off between multiple objectives is sought all six of the formulations are permissible as Pareto dominance allows two designs to be deemed equivalent, leading to the optimiser consulting further dominance relations to determine any preference between designs.

The same argument applies when only one parameter is assigned to the desirable features classification. However, this scenario is not considered here partly to maintain tractability, but also because it is unlikely to occur in practice. In the literature review conducted in Chapter 2 there was rarely a lack of parameters that the designer wanted to improve. In fact, problems associated with formulation arose due to the presence of numerous quantities of interest that could have been treated as objectives. It is therefore concluded that assigning a single performance parameter to the desirable features classification would be uncommon. If it did occur then in most cases selecting a second desirable feature from the list of relevant performance parameters should be relatively straightforward.

To determine whether the ordering of dominance relations is significant, and if any option might be recommended as the most likely to generate interesting designs, a set of experiments is conducted using both single-objective and trade-off scenarios. This demonstrates whether the significance of nesting order is consistent across the different settings in which the MDRTS algorithm might be employed.

5.2.1 Aerofoil Test Case

The experiment to determine the significance of the arrangement of dominance relations within the nested hierarchy shown in Equation 5.1 demands a richer test case than the analytic problem used in the previous chapter. The design of two-dimensional transonic aerofoil

Table 5.2 Hicks-Henne bump locations and bounds.

Surface	Location (x/chord)	Lower Bound ($\times 10^{-3}$)	Upper Bound ($\times 10^{-3}$)
Upper	0.05	-10	10
	0.2	-15	15
	0.45	-15	15
	0.7	-10	15
	0.85	-5	10
Lower	0.05	-10	10
	0.2	-15	15
	0.45	-15	15
	0.7	-10	15
	0.85	-5	10

sections, commonly employed in aerospace design optimisation research [4, 85, 97, 148, 248], meets this requirement.

The set-up used here is similar to that of Cook et al. [39] in their initial demonstration of the MDR formulation. A Reynolds number of 6.5 million is selected along with a Mach number of 0.8 and a fixed angle of attack of three degrees. The aerofoil surface is parameterised using 10 Hicks-Henne bump functions whose locations and bounds are shown in Table 5.2. Analysis is provided by the open source CFD software SU² [182] solving the RANS equations, with the RAE2822 mesh provided with the software used as a baseline (shown in Figure 5.1) that is deformed using a further program within the SU² suite. A single analysis using this mesh takes just over 5 minutes using 8 Intel Xeon 2.13 GHz CPUs.

Figure 5.2 shows contours of Mach number around the baseline RAE2822 aerofoil. This demonstrates the transonic nature of the problem, with a shockwave situated on the upper surface of the aerofoil followed by a region of separated flow. The key performance parameters are lift, drag, pitching moment, the amount of trailing edge separation and some measure of the space available inside the aerofoil for structural requirements and fuel storage. The analysis software outputs the first three of these in the form of coefficients (C_L , C_D and C_M respectively), whilst a measure of trailing edge separation is found by calculating the area under the skin-friction coefficient vs. chordwise distance curve whenever the former is negative. Following a number of other works [4, 143] the cross-sectional area, A_c , of the aerofoil is used as a measure of the available internal space.

5.2.2 Experimental Set-Up

The task of formulating suitable problems for the single-objective and trade-off scenarios is simplified by using the application framework suggested earlier in this chapter. The first

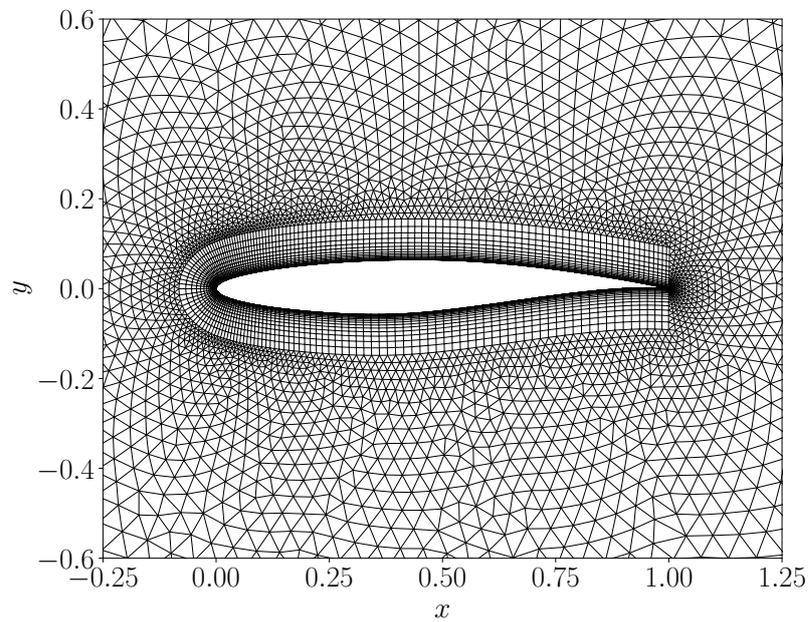


Fig. 5.1 Baseline mesh used in SU² RANS analysis.

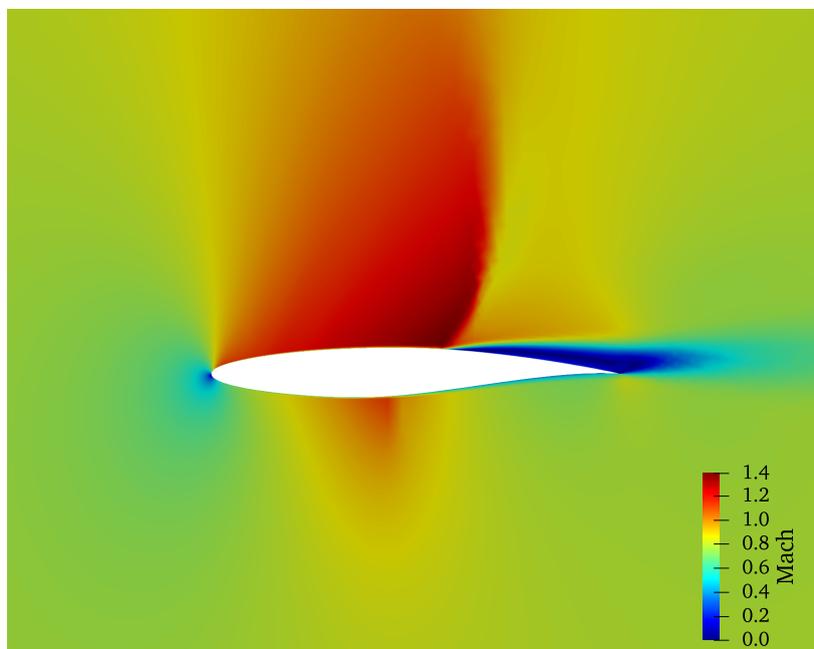


Fig. 5.2 Mach number around the baseline RAE2822 aerofoil.

A Simplifying Application Framework

requires improvement of a single performance parameter, and the aerofoil problem is therefore posed as a C_D minimisation subject to constraints applied to C_L and A_c . The lower limit for C_L is set at 0.6, whilst the internal area is required to be no less than that of an RAE2822 aerofoil with the same chord length (0.07784 m² when the chord length is 1 m). As suggested by Cook et al. [39] when developing the MDR formulation, C_M and trailing edge separation are quantities the designer wants to improve but with secondary importance compared to C_D and are therefore treated as desirable features. In the second scenario a trade-off is sought, with maximising C_L added to the objectives classification alongside minimising C_D . The C_L requirement is removed leaving the minimum A_c as the sole constraint. Again C_M and separation are assigned to the desirable features classification.

Optimisations using the dominance relation permutations in Table 5.1 are applied to these two problems, with just D and F used in the single-objective scenario and all six in the trade-off study. The MDRTS algorithm developed in the previous chapter is employed, with the ability to start from an initial set of designs, rather than a single point, retained in order to reduce the dependence of the results on the quality of the starting geometries. Each formulation is run from the same 10 sets of 20 randomly generated starting aerofoils allowing averaged results to be assessed, minimising the impact of stochastic elements of the algorithm on the conclusions drawn. The computational budget for each optimisation is 500 SU² evaluations, equivalent to around two days of continuous running time.

Whilst the dominance relations corresponding to the objectives and desirable features use the Pareto dominance criterion to select between designs, the constraints relation requires a handling methodology to ensure that aerofoils satisfying the specified limits are preferred. The penalty approach recommended by Coello [34] is employed for this purpose. An individual penalty term is calculated for each constraint as the amount of violation normalised by the limit itself, with these summed to provide an overall measure of constraint violation for a given design. Similar penalty approaches have been used in previous TS applications [82, 137], with these methods being preferable to barrier techniques, where infeasible designs are considered “tabu”, due to the severe restrictions placed on the design space. Ghisu et al. [82] noted how penalties avoid “excessive fragmentation of the design space”, whilst Kipouros et al. [137] stated that a penalty approach is necessary to navigate the highly non-linear characteristics of the search spaces that exist in aerodynamic design optimisation problems.

5.2.3 Single-Objective Results

Figure 5.3 shows the performance of designs found during the 10 runs using Formulations D and F in Table 5.1 applied in the single-objective scenario that satisfy the minimum C_L and A_c constraints and are also non-dominated in terms of C_D , C_M and separation. The optimisations employing Formulation F appear to have outperformed those using Formulation D, generating a larger number of designs that are non-dominated in terms of the objective

5.2 Arrangement Within a Nested Structure

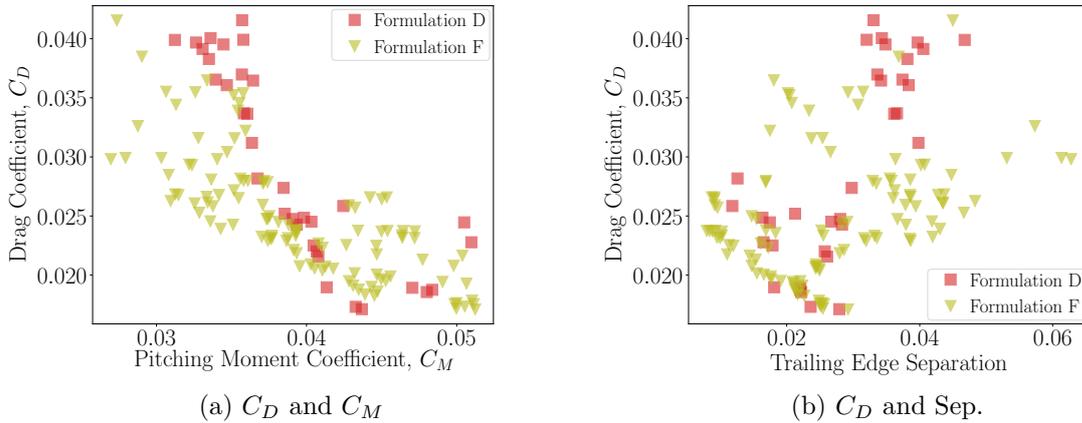


Fig. 5.3 Performance of the best designs found using Formulations D and F in Table 5.1 applied in the single-objective scenario with a minimum C_L requirement of 0.6.

and desirable features. Designs found using this permutation of the dominance relations also tend to have lower values of the two desirable features for a given level of C_D .

To further compare the two formulations a reference point approach is employed similar to that used in some of the techniques for many-objective optimisation mentioned in Chapter 3 [49, 153, 175, 206]. The reference point is given the best value of the objective and desirable features produced by the different runs applied in this scenario, considering all designs that satisfy the constraints. The Euclidean distance to this idealised design in a space containing normalised versions of the objective and desirable features is used to assess the convergence of an optimisation in two ways. Firstly, as the search progresses the minimum distance to the reference point is calculated, with the reciprocal of this value plotted to give an overall performance measure that is used to assess the speed at which an approach generates interesting designs. Secondly, at the end of the optimisation designs exhibiting performance within a given Euclidean distance of the reference point are tallied, indicating the quality of the final results.

Figure 5.4 tracks the highest values of the overall performance measure found by the searches conducted using Formulations D and F, with each of the 10 runs plotted individually and the mean shown in bold. Whilst the search using Formulation D makes good initial progress it appears to stall after around 250 evaluations, failing to produce further increases of the overall performance measure. In contrast, the aerofoils generated using Formulation F continue to exhibit improved values of this measure throughout the search.

Poor performance when employing Formulation D can also be seen in Figure 5.5 which shows the number of final aerofoils that exhibit performance within a given Euclidean distance of the idealised reference point. On average, the use of Formulation F leads to a larger number of designs exhibiting good performance in terms of the key quantities of interest than are found using Formulation D. In fact, three of the runs conducted using the latter

A Simplifying Application Framework

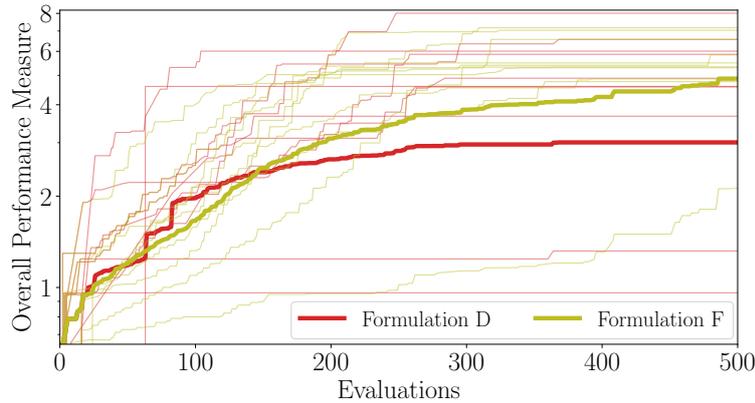


Fig. 5.4 Progression of the best overall performance measure found using Formulations D and F in Table 5.1 applied in the single-objective scenario with a minimum C_L requirement of 0.6. All 10 runs are shown as faint lines with the mean plotted in bold.

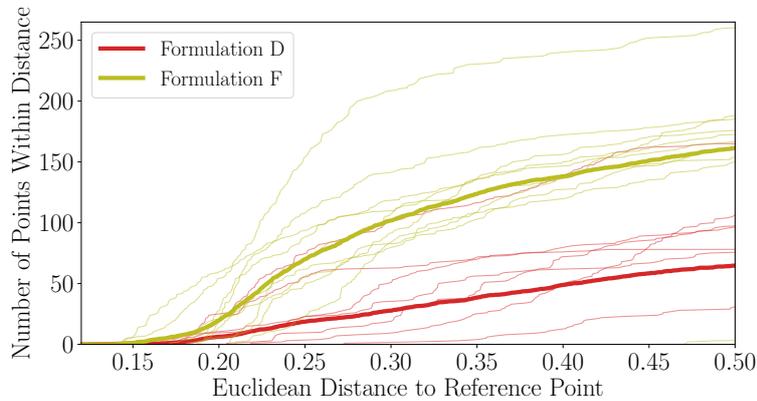


Fig. 5.5 Number of designs produced using Formulations D and F in Table 5.1 applied in the single-objective scenario with a minimum C_L requirement of 0.6 that exhibit performance within a given Euclidean distance of the idealised reference point. All 10 runs are shown as faint lines with the mean plotted in bold.

arrangement of dominance relations fail to appear on this figure, with the aerofoil exhibiting performance closest to that of the ideal design being at a Euclidean distance of greater than 0.5. In contrast, all 10 runs conducted using Formulation F are present, with just one, visible in the lower right corner of the figure, failing to produce designs with performance close to that of the reference point.

The reason for this superior performance is revealed in Figure 5.6a, which tracks the number of sufficient points found during each run using the two formulations. Whilst the searches employing Formulation F consistently find new designs that satisfy the constraints, every run using Formulation D reaches a point beyond which no new sufficient designs are generated. In some cases this occurs after around half of the computational budget has

been expended, but in the worst runs only a handful of designs that satisfy the C_L and A_c constraints are found during the entire optimisation

The remaining plots in Figure 5.6 track the performance of the aerofoils selected as centres for the pattern searches used by the MDRTS algorithm to generate new designs. They therefore provide an indication of how the searches are progressing, showing the performance of designs that the optimiser considers to be promising. Sudden large changes in performance visible on these plots are due to the intensification and diversification heuristics described in Section 4.1.2 that are employed to enhance the explorative capabilities of the TS algorithm.

Figure 5.6c shows that searches conducted using Formulation D consistently violate the C_L constraint, leading to the small number of sufficient designs seen in Figure 5.6a. Figures 5.6e and, to a lesser extent, 5.6f show that this is due to the optimiser prioritising improvements in the desirable features over satisfaction of the constraints. Designs exhibiting good performance in terms of C_M and trailing edge separation are accepted despite their C_L values being below the specified minimum. This behaviour is a result of the desirable features appearing ahead of the constraints in Formulation D (see Table 5.1). In contrast, the nested hierarchy of Formulation F prioritises satisfaction of the constraints, maintaining the focus of the optimiser on sufficient aerofoils that meet both design requirements and are therefore more likely to be of interest to the designer.

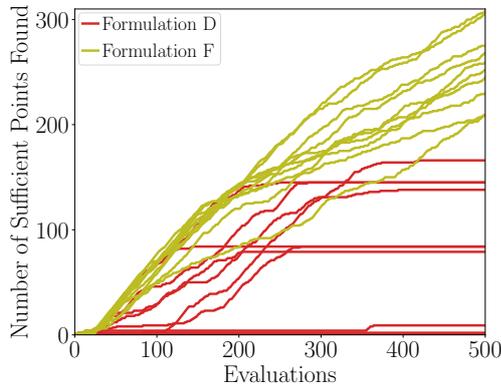
This problem of failing to focus on sufficient designs is exacerbated when the constraints become more difficult to satisfy. Figure 5.7 tracks searches conducted using a higher minimum C_L requirement of 0.7, started from the same 10 sets of initial designs. In this case searches using Formulation D, placing a higher priority on the desirable features than the constraints, are only able to find at most one sufficient design, wasting the computational budget analysing aerofoils that fail to meet the specified design requirements. Optimisations employing Formulation F, in contrast, consistently generate numerous sufficient aerofoils that are more likely to be of interest to the designer.

These results show that in a single-objective scenario the selected nesting arrangement has a significant impact on the results produced by the optimisation. To ensure other relations are considered when selecting between designs the sole objective should appear as the final element of any dominance relation ordering. Results generated using the aerofoil test case suggest that satisfaction of the constraints should appear as the first arbiter between designs to avoid wasting computational resources analysing insufficient points that are unlikely to be of interest to the designer. This leaves the desirable features to occupy the middle tier of the nested hierarchy.

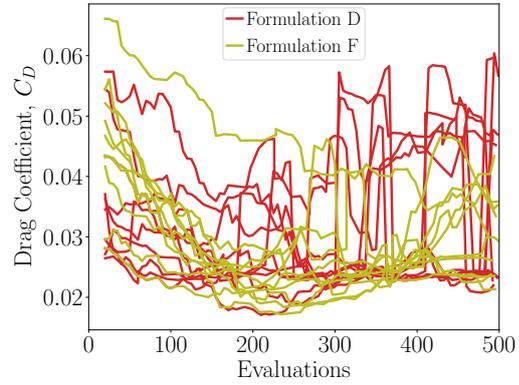
5.2.4 Trade-Off Results

Figure 5.8 shows the performance of designs on the C_L - C_D Pareto fronts produced using the six formulations in Table 5.1 applied in the trade-off scenario, considering all designs

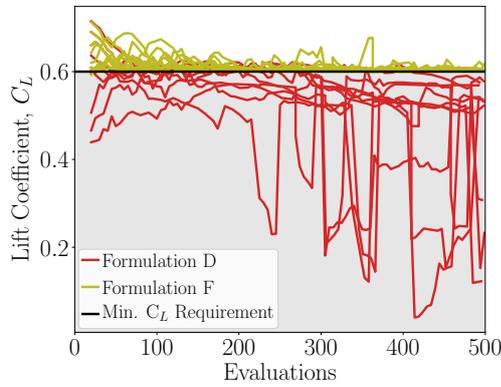
A Simplifying Application Framework



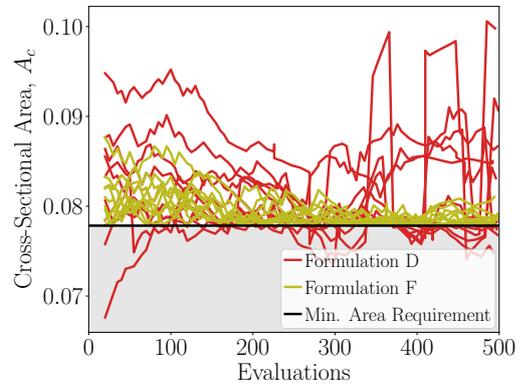
(a) Number of Sufficient Points Found



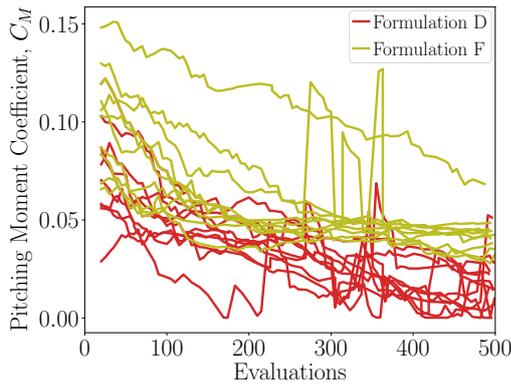
(b) C_D of Pattern Search Centres



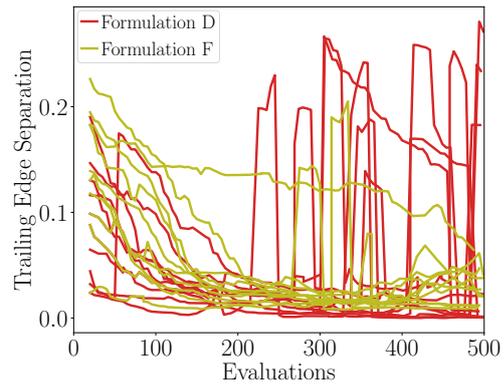
(c) C_L of Pattern Search Centres



(d) A_c of Pattern Search Centres



(e) C_M of Pattern Search Centres



(f) Sep. of Pattern Search Centres

Fig. 5.6 Progression of searches using Formulations D and F in Table 5.1 applied in the single-objective scenario with a minimum C_L requirement of 0.6. Figure 5.6a shows the number of sufficient points found and the remaining plots track performance of the pattern search centres.

5.2 Arrangement Within a Nested Structure

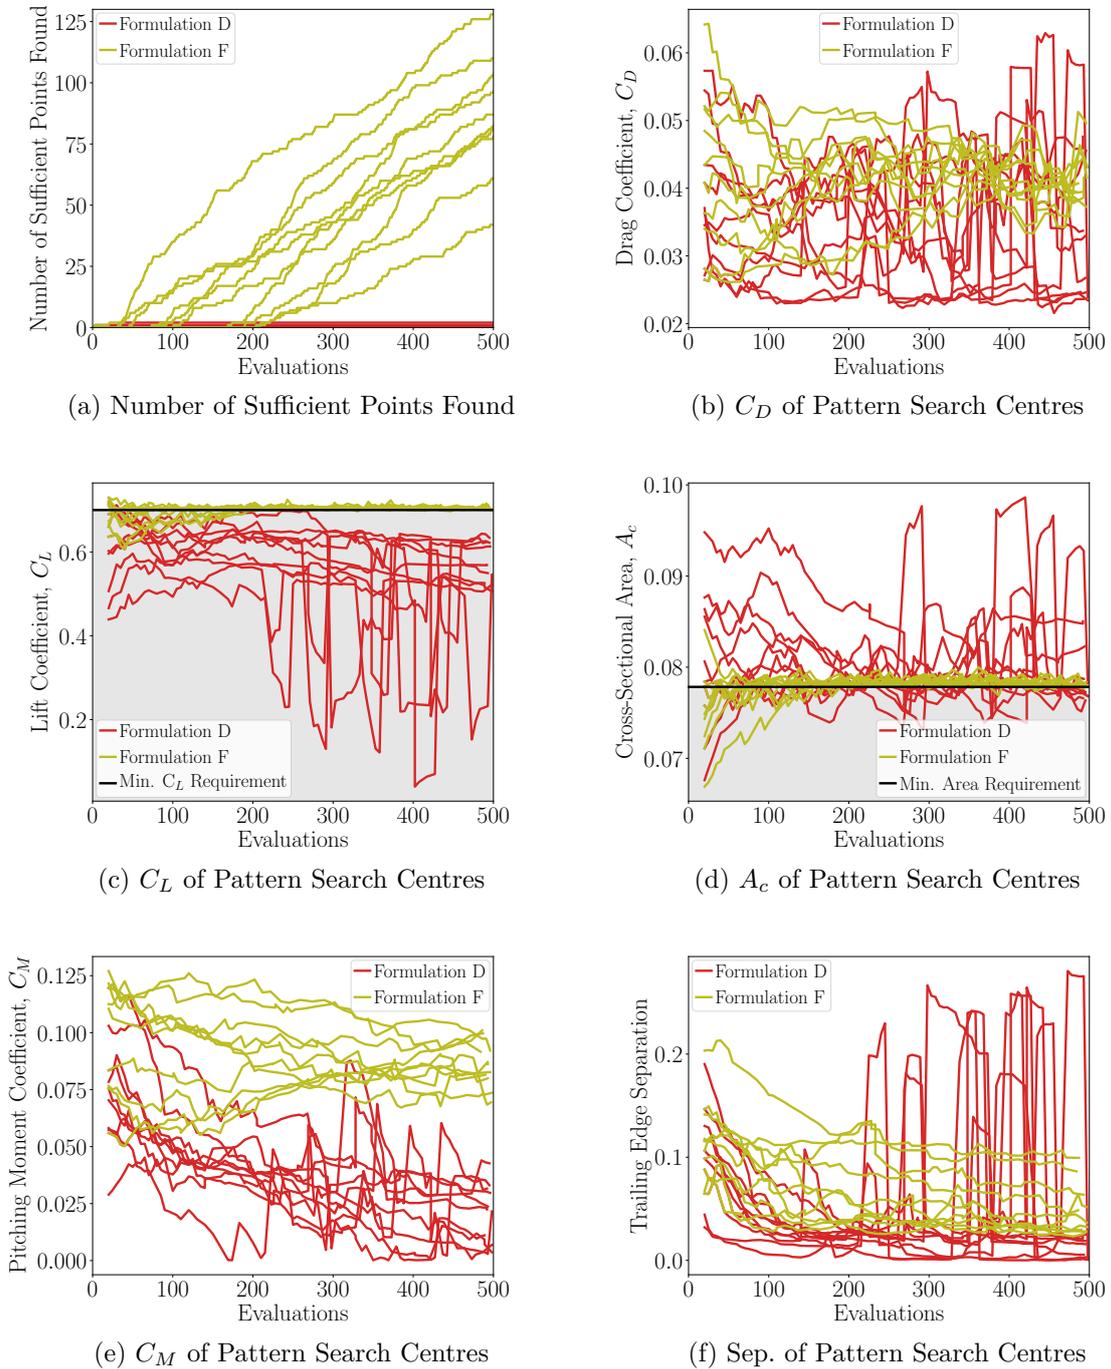


Fig. 5.7 Progression of searches using Formulations D and F in Table 5.1 applied in the single-objective scenario with a minimum C_L requirement of 0.7. Figure 5.7a shows the number of sufficient points found and the remaining plots track performance of the pattern search centres.

A Simplifying Application Framework

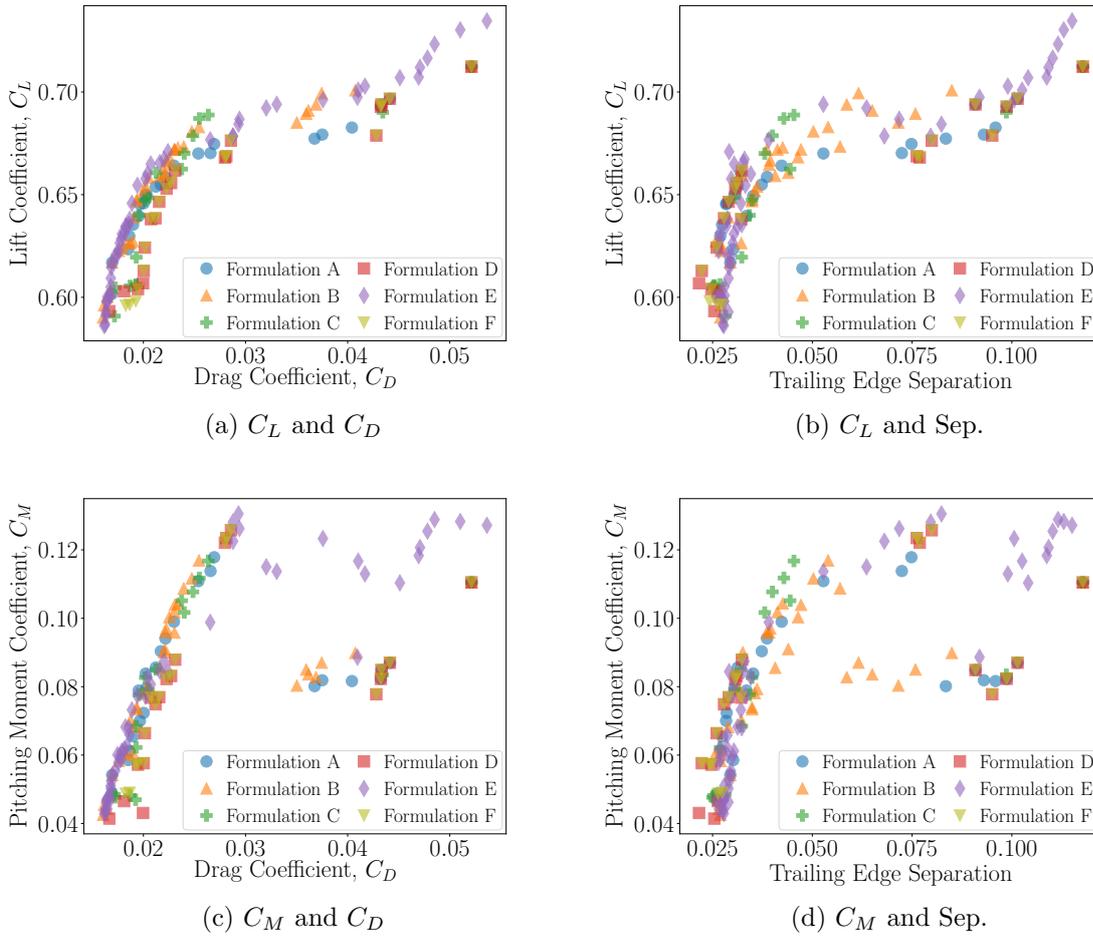


Fig. 5.8 Performance of designs on the C_L - C_D Pareto fronts found using the six formulations in Table 5.1 applied in the trade-off scenario.

generated during the 10 runs that satisfy the minimum A_c constraint. The optimisations conducted using Formulations C, D and F generate Pareto fronts that are mostly, if not completely, dominated in terms of C_L and C_D by those produced using the other three formulations. However, Figure 5.8d shows that Formulations C, D and F have managed to find designs exhibiting good performance in terms of C_M and separation. This is due to the desirable features appearing ahead of the objectives in each of these formulations, meaning improvement in C_M and separation are accepted before C_L and C_D are taken into account.

Searches conducted using Formulations A, B and E converge to similar C_L - C_D Pareto fronts, perhaps suggesting that the permutation of dominance relations within the nested hierarchy is less significant in this trade-off scenario. However, further inspection reveals a weakness in the first two of these formulations.

Figure 5.9a tracks the number of sufficient designs found using each of the different dominance relation permutations and shows that Formulations A and B have suffered a similar fate to that of Formulation D in the single-objective scenario. Each search conducted

5.3 Similarities With Other Constraint Handling Approaches

using these formulations reaches a point beyond which it fails to generate any new sufficient designs due to violation of the minimum A_c constraint visible in Figure 5.9d. Formulations A and B both place the objectives above satisfaction of the constraints in the nested hierarchy of dominance relations shown in Table 5.1. This leads to designs with good performance in terms of C_L and C_D being preferred despite their violation of the minimum A_c requirement. In contrast, searches using Formulation E, which achieve a very similar C_L - C_D Pareto front to those found using Formulations A and B (Figure 5.8a), consistently produce new sufficient designs that are more likely to be of interest to the designer.

As in the single-objective scenario, this problem of failing to find sufficient designs is exacerbated when the constraint becomes more difficult to meet. To simulate this A_c and C_L are swapped in the designation of performance parameters, with maximising A_c and minimising C_D treated as the new objectives subject to a constraint on C_L . This problem set-up, with A_c treated as an objective, is unlikely to be adopted in a practical setting, however for the purposes of this demonstration it allows the more challenging minimum C_L requirement of 0.7 to be applied and the effects on the performance of searches conducted using the formulations in Table 5.1 observed.

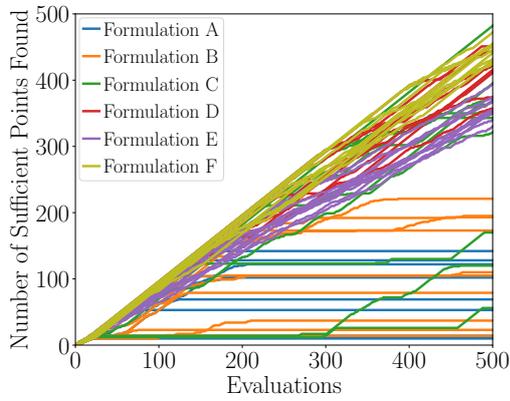
Figure 5.10a, tracking the number of sufficient designs found using the different dominance relation orderings, shows that the optimisations employing Formulations A and B fail to make any meaningful progress when applied to this new trade-off problem. Only searches conducted using Formulations E and F consistently produce aerofoils that satisfy the minimum C_L constraint, with these being the two formulations in Table 5.1 that place satisfaction of the constraints at the top of the nested hierarchy of dominance relations.

These results suggest that, as in the single-objective scenario, the selected nested arrangement has a significant impact on the final designs produced by an optimisation using MDR. In particular, placing satisfaction of the constraints at the top of the hierarchy appears to actively maintain the focus of the optimiser on sufficient designs that are more likely to be of interest to the designer. In this trade-off scenario Formulation E seems to be the better of the two formulations that do this, giving a higher priority to the objectives, increasing the likelihood of an advanced Pareto front in terms of these primary quantities of interest being produced.

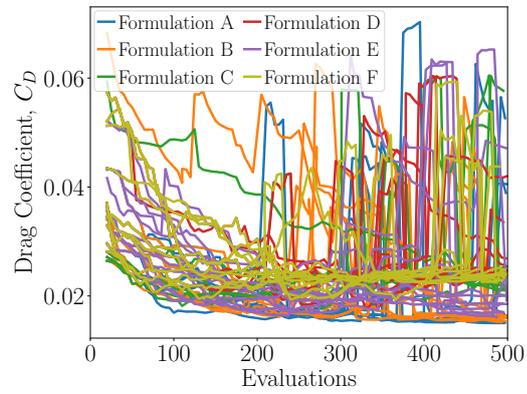
5.3 Similarities With Other Constraint Handling Approaches

Results in the previous section suggest that, irrespective of the number of objectives being used, the dominance relation associated with constraint satisfaction should appear first in the nested hierarchy shown in Equation 5.1. This agrees with the suggestion made by Cook et al. [39] when introducing the MDR formulation that constraints could be handled by setting

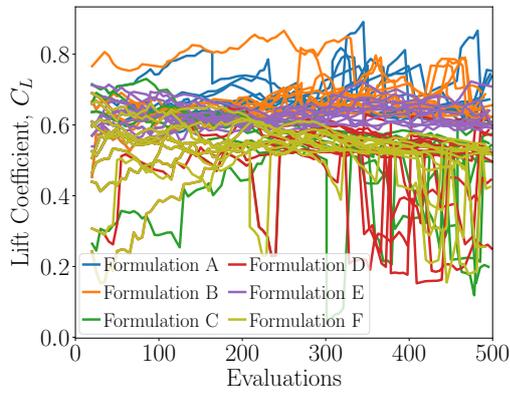
A Simplifying Application Framework



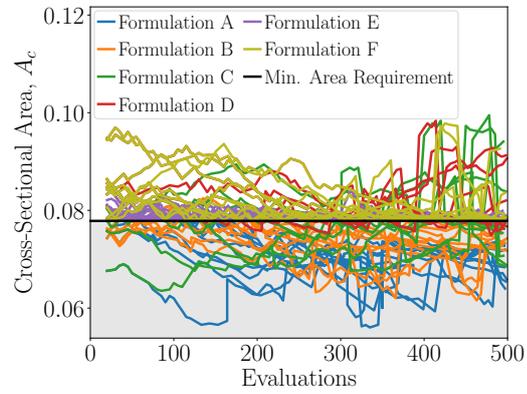
(a) Number of Sufficient Points Found



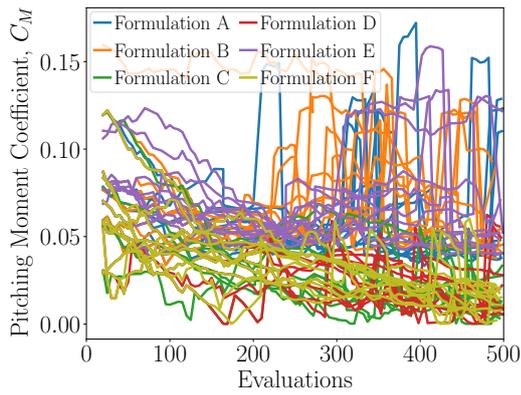
(b) C_D of Pattern Search Centres



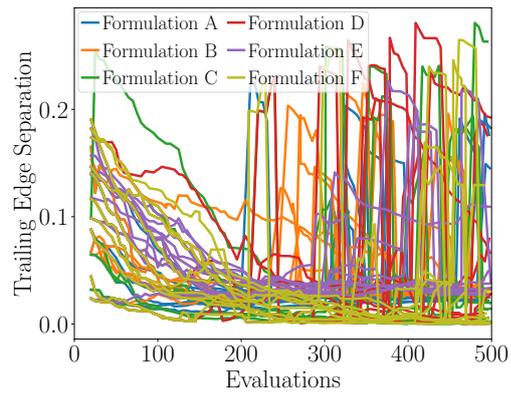
(c) C_L of Pattern Search Centres



(d) A_c of Pattern Search Centres



(e) C_M of Pattern Search Centres



(f) Sep. of Pattern Search Centres

Fig. 5.9 Progression of searches using the six formulations in Table 5.1 applied in the C_L - C_D trade-off scenario. Figure 5.9a shows the number of sufficient points found and the remaining plots track performance of the pattern search centres.

5.3 Similarities With Other Constraint Handling Approaches

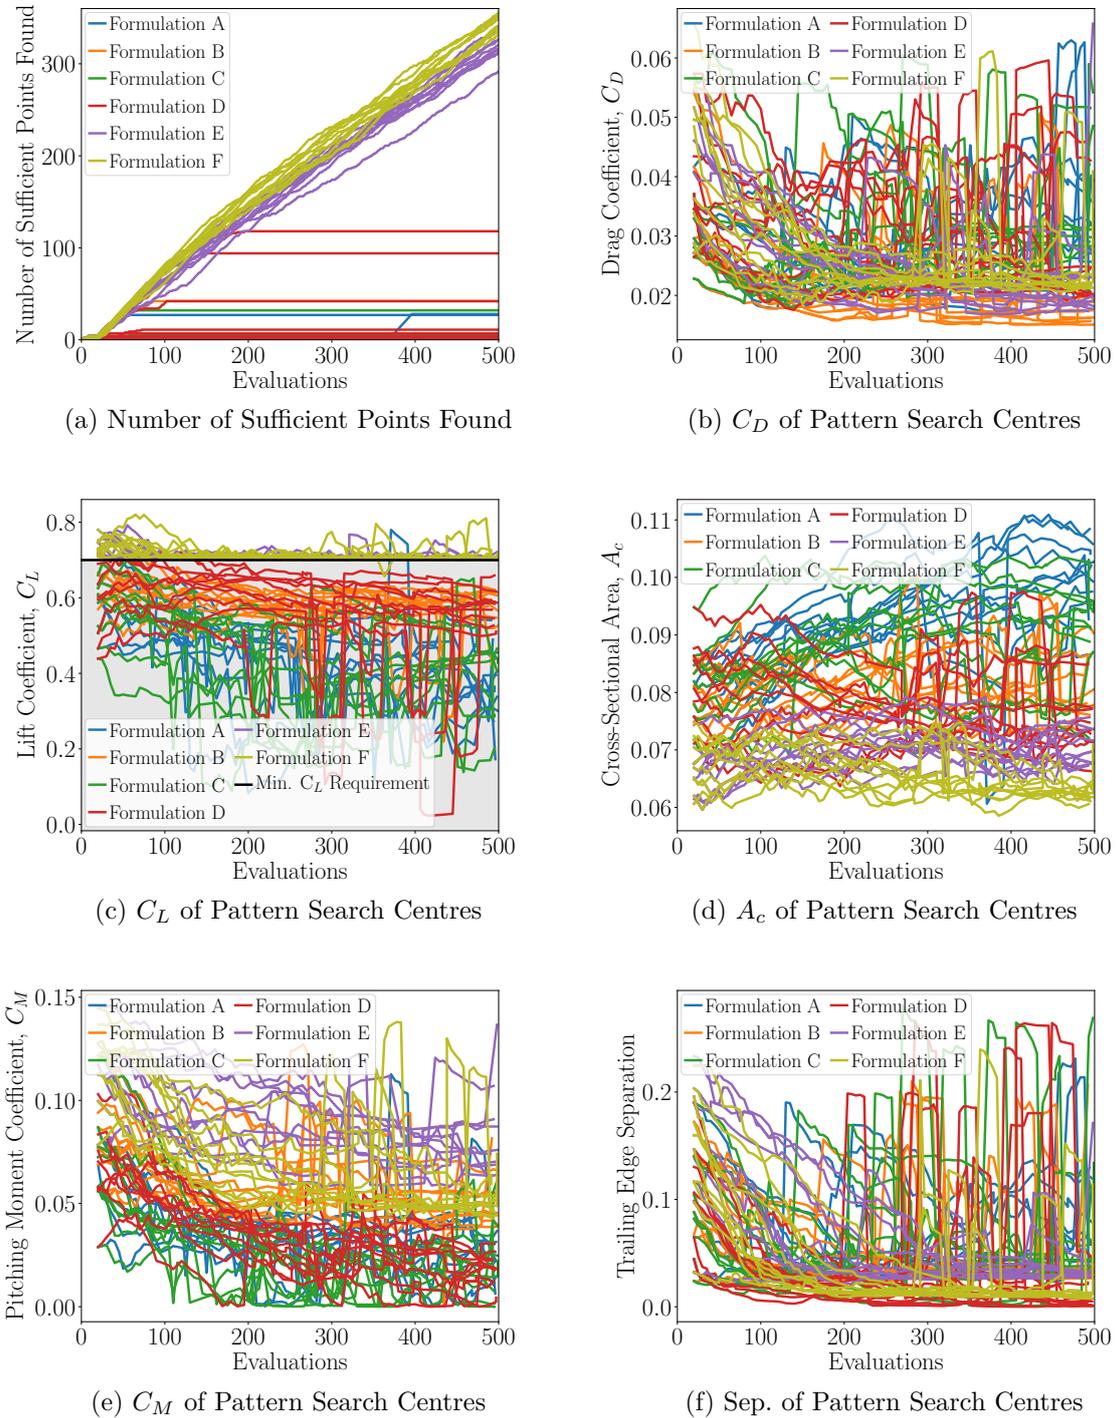


Fig. 5.10 Progression of searches using the six formulations in Table 5.1 applied in the A_c - C_D trade-off scenario. Figure 5.10a shows the number of sufficient points found and the remaining plots track performance of the pattern search centres.

the first relation to be one for which feasible designs are optimal. It is noted that many different constraint handling approaches can be included in this way, such as those presented by Mezura-Montes and Coello [174]. Cook et al. also highlighted similarities with existing approaches that use rankings to handle constraints in evolutionary algorithms, including work by Coello [33], Ho and Shimizu [103] and Ray et al. [198]. The results presented in this chapter provide evidential support for the suggestion made in the original work, and reinforce the importance of prioritising constraint satisfaction.

5.4 Summary

In this chapter a simplifying application framework has been developed that harnesses the main benefits of the MDR formulation through the introduction of one additional performance parameter classification alongside objectives and constraints. Incorporating parameters known as desirable features, this new classification allows sophisticated problem definitions to be produced that reflect what the designer actually wants more accurately than is possible using objectives and constraints alone. This is achieved without any noticeable change in complexity from the perspective of the designer tasked with setting up the optimisation.

An investigation conducted using an aerofoil test case highlights the significance of the arrangement of the dominance relations assessing performance in terms of each of these classifications within the nested hierarchy at the heart of the MDR formulation. The results suggest that some permutations are more likely than others to produce designs that are of interest to the designer. Irrespective of the number of objectives being used, satisfaction of the constraints should be placed at the top of any dominance relation hierarchy to maintain the focus of the optimiser on sufficient designs. When a single objective is employed this should appear as the final element in any nesting to ensure that other parameters are considered by the optimiser, leaving the desirable features to occupy the second spot in the hierarchy. In trade-off scenarios when multiple parameters are designated as objectives these should be placed after the constraints as the second arbiter between designs, followed by desirable features as the final dominance relation.

Whilst it cannot be concluded based on the results of one test case that these represent optimal nesting arrangements for all problems, they can act as a useful starting point when attempting to apply the MDR formulation to more complex problems, such as the preliminary design of axial compressors. The simplifying application framework developed in this chapter is a novel contribution that should enable use of the promising MDR formulation in a wider range of engineering design scenarios.

Chapter 6

Comparison to Existing Methods Using an Aerofoil Test Case

The previous two chapters discussed the implementation of the MDR formulation within a TS algorithm and the development of a simplifying framework to enable application of the resultant method to a wider range of more complex problems. Combining MDR and TS should allow the new algorithm to overcome the limitations of existing methodologies associated with formulation and understanding that were highlighted in Chapter 3. This capability is assessed in this chapter using an aerofoil test case, with results generated by the new MDRTS method compared to those found using existing techniques. Experiments are conducted in single-objective and trade-off scenarios to determine whether the performance of the new approach is consistent, with the relatively cheap computational analysis allowing multiple runs to be carried out to assess the repeatability of any benefits observed. The ability to overcome limitations associated with formulation is assessed through comparisons to alternative optimisations employing more traditional objectives-and-constraints problem formulations, whilst benchmarking against results generated by existing population-based implementations of the MDR formulation tests the efficiency of the search conducted by the new method. Towards the end of the chapter the ability of the MDRTS algorithm to produce interpretable design development information to enhance designer understanding is also demonstrated.

6.1 Experimental Set-Up

The aerofoil test case employed for the performance assessment is the same as that used in the previous chapter to investigate the significance of dominance relation ordering. The conditions, parameterisation, performance parameters and means of analysis are all as described in Section 5.2.1. The different techniques compared in this chapter are run from

the same 10 sets of 20 randomly generated starting geometries, minimising the impact of any stochastic aspects of the approaches on the conclusions drawn. Every run is also given the same computational budget of 500 SU² evaluations, corresponding to a wall clock time of around two days.

In the first part of the chapter the main goal is to assess whether the problems associated with formulation that limit existing optimisation methodologies have been successfully overcome. Designs found using the MDRTS algorithm are therefore compared to those generated by three state-of-the-art search algorithms employing more traditional objectives-and-constraints formulations. The first of these is the TS developed by Jaeggi et al. [110] that formed the basis for the new MDR implementation in Chapter 4. The second is NSGA-II [48], and the third is a PSO developed by Coello et al. [35]. Gradient-based algorithms are not included in the comparison due to their inability to handle multiple objectives and the fact that they are ill-suited to the preliminary design of axial compressors which is the primary problem of interest in this thesis [81, 225]. The settings used for the TS, GA and PSO are the same as for the analytic test case, outlined in Table 4.1, with the exception of the number of candidates randomly selected for analysis in TS, `N_SAMPLE`, which is decreased to five.

As well as assessing the effectiveness of the MDR formulation it is also important to ensure that the underlying TS mechanism is capable of performing efficient exploration of the design space. For this purpose, results generated by the MDRTS algorithm are compared to those found using the two population-based implementations of the MDR approach developed by Cook et al. [39]. These are based on the same GA and PSO used to apply the objectives-and-constraints formulations.

The MDRTS algorithm is not compared to any of the preference-based methods cited in Chapter 3 due to the difficulty associated with articulating preferences, particularly in quantitative terms, at the outset of a design process. Any results would be highly dependent on the exact preference information specified, reducing the relevance of any comparisons made to the MDRTS method that does not require this type of quantitative designer input. If the new approach is capable of producing designs that reflect the interests of the designer without requiring preference information it will have demonstrated a distinct advantage over existing preference-based methods without explicit comparison being necessary.

The priorities method of Fonseca and Fleming [72], another approach discussed in Chapter 3, is also not included in the comparison due to the likelihood of poor results being produced. The equality requirement for lower priority levels to be considered is unlikely to be met for this aerofoil test case as it would be rare for two designs to exhibit exactly the same values of a performance parameter. This means that the desirable features, or any other subsequent comparison levels, would never be considered using the priorities method. It would therefore be unlikely to produce valid or interesting solutions and the available computational budget is better spent conducting repeat runs of more promising techniques. An alternative might

be to set goals for the objective values, increasing the likelihood of lower priority levels being consulted during the optimisation. However, this raises similar challenges to those faced by the preference-based approaches discussed in the previous paragraph, with results being highly dependent on the exact goals set and values being difficult to specify before any optimisation has been carried out.

In all of the experiments conducted in this chapter constraints are applied using a penalty method. The penalty term is calculated by summing individual constraint violations normalised by the limit values themselves. In approaches employing MDR this penalty is used directly as the first dominance relation, whilst for objectives-and-constraints formulations the penalty term is added to the value of any performance parameters being treated as objectives to promote designs that satisfy the constraints.

6.2 Single-Objective Scenario

The same designation of performance parameters employed in the previous chapter when investigating the significance of dominance relation ordering in a single-objective scenario is used here. Minimising C_D is treated as the sole objective with constraints applied to C_L and A_c . The former is required to be no less than 0.6 and the latter at least that of an RAE2822 aerofoil with the same chord. Following Cook et al. [39], C_M and trailing edge separation are treated as desirable features.

Performance in terms of each parameter classification is assessed using a separate dominance relation, with these arranged within the three-tiered nested hierarchy using the results of the previous chapter. Satisfaction of the constraints is selected as the first dominance relation, followed by Pareto dominance of the desirable features and ordering of the sole objective respectively. The resultant single-objective MDR problem definition is summarised in Table 6.1 alongside two comparative formulations that are the subject of Sections 6.2.1 and 6.2.3.

6.2.1 Comparison to Traditional Formulations

When applying a more traditional objectives-and-constraints formulation to this problem C_M and trailing edge separation need to be treated as either objectives or constraints. Specifying suitable limit values for constraints applied to these quantities would be challenging, particularly given the non-standard measure for separation being employed (see Section 5.2.1). C_M and separation are therefore treated as additional objectives alongside C_D , resulting in a multi-objective problem formulation consisting of three objectives and the same two constraints as in the MDR approach applied using the penalty method. This multi-objective formulation, outlined in Table 6.1, is implemented using the TS, GA and PSO discussed earlier in this chapter.

Comparison to Existing Methods Using an Aerofoil Test Case

Table 6.1 Formulations applied to the aerofoil test case in the single-objective scenario.

	MDR	Multi-Objective	Single-Objective (Emergent)
Objectives	C_D	$C_D, C_M, \text{Sep.}$	C_D
Desirable Features	$C_M, \text{Sep.}$	-	-
Constraints	$A_c \geq 0.07784 \text{ m}^2$ $C_L \geq 0.6$	$A_c \geq 0.07784 \text{ m}^2$ $C_L \geq 0.6$	$A_c \geq 0.07784 \text{ m}^2$ $C_L \geq 0.6$ $C_M \leq 0.05$ $\text{Sep.} \leq 0.03$

Figure 6.1 shows the performance of designs produced by the MDRTS and traditional multi-objective approaches that are non-dominated in terms of C_D , C_M and separation, considering all aerofoils generated during the 10 runs that satisfy the C_L and A_c constraints. The new MDRTS algorithm produces better designs in terms of the objective and desirable features than the multi-objective GA and PSO. Aerofoils generated using the MDR formulation exhibit lower values of C_M for any given value of the objective, C_D , with a number of these designs also producing small amounts of trailing edge separation. The results generated by the multi-objective TS are closer, suggesting that in this single-objective scenario the underlying search algorithm has a larger impact on the success of an approach than the problem formulation employed. This is perhaps not surprising given the relatively small number of performance parameters being considered. The multi-objective methods are only required to resolve a three-dimensional trade-off and are able to do that relatively effectively. Despite this, the MDR formulation produces a large number of designs that are likely to be of interest to the designer, demonstrating performance that is at least as good as the traditional objectives-and-constraints approaches.

These findings are corroborated using a similar reference point approach to that employed in the previous chapter. An idealised reference point is defined as having the best performance in terms of C_D , C_M and separation found by any of the runs applied to this problem. The minimum Euclidean distance to this point in a space containing normalised versions of the objective and desirable features is tracked to determine the speed at which good designs are being produced, and the number of final designs exhibiting performance within a given Euclidean distance is tallied to show how many interesting designs are generated. In practice the reciprocal of the minimum distance is plotted to give an overall performance measure where higher values indicate better results.

Figure 6.2 tracks the best values of the overall performance measure found by each of the approaches, with all 10 runs plotted as faint lines and the average shown in bold. The

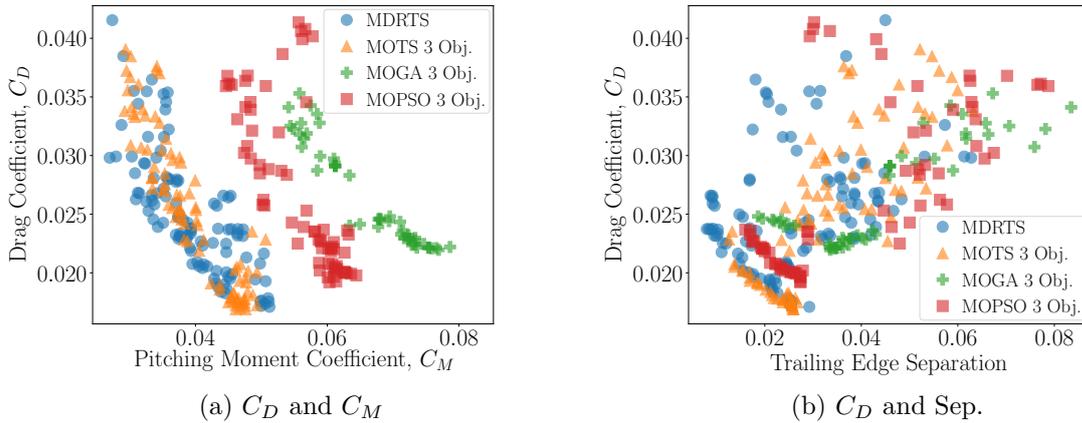


Fig. 6.1 Performance of the best designs found during 10 runs of the MDRTS algorithm and methods employing the multi-objective formulation in Table 6.1 applied to the aerofoil test case in the single-objective scenario.

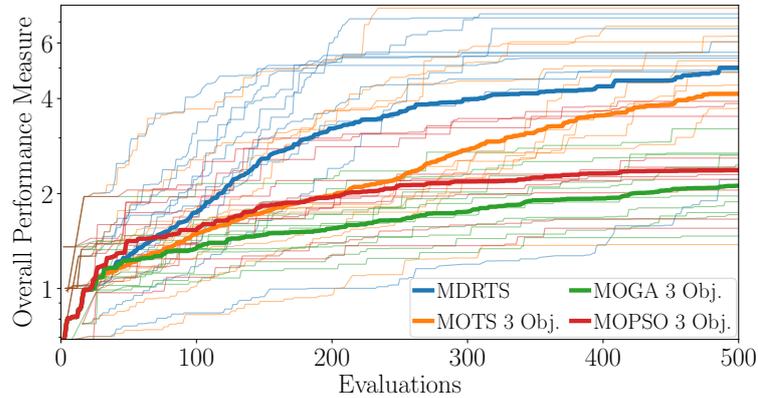


Fig. 6.2 Progression of the best overall performance measure found by the MDRTS algorithm and methods employing the multi-objective formulation in Table 6.1 applied to the aerofoil test case in the single-objective scenario. All runs are plotted as faint lines with the mean shown in bold.

two TS algorithms outperform the population-based methods, producing final designs with values of the overall performance measure that are over twice as large. On average the MDR formulation finds these aerofoils exhibiting high levels of performance using fewer calls to the SU^2 analysis routine. Results equivalent to those found using the multi-objective population-based approaches are produced after an average of 150 evaluations, a computational saving of around 70%. The MDRTS algorithm is also more efficient than the TS employing the multi-objective formulation, generating equivalent performance with an average computational saving of 32%. These results show that the sophisticated treatment of performance parameters in the MDR formulation is focussing the search more effectively, leading to efficient use of computational resources and, ultimately, a better outcome for the designer.

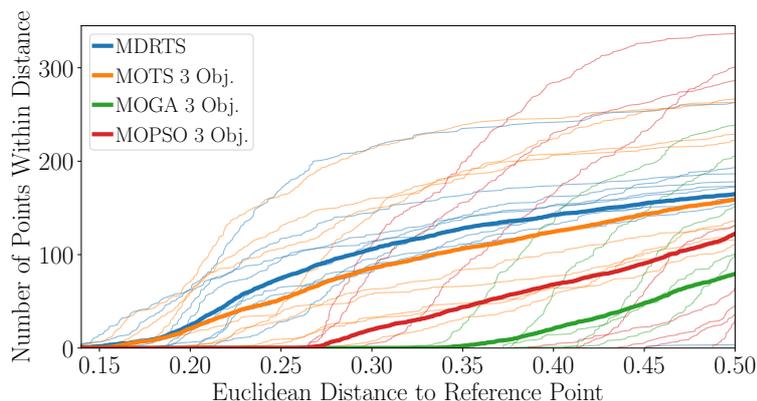


Fig. 6.3 Number of designs produced by the MDRTS algorithm and methods employing the multi-objective formulation in Table 6.1 applied to the aerofoil test case in the single-objective scenario that exhibit performance within a given Euclidean distance of the reference point. All runs are plotted as faint lines with the mean shown in bold.

The plot of the number of designs exhibiting performance within a given Euclidean distance of the reference point in Figure 6.3 also shows the two TS approaches outperforming the GA and the PSO, producing around 80 more designs that exhibit performance within a distance of 0.3 of the idealised reference on average. Again, MDRTS is seen to achieve slightly better results than the TS employing the more traditional multi-objective formulation. Of particular note on this plot is the spread of the individual runs, plotted as faint lines. For the approaches employing the multi-objective formulation these lines are scattered, suggesting that the improvement achieved is highly dependent on the set of starting designs. In contrast, the lines corresponding to the MDRTS method are closer together, with the exception of a single outlier in the bottom right corner of the figure. This indicates that the MDR formulation is able to produce designs exhibiting high levels of performance more reliably and consistently than traditional formulations.

The results presented in this section demonstrate that the MDR approach has successfully overcome the problems associated with formulation that limit traditional objectives-and-constraints techniques. Over the course of 10 runs the MDRTS algorithm generates aerofoils that are at least as good as those produced using existing methods, reaching these designs more quickly and with greater consistency.

6.2.2 Comparison to Alternative MDR Implementations

Having compared the performance of the new MDRTS method to techniques employing traditional objectives-and-constraints problem formulations the next task is to compare it to the alternative GA and PSO implementations of the MDR approach developed by Cook et al. [39]. These algorithms were employed during the analytic verification study conducted in Chapter 4 where the TS implementation was shown to produce computational savings of 18%

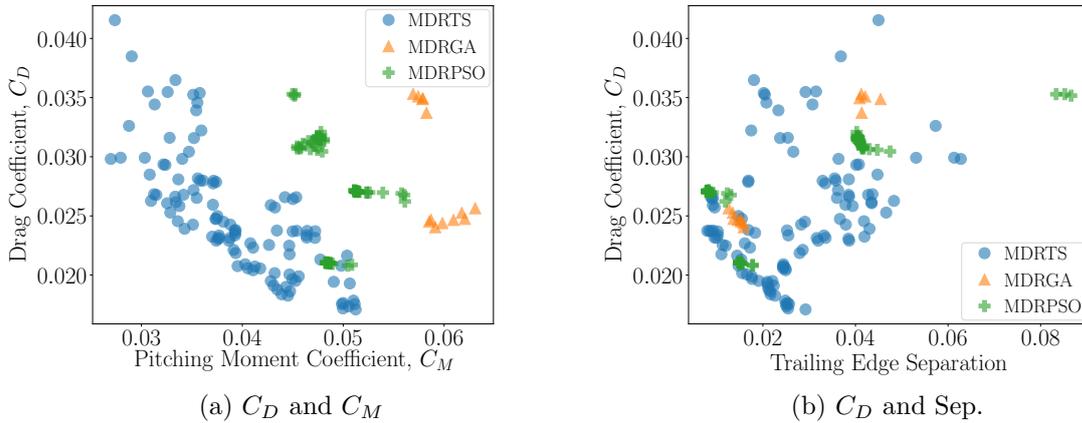


Fig. 6.4 Performance of the best designs found during 10 runs of the different MDR implementations applied to the aerofoil test case in the single-objective scenario.

and 35% compared to the PSO and GA approaches respectively. The results in this section provide a further comparison of the different search mechanisms using a more challenging test case.

Figure 6.4 shows the performance of aerofoils produced by each of the three algorithms that are non-dominated in terms of C_D , C_M and separation, taking into account all designs generated during the 10 runs that satisfy the C_L and A_c constraints. The new TS implementation produces a larger number of non-dominated designs than either the GA or the PSO, with these aerofoils also exhibiting lower values of the objective and the two desirable features.

Figure 6.5 tracks the best values of the overall performance measure found by the different methods, with this quantity calculated in the same way as the previous section using an idealised reference point given the best performance in terms of C_D , C_M and separation generated by any of the runs applied to this problem. The TS implementation again outperforms the GA and PSO approaches, producing final aerofoils with values of the overall performance measure that are over twice those found using the population-based methods. These good designs are also generated more quickly, with results equivalent to those found using the GA produced with an average computational saving of 71%, and designs with performance equal to the best aerofoils generated by the PSO found using, on average, a 67% reduction in the number of calls to the SU^2 analysis code.

The TS implementation achieves the greatest gains in the latter stages of the search. During the first 50 evaluations the three methods follow similar average trajectories in their search for improved designs. At this point the GA stalls, with the TS and PSO continuing until around 100 evaluations at which point the progress of the latter also begins to diminish. This behaviour suggests that it is in the exploitation phase that TS is outperforming the other two methods. In the early stages the optimiser focusses on exploring the design space,

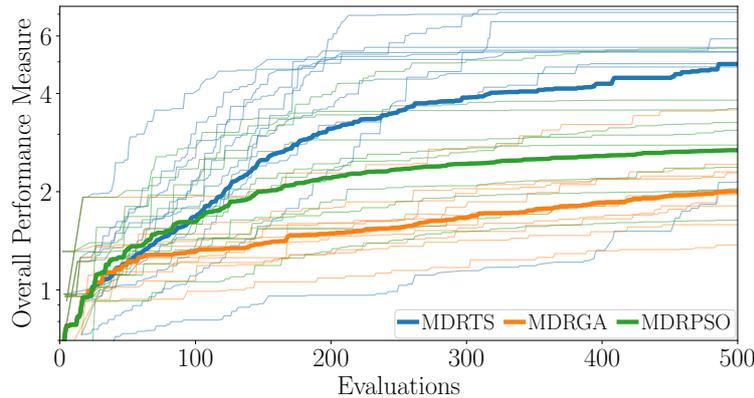


Fig. 6.5 Progression of the best overall performance measure found by the different MDR implementations applied to the aerofoil test case in the single-objective scenario. All runs are plotted as faint lines with the mean shown in bold.

attempting to locate promising regions, and at this point the population-based methods perform well. However, once the optimiser has uncovered a region of design space containing promising aerofoils the goal is to exploit this potential by generating further improvements. As discussed in Chapter 4, TS makes small incremental changes to the design vector and is therefore more adept at exploitation than the population-based GA and PSO methods that make larger changes in the hope of conducting a more global search. This limitation of GAs in particular has been noted in the literature, leading to the development of numerous hybrid methodologies where the explorative capabilities of these population-based methods are blended with the exploiting power of local approaches [133, 158, 180]. The results in Figure 6.5 suggest that the MDR formulation exposes this flaw in population-based optimisation methodologies and is possibly more effective when combined with the approach based on pattern search. The sophisticated problem definition describes the region of design space that is of interest to the designer more explicitly than when using objectives and constraints, reducing the need for exploration as the formulation itself focusses the optimiser towards the desired area. The importance of exploitation is therefore magnified, favouring algorithms such as TS that are able to generate incremental improvements within this smaller region of design space.

In Figure 6.6 the number of designs exhibiting performance within a given Euclidean distance of the idealised reference point is plotted, with each of the 10 runs shown faintly and the average highlighted in bold. The TS produces a larger number of designs exhibiting high levels of performance than either the GA or the PSO, with the lines corresponding to the new method also being closer together indicating more consistent performance across the 10 runs. In contrast, the results generated by the two population-based methods are somewhat erratic, with a handful of runs producing a large number of good designs but others leading to very few aerofoils that are likely to be of interest to the designer.

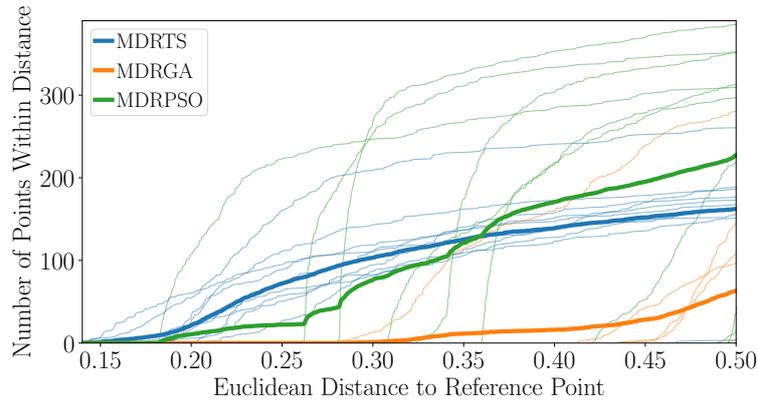


Fig. 6.6 Number of designs produced by the different MDR implementations applied to the aerofoil test case in the single-objective scenario that exhibit performance within a given Euclidean distance of the reference point. All runs are plotted as faint lines with the mean shown in bold.

The results presented in this section further demonstrate the capabilities of the new MDRTS algorithm. The improved problem definition focusses the optimiser towards a specific region of the design space, allowing the exploitive abilities of TS to consistently produce designs exhibiting high levels of performance. In contrast, the population-based implementations, more suited to exploration of the design space, stall in their search for good designs and show large variation in the quality of aerofoils that are produced.

6.2.3 Emergent Constraints

Figure 6.7 tracks the development of the desirable features, C_M and separation, during one of the MDRTS runs applied in the single-objective scenario. The performance of all of the designs generated is shown as a faint line with that of the points selected as centres for the pattern searches highlighted in bold. With the exception of the spike at around 370 evaluations, caused by the diversification heuristic used to enhance the searching capabilities of the TS algorithm (see Section 4.1.2), these plots show both parameters converging towards limit values as the search progresses.

In Figure 6.8 the progression of the desirable features during all 10 runs of the MDRTS algorithm is plotted, with just the performance of the pattern search centres shown and the mean highlighted in bold. Whilst the convergent behaviour is not always as clear as in Figure 6.8 some levelling out of these parameters is observed in most cases, with consistent convergence seen in the early stages of each optimisation. Beyond this point large changes in performance occur in a number of runs due to the intensification and diversification heuristics used by the algorithm to enhance the search for good designs.

In a traditional aerofoil optimisation [4, 97, 129] secondary parameters such as C_M and separation would not normally be treated as additional objectives, as was done in the multi-

Comparison to Existing Methods Using an Aerofoil Test Case

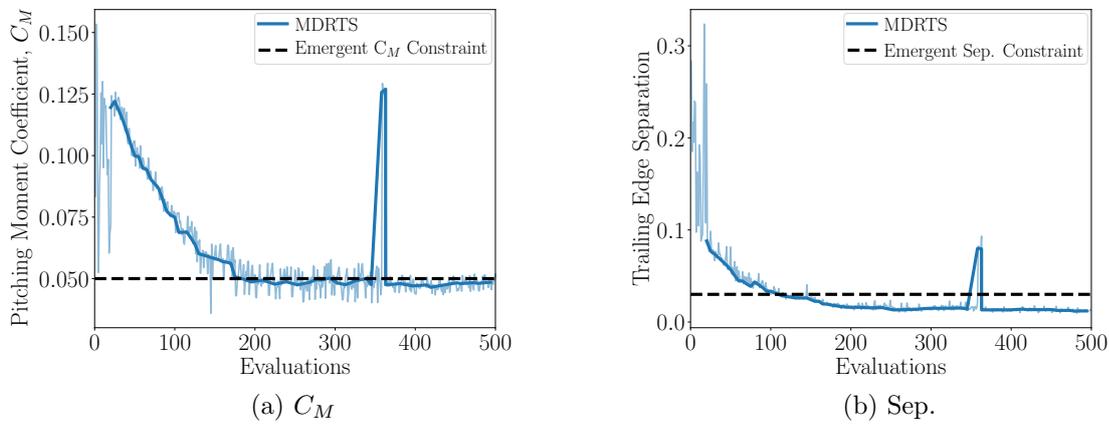


Fig. 6.7 Progression of the desirable features during one run of the MDRTS algorithm applied to the aerofoil test case in the single-objective scenario. Performance of all designs is plotted as a faint line, with that of the pattern search centres highlighted in bold.

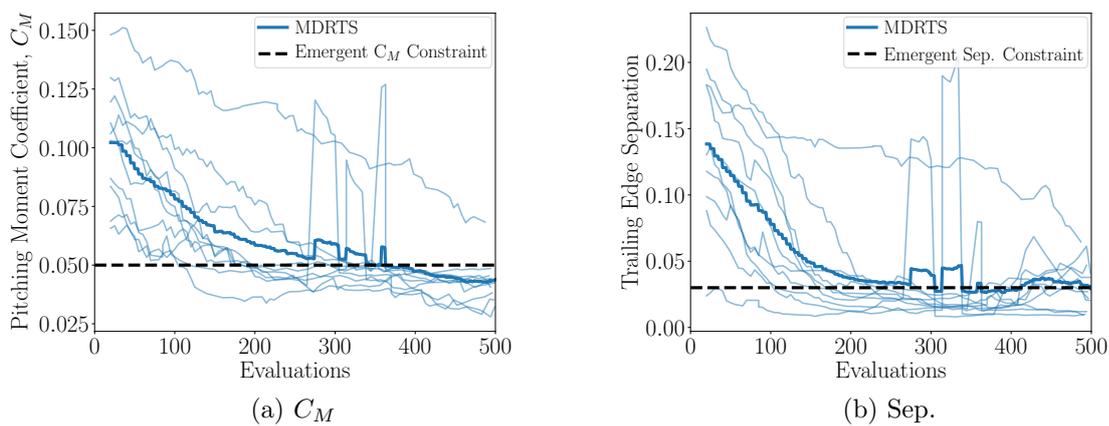


Fig. 6.8 Progression of the desirable features during all 10 runs of the MDRTS algorithm applied to the aerofoil test case in the single-objective scenario. Performance of the pattern search centres for each run is plotted as a faint line with the mean shown in bold.

objective comparison in Section 6.2.1. Instead, these quantities would be constrained to ensure that the values exhibited by final designs are “good enough”. As discussed earlier in this chapter, selecting appropriate limits for these constraints at the start of a design process can be challenging, especially when unfamiliar quantitative measures are involved such as that used to assess trailing edge separation in this work. Designating parameters as desirable features in MDRTS removes the need to specify the limit values and Figures 6.7 and 6.8 suggest that it may also allow them to emerge as outputs of the optimisation. Rather than being a potentially ill-informed input that could lead the optimiser astray, these limiting values have instead become useful by-products of the preliminary optimisation process. This behaviour, if shown to exist in other problems, could be beneficial when applying optimisation techniques in novel physical domains, such as the design of electric and hybrid aircraft, where suitable values for constraint limits are unknown. The emergent limit values could also be used to inform more detailed optimisation studies, possibly conducted using gradient-based methods, at a later stage in the design process.

Constraint limits informed by this emergent behaviour also provide a means by which to compare the new MDRTS approach to a more traditional formulation where constraints are applied to C_M and separation. Using Figure 6.8 suitable constraint limits can be defined for both quantities and applied in single-objective formulations run using the TS, GA and PSO methods already employed for comparative purposes. The selected limits, highlighted in Figures 6.7 and 6.8, are 0.05 for C_M and 0.03 for trailing edge separation, with the resultant single-objective formulation outlined in Table 6.1. The three algorithms employing this constrained problem definition are run from the same 10 sets of 20 randomly generated starting geometries and are given an equivalent computational budget of 500 SU^2 evaluations.

The results of these runs are shown in Figure 6.9, where the performance of designs that are non-dominated in terms of C_D , C_M and separation is plotted taking into account all of the aerofoils generated that satisfy the C_L and A_c constraints. Designs produced by the different methods overlap, with aerofoils generated that exhibit similar low values of the objective and the two desirable features. As was the case in Figures 6.1 and 6.4, the GA has not made as much progress as the other approaches suggesting that it may not be well suited to this aerofoil problem. Both methods using TS generate designs with lower C_D values than those produced by the PSO, with the single-objective formulation actually leading to slightly better aerofoils than the MDR approach in this regard. However, the new MDRTS algorithm produces a larger number of non-dominated designs that are likely to be of interest to the designer and generates aerofoils with lower values of the desirable features, particularly C_M , than are found using the single-objective formulation.

Figure 6.10, tracking the best values of the overall performance measure, shows that these good designs have been reached most rapidly using the single-objective TS approach. The PSO also makes fast progress towards the beginning of the optimisation but again stalls

Comparison to Existing Methods Using an Aerofoil Test Case

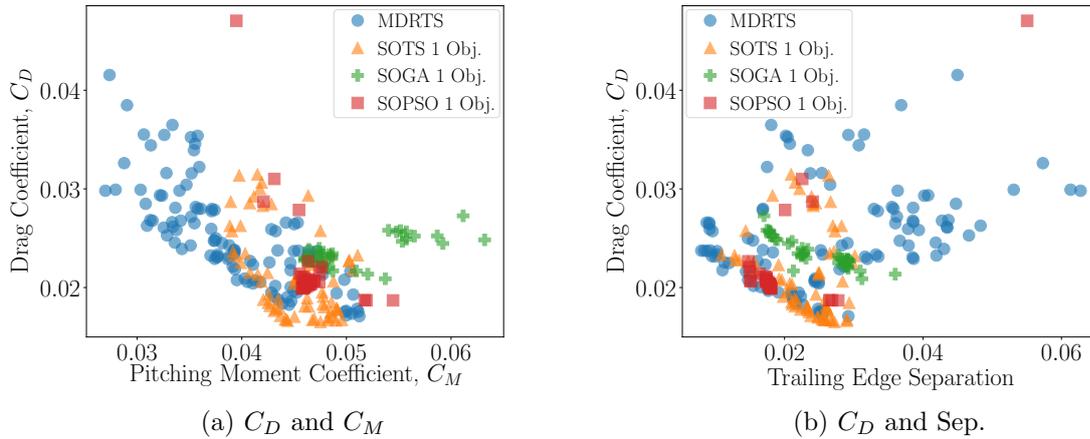


Fig. 6.9 Performance of the best designs found during 10 runs of the MDRTS algorithm and methods employing the constrained single-objective formulation in Table 6.1 applied to the aerofoil test case in the single-objective scenario.

as the search enters the exploitation phase. The new MDRTS algorithm achieves similar convergence to the single-objective TS method in the early stages and continues to improve upon the good designs found as the optimisation progresses.

The slightly slower initial convergence of the new method compared to the single-objective approaches is most likely due to the greater freedom the optimiser has when using the MDR formulation. The single-objective methods constrain C_M and separation, forcing the optimiser rapidly towards the region of design space specified by these limits. The use of MDR, in contrast, results in a more relaxed approach, informing the optimiser of the desire to minimise these quantities where possible without applying strict limits. This leads to slightly slower convergence in the early stages, but ultimately produces a large number of aerofoils that are likely to be of interest to the designer without the need to specify constraint limits at the outset of the optimisation process.

In Figure 6.11 the number of designs exhibiting performance within a given Euclidean distance of the reference are tallied. All of the single-objective formulations are capable of producing designs that exhibit performance close to that of the ideal point. The MDRTS algorithm generates designs that are at least as good as those found using these single-objective methods without the need to specify constraint limits for C_M and separation. Instead, these threshold values become outputs of the optimisation process, providing further utility to the designer.

Figure 6.12, which tracks the number of points found by each method that satisfy the C_L and A_c constraints, reveals another benefit of the new approach. On average the three single-objective methods produce around 100 fewer sufficient designs than the MDRTS algorithm, with this also evidenced in Figure 6.9 by the smaller number of non-dominated designs generated using these techniques. This is probably due to the searches conducted using the

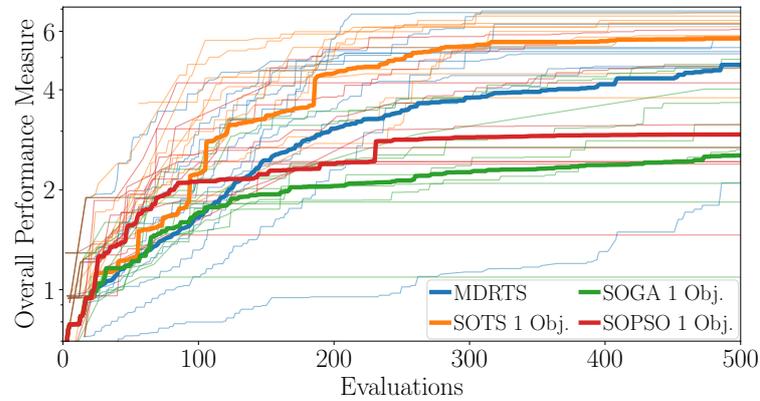


Fig. 6.10 Progression of the best overall performance measure found by the MDRTS algorithm and methods employing the constrained single-objective formulation in Table 6.1 applied to the aerofoil test case in the single-objective scenario. All runs are plotted as faint lines with the mean shown in bold.

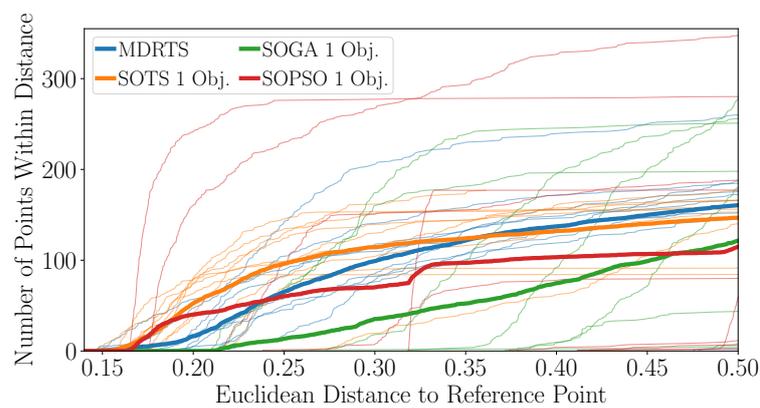


Fig. 6.11 Number of designs produced by the MDRTS algorithm and methods employing the constrained single-objective formulation in Table 6.1 applied to the aerofoil test case in the single-objective scenario that exhibit performance within a given Euclidean distance of the reference point. All runs are plotted as faint lines with the mean shown in bold.

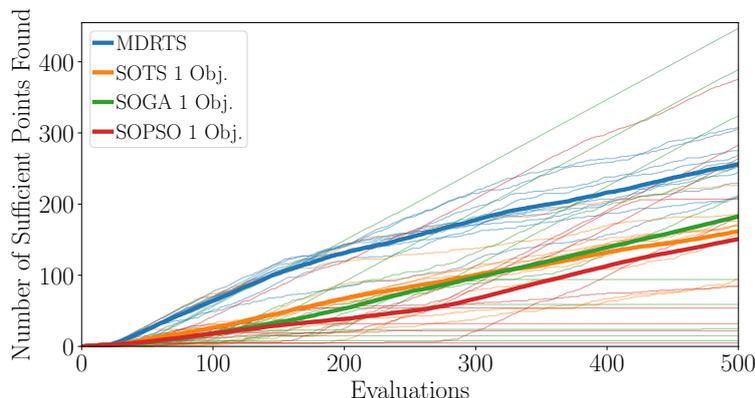


Fig. 6.12 Number of sufficient points found by the MDRTS algorithm and methods employing the constrained single-objective formulation in Table 6.1 applied to the aerofoil test case in the single-objective scenario.

constrained single-objective formulation being unable to distinguish between different types of constraint. Whilst there are strict design requirements for the minimum C_L and A_c of the final aerofoil in this scenario, no such limits exist for C_M and separation. The imposed constraints are user-defined, with the true goal being for the final designs to exhibit values of these parameters that are “good enough”. Bell et al. [15] distinguish between these two types of constraint, referring to the former as hard and the latter as soft and “process-intrinsic”. A human conducting the design process would be likely to prioritise satisfaction of the hard constraints, making progress towards any “process-intrinsic” limits only when this does not impede upon the design requirements. An automated optimisation scheme using a penalty method to apply constraints is unable to make this distinction, leading to the C_L and A_c requirements being violated in search of designs that satisfy the user-defined C_M and separation limits. The formulation employed in the new MDRTS method allows these parameters to be handled more accurately, enabling the optimiser to prioritise real constraints that need to be satisfied whilst still including secondary parameters that the designer wants to be “good enough” in a computationally efficient manner.

The results in this section demonstrate that the new MDRTS algorithm is able to produce designs that are at least as good as those generated using constrained single-objective formulations without the need to specify constraint limits that can be difficult to define before any optimisation has been carried out. There is evidence that, for this aerofoil problem at least, suitable values for these limits become outputs of the optimisation, providing learning opportunities for designers as well as potentially informing more detailed optimisation studies conducted at a later stage in the design process.

Table 6.2 Formulations applied to the aerofoil test case in the trade-off scenario.

	MDR	Four-Objective	Two-Objective (Emergent)
Objectives	C_L, C_D	$C_L, C_D, C_M, \text{Sep.}$	C_L, C_D
Desirable Features	$C_M, \text{Sep.}$	-	-
Constraints	$A_c \geq 0.07784 \text{ m}^2$	$A_c \geq 0.07784 \text{ m}^2$	$A_c \geq 0.07784 \text{ m}^2$ $C_M \leq 0.063$ $\text{Sep.} \leq 0.04$

6.3 Trade-Off Scenario

Results in the previous section show that the new MDRTS algorithm performs well when applied to the aerofoil test case in a single-objective scenario. This section seeks to determine whether good performance persists when the designer is instead interested in investigating a trade-off between numerous quantities of interest.

The experiment conducted in the previous chapter to determine the significance of nesting order found that the arrangement of performance parameter classifications most likely to produce good designs differed in this scenario. Constraint satisfaction still occupies the first position in the hierarchy, but this is now followed by the objectives and then the desirable features (these appear in the opposite order when using a single objective). Given this alternative nesting arrangement, and the inherent differences associated with investigating a trade-off between quantities rather than minimising a single performance parameter, the comparisons made in the following sections play an important role in assessing the capabilities of the new approach and determining whether the problems associated with formulation that limit current methodologies have been successfully overcome.

As in the previous chapter, a trade-off is sought between maximising C_L and minimising C_D , with these two quantities designated as objectives. C_M and trailing edge separation are again treated as desirable features, with the same minimum A_c requirement used in the single-objective comparisons applied as the sole constraint. The resultant MDR formulation is summarised in Table 6.2 alongside alternative methods that are the subject of Sections 6.3.1 and 6.3.3.

6.3.1 Comparison to Traditional Formulations

As in the single-objective scenario, the first comparison is to approaches utilising more traditional objectives-and-constraints problem formulations. Due to the difficulty of specifying suitable limit values, C_M and separation are again treated as additional objectives alongside

Comparison to Existing Methods Using an Aerofoil Test Case

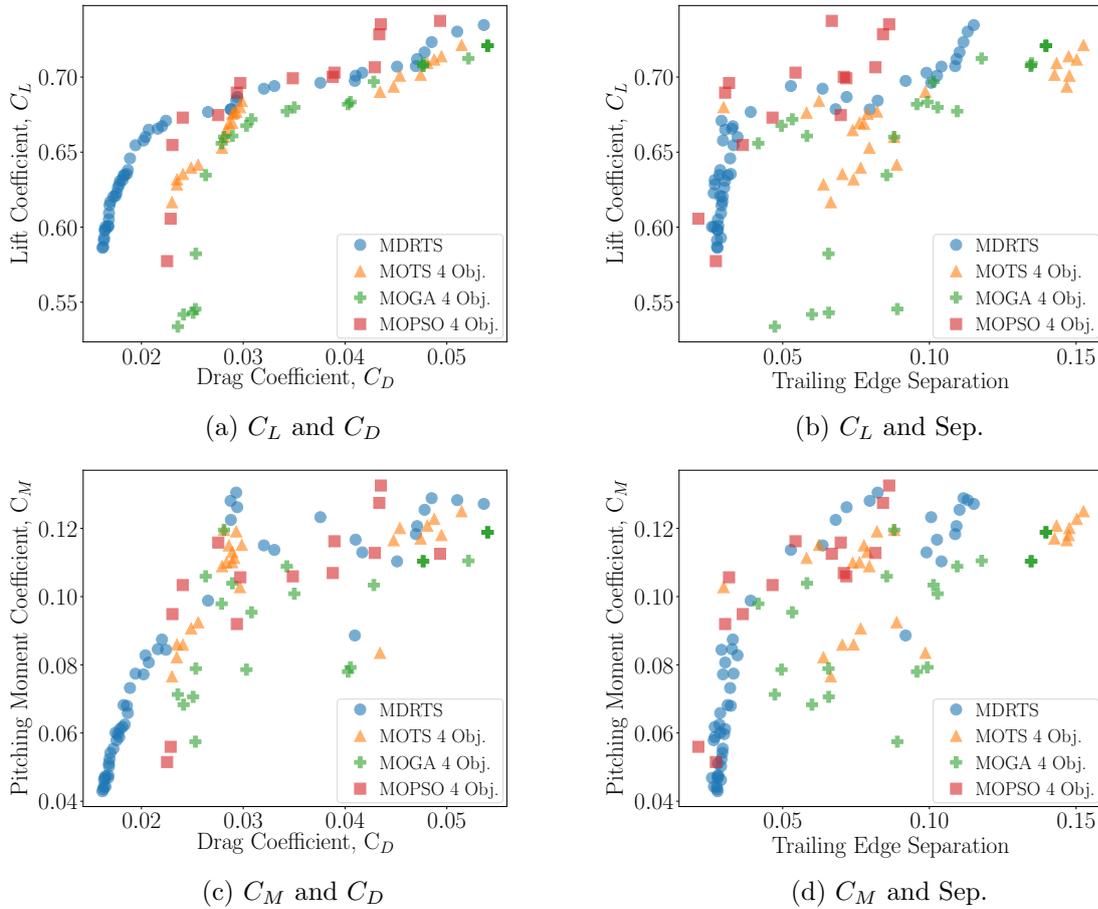


Fig. 6.13 Performance of the best designs found during 10 runs of the MDRTS algorithm and methods employing the four-objective formulation in Table 6.2 applied to the aerofoil test case in the trade-off scenario.

C_L and C_D , rather than using constraints. The resultant multi-objective problem formulation, outlined in Table 6.2, consists of four objectives and the sole constraint on A_c . This formulation is applied in the multi-objective TS, GA and PSO described earlier in the chapter, with these and the new MDRTS method run from 10 different starting populations of 20 randomly generated aerofoils with a computational budget of 500 SU^2 evaluations.

Figure 6.13 shows the performance of designs on the C_L - C_D Pareto front found by the different approaches taking into account all points generated during the 10 runs that satisfy the minimum A_c constraint. The new method, utilising the MDR formulation, produces a more advanced Pareto front than those generated using the multi-objective problem definition, particularly towards the low- C_D end. Designs on this front also exhibit the best performance in terms of the two desirable features, C_M and separation.

The difference between results produced by the new approach and the traditional methods is clearer here than in the single-objective scenario. This is possibly due to an extra

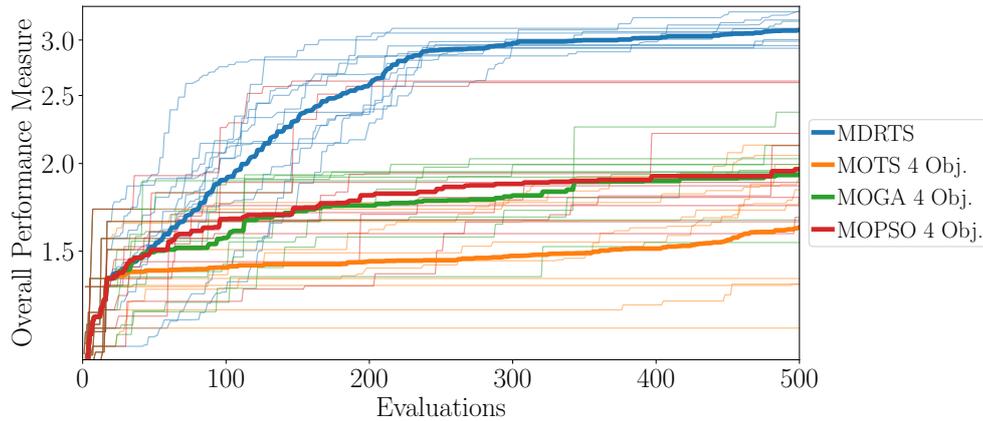


Fig. 6.14 Progression of the best overall performance measure found by the MDRTS algorithm and methods employing the four-objective formulation in Table 6.2 applied to the aerofoil test case in the trade-off scenario. All runs are plotted as faint lines with the mean shown in bold.

performance parameter being considered, resulting in the multi-objective approaches seeking to resolve a four-dimensional trade-off. The increased number of objectives negatively impacts the efficiency of the traditional methods and leads to designs exhibiting relatively poor performance being produced using the limited computational budget. The new MDRTS approach handles the extra parameter more effectively, producing designs that exhibit good performance in terms of all four quantities of interest.

Figures 6.14 and 6.15 demonstrate that this superiority is experienced consistently across the 10 runs. These plots use the same reference point approach described earlier in this chapter, with the hypothetical ideal design given the best performance in terms of C_L , C_D , C_M and separation found during all of the searches applied to this problem. The Euclidean distance of each design from this ideal point is calculated in a space containing normalised versions of these four performance parameters and is used to assess the searches conducted by the different optimisation techniques.

Figure 6.14 tracks the best values of the overall performance measure, with higher values indicating points with performance closer to that of the idealised reference point. The new approach consistently generates designs with higher values of this measure than those found using the multi-objective formulations. These good designs are also produced more quickly, with performance equivalent to the best found by the GA and PSO generated with an average computational saving of just under 80%. Values of the overall performance measure equal to the highest produced by the multi-objective TS approach are found by the MDRTS algorithm using 88% fewer calls to the SU^2 analysis code on average. These results support the conclusion drawn from Figure 6.13 that the new method is more successful than the traditional approaches at generating designs exhibiting high levels of performance in terms of all four quantities of interest.

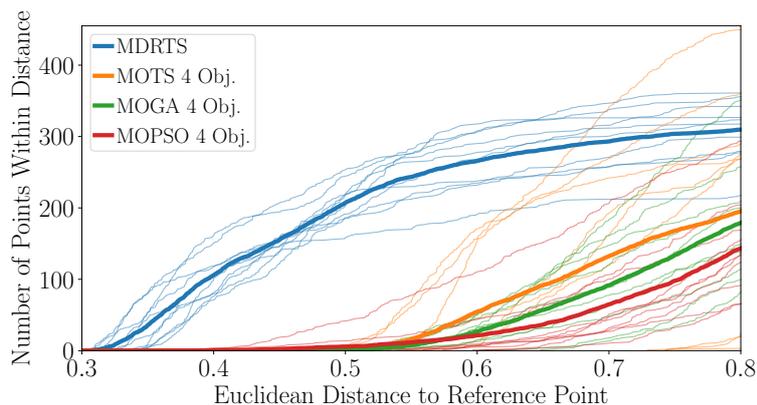


Fig. 6.15 Number of designs produced by the MDRTS algorithm and methods employing the four-objective formulation in Table 6.2 applied to the aerofoil test case in the trade-off scenario that exhibit performance within a given Euclidean distance of the reference point. All runs are plotted as faint lines with the mean shown in bold.

Figure 6.15 shows the number of designs exhibiting performance within a given Euclidean distance of the reference point. Again the new method performs well, producing a large number of designs that are likely to be of interest to the designer. Whilst the multi-objective approaches find at most one design exhibiting performance within a Euclidean distance of 0.4 of the idealised reference point the MDRTS approach finds an average of around 100 aerofoils with this level of performance. The lack of spread in the faint lines also indicates that a large number of interesting designs are produced on a consistent basis.

The new MDRTS method has been successfully applied in this trade-off scenario, producing better results than algorithms employing a more traditional multi-objective problem formulation. The difference between the approaches is more pronounced than in the single-objective scenario due to the additional performance parameter hampering progress of the multi-objective methods. The MDRTS algorithm handles the additional parameter more effectively, consistently generating a large number of designs that exhibit good performance in terms of all of the key quantities of interest.

6.3.2 Comparison to Alternative MDR Implementations

Having compared the performance of the MDRTS algorithm to approaches employing more traditional problem formulations, the next comparison is to the alternative GA and PSO implementations of the MDR approach. Figure 6.16 shows the performance of designs on the C_L - C_D trade-off curves produced by the three implementations, considering all of the aerofoils generated across the 10 runs that satisfy the minimum A_c requirement. The Pareto fronts produced by these methods are closer than those in Figure 6.13 found using the multi-objective approaches. The GA produces good designs in terms of C_L but fails to generate low values of the desirable features. The performance of the PSO is similar to that

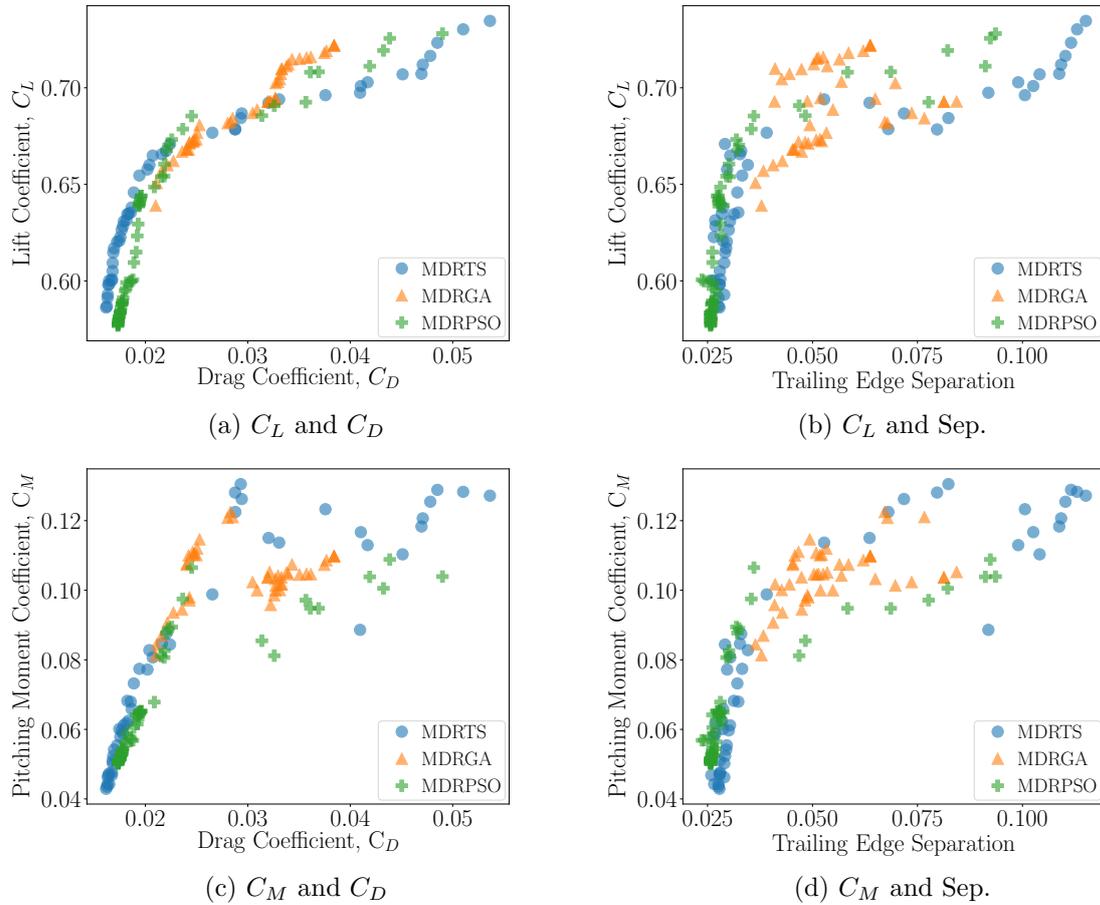


Fig. 6.16 Performance of the best designs found during 10 runs of the different MDR implementations applied to the aerofoil test case in the trade-off scenario.

of the new TS implementation, with the latter generating a slightly more advanced trade-off curve and lower values of C_M and the former producing marginally better designs in terms of trailing edge separation. In the single-objective scenario the underlying search algorithm appeared to have the greatest impact on the performance of the final designs produced. The similarity in the results in Figure 6.16 suggests that in this case it is the formulation that is more important. The new MDRTS methodology is on the correct side of this distinction in both cases, employing the best search technique in the single-objective scenario, and the best formulation for this trade-off study.

Figure 6.17, tracking the best values of the overall performance measure, is similar to Figure 6.5 produced in the single-objective scenario. Again, the GA stalls relatively early in the search, with the PSO and TS continuing to follow a similar trajectory until around 100 evaluations. After this point the progress of the PSO slows but the TS continues to produce aerofoils with increasing values of the overall performance measure. On average the MDRTS algorithm generates performance equivalent to that found by the GA implementation with a

Comparison to Existing Methods Using an Aerofoil Test Case

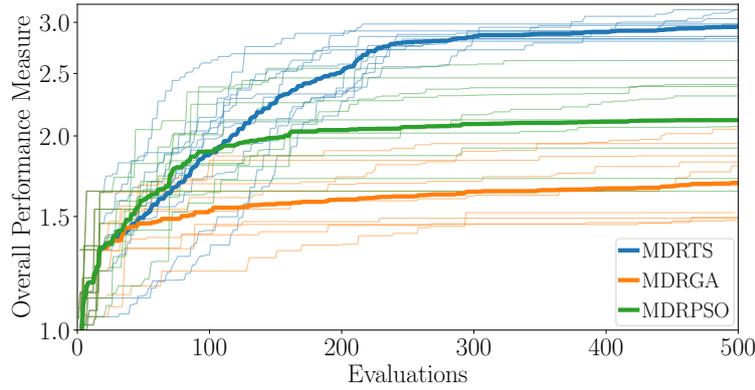


Fig. 6.17 Progression of the best overall performance measure found by the different MDR implementations applied to the aerofoil test case in the trade-off scenario. All runs are plotted as faint lines with the mean shown in bold.

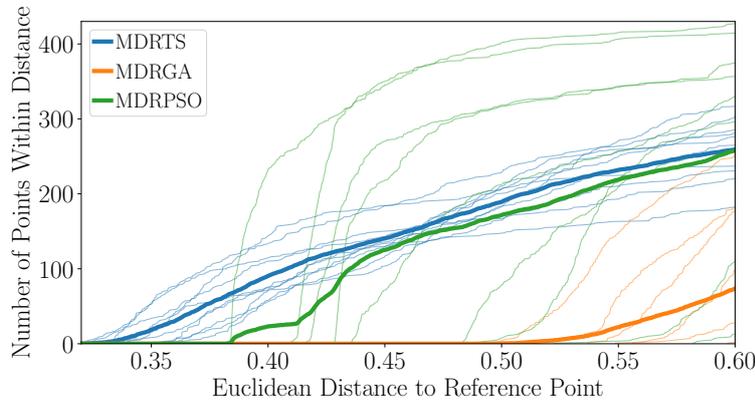


Fig. 6.18 Number of designs produced by the different MDR implementations applied to the aerofoil test case in the trade-off scenario that exhibit performance within a given Euclidean distance of the reference point. All runs are plotted as faint lines with the mean shown in bold.

computational saving of 84%, and equivalent to that produced by the PSO with a saving of 72%. This again demonstrates the exploiting capabilities of the TS method and how the MDR formulation enables this attribute to come to the fore, generating designs exhibiting high levels of performance in a computationally efficient manner.

Figure 6.18, showing the number of designs exhibiting performance within a given Euclidean distance of the reference point, highlights the ability of the new approach to consistently produce a large number of designs that are likely to be of interest to the designer. Whilst the faint lines showing individual runs of the GA and PSO implementations are spread widely, all 10 runs of the TS implementation are close together, demonstrating that the algorithm reliably generates good designs.

The new MDRTS method has performed well when applied to this aerofoil test case in a trade-off scenario. Final designs produced using the approach consistently outperform

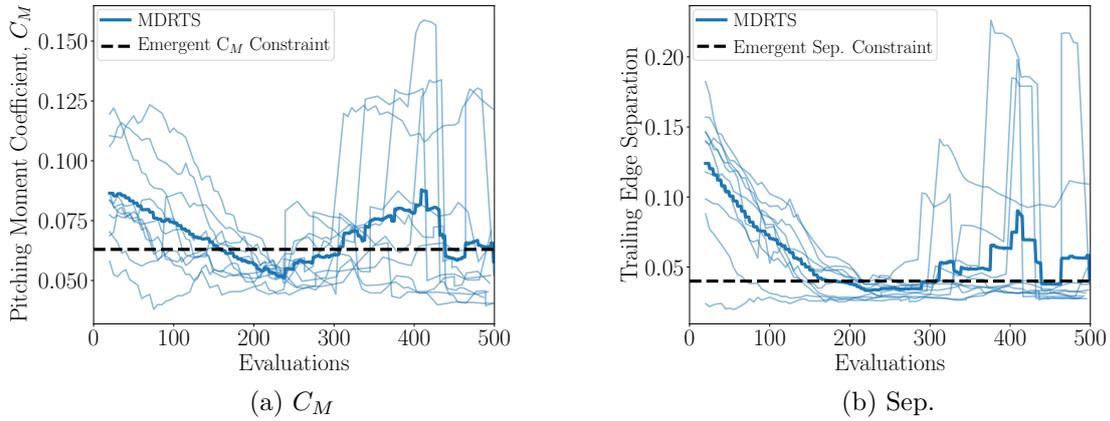


Fig. 6.19 Progression of the desirable features during 10 runs of the MDRTS algorithm applied to the aerofoil test case in the trade-off scenario. Performance of the pattern search centres for each run is plotted as a faint line with the average shown in bold.

those generated using the same formulation but different underlying search algorithms across multiple runs of the different methodologies.

6.3.3 Emergent Constraints

As in the single-objective scenario, the desirable features used by the MDRTS algorithm in this trade-off study exhibit convergent behaviour. Figure 6.19 tracks the values of C_M and separation for designs selected as pattern search centres during each of the 10 runs. Convergent behaviour is less noticeable compared to the single-objective scenario, primarily due to the larger jumps in performance observed as the algorithm intensifies and diversifies the search. A wider range of designs are also being considered as the optimiser investigates a trade-off between two quantities, as opposed to minimising a single objective. Despite this, consistent convergent behaviour is observed in the early stages of each search and suitable limit values for the desirable features can still be discerned that could potentially inform more detailed optimisation studies at a later stage of the design process.

These limits can also be used to compare the new approach to a traditional two-objective method with constraints applied to C_M and separation in addition to the existing minimum A_c requirement. The selected emergent constraint values are slightly relaxed compared to those in the single-objective scenario, accommodating the larger variations seen in Figure 6.19 and acknowledging that higher values of these quantities are likely to be considered acceptable during a trade-off study. The upper limit for C_M is set at 0.063 and that for separation is 0.04, both of which are shown in Figure 6.19.

The resultant two-objective problem formulation, outlined in Table 6.2, is applied using the TS, GA and PSO, with the performance of designs that form the C_L - C_D Pareto front generated by these approaches and the new MDRTS algorithm shown in Figure 6.20, considering all

Comparison to Existing Methods Using an Aerofoil Test Case

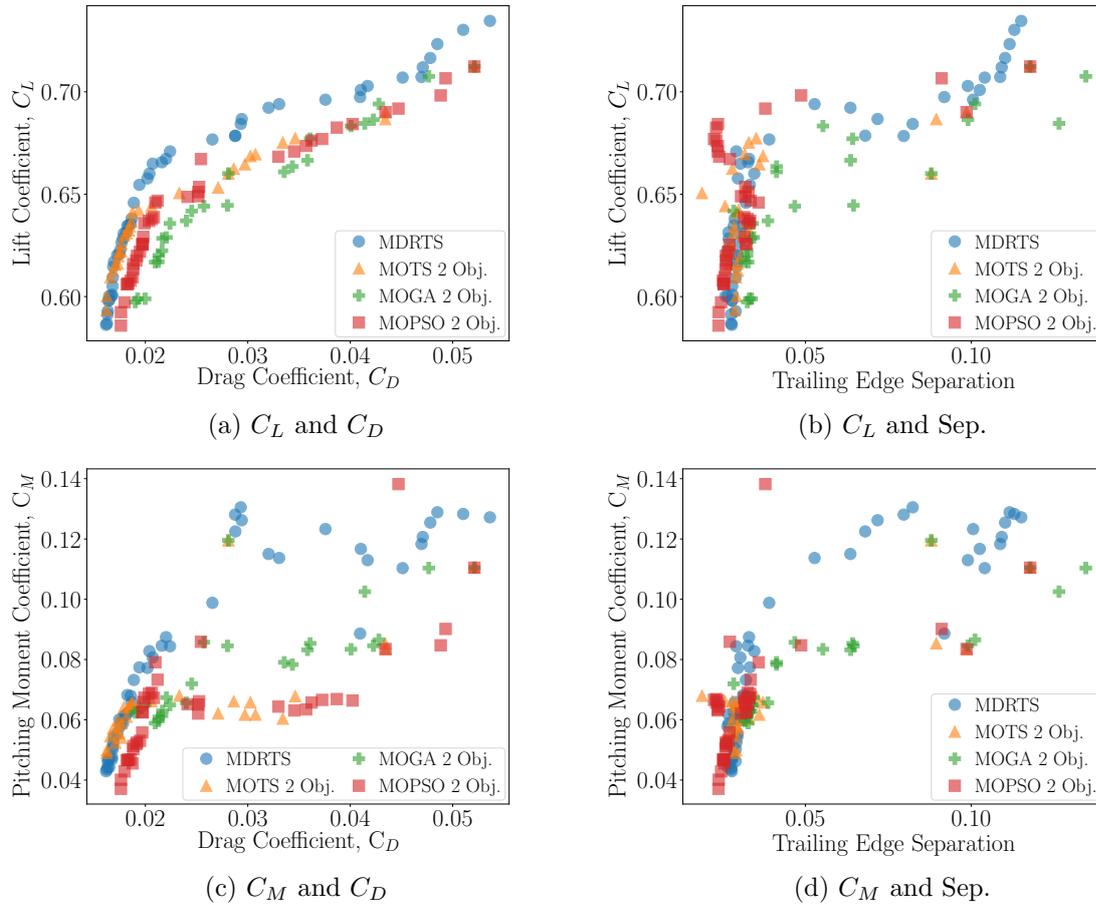


Fig. 6.20 Performance of the best designs found during 10 runs of the MDRTS algorithm and methods employing the constrained two-objective formulation in Table 6.2 applied to the aerofoil test case in the trade-off scenario.

points found during the 10 runs that satisfy the minimum A_c constraint. The trade-off curves are close together, with the two TS methods generating designs with slightly better performance in terms of C_L and C_D than those found using the population-based approaches. The Pareto front produced by the MDR formulation is more dense, consisting of a larger number of designs that are likely to be of interest to the designer. The PSO generates the best designs in terms of the desirable features, but those found by the new method also exhibit among the best performance in terms of these secondary quantities.

Figure 6.21 tracks the best values of the overall performance measure found by the different approaches. The average convergence of the two methods employing TS is similar, suggesting that the MDR formulation is capable of producing equivalent results to the more traditional two-objective method without the need to specify constraint limits for C_M and separation. The MDRTS algorithm also produces better values of the overall performance measure than the population-based approaches, surpassing the best aerofoils found by the

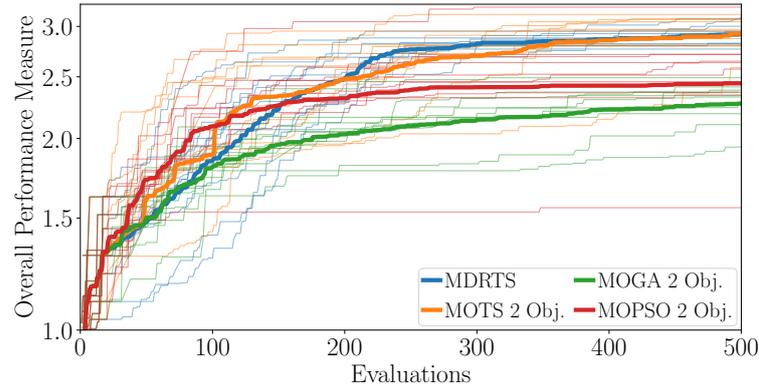


Fig. 6.21 Progression of the best overall performance measure found by the MDRTS algorithm and methods employing the constrained two-objective formulation in Table 6.2 applied to the aerofoil test case in the trade-off scenario. All runs are plotted as faint lines with the mean shown in bold.

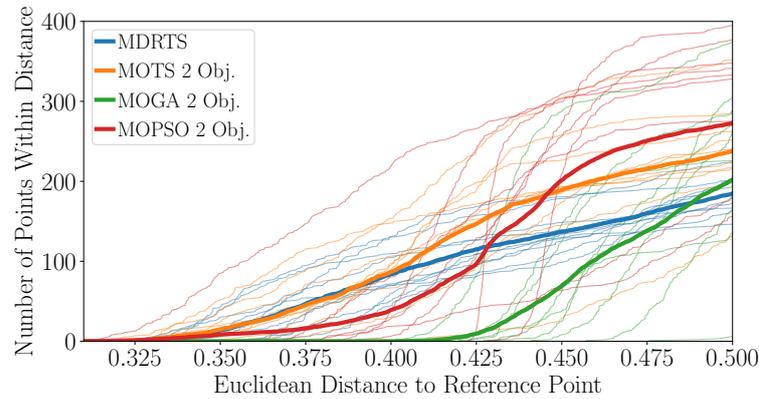


Fig. 6.22 Number of designs produced by the MDRTS algorithm and methods employing the constrained two-objective formulation in Table 6.2 applied to the aerofoil test case in the trade-off scenario that exhibit performance within a given Euclidean distance of the reference point. All runs are plotted as faint lines with the mean shown in bold.

GA using on average 73% fewer calls to the SU^2 analysis function and generating results equivalent to those produced by the PSO with an average computational saving of 59%. The superiority of the TS methods again occurs primarily in the latter stages of the search as the optimiser focusses on exploiting the potential of promising designs.

Figure 6.22, plotting the number of aerofoils exhibiting performance within a given Euclidean distance of the reference point, again shows little difference between the approaches utilising TS. These methods, along with the PSO to a lesser extent, are capable of producing designs with good values of all four quantities of interest. The new MDRTS approach achieves this without requiring constraint limits for C_M and separation to be specified.

The MDRTS algorithm, treating C_M and trailing edge separation as desirable features, produces results that are at least as good as those found using traditional methods that

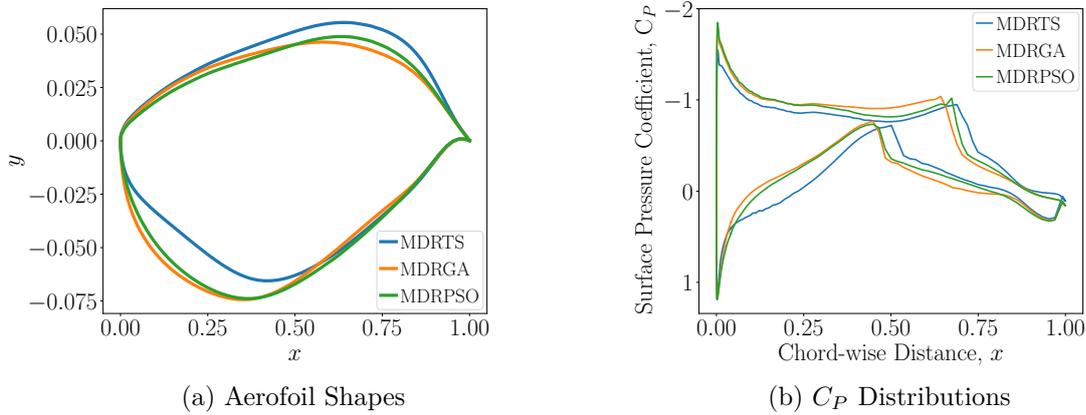


Fig. 6.23 Geometries and C_P distributions of the aerofoils exhibiting the lowest C_D values found during 10 runs of the different MDR implementations applied to the aerofoil test case in the single-objective scenario.

apply constraints to these quantities. Suitable threshold values instead become informative outputs of the optimisation process. This behaviour has been shown to be consistent across both single-objective and trade-off scenarios.

6.4 Information to Enhance Designer Understanding

So far in this chapter the new MDRTS algorithm has been compared to existing methods in terms the ability to produce designs that are likely to be of interest to the designer. This covers the formulation and, to a lesser extent, speed aspects of the motivation discussed in Chapter 3 for developing a new optimisation approach. Another motivation was the need for optimisation algorithms to fulfil the role of enhancing designer understanding and the following paragraphs seek to demonstrate the capabilities of the new approach in this regard.

Figure 6.23 shows the aerofoils exhibiting the lowest C_D values found during the 10 runs of the algorithms employing the MDR formulation applied in the single-objective scenario discussed in Section 6.2.2. The aerofoil shapes and associated surface pressure distributions are similar indicating that all three search algorithms have exploited related physical aspects of the problem to generate performance improvements. When using the new TS implementation the development of this aerofoil can be tracked, producing information to enhance designer understanding and provide assistance when attempting to determine the rationale used by the optimiser to produce the final design. To demonstrate this capability the development of the aerofoil produced by the MDRTS algorithm is tracked from the initial design through to that presented in Figure 6.23.

As discussed in previous chapters, TS generates new designs using a H&J search around an existing point. If this central point is stored as the “parent” of the new designs then reconstructing the path taken from an initial aerofoil to any that are of interest is simple,

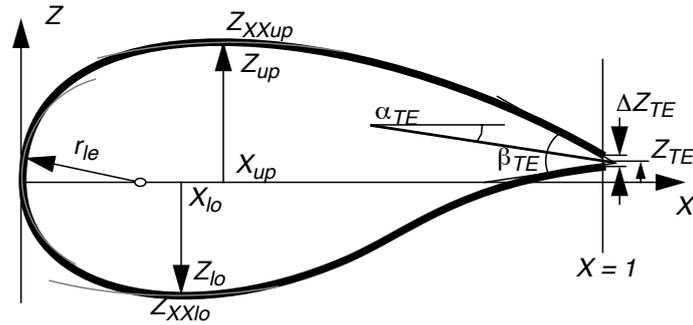


Fig. 6.24 PARSEC aerofoil parameterisation reproduced from Sobieczky [220]. Formed of 11 parameters: leading edge radius (r_{le}); upper and lower crest location (X_{up} , Z_{up} and X_{lo} , Z_{lo}) and curvature (Z_{XXup} and Z_{XXlo}); trailing edge co-ordinate (Z_{TE}), thickness (ΔZ_{TE}), direction (α_{TE}) and wedge angle (β_{TE}).

with just one variable changed at each step. Assuming adequate storage of design data during the optimisation the computational expense associated with extracting the details of this path is negligible.

Plotting the variations of the Hicks-Henne bump functions directly would be of little use to a designer attempting to enhance their understanding as it is difficult to visualise the combined effects of multiple bump functions superimposed on a base geometry. Development information presented in this way would not be interpretable. Instead, the shape of each aerofoil along the path is summarised using the PARSEC technique developed by Sobieczky [220]. This comprises 11 parameters, defined in Figure 6.24, that are easier for the designer to interpret and understand. The variation of these PARSEC parameters along the path from the initial to final designs, excluding the location and thickness of the trailing edge which were fixed during this optimisation, are shown in Figure 6.25. These plots reveal the geometric changes made by the optimiser to produce the final aerofoil.

Figure 6.25a shows that the optimiser has reduced the leading edge radius of the aerofoil, resulting in the sharper suction peak visible in Figure 6.23b. Meanwhile, the wedge angle of the trailing edge increases as the optimisation progresses (Figure 6.25f) to accommodate the aft-wards movement of the crest locations on the upper and, to a lesser extent, lower surfaces (Figure 6.25c). Moving the crest further back along the aerofoil causes the shock wave to shift in the same direction, widening the region of low pressure on the upper surface and increasing C_L . Again the effect of this change can be seen in the C_P plot of the final design in Figure 6.23b, whilst the corresponding aerofoil shape in Figure 6.23a shows the narrow nose and rear-ward location of the upper surface crest.

Both of these highlighted features may cause problems in practice. Such a narrow leading edge would prohibit the use of slats, often required to enhance low-speed performance, and the sharp suction peak is likely to cause separation of the flow, particularly at increased angles of attack [148]. Moving the shock further aft strengthens it, increasing the risk of

Comparison to Existing Methods Using an Aerofoil Test Case

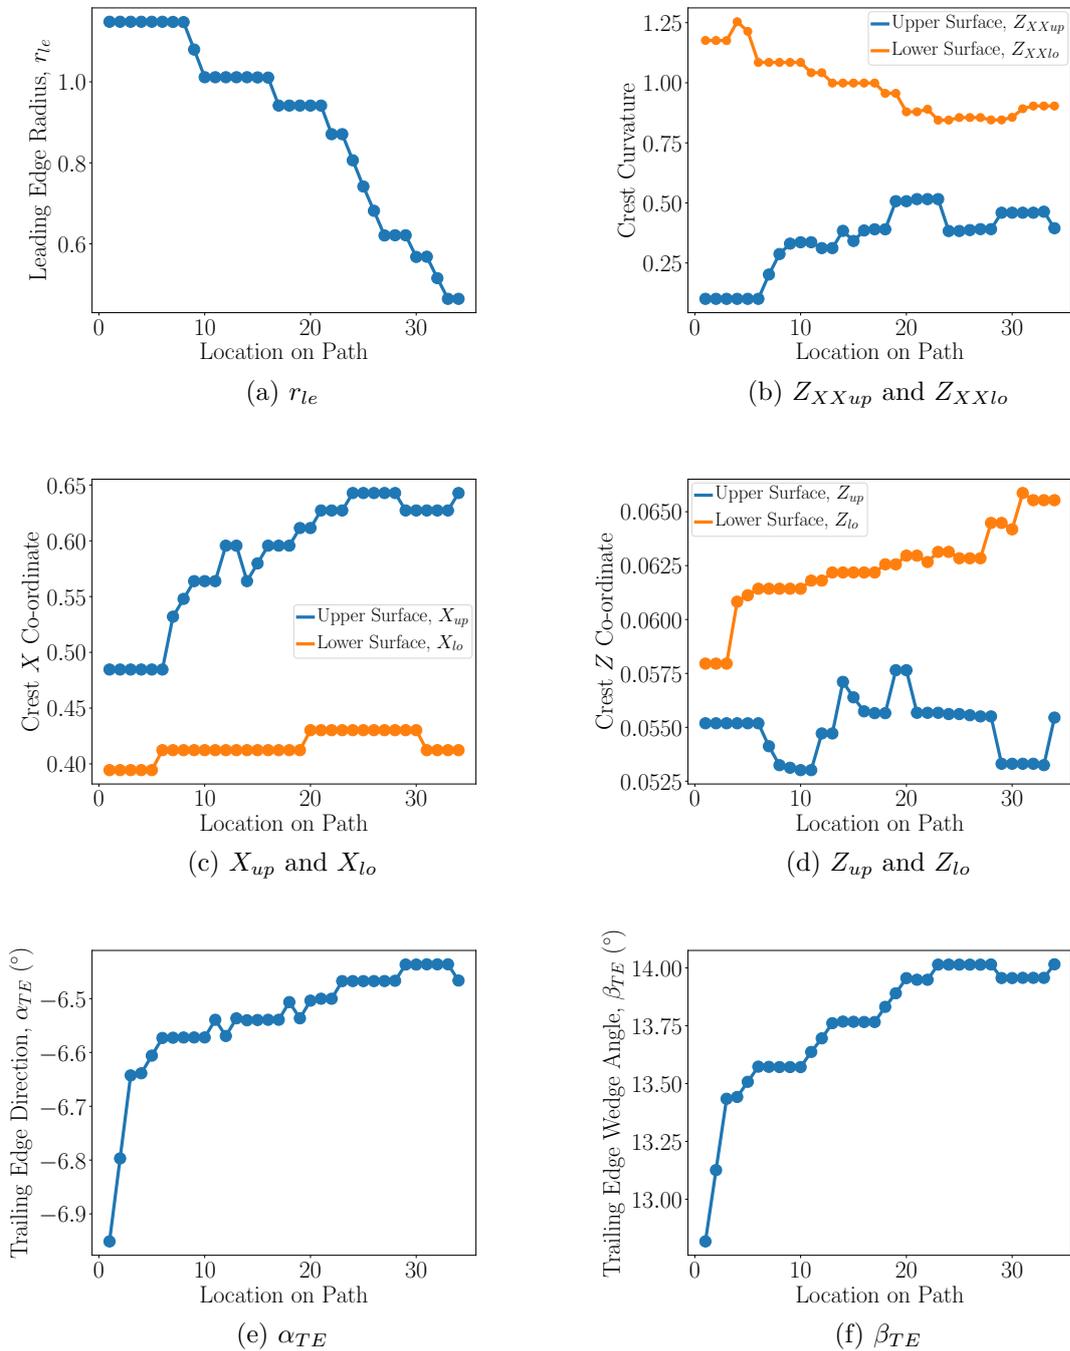


Fig. 6.25 Variation of the PARSEC parameters defined in Figure 6.24 along the development path to the aerofoil exhibiting the lowest C_D value produced by the MDRTS algorithm in the single-objective scenario.

shock-induced boundary layer separation and the onset of buffet. These aspects of the problem have not been accounted for in the analysis techniques employed here, but the use of an algorithm capable of providing development information to the designer allows them to be considered at the post-processing stage once the optimiser has explored the available design space. One possible mitigation technique for these problems, made possible by the provision of design development information, might be to select an aerofoil analysed earlier in the search along the path towards the final design. The chosen geometry would exhibit some of the performance improvements supplied by narrowing the leading edge and moving the crest further aft without going to the extremes of the final design that lead to concerns related to separation and buffet. Alternatively, the optimisation could be re-run with an additional desirable feature informing the optimiser to maintain a large leading edge radius if possible.

The interested designer could use this data to gain deeper physical understanding, but even this brief discussion highlights the ability of the new algorithm to provide rich design development information at negligible additional computational cost. This development information is available for all designs produced by the algorithm, presenting numerous opportunities for the enhancement of designer understanding. This type of information is not available from the GA or PSO implementations due to the large changes made to the design vector and the lack of ordered, interpretable movement through the design space.

Whilst the provision of design development information sufficiently addresses the limitation of existing methods being ill-equipped to fulfil the role of enhancing designer understanding, the onus remains on the designer to interpret the available data. Improved storage during the optimisation could allow this interpretation to also be carried out by the algorithm. For example, the position of the shock wave on the upper surface could be logged for each design and any change correlated with that of the crest location, allowing the optimiser to automatically generate physical justifications for the final designs produced. Implementing this technique is beyond the scope of this thesis, but could represent a fruitful avenue of further research, ultimately leading to the development of explainable design optimisation techniques that are capable of producing interpretable rationale information alongside improved designs. As well as easing the task of justifying the outputs of optimisation, this approach could potentially generate new physical insight if the optimiser produces fully explained performance improvements in an unconventional manner.

6.5 Summary

The comparisons in this chapter demonstrate the capabilities of the new MDRTS algorithm. Whether the scenario involves improving a single objective or investigating a trade-off between several quantities of interest the new approach has been shown to consistently produce designs

that are likely to be of interest to the designer in a computationally efficient manner. The MDRTS method outperforms existing approaches employing more traditional objectives-and-constraints formulations and those using the same MDR problem definition in combination with alternative search algorithms. Average computational savings of over 80% compared to these existing methods have been demonstrated in some cases.

Performance parameters treated as desirable features have also been shown to exhibit convergent behaviour, becoming additional outputs of the optimisation process. This behaviour could prove useful when attempting to apply optimisation in novel physical domains where suitable values for constraint limits are unknown. The emergent thresholds could also inform optimisations conducted at a later stage in the design process when single-objective, perhaps gradient-based, search algorithms are preferred. In this chapter the emergent limit values were used to inform additional constrained optimisation formulations, with the MDRTS algorithm shown to perform at least as well as these traditional approaches without the need to specify limit values for the constrained quantities.

As well as overcoming limitations associated with problem formulation, the capability of the new algorithm to fulfil the role of enhancing designer understanding has also been demonstrated. Rich design development information is available that enables insights to be made into the physical phenomena and rationale behind observed performance improvements. This data can assist designers, enhancing their understanding to promote creativity and innovation as well as increasing the likelihood of the performance of final designs being successfully justified and therefore accepted in an industrial setting.

Through the assessment carried out in this chapter the new MDRTS algorithm has been shown to be an effective optimisation methodology that is ready to be applied to the more challenging problem of preliminary axial compressor design.

Chapter 7

Initial Application to a Six-Stage Axial Compressor

In the previous chapter the capabilities of the new MDRTS algorithm were demonstrated using an aerofoil test case, successfully overcoming the limitations of existing methodologies associated with formulation and understanding highlighted in Chapter 3. This chapter assesses whether the technique is equally effective when applied to the more complex problem of preliminary axial compressor design using a six-stage test case. Once the effectiveness of the new approach has been verified attention in the next chapter turns to acceleration, addressing the final of the three limitations of existing methodologies that form the motivation for this thesis.

7.1 Six-Stage Axial Compressor Test Case

A six-stage axial compressor test case, based on the HPC of the three-spool turbofan engine depicted in Figure 1.1, is used to assess the MDRTS algorithm when applied to this more complex problem.

7.1.1 Performance Parameters

As discussed in Chapter 2, axial compressor design problems are considered from a purely aerodynamic perspective in this thesis. Even with this simplification numerous performance parameters need to be taken into account. The overall pressure ratio, PR , and polytropic efficiency, η_p , achieved by the machine are perhaps the most obvious, alongside an assessment of stability in the form of a surge margin, SM . The Mach number and flow angle at the compressor exit, M_{exit} and α_{exit} respectively, also need to be considered to prevent excessive pressure losses in downstream components [238], and the loading of individual blades is monitored using the de Haller number [47], DH , and the diffusion factor [157], DF . Further

performance parameters could be considered, such as the axial length or mass of the machine, but this initial study is limited to these seven quantities of interest.

The concerns of some other disciplines are acknowledged through a series of geometric constraints applied to the proposed designs. Following the recommendations of Walsh and Fletcher [238] the blade speeds at the hub and casing are limited to 350 and 400 ms^{-1} respectively to ensure acceptable stresses in the discs and blades. The aspect ratios, calculated using the axial chord, are required to lie between 1.5 and 2.8 to prevent the optimiser from producing compressors with unrealistically wide or narrow blading. Finally, to maintain compatibility with the existing engine architecture in Figure 1.1 the mean radius should not decrease through the machine.

7.1.2 Parameterisation

Following several previous works [82, 130, 185, 193] the parameterisation makes extensive use of Bézier curves [188]. These enable smooth variation of quantities through the compressor whilst simultaneously decreasing the total number of design variables required. Cubic Bézier curves, with four control points each, define the mean radius, stage pressure ratios, stage exit flow angles, rotor and stator aspect ratios, and the number of rotor and stator blades, whilst additional variables allow for alteration of the mass flow rate, inlet flow angle, blade twist and casing radius. This results in a total of 32 design variables giving the optimiser control over the compressor geometry.

As well as varying each of these parameters individually during the H&J pattern search used to generate new designs, the optimiser is also equipped to alter some of the Bézier control points in tandem. For the curves fitted to mean radius, stage pressure ratio and stage exit flow angle the optimiser has the ability to move control points in adjacent pairs or all at once. This allows the algorithm to make larger changes to the design, accelerating movement through the design space. The provision of development information to enhance designer understanding is not negatively impacted as the resultant alterations are still incremental and vary a single aspect of the design.

For this initial assessment of the MDRTS algorithm some simplification is provided by treating the casing radius as constant and by fixing the number of stages at the initial value. In Chapter 9, once the new approach has been accelerated using multi-fidelity methods, these restrictions are removed as the technique is applied to a more complex N-stage test case.

7.1.3 Open Source Analysis System

The aim of this thesis is to facilitate the use of high-fidelity optimisation in preliminary axial compressor design. An analysis system is therefore required that is capable of providing performance predictions for multi-stage machines using the RANS equations. In most existing turbomachinery optimisation studies this analysis is undertaken using proprietary codes

developed by one of the major engine manufacturers [15, 60, 79–84, 101, 113–116, 229]. However, due to the lack of industrial partnership these trusted and high-performing analysis codes are not available in this work.

Instead, the open source analysis system Multall is used. Developed by Denton [52], this features a three-dimensional RANS solver designed and optimised specifically for turbomachinery applications which has seen extensive use in axial compressor analysis and development, both for research and commercial purposes. The results have been validated through these applications and Multall therefore represents a suitable and reliable high-fidelity analysis tool for the purposes of this study.

The Multall RANS solver requires a full description of the compressor geometry in order to build a mesh and carry out the analysis. A script written by the author generates this dataset automatically, taking as inputs a combination of the design variables outlined in the previous section and parameters such as rotational speed and the inlet conditions. Meanline techniques [210] are used along with the loss models presented by Aungier [6] to generate a full compressor geometry in less than a second. This meanline script can also provide computationally cheap low-fidelity performance predictions for verification purposes.

With the exception of SM , the performance parameters discussed in Section 7.1.1 are available as direct outputs from Multall. In order for SM to be calculated knowledge of PR at the stall point is required (see Equation 2.1). Denton [52] suggests that this can be approximated by reducing the exit pressure boundary condition inputted to Multall and taking stall as being the point at which the analysis fails to converge. Whilst finding this point is feasible, it would require an iterative procedure involving several runs of the analysis code. For a single design this might be acceptable, however during optimisation hundreds of designs require analysis and carrying out multiple high-fidelity runs for each soon renders the problem computationally intractable [89]. To avoid this complication most existing studies applying optimisation to axial compressor design use empirical correlations to estimate SM [46, 60, 79, 81, 82, 84, 115, 130]. As with the high-fidelity analysis codes, these correlations are almost exclusively proprietary. However, Schweitzer and Gargaroglio [212] document a method developed at Pratt and Whitney that has been successfully used in previous work applying optimisation to the preliminary design of axial compressors [99]. Given that in optimisation it is the relative performance between designs that is of primary interest, rather than the quantitative values themselves [51, 76], the accuracy of this correlation should be sufficient for the purposes of this study.

7.1.4 Initial Geometry

Eastwood et al. [60] highlighted the importance of the starting geometry for axial compressor optimisation studies and the difficulty experienced by designers attempting to generate appropriate initial designs. The subject of how starting geometries are produced is somewhat

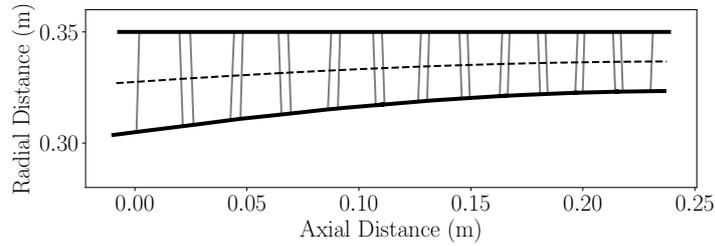


Fig. 7.1 Annulus of the initial six-stage axial compressor.

ignored in the literature, with the likely reason for this being that suitable proprietary geometries have been provided through industrial partnerships.

No such luxury is present in this work, with an initial design instead produced using a meanline generation tool called Meangen provided as part of the Multall package [52]. Overall parameters such as the number of stages, inlet conditions, rotational speed, mass flow rate and constant casing radius are chosen based on publicly available data for three-spool turbofan engines. These are fed into the Meangen program along with values for the axial chords of the blades and the desired loading, reaction and flow coefficient for each stage. The latter three parameters are non-dimensional quantities commonly used in the design of axial compressors [42]. The loading, ψ , is defined as the ratio of the stagnation enthalpy change across the rotor to the square of the blade speed, and is a measure of the work done by the stage. Flow coefficient, ϕ , is the ratio of the axial velocity through the stage to the mean blade speed, and is a measure of how quickly the flow is moving. Finally, the degree of reaction, Λ , is defined as the ratio of the static enthalpy rise across the rotor to that of the stage as a whole, with values greater than 0.5 indicating that more diffusive effort is being required of the rotor than the stator.

Meangen takes these inputs and produces a constant radius axial compressor which can be parameterised using the scheme described in Section 7.1.2 and passed through the open source analysis system. Iterative variation of the inputs, conducted by hand, results in a compressor exhibiting sufficient values of the relevant performance parameters. The geometry of this initial design is shown in Figure 7.1 with the performance outlined in Table 7.1. Whilst this compressor by no means represents a good design, it is sufficient for the purposes of this study where the aim is to demonstrate the capabilities of the new MDRTS algorithm.

7.2 Suitability of Different Algorithms

In the previous chapter the MDRTS algorithm was compared to approaches using alternative population-based search methods. Whilst the GA and PSO were able to generate good designs for the aerofoil test case they are unsuitable for application to the more complex problem of axial compressor design due to their inability to handle design spaces featuring

Table 7.1 Performance of the initial six-stage axial compressor.

η_p	0.913
PR	2.75
SM	15.2%
M_{exit}	0.334
α_{exit}	7.27°
DH_{min}	0.759
DF_{max}	0.400

large regions of infeasibility. Jaeggi et al. [110], among others [80, 82, 84, 136, 137, 141], note how the big changes made to the vector of design variables by these methods result in vast numbers of infeasible designs being generated, hampering the progress of the search. In contrast, the small incremental alterations made by TS ensure feasibility is maintained, enabling efficient and effective navigation through even the most complex design spaces.

Turbomachinery problems exhibit these challenging landscapes, with infeasibility common even in supposedly unconstrained scenarios due to geometric incompatibility and analysis code failure [10]. The former problem was discussed by Eastwood et al. [60], highlighting the difficulty associated with generating even a single compatible design to use as a starting point for compressor optimisation studies. The latter problem, of analysis code failure, is evidenced in a number of works attempting to apply optimisation to axial compressor design. Baert et al. [8, 10] constructed surrogate models using design of experiments techniques, with only 44% of the generated designs converging successfully in the first instance [8] and 59% in the second [10], resulting in a total of 283 analyses being wasted on failed designs. Brooks et al. [25] faced similar problems with 40% of their designs failing to converge, whilst Keskin and Bestle [130] reported one algorithm wasting over 11,000 function evaluations on failed designs.

To illustrate the problems faced by GAs and PSO when applied to this problem a simple demonstrative axial compressor optimisation is conducted. Both population-based methods and a TS algorithm are applied to an optimisation treating η_p as the sole objective starting from the six-stage geometry shown in Figure 7.1, with the other performance parameters constrained using a penalty method to be no worse than the values in Table 7.1. The meanline code provides cheap performance predictions, with each algorithm run 10 times using a budget of 1000 analyses.

Figure 7.2 tracks the number of successful geometries generated by the different algorithms, with the data plotted both as raw numbers and as a percentage of the total evaluations attempted. A successful evaluation is one for which a compressor geometry is produced, making this purely a test of the ability of an algorithm to avoid code failure rather than accounting for any infeasibility due to the violation of constraints. The approach based

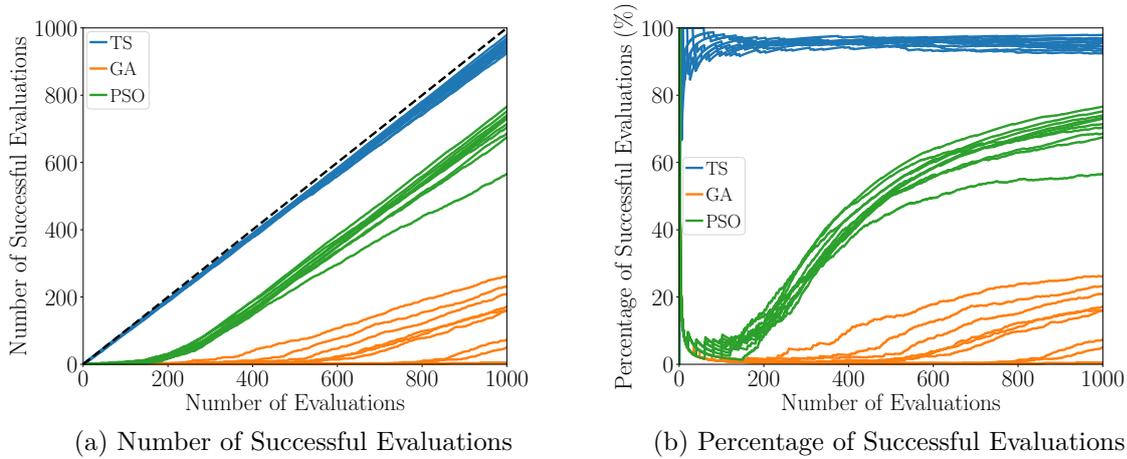


Fig. 7.2 Number of successful geometries generated by the TS, GA and PSO applied to a meanline single-objective axial compressor optimisation.

on pattern search consistently generates successful designs, with only around 5-10% of the available budget wasted attempting to analyse failed geometries. This success rate is independent of the computational budget. In contrast, the population-based methods spend a significant proportion of the available budget attempting to analyse machines that fail to converge. In one instance the GA is unable to generate a single successful geometry, with even the best results only managing to allocate 25% of the budget to convergent machines. The PSO fares slightly better, but still requires a computational budget of over 1000 analyses before it can be expected to waste less than a quarter on failing designs.

The performance of population-based methods is likely to be even worse when applied to full-scale high-fidelity optimisation problems. This demonstration employed a meanline analysis code that is more robust than iterative three-dimensional RANS analysis, with the constraint violation of designs that do converge also not taken into account.

When attempting to facilitate the use of high-fidelity analysis at the preliminary design stage wasting such a high proportion of the computational budget is unacceptable. The two population-based methods are therefore not applied to the axial compressor test case, with the MDRTS approach instead compared to TS methods using alternative problem formulations.

7.3 Single-Objective Scenario

In the previous two chapters a distinction has been made between scenarios where a single performance parameter is improved and those in which a trade-off is sought between multiple quantities of interest. Both are used to assess the performance of the MDRTS algorithm

when applied to the six-stage axial compressor test case, with results from the single-objective scenario presented in this section.

7.3.1 Problem Formulations

The simplifying application framework developed in Chapter 5 eases the task of generating a suitable MDR formulation for the axial compressor design problem. Without this framework the designer would be faced with seven performance parameters (η_p , PR , SM , DH , DF , M_{exit} and α_{exit}) that could be treated using any number of dominance relations arranged in any order. Instead, each parameter simply needs to be assigned to the objectives, desirable features or constraints classifications. For this first scenario η_p is selected as the sole objective, following several previous axial compressor optimisation studies [100, 104, 113, 115, 121, 231]. PR and SM are treated as constraints, with lower limits for these quantities likely to be specified in the design requirements. The recommended limits for DH [47] and DF [157] are also applied to ensure acceptable loading of individual blades.

The treatment of M_{exit} and α_{exit} is one of the main motivators for applying the MDR formulation to this axial compressor design problem. Using current methods these quantities would be treated as constraints, with upper limits such as those suggested by Walsh and Fletcher [238] applied, resulting in a single-objective problem formulation. However, this approach fails to accommodate the desire for these quantities to be minimised to prevent excessive pressure losses in downstream components [238]. A possible alternative is to treat them as additional objectives, alongside η_p , resulting in a three-objective formulation. Whilst this informs the optimiser of the desire for these quantities to be minimised it also treats M_{exit} and α_{exit} with an equal level of importance as η_p . The MDR formulation enables a more accurate representation, with M_{exit} and α_{exit} classified as desirable features. This informs the optimiser that these values should be minimised where possible, but with a lower importance compared to the main objective of increasing η_p .

The MDR formulation and the two traditional alternatives are outlined in Table 7.2, where limit values for the constraints are specified. Also present in this table is a fourth formulation employing emergent limits that is the subject of Section 7.3.3. Methods based on preferences and priorities are again excluded from comparison for the reasons discussed in the previous chapter. For the formulations treating M_{exit} and α_{exit} as additional objectives and desirable features, the upper limits on these quantities are retained. Whilst the designer does want to minimise M_{exit} and α_{exit} they also require these values to be below the thresholds suggested by Walsh and Fletcher [238]. It therefore makes sense to apply these constraints even when instructing the optimiser to minimise the quantities if and where possible.

As in the aerofoil test case, constraints are imposed using a penalty method, with individual terms calculated as the amount of constraint violation normalised by the limit value itself. The geometric constraints described in Section 7.1 are applied using a barrier

Initial Application to a Six-Stage Axial Compressor

Table 7.2 Formulations applied to the six-stage axial compressor test case in the single-objective scenario.

	MDR	Multi-Objective	Single-Objective (Standard)	Single-Objective (Emergent)
Objectives	η_p	$\eta_p, M_{exit}, \alpha_{exit}$	η_p	η_p
Desirable Features	M_{exit}, α_{exit}	-	-	-
Constraints	$PR \geq 2.75$ $DF \leq 0.6$ $DH \geq 0.72$ $SM \geq 15.0\%$ $M_{exit} \leq 0.35$ $\alpha_{exit} \leq 10.0^\circ$	$PR \geq 2.75$ $DF \leq 0.6$ $DH \geq 0.72$ $SM \geq 15.0\%$ $M_{exit} \leq 0.35$ $\alpha_{exit} \leq 10.0^\circ$	$PR \geq 2.75$ $DF \leq 0.6$ $DH \geq 0.72$ $SM \geq 15.0\%$ $M_{exit} \leq 0.35$ $\alpha_{exit} \leq 10.0^\circ$	$PR \geq 2.75$ $DF \leq 0.6$ $DH \geq 0.72$ $SM \geq 15.0\%$ $M_{exit} \leq 0.322$ $\alpha_{exit} \leq 1.0^\circ$

method, with the optimiser prevented from selecting designs that violate these limits as centres for the pattern searches used to generate new designs.

Each of the formulations in Table 7.2 is applied to the six-stage axial compressor test case using the initial geometry described in Section 7.1.4, a computational budget of 500 Multall analyses and the algorithmic parameters in Table 4.1. A single Multall analysis of the six-stage machine takes around 30 minutes on an Intel Xeon 2.13 GHz CPU, with this increasing to 50 minutes when eight designs are analysed simultaneously. Using this parallelisation strategy the overall time taken for each optimisation is around 52 hours.

7.3.2 Results

Figure 7.3 shows the performance of machines found using the different formulations that satisfy all of the constraints and are also non-dominated in terms of η_p , M_{exit} and α_{exit} . The designs exhibiting the highest values of η_p are generated using the single-objective approach. However, due to the lack of incentive for an optimiser using this formulation to reduce M_{exit} and α_{exit} , these high- η_p machines have poor values of both quantities. The multi-objective approach successfully produces designs with lower M_{exit} and α_{exit} values, but this comes at the expense of performance in terms of η_p . This is perhaps not unexpected, as improvement in the secondary quantities is unlikely to be achieved without a trade-off with another aspect of performance.

The results produced by the new MDRTS algorithm appear to form two groups distinguished by their values of α_{exit} . This is due to the non-dominated filtering applied to the results, with the gap between the two groups actually being well populated when the performance of all generated designs is plotted. The machines exhibiting low values of

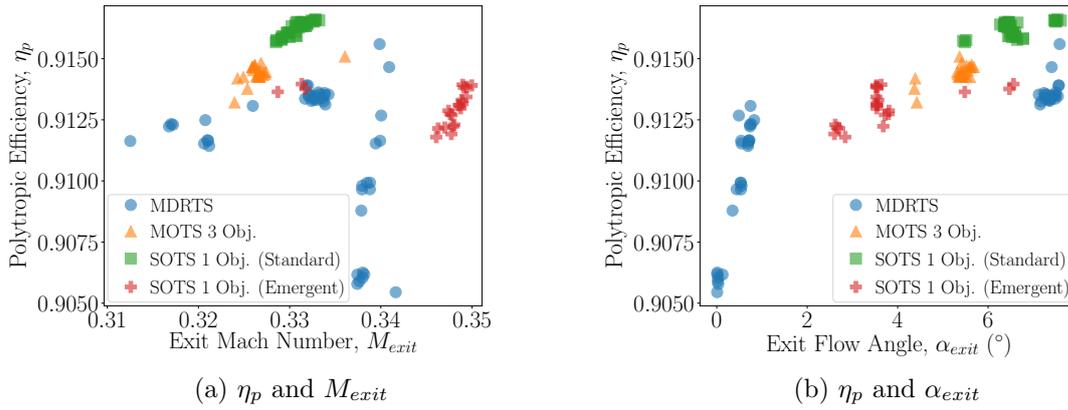


Fig. 7.3 Performance of the best designs found by the formulations in Table 7.2 applied to the six-stage axial compressor test case in the single-objective scenario.

α_{exit} are more likely to be of interest to the designer, with some of these also having lower M_{exit} values than are achieved using the traditional single- and multi-objective formulations. However, following the trend set by the multi-objective results, achieving these low values of the desirable features comes at the expense of a further reduction in η_p . This η_p value is calculated by considering the compressor in isolation, and it may be that the overall efficiency of an engine incorporating one of these compressors with reduced M_{exit} and α_{exit} would actually be higher due to reduced losses in downstream components. Determining whether this is true would, as a minimum, require an integrated optimisation scheme that includes the effects of lower M_{exit} and α_{exit} values on the other components in the gas turbine. Even taking into account the lower values of η_p , the MDR formulation proves effective, generating interesting options for the designer to consider.

These conclusions are supported by the reference point method used in previous chapters. The reference point is a hypothetical idealised design given the best performance in terms of η_p , M_{exit} and α_{exit} found by any of the formulations applied to this problem. Figure 7.4 tracks the best values of the overall performance measure found by each formulation, with this calculated as the reciprocal of the minimum distance to the idealised reference point in a space containing normalised versions of the objective and the two desirable features. Whilst the MDRTS algorithm eventually finds designs with values of the overall performance measure that are over twice as large as those generated by the next best approach, it takes around 125 evaluations to make significant improvements. This is due to a delay in reducing the value of α_{exit} . One of the Bézier curves in the parameterisation is fitted to the flow angle at the exit of each stage, which means the only way to significantly vary α_{exit} is to alter the design variables corresponding to the control points of this curve. During the H&J iterations TS takes a random sample of the design variable changes, moving on if some improvement is found. It can therefore take a few iterations before the parameter affecting α_{exit} is selected

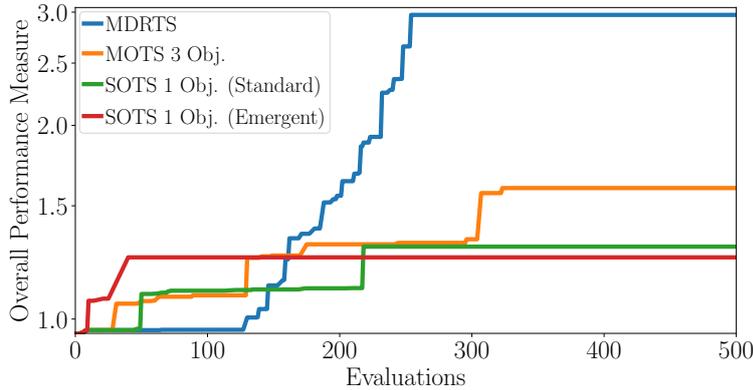


Fig. 7.4 Progression of the best overall performance measure found by the formulations in Table 7.2 applied to the six-stage axial compressor test case in the single-objective scenario.

by the optimiser, at which point it is varied a number of times in succession leading to the rapid increase in performance observed in Figure 7.4.

This behaviour shows that the random sampling suggested by Jaeggi et al. [110] to accelerate TS sometimes results in improvement opportunities being missed. However, the speed-up provided outweighs any missed opportunities as the alternative is to analyse over 60 designs per iteration, allowing fewer than 10 steps to be made through the design space using the limited computational budget available here. Randomly sampling eight designs at a time allows the MDRTS algorithm to make 34 steps, suggesting that the heuristic is worthwhile despite the delayed improvement seen in Figure 7.4. Attempts to accelerate the MDRTS algorithm are made in the following chapter using multi-fidelity techniques, overcoming this computational cost problem.

Figure 7.5, plotting the number of designs exhibiting performance within a given Euclidean distance of the reference point, shows that MDRTS produces a larger number of designs that are likely to be of interest to the designer. Whilst this suggests that the MDRTS algorithm performs significantly better than the traditional methods it should be considered in conjunction with Figure 7.3, which shows that much of the improved performance is in terms of the desirable features at the expense of η_p .

Overall, the MDRTS algorithm performs well, producing a large number of designs exhibiting high levels of performance when given the same computational budget as traditional methods. It has done so using a more accurate representation of what the designer actually wants from the optimisation.

7.3.3 Emergent Constraints

An alternative way to achieve lower levels of M_{exit} and α_{exit} using traditional methods might be to set aspirational constraint limits for these quantities. Through the use of penalty functions the optimiser would try to reach these aspirational limits, resulting in lower values

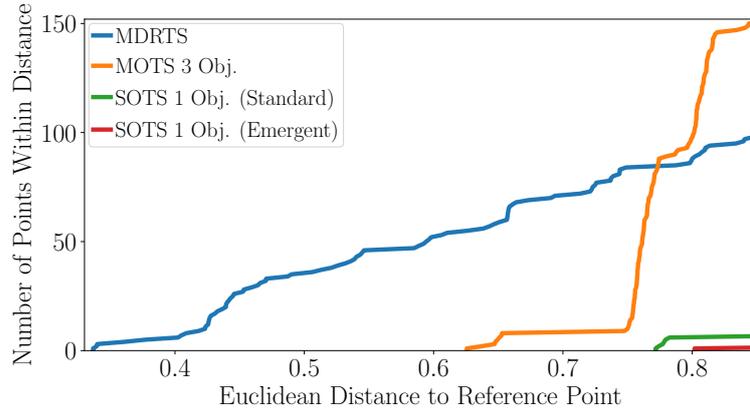


Fig. 7.5 Number of designs produced by the formulations in Table 7.2 applied to the six-stage axial compressor test case in the single-objective scenario that exhibit performance within a given Euclidean distance of the reference point.

being exhibited by the machines produced during the search. Selection of suitable constraint values can be informed by the results of the MDRTS algorithm, as limits are seen to emerge from the optimisation just as they did in the aerofoil test case in Chapter 6.

Figure 7.6 tracks the key performance parameters for all of the designs generated during the MDRTS optimisation, with values for the machines selected as pattern search centres highlighted in bold. The values of α_{exit} exhibit convergent behaviour, approaching an emergent limit of zero after around 300 evaluations. M_{exit} appears to converge towards 0.32 before suddenly increasing for the remainder of the search. Figure 7.6c reveals that η_p decreases in a similarly sudden fashion at the same point, with the optimiser accepting designs with lower η_p and higher M_{exit} as it seeks to further reduce α_{exit} . Rather than being a limitation of the MDR formulation itself, this behaviour instead highlights a weakness of the Pareto dominance criterion used to assess the desirable features in that it allows relatively small improvements in α_{exit} to be accepted despite the large increase in M_{exit} . Alternative dominance criteria could be used to overcome this limitations [12], with the ability of the MDR formulation to accommodate any type of dominance relation being a key strength [39]. However incorporating and evaluating alternatives to Pareto dominance is beyond the scope of this study.

When using the emergent behaviour in Figure 7.6 to define suitable aspirational limit values it is sensible to take this weakness of Pareto dominance into account. The limit values that emerge after around 200 evaluations are therefore selected, as these correspond to designs with higher η_p and lower M_{exit} that are more likely to be of interest to the designer. Limits of 0.322 for M_{exit} and 1.0° for α_{exit} are chosen, as shown in Figure 7.6.

These emergent constraints are applied in a second single-objective optimisation, outlined in Table 7.2, using the same starting point and computational budget as the MDRTS approach, with the performance shown in Figures 7.3, 7.4 and 7.5. Whilst some reduction in

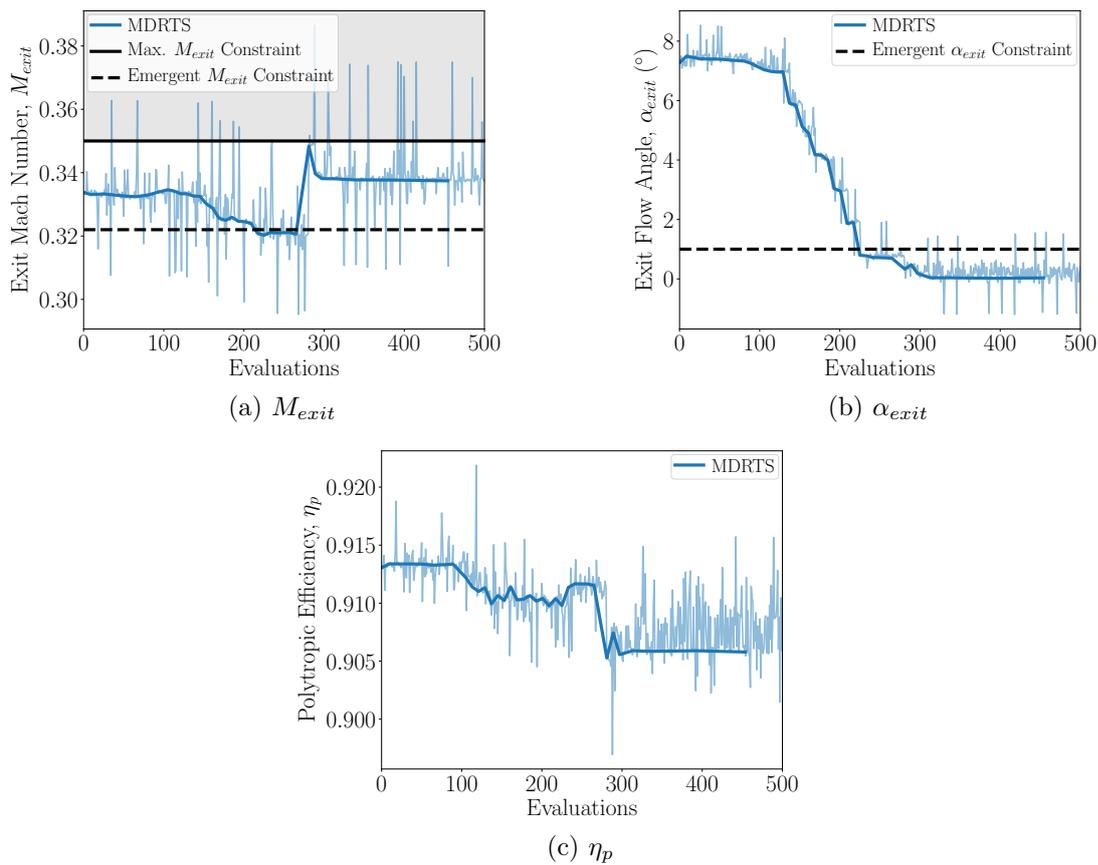


Fig. 7.6 Progression of the objective and desirable features during the search using the MDR formulation in Table 7.2 applied to the six-stage axial compressor test case in the single-objective scenario. Performance of all designs is plotted as a faint line with that of the pattern search centres highlighted in bold.

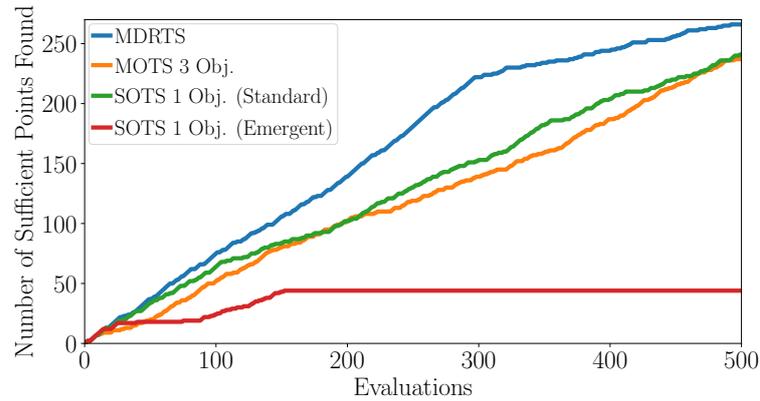


Fig. 7.7 Number of sufficient points found by the formulations in Table 7.2 applied to the six-stage axial compressor test case in the single-objective scenario.

α_{exit} can be seen in Figure 7.3b the lowest M_{exit} values are higher than those found using standard constraints when there was no incentive for the optimiser to reduce this performance parameter. Moreover, neither of the aspirational constraints have been satisfied during the optimisation and it fails to produce designs exhibiting performance comparable to that found using the MDRTS algorithm.

The reason for this poor performance is revealed in Figure 7.7 which tracks the number of sufficient points generated by each of the methods applied to this problem. A sufficient point is any that satisfies the constraints applied in the MDR formulation and does not take into account the aspirational emergent values shown in Figure 7.6. The single-objective method employing the emergent limit values reaches a point beyond which it fails to generate any new sufficient designs. Figure 7.8, which tracks the progression of the desirable features during searches conducted using the MDRTS approach and that employing the emergent limit values, shows that this is due to the latter improving α_{exit} at the expense of M_{exit} , violating even the higher standard constraint limit to a point from which it is unable to return using the available computational budget.

This highlights a limitation of penalty methods when used in conjunction with aspirational limit values, as large violations of one constraint can be accepted in order to make progress towards satisfying others. Alternative constraint handling techniques are not viable for complex problems such as axial compressor optimisation in which the ability of the optimiser to generate new designs is already severely restricted by the fragmented nature of the design space [82, 137]. The more relaxed approach adopted by the MDRTS algorithm informs the optimiser of the desire for these quantities to be minimised without applying additional restrictions to the movement of the optimiser. This reduces the risk of large detriments in performance being accepted in order to improve one parameter as more aspects of the design are included in the decision-making process.

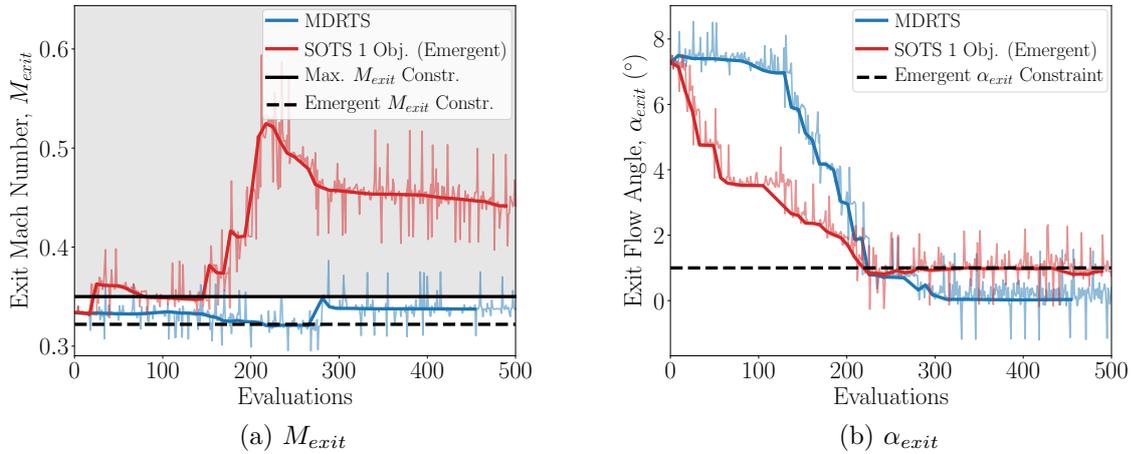


Fig. 7.8 Progression of M_{exit} and α_{exit} during searches using the MDR formulation and the emergent constraint limits in Table 7.2 applied to the six-stage axial compressor test case in the single-objective scenario. Performance of all designs is plotted as a faint line with that of the pattern search centres highlighted in bold.

7.4 Trade-Off Scenario

Results in the previous section demonstrate the capabilities of the MDRTS algorithm when applied to the preliminary design optimisation of an axial compressor in a single-objective scenario. In this section a trade-off study is conducted to determine whether similar performance can be expected when employing the new method for this alternative purpose.

7.4.1 Problem Formulations

In the trade-off scenario SM is treated as a second objective alongside η_p , aligning with a number of multi-objective studies conducted in the literature [79–83, 99, 193]. The constraints on PR , DF and DH are retained from the single-objective scenario, as is the lower limit on SM to ensure the optimiser focusses on designs with sufficient stability that are more likely to be of interest to the designer.

As in the previous section the MDRTS approach treats M_{exit} and α_{exit} as desirable features, with this formulation compared to two alternative multi-objective methods. The first treats the exit conditions as additional objectives alongside η_p and SM , resulting in the optimiser attempting to resolve a four-objective trade-off. The second simply applies the upper limit constraints to M_{exit} and α_{exit} , requiring them to be less than 0.35 and 10° respectively. These maximum thresholds are also applied as constraints in the four-objective and MDR formulations to promote acceptable final values. As discussed in Section 7.3.1 this sensibly incorporates the desires of the designer expressed by Walsh and Fletcher [238].

The resultant formulations, including the constraint values used, are outlined in Table 7.3 alongside an additional two-objective approach employing emergent constraint limits that

Table 7.3 Formulations applied to the six-stage axial compressor test case in the trade-off scenario.

	MDR	Four-Objective	Two-Objective (Standard)	Two-Objective (Emergent)
Objectives	η_p, SM	$\eta_p, SM, M_{exit}, \alpha_{exit}$	η_p, SM	η_p, SM
Desirable Features	M_{exit}, α_{exit}	-	-	-
Constraints	$PR \geq 2.75$ $DF \leq 0.6$ $DH \geq 0.72$ $SM \geq 15.0\%$ $M_{exit} \leq 0.35$ $\alpha_{exit} \leq 10.0^\circ$	$PR \geq 2.75$ $DF \leq 0.6$ $DH \geq 0.72$ $SM \geq 15.0\%$ $M_{exit} \leq 0.35$ $\alpha_{exit} \leq 10.0^\circ$	$PR \geq 2.75$ $DF \leq 0.6$ $DH \geq 0.72$ $SM \geq 15.0\%$ $M_{exit} \leq 0.35$ $\alpha_{exit} \leq 10.0^\circ$	$PR \geq 2.75$ $DF \leq 0.6$ $DH \geq 0.72$ $SM \geq 15.0\%$ $M_{exit} \leq 0.327$ $\alpha_{exit} \leq 4.2^\circ$

is the subject of Section 7.4.3. Constraints are again applied using a penalty method, and the same set of geometric limits are used to ensure realistic final compressors are produced.

7.4.2 Results

Figure 7.9 shows the performance of designs that form the η_p - SM Pareto fronts generated by the different formulations when using a computational budget of 500 Multall evaluations, just considering designs that satisfy the MDR constraints. The two-objective approach employing standard limit values, representing the current state-of-the-art, produces the most advanced and widespread Pareto front. However, the designs that make up this trade-off curve exhibit poor performance in terms of M_{exit} and α_{exit} . Treating these parameters solely using constraints allows the optimiser to accept high values, with no attempt made to minimise them “to prevent excessive downstream pressure loss” [238]. The four-objective approach, treating the exit conditions as additional objectives alongside η_p and SM , generates a number of designs exhibiting better performance in terms of the secondary quantities. However, this comes at the expense of η_p and SM , a result which is perhaps unsurprising given that the exit conditions being treated as objectives in this formulation are known to be negatively correlated with these parameters [81]. Treating M_{exit} and α_{exit} as desirable features in the new MDRTS algorithm results in designs that perform well in terms of all four quantities of interest. The η_p - SM Pareto front is close to that generated using the two-objective formulation and the constituent designs exhibit lower values of both M_{exit} and α_{exit} . The MDR formulation appears to enable the optimiser to focus on a region of the design space that is more likely to be of interest to the designer, resulting in a narrower

Initial Application to a Six-Stage Axial Compressor

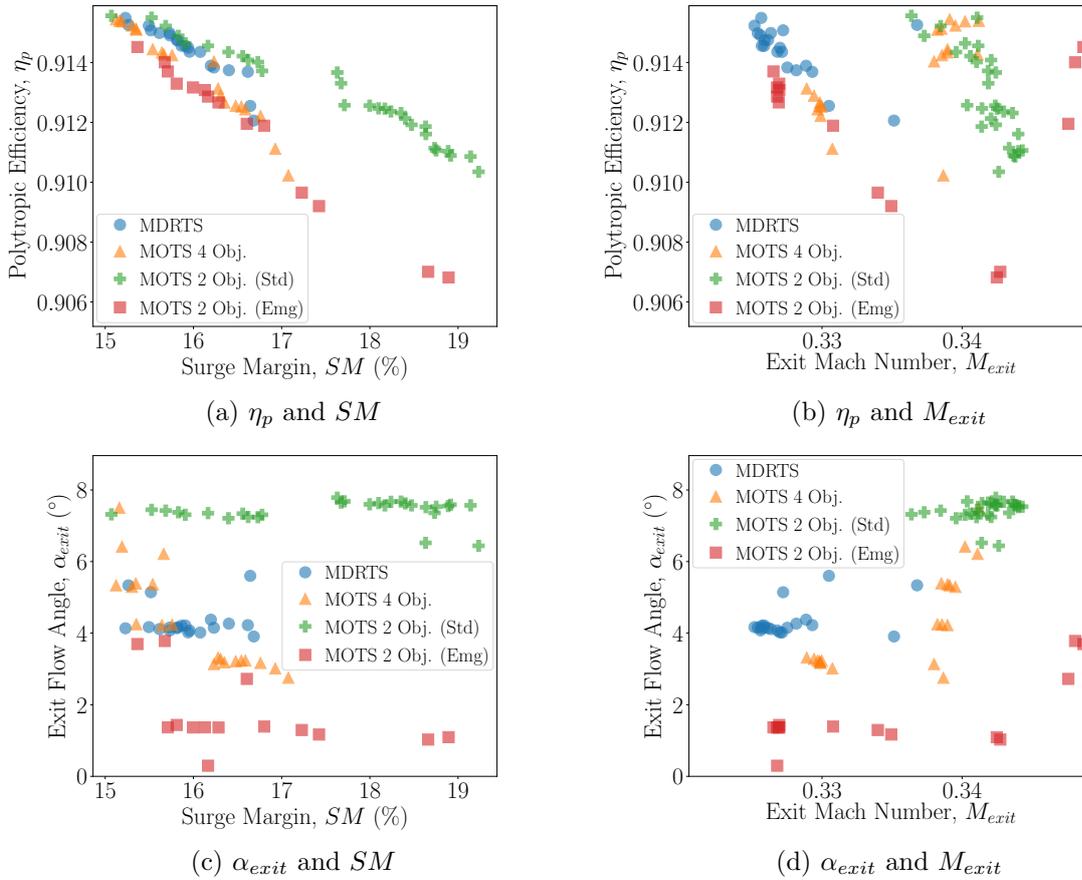


Fig. 7.9 Performance of the best designs found by the formulations in Table 7.3 applied to the six-stage axial compressor test case in the trade-off scenario.

Pareto front than that produced using the traditional approaches and more efficient use of the available computational budget.

Figure 7.10 tracks the best values of the overall performance measure found by each method, with this again defined as the reciprocal of the shortest Euclidean distance to an idealised reference point given the best performance in terms of η_p , SM , M_{exit} and α_{exit} , calculated in a space containing normalised versions of these four quantities. The MDRTS method generates designs with good overall performance using fewer Multall evaluations than the other methods. Values equal to the best found using the two-objective formulation applying standard constraints are produced with a computational saving of over 80%. The new algorithm also outperforms the four-objective approach in terms of this overall performance measure up until the last 10% of the search, with the improved performance of the alternative method probably a result of the lower values of α_{exit} visible in Figure 7.9.

Figure 7.11, showing the number of designs exhibiting performance within a given Euclidean distance of the reference point, reveals that the new approach generates a larger number of designs with good values of all four quantities of interest than the traditional

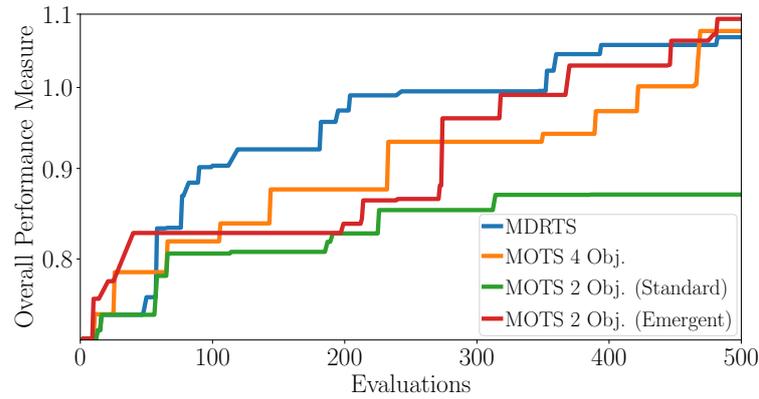


Fig. 7.10 Progression of the best overall performance measure found by the formulations in Table 7.3 applied to the six-stage axial compressor test case in the trade-off scenario.

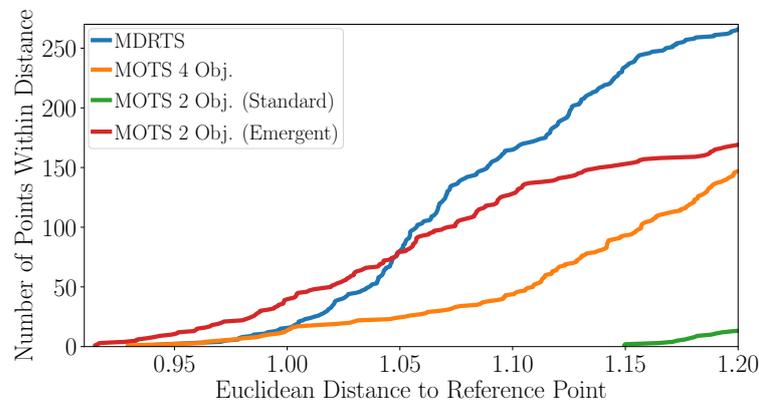


Fig. 7.11 Number of designs produced by the formulations in Table 7.3 applied to the six-stage axial compressor test case in the trade-off scenario that exhibit performance within a given Euclidean distance of the reference point.

methods. The four-objective approach also appears to have produced plenty of interesting designs, however, viewed in conjunction with Figure 7.9 this good performance is seen to be weighted towards the desirable features, with the η_p -SM Pareto front generated using this method almost entirely dominated by those produced by the alternative approaches. The MDRTS algorithm, in contrast, achieves good performance in terms of the desirable features whilst maintaining the focus of the optimiser on advancing the η_p -SM Pareto front.

The MDR formulation has been successfully applied to this trade-off problem, overcoming the limitations of current methods through more accurate representation of the desires of the designer for the performance parameters being considered. Using the same computational budget the new approach generates designs that are more likely to be of interest to the designer, with good performance in terms of the secondary parameters as well as the main objectives.

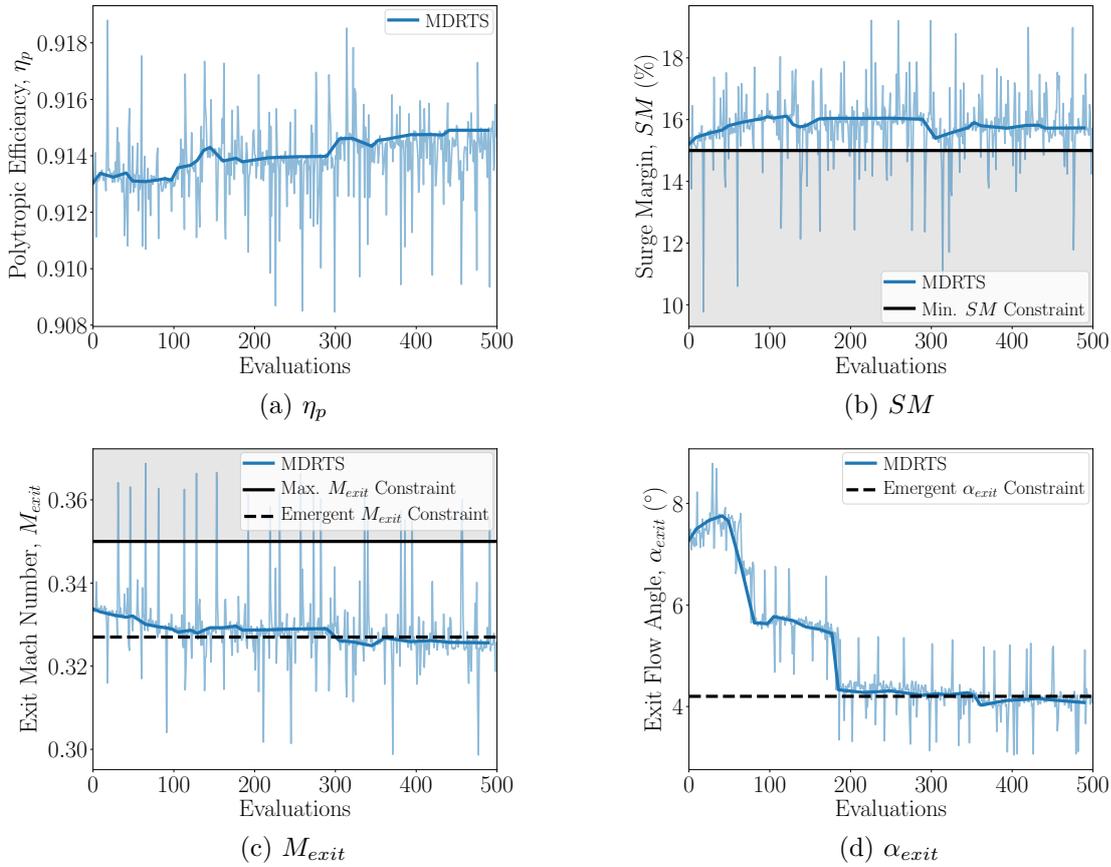


Fig. 7.12 Progression of the objectives and desirable features during the search using the MDR formulation in Table 7.3 applied to the six-stage axial compressor test case in the trade-off scenario. Performance of all designs is plotted as a faint line with that of the pattern search centres highlighted in bold.

7.4.3 Emergent Constraints

As in the single-objective scenario and the aerofoil test case in the previous chapter, limit values for the performance parameters classified as desirable features emerge as outputs of the MDRTS optimisation. This can be seen in Figure 7.12 which tracks the key quantities of interest for all of the designs generated during the search. Whilst η_p and SM continue to vary, the desirable features converge as the optimisation progresses. The emergent limit values can be used to inform an alternative two-objective constrained optimisation using aspirational values for the M_{exit} and α_{exit} constraints. Suitable limits of 0.327 and 4.2° are selected using Figures 7.12c and 7.12d respectively.

The goal of applying this emergent constraints formulation, outlined in Table 7.3, is to determine whether the same portion of design space can be accessed using a traditional constrained approach with these lower, aspirational limits as has been found by the MDRTS algorithm. The results in Figures 7.9, 7.10 and 7.11 show that this is not possible. Although

the formulation applying the emergent constraints produces good designs in terms of M_{exit} and α_{exit} it fails to achieve the same advancement of the η_p - SM Pareto front demonstrated by the methods employing standard constraint limits and the MDR formulation. A number of the best designs also fail to satisfy the specified emergent constraints.

Figure 7.10, tracking the overall performance measure, shows that the method using emergent constraints manages to keep pace with the other formulations, and the tally of the number of designs exhibiting performance within a given distance of the reference point in Figure 7.11 suggests that it also produces the largest number of interesting designs. However, the performance of these machines is heavily weighted towards the desirable features, with less improvement seen in the main objectives of η_p and SM .

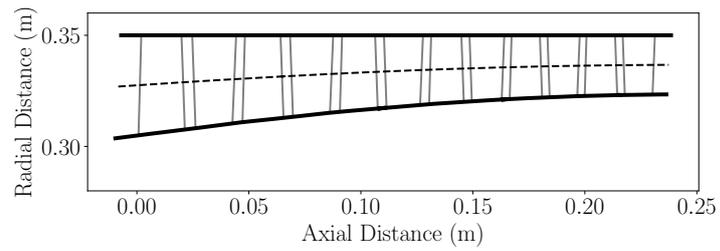
As in the single-objective scenario, the optimiser using this constrained formulation is hampered by an inability to distinguish between the importance of different constraints. The penalty approach results in reductions in η_p and SM being accepted as the optimiser pursues satisfaction of the aspirational constraints applied to the exit conditions. This leads to wasted computational effort, with fewer sufficient designs generated and a relatively poor η_p - SM Pareto front being produced. The new MDRTS algorithm allows for a distinction between requirements that must be satisfied for a design to be considered acceptable and parameters that the designer wants to be minimised if possible but with secondary importance compared to the main objectives. This gives the optimiser more freedom as it moves through the design space, ultimately producing a larger number of designs that are likely to be of interest to the designer and making better use of the available computational budget.

7.5 Information to Enhance Designer Understanding

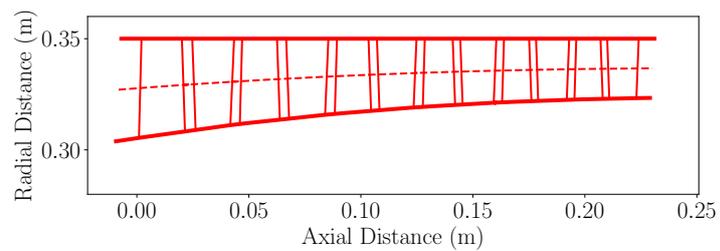
The results presented so far in this chapter demonstrate the new MDRTS algorithm successfully overcoming the formulation limitations that hamper existing optimisation methodologies when applied to the preliminary design of axial compressors. This section highlights the information provided by the optimiser to enhance designer understanding, addressing the second limitation discussed in Chapter 3.

The pattern search used to generate new designs results in rich development information being available for all of the compressors generated using the MDRTS approach. To demonstrate this, information for the design exhibiting the highest value of η_p produced when the new algorithm is applied in the single-objective scenario is presented. The annulus of this machine is shown in Figure 7.13 alongside the initial geometry for comparative purposes.

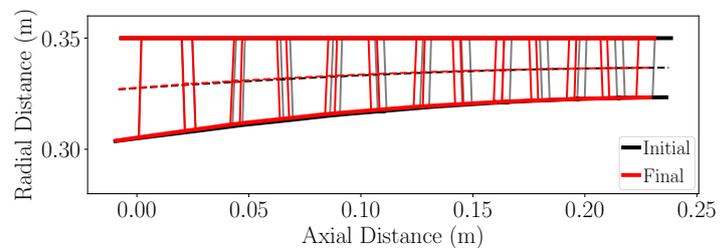
This axial compressor problem is more complex than the aerofoil test case used in the previous chapter, resulting in a large amount of information concerning the development of a particular design. Whilst this provides additional opportunities for enhancing designer understanding, it also makes presentation challenging. One way to view the development



(a) Initial Design



(b) Final Design



(c) Initial and Final Designs

Fig. 7.13 Annuli of the initial design and that exhibiting the highest value of η_p found by the MDRTS algorithm applied to the six-stage axial compressor test case in the single-objective scenario.

7.5 Information to Enhance Designer Understanding

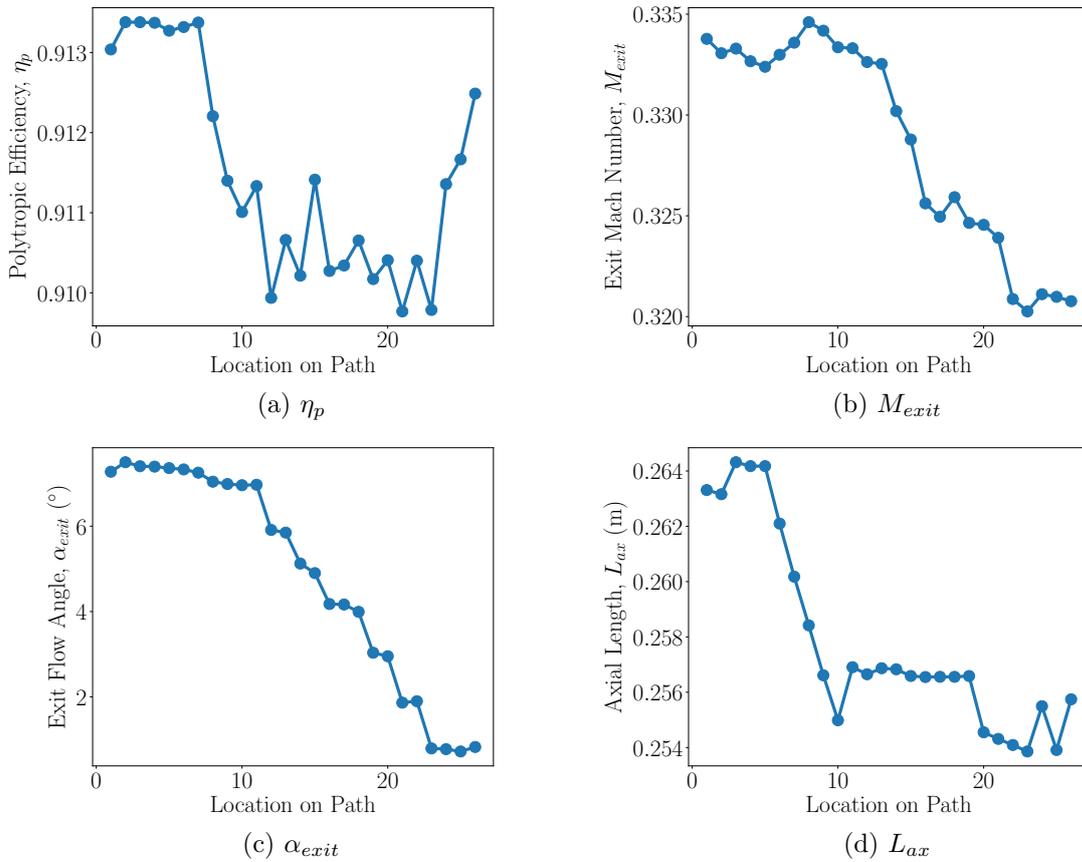


Fig. 7.14 Changes to overall performance parameters during development of the design exhibiting the highest value of η_p found by the MDRTS algorithm applied to the six-stage axial compressor test case in the single-objective scenario.

history is in Table 7.4, where the changes made by the optimiser to generate the final design are tracked alongside the impact each has on the key performance parameters. The right-hand columns inform the designer that the optimiser has primarily improved the desirable features, rather than η_p . This can also be seen in Figure 7.14 tracking the overall performance parameters along the path. The main objective only begins to receive attention in the latter stages once the desirable features have settled on their emergent values, perhaps suggesting that further improvements could be achieved by granting the optimiser additional computational budget to continue the search. Table 7.4 also informs the designer that the main variables altered by the optimiser are Bézier control points effecting the stage exit flow angles and the blade aspect ratios. The former are reduced in order to achieve lower values of α_{exit} , whilst increasing the latter results in a reduction in the axial length, L_{ax} , of the compressor (Figure 7.14d).

Figure 7.15 tracks the variation of the annulus area and a number of stage performance parameters along the development path. This includes the Mach number and flow angle at

Initial Application to a Six-Stage Axial Compressor

Table 7.4 Variable changes and their impact on the key performance parameters during development of the design exhibiting the highest value of η_p found by the MDRTS algorithm applied to the six-stage axial compressor test case in the single-objective scenario. Green indicates improvement, red indicates worsening. CP refers to a Bézier curve control point and AR to the blade aspect ratio.

Step	Variable(s) Changed	Change Made	η_p	M_{exit}	α_{exit} (°)
1	-	-	0.913039	0.3338	7.27
2	2nd and 3rd Stage α_{exit} CPs	Decreased	0.913380	0.3331	7.50
3	3rd Stator AR CP	Decreased	0.913379	0.3333	7.40
4	2nd and 3rd Stage α_{exit} CPs	Decreased	0.913372	0.3327	7.39
5	Blade Twist	Decreased	0.913273	0.3324	7.36
6	4th Stator AR CP	Increased	0.913319	0.3330	7.33
7	4th Stator AR CP	Increased	0.913375	0.3336	7.25
8	2nd Mean Radius CP	Increased	0.912206	0.3346	7.04
9	2nd Mean Radius CP	Increased	0.911400	0.3342	6.99
10	1st Mean Radius CP	Increased	0.911009	0.3334	6.96
11	4th Stator AR CP	Decreased	0.911333	0.3333	6.97
12	All Stage α_{exit} CPs	Decreased	0.909937	0.3326	5.91
13	1st and 2nd Stage α_{exit} CPs	Increased	0.910660	0.3325	5.85
14	3rd and 4th Stage α_{exit} CPs	Decreased	0.910214	0.3302	5.12
15	2nd Stage PR CP	Increased	0.911413	0.3288	4.90
16	3rd and 4th Stage α_{exit} CPs	Decreased	0.910273	0.3256	4.18
17	2nd N° Rotor Blades CP	Increased	0.910341	0.3250	4.16
18	4th N° Stator Blades CP	Decreased	0.910654	0.3259	3.99
19	4th Stage α_{exit} CP	Decreased	0.910172	0.3247	3.03
20	4th Rotor AR CP	Increased	0.910407	0.3246	2.95
21	All Stage α_{exit} CPs	Decreased	0.909768	0.3239	1.87
22	2nd Stage PR CP	Increased	0.910401	0.3209	1.90
23	All Stage α_{exit} CPs	Decreased	0.909786	0.3203	0.79
24	1st Mean Radius CP	Decreased	0.911357	0.3211	0.78
25	1st Stator AR CP	Increased	0.911667	0.3210	0.72
26	2nd Mean Radius CP	Decreased	0.912487	0.3208	0.83

the exit of each stage as well as values of ψ , ϕ and Λ , three quantities defined in Section 7.1.4. Figure 7.15d shows that whilst the annulus area does vary along the development path the final values are actually similar to the initial design. A slight increase in ψ can be seen in all stages, with the additional work reducing the value of M_{exit} . The early stages of the machine show a rise in ϕ , corresponding to the reduction in annulus area observed in Figure 7.15d. However, this rise is negated towards the end of the development path as the value of η_p begins to increase.

The parameter exhibiting the largest variation is Λ , shown in Figure 7.15b, with moves to higher values placing a greater pressure-rise requirement on the rotors particularly noticeable in the final two stages of the machine. The increased diffusion can be seen in Figure 7.16, which tracks the values of DH and DF as well as the aspect ratios for the different blade rows. The rotor DH values generally decrease and the DF values increase, most notably in stage six, indicating higher loading on these blade rows. Conversely, the values for the stators reveal offloading of these blades as the optimiser seeks to reduce α_{exit} . The increased diffusion requirements in the rotors also correlate with higher aspect ratios in the final two stages.

It would probably be possible to write several pages discussing this development information and the physical phenomena behind the observed performance improvements. However, the focus of this thesis is on facilitating the use of optimisation in the design process, rather than attempting to suggest improvements to axial compressors themselves. More importantly, the open source analysis codes employed, in particular the SM correlation, are not accurate enough to validate that type of discussion. The amount of information available and the brief comments that can be made serve the purpose of demonstrating the ability of the MDRTS algorithm to fulfil the role of enhancing designer understanding. Rich information regarding the development history of final designs is generated to assist designers in justifying the observed performance improvements and increase their knowledge of the underlying problem to promote creativity and innovation. This data is provided alongside the enhanced searching capabilities demonstrated earlier in this chapter and at no additional computational cost.

7.6 Summary

In this chapter an initial application of the MDRTS algorithm to the preliminary design optimisation of an axial compressor has been successfully carried out. The new approach produces designs exhibiting high levels of performance through better handling of the large number of relevant performance parameters. The ability to provide rich design development information has also been demonstrated, equipping the algorithm for the role of enhancing designer understanding. This has all been achieved using open source analysis methods, representing a significant departure from the norm in turbomachinery optimisation research.

Initial Application to a Six-Stage Axial Compressor

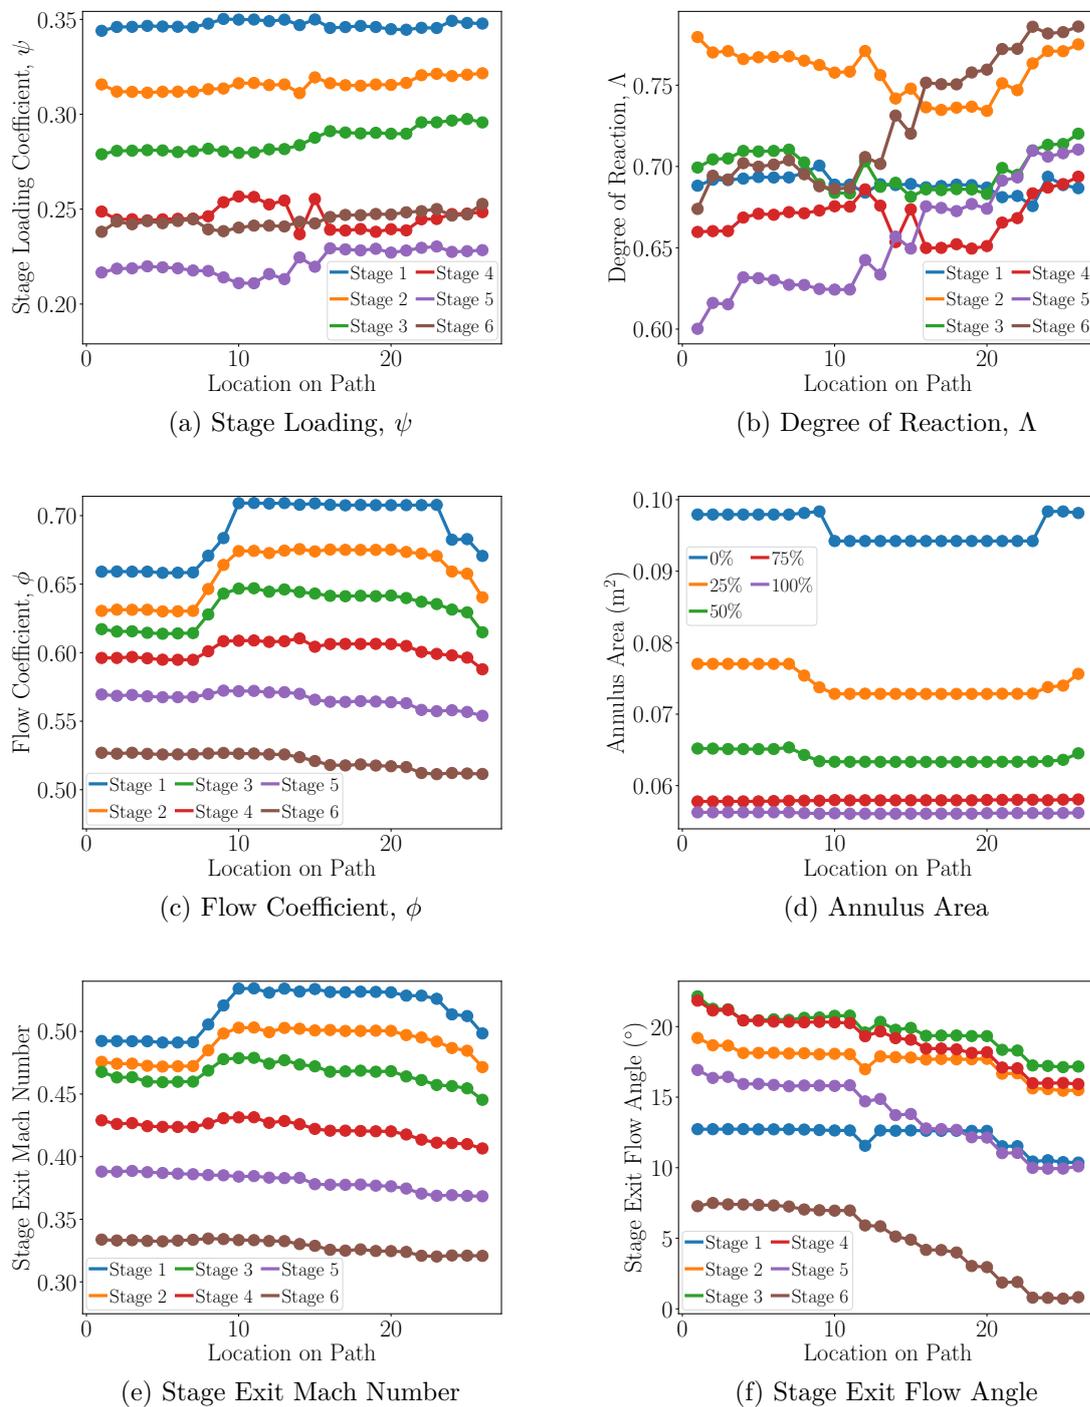


Fig. 7.15 Changes to stage parameters during development of the design exhibiting the highest value of η_p found by the MDRTS algorithm applied to the six-stage axial compressor test case in the single-objective scenario.

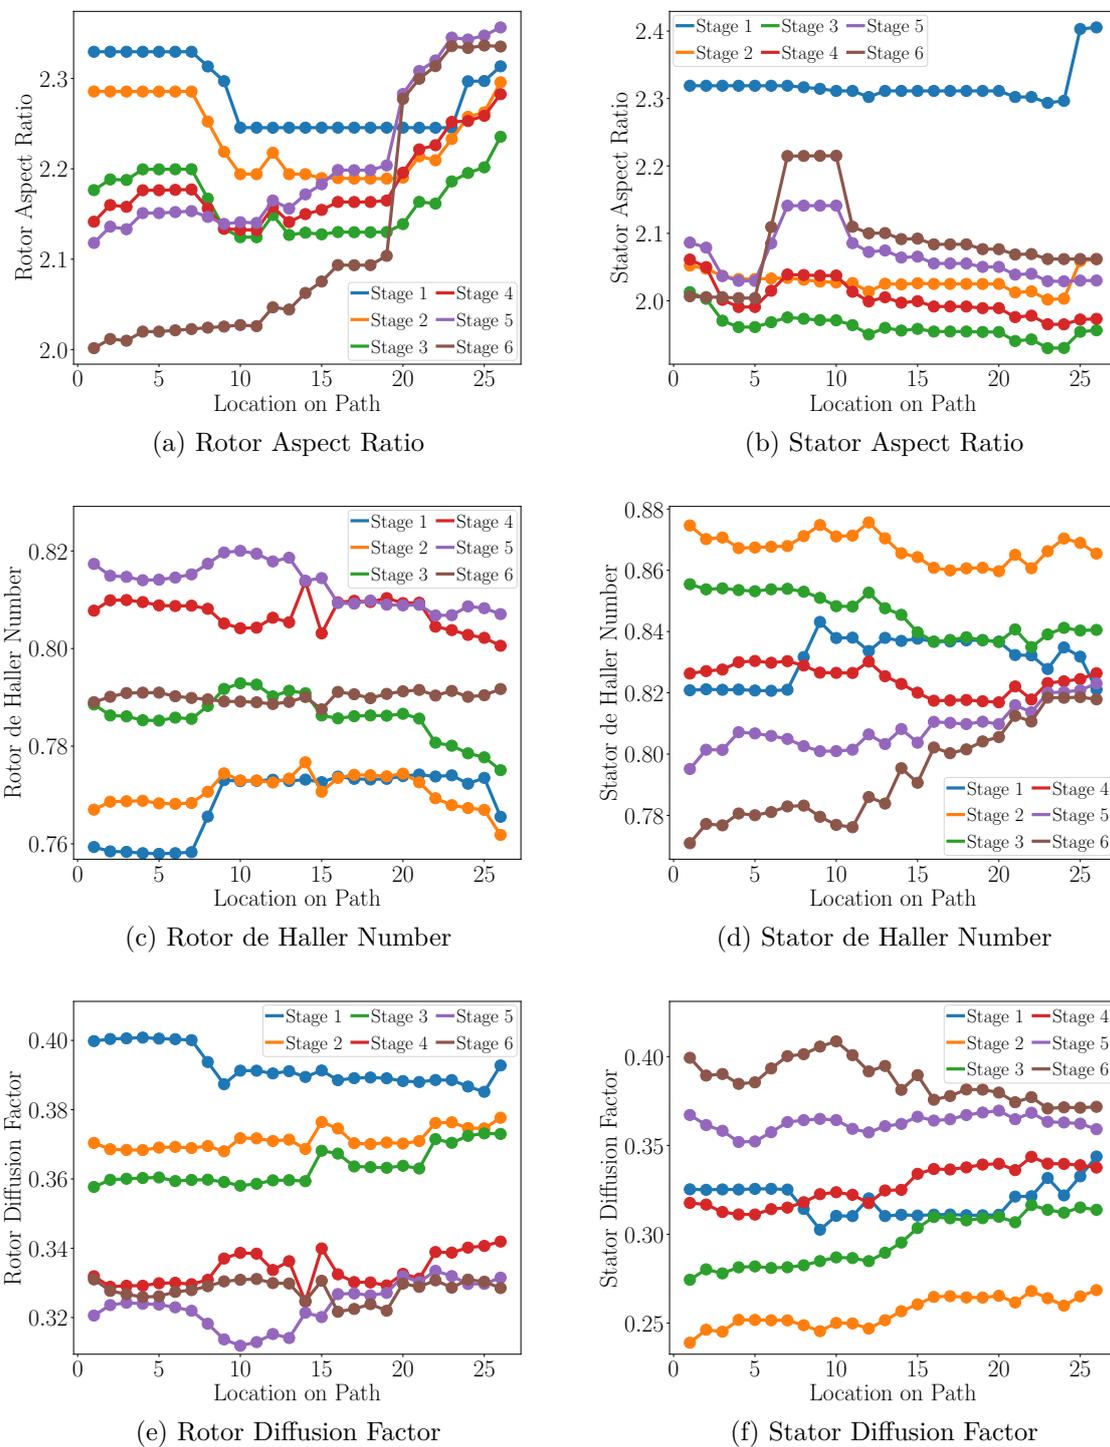


Fig. 7.16 Changes to blade parameters during development of the design exhibiting the highest value of η_p found by the MDRTS algorithm applied to the six-stage axial compressor test case in the single-objective scenario.

Initial Application to a Six-Stage Axial Compressor

The MDRTS algorithm developed in this thesis has been shown to overcome two of the three problems that limit existing optimisation methodologies when applied to the preliminary design of axial compressors. All that remains is to accelerate the approach using multi-fidelity methods, which is the subject of the following chapter.

Chapter 8

Multi-Fidelity Acceleration

In the preceding chapters a new methodology has been developed to facilitate the use of optimisation in the preliminary design of axial compressors. The MDRTS algorithm overcomes problems associated with formulation and understanding that limit existing approaches by enabling more accurate handling of the relevant performance parameters and providing development information to assist with the physical justification of final designs.

However, in engineering design time is money [53, 92, 115, 129, 205] and for any optimisation algorithm to be useful in an industrial setting it needs to be fast, making efficient use of the available computational budget [214]. In this chapter the MDRTS methodology is accelerated to meet this goal and overcome the final limitation discussed in Chapter 3. This is achieved using multi-fidelity techniques, with computationally cheap analysis codes employed to speed-up the search. The following sections detail the incorporation of different analysis fidelities within the MDRTS algorithm and assess the performance of the resultant approach using analytic, aerofoil and six-stage axial compressor test cases.

8.1 Incorporating Multiple Fidelities Within Optimisation Using MDR

The nested hierarchy of dominance relations at the heart of the MDR formulation, shown in Equation 3.4, lends itself naturally to a multi-fidelity approach. In the single-fidelity format each dominance relation compares designs using different performance parameters. In a multi-fidelity context this can be extended to allow dominance relations to use different analysis codes to compare designs. For example, the first dominance relation, \preceq_1 , could compare designs using a low-fidelity model, with the second, \preceq_2 , using high-fidelity analysis. The optimiser would then be seeking good designs according to the high-fidelity (the ultimate goal of the optimisation) within the reduced set that perform well according to the low-fidelity. This corresponds to a multi-level multi-fidelity method [189], with the low-fidelity performing

an initial sift of potential designs and only the most promising progressing for assessment using the more computationally expensive high-fidelity code. Any number of fidelity levels could be introduced in this way, and there is no requirement for the same performance parameters to be used in each dominance relation. This could allow different analysis codes to be used to assess different aspects of a design, a potentially useful feature in scenarios where a low-fidelity model provides accurate approximations of certain performance parameters but not others [229].

8.1.1 Analysis on a “Need-to-Know” Basis

This concept for including information from different analysis codes in a MDR formulation does not in itself provide any acceleration of the search. In the current algorithm, depicted in Figure 4.1, analysis occurs immediately after a candidate is selected, meaning that all designs have the information required by each of the different dominance relations. In a multi-fidelity context this procedure cannot provide any search acceleration, as all of the designs are analysed using all of the fidelity levels irrespective of their performance. Instead, analysis should be conducted on a “need-to-know” basis, with a particular fidelity only applied when that level of information is required to select between designs. If a design can be ruled out using a low-fidelity model then there should be no need to analyse it using more accurate and computationally expensive methods.

To achieve this goal analysis is moved inside the function used to find non-dominated designs. A flowchart for the MFMDRTS (Multi-Fidelity MDRTS) algorithm facilitating analysis on a “need-to-know” basis is shown in Figure 8.1, with a detailed view of the new selection procedure depicted in Figure 8.2. When selecting between designs the new function takes a set of input points and starts by determining whether any have yet to be analysed using the fidelity required by the first dominance relation. The appropriate analysis is carried out followed by an update to the rankings associated with this dominance relation to reflect the new information (see Section 4.2). The input designs are then sorted based on their rankings for this dominance relation, with only those that are non-dominated progressing to be considered using the next dominance relation. The algorithm again checks whether the remaining designs have the necessary information to be assessed using the new dominance relation and carries out analysis of any that need it. In this way the optimiser finds non-dominated designs whilst only conducting analysis when that level of information is required to make a decision. This process is efficient, incurring no additional computational overhead compared to the procedures for calculating rankings and finding non-dominated designs in the single-fidelity algorithm. It also enables the use of different fidelities to provide speed-up, as designs that are discounted at an early dominance relation using low-fidelity information need never be analysed at a more computationally expensive higher-fidelity level.

8.1 Incorporating Multiple Fidelities Within Optimisation Using MDR

All types of low-fidelity model can be accommodated for this purpose, including those based on coarsened grids, relaxed convergence criteria and reduced physics models [69].

Similar multi-level approaches have been applied to the optimisation of axial compressors by Kipouros et al. [137] and Poehlmann and Bestle [193]. The former altered the shape of blades using MOTS, accelerating the procedure by increasing the number of iterations carried out by the CFD solver every 200 optimisation steps. The latter attempted sequential optimisations, with the outputs of an initial search using meanline analysis acting as inputs for a more expensive throughflow stage. Some success was reported in both cases, however the authors also highlighted the need for feedback between the different levels. Without this feedback results generated by the low-fidelity model can mask potentially promising regions of the high-fidelity design space, and may continually drive the optimiser towards designs that perform well according to the low-fidelity but poorly when analysed using a more accurate model [160].

8.1.2 Avoiding Low-Fidelity Masking

To avoid masking in multi-level methods results from the higher-fidelity levels need to be used to inform less accurate models of the “true” performance of designs sent for high-fidelity analysis. One way this can be achieved is through additive, multiplicative or hybrid corrections applied to the low-fidelity codes, as shown in Equation 8.1 [69].

$$\begin{aligned} \text{additive} \quad y_{corr}(\mathbf{x}) &= y_{LF}(\mathbf{x}) + \delta(\mathbf{x}) \\ \text{multiplicative} \quad y_{corr}(\mathbf{x}) &= \mu(\mathbf{x})y_{LF}(\mathbf{x}) \\ \text{hybrid} \quad y_{corr}(\mathbf{x}) &= \mu(\mathbf{x})y_{LF}(\mathbf{x}) + \delta(\mathbf{x}) \end{aligned} \tag{8.1}$$

A response surface built using high-fidelity data can be fitted to the correction parameters (δ for additive, μ for multiplicative, and both for hybrid) and applied to the output of the low-fidelity code at that point, y_{LF} , to generate an improved performance prediction, y_{corr} . This approach, used by the likes of Jarrett and Ghisu [113] and Rumpfkeil and Beran [207], reduces the risk of masking as the corrected low-fidelity code reflects the “true” performance with increasing accuracy as the optimisation progresses and more high-fidelity training data becomes available.

To facilitate the use of corrective feedback a number of additions are made to the procedure for finding non-dominated designs. Firstly, the optimiser is forced to cycle through every dominance relation even if only one design remains. This ensures that any design considered to be non-dominated is assessed using all of the different analysis codes, preventing the optimisation from converging without ever carrying out high-fidelity analysis of the design it considers to be optimal. It also means that high-fidelity information is available for all

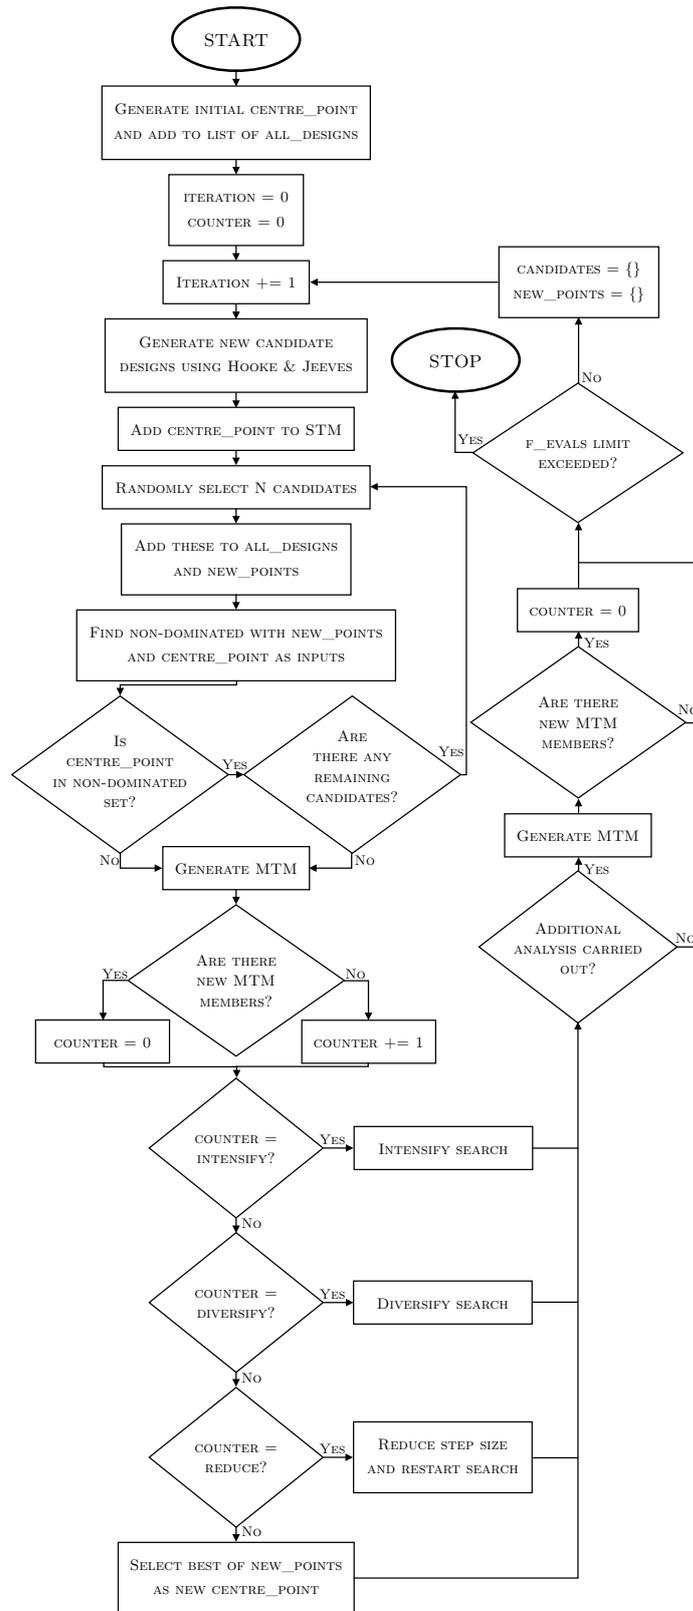


Fig. 8.1 Flowchart for the MFMDRTS algorithm.

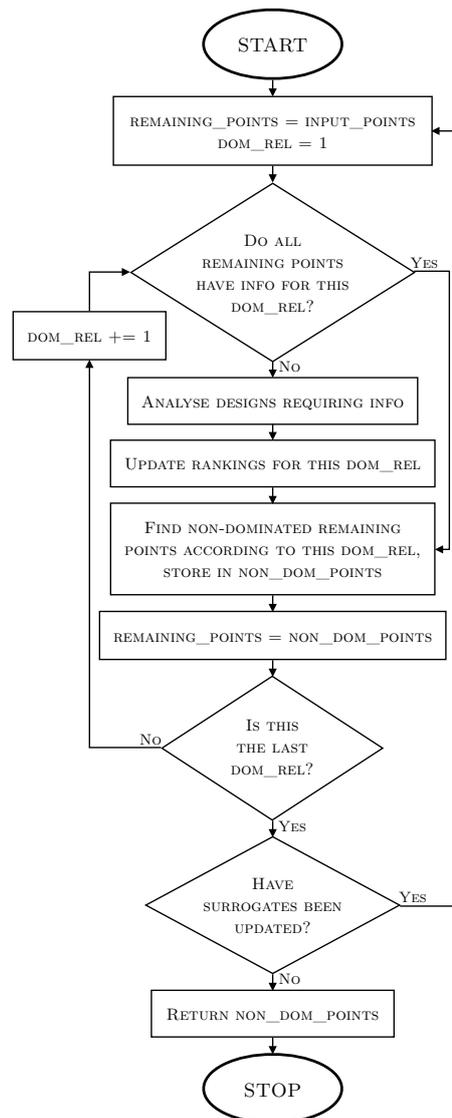


Fig. 8.2 The procedure used by the MFMDRTS algorithm to find non-dominated designs.

designs selected as pattern search centres, resulting in the provision of the same amount of design development data to enhance understanding as the single-fidelity algorithm.

Secondly, the optimiser is instructed to re-find the non-dominated set if new training data is generated for any response surfaces. Designs predicted to perform well by the low-fidelity will eventually be analysed using the high-fidelity code due to the action of the first masking-avoidance feature. This produces additional training data for the response surfaces used to predict the performance of these promising points, potentially improving their accuracy. Re-finding the non-dominated designs using this updated training set ensures that the most accurate information has been used in the decision-making process. This is particularly important when a design predicted to perform well by the corrected low-fidelity model actually performs poorly when analysed using the high-fidelity code. Without the requirement to re-find the set of non-dominated designs this point would be considered non-dominated despite poor high-fidelity performance. Repeating the process using up-to-date response surfaces allows the poor performance to be taken into account, altering the designs that progress to the high-fidelity dominance relations. The feedback loop also reduces the likelihood of promising designs being missed.

These requirements to always consider every dominance relation and to re-find the non-dominated set until the most up-to-date response surfaces have been employed are visible in Figure 8.2. They should reduce the risk of low-fidelity masking and enable the use of surrogate-based correction techniques within the MFMDRTS algorithm. Data-fit response surfaces can also be incorporated as additive correction approaches where the low-fidelity model, y_{LF} , outputs zero at all points.

One final change is in the calculation of rankings for correction-based low-fidelity models. The rankings for dominance relations employing analysis codes without correction are calculated by counting the number of other designs that dominate a particular design according to that dominance relation. When using correction-based low-fidelity approaches it is more important to compare the output to existing high-fidelity data, rather than to other corrected performance predictions that may be erroneous due to inaccuracies in the response surface fitted to the correction parameter. Therefore, for dominance relations that use correction-based analysis the rankings are calculated as the number of points with high-fidelity data (the data that is being corrected to) that dominate the predicted performance of a given design. This method should minimise the impact of inaccuracies in correction response surfaces as comparisons are made between low- and high-fidelity data rather than between the low-fidelity values themselves. Coupled with the alterations to the decision procedure, this new ranking scheme should allow the algorithm to effectively handle all types of low-fidelity model without masking promising regions of the design space.

8.2 Similarities With Existing Multi-Level Approaches

The MFMDRTS algorithm developed in the previous section has similarities with a number of alternative multi-level methodologies in the literature. These are reviewed in this section to give the new approach context as a technique for incorporating multiple fidelities into an optimisation routine.

8.2.1 Sequential Approaches

The simplest multi-level techniques are sequential approaches in which the output of a low-fidelity optimisation is used as the input to a high-fidelity one [139, 145, 193, 218, 226]. The main problem facing these methods, highlighted in the previous section and by Liu et al. [160], is that of low-fidelity masking. Promising regions often fail to match up between the low- and high-fidelity design spaces, resulting in the low-fidelity optimisation producing poor inputs for the high-fidelity stage. In the most extreme cases, experienced by Poehlmann and Bestle [193], the outputs of the low-fidelity optimisation may even fail to converge when analysed using the high-fidelity code.

Liu et al. [160] attempted to circumvent this problem by using data-mining approaches to generate input designs for the high-fidelity optimisation from the low-fidelity results. El-Beltagy and Keane [62] opted to blend the two optimisations by assigning probabilities to the fidelity level employed for analysis, with use of the high-fidelity becoming increasingly likely as the search progresses. The approach of Liu et al. showed promise, despite a complex implementation, but the mixed methods of El-Beltagy and Keane provided minimal improvement over the basic sequential strategy.

Rather than conducting two independent optimisations in series, the new MFMDRTS approach carries out a single search. At all times the design space being considered is that of the high-fidelity code, with the low-fidelity used to accelerate search within this space. The problem of masking should be counteracted by the techniques described in Section 8.1.2.

8.2.2 Inexact Pre-Evaluation

An existing method that is perhaps more similar to the new algorithm is inexact pre-evaluation (IPE). This technique, originally proposed by Giannakoglou et al. [86], reduces the number of high-fidelity analyses required by first assessing performance using a data-fit response surface. Only designs that are predicted to perform well are nominated for high-fidelity analysis, reducing the overall computational cost of the algorithm. Whilst the original work used RBF response surfaces to accelerate a GA, the approach has since been extended to include Kriging [65] and PSO [195] as well as being applied by several other authors [59, 127, 246].

When used in conjunction with data-fit low-fidelity models the MFMDRTS algorithm closely resembles IPE. Designs must first pass through any dominance relations based on the

low-fidelity response surfaces before being considered for high-fidelity analysis, just as when using IPE. However, the new method is more flexible as it accommodates not only the use of several different types of low-fidelity model but also the more sophisticated MDR formulation that could allow additional parameters to be assessed using the high-fidelity that are not approximated by a response surface, something that is not possible when using IPE.

8.2.3 Simultaneous Search Methods

Another popular multi-level approach is to conduct separate simultaneous low- and high-fidelity searches, with promising designs passed between the two. Examples of this approach include the injection-island GA developed by Eby et al. [61] and that of Kampolis and Giannakoglou [123]. These methods not only provide speed-up compared to single-fidelity optimisation but also enable different objectives, parameterisations and search algorithms to be employed by the different levels. Whilst the new multi-fidelity method developed in this chapter is unable to accommodate different search algorithms or parameterisations, it does allow for different performance parameters to be used by different analysis codes. This is achieved without the added complexity of conducting separate simultaneous optimisations and also removes the need for designers to specify the proportion of the computational budget assigned to each fidelity.

8.2.4 Different Performance Parameters for Different Fidelity Levels

The ability to accept different performance parameters for different fidelity levels is similar to the approach developed by Bahrami et al. [11] in which the user sets a series of targets for the low-fidelity optimisation. The results of this search are analysed using a high-fidelity code, with the outputs informing either an update of the targets or an advancement to a more complex parameterisation. Whilst the current form of the new algorithm cannot accommodate variations in parameterisation, it does have the ability to assign different targets and performance parameters in the low- and high-fidelity dominance relations.

8.2.5 Other Approaches Conducting Analysis on a “Need-to-Know” Basis

Approaches developed by Li et al. [152] and Shu et al. [215] employ a similar philosophy of conducting analysis on a “need-to-know” basis. These utilise the uncertainty property of Kriging to determine the best and worst possible performance of designs and only analyse those whose dominance status would change if performance was taken to these extreme values. This results in efficient use of the available computational budget as analysis is only carried out when additional information is required for a decision to be made. It is expected that

adopting the philosophy of conducting analysis on a “need-to-know” basis will lead to similar search accelerations in the new MFMDRTS algorithm.

8.2.6 A Distinct and Flexible Multi-Fidelity Methodology

Whilst the MFMDRTS approach has similarities with a number of previously developed multi-level methods, it also has features that set it apart as a novel technique for incorporating codes of differing fidelity into an optimisation. The inherent flexibility allows several models of any type to be included, with no requirement for the same performance parameters to be used across the different levels. The feedback loop from the high- to low-fidelity levels should guard against low-fidelity masking, a problem which has limited previous multi-level approaches.

8.3 Application to Test Cases

To assess the effectiveness of the new MFMDRTS algorithm it is applied to each of the three test cases employed in the preceding chapters. The analytic problem from Chapter 4 is used to verify the philosophy of providing acceleration by conducting analysis on a “need-to-know” basis, as well as to compare two different response surface construction techniques. The aerofoil test case then enables an initial assessment of the effectiveness of the methodology, with the relatively cheap computational analysis allowing the repeatability of any benefits to be determined. Finally, the MFMDRTS algorithm is applied to the six-stage axial compressor test case introduced in the previous chapter, with a variety of low-fidelity models employed. The results demonstrate whether search using the new algorithm has been accelerated to overcome the third and final problem outlined in Chapter 3 that limits existing optimisation methodologies when applied to the preliminary design of axial compressors.

8.3.1 Analytic Test Case

The analytic test case developed by Cook et al. [39] was used in Chapter 4 to verify the initial TS implementation of the MDR formulation. Here it is used to assess the validity of providing acceleration by conducting analysis on a “need-to-know” basis. As discussed in Section 4.4.1, the problem consists of three performance parameters, f_1 , f_2 and f_3 , defined in Equation 4.1, and four design variables each in the range $[-5, 5]$. The goal is to find the point that minimises f_3 within the Pareto front defined by minimising f_1 and f_2 .

8.3.1.1 Problem Formulations

The single-fidelity MDR formulation for this problem is outlined in Table 8.1, with Pareto dominance of f_1 and f_2 used as the first dominance relation, \preceq_1 , and ordering of f_3 as

Multi-Fidelity Acceleration

Table 8.1 Formulations applied to the multi-fidelity analytic test case.

		Single- Fidelity	Separate f_1/f_2 and f_3	RBF Response Surface	GP Response Surface
\preceq_1	Fidelity Parameters	Direct (all) f_1, f_2	Direct (sep.) f_1, f_2	RBF f_1, f_2	GP f_1, f_2
\preceq_2	Fidelity Parameters	Direct (all) f_3	Direct (sep.) f_3	Direct (sep.) f_1, f_2	Direct (sep.) f_1, f_2
\preceq_3	Fidelity Parameters	-	-	Direct (sep.) f_3	Direct (sep.) f_3

the second, \preceq_2 . Two multi-fidelity variations of this formulation are used to verify the new approach developed in this chapter. Firstly, the calculation of f_1 and f_2 is separated from that of f_3 . Evaluating each of these objectives is assumed to be equally expensive, so retrieving values for f_1 and f_2 costs 0.67 units and calculating f_3 costs 0.33 units. This formulation, entitled “Separate f_1/f_2 and f_3 ” in Table 8.1, assesses the ability to provide acceleration by conducting analysis on a “need-to-know” basis. The number of points for which f_3 is calculated should be reduced, leading to acceleration of the search, as there is no need to calculate f_3 for points exhibiting poor values of f_1 and f_2 . Given the small difference in computational cost between the two levels the acceleration is likely to be marginal, but should still be apparent if the philosophy of conducting analysis on a “need-to-know” basis has been implemented effectively.

The other two formulations in Table 8.1 seek further acceleration by employing data-fit low-fidelity models. Response surfaces constructed using all previously analysed points provide approximations for f_1 and f_2 . An additional dominance relation utilising these computationally cheap predictions is included ahead of those present in the previous two formulations, sifting designs before resorting to direct calculation of any objective values. This should lead to further acceleration as computational effort will not be wasted calculating objective values for points predicted to perform poorly by the response surfaces.

A low-fidelity model is not used for f_3 due to an increased risk of masking. When ordering a single performance parameter only one design is considered non-dominated. Using a low-fidelity assessment for this parameter would therefore result in just a single design progressing to the dominance relations employing the true objective values, with the absolute predictive accuracy of the response surface playing an important role in determining which design this should be. This reliance on the accuracy of the response surface increases the likelihood of low-fidelity masking, therefore dominance relations seeking acceleration through the use of low-fidelity models are limited to those that trade-off performance parameters against one another where the absolute accuracy of the response surfaces is less important.

8.3.1.2 Response Surface Construction Techniques

Gaussian Process (GP) and Radial Basis Function (RBF) response surface construction techniques are employed, with both implemented using pySOT [66].

GPs, also known as Kriging models, were initially developed by mining engineer Krige [140] and introduced to the engineering design community by Sacks et al. [208]. This technique treats data points as random variables, interpolating between them by considering their correlation. This results in the output of a GP response surface being itself a random variable, with a mean quantifying the expected function value and a variance giving an indication of confidence which decreases further away from known data points. Details of the mathematics underpinning GPs can be found in the work of Forrester and Keane [73]. GPs are popular in multi-fidelity optimisation studies [107, 108, 120, 125, 200, 217] due to their ability to model complex functions using a relatively small amount of data. However, this comes at the cost of complexity and construction time, with calculation of the model parameters involving an optimisation routine that can become prohibitively expensive as dimensionality increases.

RBFs were first introduced by Broomhead and Lowe [26] and use a weighted sum of basis functions (or “kernels”). These range from linear examples to splines and Gaussian exponentials capable of modelling more complex landscapes. RBFs have lower construction costs as no optimisation is required to determine the model parameters. The approach has found successful employment in several multi-fidelity optimisation studies [95, 113, 115, 124, 165, 166, 180, 199, 244].

Instead of generating an initial dataset using design of experiments techniques, points analysed previously during the search are used as training data. This ensures that progress is made from the outset rather than significant computational effort being expended before any optimisation has been carried out. However, at the start of the optimisation there will be insufficient data to construct the response surfaces. In this situation all designs are assumed to have the same performance level, making them incomparable according to the dominance relation using the data-fit low-fidelity model. As a result this dominance relation is effectively bypassed, with the optimiser proceeding in “single-fidelity mode”, making progress through the design space whilst building the set of training data.

8.3.1.3 Experimental Details

Each of the formulations in Table 8.1 is run from the same 20 sets of 50 randomly generated starting points using a computational budget equivalent to 1000 calculations of all three objectives. The computational cost of training the response surfaces is not taken into account as it is assumed to be negligible compared to the high-fidelity analysis employed in the engineering problems that are of primary interest. However, this training cost is discussed in Section 8.3.1.5.

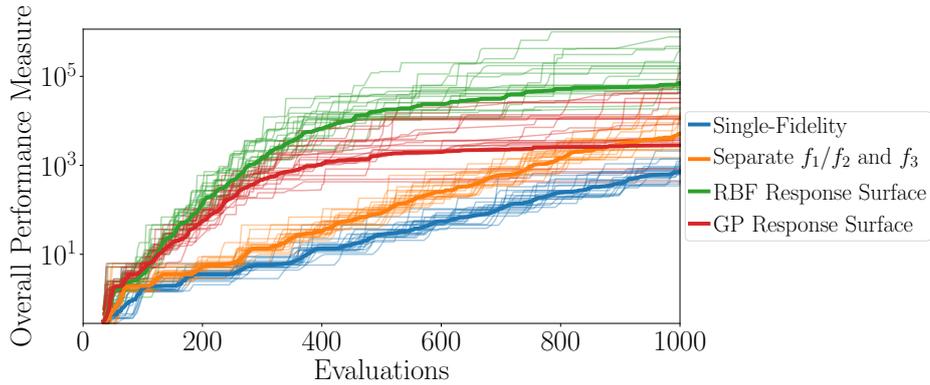


Fig. 8.3 Progression of the best overall performance measure found using the formulations in Table 8.1 applied to the multi-fidelity analytic test case. All runs are shown as faint lines with the mean plotted in bold.

8.3.1.4 Results

Figure 8.3 tracks the best values of the overall performance measure found using each formulation. This is calculated by taking the reciprocal of the shortest Euclidean distance in design space to the known optimum situated at $[3, 0, 0, 0]$, with higher values indicating better performance. Individual runs are plotted as faint lines, with the mean shown in bold.

Separating the calculation of the three objective functions leads to acceleration of the search. Points closer to the optimum are found using fewer evaluations, with performance equivalent to that generated by the single-fidelity method produced with an average computational saving of 30%. The multi-fidelity approach also finds designs with values of the overall performance measure that are an order of magnitude higher than those generated using the single-fidelity formulation once the full computational budget has been expended. This validates the philosophy of providing speed-up by conducting analysis on a “need-to-know” basis and verifies the implementation of that philosophy within the MFMDRTS algorithm described earlier in this chapter.

The two formulations employing low-fidelity models provide further acceleration. The RBF method generates an average computational saving of 74% compared to the single-fidelity approach, with the reduction in cost associated with using GP response surfaces slightly lower at 60%. The final values of the overall performance measure generated using the RBF method are over two orders of magnitude higher than those found by the single-fidelity algorithm. The use of response surfaces also provides additional speed-up compared to separating the calculation of the different objectives, with the RBF approach finding points equivalent to the best produced by the alternative multi-fidelity formulation with an average computational saving of 58%. These results demonstrate the ability of the MFMDRTS algorithm to provide speed-up through successful handling of data-fit low-fidelity models.

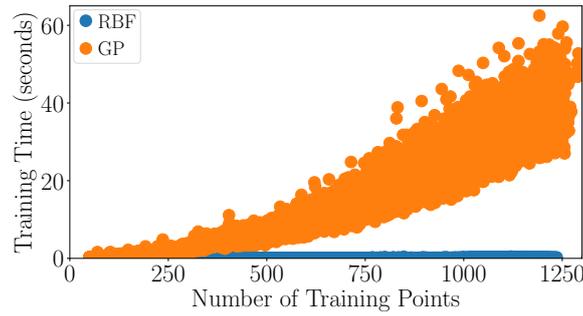


Fig. 8.4 Time taken to train the RBF and GP response surfaces during the optimisations applied to the multi-fidelity analytic test case.

8.3.1.5 Comparing Response Surface Methodologies

Both response surface techniques provide acceleration compared to the single-fidelity formulation, however the RBF approach appears to have outperformed the GP method as the search moves into the exploitation phase. Whilst the former produces points that are up to three orders of magnitude closer to the optimum than those found using the single-fidelity method, the latter stalls, ultimately generating less improvement on average than separating the calculation of the different objectives. This suggests that the RBF technique is better suited to modelling the landscape of this analytic problem than the GP.

A more general comparison can be made in terms of the computational cost of training the different models. Figure 8.4 plots the time taken to train all of the different response surfaces constructed throughout the searches depicted in Figure 8.3 against the number of training points. Whilst the training cost of the RBF model remains consistently low, at less than one second, the time taken to construct the GP response surfaces increases exponentially with the number of training points available. The resultant computational expense for training a single response surface is still small when compared to the cost of even a single high-fidelity analysis in an engineering design problem. However, within the iterative context of optimisation response surfaces often need to be trained several hundred times, rendering the expense of training the GP seen in Figure 8.4 unacceptable. The RBF response surface methodology is therefore selected for use in the remaining test cases in order to retain training costs that can be considered negligible compared to the time taken for high-fidelity analysis.

8.3.1.6 Summary

The results presented in this section validate the philosophy of providing acceleration by conducting analysis on a “need-to-know” basis and verify the implementation of this approach within the MFMDRTS algorithm described earlier in this chapter. For the analytic test case reductions in computational cost of 30% have been demonstrated by separating the

calculation of the different objectives, with this increasing to as much as 74% when additional data-fit low-fidelity models are employed. Comparing two response surface construction techniques also informs the selection of an RBF methodology for the remaining test cases.

8.3.2 Aerofoil Test Case

As in previous chapters the aerofoil test case is used to assess the effectiveness of the approach when applied to a more challenging engineering design problem, with the relatively cheap computational analysis allowing the repeatability of any benefits to be determined. The conditions and parameterisation are the same as those used in Chapters 5 and 6. The problem is transonic, with the magnitudes of 10 Hicks-Henne bump functions acting as design variables and the key performance parameters being lift, drag, pitching moment, the amount of trailing edge separation, and a measure of the space available inside the aerofoil for structural requirements and fuel storage.

8.3.2.1 Multi-Fidelity Analysis

Analysis is again carried out using the open source CFD software SU² [182], with solution of the RANS equations selected as the high-fidelity using the same RAE2822 mesh as in Chapters 5 and 6. Solution of the Euler equations, also available through SU², provides a physics-based low-fidelity model, with suitable inviscid meshes constructed for each aerofoil using Gmsh [77]. Including the time taken for mesh generation, the computational cost of a single low-fidelity Euler analysis is around 6% of that of the high-fidelity RANS solver, which is approximately 20 seconds when run on 8 Intel Xeon 2.13 GHz CPUs. This computational saving comes at the expense of accuracy, with the Euler model unable to account for viscosity and the associated separation effects that can have a significant impact on the overall aerofoil performance. Values of C_L , C_D and C_M are available directly from both solvers, and a measure for the amount of trailing edge separation is given by calculating the area under the skin-friction coefficient vs. chordwise distance curve whenever the former is negative. The cross-sectional area, A_c , provides an assessment of the internal space available.

8.3.2.2 Problem Formulations

Throughout this thesis a distinction has been made between scenarios in which the designer primarily wishes to improve a single objective and those where they are interested in investigating a trade-off between two or more quantities. Both scenarios are used to assess the accelerating capability of the MFMDRTS algorithm.

As in Chapters 5 and 6, the MDR formulation in the single-objective scenario, outlined in Table 8.2, treats C_D as the objective and C_M and separation as desirable features, with constraints applied requiring C_L to be greater than 0.6 and A_c to be less than that of an

Table 8.2 Formulations applied to the multi-fidelity aerofoil test case in the single-objective scenario. $A_{min} = 0.07784 \text{ m}^2$.

		Single-Fidelity	Separate Area Calc.	Response Surface	Corrected Euler
\succ_1	Fidelity Parameters	RANS $C_L \geq 0.6,$ $A_c \geq A_{min}$	Analytic $A_c \geq A_{min}$	Analytic $A_c \geq A_{min}$	Analytic $A_c \geq A_{min}$
\succ_2	Fidelity Parameters	RANS $C_M, \text{Sep.}$	RANS $C_L \geq 0.6$	Data-Fit RBF $C_L \geq 0.6$	Corr. Euler $C_L \geq 0.6$
\succ_3	Fidelity Parameters	RANS C_D	RANS $C_M, \text{Sep.}$	Data-Fit RBF $C_M, \text{Sep.}$	Corr. Euler $C_M, \text{Sep.}$
\succ_4	Fidelity Parameters	-	RANS C_D	RANS $C_L \geq 0.6$	RANS $C_L \geq 0.6$
\succ_5	Fidelity Parameters	-	-	RANS $C_M, \text{Sep.}$	RANS $C_M, \text{Sep.}$
\succ_6	Fidelity Parameters	-	-	RANS C_D	RANS C_D

RAE2822 aerofoil of the same chord. In the trade-off scenario, shown in Table 8.3, the C_L constraint is removed and this quantity is instead treated as an additional objective.

Three multi-fidelity variations of these formulations are applied in both scenarios. The first, labelled ‘‘Separate Area Calc.’’ in Tables 8.2 and 8.3, recognises that A_c can be calculated analytically, allowing the minimum A_c constraint to be assessed before resorting to computationally expensive RANS analysis. This should accelerate the search as any insufficient designs will be ruled out at the first dominance relation, meaning valuable computational budget is not wasted conducting high-fidelity analysis of designs that are unlikely to be of interest to the designer.

The other two multi-fidelity formulations seek further acceleration through the use of low-fidelity models. Two types are employed: a response surface fitted to data from all previously analysed points, similar to that used in the analytic study, and the physics-based Euler model combined with an additive correction function that again uses previously analysed points as training data. RBF response surfaces are employed in both cases following the results of the analytic test case. For these formulations, respectively labelled ‘‘Response Surface’’ and ‘‘Corrected Euler’’ in Tables 8.2 and 8.3, acceleration is provided by replicating the high-fidelity dominance relations. The low-fidelity models act as filters, ensuring only designs predicted to perform well are considered at the high-fidelity levels, leading to a reduction in the number of expensive analyses required. In the trade-off scenario this results in two additional dominance relations, one for low-fidelity assessment of the objectives and another

Multi-Fidelity Acceleration

Table 8.3 Formulations applied to the multi-fidelity aerofoil test case in the trade-off scenario. $A_{min} = 0.07784 \text{ m}^2$.

		Single-Fidelity	Separate Area Calc.	Response Surface	Corrected Euler
\preceq_1	Fidelity Parameters	RANS $A_c \geq A_{min}$	Analytic $A_c \geq A_{min}$	Analytic $A_c \geq A_{min}$	Analytic $A_c \geq A_{min}$
\preceq_2	Fidelity Parameters	RANS C_L, C_D	RANS C_L, C_D	Data-Fit RBF C_L, C_D	Corr. Euler C_L, C_D
\preceq_3	Fidelity Parameters	RANS $C_M, \text{Sep.}$	RANS $C_M, \text{Sep.}$	Data-Fit RBF $C_M, \text{Sep.}$	Corr. Euler $C_M, \text{Sep.}$
\preceq_4	Fidelity Parameters	-	-	RANS C_L, C_D	RANS C_L, C_D
\preceq_5	Fidelity Parameters	-	-	RANS $C_M, \text{Sep.}$	RANS $C_M, \text{Sep.}$

for the desirable features. In the single-objective scenario there are additional levels for the minimum C_L requirement and desirable features, however the objective is not assessed using the low-fidelity models for the same reasons that f_3 was not approximated using a response surface in the analytic problem. Doing so would result in just a single design progressing to the high-fidelity dominance relations, increasing the risk of low-fidelity masking.

As discussed in Section 8.3.1.2, using previously analysed points as training data results in insufficient information being available to construct the response surfaces in the early stages of the optimisation. For the data-fit method all designs are assumed to have the same performance, with the associated dominance relations effectively bypassed to allow the optimiser to proceed in “single-fidelity mode”. If a correcting response surface cannot be built for the physics-based low-fidelity model the correction from the nearest point in design space with high-fidelity data is used. This ensures that progress is made from the outset whilst taking advantage of the available physics-based information to carry out some sifting of designs before resorting to high-fidelity analysis.

8.3.2.3 Experimental Details

Each of the formulations in Tables 8.2 and 8.3 is run from 10 randomly generated sets of 20 starting geometries, allowing conclusions to be drawn about the repeatability of any benefits observed. The computational budget is equivalent to 500 RANS evaluations with a single Euler run costing 0.06 high-fidelity analysis units. The computational expense associated with the analytic A_c calculations and training the RBF response surfaces is negligible compared to the cost of a single RANS analysis. Each optimisation takes around two days to complete using 8 Intel Xeon 2.13 GHz CPUs.

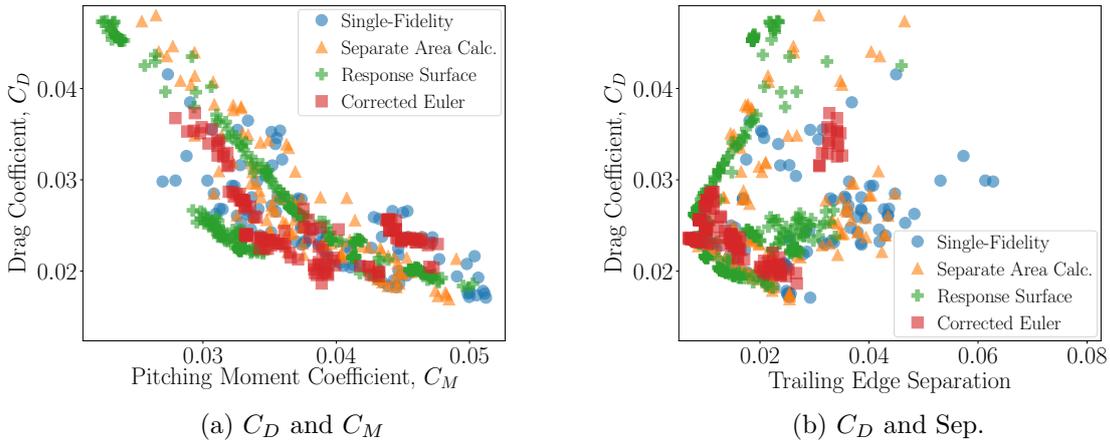


Fig. 8.5 Performance of the best designs found during 10 runs of the formulations in Table 8.2 applied to the multi-fidelity aerofoil test case in the single-objective scenario.

8.3.2.4 Single-Objective Results

Figure 8.5 shows the performance of designs found during the 10 runs of the formulations in Table 8.2 applied in the single-objective scenario that satisfy the C_L and A_c constraints and are non-dominated in terms of C_D , C_M and separation. The results produced by the different methods overlap, with the performance of aerofoils generated by the Single-Fidelity and Separate Area Calculation techniques being particularly similar. The approaches employing low-fidelity models in general find a larger number of non-dominated designs, with some of these having low values of the objective and desirable features.

Clearer distinctions can be made between the different methods using the reference point approach employed in previous chapters. The idealised reference point is given the best performance in terms of C_D , C_M and separation found by all of the runs applied to this problem, with the Euclidean distance from this reference to the points found by each search calculated in a space containing normalised versions of these three parameters. Figure 8.6 tracks the best values of an overall performance measure defined as the reciprocal of the minimum distance to this ideal point, with higher values indicating better performance. Individual runs are plotted as faint lines with the average shown in bold.

Separating the area calculation from the RANS analysis results in acceleration of the search, with equivalent final performance to that of the Single-Fidelity method produced with an average computational saving of 47%. Including low-fidelity models leads to further acceleration, particularly during the early stages of the search. The average computational saving achieved using the Response Surface and Corrected Euler formulations is 66% and 55% respectively. As the optimisation progresses the magnitude of acceleration decreases, as predicted by the hypothetical market share curves in Figure 3.1 discussed by Rubbert [205].

Multi-Fidelity Acceleration

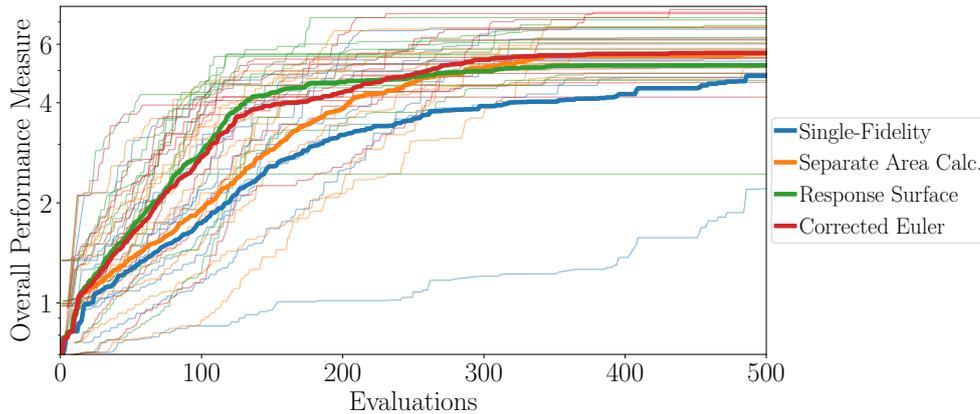


Fig. 8.6 Progression of the best overall performance measure found by the formulations in Table 8.2 applied to the multi-fidelity aerofoil test case in the single-objective scenario. All runs are plotted as faint lines with the average shown in bold.

This is probably due to the search moving between the exploration and exploitation phases, relying more heavily upon high-fidelity analysis in the latter to distinguish between designs.

Figure 8.7 shows the number of designs found by each method that exhibit performance within a given Euclidean distance of the idealised reference point, again with individual runs plotted faintly and the average shown in bold. The same conclusions can be drawn from this plot, with the separated A_c calculation producing higher numbers of interesting designs and the use of low-fidelity models improving the search further still.

There is little to choose between the Response Surface and Corrected Euler approaches in Figure 8.6, but Figure 8.7 appears to show the former outperforming the latter. These results disagree with previous research suggesting that the use of physics-based low-fidelity models is beneficial [69]. The most likely reason for this is the lack of physical modelling in the Euler code leading to erroneous predictions that fail to assist the optimiser. One of the features of the transonic aerofoil problem is flow separation, a phenomenon that the Euler equations are incapable of modelling due to the absence of viscosity. This leads to the physics-based technique being unable to provide values for trailing edge separation, with erroneous predictions for other performance parameters also likely. In this instance the acceleration provided by the Corrected Euler formulation is equivalent to that generated using the Response Surface approach, suggesting that any errors in the Euler predictions are not significant enough to lead the optimiser astray. This is possibly due to the goal of the optimisation being to minimise C_D , with the optimiser reducing the amount of trailing edge separation to limit the C_D increase associated with separated flow regions. With less separation present the inability of the Euler code to predict this feature is not as important, resulting in performance similar to that generated using a data-fit response surface.

This poor performance is not evidence of the techniques developed in Section 8.1.2 to avoid low-fidelity masking being insufficient. Previous work utilising a multi-fidelity version

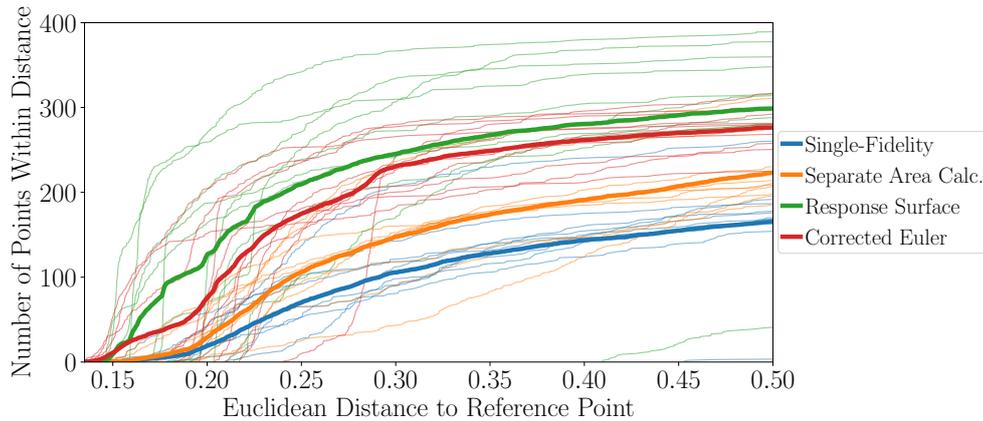


Fig. 8.7 Number of designs produced by the formulations in Table 8.2 applied to the multi-fidelity aerofoil test case in the single-objective scenario that exhibit performance within a given Euclidean distance of the reference point. All runs are plotted as faint lines with the mean shown in bold.

of Kriging [229, 243] showed that poor physics-based low-fidelity predictions can hamper any multi-fidelity optimiser. Instead these results suggest a need for caution when employing physics-based low-fidelity models. If the expected acceleration is to be achieved the selected codes need to provide sufficiently accurate performance predictions for the problem under consideration.

8.3.2.5 Trade-Off Results

The performance of designs on the C_L - C_D Pareto fronts found by the formulations in Table 8.3 applied in the trade-off scenario are shown in Figure 8.8, considering all designs generated across the 10 runs that satisfy the minimum A_c requirement. All four methods generate equivalent Pareto fronts in terms of C_L and C_D . The multi-fidelity formulations produce some improvement in terms of the desirable features, with this particularly evident in Figure 8.8b where designs found by the Response Surface and, to a lesser extent, Corrected Euler approaches exhibit lower values of trailing edge separation.

As in the single-objective scenario, the acceleration provided by the multi-fidelity methods is visible when tracking the best values of the overall performance measure found by the different formulations in Figure 8.9. This is again defined as the reciprocal of the minimum distance to an idealised reference point given the best performance in terms of C_L , C_D , C_M and separation found by any of the approaches applied to this problem, calculated in a space containing normalised versions of these quantities. The speed-up is most evident during the first 200 evaluations, with the use of low-fidelity models providing additional acceleration compared to separating the calculation of A_c . After this point the Single-Fidelity, Separate Area Calculation, and Response Surface methods converge as the optimiser moves

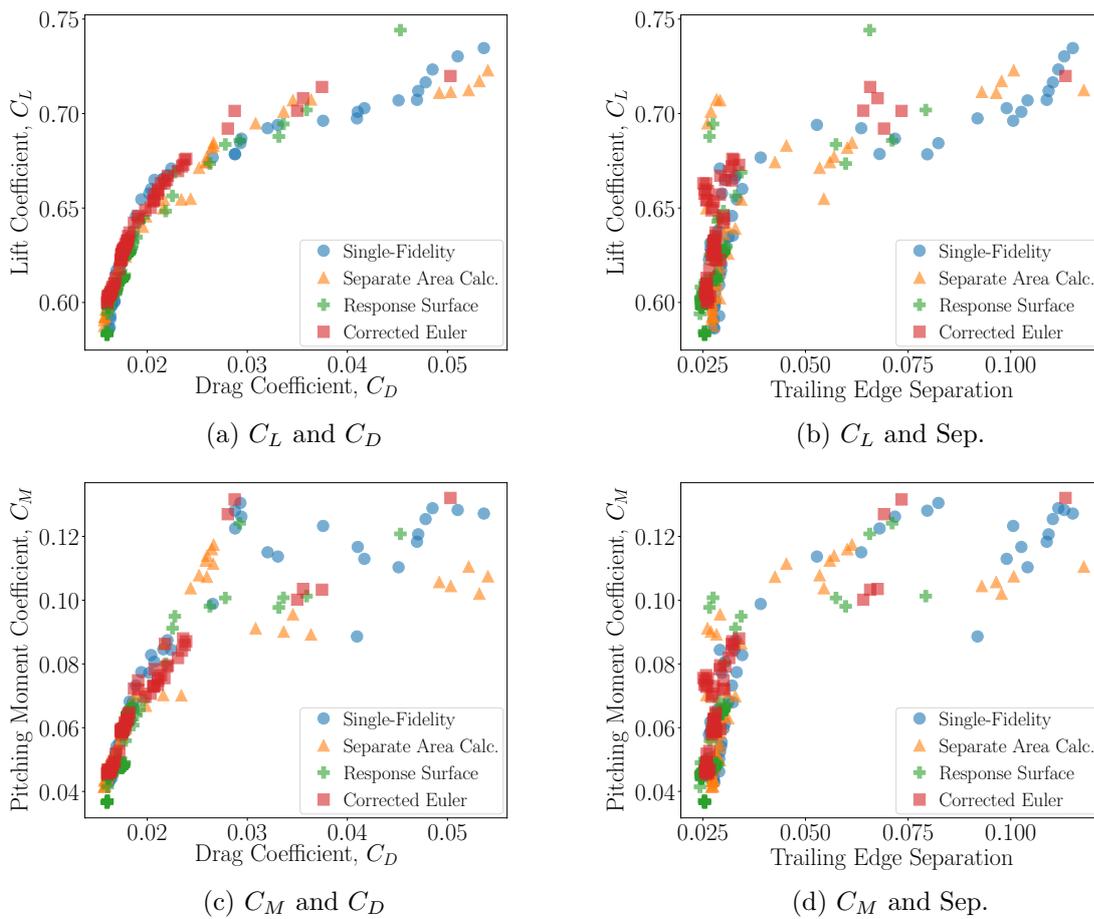


Fig. 8.8 Performance of the best designs found during 10 runs of the formulations in Table 8.3 applied to the multi-fidelity aerofoil test case in the trade-off scenario.

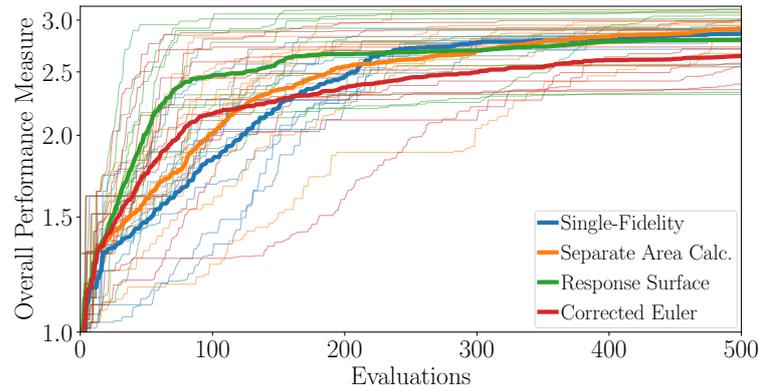


Fig. 8.9 Progression of the best overall performance measure found by the formulations in Table 8.3 applied to the multi-fidelity aerofoil test case in the trade-off scenario. All runs are plotted as faint lines with the average shown in bold.

into the exploitation phase, relying more heavily on accurate high-fidelity analysis. This convergence happens sooner than in the single-objective scenario, suggesting that high-fidelity information is more important for this trade-off problem. This is perhaps unsurprising as in a trade-off study more designs will be considered non-dominated and therefore passed through the hierarchy of dominance relations to the high-fidelity analysis levels. Separating the calculation of A_c leads to higher final values of the overall performance measure than the Single-Fidelity approach in 7 of the 10 runs, with values equivalent to the latter found with an average computational saving of 44%. The Response Surface methodology also outperforms the Single-Fidelity in 7 runs, achieving a larger average saving of 70%.

Figure 8.10, showing the number of designs exhibiting performance within a given Euclidean distance of the idealised reference point, demonstrates that the multi-fidelity approaches successfully use the available computational budget to produce large numbers of designs that are likely to be of interest to the designer. Again, the use of a response surface provides additional enhancements compared to just separating the calculation of A_c .

In Figure 8.9 the Corrected Euler method appears to stall after around 100 evaluations, and the results in Figure 8.10 suggest that on average a higher number of interesting designs have been produced by separate calculation of A_c . As in the single-objective scenario these poor results are most likely due to the inability of the Euler analysis to model separation effects. This appears to cause more problems in this trade-off scenario, possibly due to the optimiser attempting to maximise C_L . The resultant aerofoils exhibit greater curvature on their suction surfaces, leading to a higher probability of separation occurring and playing a significant role in the flow. The accuracy of the Euler method is low when analysing these designs, hampering the progress of the optimiser and leading to the poor performance seen in Figures 8.8, 8.9 and 8.10. Again, this is not evidence of a failure of the techniques for

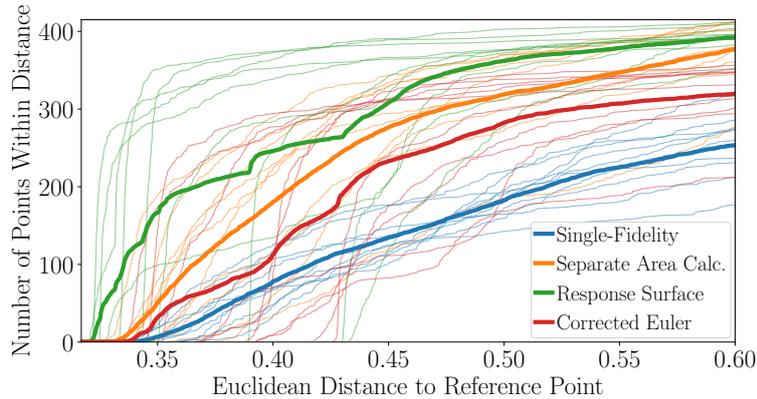


Fig. 8.10 Number of designs produced by the formulations in Table 8.3 applied to the multi-fidelity aerofoil test case in the trade-off scenario that exhibit performance within a given Euclidean distance of the reference point. All runs are plotted as faint lines with the average shown in bold.

avoiding low-fidelity masking, but instead highlights the need for careful consideration when selecting appropriate physics-based low-fidelity codes.

8.3.2.6 Summary

The results presented in this section demonstrate the accelerating capabilities of the MFMDRTS algorithm when applied to a more realistic engineering design problem. Average computational savings of around 45% are achieved utilising analytic A_c calculations, with this rising as high as 70% when data-fit response surfaces are employed. These savings are repeatable and experienced in both single-objective and trade-off scenarios. The approach successfully handles physics-based low-fidelity models, however the results highlight the need for care when employing this type of code for acceleration.

8.3.3 Six-Stage Axial Compressor Test Case

Having carried out an initial assessment of the acceleration provided by the new MFMDRTS approach it is ready to be applied to the preliminary design optimisation of an axial compressor. The six-stage test case from Chapter 7 is employed for this purpose, with the parameterisation and starting design detailed in Sections 7.1.2 and 7.1.4 respectively.

8.3.3.1 Multi-Fidelity Analysis

High-fidelity analysis is again provided by the open source turbomachinery CFD package Multall [52]. This includes a lower-fidelity axisymmetric throughflow method in addition to the three-dimensional RANS code employed in Chapter 7. The meanline script used to generate compressor geometries from the design variables can also act as a low-fidelity

performance prediction tool. The open source analysis suite therefore contains three distinct fidelity levels. The cheap meanline code relying on empirical loss models and calculations at the mean radius, a slightly more expensive throughflow code taking into account radial effects, and a three-dimensional RANS analysis providing the most accurate performance predictions available in this study. Each throughflow analysis takes around 45 seconds, a computational saving of 97.5% compared to that of a RANS run, whilst the meanline code takes $O(0.1)$ seconds.

As in the previous chapter, seven performance parameters are considered for this test case: η_p , PR , SM , DH , DF , M_{exit} and α_{exit} . Approximations for all but SM are provided directly by each of the analysis codes, with the RANS and meanline fidelity levels generating SM predictions using the correlation of Schweitzer and Garberoglio [212]. The throughflow code is unable to reliably use this correlation due to an inherent under-prediction of losses in the machine, meaning it cannot provide SM predictions and also produces consistently high values of η_p and M_{exit} . The correction procedure should overcome these shortcomings but the lack of SM prediction is likely to limit optimisations employing the throughflow model.

8.3.3.2 Problem Formulations

As in the aerofoil test case, two scenarios are considered for the six-stage axial compressor problem, one where a single objective is being improved and another in which a trade-off is sought between two performance parameters. For the former, η_p is selected as the objective with M_{exit} and α_{exit} treated as desirable features. Constraints are applied to the remaining performance parameters, with additional upper limits set for M_{exit} and α_{exit} as in Chapter 7. The trade-off study retains the same desirable features and constraints, but treats SM as an additional objective alongside η_p . These performance parameter classifications are summarised in Table 8.4, with the corresponding single-fidelity MDR formulations outlined in Tables 8.5 and 8.6.

Four approaches are used to assess whether the multi-fidelity method developed in this chapter is capable of accelerating the search in these two scenarios. The first three employ different low-fidelity models, replicating the high-fidelity dominance relations to act as filters that allow only the most promising designs to progress for computationally expensive high-fidelity analysis. In the trade-off study all three high-fidelity dominance relations are replicated, whereas in the single-objective scenario only the constraints and desirable features are treated using the low-fidelity models, as in the analytic and aerofoil test cases, to limit the risk of masking.

The low-fidelity models employed are a RBF response surface fitted directly to high-fidelity data points, labelled “Response Surface” in Tables 8.5 and 8.6, and meanline and throughflow physics-based codes in conjunction with additive correction functions, labelled “Corrected Meanline” and “Corrected Throughflow” respectively. The response surfaces

Multi-Fidelity Acceleration

Table 8.4 Performance parameter classifications for the multi-fidelity six-stage axial compressor test case.

	Single-Objective Scenario	Trade-Off Scenario
Objectives	η_p	η_p, SM
Desirable Features	M_{exit}, α_{exit}	M_{exit}, α_{exit}
Constraints	$PR \geq 2.75$ $DF \leq 0.6$ $DH \geq 0.72$ $SM \geq 15.0\%$ $M_{exit} \leq 0.35$ $\alpha_{exit} \leq 10.0^\circ$	$PR \geq 2.75$ $DF \leq 0.6$ $DH \geq 0.72$ $SM \geq 15.0\%$ $M_{exit} \leq 0.35$ $\alpha_{exit} \leq 10.0^\circ$

Table 8.5 Formulations applied to the multi-fidelity six-stage axial compressor test case in the single-objective scenario. The performance parameters assigned to each classification are shown in Table 8.4.

		Single-Fidelity	Response Surface	Corrected Meanline	Corrected Throughflow	Corrected ML and TF
\sphericalangle_1	Fidelity Class	RANS Constraints	Data-Fit RBF Constraints	Corr. ML Constraints	Corr. TF Constraints	Corr. ML Constraints
\sphericalangle_2	Fidelity Class	RANS Des. Feat.	Data-Fit RBF Des. Feat.	Corr. ML Des. Feat.	Corr. TF Des. Feat.	Corr. ML Des. Feat.
\sphericalangle_3	Fidelity Class	RANS Objectives	RANS Constraints	RANS Constraints	RANS Constraints	Corr. TF Constraints
\sphericalangle_4	Fidelity Class	-	RANS Des. Feat.	RANS Des. Feat.	RANS Des. Feat.	Corr. TF Des. Feat.
\sphericalangle_5	Fidelity Class	-	RANS Objectives	RANS Objectives	RANS Objectives	RANS Constraints
\sphericalangle_6	Fidelity Class	-	-	-	-	RANS Des. Feat.
\sphericalangle_7	Fidelity Class	-	-	-	-	RANS Objectives

Table 8.6 Formulations applied to the multi-fidelity six-stage axial compressor test case in the trade-off scenario. The performance parameters assigned to each classification are shown in Table 8.4.

		Single-Fidelity	Response Surface	Corrected Meanline	Corrected Throughflow	Corrected ML and TF
\preceq_1	Fidelity Class	RANS Constraints	Data-Fit RBF Constraints	Corr. ML Constraints	Corr. TF Constraints	Corr. ML Constraints
\preceq_2	Fidelity Class	RANS Objectives	Data-Fit RBF Objectives	Corr. ML Objectives	Corr. TF Objectives	Corr. ML Objectives
\preceq_3	Fidelity Class	RANS Des. Feat.	Data-Fit RBF Des. Feat.	Corr. ML Des. Feat.	Corr. TF Des. Feat.	Corr. ML Des. Feat.
\preceq_4	Fidelity Class	-	RANS Constraints	RANS Constraints	RANS Constraints	Corr. TF Constraints
\preceq_5	Fidelity Class	-	RANS Objectives	RANS Objectives	RANS Objectives	Corr. TF Objectives
\preceq_6	Fidelity Class	-	RANS Des. Feat.	RANS Des. Feat.	RANS Des. Feat.	Corr. TF Des. Feat.
\preceq_7	Fidelity Class	-	-	-	-	RANS Constraints
\preceq_8	Fidelity Class	-	-	-	-	RANS Objectives
\preceq_9	Fidelity Class	-	-	-	-	RANS Des. Feat.

are again constructed using previously analysed points, with the same tactics employed if sufficient training data is unavailable. For the data-fit method all designs are considered equivalent, whilst in the case of the physics-based low-fidelities the correction at the nearest high-fidelity data point in design space is used.

The final formulation, labelled “Corrected ML and TF” in Tables 8.5 and 8.6, exploits the capability of the MFMDRTS algorithm to handle multiple fidelity levels, with both meanline and throughflow analysis codes employed. The meanline data is corrected to the throughflow results, and the throughflow is in turn corrected to match the high-fidelity RANS predictions, both through RBF response surfaces fitted to additive correction functions. This formulation models the traditional axial compressor design process presented by Gallimore [76], passing designs sequentially through the analysis levels when further information is required and reserving computationally expensive high-fidelity analysis for only the most promising designs.

8.3.3.3 Experimental Details

Each of the formulations in Tables 8.5 and 8.6 is applied to the six-stage axial compressor test case using the initial design discussed in Section 7.1.4 and a computational budget equivalent to 500 high-fidelity RANS analyses. The computational costs associated with meanline analysis and constructing the RBF response surfaces are considered to be negligible compared to a single high-fidelity Multall run.

8.3.3.4 Single-Objective Results

Figure 8.11 shows the performance of designs generated by each of the formulations applied in the single-objective scenario that satisfy the constraints and are also non-dominated in terms of η_p , M_{exit} and α_{exit} . In general the multi-fidelity formulations produce machines exhibiting higher η_p and lower values of both desirable features than are found using the Single-Fidelity approach. This is particularly evident in Figure 8.11b showing the low values of α_{exit} achieved by the methods employing physics-based low-fidelity models.

The acceleration provided by the multi-fidelity formulations is visible in Figure 8.12 which tracks the best values of the overall performance parameter found by the different approaches. This is again defined as the reciprocal of the minimum distance to an idealised reference point given the best values of η_p , M_{exit} and α_{exit} found by all of the formulations applied to this problem, calculated in a space containing normalised versions of these parameters. Using physics-based low-fidelity models leads to good designs being generated with fewer calls to the high-fidelity analysis code. The Corrected Meanline formulation achieves performance equivalent to the best found by the Single-Fidelity in 80 fewer evaluations, a computational saving of 31%, whilst including the throughflow code as well in the Corrected ML and TF method leads to a slightly slower initial search but higher final saving of 37%. The Corrected

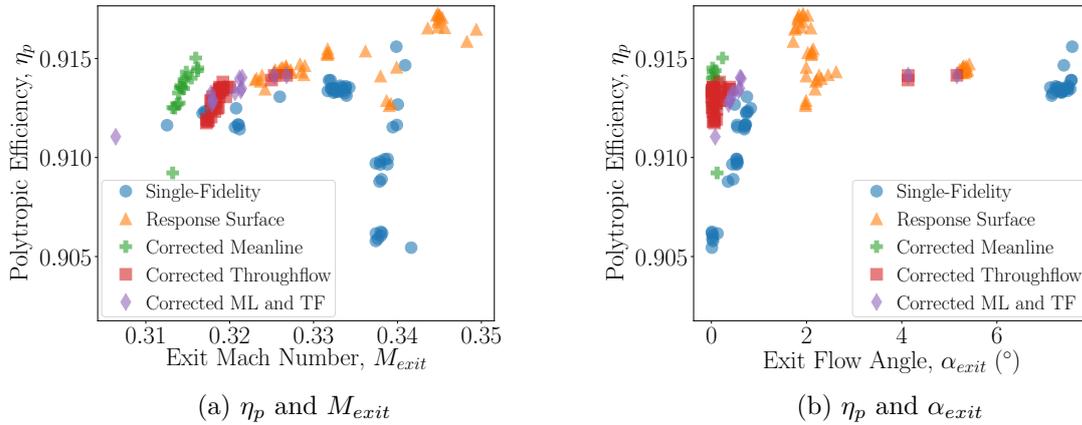


Fig. 8.11 Performance of the best designs found by the formulations in Table 8.5 applied to the multi-fidelity six-stage axial compressor test case in the single-objective scenario.

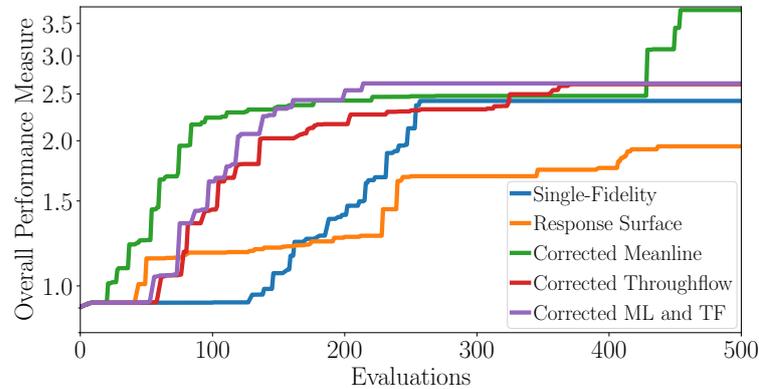


Fig. 8.12 Progression of the best overall performance measure found by the formulations in Table 8.5 applied to the multi-fidelity six-stage axial compressor test case in the single-objective scenario.

Throughflow formulation achieves similar acceleration in the early stages before slowing to a point where it is temporarily overtaken by the Single-Fidelity approach.

In Figure 8.13, which tracks the number of designs exhibiting performance within a given Euclidean distance of the idealised reference point, the methods employing physics-based low-fidelity codes are also seen to produce a higher number of designs that are likely to be of interest to the designer. Both this plot and Figure 8.12 suggest that the multi-fidelity technique is successfully accelerating search using the MDR formulation.

In contrast to the approaches employing physics-based low-fidelity codes, the Response Surface formulation performs relatively poorly. Despite good initial progress the data-fit technique stalls and results in designs with higher values of the desirable features than those found using the alternative multi-fidelity formulations. This suggests that the data-fit approach is ill-suited to the exploitation phase of the optimisation. In the early stages, when the optimiser is primarily exploring the design space, the Response Surface formulation does

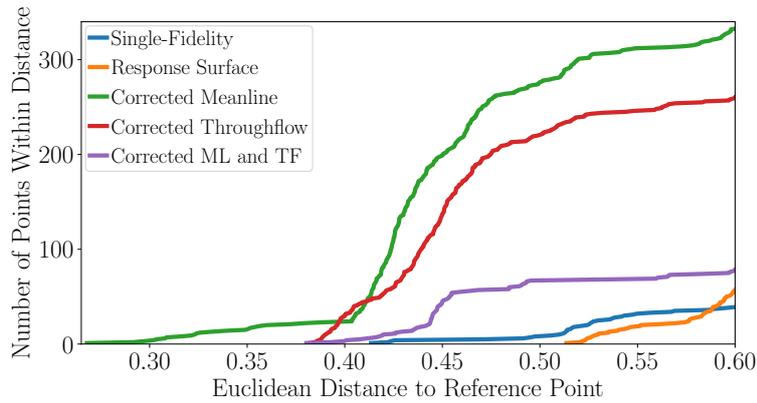


Fig. 8.13 Number of designs produced by the formulations in Table 8.5 applied to the multi-fidelity six-stage axial compressor test case in the single-objective scenario that exhibit performance within a given Euclidean distance of the reference point.

well, producing designs exhibiting good overall performance relatively quickly. However, as the search progresses the data-fit low-fidelity model appears to lack the accuracy required to improve promising designs further. This absence of effective exploitation has been experienced in Efficient Global Optimisation (EGO) [68], a popular optimisation approach developed by Jones et al. [120] that makes extensive use of response surfaces. The results in Figures 8.11, 8.12 and 8.13 suggest that good physics-based models can improve the exploiting capabilities of the MFMDRTS algorithm.

The three physics-based low-fidelity techniques produce similar results, although the Corrected Meanline method appears to have slightly outperformed the other two formulations. It might be expected that the higher fidelity of the throughflow analysis would result in more accurate approximations and therefore better optimisation performance. However, it appears that any accuracy improvement is outweighed by the increased computational cost. The inability of the throughflow method to provide accurate loss estimations may also be contributing to the relatively poor performance, and utilising an improved middle fidelity within this framework would probably lead to better results. Nevertheless, the fact that the Corrected ML and TF formulation produces designs that are likely to be of interest to the designer demonstrates the ability of the new algorithm to effectively handle multiple fidelity levels.

8.3.3.5 Trade-Off Results

Figure 8.14 shows the performance of designs that form the η_p -SM Pareto fronts found by the formulations in Table 8.6 applied in the trade-off scenario, only considering machines that satisfy the constraints. The most advanced Pareto front is generated by the Corrected ML and TF formulation, utilising both meanline and throughflow low-fidelity levels. However, these designs exhibit relatively high values of M_{exit} . The Corrected Throughflow method also

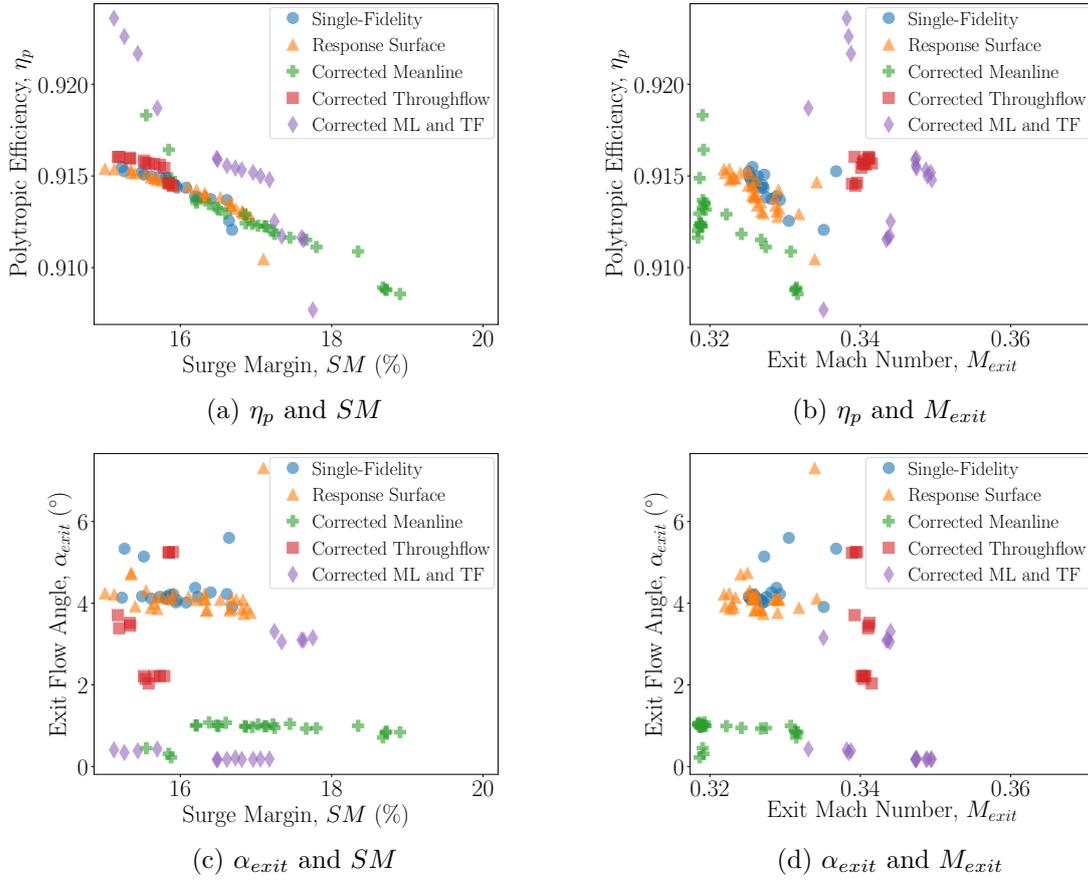


Fig. 8.14 Performance of the best designs found by the formulations in Table 8.5 applied to the multi-fidelity six-stage axial compressor test case in the trade-off scenario.

produces designs with high M_{exit} values, suggesting that modelling errors in the throughflow code may be hampering the optimiser as it attempts to minimise this quantity. The Corrected Meanline formulation produces a Pareto front that is more advanced than that generated by the Single-Fidelity approach, particularly at the extreme points of highest η_p and SM . Designs found using this method also exhibit some of the lowest values of both desirable features. As was the case in the single-objective scenario, the Response Surface formulation produces limited improvement compared to the Single-Fidelity approach, providing further evidence that physics-based low-fidelity information is required to ensure acceleration persists into the exploitation phase.

Figure 8.15 tracks the best values of the overall performance measure found during each search. This is again defined as the reciprocal of the minimum Euclidean distance to an idealised reference point given the best performance in terms of η_p , SM , M_{exit} and α_{exit} found by any of the formulations applied in this scenario, calculated in a space containing normalised versions of these parameters. The greatest acceleration is provided by the Corrected Meanline formulation, producing performance equivalent to the best generated by

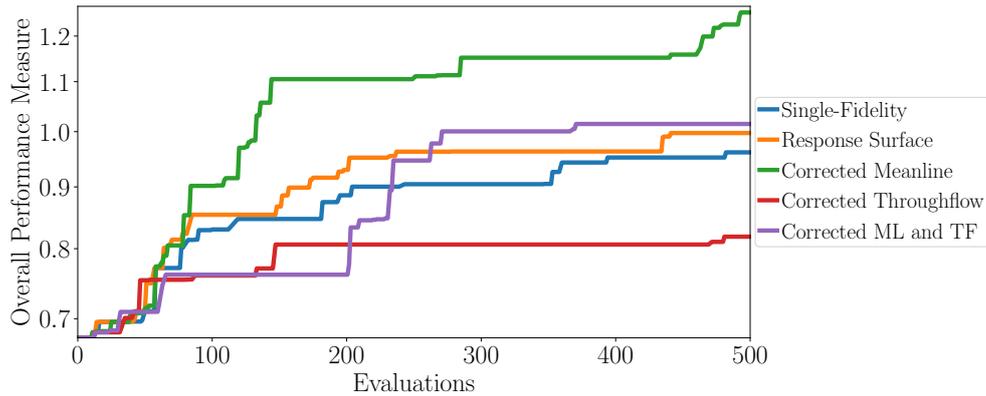


Fig. 8.15 Progression of the best overall performance measure found by the formulations in Table 8.6 applied to the multi-fidelity six-stage axial compressor test case in the trade-off scenario.

the Single-Fidelity approach with a computational saving of 75%. Moreover, this multi-fidelity method continues to find improved designs throughout the search, making good use of the available computational budget to evaluate the trade-off that is of interest. The Response Surface and Corrected ML and TF formulations also provide respective savings of 51% and 45% compared to the Single-Fidelity method. Values of the overall performance measure generated by the Corrected Throughflow approach are low due to the poor M_{exit} values found using this method.

Figure 8.16, plotting the number of designs exhibiting performance within a given Euclidean distance of the reference point, shows similar trends. The Corrected Meanline method generates a larger number of interesting designs than the Single-Fidelity approach, with the Response Surface and Corrected ML and TF formulations also providing improvement but to a lesser extent. The Corrected Throughflow method again performs poorly, primarily due to the high values of M_{exit} adding to the distance of the generated designs from the reference point.

As in the single-objective scenario, the Corrected Meanline formulation produces the best performance, suggesting that any accuracy improvements provided by the throughflow code do not outweigh the associated increase in computational cost.

8.3.3.6 Summary

The results in this section confirm that the accelerating capabilities of the MFMDRTS algorithm persist when applied to the preliminary design optimisation of axial compressors. The approach is able to successfully utilise the trusted physics-based low-fidelity models available to axial compressor designers to produce up to a 75% saving in computational cost compared to the single-fidelity method. Acceleration is observed across both single-objective and trade-off scenarios, with the greatest speed-up achieved using the meanline low-fidelity

8.4 Potential Applications of the Flexible MFMDRTS Algorithm

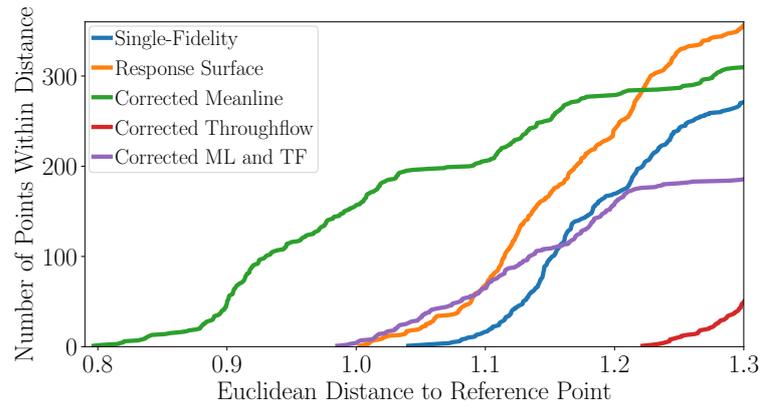


Fig. 8.16 Number of designs produced by the formulations in Table 8.5 applied to the multi-fidelity six-stage axial compressor test case in the trade-off scenario that exhibit performance within a given Euclidean distance of the reference point.

code. Whilst some limitations of the throughflow analysis technique have been exposed it is expected that employing the more adequate codes available in an industrial setting would lead to further improvement. The importance of using physics-based low-fidelity codes when appropriate has also been highlighted, with data-fit response surfaces apparently unable to provide sufficient accelerating assistance during the exploitation phase of the search.

8.4 Potential Applications of the Flexible MFMDRTS Algorithm

The results generated using the analytic, aerofoil and axial compressor test cases demonstrate that the goal of accelerating the MDRTS algorithm has been achieved. Through the use of low-fidelity models the computational budget required to produce performance equivalent to that found using the single-fidelity method has been reduced by as much as 75%. In the process a novel and inherently flexible multi-fidelity optimisation methodology has been developed with potential that goes beyond computational speed-up.

For example, the MFMDRTS algorithm is able to assess different performance parameters using different analysis codes. This attribute could be used to avoid the problems seen in the aerofoil and axial compressor test cases where physics-based low-fidelity codes were tasked with approximating performance parameters outside the scope of the phenomena being modelled. The technique could also be applied, without modification, to multidisciplinary and multipoint optimisation problems. In a multidisciplinary context performance could be assessed using an analysis code from one discipline at the first dominance relation, with analysis from another used at the next dominance relation, and so on. For multipoint studies each dominance relation could assess performance at a different operating condition.

In addition to facilitating acceleration through the inclusion of multiple fidelities, the algorithm developed in this chapter could in this way enable the application of the promising MDR formulation to alternative problem scenarios. The flexibility of the approach suggests that it has potential for future development with applications not limited to the preliminary design of axial compressors that formed the motivation for this work. All of these additional features come alongside the capabilities of the single-fidelity version of the algorithm developed and assessed in Chapters 4-7. This includes sophisticated problem definitions enabled by the MDR formulation and the ability to provide interpretable design development information to enhance designer understanding.

8.5 Summary

The MDRTS algorithm has been successfully accelerated using multi-fidelity methods. In this chapter the development of the improved implementation, incorporating the use of different analysis codes, has been detailed alongside results demonstrating the speed-up provided. This overcomes the last of the three problems highlighted in Chapter 3 that limit existing optimisation methodologies when applied to the preliminary design of axial compressors. The MFMDRTS algorithm enables sophisticated problem definitions to be used that adequately represent the desires of the designer, provides interpretable development information to help designers to determine the physical reasoning behind observed performance improvements, and efficiently exploits the available computational budget using multi-fidelity techniques.

Chapter 9

Application to an N-Stage Axial Compressor

As a final demonstration and assessment of the capabilities of the new MFMDRTS algorithm it is applied in this chapter to a more complex axial compressor design problem in which the number of stages in the machine is allowed to vary. This test case represents a significant challenge, with the optimiser given greater control over the configuration and geometry of the compressor. Using an increased computational budget compared to previous applications in Chapters 7 and 8, the ability of the new algorithm to efficiently produce designs that are likely to be of interest to the designer is assessed through comparisons to existing approaches employing more traditional objectives-and-constraints problem formulations.

9.1 N-Stage Axial Compressor Test Case

Jarrett and Ghisu [115] highlighted the importance of balancing configuration and refinement during the axial compressor design process. Configuration primarily refers to the number of stages in the machine, with this overall layout having a significant impact on the final performance. In a traditional axial compressor design process the configuration is often fixed at the preliminary design stage using a combination of designer experience and computationally cheap, low-fidelity meanline analysis techniques [76, 210]. However, there is potential for improvement if alterations to the configuration can be included in an optimisation scheme by treating the number of stages as an additional design variable [82, 83, 113, 115].

The benefits of this approach come with increased complexity due to the discrete design variable introduced to the problem. Previous work overcame this complexity by restricting the formulation to a single objective [113, 115] or employing low-fidelity analysis techniques to reduce the computational cost [82, 83]. In this chapter the new MFMDRTS algorithm is used to give the optimiser control over the configuration of the machine without either

of these restrictions being necessary. The MDR formulation allows a larger number of performance parameters to be handled effectively, whilst the multi-fidelity technique ensures that high-fidelity RANS analysis can be used with reasonable overall runtimes.

The simpler axial compressor test case employed in Chapters 7 and 8 was modelled on the HPC from the three-spool turbofan engine shown in Figure 1.1. To ensure adequate differentiation the IPC is used as a basis in this chapter, with a lower rotational speed and higher *PR*. A similar test case is used by Ghisu et al. [82, 83] allowing for some comparisons to be made.

9.1.1 Parameterisation

The parameterisation for this N-stage axial compressor test case is the same as that used for the six-stage problem described in Section 7.1.1, with the number of stages treated as an additional design variable. The optimiser is also given greater control of the annulus shape by removing the requirement for the casing radius to remain constant throughout the machine. The casing radius design variable is replaced by four parameters altering control points of a Bézier curve fitted to the cross-sectional area of the annulus. The final parameterisation consists of 36 variables that allow the optimiser to alter the configuration and geometry of the machine. As with the six-stage test case some of the Bézier control points can be changed simultaneously, with the optimiser able to alter variables corresponding to the mean radius, annulus area, stage pressure ratio, and stage exit flow angle in adjacent pairs or all at once.

Including the number of stages in this parameterisation and removing restrictions on the annulus shape results in a significantly more challenging problem than that of the six-stage axial compressor test case used in Chapters 7 and 8. The design space is larger and contains step changes in performance due to the presence of an integer design variable. This N-stage test case therefore represents a sufficiently difficult final problem to assess and demonstrate the capabilities of the new algorithm developed in this thesis.

9.1.2 Analysis

The open source analysis system discussed in Chapters 7 and 8 is also used for this N-stage problem. High-fidelity RANS CFD is employed to ensure improvement opportunities are not missed and to increase the likelihood of defects being discovered at the preliminary design stage when there is ample time and scope for their rectification [28, 76]. In the previous chapter a number of low-fidelity techniques were used to accelerate the algorithm when applied to the six-stage axial compressor test case. The corrected meanline low-fidelity code produced the greatest acceleration and is therefore used again in this chapter to demonstrate the achievable speed-up on the more challenging design problem.

9.1.3 Performance Parameters

The performance parameters relevant to this test case are the same as for the six-stage machine. These include η_p , PR , SM , and the two measures of blade diffusion, DH and DF . The exit conditions, M_{exit} and α_{exit} , also still need to be considered to prevent excessive pressure losses in downstream components [238]. For the HPC the main downstream component was the combustion chamber, however in this test case based on the IPC M_{exit} and α_{exit} need to be small to avoid separation in the duct linking the two compressors (see Figure 1.1). Some researchers have attempted to allow higher values of M_{exit} and α_{exit} by incorporating this duct into a single optimisation routine [82, 83, 115]. However, the added complexity associated with extending the realm of analysis is not considered worthwhile for this study where the goal is to demonstrate the capabilities of a new optimisation methodology rather than suggest improvements to axial compressors themselves.

When allowing the number of stages to vary it is important to include a measure of weight in the problem definition. Additional stages add mass to the machine, potentially leading to a reduction in the overall efficiency of the engine if used to power an aircraft that has to carry the extra weight. In lieu of a more accurate mass estimation the axial length, L_{ax} , of the machine is used, with shorter values corresponding to reduced weight [117, 185].

As in the six-stage axial compressor test case a number of geometric constraints are imposed. Maximum values for the blade speed at the hub and casing are set at 350 and 500 ms^{-1} respectively to ensure stresses in the discs and blades remain manageable [238]. To promote realistic blading the aspect ratios for all rows are restricted to values between 1.5 and 3.0. Finally, the mean radius is required to decrease throughout the machine to ensure compatibility with the existing engine architecture shown in Figure 1.1. These geometric constraints are again applied using a barrier method, with the optimiser prevented from selecting designs violating these limits as centres for the patterns searches used to generate new designs. Other constraints applied to the eight performance parameters discussed above are implemented using a penalty method, with individual penalty terms calculated as the amount of violation normalised by the limit value itself.

9.1.4 Initial Geometry

An initial geometry is generated using the techniques outlined in Section 7.1.4. Publicly available data informs suitable values for the number of stages, inlet conditions, rotational speed, mass flow rate and hub radius, with these fed into the Meangen program alongside the axial chords of the blades and values of ψ , ϕ and Λ for each stage. The resultant machine is parameterised and passed through the meanline generation script and Multall analysis system, with iterative variation of the input parameters leading to an eight-stage starting design whose geometry and performance are shown in Figure 9.1 and Table 9.1 respectively.

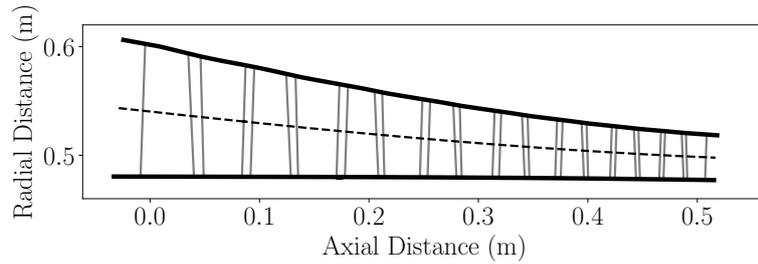


Fig. 9.1 Annulus of the initial eight-stage axial compressor used for the N-stage test case.

Table 9.1 Performance of the initial eight-stage axial compressor used for the N-stage test case.

η_p	0.875
PR	7.05
SM	15.04%
M_{exit}	0.347
α_{exit}	4.27°
DH_{min}	0.746
DF_{max}	0.441
L_{ax}	0.570 m

As in Chapter 7 this by no means represents a good axial compressor design, but is a sufficient starting point for the optimisations conducted in this chapter to demonstrate the capabilities of the new MFMDRTS algorithm.

9.1.5 Experimental Details

TS is applied in combination with more traditional objectives-and-constraints formulations throughout this chapter for comparative purposes due to the poor performance of alternative search methods such as GAs and PSO when applied to the preliminary design of axial compressors, as demonstrated in Section 7.2. Each run is given a larger computational budget equivalent to 1000 high-fidelity analyses, ensuring adequate time to explore the more complex design space. This increase is in line with similar previous work in which the optimiser was given control over the number of stages in the machine [113, 115]. Due to the additional stages in the machines being analysed the Multall analysis takes longer than for the six-stage test case, with a single run taking around 50 minutes on an Intel Xeon 2.13 GHz CPU and eight parallel analyses requiring approximately 90 minutes. The overall time taken for each optimisation in this more complex test case is therefore around eight days.

Table 9.2 Formulations applied to the N-stage axial compressor test case in the single-objective scenario.

	MDRTS & MFMDRTS	Multi-Objective	Single-Objective (Standard)	Single-Objective (Emergent)
Objectives	η_p	$\eta_p, M_{exit}, \alpha_{exit}, L_{ax}$	η_p	η_p
Desirable Features	$M_{exit}, \alpha_{exit}, L_{ax}$	-	-	-
Constraints	$PR \geq 7.0$ $DF \leq 0.6$ $DH \geq 0.72$ $SM \geq 15.0\%$ $M_{exit} \leq 0.35$ $\alpha_{exit} \leq 10.0^\circ$	$PR \geq 7.0$ $DF \leq 0.6$ $DH \geq 0.72$ $SM \geq 15.0\%$ $M_{exit} \leq 0.35$ $\alpha_{exit} \leq 10.0^\circ$	$PR \geq 7.0$ $DF \leq 0.6$ $DH \geq 0.72$ $SM \geq 15.0\%$ $M_{exit} \leq 0.35$ $\alpha_{exit} \leq 10.0^\circ$	$PR \geq 7.0$ $DF \leq 0.6$ $DH \geq 0.72$ $SM \geq 15.0\%$ $M_{exit} \leq 0.311$ $\alpha_{exit} \leq 0.75^\circ$ $L_{ax} \leq 0.597 \text{ m}$

9.2 Single-Objective Scenario

Throughout this thesis a distinction has been made between single-objective and trade-off scenarios. In the former designers primarily wish to improve a single performance parameter, whereas in the latter they are interested in investigating a trade-off between two or more quantities of interest. The performance of the new MFMDRTS algorithm applied to the N-stage axial compressor test case is assessed using both scenarios, with results for the first presented in this section.

9.2.1 Problem Formulations

As with the six-stage machine in Chapter 7 and previous work in which the optimiser varied the number of stages [113, 115], η_p is selected as the sole objective. Using the simplifying application framework developed in Chapter 5, M_{exit} , α_{exit} and L_{ax} are assigned to the desirable features classification as the designer wants to improve these but with a lower importance compared to η_p . The remaining parameters are treated using constraints, as outlined in Table 9.2, with upper limits for M_{exit} and α_{exit} applied in addition to their treatment as desirable features for the reasons discussed in Section 7.3.1. No constraint is applied to L_{ax} as specifying an appropriate limit value before carrying out any optimisation would be difficult and may artificially impede the progress of the optimiser through the design space. As in the previous chapter the multi-fidelity MDR formulation seeks to accelerate the search by replicating the dominance relations corresponding to the constraints and desirable features, using the corrected meanline low-fidelity model for analysis.

Performance using this MDR formulation is compared to that generated by three alternative objectives-and-constraints methods. The first is a multi-objective approach in which the three desirable features are treated as additional objectives alongside η_p , resulting in a four-objective trade-off being sought by the optimiser. The second removes the desirable features from the problem, retaining η_p as the sole objective and treating M_{exit} and α_{exit} using the upper limit constraints applied in the MDR formulation. Finally, a second single-objective formulation, discussed in Section 9.2.3, is applied using constraint limits that emerge as outputs of the two MDR approaches.

Each of the formulations in Table 9.2 is applied to the N-stage axial compressor design problem using the eight-stage starting geometry depicted in Figure 9.1 and a computational budget equivalent to 1000 high-fidelity analyses.

9.2.2 Results

Figure 9.2 shows the performance of designs found using each of the formulations in Table 9.2 that satisfy the constraints applied in the MDR problem definition and are also non-dominated in terms of η_p , M_{exit} , α_{exit} and L_{ax} . Designs with the highest η_p are found using the single-objective approach applying standard constraints. However, these machines exhibit poor performance in terms of the desirable features, particularly α_{exit} . The multi-objective formulation obtains good designs in terms of α_{exit} and L_{ax} , but these have poor values of η_p and M_{exit} due to the optimiser being unable to efficiently handle the four-objective trade-off. The single-fidelity MDRTS algorithm achieves better values of M_{exit} and α_{exit} than the traditional methods, whilst ensuring the resultant machines maintain high values of η_p . The multi-fidelity version produces similar results, but with higher values of η_p and a larger number of designs that are non-dominated in terms of the four key quantities of interest. The machines found by the MDR formulations exhibit higher values of L_{ax} than those generated using the multi-objective approach. There appears to be a trade-off between L_{ax} and M_{exit} , with the MFMDRTS algorithm uncovering designs along a wider extent of this trade-off than the single-fidelity version.

The same plots are repeated in Figure 9.3 with marker shape used to indicate the number of stages. Only designs that are non-dominated in terms of η_p , M_{exit} , α_{exit} and L_{ax} are plotted, meaning that for some approaches all of the results exhibit the same number of stages. For example, the best designs found by the MDRTS algorithm all have nine stages, with the additional blade rows allowing for low M_{exit} and high η_p values to be achieved at the expense of L_{ax} . In contrast, the multi-objective approach produces low values of L_{ax} by retaining the eight-stage configuration of the initial design.

Figure 9.4 tracks the best values of the overall performance measure found by the different formulations. This is defined as the reciprocal of the minimum Euclidean distance to an idealised reference point that is given the best performance in terms of η_p , M_{exit} , α_{exit} and

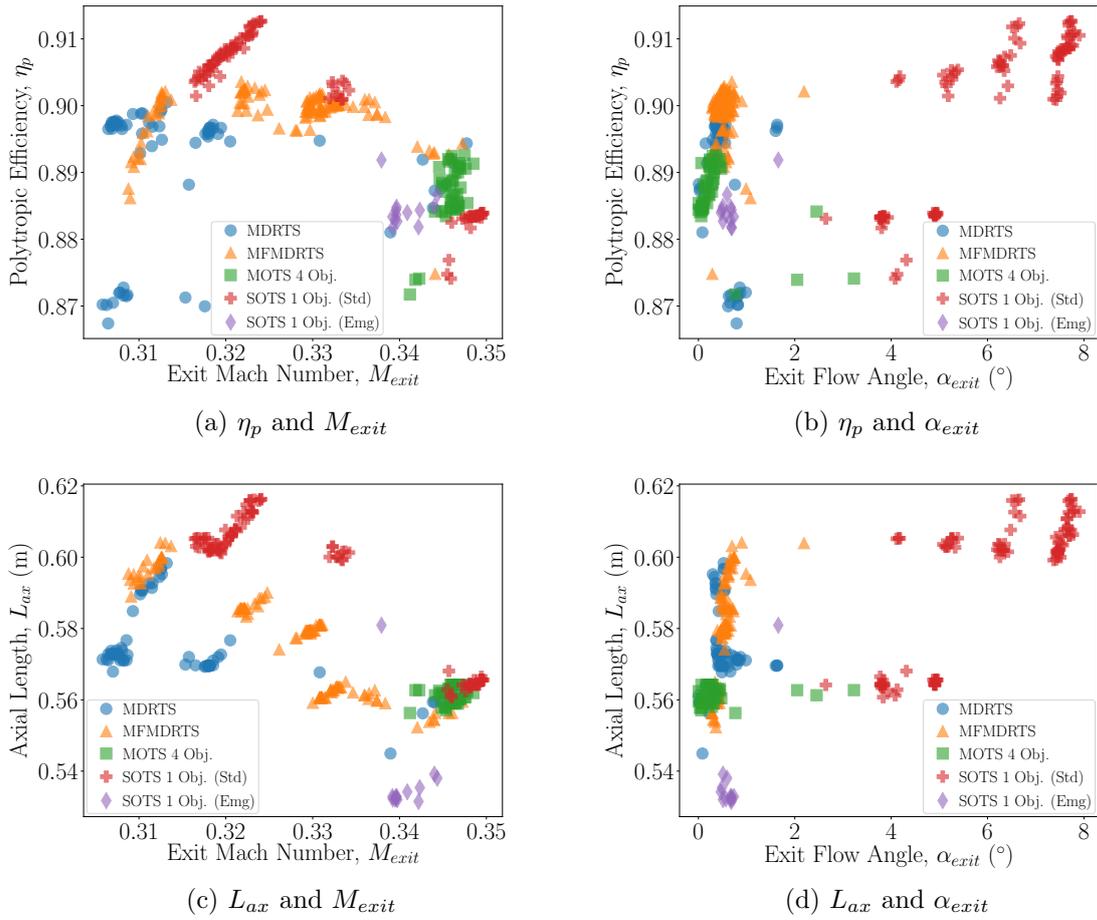


Fig. 9.2 Performance of the best designs found by the formulations in Table 9.2 applied to the N-stage axial compressor test case in the single-objective scenario.

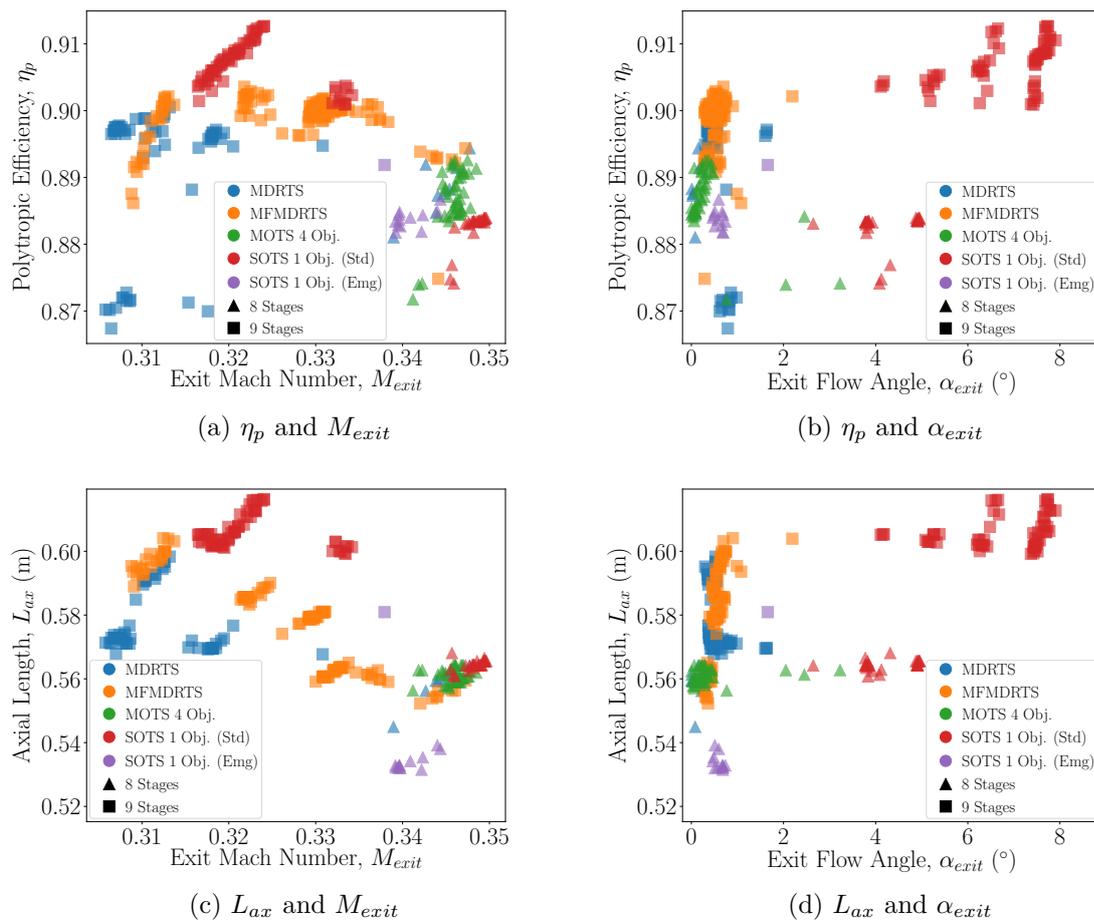


Fig. 9.3 Performance of the best designs found by the formulations in Table 9.2 applied to the N-stage axial compressor test case in the single-objective scenario with the number of stages indicated.

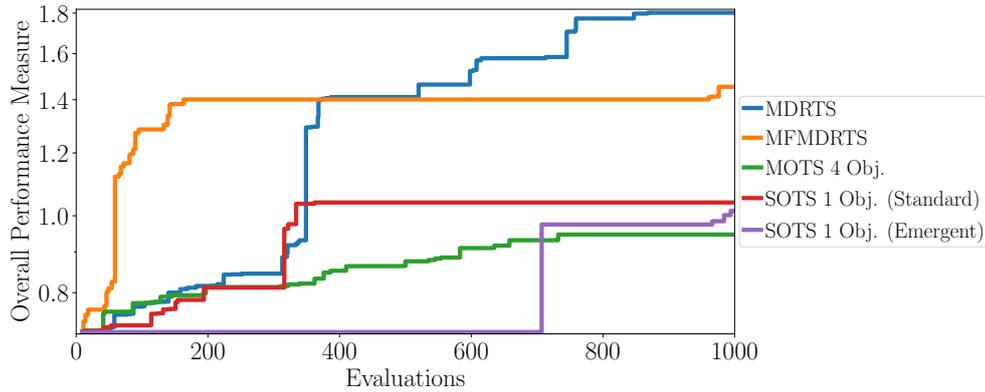


Fig. 9.4 Progression of the best overall performance measure found by the formulations in Table 9.2 applied to the N-stage axial compressor test case in the single-objective scenario.

L_{ax} found by any of the formulations applied to this problem, with the distance calculated in a space containing normalised versions of these four key parameters. Both the single- and multi-fidelity MDRTS algorithms generate designs with higher values of this overall performance measure than the traditional methods. The former finds designs equivalent to those produced by the multi-objective formulation using half as many calls to the high-fidelity analysis routine. The MDRTS algorithm fails to offer speed-up compared to the single-objective method employing standard constraints, but does continue to produce designs with improved overall performance throughout the search, with final values almost twice the best produced by the single-objective formulation. The multi-fidelity method provides additional acceleration, with computational savings of 92% compared to the multi-objective approach and 84% compared to the single-objective formulation employing standard constraints. For the first 350 evaluations the acceleration provided by the MFMDRTS algorithm compared to the single-fidelity version is evident, with a computational saving of around 75% achieved in these early stages. As the search progresses into the exploitation phase, however, the single-fidelity MDRTS method begins to outperform the multi-fidelity approach. This is probably due to the latter utilising the available budget to more thoroughly explore the trade-off between M_{exit} and L_{ax} , producing a wider range of interesting candidates for the designer to consider.

Similar conclusions can be drawn from Figure 9.5 which plots the number of designs exhibiting performance within a given Euclidean distance of the idealised reference point. Both MDRTS methods produce a larger number of designs with performance closer to that of the idealised machine than the traditional approaches. Whilst the single-fidelity method finds designs closer to the ideal point overall, the multi-fidelity method generates a higher total number of designs that are likely to be of interest to the designer.

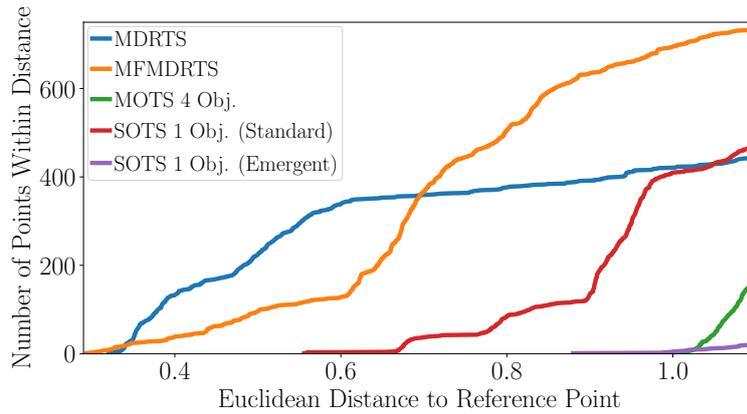


Fig. 9.5 Number of designs produced by the formulations in Table 9.2 applied to the N-stage axial compressor test case in the single-objective scenario that exhibit performance within a given Euclidean distance of the reference point.

9.2.3 Emergent Constraints

Figure 9.6 tracks the development of the desirable features during the searches conducted by the single- and multi-fidelity MDRTS approaches. As in previous chapters, the sudden changes in performance are due to the intensification and diversification heuristics used by the algorithm to enhance the search (see Section 4.1.2). The value of α_{exit} converges to a similar quantity in both cases, informing a suitable emergent constraint limit of 0.75° . Convergence of M_{exit} is more obvious in the single-fidelity approach as the multi-fidelity algorithm spends significant time evaluating the trade-off between this quantity and L_{ax} . The lower emergent value of 0.311 from the single-fidelity method is used for comparative purposes as the designer ideally wants to minimise this quantity. The plot for L_{ax} shows some convergence towards the latter stages, but is impacted in both cases by the increase in the number of stages. Despite the lack of a clear emergent constraint value this behaviour still enables a suitable upper limit of 0.597 m to be defined for L_{ax} that ensures the size of the compressor remains reasonable whilst giving the optimiser sufficient freedom to add stages if required.

These emergent constraint limits are applied within an additional single-objective formulation, outlined in Table 9.2, to determine whether traditional methods are capable of uncovering the same portion of design space found using the MDRTS technique. The results in Figures 9.2-9.5 suggest that this is not possible. Designs produced by the single-objective method employing emergent constraints exhibit good values of α_{exit} and L_{ax} , but among the worst performance in terms of η_p and M_{exit} . The most likely reason for this, as in previous chapters, is an inability to distinguish between constraints. Figure 9.7 tracks the number of designs found by the searches conducted using each of the different methods in Table 9.2 that satisfy the constraints applied in the MDR formulation. The single-objective approach using

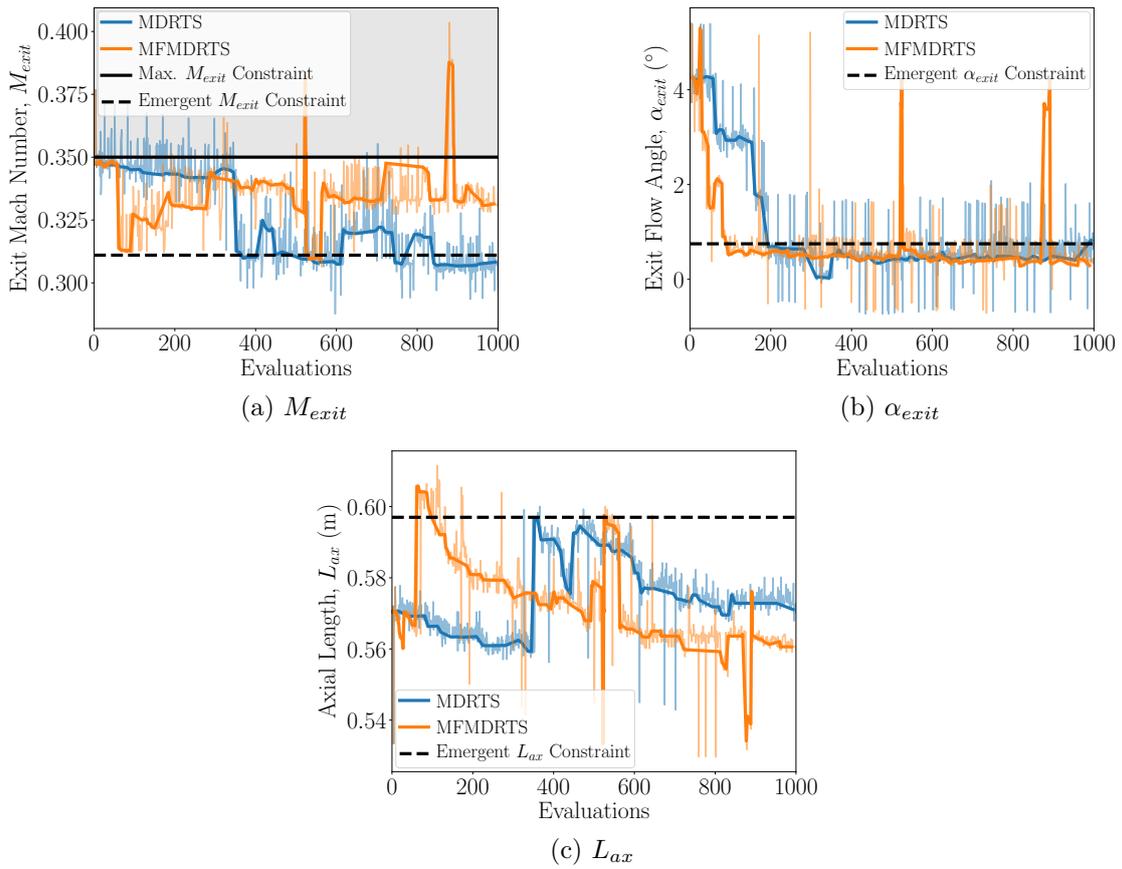


Fig. 9.6 Progression of the desirable features during searches using the single- and multi-fidelity MDR formulations applied to the N-stage axial compressor test case in the single-objective scenario. Performance of all designs is plotted as a faint line with that of the pattern search centres highlighted in bold.

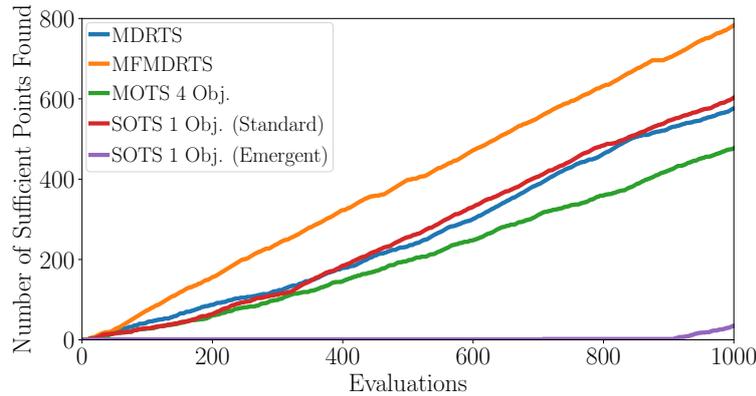


Fig. 9.7 Number of sufficient points found by the formulations in Table 9.2 applied to the N-stage axial compressor test case in the single-objective scenario.

emergent constraint limits spends over 90% of the available budget analysing insufficient designs. Tracking the development of M_{exit} and α_{exit} in Figure 9.8 reveals that this is primarily due to the upper limit for the former quantity being breached as the optimiser seeks to satisfy the aspirational constraint applied to the latter. The use of a penalty method means the optimiser is unable to recognise the problems associated with violating this upper limit, leading to the majority of the available budget being wasted analysing designs that are unlikely to be of interest to the designer. Alternative constraint handling methodologies could overcome this problem, but are also likely to inhibit the progress of the optimiser through the design space [82, 137]. The sophisticated MDR problem definition allows the desires of the designer to be handled more accurately, with fewer restrictions placed on the movement of the optimiser. This results in a larger number of potentially interesting designs being generated for the designer to consider. As a by-product, emergent constraint values are produced that could be employed in more detailed optimisation studies at a later stage in the design process.

9.2.4 Summary

The new MDRTS algorithm has been successfully applied to the N-stage axial compressor test case in this single-objective scenario. Designs produced by the new method exhibit better overall performance than those found using traditional methods. The multi-fidelity version of the new algorithm provides computational savings of over 90% compared to existing approaches, and enables a more thorough exploration of the available design space, focussing on the trade-off between M_{exit} and L_{ax} .

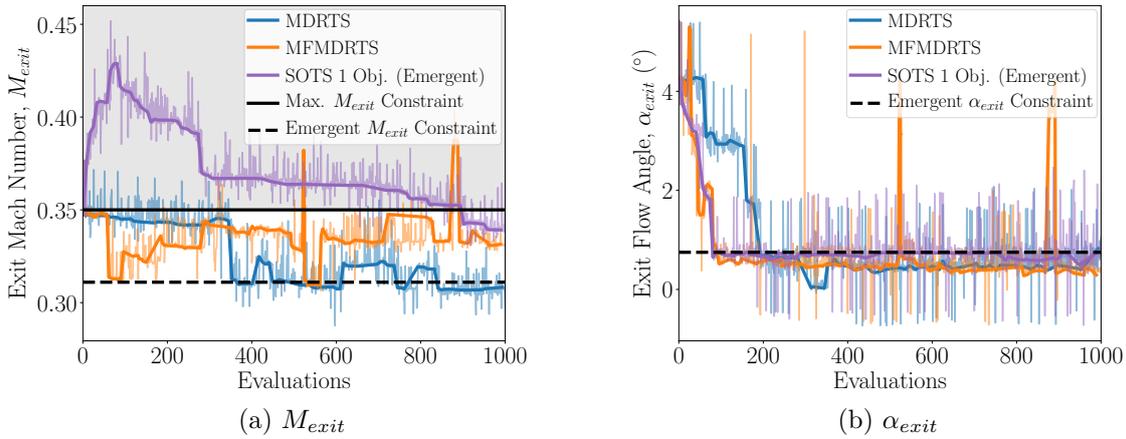


Fig. 9.8 Progression of M_{exit} and α_{exit} during searches using the single- and multi-fidelity MDR formulations and the emergent constraints method applied to the N-stage axial compressor test case in the single-objective scenario. Performance of all designs is plotted as a faint line with that of the pattern search centres highlighted in bold.

9.3 Trade-Off Scenario

Having demonstrated the capabilities of the new algorithm when applied to the N-stage axial compressor test case in a single objective scenario, in this section the same comparisons are made for a situation in which the designer wishes to investigate a trade-off between two or more quantities of interest.

9.3.1 Problem Formulations

As in Chapter 7 the trade-off scenario treats SM as an additional objective alongside η_p . This follows similar work by Ghisu et al. [83] in which the optimiser was given control over the number of stages in the compressor. The MDR formulation retains the same desirable features and objectives as the single-objective scenario and multi-fidelity acceleration is again sought using the corrected meanline model with the dominance relations corresponding to the constraints, objectives and desirable features replicated.

The new algorithm is compared to three traditional objectives-and-constraints formulations. The first treats the desirable features as objectives alongside η_p and SM , resulting in a five-objective problem definition. The second retains just two objectives, treating M_{exit} and α_{exit} using the upper limits suggested by Walsh and Fletcher [238], and removing L_{ax} from the problem entirely. Finally, an additional two-objective optimisation, discussed in Section 9.3.3, uses emergent constraints limits.

Each of these formulations, outlined in Table 9.3, is applied to the N-stage axial compressor test case using the initial geometry depicted in Figure 9.1 and a computational budget of 1000 Multall analyses.

Application to an N-Stage Axial Compressor

Table 9.3 Formulations applied to the N-stage axial compressor test case in the trade-off scenario.

	MDRTS & MFMDRTS	Five-Objective	Two-Objective (Standard)	Two-Objective (Emergent)
Objectives	η_p, SM	$\eta_p, SM, M_{exit}, \alpha_{exit}, L_{ax}$	η_p, SM	η_p, SM
Desirable Features	$M_{exit}, \alpha_{exit}, L_{ax}$	-	-	-
Constraints	$PR \geq 7.0$ $DF \leq 0.6$ $DH \geq 0.72$ $SM \geq 15.0\%$ $M_{exit} \leq 0.35$ $\alpha_{exit} \leq 10.0^\circ$	$PR \geq 7.0$ $DF \leq 0.6$ $DH \geq 0.72$ $SM \geq 15.0\%$ $M_{exit} \leq 0.35$ $\alpha_{exit} \leq 10.0^\circ$	$PR \geq 7.0$ $DF \leq 0.6$ $DH \geq 0.72$ $SM \geq 15.0\%$ $M_{exit} \leq 0.35$ $\alpha_{exit} \leq 10.0^\circ$	$PR \geq 7.0$ $DF \leq 0.6$ $DH \geq 0.72$ $SM \geq 15.0\%$ $M_{exit} \leq 0.335$ $\alpha_{exit} \leq 1.0^\circ$ $L_{ax} \leq 0.602 \text{ m}$

9.3.2 Results

Figure 9.9 shows the performance of designs found by each of the formulations applied in this trade-off scenario that are on the η_p - SM Pareto front, just considering machines satisfying the constraints applied in the MDR formulation. Perhaps surprisingly, the five-objective method produces designs exhibiting among the best performance in terms of η_p and SM . However, these machines have poor values of the desirable features, particularly α_{exit} and L_{ax} , demonstrating the inability of the traditional method to effectively handle the large number of performance parameters. The two-objective approach employing standard constraints generates an η_p - SM Pareto front that is entirely dominated by designs found using the five-objective formulation. Machines on this front also exhibit high values of at least one of the desirable features and are therefore unlikely to be of interest to the designer.

The single-fidelity version of the MDRTS algorithm produces a Pareto front that is equally as advanced as the best found by the traditional methods, but with constituent designs that have lower values of the desirable features. This is particularly true of α_{exit} and L_{ax} , demonstrating the ability of the new method to more effectively handle a larger number of relevant performance parameters. The multi-fidelity version of the new approach produces a Pareto front that is dominated by those generated using the single-fidelity MDRTS algorithm and the traditional five-objective formulation, with designs on this trade-off curve also exhibiting relatively high values of M_{exit} . This is probably due to the search focussing on designs with small values of L_{ax} , seen in Figures 9.9e and 9.9f, which have lower η_p but higher SM compared to those found using the single-fidelity MDRTS method and the traditional five-objective formulation. This highlights the weakness of the Pareto dominance criterion

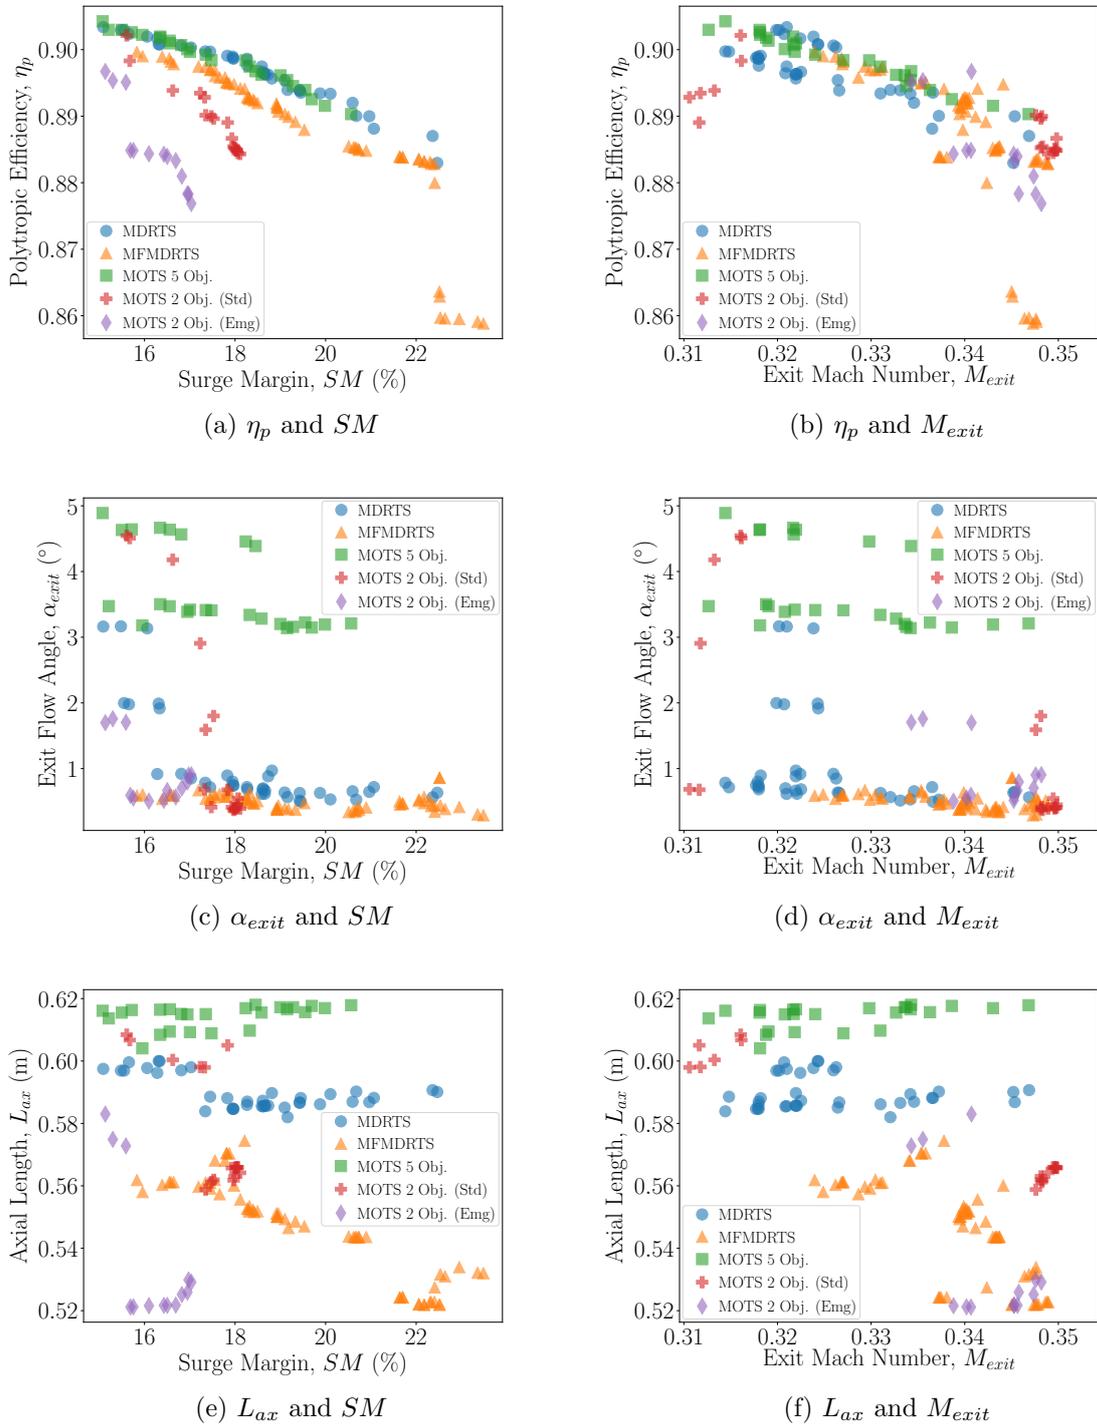


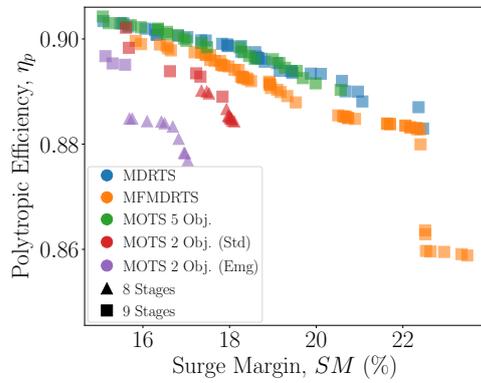
Fig. 9.9 Performance of the best designs found by the formulations in Table 9.3 applied to the N-stage axial compressor test case in the trade-off scenario.

that was discussed in Chapter 6, with improvements in L_{ax} accepted despite the associated designs exhibiting high values of M_{exit} . A more sophisticated dominance criterion, such as α -dominance [109], could encourage the optimiser to make similar reductions to all of the desirable features, rather than focussing on just one or two. Although outside the scope of this work, the MDR methodology, and by extension the new MDRTS algorithm, is fully capable of handling alternative dominance criteria and the inclusion of such techniques could be an avenue of further research.

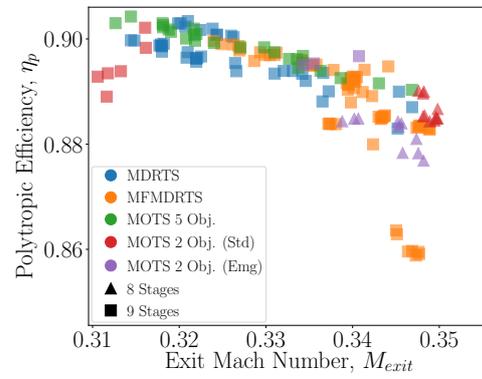
Figure 9.10 repeats the plots shown in Figure 9.9 using alternative marker shapes to indicate the configuration selected by the optimiser. Only designs that are non-dominated in terms of the objectives and desirable features are plotted meaning results produced by some formulations all exhibit the same number of stages. As in the single-objective scenario an additional stage allows for increased η_p and SM as well as lower values of M_{exit} . This increase in the number of stages counters the conclusions of Ghisu et al. [83] in which machines with fewer stages produced performance improvements. In that study M_{exit} was allowed to increase due to the inclusion of the inter-compressor duct in the optimisations scheme and it appears that the desire to reduce this parameter is the primary reason behind the higher number of stages in the results presented here.

The Pareto front produced by the MFMDRTS algorithm contains a step change in η_p as the value of SM passes around 22.3%. This is unexpected, and Figure 9.10a shows that it is not due to a change in the number of stages. Instead, the low values of η_p are a result of an error in the analysis procedure. Figure 9.11 shows the η_p - SM Pareto front alongside a plot of η_p against the number of iterations at which the CFD analysis for that design converged. Machines to the right of the step in η_p all converge after 3000 iterations, the upper limit set for the CFD analysis. Despite reaching the iteration limit these analyses have been treated as successful due to a sufficiently low value of the residual used to assess convergence. These results indicate that this acceptance was a mistake, as inaccurate performance prediction appears to have been provided for these designs. This error, with analysis results being accepted despite reaching the maximum number of iterations, was not observed in any of the other results reported in this thesis and does not effect the conclusions drawn about the performance of the single- or multi-fidelity versions of the MDRTS algorithm.

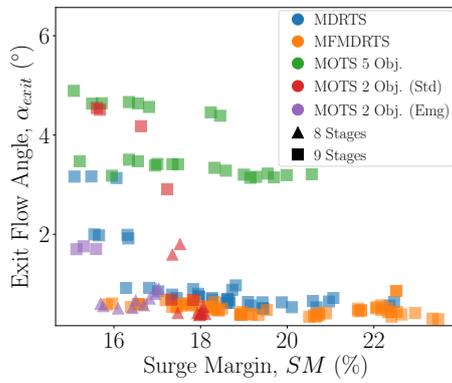
Figure 9.12 tracks the best values of the overall performance measure found by the different methods, with this again defined as the reciprocal of the minimum distance to an idealised reference point given the best performance in terms of the two objectives and three desirable features found by the different approaches, with this distance calculated in a space containing normalised versions of these parameters. The single-fidelity MDRTS algorithm produces designs with higher values of the overall performance measure than the traditional methods using fewer calls to the high-fidelity analysis function. Performance equivalent to the best found using the five-objective approach is generated using 73% fewer Multall evaluations



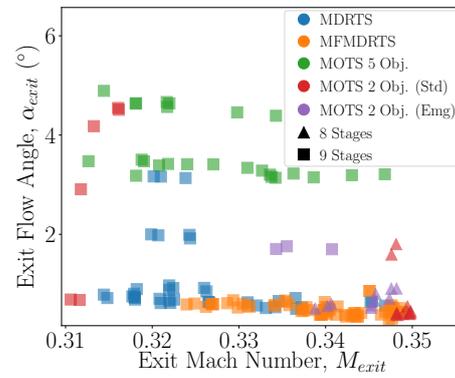
(a) η_p and SM



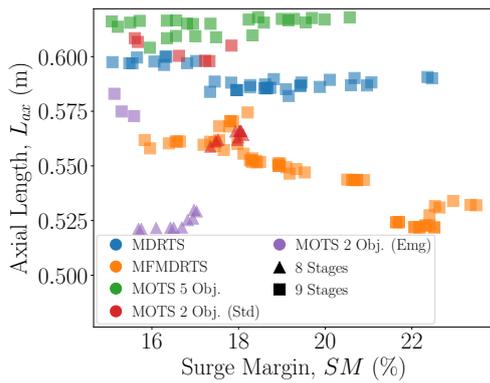
(b) η_p and M_{exit}



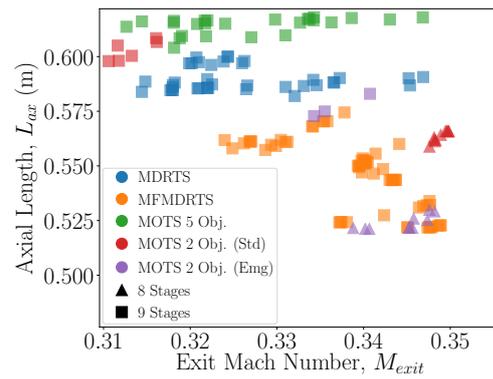
(c) α_{exit} and SM



(d) α_{exit} and M_{exit}



(e) L_{ax} and SM



(f) L_{ax} and M_{exit}

Fig. 9.10 Performance of the best designs found by the formulations in Table 9.3 applied to the N-stage axial compressor test case in the trade-off scenario with the number of stages indicated.

Application to an N-Stage Axial Compressor

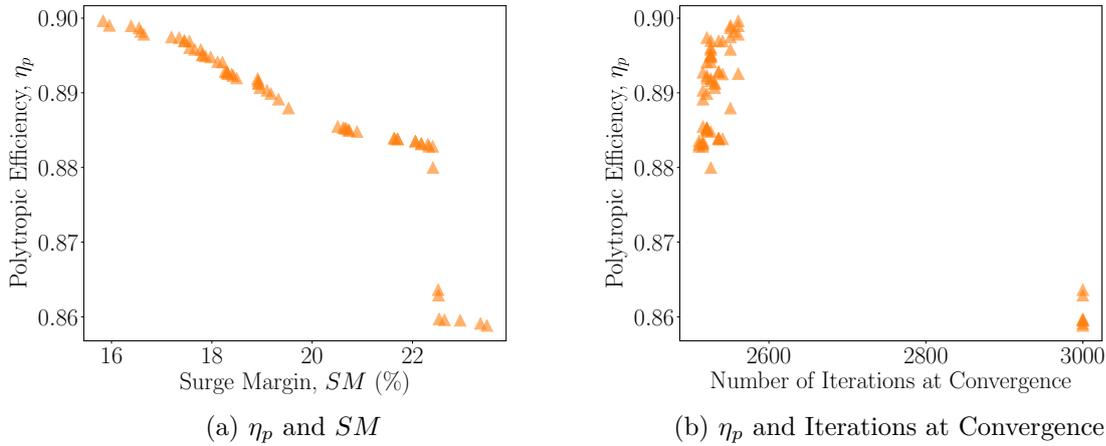


Fig. 9.11 Performance of designs on the η_p - SM Pareto front found by the MFMDRTS approach applied to the N-stage axial compressor test case in the trade-off scenario and the number of iterations taken for the CFD analysis to converge.

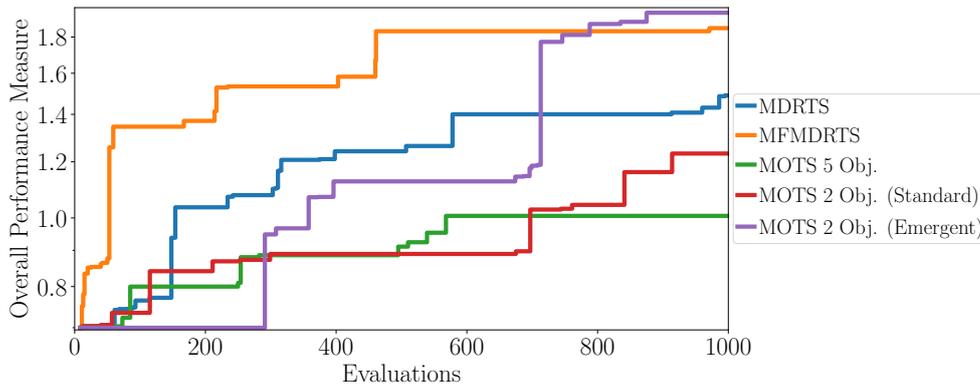


Fig. 9.12 Progression of the best overall performance measure found by the formulations in Table 9.3 applied to the N-stage axial compressor test case in the trade-off scenario.

and a computational saving of 56% is achieved compared to the two-objective approach employing standard constraints. The MFMDRTS algorithm provides further acceleration, with higher computational savings of 91% compared to the five-objective method and 94% compared to the standard two-objective approach. The multi-fidelity technique also produces performance equivalent to that generated by the single-fidelity MDRTS algorithm with a 78% reduction in the number of high-fidelity analyses.

Figure 9.13, tracking the number of designs found by the different approaches that exhibit performance within a given Euclidean distance of the reference point, supports these conclusions, with the multi-fidelity algorithm producing a larger number of designs that are likely to be of interest to the designer than the single-fidelity MDRTS method. Both approaches produce many more designs exhibiting high levels of overall performance than the two traditional formulations.

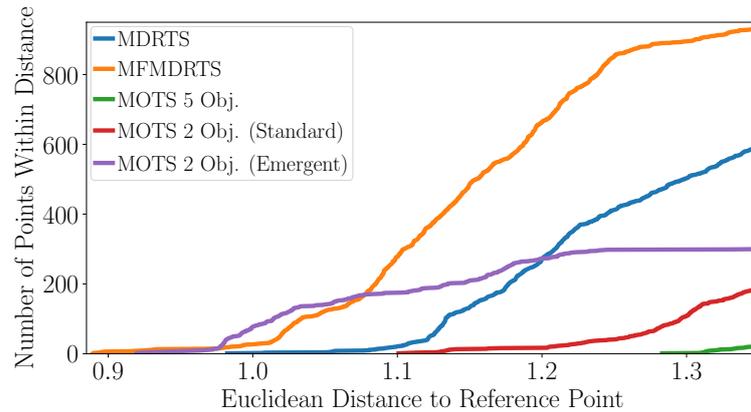


Fig. 9.13 Number of designs produced by the formulations in Table 9.3 applied to the N-stage axial compressor test case in the trade-off scenario that exhibit performance within a given Euclidean distance of the reference point.

9.3.3 Emergent Constraints

Figure 9.14 tracks the development of the desirable features during the single- and multi-fidelity MDRTS searches applied in this scenario. In both cases α_{exit} converges to similar values, allowing an informed choice of 1.0° to be made for the emergent constraint. The values of M_{exit} do not converge as obviously, possibly due to SM being treated as an objective when it is known to be negatively correlated to M_{exit} [81]. The opposing desires for these quantities results in a continuous trade-off between the objectives and desirable features, producing the oscillatory behaviour observed in Figure 9.14a. Despite the lack of convergence this plot can still inform a suitable limit value by selecting the mean of the oscillations, with a value of 0.335 chosen in this case. The progression of L_{ax} follows a similar trajectory to that observed in the single-objective scenario, decreasing steadily following a step increase when the ninth stage is added by the optimiser. As before, these results are used to define an upper limit of 0.602 m for L_{ax} that suitably restricts the length of the compressor whilst affording the optimiser sufficient freedom to add stages if necessary.

These emergent constraint values informed by the single- and multi-fidelity MDRTS methods are applied in an additional two-objective optimisation, outlined in Table 9.3, with the results shown in Figures 9.9-9.13. As in the single-objective scenario, this approach produces designs exhibiting poor performance compared to the alternative methods. The Pareto front is dominated in terms of η_p and SM , with the constituent designs also producing high values of M_{exit} . Performance in terms of the other two desirable features, α_{exit} and L_{ax} , is comparable to that produced by the MFMDRTS algorithm, with the low values of these parameters leading to the high overall performance measure in Figure 9.12 and a large number of designs close to the idealised reference point in Figure 9.13.

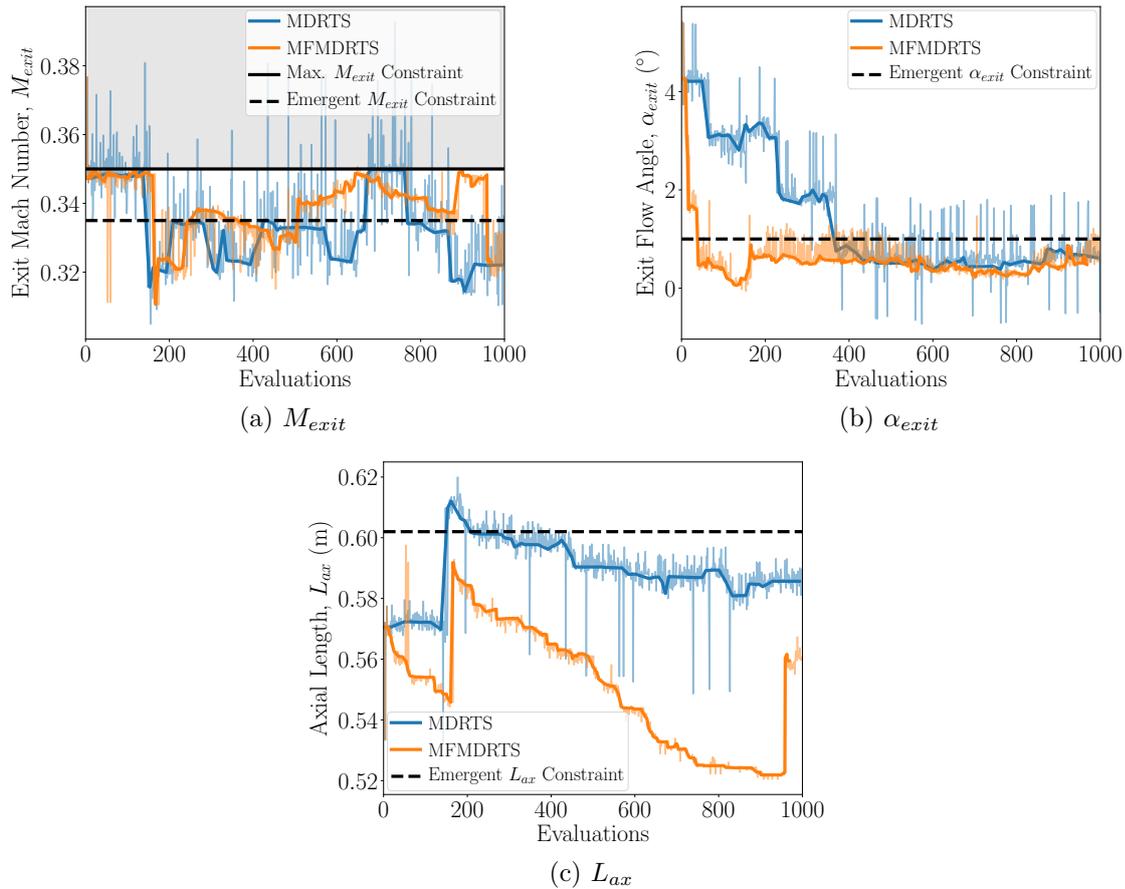


Fig. 9.14 Progression of the desirable features during searches using the single- and multi-fidelity MDR formulations applied to the N-stage axial compressor test case in the trade-off scenario. Performance of all designs is plotted as a faint line with that of the pattern search centres highlighted in bold.

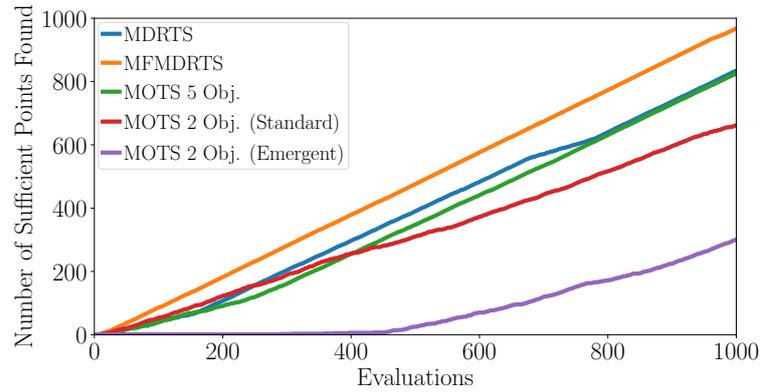


Fig. 9.15 Number of sufficient points found by the formulations in Table 9.3 applied to the N-stage axial compressor test case in the trade-off scenario.

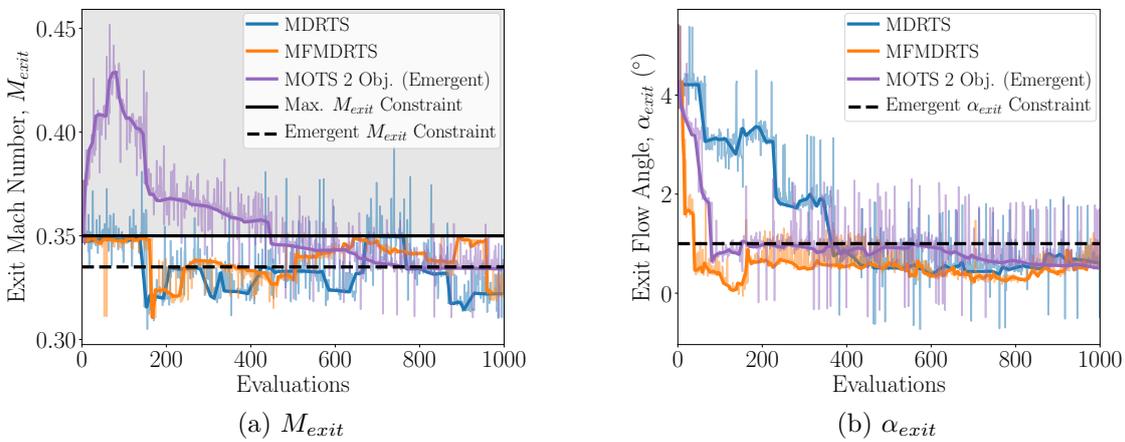


Fig. 9.16 Progression of M_{exit} and α_{exit} during searches using the single- and multi-fidelity formulations and the emergent constraints method applied to the N-stage axial compressor test case in the trade-off scenario. Performance of all designs is plotted as a faint line with that of the pattern search centres highlighted in bold.

The likely reason for the poor performance is again an inability to distinguish between the importance of different constraints, leading to a small number of sufficient designs being generated. This is shown in Figure 9.15 which tracks the number of designs satisfying the constraints applied in the MDR formulation, with the method employing emergent limit values producing fewer sufficient designs than the other formulations. Plots tracking M_{exit} and α_{exit} in Figure 9.16 reveal that this is again due to the optimiser violating the upper limit on the former parameter as it seeks to satisfy the aspirational emergent constraint applied to the latter. The MDR formulation enables these parameters to be handled more effectively, expressing the desire for them to be minimised without causing a large proportion of the available computational budget to be wasted analysing insufficient designs that are unlikely to be of interest to the designer.

9.3.4 Summary

The new MDRTS algorithm has been successfully applied to the N-stage axial compressor test case in a trade-off scenario. Designs produced exhibit better overall performance than those generated by traditional formulations, with the multi-fidelity version of the algorithm using low-fidelity meanline analysis to produce computational savings of over 90%.

9.4 Handling a Variable Number of Stages

Varying the number of stages during an optimisation is a non-trivial task as alterations to the fundamental architecture of the machine could result in numerous analysis code failures, potentially hampering the search if not handled with care. This complexity is probably one of the main reasons why the number of stages is fixed in most axial compressor optimisation studies in the literature. Both the single- and multi-fidelity versions of the MDRTS algorithm appear to have handled this challenging problem effectively, adopting the additional design parameter without any modification to the underlying method being required. A number of the final designs presented in the previous two sections exhibit more stages than the starting design, demonstrating the ability of the optimiser to handle the step change in performance. Several machines generated during the single- and multi-fidelity searches applied in both the single-objective and trade-off scenarios were also found to vary the number of stages more than once along their development path.

The suspected reason behind this ability to alter the number of stages with apparent ease is the procedure used to generate an axial compressor from the vector of design variables. Rather than making changes directly to the geometry itself, the approach adopted here instead varies the inputs to a meanline compressor generation script. This results in a more robust system that is less likely to fail when the number of stages changes, as can be seen in the results produced by the single- and multi-fidelity MDRTS algorithms. Across the two scenarios these algorithms attempted to vary the number of stages in the machine 863 times. None of these attempted geometry generations failed, with the meanline code successfully producing a set of input data for Multall on each occasion. Due to the sampling carried out few of these designs were actually selected for analysis, but the fact that a geometry was successfully generated in every case demonstrates the ability of the approach to effectively handle the number of stages as a design variable.

Whilst outright failure was not a problem, 460 of these 863 compressors generated by altering the number of stages did violate the aspect ratio constraints, suggesting that these user-defined limits may be preventing more regular alterations being made to the configuration of the machine. It appears that reducing the number of stages causes the greatest problems, with 79% of the designs violating the geometric constraints produced in this scenario. The constraint violation is probably caused by variations to the axial location of control points

for the Bézier curves used to alter the aspect ratios. An enhanced parameterisation that is less dependent on the number of stages in the machine could lead to further improvements in the ability of the new algorithm to vary the configuration.

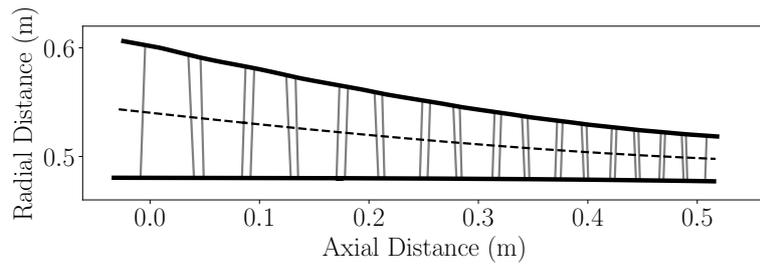
9.5 Information to Enhance Designer Understanding

The results in Sections 9.2 and 9.3 demonstrate the ability of the new algorithm to overcome the problems associated with formulation and computational cost that limit existing optimisations methodologies when applied to this challenging N-stage axial compressor test case. As an additional benefit the new method is also fully equipped to enhance designer understanding. Rich development information is available to help designers determine the physical reasoning behind the observed performance improvements, allowing final designs to be justified as well as improving understanding of the underlying problem to support creativity and innovation. In Section 7.5 the availability of this design development information was demonstrated for the single-fidelity version of the algorithm. In this section it is shown that similar data can be provided by the MFMDRTS method alongside the accelerated searching capabilities. Information is available for all designs produced by this method, but the machine exhibiting the highest values of η_p produced in the single-objective scenario is selected as an example. The annulus of this nine-stage machine is shown in Figure 9.17 alongside the eight-stage starting geometry.

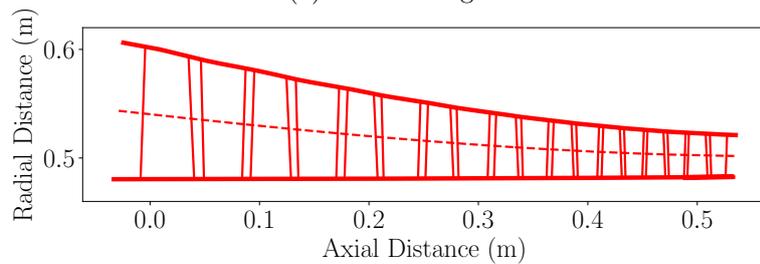
Development information is generated by stepping along the path taken by the optimiser to reach a design from the initial point. The pattern search technique used to produce new designs means that only one variable is changed between points on this path, leading to interpretable development information that can help the designer gain physical insight into how the observed performance improvements have been achieved. One form for this development information is presented in Table 9.4, with the variable changed at each step logged alongside the impact on the primary quantities of interest. The optimiser focusses on decreasing L_{ax} , with only five of the 23 steps along the path not resulting in a reduction of this desirable feature. However, the tabular format masks the increase in L_{ax} that occurs moving between designs eight and nine as an additional stage is added, a change that is visible when the values of some overall performance parameters are tracked in Figure 9.18.

Table 9.4 also shows that the optimiser is primarily altering the rear stages of the machine. Only two design variable changes impact the first few stages, with these seen to remain almost identical to the starting geometry in Figure 9.17. This is perhaps unsurprising given that two of the desirable features are exit conditions, with the optimiser improving these values by varying aspects of the compressor that are closer to the exit.

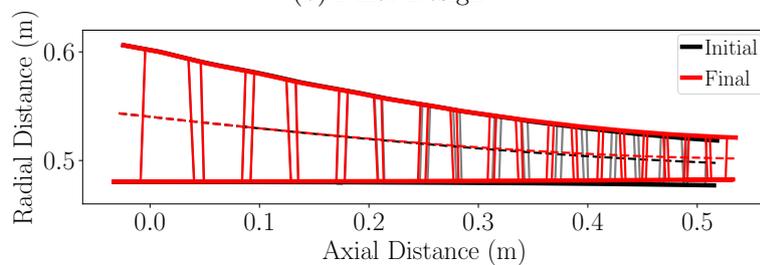
The plots tracking overall performance in Figure 9.18, as well as the progression of individual stage parameters in Figure 9.19, reveal a number of key changes along the path to



(a) Initial Design



(b) Final Design



(c) Initial and Final Designs

Fig. 9.17 Annuli of the initial design and that exhibiting the highest value of η_p found by the MFMDRTS algorithm applied to the N-stage axial compressor test case in the single-objective scenario.

9.5 Information to Enhance Designer Understanding

Table 9.4 Variable changes and their impact on the key performance parameters during development of the design exhibiting the highest value of η_p found by the MFMDRTS algorithm applied to the N-stage axial compressor test case in the single-objective scenario. Green indicates improvement, red indicates worsening. CP refers to a Bézier curve control point and AR to the blade aspect ratio.

Step	Variable(s) Changed	Change Made	η_p	M_{exit}	α_{exit} (°)	L_{ax} (m)
1	-	-	0.87533	0.34728	4.273	0.57049
2	Blade Twist	Increased	0.87561	0.34789	4.244	0.57049
3	3rd and 4th Stage α_{exit} CPs	Decreased	0.87556	0.34725	3.111	0.57045
4	3rd Stage α_{exit} CP	Decreased	0.87559	0.34691	3.117	0.57036
5	2nd and 3rd Area CPs	Decreased	0.88075	0.34863	2.801	0.56801
6	3rd and 4th Stage α_{exit} CPs	Decreased	0.88201	0.34821	1.602	0.56796
7	4th Mean Radius CP	Increased	0.88182	0.34749	1.499	0.56664
8	2nd Stage α_{exit} CP	Decreased	0.88160	0.34745	1.515	0.56653
9	N° Stages	Increased	0.89925	0.31356	2.026	0.60577
10	Blade Twist	Decreased	0.90032	0.31347	2.083	0.60577
11	3rd Stator AR CP	Increased	0.90062	0.31284	2.123	0.60416
12	4th N° Stator Blades CP	Decreased	0.90057	0.31295	2.045	0.60414
13	1st N° Stator Blades CP	Decreased	0.90056	0.31287	2.031	0.60413
14	4th Stage α_{exit} CP	Decreased	0.90066	0.31246	0.899	0.60411
15	4th N° Rotor Blades CP	Decreased	0.90113	0.31303	0.904	0.60390
16	4th Stator AR CP	Increased	0.90195	0.31260	0.743	0.59986
17	4th Area CP	Decreased	0.90111	0.32743	0.514	0.59751
18	4th Rotor AR CP	Increased	0.90202	0.32497	0.602	0.59205
19	2nd Stage α_{exit} CP	Increased	0.90185	0.32537	0.588	0.59217
20	4th Stator AR CP	Increased	0.90224	0.32366	0.542	0.58853
21	3rd Stator AR CP	Increased	0.90265	0.32334	0.516	0.58700
22	3rd Stator AR CP	Increased	0.90241	0.32222	0.523	0.58557
23	4th N° Stator Blades CP	Increased	0.90199	0.32177	0.618	0.58556
24	Blade Twist	Decreased	0.90358	0.32179	0.692	0.58556

Application to an N-Stage Axial Compressor

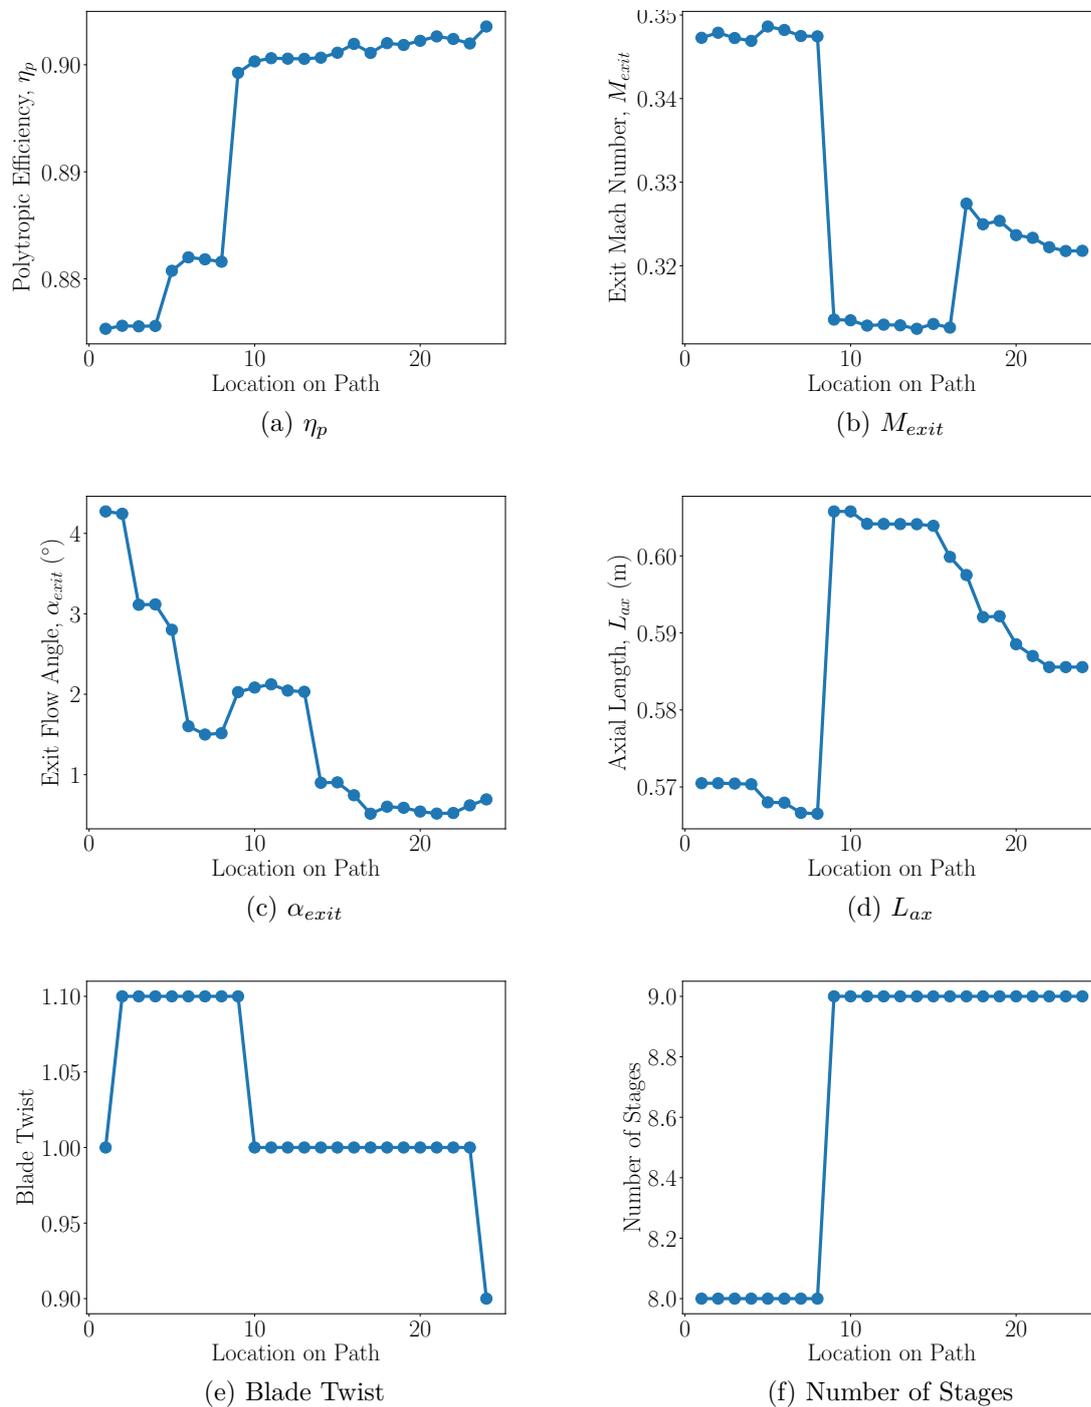


Fig. 9.18 Changes to overall parameters during development of the design exhibiting the highest value of η_p found by the MFMDRTS algorithm applied to the N-stage axial compressor test case in the single-objective scenario.

the final design. The most obvious is the move to nine stages, appropriately made between designs eight and nine. This leads to an increase in η_p and reduction in M_{exit} at the expense of L_{ax} . Figures 9.19a and 9.19b show that ψ and Λ for stage eight drop when the new stage is added downstream, whilst the flow angles at the exits of stages six, seven and eight increase, resulting in smoother variation of this quantity through the machine.

Another significant change occurs between designs 16 and 17 as the annulus area is reduced towards the rear of the machine. This leads to lower values of both α_{exit} and L_{ax} , but an increase in M_{exit} , highlighting the trade-off between these parameters. The area change is seen in Figure 9.19d, with the order of the lines reversing at this point as the hub radius switches to increasing through the machine rather than decreasing as before.

At a more detailed level, Figure 9.20 shows how the aspect ratios and diffusion measures for the rotors and stators develop along the path from the initial to final designs. The rotor aspect ratios increase between designs 17 and 18, leading to lower values of M_{exit} and L_{ax} at the expense of performance in terms of η_p and α_{exit} . The aspect ratios of the stators in all but the first two stages of the machine also increase towards the end of the optimisation as the algorithm seeks to reduce L_{ax} whilst maintaining the benefits of having an additional stage. These increased aspect ratios do not appear to significantly impact the values of DF or DH for these rows. In reality some link would be expected, especially for such large variations in aspect ratio, suggesting that a weakness in the analysis code may have been exposed. This is important to be aware of, and has been highlighted by the design development information made available through the use of the TS algorithm. The ability of the MDR formulation to effectively handle a large number of performance parameters also means that remedying this situation is a simple task for the designer. All that is required is to add an additional desirable feature to the problem definition informing the optimiser to reduce the aspect ratios where possible.

Further observations could be made and insight gained from this rich development information, but the purposes of this section have been fulfilled. The presence of information to enhance designer understanding has been demonstrated, with this data available from the MFMDRTS algorithm just as it was when using the single-fidelity method in Chapter 7. This development information comes at no additional computational cost and helps the designer to determine the physical justification for the observed performance improvements as well as increasing their knowledge of the problem, potentially leading to creativity and innovation.

9.6 Summary

In this chapter the new MDRTS algorithm has been used to facilitate the high-fidelity preliminary design optimisation of an N-stage axial compressor. Given authority over the annulus geometry and number of stages, the optimiser is able to consider numerous quantities

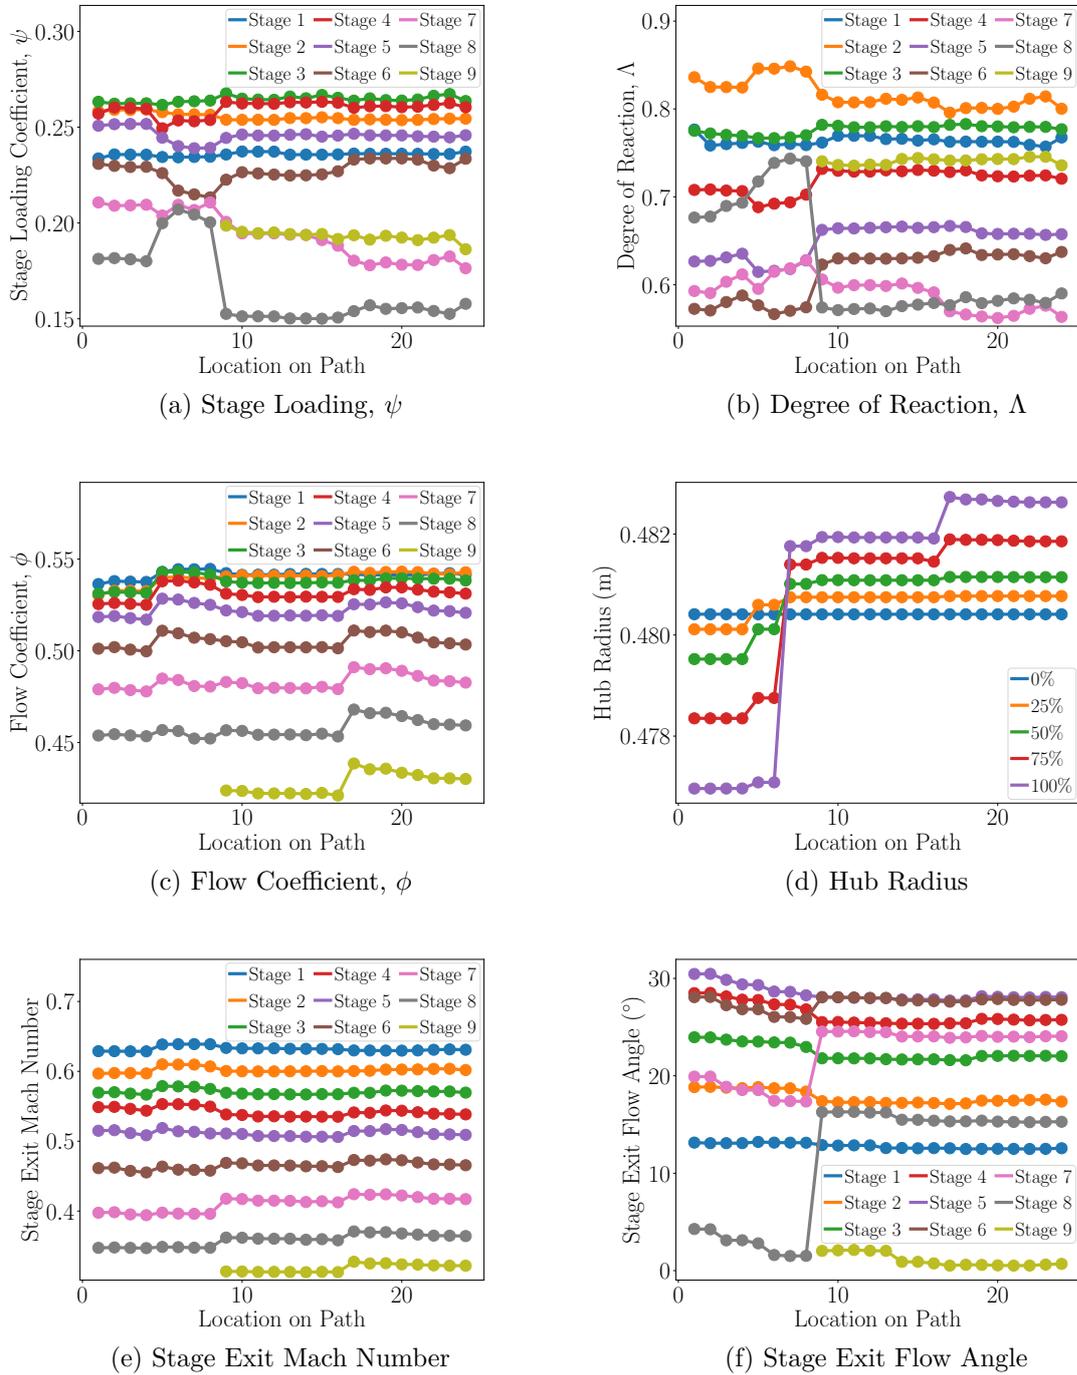


Fig. 9.19 Changes to stage parameters during development of the design exhibiting the highest value of η_p found by the MFMDRTS algorithm applied to the N-stage axial compressor test case in the single-objective scenario.

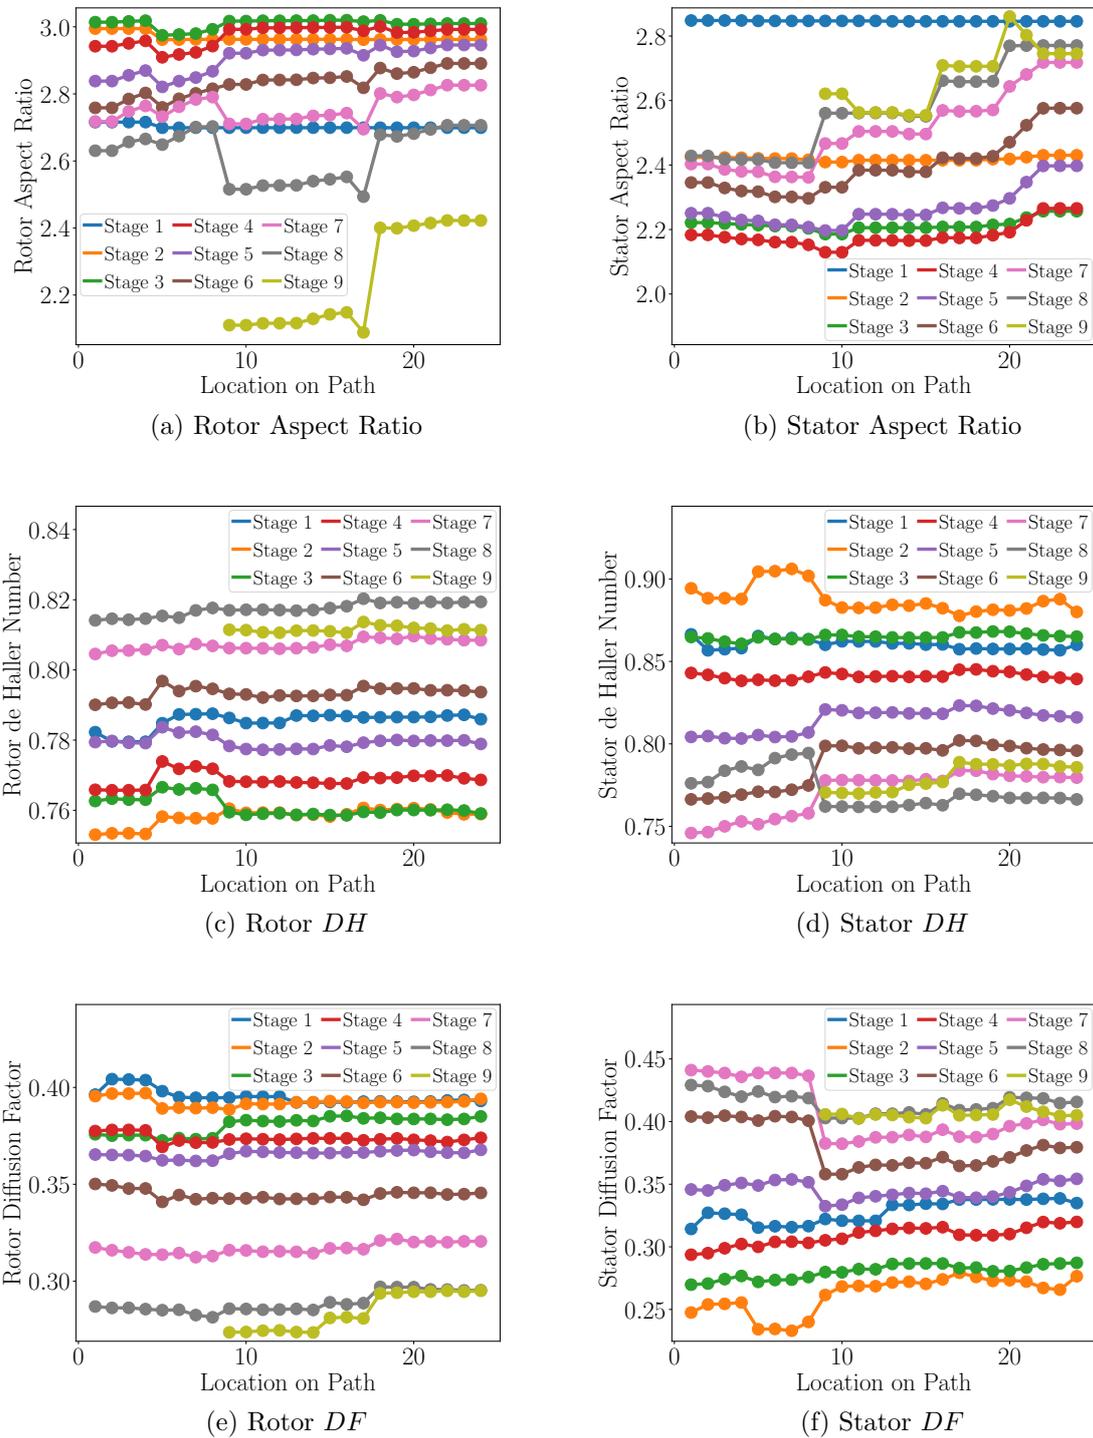


Fig. 9.20 Changes to blade parameters during development of the design exhibiting the highest value of η_p found by the MFMDRTS algorithm applied to the N-stage axial compressor test case in the single-objective scenario.

Application to an N-Stage Axial Compressor

of interest in an efficient and effective manner. The use of physics-based low-fidelity codes accelerates the optimisation, reducing the computational cost by over 90% compared to existing methodologies. Moreover, the presence of design development information for the final designs increases the utility of the optimisation, providing learning opportunities that promote innovation and increase the likelihood of final designs being accepted in an industrial context.

Chapter 10

Conclusion

The primary goal of this thesis has been to facilitate the use of optimisation in the preliminary aerodynamic design of axial compressors by developing an improved methodology that fulfills the following criteria:

- Is capable of efficiently handling a large number of performance parameters, using a problem formulation that accurately reflects the desires of the designer.
- Is fully equipped for the role of enhancing designer understanding.
- Is computationally efficient, making minimal use of expensive high-fidelity analysis whilst retaining sufficient accuracy to ensure that the final designs produced are useful.

Results in the preceding chapters demonstrate that each of these criteria have been met. The MDR formulation enables sophisticated problem definitions to be used, incorporating numerous quantities of interest in a way that accurately reflects the desires of the designer whilst simultaneously ensuring a focussed search that efficiently exploits the available computational budget. The TS algorithm provides rich development information to help designers determine the physical reasoning behind the observed performance improvements, as well as increasing their knowledge of the underlying problem to promote creativity and innovation. Finally, a multi-fidelity technique accelerates the combined MDRTS algorithm, with the inclusion of cheap low-fidelity analysis producing computational savings of over 90% compared to existing methods.

The potential improvement to the design process provided by the new MFMDRTS algorithm is shown schematically in Figure 10.1 using the market share curves discussed by Rubbert [205]. Employing the new method could lead to commercially optimal designs that have greater sales potential for the company as well as exhibiting improved technical performance. These designs would also be reached more quickly, leaving time for further research and development work that could be assisted by the information provided by the new algorithm.

Conclusion

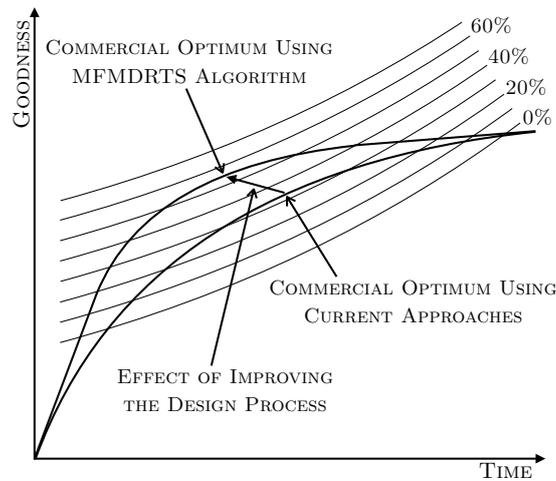


Fig. 10.1 Potential improvement to the design process offered by the MFMDRTS algorithm.

In addition to this high-level achievement the following contributions have been made:

- A TS implementation of the MDR formulation.
- A simplifying framework that facilitates the application of the MDR formulation to more challenging problems without increasing complexity from the perspective of the designer.
- A first application of the MDR formulation to the optimisation of axial compressors, enabling more accurate handling of the relevant performance parameters.
- Use of open source analysis codes in the preliminary design optimisation of axial compressors.
- Extension of the MDR technique to incorporate multiple analysis fidelities, including measures to guard against low-fidelity masking, resulting in a novel multi-level multi-fidelity optimisation methodology.

The following sections outline the main conclusions of the thesis before discussing potential avenues of further research.

10.1 Tabu Search Using Multiple Dominance Relations

The first step towards the improved methodology facilitating the use of optimisation in the preliminary aerodynamic designs of axial compressors was to develop a TS implementation of the MDR formulation. The resultant MDRTS algorithm was shown to outperform both traditional multi-objective TS methods and previous population-based implementations of the MDR formulation when applied to an analytic test problem. The new approach found

designs around two orders of magnitude closer to the theoretical optimum on average than the traditional formulations, and generated performance equivalent to that produced by GA and PSO implementations of the MDR technique with average computational savings of 35% and 18% respectively. When given the same computational budget the new algorithm also generated points three times closer to the known optimum than the PSO and an order of magnitude closer than the GA. These results demonstrate the synergy that exists between the formulation and the search algorithm, with the hierarchical problem definition focussing the optimiser on a smaller region of design space allowing the exploiting capabilities of TS to come to the fore.

10.2 A Simplifying Application Framework

The infancy of the MDR approach necessitated the development of a simplifying framework to enable application of the new formulation to more challenging problems such as the preliminary design of axial compressors. Introducing a third performance parameter classification, known as desirable features, alongside traditional objectives and constraints enables the main benefits of the MDR formulation to be realised without increasing complexity from the perspective of the designer. Separate dominance relations are used to select between designs based on these three classifications.

An experiment conducted using an aerofoil test case informed suggestions for the ordering of these dominance relations within the nested hierarchy that is most likely to produce designs that are of interest to the designer. In both single-objective and trade-off scenarios constraints should appear as the first dominance relation to ensure the optimiser focusses on designs that satisfy the specified limits. When using a single objective this must appear as the final dominance relation in the hierarchy, otherwise subsequent dominance relations will never be consulted during the optimisation. The desirable features then take the second spot in the hierarchy. When multiple objectives are specified in a trade-off scenario the order of the objectives and desirable features should be reversed to place a higher priority on improving the primary quantities of interest.

The simplifying application framework, comprising the additional desirable features performance parameter classification and the dominance relation orderings suggested by the aerofoil test case, enables the use of MDR without increasing complexity from the perspective of the designer compared to traditional objectives-and-constraints approaches. All that is required is for the relevant performance parameters to be assigned to the objectives, constraints or desirable features classifications, and the necessary constraint limits specified. This framework could facilitate the application of the promising MDR formulation to a wider range of more challenging problems, potentially accelerating the uptake of the technique within the engineering design optimisation community.

10.3 Comparison to Existing Methods Using an Aerofoil Test Case

Results generated by the MDRTS algorithm were compared to those produced using existing methods when applied to an aerofoil test case. In both single-objective and trade-off scenarios the new approach consistently outperformed those employing more traditional objectives-and-constraints formulations as well as previously developed population-based implementations of the MDR approach.

In the single-objective scenario the MDRTS algorithm generated average computational savings of 30% compared to a multi-objective TS, and 70% compared to multi-objective GA and PSO approaches. Similar savings of 71% and 67% were observed on average across 10 runs compared to alternative population-based implementations of the MDR formulation, with the new algorithm achieving final values of the overall performance measure over twice as large as the best generated by the existing methods. Even greater computational savings were seen in the trade-off scenario where the traditional formulations struggled to effectively handle the additional performance parameter. The MDRTS algorithm found designs equivalent to those generated by a multi-objective TS method with an average reduction in the number of high-fidelity analyses required of 88%, with savings of just under 80% compared to multi-objective GA and PSO approaches. Similar reductions in computational cost were also demonstrated compared to GA and PSO implementations of the MDR formulation, with average savings of 84% and 72% respectively.

Convergent behaviour of the performance parameters designated as desirable features was also highlighted in both scenarios, with emergent values for these quantities potentially informing suitable constraint limits for use at a later, more detailed stage of the design process. The MDR formulation and simplifying application framework allow limit values that were previously difficult to define to instead become outputs of the optimisation process itself. These emergent constraint limits informed further comparative optimisations, demonstrating that the MDRTS algorithm is capable of producing equivalent performance to traditional constrained methods without the need to specify limit values at the outset of the optimisation.

The provision of design development information was also demonstrated using the aerofoil test case, highlighting the ability of the new algorithm to effectively fulfil the role of enhancing designer understanding. This should increase the likelihood of final designs generated using the optimisation methodology being accepted in an industrial setting as the physical reasoning behind the observed performance improvements can be determined more easily. Creativity and innovation are also promoted through increased knowledge of the underlying problem.

10.4 Initial Six-Stage Axial Compressor Application

A six-stage axial compressor test case provided an initial demonstration of the capabilities of the new MDRTS algorithm when applied to a more challenging design problem. The failure of two population-based search algorithms applied to this test case was shown, highlighting the large proportion of computational budget that can be wasted generating physically incompatible geometries. In contrast, TS consistently produced viable designs, efficiently and effectively navigating the complex design space.

The MDRTS algorithm, applied to this problem using the simplifying framework developed earlier in the thesis, produced designs exhibiting better performance than those found by approaches employing more traditional objectives-and-constraints problem formulations. This was observed in both single-objective and trade-off scenarios, with the ability to handle performance parameters in a way that more accurately reflected the desires of the designer leading to designs that were more likely to be of interest to the designer. In the trade-off scenario computational savings of over 80% were observed compared to existing methods. These results were produced using a high-fidelity analysis system consisting entirely of open source codes, the first time known to the author that this has been achieved in a preliminary axial compressor design optimisation study. The availability of rich design development information was also demonstrated for this more complex problem.

This test case showed that the new MDRTS algorithm had addressed the problems associated with formulation and understanding that limit existing optimisation methodologies when applied to the preliminary design of axial compressors.

10.5 Multi-Fidelity Acceleration

To fulfil the final criterion for the new methodology and facilitate the use of the MDRTS algorithm in an industrial setting the approach was accelerated using multi-fidelity techniques. The new implementation allows different fidelities to be used for different dominance relations, providing computational acceleration by conducting analysis on a “need-to-know” basis, with a particular code only consulted when that level of information is required to make a decision. The resultant algorithm represents a novel multi-level multi-fidelity optimisation methodology that is able to incorporate any type of low-fidelity model, with additions made to reduce the risk of low-fidelity masking.

The computational speed-up provided by the multi-fidelity technique was demonstrated using the analytic, aerofoil and six-stage axial compressor test cases. For the analytic problem separating out the calculation of the different objectives produced an average computational saving of 30%, with this increasing to 74% when low-fidelity models based on response surfaces were employed. In the aerofoil test case consistent computational savings were observed

Conclusion

across multiple runs in both single-objective and trade-off scenarios, with analytic calculation of A_c leading to around a 45% reduction in the number of RANS analyses required, and the inclusion of RBF response surfaces resulting in a larger computational saving of between 66% and 70%. For the six-stage axial compressor problem 31% savings were demonstrated in the single-objective scenario when employing a meanline low-fidelity code, with this rising to 75% when applied in the trade-off study.

The multi-level multi-fidelity technique is able to incorporate more than two levels of fidelity. In the turbomachinery field this allows both meanline and throughflow methods to be employed to accelerate RANS-based optimisation for the first time, with the resultant problem formulation modelling the traditional axial compressor design process. Machines produced using this three-level formulation exhibited worse performance than expected due to inadequacies in the open source throughflow code. However, computational savings of 37% and 45% were still observed compared to the single-fidelity approach in the single-objective and trade-off scenarios respectively. It is expected that further savings would be achieved if a more sophisticated throughflow methodology was employed.

Results from the aerofoil design test case highlighted the importance of careful selection of physics-based low-fidelity models. The inability of the Euler method to model regions of separated flow severely impacts the accelerating capabilities of that low-fidelity code. However, results from the axial compressor test case demonstrated the potential utility of well chosen physics-based models, providing improvement over data-fit response surfaces as the search moves into the exploitation phase.

10.6 Application to an N-Stage Axial Compressor

As a final assessment and demonstration of the capabilities of the new method it was applied to a more complex axial compressor design problem in which the optimiser was given control over the number of stages in the machine. Adding and removing stages during an optimisation is a non-trivial task, potentially resulting in numerous analysis code failures severely restricting the movement of the optimiser through the fragmented design space. This problem therefore acted as a suitable benchmark to assess the ability of the new algorithm to facilitate the use of optimisation in the preliminary aerodynamic design of axial compressors.

Across both single-objective and trade-off scenarios the MFMDRTS method produced designs exhibiting high levels of performance without needing to reduce the number of performance parameters being considered or the fidelity of the analysis. In the single-objective scenario the single-fidelity MDRTS algorithm generated performance equivalent to that found using a traditional multi-objective formulation with a computational saving of over 50%, also finding designs with values of the overall performance measure that were twice as high as those found using a traditional single-objective approach. The multi-fidelity

version of the new algorithm, employing a corrected meanline low-fidelity code, generated further computational savings, equating to 92% and 84% compared to the traditional single- and multi-objective formulations respectively. This multi-fidelity approach also showed a 75% reduction in the number of calls to the high-fidelity analysis compared to the single-fidelity MDRTS algorithm in the early stages of the search. Similar results were seen in the trade-off scenario, with the MDRTS algorithm generating savings of 73% and 56% compared to approaches employing traditional five- and two-objective formulations. These savings increased to 91% and 94% when the multi-fidelity version of the new method was employed, with this approach also generating a computational saving of 78% compared to the single-fidelity MDRTS algorithm.

This performance when applied to such a challenging test case highlights the potential of the new algorithm and demonstrates that an improved methodology facilitating the use of optimisation in the preliminary aerodynamic design of axial compressors has successfully been developed.

10.7 Suggestions for Further Research

Whilst the key goal of this thesis has been achieved in the development of the MFMDRTS algorithm, a number of potential avenues for further research have emerged.

Firstly, despite the merits of employing open source analysis codes, applying this technique in conjunction with more robust and high-performing proprietary analysis programs could be worthwhile. It is expected that similar improvement over current methods would be observed, and in the case of a more sophisticated throughflow methodology even better performance may be possible.

There is also no aspect of the final algorithm that limits applications to the design of axial compressors. Given the performance observed throughout this thesis improvements can be expected when applying the algorithm to alternative problems, both within the turbomachinery field and in the wider aerospace and engineering design communities. Further investigation into the role of emergent constraints in these alternative problem settings is also required to determine whether similar convergent properties exist and to explore potential uses of any emergent values in more depth.

As well as applying the algorithm in alternative physical domains it should also be possible to use it to tackle different types of optimisation problem. As suggested in Chapter 8, the multi-fidelity implementation of the MDR methodology could be used to conduct multidisciplinary and multipoint optimisation studies without any modification being required. There is also a need to extend the current algorithm to incorporate techniques for optimisation under uncertainty. If the methodology is shown to be effective in these scenarios then an

Conclusion

inherent flexibility will have been demonstrated, with the resultant algorithm capable of handling several diverse types of optimisation problem within a single formulation.

Another aspect of the flexibility of the new approach that merits further attention is the ability to incorporate alternatives to the Pareto dominance criterion. These could be used within individual dominance relations to enable an even greater number of performance parameters to be effectively handled by the sophisticated MDR formulation.

Finally, when discussing the importance of understanding in the engineering design process links were drawn to the ideas of explainability and interpretability in the fields of ML and AI. Formalising these links could represent a fruitful avenue of further research culminating in the development of explainable design optimisation techniques that are capable of automatically generating interpretable justifications for the final designs they produce.

References

- [1] ADADI, A. AND BERRADA, M. (2018) Peeking Inside the Black-Box: A Survey on Explainable Artificial Intelligence (XAI). *IEEE Access*, 6:52138–52160. doi:10.1109/ACCESS.2018.2870052.
- [2] AIRBUS (2019). Global Market Forecast 2019-2038. URL: <https://www.airbus.com/aircraft/market/global-market-forecast.html>. Accessed: 23/10/2020.
- [3] AIRBUS (2020). ZEROe: Towards the world’s first zero-emission commercial aircraft. URL: <https://www.airbus.com/innovation/zero-emission/hydrogen/zeroe.html>. Accessed: 23/10/2020.
- [4] AMRIT, A., LEIFSSON, L. AND KOZIEL, S. (2018) Multi-fidelity aerodynamic design trade-off exploration using point-by-point Pareto set identification. *Aerospace Science and Technology*, 79:399–412. doi:10.1016/j.ast.2018.05.023.
- [5] ANTOULAS, A.C. (2005) *Approximation of Large-Scale Dynamical Systems*. SIAM, Philadelphia, PA, USA.
- [6] AUNGIER, R.H. (2003) *Axial-Flow Compressors: A Strategy for Aerodynamic Design and Analysis*. ASME Press, New York, NY, USA.
- [7] AZABI, Y., SAVVARIS, A. AND KIPOUROS, T. (2018) Initial Investigation of Aerodynamic Shape Design Optimisation for the Aegis UAV. *Transportation Research Procedia*, 29:12–22. doi:10.1016/j.trpro.2018.02.002.
- [8] BAERT, L., BEUCAIRE, P., LEBORGNE, M., SAINVITU, C. AND LEPOT, I. (2017) Tackling Highly Constrained Design Problems: Efficient Optimisation of a Highly Loaded Transonic Compressor. In *Turbo Expo: Power for Land, Sea, and Air*. Charlotte, NC, USA. doi:10.1115/GT2017-64610.
- [9] BAERT, L., CHÉRIÈRE, E., SAINVITU, C., LEPOT, I., NOUVELLON, A. AND LEONARDON, V. (2020) Aerodynamic Optimization of the Low-Pressure Turbine Module: Exploiting Surrogate Models in a High-Dimensional Design Space. *Journal of Turbomachinery*, 142(3):031005. doi:10.1115/1.4046232.
- [10] BAERT, L., DUMEUNIER, C., LEBORGNE, M., SAINVITU, C. AND LEPOT, I. (2018) Agile SBO Framework Exploiting Multisimulation Data: Optimising Efficiency and Stall Margin of a Transonic Compressor. In *Turbo Expo: Power for Land, Sea, and Air*. doi:10.1115/GT2018-76639.
- [11] BAHRAMI, S., TRIBES, C., DEVALS, C., VU, T.C. AND GUIBAULT, F. (2016) Multi-fidelity shape optimization of hydraulic turbine runner blades using a multi-objective mesh adaptive direct search algorithm. *Applied Mathematical Modelling*, 40(2):1650–1668. doi:10.1016/j.apm.2015.09.008.

References

- [12] BATISTA, L.S., CAMPELO, F., GUIMARÃES, F.G. AND RAMÍREZ, J.A. (2011) A comparison of dominance criteria in many-objective optimization problems. In *2011 IEEE Congress of Evolutionary Computation (CEC)*, pp. 2359–2366. doi:10.1109/CEC.2011.5949909.
- [13] BECKER, K., LAWRENZ, M., VOSS, C. AND MÖNIG, R. (2008) Multi-Objective Optimization in Axial Compressor Design Using a Linked CFD-Solver. In *Turbo Expo: Power for Land, Sea, and Air*, pp. 2533–2542. Berlin, Germany. doi:10.1115/GT2008-51131.
- [14] BELIE, G. (2002) Non-Technical Barriers to Multidisciplinary Optimization in the Aerospace Industry. In *9th AIAA/ISSMO Symposium on Multidisciplinary Analysis and Optimization*. Atlanta, GA, USA. doi:10.2514/6.2002-5439.
- [15] BELL, T.A., JARRETT, J.P. AND CLARKSON, P.J. (2008) Exploring the Effects of Removing Process-Intrinsic Constraints on Gas Turbine Design. *Journal of Propulsion and Power*, 24(4):751–762. doi:10.2514/1.34092.
- [16] BENINI, E. (2004) Three-Dimensional Multi-Objective Design Optimization of a Transonic Compressor Rotor. *Journal of Propulsion and Power*, 20(3):559–565. doi:10.2514/1.2703.
- [17] BERGH, J., SNEDDEN, G. AND DUNN, D. (2020) Optimization of Non-axisymmetric Endwall Contours for the Rotor of a Low Speed, 12-Stage Research Turbine With Unshrouded Blades—Optimization and Experimental Validation. *Journal of Turbomachinery*, 142(4):041006. doi:10.1115/1.4045988.
- [18] BERGHAMMER, R., FRIEDRICH, T. AND NEUMANN, F. (2012) Convergence of set-based multi-objective optimization, indicators and deteriorative cycles. *Theoretical Computer Science*, 456:2–17. doi:10.1016/j.tcs.2012.05.036.
- [19] BIRAN, O. AND COTTON, C. (2017) Explanation and Justification in Machine Learning: A Survey. In *IJCAI Workshop on Explainable AI (XAI)*. Melbourne, Australia.
- [20] BRACEWELL, R., WALLACE, K., MOSS, M. AND KNOTT, D. (2009) Capturing design rationale. *CAD Computer-Aided Design*, 41(3):173–186. doi:10.1016/j.cad.2008.10.005.
- [21] BRADNER, E., IORIO, F. AND DAVIS, M. (2014) Parameters tell the design story: Ideation and abstraction in design optimization. In *Symposium on Simulation for Architecture and Urban Design*. Tampa, FL, USA.
- [22] BRANDVIK, T. AND PULLAN, G. (2011) An Accelerated 3D Navier–Stokes Solver for Flows in Turbomachines. *Journal of Turbomachinery*, 133(2):021025. doi:10.1115/1.4001192.
- [23] BRANKE, J., KAUSLER, T. AND SCHMECK, H. (2001) Guidance in evolutionary multi-objective optimization. *Advances in Engineering Software*, 32(6):499–507. doi:10.1016/S0965-9978(00)00110-1.
- [24] BRIGGS, W.L., HENSON, V.E. AND MCCORMICK, S.F. (2000) *A Multigrid Tutorial*. SIAM, Philadelphia, PA, USA, 2nd edition.
- [25] BROOKS, C., FORRESTER, A., KEANE, A. AND SHAHPAR, S. (2011) Multi-fidelity Design Optimisation of a Transonic Compressor Rotor. In *9th European Conference on Turbomachinery Fluid Dynamics and Thermodynamics*. Istanbul, Turkey.

-
- [26] BROOMHEAD, D.S. AND LOWE, D. (1988) Multivariable Functional Interpolation and Adaptive Networks. *Complex Systems*, 2:321–355. doi:10.1126/science.1179047.
- [27] BRYSON, D.E. AND RUMPFKEIL, M.P. (2018) Multifidelity Quasi-Newton Method for Design Optimization. *AIAA Journal*, 56(10):4074–4086. doi:10.2514/1.J056840.
- [28] BRYSON, D.E. AND RUMPFKEIL, M.P. (2019) Aerostructural Design Optimization Using a Multifidelity Quasi-Newton Method. *Journal of Aircraft*, 56(5):2019–2031. doi:10.2514/1.C035152.
- [29] CALVERT, W.J. AND GINDER, R.B. (1999) Transonic fan and compressor design. *Proceedings of the Institution of Mechanical Engineers, Part C: Journal of Mechanical Engineering Science*, 213(5):419–436. doi:10.1243/0954406991522671.
- [30] CHAND, S. AND WAGNER, M. (2015) Evolutionary many-objective optimization: A quick-start guide. *Surveys in Operations Research and Management Science*, 20(2):35–42. doi:10.1016/j.sorms.2015.08.001.
- [31] CHATTERJEE, A. (2000) An introduction to the proper orthogonal decomposition. *Current Science*, 78(7):808–817.
- [32] CHATTERJEE, T., CHAKRABORTY, S. AND CHOWDHURY, R. (2019) A Critical Review of Surrogate Assisted Robust Design Optimization. *Archives of Computational Methods in Engineering*, 26(1):245–274. doi:10.1007/s11831-017-9240-5.
- [33] COELLO, C.A.C. (2000) Handling preferences in evolutionary multiobjective optimization: a survey. In *Proceedings of the 2000 Congress on Evolutionary Computation (CEC00)*, volume 1, pp. 30–37. doi:10.1109/CEC.2000.870272.
- [34] COELLO, C.A.C. (2002) Theoretical and numerical constraint-handling techniques used with evolutionary algorithms: a survey of the state of the art. *Computer Methods in Applied Mechanics and Engineering*, 191(11):1245–1287. doi:10.1016/S0045-7825(01)00323-1.
- [35] COELLO, C.A.C., PULIDO, G.T. AND LECHUGA, M.S. (2004) Handling multiple objectives with particle swarm optimization. *IEEE Transactions on Evolutionary Computation*, 8(3):256–279. doi:10.1109/TEVC.2004.826067.
- [36] COLSON, B., MARCOTTE, P. AND SAVARD, G. (2007) An overview of bilevel optimization. *Annals of Operations Research*, 153:235–256. doi:10.1007/s10479-007-0176-2.
- [37] CONN, A.R., GOULD, N.I.M. AND TOINT, P.L. (2000) *Trust-Region Methods*. SIAM, Philadelphia, PA, USA.
- [38] COOK, L.W. AND JARRETT, J.P. (2018) Optimization Using Multiple Dominance Criteria for Aerospace Design Under Uncertainty. *AIAA Journal*, 56(12):4965–4976. doi:10.2514/1.J056951.
- [39] COOK, L.W., WILLCOX, K.E. AND JARRETT, J.P. (2020) Design optimization using multiple dominance relations. *International Journal for Numerical Methods in Engineering*, 121(11):2481–2502. doi:10.1002/nme.6316.
- [40] COOPER, T., REAGAN, I., PORTER, C. AND PRECOURT, C. (2019) *Global Fleet & MRO Market Forecast Commentary 2019-2029*. Oliver Wyman.

References

- [41] CORNE, D.W. AND KNOWLES, J.D. (2007) Techniques for highly multiobjective optimisation. In *Proceedings of the 9th annual conference on Genetic and evolutionary computation - GECCO '07*, pp. 773–780. doi:10.1145/1276958.1277115.
- [42] CUMPSTY, N.A. (1989) *Compressor Aerodynamics*. Longman Scientific and Technical, Harlow, Essex, UK.
- [43] CUMPSTY, N.A. AND GREITZER, E.M. (2004) Ideas and Methods of Turbomachinery Aerodynamics: A Historical View. *Journal of Propulsion and Power*, 20(1):15–26. doi:10.2514/1.9176.
- [44] CVETKOVIC, D. AND PARMEE, I.C. (2002) Preferences and their application in evolutionary multiobjective optimization. *IEEE Transactions on Evolutionary Computation*, 6(1):42–57. doi:10.1109/4235.985691.
- [45] DASKILEWICZ, M.J. AND GERMAN, B.J. (2013) Nondomination-Level Coordinate System for Parameterizing and Exploring Pareto Frontiers. *AIAA Journal*, 51(8):1946–1959. doi:10.2514/1.J052217.
- [46] DAY, I.J. (2015) Stall, Surge, and 75 Years of Research. *Journal of Turbomachinery*, 138(1):011001. doi:10.1115/1.4031473.
- [47] DE HALLER, P. (1953) Das Verhalten von Tragflügelgittern in Axialverdichtern und im Windkanal. *Brennstoff-WärmeKraft*, 5(10):333–337.
- [48] DEB, K., PRATAP, A., AGARWAL, S. AND MEYARIVAN, T. (2002) A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Transactions on Evolutionary Computation*, 6(2):182–197. doi:10.1109/4235.996017.
- [49] DEB, K., SUNDAR, J., RAO N., U.B. AND CHAUDHURI, S. (2006) Reference Point Based Multi-Objective Optimization Using Evolutionary Algorithms. *International Journal of Computational Intelligence Research*, 2(3):273–286.
- [50] DEL ROSARIO, Z., CONSTANTINE, P. AND IACCARINO, G. (2017) Developing Design Insight Through Active Subspaces. In *19th AIAA Non-Deterministic Approaches Conference*. Grapevine, TX, USA. doi:10.2514/6.2017-1090.
- [51] DENTON, J.D. (2010) Some Limitations of Turbomachinery CFD. In *Turbo Expo: Power for Land, Sea, and Air*, pp. 735–745. Glasgow, UK. doi:10.1115/GT2010-22540.
- [52] DENTON, J.D. (2017) Multall—An Open Source, Computational Fluid Dynamics Based, Turbomachinery Design System. *Journal of Turbomachinery*, 139(12):121001. doi:10.1115/1.4037819.
- [53] DENTON, J.D. AND DAWES, W.N. (1998) Computational fluid dynamics for turbomachinery design. *Proceedings of the Institution of Mechanical Engineers, Part C: Journal of Mechanical Engineering Science*, 213(2):107–124. doi:10.1243/0954406991522211.
- [54] DEPARTMENT FOR BUSINESS, ENERGY AND INDUSTRIAL STRATEGY (2020). Energy Trends: March 2020. URL: https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/875381/Energy_Trends_March_2020.pdf. Accessed: 23/10/2020.
- [55] DINH, C.T., MA, S.B. AND KIM, K.Y. (2017) Aerodynamic Optimization of a Single-Stage Axial Compressor with Stator Shroud Air Injection. *AIAA Journal*, 55(8):2739–2754. doi:10.2514/1.J055909.

-
- [56] DORCA-LUQUE, J.M. AND PERROT, V. (2005) Multi-disciplinary System Design Optimisation of a 2-Stage Axial Compressor for Aerodynamic Performance, Production Cost and Dynamic Stability. Technical report, Massachusetts Institute of Technology, No. 16.888 MSDO.
- [57] DORNBERGER, R., BUCHE, D. AND STOLL, P. (2000) Multidisciplinary Optimization in Turbomachinery Design. In *European Congress on Computational Method in Applied Sciences and Engineering (ECCOMAS)*. Barcelona, Spain.
- [58] DU, X. AND LEIFSSON, L. (2019) Optimum aerodynamic shape design under uncertainty by utility theory and metamodeling. *Aerospace Science and Technology*, 95:105464. doi:10.1016/j.ast.2019.105464.
- [59] DUVIGNEAU, R. AND PRAVEEN, C. (2007) Meta-Modeling for Robust Design and Multi-Level Optimization. In *42e Colloque d'Aérodynamique Appliquée, Couplages et Optimisation Multidisciplinaires*. Sophia Antipolis, France.
- [60] EASTWOOD, J.P., GHISU, T. AND JARRETT, J.P. (2013) Tailored Seeding Geometries for the Multi-Fidelity Design of Compression Systems. In *54th AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*. Boston, MA, USA. doi:10.2514/6.2013-1600.
- [61] EBY, D., AVERILL, R.C., PUNCH, W.F. AND GOODMAN, E.D. (1998) Evaluation of injection island GA performance on flywheel design optimisation. In I.C. Parmee (ed.), *Adaptive Computing in Design and Manufacture*. Springer, London. doi:10.1007/978-1-4471-1589-2_10.
- [62] EL-BELTAGY, M.A. AND KEANE, A.J. (1999) A comparison of various optimization algorithms on a multilevel problem. *Engineering Applications of Artificial Intelligence*, 12(5):639–654. doi:10.1016/S0952-1976(99)00033-0.
- [63] ELHAM, A., VAN TOOREN, M.J.L. AND SOBIESZCZANSKI-SOBIESKI, J. (2014) Bilevel Optimization Strategy for Aircraft Wing Design Using Parallel Computing. *AIAA Journal*, 52(8):1770–1783. doi:10.2514/1.J052696.
- [64] ELLBRANT, L., ERIKSSON, L.E. AND MÅRTENSSON, H. (2013) Balancing Efficiency and Stability in the Design of Transonic Compressor Stages. In *Turbo Expo: Power for Land, Sea, and Air*. San Antonio, TX, USA. doi:10.1115/GT2013-94838.
- [65] EMMERICH, M.T.M., GIANNAKOGLU, K.C. AND NAUJOKS, B. (2006) Single- and multiobjective evolutionary optimization assisted by Gaussian random field metamodels. *IEEE Transactions on Evolutionary Computation*, 10(4):421–439. doi:10.1109/TEVC.2005.859463.
- [66] ERIKSSON, D., BINDEL, D. AND SHOEMAKER, C.A. (2019) pySOT and POAP: An event-driven asynchronous framework for surrogate optimization. *arXiv preprint arXiv:1908.00420*.
- [67] FARINA, M. AND AMATO, P. (2004) A fuzzy definition of “optimality” for many-criteria optimization problems. *IEEE Transactions on Systems, Man, and Cybernetics - Part A: Systems and Humans*, 34(3):315–326. doi:10.1109/TSMCA.2004.824873.
- [68] FENG, Z., ZHANG, Q., ZHANG, Q., TANG, Q., YANG, T. AND MA, Y. (2015) A multiobjective optimization based framework to balance the global exploration and local exploitation in expensive optimization. *Journal of Global Optimization*, 61(4):677–694. doi:10.1007/s10898-014-0210-2.

References

- [69] FERNÁNDEZ-GODINO, M.G., PARK, C., KIM, N.H. AND HAFTKA, R.T. (2019) Issues in Deciding Whether to Use Multifidelity Surrogates. *AIAA Journal*, 57(5):2039–2054. doi:10.2514/1.J057750.
- [70] FISCHER, G.R., KIPOUROS, T. AND SAVILL, A.M. (2014) Multi-objective optimisation of horizontal axis wind turbine structure and energy production using aerofoil and blade properties as design variables. *Renewable Energy*, 62:506–515. doi:10.1016/j.renene.2013.08.009.
- [71] FONSECA, C.M. AND FLEMING, P.J. (1993) Genetic Algorithms for Multiobjective Optimization: Formulation, Discussion and Generalization. In *Genetic Algorithms: Proceedings of the Fifth International Conference*. San Mateo, CA, USA.
- [72] FONSECA, C.M. AND FLEMING, P.J. (1998) Multiobjective optimization and multiple constraint handling with evolutionary algorithms. I. A unified formulation. *IEEE Transactions on Systems, Man, and Cybernetics - Part A: Systems and Humans*, 28(1):26–37. doi:10.1109/3468.650319.
- [73] FORRESTER, A.I.J. AND KEANE, A.J. (2009) Recent advances in surrogate-based optimization. *Progress in Aerospace Sciences*, 45:50–79. doi:10.1016/j.paerosci.2008.11.001.
- [74] FRIEDRICH, T., KROEGER, T. AND NEUMANN, F. (2011) Weighted Preferences in Evolutionary Multi-objective Optimization. In *Advances in Artificial Intelligence Lecture Notes in Computer Science vol 7106*. doi:10.1007/978-3-642-25832-9_30.
- [75] GAIER, A., ASTEROTH, A. AND MOURET, J.B. (2017) Aerodynamic Design Exploration through Surrogate-Assisted Illumination. In *18th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*. Denver, CO, USA. doi:10.2514/6.2017-3330.
- [76] GALLIMORE, S.J. (1999) Axial flow compressor design. *Proceedings of the Institution of Mechanical Engineers, Part C: Journal of Mechanical Engineering Science*, 213(5):437–449. doi:10.1243/0954406991522680.
- [77] GEUZAIN, C. AND REMACLE, J.F. (2009) Gmsh: A 3-D finite element mesh generator with built-in pre- and post-processing facilities. *International Journal for Numerical Methods in Engineering*, 79:1309–1331. doi:10.1002/nme.2579.
- [78] GHISU, T., JARRETT, J.P. AND PARKS, G.T. (2011) Robust design optimization of airfoils with respect to ice accretion. *Journal of Aircraft*, 48(1):287–304. doi:10.2514/1.C031100.
- [79] GHISU, T., PARKS, G. AND JARRETT, J. (2008) Integrated Design Optimization of Gas Turbine Compression Systems. In *4th AIAA Multidisciplinary Design Optimization Specialists Conference*. Schaumburg, IL, USA. doi:10.2514/6.2008-1979.
- [80] GHISU, T., PARKS, G.T., JAEGGI, D.M., JARRETT, J.P. AND CLARKSON, P.J. (2010) The benefits of adaptive parametrization in multi-objective Tabu Search optimization. *Engineering Optimization*, 42(10):959–981. doi:10.1080/03052150903564882.
- [81] GHISU, T., PARKS, G.T., JARRETT, J.P. AND CLARKSON, P.J. (2006) Multi-Objective Optimisation of Aero-Engine Compressors. In *3rd International Conference on the Future of Gas Turbine Technology*. Brussels, Belgium.

-
- [82] GHISU, T., PARKS, G.T., JARRETT, J.P. AND CLARKSON, P.J. (2011) An Integrated System for the Aerodynamic Design of Compression Systems—Part I: Development. *Journal of Turbomachinery*, 133(1):011011. doi:10.1115/1.4000534.
- [83] GHISU, T., PARKS, G.T., JARRETT, J.P. AND CLARKSON, P.J. (2011) An Integrated System for the Aerodynamic Design of Compression Systems—Part II: Application. *Journal of Turbomachinery*, 133(1):011012. doi:10.1115/1.4000535.
- [84] GHISU, T., PARKS, G.T., JARRETT, J.P. AND CLARKSON, P.J. (2011) Robust Design Optimization of Gas Turbine Compression Systems. *Journal of Propulsion and Power*, 27(2):282–295. doi:10.2514/1.48965.
- [85] GHOREISHI, S.F. AND ALLAIRE, D. (2019) Multi-information source constrained Bayesian optimization. *Structural and Multidisciplinary Optimization*, 59:977–991. doi:10.1007/s00158-018-2115-z.
- [86] GIANNAKOGLU, K.C., GIOTIS, A.P. AND KARAKASIS, M.K. (2001) Low-cost genetic optimization based on inexact pre-evaluations and the sensitivity analysis of design parameters. *Inverse Problems in Engineering*, 9(4):389–412. doi:10.1080/174159701088027771.
- [87] GILPIN, L.H., BAU, D., YUAN, B.Z., BAJWA, A., SPECTER, M. AND KAGAL, L. (2018) Explaining Explanations: An Overview of Interpretability of Machine Learning. In *IEEE 5th International Conference on Data Science and Advanced Analytics (DSAA)*, pp. 80–89. doi:10.1109/DSAA.2018.00018.
- [88] GLOVER, F. AND LAGUNA, M. (1997) *Tabu Search*. Springer, Boston, MA, USA.
- [89] GOINIS, G. AND NICKE, E. (2016) Optimizing Surge Margin and Efficiency of a Transonic Compressor. In *Turbo Expo: Power for Land, Sea, and Air*. Seoul, South Korea. doi:10.1115/GT2016-57896.
- [90] GOLDBERG, D.E. (1989) *Genetic Algorithms in Search, Optimization, and Machine Learning*. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA.
- [91] GORYACHKIN, E., POPOV, G., BATURIN, O. AND KOLMAKOVA, D. (2015) Three-Stage Low Pressure Compressor Modernization by Means of Optimization Methods. In *Turbo Expo: Power for Land, Sea, and Air*. Montreal, Quebec, Canada. doi:10.1115/GT2015-43384.
- [92] GRAMATYKA, J.J., ERES, M.H., SCANLAN, J.P., MOSS, M., HOLLOWAY, P. AND BATES, R. (2016) A Probabilistic Multi-Fidelity Aero-Engine Preliminary Design Optimization Framework: Technical and Commercial Perspectives. In *17th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*. Washington, D.C., USA. doi:10.2514/6.2016-4289.
- [93] GRATTON, T., GHISU, T., PARKS, G., CAMBULI, F. AND PUDDU, P. (2018) Optimization of blade profiles for the Wells turbine. *Ocean Engineering*, 169:202–214. doi:10.1016/j.oceaneng.2018.08.066.
- [94] GUIDOTTI, R., MONREALE, A., RUGGIERI, S., TURINI, F., GIANNOTTI, F. AND PEDRESCHI, D. (2018) A Survey of Methods for Explaining Black Box Models. *ACM Computing Surveys*, 51(5):93. doi:10.1145/3236009.
- [95] GUTMANN, H. (2001) A Radial Basis Function Method for Global Optimization. *Journal of Global Optimization*, 19(3):201–227. doi:10.1023/A:1011255519438.

References

- [96] HAN, S.Q., SONG, W.P., HAN, Z.H., LI, S.B. AND LIN, Y.F. (2019) Hybrid inverse/optimization design method for rigid coaxial rotor airfoils considering reverse flow. *Aerospace Science and Technology*, 95:105488. doi:10.1016/j.ast.2019.105488.
- [97] HE, X., LI, J. AND MADER, C.A., YILDIRIM, A. AND MARTINS, J.R.R.A. (2019) Robust aerodynamic shape optimization—From a circle to an airfoil. *Aerospace Science and Technology*, 87:48–61. doi:10.1016/j.ast.2019.01.051.
- [98] HE, Y., SUN, J., SONG, P., WANG, X. AND USMANI, A.S. (2020) Preference-driven Kriging-based multiobjective optimization method with a novel multipoint infill criterion and application to airfoil shape design. *Aerospace Science and Technology*, 96:105555. doi:10.1016/j.ast.2019.105555.
- [99] HE, Y., SUN, J., SONG, P., WANG, X. AND XU, D. (2018) Development of a Multi-Objective Preliminary Design Optimization Approach for Axial Flow Compressors. In *Turbo Expo: Power for Land, Sea, and Air*. Oslo, Norway. doi:10.1115/GT2018-76155.
- [100] HEARSEY, R.M. (1989) Numerical Optimization of Axial Compressor Designs. In *Turbo Expo: Power for Land, Sea, and Air*. Toronto, Ontario, Canada. doi:10.1115/89-GT-14.
- [101] HENDLER, M., EXTRA, S., LOCKAN, M., BESTLE, D. AND FLASSIG, P. (2017) Compressor Design in the Context of Holistic Aero Engine Design. In *18th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*. Denver, CO, USA. doi:10.2514/6.2017-3334.
- [102] HERSHEL, M., DIESTELKÄMPER, R. AND BEN LAHMAR, H. (2017) A survey on provenance: What for? What form? What from? *VLDB Journal*, 26(6):881–906. doi:10.1007/s00778-017-0486-1.
- [103] HO, P.Y. AND SHIMIZU, K. (2007) Evolutionary constrained optimization using an addition of ranking method and a percentage-based tolerance value adjustment scheme. *Information Sciences*, 177(14):2985–3004. doi:10.1016/j.ins.2007.01.011.
- [104] HOLT, G. AND BASSLER, S. (1991) Preliminary Design of Axial Compressors Using Artificial Intelligence and Numerical Optimization Techniques. In *Turbo Expo: Power for Land, Sea, and Air*. Orlando, FL, USA. doi:10.1115/91-GT-334.
- [105] HOOKE, R. AND JEEVES, T.A. (1961) “Direct Search” Solution of Numerical and Statistical Problems. *Journal of the ACM*, 8(2):212–229. doi:10.1145/321062.321069.
- [106] HOWELL, A.R. (1945) Fluid Dynamics of Axial Compressors. *Proceedings of the Institution of Mechanical Engineers*, 153(1):441–452. doi:10.1243/PIME_PROC_1945_153_049_02.
- [107] HOYLE, N., BRESSLOFF, N.W. AND KEANE, A.J. (2006) Design Optimization of a Two-Dimensional Subsonic Engine Air Intake. *AIAA Journal*, 44(11):2672–2681. doi:10.2514/1.16123.
- [108] HUANG, D., ALLEN, T.T., NOTZ, W.I. AND MILLER, R.A. (2006) Sequential kriging optimization using multiple-fidelity evaluations. *Structural and Multidisciplinary Optimization*, 32(5):369–382. doi:10.1007/s00158-005-0587-0.
- [109] IKEDA, K., KITA, H. AND KOBAYASHI, S. (2001) Failure of Pareto-based MOEAs: does non-dominated really mean near to optimal? In *Proceedings of the 2001 Congress on Evolutionary Computation*, volume 2, pp. 957–962. doi:10.1109/CEC.2001.934293.

-
- [110] JAEGLI, D., PARKS, G., KIPOUROS, T. AND CLARKSON, P. (2008) The development of a multi-objective Tabu Search algorithm for continuous optimisation problems. *European Journal of Operational Research*, 185(3):1192–1212. doi:10.1016/j.ejor.2006.06.048.
- [111] JAMESON, A. (1988) Aerodynamic design via control theory. *Journal of Scientific Computing*, 3(3):233–260. doi:10.1007/BF01061285.
- [112] JAMESON, A. AND VASSBERG, J. (2001) Computational fluid dynamics for aerodynamic design - Its current and future impact. In *39th Aerospace Sciences Meeting and Exhibit*. Reno, NV, USA. doi:10.2514/6.2001-538.
- [113] JARRETT, J. AND GHISU, T. (2012) An Approach to Multi-fidelity Optimization of Aeroengine Compression Systems. In *14th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*. Indianapolis, IN, USA. doi:10.2514/6.2012-5634.
- [114] JARRETT, J.P., DAWES, W.N. AND CLARKSON, P.J. (2006) An Approach to Integrated Multi-Disciplinary Turbomachinery Design. *Journal of Turbomachinery*, 129(3):488–494. doi:10.1115/1.2472416.
- [115] JARRETT, J.P. AND GHISU, T. (2015) Balancing Configuration and Refinement in the Design of Two-Spool Multistage Compression Systems. *Journal of Turbomachinery*, 137(9):091008. doi:10.1115/1.4030051.
- [116] JARRETT, J.P., GHISU, T. AND PARKS, G.T. (2009) On the Coupling of Designer Experience and Modularity in the Aerothermal Design of Turbomachinery. *Journal of Turbomachinery*, 131(3):031018. doi:10.1115/1.2992513.
- [117] JESCHKE, P., KURZKE, J., SCHABER, R. AND RIEGLER, C. (2002) Preliminary Gas Turbine Design Using the Multidisciplinary Design System MOPEDES. In *Turbo Expo: Power for Land, Sea, and Air*, pp. 75–84. Amsterdam, The Netherlands. doi:10.1115/GT2002-30496.
- [118] JOLY, M., SARKAR, S. AND MEHTA, D. (2019) Machine Learning Enabled Adaptive Optimization of a Transonic Compressor Rotor With Precompression. *Journal of Turbomachinery*, 141(5):051011. doi:10.1115/1.4041808.
- [119] JOLY, M.M., VERSTRAETE, T. AND PANIAGUA, G. (2014) Integrated multifidelity, multidisciplinary evolutionary design optimization of counterrotating compressors. *Integrated Computer-Aided Engineering*, 21(3):249–261. doi:10.3233/ICA-140463.
- [120] JONES, D.R., SCHONLAU, M. AND WELCH, W.J. (1998) Efficient Global Optimization of Expensive Black-Box Functions. *Journal of Global Optimization*, 13:455–492. doi:10.1023/a:1008306431147.
- [121] JÚNIOR, F.D.S., JOSÉ, R. AND BARBOSA, J.R. (2005) Single Objective Optimization of a Multi-Stage Compressor Using a Gradient-Based Method Applied to a Specially Developed Design Computer Code. In *6th World Congress of Structural and Multidisciplinary Optimization*. Rio de Janeiro, Brazil.
- [122] KAMENIK, J., VOUTCHKOV, I., TOAL, D.J.J., KEANE, A.J., HÖGNER, L., MEYER, M. AND BATES, R. (2018) Robust Turbine Blade Optimization in the Face of Real Geometric Variations. *Journal of Propulsion and Power*, 34(6):1479–1493. doi:10.2514/1.B37091.

References

- [123] KAMPOLIS, I.C. AND GIANNAKOGLU, K.C. (2008) A multilevel approach to single- and multiobjective aerodynamic optimization. *Computer Methods in Applied Mechanics and Engineering*, 197(33):2963–2975. doi:10.1016/j.cma.2008.01.015.
- [124] KAMPOLIS, I.C., ZYINARIS, A.S., ASOUTI, V.G. AND GIANNAKOGLU, K.C. (2007) Multilevel optimization strategies based on metamodel-assisted evolutionary algorithms, for computationally expensive problems. In *IEEE Congress on Evolutionary Computation*. Singapore. doi:10.1109/CEC.2007.4425008.
- [125] KANAZAKI, M., TAKAGI, H. AND MAKINO, Y. (2013) Mixed-fidelity efficient global optimization applied to design of supersonic wing. *Procedia Engineering*, 67:85–99. doi:10.1016/j.proeng.2013.12.008.
- [126] KANG, Y.S., PARK, T.C., YANG, S.S., LEE, S.I. AND LEE, D.H. (2012) Multi Disciplinary Design Optimization and Performance Evaluation of a Single-Stage Transonic Axial Compressor. In *Turbo Expo: Power for Land, Sea, and Air*, pp. 361–369. Copenhagen, Denmark. doi:10.1115/GT2012-69252.
- [127] KARAKASIS, M.K. AND GIANNAKOGLU, K.C. (2006) On the use of metamodel-assisted, multi-objective evolutionary algorithms. *Engineering Optimization*, 38(8):941–957. doi:10.1080/03052150600848000.
- [128] KEANE, A.J. (2009) Comparison of Several Optimization Strategies for Robust Turbine Blade Design. *Journal of Propulsion and Power*, 25(5):1092–1099. doi:10.2514/1.38673.
- [129] KEANE, A.J. AND NAIR, P.B. (2005) *Computational Approaches for Aerospace Design: The Pursuit of Excellence*. John Wiley & Sons, Ltd, London, UK.
- [130] KESKIN, A. AND BESTLE, D. (2006) Application of multi-objective optimization to axial compressor preliminary design. *Aerospace Science and Technology*, 10(7):581–589. doi:10.1016/j.ast.2006.03.007.
- [131] KESKIN, A., SWOBODA, M., FLASSIG, P.M., DUTTA, A.K. AND BESTLE, D. (2008) Accelerated Industrial Blade Design Based on Multi-Objective Optimization Using Surrogate Model Methodology. In *Turbo Expo: Power for Land, Sea, and Air*, pp. 2339–2349. Berlin, Germany. doi:10.1115/GT2008-50506.
- [132] KHELGHATIBANA, M., TRÉPANIÉ, J.Y., TRIBES, C. AND NICHOLS, J. (2014) Multi-Objective and Multi-Point Aerodynamic Optimization of Transonic Fan Blades. In *ASME International Mechanical Engineering Congress and Exposition*. Montreal, Quebec, Canada. doi:10.1115/IMECE2014-39079.
- [133] KIM, H. AND LIOU, M.S. (2014) Adaptive directional local search strategy for hybrid evolutionary multiobjective optimization. *Applied Soft Computing Journal*, 19:290–311. doi:10.1016/j.asoc.2014.02.019.
- [134] KIM, J.H., CHOI, K.J. AND KIM, K.Y. (2011) Optimization of a Transonic Axial Compressor Considering Interaction of Blade and Casing Treatment to Improve Operating Stability. In *Turbo Expo: Power for Land, Sea, and Air*, pp. 1129–1138. Vancouver, BC, Canada. doi:10.1115/GT2011-45404.
- [135] KIM, Y., LEE, S. AND YEE, K. (2018) Variable-Fidelity Optimization of Film-Cooling Hole Arrangements Considering Conjugate Heat Transfer. *Journal of Propulsion and Power*, 34(5):1140–1151. doi:10.2514/1.B36880.

-
- [136] KIPOUROS, T., JAEGGI, D., DAWES, W.N., PARKS, G. AND SAVILL, M. (2005) Multi-objective Optimisation of Turbomachinery Blades Using Tabu Search. In *Evolutionary Multi-Criterion Optimization (EMO) Lecture Notes in Computer Science, vol 3410.*, pp. 897–910. doi:10.1007/978-3-540-31880-4_62.
- [137] KIPOUROS, T., JAEGGI, D.M., DAWES, W.N., PARKS, G.T., SAVILL, A.M. AND CLARKSON, P.J. (2008) Biobjective Design Optimization for Axial Compressors Using Tabu Search. *AIAA Journal*, 46(3):701–711. doi:10.2514/1.32794.
- [138] KONTOGIANNIS, S.G., DEMANGE, J., SAVILL, A.M. AND KIPOUROS, T. (2020) A comparison study of two multifidelity methods for aerodynamic optimization. *Aerospace Science and Technology*, 97:105592. doi:10.1016/j.ast.2019.105592.
- [139] KOZIEL, S. AND LEIFSSON, L. (2013) Surrogate-Based Aerodynamic Shape Optimization by Variable-Resolution Models. *AIAA Journal*, 51(1):94–106. doi:10.2514/1.J051583.
- [140] KRIGE, D.G. (1951) A Statistical Approach to Some Basic Mine Valuation Problems on the Witwatersrand. *Journal of the Southern African Institute of Mining and Metallurgy*, 52(6):119–139.
- [141] KROO, I. (2004) Aeronautical Applications of Evolutionary Design. In *Von Karman Institute Lecture Series on Optimization Methods & Tools for Multicriteria/Multidisciplinary Design*. Rhode-Saint-Genèse, Belgium.
- [142] LEE, J. (1997) Design Rationale Systems: Understanding the Issues. *IEEE Intelligent Systems*, 12(03):78–85. doi:10.1109/64.592267.
- [143] LEIFSSON, L.T., KOZIEL, S., TESFAHUNEGN, Y.A., HOSDER, S. AND GRAMANZINI, J.R. (2015) Application of Physics-Based Surrogate Models to Benchmark Aerodynamic Shape Optimization Problems. In *53rd AIAA Aerospace Sciences Meeting*. Kissimmee, FL, USA. doi:10.2514/6.2015-0265.
- [144] LEJON, M., GRÖNSTEDT, T., GLODIC, N., PETRIE-REPAR, P., GENRUP, M. AND MANN, A. (2017) Multidisciplinary Design of a Three Stage High Speed Booster. In *Turbo Expo: Power for Land, Sea, and Air*. Charlotte, NC, USA. doi:10.1115/GT2017-64466.
- [145] LEUSINK, D., ALFANO, D. AND CINNELLA, P. (2015) Multi-fidelity optimization strategy for the industrial aerodynamic design of helicopter rotor blades. *Aerospace Science and Technology*, 42:136–147. doi:10.1016/j.ast.2015.01.005.
- [146] LI, B., GU, C., LI, X. AND LIU, T. (2016) Numerical optimization for stator vane settings of multi-stage compressors based on neural networks and genetic algorithms. *Aerospace Science and Technology*, 52:81–94. doi:10.1016/j.ast.2016.02.024.
- [147] LI, B., LI, J., TANG, K. AND YAO, X. (2015) Many-Objective Evolutionary Algorithms: A Survey. *ACM Computing Surveys*, 48(1):13. doi:10.1145/2792984.
- [148] LI, J., BOUHLEL, M.A. AND MARTINS, J.R.R.A. (2019) Data-Based Approach for Fast Airfoil Analysis and Optimization. *AIAA Journal*, 57(2):581–596. doi:10.2514/1.J057129.
- [149] LI, J. AND CAI, J. (2020) Massively Multipoint Aerodynamic Shape Design via Surrogate-Assisted Gradient-Based Optimization. *AIAA Journal*, 58(5):1949–1963. doi:10.2514/1.J058491.

References

- [150] LI, K., CHEN, R., MIN, G. AND YAO, X. (2018) Integration of Preferences in Decomposition Multiobjective Optimization. *IEEE Transactions on Cybernetics*, 48(12):3359–3370. doi:10.1109/TCYB.2018.2859363.
- [151] LI, L., YUAN, T., LI, Y., YANG, W. AND KANG, J. (2019) Multidisciplinary Design Optimization Based on Parameterized Free-Form Deformation for Single Turbine. *AIAA Journal*, 57(5):2075–2087. doi:10.2514/1.J057819.
- [152] LI, M., LI, G. AND AZARM, S. (2008) A Kriging Metamodel Assisted Multi-Objective Genetic Algorithm for Design Optimization. *Journal of Mechanical Design*, 130(3):031401. doi:10.1115/1.2829879.
- [153] LI, M., YANG, S. AND LIU, X. (2015) Bi-goal evolution for many-objective optimization problems. *Artificial Intelligence*, 228:45–65. doi:10.1016/j.artint.2015.06.007.
- [154] LI, Z. AND LIU, Y. (2018) Optimization of rough transonic axial compressor. *Aerospace Science and Technology*, 78:12–25. doi:10.1016/j.ast.2018.03.031.
- [155] LI, Z., LIU, Y. AND AGARWAL, R.K. (2018) Robust Optimization Design of Single-Stage Transonic Axial Compressor Considering the Manufacturing Uncertainties. In *Turbo Expo: Power for Land, Sea, and Air*. Oslo, Norway. doi:10.1115/GT2018-75415.
- [156] LI, Z. AND ZHENG, X. (2017) Review of design optimization methods for turbomachinery aerodynamics. *Progress in Aerospace Sciences*, 93:1–23. ISSN 0376-0421. doi:10.1016/j.paerosci.2017.05.003.
- [157] LIEBLEIN, S., SCHWENK, F.C. AND BRODERICK, R.L. (1953) Diffusion factor for estimating losses and limiting blade loadings in axial flow compressor blade elements. Technical report, NACA.
- [158] LIM, H.W. AND KIM, H. (2019) Multi-objective airfoil shape optimization using an adaptive hybrid evolutionary algorithm. *Aerospace Science and Technology*, 87:141–153. doi:10.1016/j.ast.2019.02.016.
- [159] LIPTON, Z.C. (2018) The mythos of model interpretability. *Communications of the ACM*, 61(10):35–43. doi:10.1145/3233231.
- [160] LIU, B., KOZIEL, S. AND ZHANG, Q. (2016) A multi-fidelity surrogate-model-assisted evolutionary algorithm for computationally expensive optimization problems. *Journal of Computational Science*, 12:28–37. doi:10.1016/j.jocs.2015.11.004.
- [161] LUO, C., SONG, L., LI, J. AND FENG, Z. (2012) A Study on Multidisciplinary Optimization of an Axial Compressor Blade Based on Evolutionary Algorithms. *Journal of Turbomachinery*, 134(5):054501. doi:10.1115/1.4003817.
- [162] LUO, J. AND LIU, F. (2015) Multi-Objective Optimization of a Transonic Compressor Rotor by Using an Adjoint Method. *AIAA Journal*, 53(3):797–801. doi:10.2514/1.J053436.
- [163] MA, C., SU, X. AND YUAN, X. (2016) An Efficient Unsteady Adjoint Optimization System for Multistage Turbomachinery. *Journal of Turbomachinery*, 139(1):011003. doi:10.1115/1.4034185.
- [164] MALEKI, H. (2008) Multi-Objective Optimization of an Axial Compressor. In *Fluids Engineering Division Summer Meeting*, pp. 135–139. Jacksonville, FL, USA. doi:10.1115/FEDSM2008-55015.

-
- [165] MARCH, A. AND WILLCOX, K. (2012) Constrained multifidelity optimization using model calibration. *Structural and Multidisciplinary Optimization*, 46(1):93–109. doi:10.1007/s00158-011-0749-1.
- [166] MARCH, A. AND WILLCOX, K. (2012) Provably Convergent Multifidelity Optimization Algorithm Not Requiring High-Fidelity Derivatives. *AIAA Journal*, 50(5):1079–1089. doi:10.2514/1.J051125.
- [167] MARTIN, I., HARTWIG, L. AND BESTLE, D. (2019) A multi-objective optimization framework for robust axial compressor airfoil design. *Structural and Multidisciplinary Optimization*, 59(6):1935–1947. doi:10.1007/s00158-018-2164-3.
- [168] MARTINS, J.R.R.A. AND LAMBE, A.B. (2013) Multidisciplinary Design Optimization: A Survey of Architectures. *AIAA Journal*, 51(9):2049–2075. doi:10.2514/1.J051895.
- [169] MASSARDO, A. AND SATTA, A. (1990) Axial Flow Compressor Design Optimization: Part I—Pitchline Analysis and Multivariable Objective Function Influence. *Journal of Turbomachinery*, 112(3):399–404. doi:10.1115/1.2927673.
- [170] MASSARDO, A., SATTA, A. AND MARINI, M. (1990) Axial Flow Compressor Design Optimization: Part II—Throughflow Analysis. *Journal of Turbomachinery*, 112(3):405–410. doi:10.1115/1.2927674.
- [171] MCCULLERS, L.A. (1984) Aircraft Configuration Optimization Including Optimized Flight Profiles. Technical report, NASA Langley Research Center Recent Experiences in Multidisciplinary Analysis and Optimization, Part 1.
- [172] MESSAC, A. (1996) Physical programming - Effective optimization for computational design. *AIAA Journal*, 34(1):149–158. doi:10.2514/3.13035.
- [173] MESSAC, A. (2000) From Dubious Construction of Objective Functions to the Application of Physical Programming. *AIAA Journal*, 38(1):155–163. doi:10.2514/2.936.
- [174] MEZURA-MONTES, E. AND COELLO, C.A.C. (2011) Constraint-handling in nature-inspired numerical optimization: Past, present and future. *Swarm and Evolutionary Computation*, 1(4):173–194. doi:10.1016/j.swevo.2011.10.001.
- [175] MOHAMMADI, A., OMIDVAR, M.N., LI, X. AND DEB, K. (2014) Integrating user preferences and decomposition methods for many-objective optimization. In *2014 IEEE Congress on Evolutionary Computation (CEC)*, pp. 421–428. doi:10.1109/CEC.2014.6900595.
- [176] MOLINARI, M. AND DAWES, W.N. (2006) Review of evolution of compressor design process and future perspectives. *Proceedings of the Institution of Mechanical Engineers, Part C: Journal of Mechanical Engineering Science*, 220(6):761–771. doi:10.1243/09544062JMES298.
- [177] MOUSTAPHA, M. AND SUDRET, B. (2019) Surrogate-assisted reliability-based design optimization: a survey and a unified modular framework. *Structural and Multidisciplinary Optimization*, 60(5):2157–2176. doi:10.1007/s00158-019-02290-y.
- [178] NGUYEN, A.Q., WAGNER, M. AND NEUMANN, F. (2014) User Preferences for Approximation-Guided Multi-objective Evolution. In *Simulated Evolution and Learning Lecture Notes in Computer Science vol 8886*. doi:10.1007/978-3-319-13563-2_22.

References

- [179] NING, T., GU, C., LI, X. AND LIU, T. (2016) Three-Dimensional Aerodynamic Optimization of a Multi-Stage Axial Compressor. In *Turbo Expo: Power for Land, Sea, and Air*. Seoul, South Korea. doi:10.1115/GT2016-57626.
- [180] ONG, Y.S., NAIR, P.B. AND KEANE, A.J. (2003) Evolutionary Optimization of Computationally Expensive Problems via Surrogate Modeling. *AIAA Journal*, 41(4):687–696. doi:10.2514/2.1999.
- [181] OYAMA, A. AND LIOU, M.S. (2002) Multiobjective Optimization of a Multi-Stage Compressor Using Evolutionary Algorithm. In *38th AIAA/ASME/SAE/ASEE Joint Propulsion Conference & Exhibit*. Indianapolis, IN, USA. doi:10.2514/6.2002-3535.
- [182] PALACIOS, F., A., J., DURAISAMY, K., COLONNO, M., HICKEN, J., ARANAKE, A., CAMPOS, A., COPELAND, S., ECONOMON, T., LONKAR, A., LUKACZYK, T. AND TAYLOR, T. (2013) Stanford University Unstructured (SU²): An open-source integrated computational environment for multi-physics simulation and design. In *51st AIAA Aerospace Sciences Meeting*. Grapevine, TX, USA. doi:10.2514/6.2013-287.
- [183] PANCHENKO, Y., MOUSTAPHA, H., MAH, S., PATEL, K., DOWHAN, M.J. AND HALL, D. (2002) Preliminary Multi-Disciplinary Optimization in Turbomachinery Design. In *Symposium on “Reduction of Military Vehicle Acquisition Time and Cost Through Advanced Modeling and Visual Simulation”*. Paris, France.
- [184] PAPAGEORGIOU, A. AND ÖLVANDER, J. (2017) The role of multidisciplinary design optimization (MDO) in the development process of complex engineering products. In *Proceedings of the International Conference on Engineering Design (ICED)*, pp. 109–118. Vancouver, Canada.
- [185] PARK, K., TURNER, M.G., SIDDAPPAJI, K., DEY, S. AND MERCHANT, A. (2011) Optimization of a 3-Stage Booster: Part 1—The Axisymmetric Multi-Disciplinary Optimization Approach to Compressor Design. In *Turbo Expo: Power for Land, Sea, and Air*, pp. 1413–1422. Vancouver, British Columbia, Canada. doi:10.1115/GT2011-46569.
- [186] PARKINSON, A.R., BALLING, R. AND HEDENGREN, J.D. (2018) *Optimization Methods for Engineering Design*. Brigham Young University, 2nd edition.
- [187] PASQUALE, D., PERSICO, G. AND REBAY, S. (2013) Optimization of Turbomachinery Flow Surfaces Applying a CFD-Based Throughflow Method. *Journal of Turbomachinery*, 136(3):031013. doi:10.1115/1.4024694.
- [188] PASTVA, T.A. (1998) *Bézier Curve Fitting*. Master’s thesis, Naval Postgraduate School, Monterey, CA, USA.
- [189] PEHERSTORFER, B., WILLCOX, K. AND GUNZBURGER, M. (2018) Survey of Multifidelity Methods in Uncertainty Propagation, Inference, and Optimization. *SIAM Review*, 60(3):550–591. doi:10.1137/16M1082469.
- [190] 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. doi:10.1115/1.4043252.
- [191] PHILLIPS, S. AND JARRETT, J.P. Assessing the Significance of Nesting Order in Optimization Using Multiple Dominance Relations. In *AIAA AVIATION 2020 Forum*. doi:10.2514/6.2020-3134.

-
- [192] PHILLIPS, S. AND JARRETT, J.P. Enhancing Designer Understanding by Combining Multiple Dominance Relations and Tabu Search. In *AIAA Scitech 2020 Forum*. Orlando, FL, USA. doi:10.2514/6.2020-0161.
- [193] POEHLMANN, F. AND BESTLE, D. (2012) Multi-Objective Compressor Design Optimization Using Multi-Design Transfer Between Codes of Different Fidelity. In *Turbo Expo: Power for Land, Sea, and Air*, pp. 2001–2010. Copenhagen, Denmark. doi:10.1115/GT2012-68577.
- [194] POPOV, G., GORIACHKIN, E., KOLMAKOVA, D. AND NOVIKOVA, Y. (2016) Multicriteria Optimization of Axial Low Pressure Compressor of Gas Turbine Power Plant. In *Turbo Expo: Power for Land, Sea, and Air*. Seoul, South Korea. doi:10.1115/GT2016-57856.
- [195] PRAVEEN, C. AND DUVIGNEAU, R. (2009) Low cost PSO using metamodels and inexact pre-evaluation: Application to aerodynamic shape design. *Computer Methods in Applied Mechanics and Engineering*, 198:1087–1096. doi:10.1016/j.cma.2008.11.019.
- [196] PUENTE, R., CORRAL, R. AND PARRA, J. (2018) Comparison between aerodynamic designs obtained by human driven and automatic procedures. *Aerospace Science and Technology*, 72:443–454. doi:10.1016/j.ast.2017.11.011.
- [197] RACHMAWATI, L. AND SRINIVASAN, D. (2006) Preference Incorporation in Multi-objective Evolutionary Algorithms: A Survey. In *2006 IEEE International Conference on Evolutionary Computation*, pp. 962–968. doi:10.1109/CEC.2006.1688414.
- [198] RAY, T., SINGH, H.K., ISAACS, A. AND SMITH, W. (2009) Infeasibility driven evolutionary algorithm for constrained optimization. In E. Mezura-Montes (ed.), *Constraint-Handling in Evolutionary Optimization. Studies in Computational Intelligence, vol 198.*, pp. 145–165. Springer, Berlin, Heidelberg. doi:10.1007/978-3-642-00619-7_7.
- [199] REGIS, R.G. (2014) Constrained optimization by radial basis function interpolation for high-dimensional expensive black-box problems with infeasible initial points. *Engineering Optimization*, 46(2):218–243. doi:10.1080/0305215X.2013.765000.
- [200] REGIS, R.G. (2016) Trust regions in Kriging-based optimization with expected improvement. *Engineering Optimization*, 48(6):1037–1059. doi:10.1080/0305215X.2015.1082350.
- [201] REITENBACH, S., SCHNÖS, M., BECKER, R.G. AND OTTEN, T. (2015) Optimization of Compressor Variable Geometry Settings Using Multi-Fidelity Simulation. In *Turbo Expo: Power for Land, Sea, and Air*. Montreal, Quebec, Canada. doi:10.1115/GT2015-42832.
- [202] RODRIGUES, S.S. AND MARTA, A.C. (2020) Adjoint-based shape sensitivity of multi-row turbomachinery. *Structural and Multidisciplinary Optimization*, 61(2):837–853. doi:10.1007/s00158-019-02386-5.
- [203] ROUSE, W.B. (1986) A note on the nature of creativity in engineering: Implications for supporting system design. *Information Processing and Management*, 22(4):279–285. doi:10.1016/0306-4573(86)90026-9.
- [204] ROUSE, W.B. (1986) On the value of information in system design: A framework for understanding and aiding designers. *Information Processing and Management*, 22(3):217–228. doi:10.1016/0306-4573(86)90054-3.

References

- [205] RUBBERT, P.E. (1994) AIAA Wright Brothers Lecture: CFD and the Changing World of Airplane Design. In *19th Congress of the International Council of the Aeronautical Sciences*. Anaheim, CA, USA.
- [206] RUIZ, A.B., SABORIDO, R. AND LUQUE, M. (2015) A preference-based evolutionary algorithm for multiobjective optimization: the weighting achievement scalarizing function genetic algorithm. *Journal of Global Optimization*, 62(1):101–129. doi:10.1007/s10898-014-0214-y.
- [207] RUMPFKEIL, M.P. AND BERAN, P. (2017) Construction of Dynamic Multifidelity Locally Optimized Surrogate Models. *AIAA Journal*, 55(9):3169–3179. doi:10.2514/1.J055834.
- [208] SACKS, J., WELCH, W.J., MITCHELL, T.J. AND WYNN, H.P. (1989) Design and Analysis of Computer Experiments. *Statistical Science*, 4(4):409–435. doi:10.1214/ss/1177012413.
- [209] SAMAD, A. AND KIM, K.Y. (2009) Surrogate Based Optimization Techniques for Aerodynamic Design of Turbomachinery. *International Journal of Fluid Machinery and Systems*, 2(2):179–188. doi:10.5293/ijfms.2009.2.2.179.
- [210] SARAVANAMUTTOO, H.I.H., COHEN, H. AND ROGERS, G.F.C. (2009) *Gas Turbine Theory*. Prentice Hall/Pearson Education, Harlow, Essex, UK, 6th edition.
- [211] SCHEMMANN, C., GELLER, M. AND KLUCK, N. (2018) A Multi-Fidelity Sampling Method for Efficient Design and Optimization of Centrifugal Compressor Impellers. In *Turbo Expo: Power for Land, Sea, and Air*. Oslo, Norway. doi:10.1115/GT2018-75160.
- [212] SCHWEITZER, J.K. AND GARBEROGLIO, J.E. (1984) Maximum loading capability of axial flow compressors. *Journal of Aircraft*, 21(8):593–600. doi:10.2514/3.45028.
- [213] SCHÜTZE, O., CUATE, O., MARTÍN, A., PEITZ, S. AND DELLNITZ, M. (2020) Pareto Explorer: a global/local exploration tool for many-objective optimization problems. *Engineering Optimization*, 52(5):832–855. doi:10.1080/0305215X.2019.1617286.
- [214] SHAHPAR, S. (2011) Challenges to overcome for routine usage of automatic optimisation in the propulsion industry. *The Aeronautical Journal*, 115(1172):615–625. doi:10.1017/S0001924000006308.
- [215] SHU, L., JIANG, P., ZHOU, Q. AND XIE, T. (2019) An online variable-fidelity optimization approach for multi-objective design optimization. *Structural and Multidisciplinary Optimization*, 60(3):1059–1077. doi:10.1007/s00158-019-02256-0.
- [216] SILLER, U., VOSS, C. AND NICKE, E. (2009) Automated Multidisciplinary Optimization of a Transonic Axial Compressor. In *47th AIAA Aerospace Sciences Meeting including The New Horizons Forum and Aerospace Exposition*. Orlando, FL, USA. doi:10.2514/6.2009-863.
- [217] SIMPSON, T.W., MAUERY, T.M., KORTE, J. AND MISTREE, F. (2001) Kriging models for global approximation in simulation-based multidisciplinary design optimization. *AIAA Journal*, 39(12):2233–2241. doi:10.2514/3.15017.
- [218] SINGH, G. AND GRANDHI, R.V. (2010) Mixed-Variable Optimization Strategy Employing Multifidelity Simulation and Surrogate Models. *AIAA Journal*, 48(1):215–223. doi:10.2514/1.43469.

- [219] SINHA, A., MALO, P. AND DEB, K. (2018) A Review on Bilevel Optimization: From Classical to Evolutionary Approaches and Applications. *IEEE Transactions on Evolutionary Computation*, 22(2):276–295. doi:10.1109/TEVC.2017.2712906.
- [220] SOBIECZKY, H. (1999) Parametric Airfoils and Wings. In K. Fujii and G.S. Dulikravich (eds.), *Recent Development of Aerodynamic Design Methodologies. Notes on Numerical Fluid Mechanics (NNFM)*, vol. 65. Vieweg+Teubner Verlag. doi:10.1007/978-3-322-89952-1_4.
- [221] SONG, L., LUO, C., LI, J. AND FENG, Z. (2012) Automated multi-objective and multidisciplinary design optimization of a transonic turbine stage. *Proceedings of the Institution of Mechanical Engineers, Part A: Journal of Power and Energy*, 226(2):262–276. doi:10.1177/0957650911425005.
- [222] SONG, W., ZHANG, Y., CHEN, H. AND DENG, K. (2019) Transonic Compressor Blade Optimization Integrated With Circumferential Groove Casing Treatment. *Journal of Turbomachinery*, 141(3):031015. doi:10.1115/1.4041699.
- [223] SPALART, P.R. AND VENKATAKRISHNAN, V. (2016) On the role and challenges of CFD in the aerospace industry. *Aeronautical Journal*, 120(1223):209–232. doi:10.1017/aer.2015.10.
- [224] SZYKMAN, S., SRIRAM, R.D. AND REGLI, W.C. (2001) The Role of Knowledge in Next-generation Product Development Systems. *Journal of Computing and Information Science in Engineering*, 1(1):3–11. doi:10.1115/1.1344238.
- [225] TAGHAVI-ZENOUEZ, R. AND AFZALI, S.M. (2008) Preliminary design optimization of profile losses in multi-stage axial compressors based on complex method. *Proceedings of the Institution of Mechanical Engineers, Part G: Journal of Aerospace Engineering*, 222(7):951–961. doi:10.1243/09544100JAERO372.
- [226] TESFAHUNEGN, Y.A., KOZIEL, S. AND LEIFSSON, L.T. (2015) Surrogate-Based Airfoil Design with Multi-Level Optimization and Adjoint Sensitivity. In *53rd AIAA Aerospace Sciences Meeting*. Kissimmee, FL, USA. doi:10.2514/6.2015-0759.
- [227] THE BOEING COMPANY (2020). Commercial Market Outlook 2020-2039. URL: <https://www.boeing.com/commercial/market/commercial-market-outlook>. Accessed: 23/10/2020.
- [228] TOAL, D.J.J. AND KEANE, A.J. (2011) Efficient Multipoint Aerodynamic Design Optimization Via Cokriging. *Journal of Aircraft*, 48(5):1685–1695. doi:10.2514/1.C031342.
- [229] TOAL, D.J.J., KEANE, A.J., BENITO, D., DIXON, J.A., YANG, J., PRICE, M., ROBINSON, T., REMOUCHAMPS, A. AND KILL, N. (2014) Multifidelity Multidisciplinary Whole-Engine Thermomechanical Design Optimization. *Journal of Propulsion and Power*, 30(6):1654–1666. doi:10.2514/1.B35128.
- [230] TRAPANI, G., KIPOUROS, T. AND SAVILL, A.M. (2012) The Design of Multi-Element Airfoils Through Multi-Objective Optimization Techniques. *CMES - Computer Modelling in Engineering & Sciences*, 88(2):107–140. doi:10.3970/cmes.2012.088.107.
- [231] TUCCILLO, R. (1989) A Proposal for Optimized Design of Multistage Compressors. In *Turbo Expo: Power for Land, Sea, and Air*. Toronto, Ontario, Canada. doi:10.1115/89-GT-34.

References

- [232] TURNER, M.G., PARK, K., SIDDAPPAJI, K., DEY, S., GUTZWILLER, D.P., MERCHANT, A. AND BRUNA, D. (2010) Framework for Multidisciplinary Optimization of Turbomachinery. In *Turbo Expo: Power for Land, Sea, and Air*, pp. 623–631. Glasgow, UK. doi:10.1115/GT2010-22228.
- [233] VASSBERG, J. AND JAMESON, A. (2016) Industrial applications of aerodynamic shape optimization. In *Von Karman Institute Lecture Series on Introduction to Optimization and Multidisciplinary Design*.
- [234] VASSBERG, J. AND JAMESON, A. (2016) Theoretical Background for Aerodynamic Shape Optimization. In *Von Karman Institute Lecture Series on Introduction to Optimization and Multidisciplinary Design*.
- [235] VELASQUEZ, M. AND HESTER, P.T. (2013) An Analysis of Multi-Criteria Decision Making Methods. *International Journal of Operations Research*, 10(2):56–66.
- [236] VELLIDO, A. (2020) The importance of interpretability and visualization in machine learning for applications in medicine and health care. *Neural Computing and Applications*, 32:18069–18083. doi:10.1007/s00521-019-04051-w.
- [237] VITALE, S., PINI, M. AND COLONNA, P. (2020) Multistage Turbomachinery Design Using the Discrete Adjoint Method Within the Open-Source Software SU2. *Journal of Propulsion and Power*, 36(3):465–478. doi:10.2514/1.B37685.
- [238] WALSH, P.P. AND FLETCHER, P. (2004) *Gas Turbine Performance*. John Wiley & Sons, Ltd, 2nd edition.
- [239] WALTHER, B. AND NADARAJAH, S. (2013) Constrained Adjoint-Based Aerodynamic Shape Optimization of a Single-Stage Transonic Compressor. *Journal of Turbomachinery*, 135(2):021017. doi:10.1115/1.4007502.
- [240] WALTHER, B. AND NADARAJAH, S. (2015) Adjoint-Based Constrained Aerodynamic Shape Optimization for Multistage Turbomachines. *Journal of Propulsion and Power*, 31(5):1298–1319. doi:10.2514/1.B35433.
- [241] WALTHER, B. AND NADARAJAH, S. (2015) Optimum Shape Design for Multirow Turbomachinery Configurations Using a Discrete Adjoint Approach and an Efficient Radial Basis Function Deformation Scheme for Complex Multiblock Grids. *Journal of Turbomachinery*, 137(8):081006. doi:10.1115/1.4029550.
- [242] WANG, L., DISKIN, B., BIEDRON, R.T., NIELSEN, E.J. AND BAUCHAU, O.A. (2020) Evaluation of High-Fidelity Multidisciplinary Sensitivity-Analysis Framework for Multipoint Rotorcraft Optimization. *Journal of Aircraft*, 57(5):830–842. doi:10.2514/1.C035748.
- [243] WANKHEDE, M.J., BRESSLOFF, N.W. AND KEANE, A.J. (2013) Efficient Strategy for Low NO_x Combustor Design in the Spatial Domain Using Multi-Fidelity Solutions. In *Turbo Expo: Power for Land, Sea, and Air*. San Antonio, TX, USA. doi:10.1115/GT2013-94317.
- [244] WILD, S.M., REGIS, R.G. AND SHOEMAKER, C.A. (2008) ORBIT: Optimization by Radial Basis Function Interpolation in Trust-Regions. *SIAM Journal on Scientific Computing*, 30(6):3197–3219. doi:10.1137/070691814.

-
- [245] YANG, J., LUO, J., XIONG, J. AND LIU, F. (2016) Aerodynamic Design Optimization of the Last Stage of a Multi-Stage Compressor by Using an Adjoint Method. In *Turbo Expo: Power for Land, Sea, and Air*. Seoul, South Korea. doi:10.1115/GT2016-56893.
- [246] ZHOU, Z., ONG, Y.S., NAIR, P.B., KEANE, A.J. AND LUM, K.Y. (2007) Combining global and local surrogate models to accelerate evolutionary optimization. *IEEE Transactions on Systems, Man and Cybernetics Part C: Applications and Reviews*, 37(1):66–76. doi:10.1109/TSMCC.2005.855506.
- [247] ZITZLER, E., THIELE, L. AND BADER, J. (2010) On set-based multiobjective optimization. *IEEE Transactions on Evolutionary Computation*, 14(1):58–79. doi:10.1109/TEVC.2009.2016569.
- [248] ZUHAL, L.R., PALAR, P.S. AND SHIMOYAMA, K. (2019) A comparative study of multi-objective expected improvement for aerodynamic design. *Aerospace Science and Technology*, 91:548–560. doi:10.1016/j.ast.2019.05.044.