


Patching a Fuselage

Application of Straight Fiber Variable Stiffness Laminates in Large Aerospace Structures

T.J.G. Roepman

Technische Universiteit Delft



Patching a Fuselage

Application of Straight Fiber Variable Stiffness Laminates in Large Aerospace Structures

by

T.J.G. Roepman

to obtain the degree of Master of Science
at the Delft University of Technology.

Student number: 4367944
Project duration: September, 2020 – July, 2021
Supervisors: Dr. J.M.J.F. Van Campen, TU Delft, Assistant Professor
Dr. C. Kassapoglou, TU Delft, Professor

This thesis is confidential and cannot be made public until July 2, 2023.

"The ideal engineer is a composite ...

*He is not a scientist,
he is not a mathematician,
he is not a sociologist or a writer;*

*but he may use the knowledge and techniques
of any or all of these disciplines
in solving engineering problems."*

-Nathan W. Dougherty

DELFT UNIVERSITY OF TECHNOLOGY
FACULTY OF AEROSPACE ENGINEERING
DEPARTMENT OF AEROSPACE STRUCTURES AND MATERIALS

Graduation Committee

Dated: July 2, 2021

Chair holder:

Prof. dr. C. Kassapoglou

Committee members:

Dr. ir. J.M.J.F. van Campen

Dr. ir. D.M.J. Peeters

Dr. ir. R. Vos

Preface

This report is the conclusion to nearly a year of challenging, but fun, thesis work. In our quest to a more sustainable future, I believe we should not avoid radical new technologies because they seem hard or incorporate features we do not yet understand properly. Moreover, I think these technologies are most deserving to study because they are challenging, require out of the box thinking, and demand a combination of several engineering fields. This is where straight fiber variable stiffness laminates are; They dare us to search in a complex design space and find opportunities to save weight.

Challenging research does not come cheap, and I am glad to be supported along the way. I would like to thank IJsselmuiden for sharing his lamination parameters results to verify mine. Sporting together with teammates has been a welcome change to what would otherwise be intense days of studying. It is only after we started training again that I noticed how much this actually helped. I have had fun talking with Martijn during my Thesis. Thank you for sharing your work and your insights on my subject, our weekly conversation were a nice change. Furthermore, the discussion and support I received from my predecessor Mohamed were greatly appreciated. He allowed me to put this work in a broad perspective and provide valuable advice on the global interpretation. Also, I am thankful to Roelof Jan for our nice conversations in which we always simplified each other's problem a bit too much. For some reason we always conclude that things can be neglected, a dimension should be a shrink fit, or that the answer was halfway of what we both thought. This pandemic required me to work from home and I am grateful to my roommates for the way we supported each other during these times; To Kenneth for making us all feel a bit younger sometimes; To Olaf for organizing our nice house and bringing us together; To Hielke for his thorough interest in my oddly shaped buckled cylinders; And to Hans for all the rubber ducking we did trying to optimize my fictive plates. I am thankful to my stepparents and parents for their support throughout this endeavour. Although I do not expected you to understand everything, your questions and endless interest have kept my eyes on the ball. Lastly, I am especially grateful to my two supervisors Julien and Christos who have listened to me weekly for almost a year. Your support, critical questioning, and helpful thinking allowed me to up my game. I have enjoyed every experience you shared Christos; during our weekly meetings and during your courses. And Julien, I am glad to have followed your advice from two years ago to take additional courses at other faculties. They were extremely helpful during this work.

Before we take a deep dive let me substantiate the quote at the start of this report. This thesis required a diverse combination of knowledge and technologies. I like to think out of the box, reformulate the problem, and discuss it with others because each of these things might provide new insights. It is the combination of all these things that makes engineering exciting.

Thore Roepman
Delft, June 2021

Abstract

Carbon fiber reinforced polymer are rapidly adopted in the aerospace industry because of their high specific property and their tailoring ability. Straight fiber variable stiffness (SFVS) laminates tailor the fiber orientation locally by introducing smaller patches that only cover certain parts of the structure. This allows the stiffness to be tailored to specific needs of each region. Moreover, load paths within the laminate can be created to redirect stress away from areas prone to failure and increase the performance of the laminate as a whole.

The robustness of SFVS laminates is guaranteed through blending guidelines that describe the cohesion between regions. However, these guidelines are complex to evaluate and a complaint design that maximizes the structural performance cannot be easily created. Additionally, a global interpretation of the patch design is needed to guarantee manufacturability using current technologies.

In this thesis a novel, automated framework for the design of manufacturable laminate is developed. The framework uses four subsequent stages that each use a different stiffness representation and optimizer to gradually introduce the manufacturing constraints.

First the stiffness of individual regions in the laminate is optimized using a gradient descent optimizer. The buckling performance is predicted using a Kriging surrogate model that is fitted using samples generated via finite element analysis. Samples points are chosen in a latin hypercube fashion to cover the entire considered sample space. The sample space is reduced progressively with every iteration to improve the detail of the surrogate around the solution. Then, the optimal stiffness distribution is converted on a per region basis to stacking sequences using a genetic algorithm. Multiple designs are generated by randomly choosing second-best individuals for some regions to increase the number of designs considered in the next stages.

In the third and fourth stages a cellular automaton and network theory are used to introduce cohesion between regions and comply with current manufacturing standards. The cellular automaton uses different phases to remove butted edges within a cell based on its neighborhood. A set of rules is used to create a blended design while preserving the buckling improvements as much as possible. A novel connection search algorithm is used to connect individual regions together and create global patches. Graph theory is used to find cycles in the dependency order of the patches and ensure a manufacturable patch design where all patches can be stacked monotonically. The result is laminate that consists of several patches with different orientation that should be stacked in a specific order.

Plates in uniaxial compression and cylinders in bending are designed using the framework to increase their buckling performance. Load redistribution using SFVS laminates in the plates increases the buckling performance by 42% over a constant stiffness design. An improved region layout using more regions perpendicular to the load shows the largest improvements. The cylinder accomplishes a 30% structural improvement over the baseline constant stiffness design. That is one percent higher than the maximum reported in literature for fiber steering. Unlike what is expected, increasing the number of regions in the plate results in lower buckling performance because the algorithms struggle to converge. This is not the case for cylinders because an increase in the number of unique regions from 10 to 15 did not yield further structural improvement.

The mentioned improvements allow weight savings of 11% for plates and 7% for cylinders, which are representative for skin panels and fuselage sections. When the additional cost of manufacturing and the cost-to-weight ratio for airlines are considered, it is already worthwhile to use these laminates in airplanes. Further manufacturing improvements and fuel price prospects allow for even more weight savings, but do require the algorithms created as part of this thesis to be developed further.

Contents

Preface	vii
Abstract	ix
Acronyms	xv
Nomenclature	xvii
Diagram Legend	xix
1 Introduction	1
I Literature Study	3
2 Design of Variable Stiffness Structures	7
2.1 Variable Stiffness Laminates	7
2.1.1 Fiber Steered	7
2.1.2 Straight Fiber Variable Stiffness	8
2.2 Multi-Step Design Methods	9
2.3 Optimal Local Stiffness Design using Lamination Parameters	9
2.3.1 Feasible Design Region	10
2.3.2 ABD Matrix	10
2.3.3 Conversion to Stacking Sequence.	11
2.4 Variable Stiffness Benchmarks from Literature	11
2.4.1 High Fidelity Theoretical Plate	12
2.4.2 SFVS Demonstrator	12
2.4.3 A Fiber Steered Hull	13
3 Straight Fiber Variable Stiffness Laminates	15
3.1 Laminate Guidelines	15
3.2 Blending Methods	16
3.2.1 Laminate Guide Formulation.	17
3.2.2 Cellular Automaton	17
4 Practical Aerospace Structures	19
4.1 Economics of Aerospace.	19
4.1.1 Cost Estimation Frameworks	20
4.1.2 Costs in Composite Manufacturing	20
4.2 Adoption of Carbon Composites in Industry.	21
5 Research for a Higher Adoption of SFVS Laminates	23
5.1 Objective	23
5.2 Questions	23
II A Blended Design Framework	25
6 Buckling Prediction using Commercial Finite Element Analysis	29
6.1 Buckling Analysis	29
6.1.1 Eigenvalue Problem	30
6.1.2 Preconditioning	30
6.2 Stiffness Modelling	30
6.2.1 Element Setup	30
6.2.2 Variable Stiffness	31

7	Optimal Design of a Variable Stiffness Laminates for Maximum Buckling Load	33
7.1	Multi-Step Method	33
7.2	Sampling Assisted Local Stiffness Design	34
7.2.1	Design of Experiments	36
7.2.2	Surrogate Modeling	38
7.2.3	Constraint Gradient Descent Optimization	38
7.3	Stacking Sequence Retrieval	39
7.3.1	Genetic Algorithm	39
7.3.2	Alternative Designs	40
7.3.3	Lamination Parameter Sensitivity	40
8	Local Blending Using a Cellular Automaton	43
8.1	CA Setup	44
8.1.1	Cell and Neighborhood	44
8.1.2	Evolution Scheme	45
8.2	Foundation of Patches	46
8.3	Elimination of Butted Edges	47
9	Global Patch Interpretation	51
9.1	Manufacturable Laminates	51
9.2	Algorithm Overview	52
9.3	Connection Search	53
9.3.1	Nodes and Connections	53
9.3.2	Grouping of Nodes in a Patch	54
9.3.3	Stacking Order of Patches	56
9.3.4	Depth First Search	57
9.4	Fixing Strategy	60
9.5	Convergence	61
10	Qualitative Cost Comparison	63
10.1	Number of Patches	63
10.2	Patch Shape Efficiency	63
10.3	Patch Complexity	64
11	Framework Validation	69
11.1	Local Stiffness using Lamination Parameters	69
11.1.1	Comparison of Modelling Techniques	69
11.1.2	Optimization	71
11.2	Stacking Sequence Retrieval	72
12	Results	75
12.1	Clamped Plate with Uniform Region Layout	76
12.2	Clamped Plate with Varying Cross Axis Regions	78
12.3	Cylinder in Bending	80
12.4	Framework Performance	82
12.5	Weight Savings	85
12.6	Cost	86
13	Conclusion	87
14	Discussion and Recommendations	89
	Bibliography	91
A	Detection of Blended Edges	95
A.1	Detection of Blended Edges	95

B	Minimum Cutting Size	97
B.1	Minimum Manufacturing Length	97
B.2	Neighbor length.	97
B.3	Local Minimum Cut Length Compliance Procedure	99
B.4	Results	100
C	Additional Design Results	103
C.1	Patch Metrics of the 5xM Plate	103
C.2	Results for the Mx5 Plate	103
C.3	Global Interpretation Tree of Cylinder Results	104
C.4	Cost and Weight of cylinders.	104
D	Determination of the Number of Samples	109
E	Cylinder Optimization Convergence	111

Acronyms



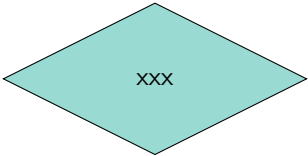


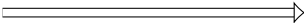
Notation	Description
AFP	Automated Fiber Placement.
ATL	Automated Tape Laying.
CA	Cellular Automaton.
CS	Constant Stiffness.
DOC	Direct Operating Cost.
GA	Genetic Algorithm.
GI	Global Interpretation.
LHS	Latin Hypercube Sampling.
LP	Lamination Parameter.
SFVS	Straight Fiber Variable Stiffness.
VS	Variable Stiffness.

Nomenclature

Notation	Description	Unit
ν_{12}	Poisson's ratio between the first and second direction.	-
θ, θ_i	Fiber orientation of the ply, subscript denotes the ply's level	°
$\Gamma_1, \Gamma_2, \Gamma_3, \Gamma_4, \Gamma_5$	Material property matrices to be used in conjunction with the lamination parameters	MPa/mm
a, b	The width and height of a plate respectively	mm
A, B, D	The extensional, coupling, and flexural stiffness matrices of the laminate respectively	GPa
AR	Aspect ratio	-
C	Patch complexity marked by the total number of corners	-
$E_{11}, E_{22}, E_{12}, E_{13}, E_{23}$	Material stiffness properties	GPa
h	Thickness of the entire laminate	mm
m, n, i, j, k	General counter values	-
M, N, I, J, K	General maximum counter values	-
$(M)NAR$	(Mean) Normalized Aspect Ratio	-
t, t_{ply}	Thickness of a single ply	mm
$Q_{xx}, Q_{yy}, Q_{ss}, Q_{xy}$	Lamina stiffness	GPa
U_1, U_2, U_3, U_4, U_5	Material invariants	GPa
V_1, V_2, V_3, V_4	Extensional lamination parameters (always normalized by the thickness)	-
W_1, W_2, W_3, W_4	Flexural lamination parameters (always normalized by the thickness)	-
X	Design Variable	-
Y	Design Objective	-
z	Through the thickness coordinate within the laminate	mm
$\square^*, \bar{\square}, \tilde{\square}$	Denotes the normalization of the variable to some entity	-

Diagram Legend

This report contains a higher number of flow diagram. They have been standardized using the below presented legend.

	Input or output of flow diagram
	Describes a process or the manipulation of variables
	Junction in a flow diagram based on the statement in side the deci- sion node
	Indicates that this process is iterative or causes multiple branches in the algorithm
	Indicates the flow of processes
	Indicates the parallel/branching flow of processes or data; For ex- ample multiple designs are considered

Introduction

Carbon fiber composite structures are gaining traction in the aerospace industry. New airplane design incorporate carbon composite structure in large parts of the wing and fuselage. The benefits are evident; Increased structural performance leads to a reduction in fuel and operating costs for airliners. However, around the first large-scale application of carbon composite structures in a commercial airplane by Boeing [21], Hyer and Lee introduced the idea of tailoring fiber orientation throughout a laminate to increase its global performance. [17] The ability to locally tailor the stiffness is an additional design freedom to engineers, but also makes the design inherently more complex[36]; Efficient search algorithms are needed to find an optimum when a higher number of design variables are used. [34] When designed correctly, variable stiffness (VS) design can further decrease the mass of structures and thereby decrease the fuel consumption of modern airplanes. [32]

In this thesis, a alternative way to achieve variable stiffness is explored via laminates that introduce new plies where they are needed and drops them when they become obsolete. A change in stacking sequences brings about a change in stiffness that can be used to create load paths in a structure. Such laminates are called straight fiber variable stiffness(SFVS) laminates because they use straight fibers to create variable stiffness. [4]. Their counter parts use curved fiber paths which are generally harder to manufacture. Still the adoption of SFVS laminates is limited due to the complexity associated with the design. [50]. Even though an experimental campaign has yielded buckling improvements of over 80%. [12]

Further research into design methods for these type of laminates is performed in this work. New methods that combine several optimization techniques and mathematical representations have the potential to improve these design techniques. The objective is to create an integrated end-to-end design methodology that takes care of the local stiffness optimization and creates a blended design cohort the broadest blending guidelines. Such a design has the potential to increase structural performance comparable to the results found in the experimental campaign.

Application of such a framework to practical large aerospace structure is not trivial and requires computational efficient optimization techniques that make use of finite element analysis. Also, the framework should be automated to create consistent and scalable results.

This work is split into three main parts. First a literature review is conducted in Part I which contains chapter 2 to 4. Several subjects including design methods for variable stiffness and aerospace structures are discussed. In part two the framework developed as part of this thesis is introduced. Each chapter covers a different aspect of the framework which itself is detailed at start of Part II and in chapter 7. Then, Equation 10.3 discusses the validation and design results of the framework's application literature benchmarks. The report is concluded by the conclusion and discussion in chapter 13 and chapter 14.



Literature Study

The framework developed here consists of several individual parts that are based on other research. Extensive research on the state-of-the-art of optimization techniques was conducted throughout the thesis. This part of the report aims to discuss the individual research that shape the underlying principles for this work or has served as inspiration.

The literature study is divided in three parts that each introduce a different part of the background knowledge that was required to design the framework. The first chapter, chapter 2, introduces the current state of art in variable stiffness design by discussing multi-step processes and other variables stiffness techniques such as fiber steering. Then, in chapter 3, straight fiber variable stiffness (SFVS) laminates are discussed including, but not limited to, the different blending guidelines and the current best blending techniques. Lastly, in chapter 4, the current adoption of variable stiffness laminates in the aerospace industry is discussed.

Design of Variable Stiffness Structures

Similar to regular carbon composite structures, the design of variable stiffness laminates is a trinity between manufacturing, materials, and design methods. The limitation of each step should be understood by the other to design an optimal laminate. [36]

Currently, researchers are pursuing two ways to produce variable stiffness structures that are discussed in the next section. Then in the second section an introduction to the multi-step methods used to design these laminates is given. section 2.3 discusses the first step in the processes that is generally used. To compare the outcome of the algorithm, benchmark problems are introduced in the last section of the chapter.

2.1. Variable Stiffness Laminates

As mentioned in the introduction, Hyer and Lee were one of the first to think of the concept of changing the orientation of fibers inside a laminate. [27] Changing the orientation of fibers in the laminates changes stiffness locally and redirects stress to parts of the structure that can better handle these loads. As IJsselmuiden et al. points out in their work, [20] load redistribution is an important to the increase in buckling load. Redistribution of the load to the sides of the plate increases the bifurcation load with more than 180% in some cases.

The fiber orientation cannot abruptly change in a laminate. Such a laminate would have no practical rigidity and can merely exist in computer analysis. Continuity within the laminate is needed to transfer loads and provide rigidity. Currently, two techniques to obtain variable stiffness within a laminate are investigated: Fiber Steering and Straight Fiber Variable Stiffness. These techniques use a different mechanism to obtain their variety and discussed in the next two sections respectively.

2.1.1. Fiber Steered

Laminates are called fiber steered designs when they consist of continuous fibers that change orientation continuously and follow a curved path. By changing the orientation of the fiber, stiffness properties can differ within the laminate while preserving load carrying capabilities through the continuous fiber. Fiber paths are constructed based on the local orientation at control points along the path. The desired orientation determines the path's gradient and specifies how it should continue. [27, 36]

Design of these paths is done with control points along the structure and linearly interpolating the orientation to obtain a model suitable for structural evaluation. The exact fiber paths are constructed by integrating the curve between control points. An example of a parameterization using control points often used in cylinders is shown in Figure 2.1. [6, 40]

Fiber steered designs require complex machines to precisely align fibers along the path. Current automated tape laying (ATL) and automated fiber placement (AFP) technology produce acceptable results but still introduce manufacturing artifacts such as overlaps and gaps. Because the tows and fiber strands used in these processes have a minimum width the curvature of the path is limited. Limits on path curvature should be considered throughout the design and set an upper limit on the difference in




Figure 2.1: Points along a cylinder to control the curvature of the fiber path as used by Blom. Picture original published by Blom. [5]

the stiffness. Paths that require a too tight curve will have wrinkles on the inside and cause fiber build up. [36]

Advances such as continuous tow shearing and variable extrusion speed allow for a better control of the fiber and can decrease the curvature to 30mm. [25] More precise control and simulation of the resulting laminate have also led to reductions in overlaps and other manufacturing increase the yield of this type of panels. [44]

2.1.2. Straight Fiber Variable Stiffness

Another way to obtain variable stiffness is by dropping plies when they are no longer needed or might negatively influence the structural performance. Such laminates are called straight fiber variable stiffness laminates because, in comparison to fiber steered designs, they are made using straight fibers only. By starting and dropping layers in certain regions of the laminate stiffness changes can be incorporated in the design without loss of the cohesion between these regions. [47] Guidelines enforce cohesion and limit abrupt stiffness changes ensuring a robust laminate.

Production of SFVS laminates does not require complex machines and can be done by hand. Manufacturing involves stacking patches of carbon fiber in a specific pattern and order. This can be done using a pick and place machine, by hand, or via an automated process. The simplicity makes the manufacturing both simple and low cost.

However, as the complexity of the patches grows it will be harder to construct the laminate precisely. Moreover, current design methods of SFVS laminates can require two patches to become woven, meaning that technicians have to lift an already stacked ply to weave another under it. There does not exist a machine that can currently do this to the best knowledge of the author. [27]

Straight fiber variable stiffness laminates provide a large design space and these laminates are generally hard to design. Engineers have to guarantee a certain degree of continuity between regions for the laminate to be stable. Although guidelines exist, it is hard to find a solution to a large problem that satisfies these guidelines and optimizes the performance. A stable design framework for these types of laminates is not yet present. [12] Details of the technique are further discussed in chapter 3.

2.2. Multi-Step Design Methods

Designing variable stiffness structure requires finding a local stiffness such that the loads are distributed globally while satisfying several strength, stability and manufacturing constraints. [20] Optimization of the structural performance involves solving a large optimization problem with a high number of design variables. [37]

Due to the complexity of SFVS and fiber steering constraints, it is often beneficiary to explore the optimum first without constraints and then reduce it to a manufacturable laminate. This multi-step approach first tries to find the optimum stiffness distribution and then tries to find a laminate that best fits this stiffness distribution. This breaks the problem into sub problems that can utilize their own optimization strategies. [34, 36]

A structure is generally parameterized by dividing it into sections with constant stiffness or using a linear interpolation of the stiffness between control points. The fiber orientation for each of these design points should be determined to optimize structural performance. Manufacturing consideration such as maximum fiber curvature limit the design space some, but the remainder is too large for simple enumeration. [10] Depending on the exact formulation the problem might be separable into a subproblem per region such that a smaller local optimization leads to a global optimum. [20]

Efficient parameterization of the design aims to reduce the number of variables to a smaller set that represents the entire design space. Additionally, parameterization can cause constraints to be automatically satisfied or to make the design space convex, a property that most optimizers need. Often laminate designs employ this trick on default: By only considering balanced symmetric laminates a large part of the ABD matrix can be neglected. Other parameterization such lamination parameters and polar parameters can be used to factor out trigonometric functions from the stiffness computation. This has the benefit that stiffness can be interpolated linearly between regions. [13]

A method that effectively searches the design space to minimize, or maximize, an objective is called an optimizer. Optimization problems are either unconstrained or constrained where the latter also deals with constraints imposed on the problem. The mathematical formulation is:

$$\begin{aligned} & \min_{\mathbf{x}} f(\mathbf{x}) \\ & \text{s.t.} \\ & \quad g_i \mathbf{x} = 0 \\ & \quad h_j \mathbf{x} \geq 0 \end{aligned} \tag{2.1}$$

where \mathbf{x} are the design variables, $f(\cdot)$ the objective to minimize, $g_i(\cdot)$ the equality constraints, and $h_j(\cdot)$ the inequality constraints. A naive optimization method that quickly increases in computational effort is a enumeration of all design combinations. Gradient descent optimizers use additional gradient information to reach an optimum. Although this type is fast it is vulnerable for local optima. [34] A third type use stochastic functions to overcome these local optima. One particular stochastic type is the genetic algorithm. [48]

The exact number of steps a multi-step method uses differs per implementation but often at least two steps can be identified: First the local stiffness is optimized to produce a design that has the highest structural performance, but is infeasible to produce; then a second step reduces this optimum to a manufacturable design that satisfies the manufacturing constraints. Converting between different representations of the structures, such as lamination parameters and stacking sequences, can introduces additional steps. [34]

The first step, local optimization, is often formulated independent of the chosen variable stiffness type. Lamination parameters are commonly used to reduce the number of variables in this step as is discussed in the next section. The second step depends on the choice between SFVS laminates or fiber steered laminates. A method developed by Van den Oord and Van Campen et al. is discussed in the next chapter, chapter 3.

2.3. Optimal Local Stiffness Design using Lamination Parameters

Commonly, the optimum stiffness distribution of structure is found using a parameterization that abstracts the fiber angles. Such a parameterization can reduce the number of variables significantly

because independent plies are no longer considered. Also variable thickness can be implemented continuously through a thickness variable. Subsequent steps are needed to construct the exact plies.

A common parameterization uses the lamination parameters (LPs). They are straightforwardly defined as

$$\{V_1, V_2, V_3, V_4\} = \int_{-\frac{1}{2}}^{\frac{1}{2}} \{\cos 2\theta(\bar{z}), \sin 2\theta(\bar{z}), \cos 4\theta(\bar{z}), \sin 4\theta(\bar{z})\} d\bar{z} \quad (2.2)$$

$$\{W_1, W_2, W_3, W_4\} = 12 \int_{-\frac{1}{2}}^{\frac{1}{2}} \bar{z}^2 \{\cos 2\theta(\bar{z}), \sin 2\theta(\bar{z}), \cos 4\theta(\bar{z}), \sin 4\theta(\bar{z})\} d\bar{z} \quad (2.3)$$

where \bar{z} is the normalized z location of the ply in the laminate and $\theta(\bar{z})$ is the local ply angle. [13] There exist four more parameters that couple the in-plane and out-of-plane behavior but they are neglected here. For conventional laminates they are equal to zero. Also when the laminates are limited to be balanced and symmetric, $V_1 = V_3 = 0$ and $W_1 \approx W_3 \approx 0$ can be assumed.

Lamination parameters have been used by research in the design of variable stiffness structures. Constraints that are physically imposed on the design can be translated to the lamination parameters space. Such constraints include the failure criteria and robustness rules. [3, 18] The feasible design space and relation to laminate stiffness are discussed in the next two sections.

2.3.1. Feasible Design Region

As implied by Equation 2.2 and 2.3, there exists a strong correlation within lamination parameters. As a result the design space is limited. A complete formulation that includes all LPs has not been found. For laminates where the extensional and flexural properties are decoupled, most symmetric laminates, a closed-form limit exists:

$$\begin{aligned} V_3 &\geq 2V_1^2 - 1 \quad \wedge \quad V_3 \leq 1 \\ W_3 &\geq 2W_1^2 - 1 \quad \wedge \quad W_3 \leq 1 \end{aligned}$$

When laminates are also required to be balanced the upper limit reduces further to the one presented in Figure 2.2.




Figure 2.2: Feasible domain for in-plane lamination parameters of balanced symmetric laminates with k unconstrained ply angles. Images is adopted from Van Campen and Gürdal. [46]

IJsselmuiden et al. used a convex hull approximation to linearize the feasible domain. The resulting set of inequalities is used in the optimization process. The convex hull is generated by sampling all the angle combinations and it does not necessarily span the full design space. For example, it flattens the parabolic insets at the upper limit and quantizes the lower curvature. [20]

2.3.2. ABD Matrix

The advantage of using lamination parameters in the optimization process is linear relation of the variable to the stiffness. The extensional and flexural matrices, respectively A and D are defined as:

$$A = h(\Gamma_0 + V_1\Gamma_1 + V_2\Gamma_2 + V_3\Gamma_3 + V_4\Gamma_4) \quad (2.4)$$

and

$$D = \frac{h^3}{12} (\Gamma_0 + W_1 \Gamma_1 + W_2 \Gamma_2 + W_3 \Gamma_3 + W_4 \Gamma_4) \quad (2.5)$$

where h is the laminate thickness and Γ_i are the material properties. They are computed using the material variants

$$\begin{aligned} U_1 &= \frac{3Q_{xx} + 3Q_{yy} + 2Q_{xy} + 4Q_{ss}}{8} \\ U_2 &= \frac{Q_{xx} - Q_{yy}}{2} \\ U_3 &= \frac{Q_{xx} + Q_{yy} - 2Q_{xy} - 4Q_{ss}}{8} \\ U_4 &= \frac{Q_{xx} + Q_{yy} + 6Q_{xy} - 4Q_{ss}}{8} \\ U_5 &= \frac{Q_{xx} + Q_{yy} - 2Q_{xy} + 4Q_{ss}}{8} \end{aligned} \quad (2.6)$$

to produce the Γ_i matrices[13, 38]

$$\begin{aligned} \Gamma_0 &= \begin{bmatrix} U_1 & U_4 & 0 \\ U_4 & U_1 & 0 \\ 0 & 0 & U_5 \end{bmatrix}, \quad \Gamma_1 = \begin{bmatrix} U_2 & 0 & 0 \\ 0 & -U_2 & 0 \\ 0 & 0 & 0 \end{bmatrix}, \quad \Gamma_2 = \frac{1}{2} \begin{bmatrix} 0 & 0 & U_2 \\ 0 & 0 & U_2 \\ U_2 & U_2 & 0 \end{bmatrix}, \\ \Gamma_3 &= \begin{bmatrix} U_3 & -U_3 & 0 \\ -U_3 & U_3 & 0 \\ 0 & 0 & -U_3 \end{bmatrix}, \quad \Gamma_4 = \begin{bmatrix} 0 & 0 & U_3 \\ 0 & 0 & -U_3 \\ U_3 & -U_3 & 0 \end{bmatrix} \end{aligned} \quad (2.7)$$

In optimization problems the Γ_i matrices only have to be computed once. During optimization, the stiffness is computed based on the lamination parameters only. Additionally, the gradient of the stiffness with respect to the design variable is linear and thus does not require gradient computation at every iteration of the optimizer. Gradient descent optimizer can take advantage of this ability to quickly converge to the optimum. [20]

2.3.3. Conversion to Stacking Sequence

Although lamination parameters provide a full representation of the laminate stiffness, they do not contain information on the stacking sequence. There exists no direct conversion from lamination parameters to a discrete set of orientations. It has been proven that for a laminate with variable thickness plies a solution can be found, but in reality continuous variable thickness plies do not exist. [14]

During the early uses of lamination parameter a backward conversion based on layerwise optimization was used. This branch-and-bound optimization methodology started at the outer layer and picked the orientation for every layer that best approaches the desired lamination parameters. However, this method struggles to find a good solution in some cases because the design space is reduced after every layer. [29]

Als genetic algorithms have been used to find a set of orientations that best optimizes a structural objective. Genetic algorithms are a good fit here because they are suitable for integer problems and do not get stuck in a local optimum unlike the layerwise optimization. [4, 47] Their application in the conversion of lamination parameters to stacking sequences is not well understood. However, it is expected that they are applicable for such problems.

2.4. Variable Stiffness Benchmarks from Literature

Variable stiffness designs that use carbon fiber orientation tailoring have been researched extensively. Most researchers have limited themselves to simple shapes such as plates[12, 20, 39], cylinders[5, 24, 35], or cutouts[7, 51]. The highest experimentally recorded improvement made by SFVS laminates is over 80% in buckling load and close to 100% in stiffness by the SFVS demonstrator[12].

Besides showing the potential of VS laminates, these designs can be used as benchmarks in the development of a SFVS framework. It is expected that the framework developed here can generate

designs with similar performance to the laminates presented in literature. Some of the benchmarks might only be applicable to parts of the framework. For example, some works do not generate a producible laminate and derive just the optimal lamination parameters. In these cases those parameters can be compared to the results obtained in this thesis.

Three cases have been selected to serve as benchmark in this framework. These have been chosen such that a comparison to their respective constant stiffness and optimum design can be made. They are presented in the following three sections.

2.4.1. High Fidelity Theoretical Plate

The first step in the method developed as part of this thesis optimizes the lamination parameters to maximize the structural performance. Because this is the starting point the output should be of a high quality. Generally, it is hard to converge a large number of variables to their optimum using a stochastic optimization process such as the one used here. By comparing the output of this stage to a benchmark, the first step in the framework can be validated.

IJsselmuiden et al.[20] developed a gradient descent optimizer to maximize the buckling load of a plate by means of its lamination parameters. Several cases were optimized, but the uniaxial compression with straight simply supported edges was chosen as a benchmark. The optimal variable stiffness design uses 400 elements and increases the buckling load more than twofold over the constant stiffness (CS) design.

The use of this design as benchmarks allows validation of the framework on a large scale problem that was solved using a different optimization technique. The details of the benchmark are presented in Table 2.1 and a layout of the plate is shown in Figure 2.3.

Table 2.1: Problem properties of the plate optimized using lamination parameters by IJsselmuiden et al.. [20] Note that the loads have been made dimensionless according to $N_{cr} = N_{cr} a^2 / E_1 h^3$

Property	Value	Material Property	Value
width	500 mm	E_{11}	181 GPa
height	500 mm	E_{22}	10.3 GPa
Constant Stiffness Load	1.73 [-]	E_{12}	7.17 GPa
Optimum Theoretical Design Load	3.99 [-]	ν_{12}	0.28 [-]
Plate thickness, h	1.524 mm	$E_{13} \approx E_{23}$	7.17 GPa




Figure 2.3: Simple supported square plate in uniaxial compression. Buckling loads are normalized with respect to the plate dimensions, hence no dimensions are given.




Figure 2.4: Clamped non square plate in uniaxial compression.

2.4.2. SFVS Demonstrator

As part of his MSc. Thesis Goma[12] designed, manufactured, and tested the first SFVS panel. The panel has a theoretical buckling load improvement of 81% and increased stiffness of 94%, while using only 5 by 5 unique regions. The algorithm that was used to blend this panel was originally developed by Van den Oord[50] and will be explained in more detail in the upcoming chapter.

This panel is one of the few designs reported in literature that uses SFVS to increase its buckling performance. The laminates designed as part of this thesis use equivalent techniques, but are produced using an automated process. It should be noted here that the algorithm did not adhere to all the blending guidelines and some blending defects were present.

The design is loaded in uniaxial compression, is clamped at the ends, and is supported using a knife edge on the sides. After blending this design had a theoretical buckling load of 132.3 kN, which is comparable the experimental measurements. The design is schematically drawn in Figure 2.4 and the benchmark properties are given in Table 2.2. [12]

Table 2.2: Problem properties of the demonstrator panel designed by van Van Campen and tested by Gomma. [12, 52]

Property	Value	Material Property	Value
width	600 mm	E_{11}	155 GPa
height	400 mm	E_{22}	7.8 GPa
Constant Stiffness Load	74.0 kN	E_{12}	5.5 GPa
Optimum Theoretical Design Load	132.3 [-]	ν_{12}	0.27 [-]
Ply thickness	0.156 mm	E_{13}	4.7 GPa
# plies	32 [-]	E_{23}	2.9 GPa

2.4.3. A Fiber Steered Hull

Airplane manufactures are now also using carbon fiber in large components such as the fuselage. [15] These large structures have been the subject of research efforts to reduce weight. Blom et al. tailored the circumferential stiffness of a cylinder to redistributed stress induced by bending. A significant improvement of more than 17% and 27% was achieved when material failure was considered and neglected respectively[6].

The cylinder used a fiber steered laminate to obtain variable stiffness. To the best knowledge of the author, no attempt has been made in literature to utilize SFVS laminates to optimize the buckling load of cylinders in bending. It is therefore that the designs created by Blom, are the current optima and can serve as benchmark.

Table 2.3: Problem properties of cylinder reference case as designed Blom. [5]

Property	Value	Material Property	Value
Radius	305 mm	E_{11}	129 GPa
Length	813 mm	E_{22}	9.1 GPa
Constant Stiffness Load	627 kNm	E_{12}	5.3 GPa
Fiber Steered Load	809 kNm	ν_{12}	0.32 [-]
# layers	24	$E_{13} \approx E_{23}$	5.3 GPa
ply thickness	0.183 mm		

Often, literature uses another benchmark that is not used in this thesis: The horseshoe panel design[43, 50]. The horseshoe benchmark defines loads for every region separately. The goal of the benchmark is to design a structure that minimizes weight but fulfills buckling constraints per region. However, load redistribution between regions is not taken into account nor is global buckling of the structure. The lack of global optimization causes the benchmark to be of less interest and hence this load case is not used in this thesis.

3

Straight Fiber Variable Stiffness Laminates

An elegant and simple way to introduce variable stiffness in a laminate is by dropping old and introducing new plies where other properties are desired. The resulting differences in stacking sequence consequently cause a region to have different stiffness and strength. Moreover, this technique does not require automated processes or complex fiber paths thus making the technique accessible.

Similar to fiber steering, the amount of variance in stiffness is limited by practical constraints; in the case of SFVS laminates this is the rigidity and robustness of a laminate. It is not practical to change the entire layup from one region to another because such a laminate would only be held together by its matrix. Since such a laminate would certainly fail when loaded a certain amount of rigidity is needed. [47] This rigidity is guaranteed through the use of blending rules as further discussed in the next section, section 3.1.

Creating a design that satisfies these guidelines is not a straightforward process due to the high number and formulation of the blending constraints. Research has proposed several blending techniques which often scale exponentially rendering impractical for high detail large structures. [43] Promising methods are further explored in section 3.2.

3.1. Laminate Guidelines

Structural rigidity of the laminate is preserved by complying with guidelines specifically developed for this purpose. Each guideline introduces a set of rules describing how plies are allowed to change from one region to another. Thereby they guarantee a certain amount of continuity between regions and ensure a robust laminate. A drawback of the enforced continuity is the limited variability that can be designed into the laminate. To overcome this problem more complex guidelines have been introduced that further increase the allowed variability at the downside of being more complex to design for. [4, 47]

Initially drafted by Adams et al., the "outer-and-inner"-rule allows plies to be dropped at the symmetry line or outside of the laminate. An example of a valid laminate is shown in Figure 3.1. This technique is also often used to create build ups around features such as cutouts. Although the guideline is simple to design for, the lack of the ability to drop plies in the center limits the amount of variability the engineer can introduce. [4]




Figure 3.1: Visualization of allowed ply drops according to the "outer-and-inner"-rule. Modified from Van Campen et al.. [47]

The next logical step is to allow dropping of plies at a levels as described by the general blending rule,

of which an example is given in Figure 3.2:

We consider two adjacent panels completely blended if all the layers from the thinner panel continue in the thicker one regardless of their position along the thickness of the laminate.

- Van Campen et al.[47]




Figure 3.2: Generalized blending requires all plies from the thinner region to continue into the thicker region. Modified from Van Campen et al.. [47]

This rule allows plies at any position of the thicker region to be dropped. This guideline can be evaluated using switch genomes as further explained in section 3.2. The guideline does require a change in the thickness if the stiffness should change. Variable thickness complicates the initial local stiffness optimization. [47]

Both of the previous guidelines require the structure to incorporate variable thickness to introduce variable stiffness. The relaxed generalized blending guideline overcomes this and allows for the most stiffness variability at the expense of a complex evaluation. Van Campen et al. introduced the guideline as:

We consider two adjacent panels completely blended if there are no dropped edges in physical contact.

- Van Campen et al.[47]




Figure 3.3: Relaxed generalized blending prohibits the presence of butted edges in the laminate. Modified from Van Campen et al.. [47]

An example of the correct application is given in Figure 3.3. This guideline allows the highest variability in the laminate and does not require a variable thickness of the laminate. Therefore, the thickness can either be neglected in the local stiffness optimization or two adjacent regions with same thickness can still have a different stiffness. Ultimately this guideline provides the broadest design space for engineers but has not been researched extensively due to the complex evaluation of the rule. [47]

3.2. Blending Methods

To convert a local stiffness design into a global blended design several methods have been introduced in literature. These methods try to match a global laminate with a local stiffness distribution while complying with one of the guidelines. The performance of this stage can be measured by the loss in structural performance from the local optimum which is often coupled to the size of the problem. Highly detailed designs have a high number of variables and an exponentially growing number of solutions, thus requiring a high performance algorithm.

Recent research using integer linear programs has been centered around the horseshoe model and compliance with the general guidelines. No attempts have been made with regards to plates and cylinders or compliance with the relaxed blending guidelines yet. [31, 36]

In the next two sections two blending methods that incorporate relaxed blending are more thoroughly examined. First an approach based on a genetic algorithm is discussed and then a cellular automaton that aims to overcome the exponential growth of the problem.

3.2.1. Laminate Guide Formulation

Besides using a genetic algorithm to convert lamination parameters (LPs) to stacking sequences, it can be extended to blend a panel. Van Campen et al. used three chromosomes to blend a single interface using a genetic algorithm (GA). Although the algorithm was able to design an interface that complies with the relaxed blending guidelines, the algorithm requires a repair strategy making it inapplicable on its own. [47]

The first of the chromosome in the individual encodes the orientation of the guide. An integer number indicates the orientation of that ply in the laminate. Evolution through the GA changes the orientation of some parts of the guide through operators like crossover or permutation. The individuals that yield the best match to the required lamination parameters (LPs) are likeliest to proceed and modified using the operators to form the next population. Every iteration elite individuals proceed directly to the new population to avoid loss of the best solution.

Accompanying the orientation chromosome are two switch guides that mark the presence of a ply in a region. These switch guides take the value 0 or 1 where 1 marks the presence of the ply and 0 the absence. An example of a complete individual is given in Figure 3.4. Compliance with the relaxed blending rule is checked by identifying the presence of mismatched plies in both switch chromosomes between two continuous plies. Such an occurrence results in two dropped edges that are in direct contact and would violate the blending rule. A penalty function is used to prevent these individuals from evolving into future generations.




Figure 3.4: The guide chromosome, top, encodes the orientation of the plies. Switch chromosomes define the presence of a ply from the guide chromosome in each region. Evaluation of the butted edges is based on evaluation of the discontinuous plies in the switch chromosome. Diagram originally published by Van Campen et al.. [47]

The crossover operator has been identified as the problem because it has a high likelihood of introducing butted edges. As a consequence the method has convergence problems for two regions. This shortcoming would only be exaggerated for problems with more regions. A repair strategy is needed to make the method suitable for large structures. [47]

3.2.2. Cellular Automaton

Van den Oord[50] designed a multi-phase cellular automaton (CA) initially meant to repair the defects left after the GA previously discussed. However, the cellular automaton (CA) can also blend a set of stacking sequence that were derived using other algorithms, including those that did not have blending constraints. The method does also work for large structures because it scales linearly in time instead of exponentially.

Linear scaling is an intrinsic property of cellular automata and is caused by the local evaluation. In every iteration each cell is changed based on a set of rules that only take into account the direct environment. As the problems grows the number of cell evaluation grows linear, however the evaluation time per cell does not change. It is up to the designer of the algorithm to specify suitable rules and set neighborhood limits such that the algorithm converges to a global optimum. [50]

The implementation of Van den Oord used a single cell for every region in the laminate. For every iteration the cell evolved based on the direct neighborhood consisting of the adjoint cells. During each of the three stages a different set of rules was used to gradually introduce the different manufacturing constraints. The separate stages are visualized in Figure 3.5 and are further discussed later in this section. To prevent the CA from continuously switching between two solutions a checkerboard evolution

order was used. This meant that two adjoint cells cannot evolve during the same iteration. When cells that evolve at the same time are marked with the same color a checkerboard emerges. The used evolution order for the horseshoe panel is presented in Figure 3.6. [50]

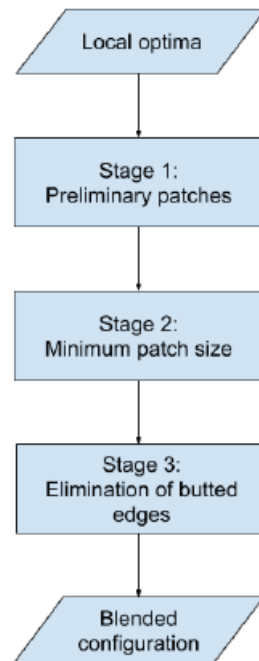


Figure 3.5: Flow diagram of the CA blending algorithm developed by van den Oord and van Campen. Diagram originally published by Van den Oord. [50]




Figure 3.6: Parameterization and representation of the updating scheme for design with complex patches. Diagram adopted from Van den Oord. [50]

During the first stages of the CA a set of rules is used to form preliminary patches. Based on neighboring cells, a cell changes its orientation to join groups of similar oriented neighbors. Changes are prioritized such that the total amount of change is minimized and such that plies close to the symmetry plane are changed first. This postpones large changes at the outside of the laminate which can have significant impact on the buckling load.

The second step is to make the laminate compliant with minimum cutting constraints imposed by modern ATL and AFP machines. Preliminary patches are combined such that the patch satisfies a minimum size equal to the minimum cut length in every direction independent of the patch' orientation. To this end, the algorithm does not use a CA approach and iterates first over X and then Y to check every region individually.

Lastly, butted edges are removed by changing one layer in every cell to the orientation that minimizes the number of butted edges locally. Enumeration of all the options leads to a change of one layer in the cell for which the total amount of butted edges is minimum. This forms the output of the algorithm. [50]

The resulting CA has a linear solution time and is therefore suitable for use in large structures. In direct comparison on the horseshoe benchmark a solution with equivalent weight as in literature[50] was found which validates the effectiveness of the algorithm. The algorithm does provide global interpreted solutions for 2D laminates making it suitable for some applications.

However, the algorithm does not guarantee a global patch solution for problems with a 2D grid of regions and manual interpretation of the result is needed. This prevents large scale application of the algorithm since this becomes infeasible for complex structures. Also in some cases the algorithm is not able to remove all butted edges and comply with minimum size constraints due to the locality of the butted edges removal process.

Practical Aerospace Structures

Research deals with simplified structural models that should provide an accurate representation of practical aerospace structures. Simply supported plates loaded in uniaxial compression represent skin panels of wings or fuselages[12, 19], perfect cylinders or elliptical shells represent rocket stages or fuselage section[6, 24, 45], and circular cutouts represent doors or access holes. [7] Simplification leads to perfect shapes that can be modelled using analytical expression and does not require time-consuming finite element modelling.

Each of these researches aims to increase structural performance, usually at the cost of a more complex design, to decrease the weight of the structure. The weight of the airframe ties directly into the cost associated with flying as is further discussed in section 4.1. Besides weight reduction, research provides a better insight into the behavior of composites which is necessary for adoption in the industry, as discussed in section 4.2.

4.1. Economics of Aerospace

It should be clear that weight savings have other benefits such as reduction of fuel consumption and pollution, but the primary rationale for airlines is the decreased cost of flying. The direct operating cost(DOC) of an airline encompasses all the cost associated with operating an airplane. A split of the cost is presented in Figure 4.1.




Figure 4.1: Split of Direct Operating Cost as approximated by Boeing during the ATCAS program. Reference data is taken from Ilcewicz et al.[21].




Figure 4.2: Approximate division of Ownership Cost as defined during the ATCAS program[21].

Fuel makes up a quarter of the operating cost of airplanes and a slight reduction in this cost is significant at the scale of commercial air travel. The amount of fuel that is required strongly depends on the weight and payload of an airplane. It is therefore that a small reduction in weight caused by a more efficient structure leads to a reduction of the cost for airliners.

Nevertheless, these weight savings come at a price. Complicated designs require complicated manufacturing technologies that increase the cost associated with manufacturing. The impact of the cost increase translates to the airlines via the ownership cost that is further divided in Figure 4.2. Here it can be seen that approximately 51% of the cost is made up by the airframe to which the structural components belong. The exact decomposition of that figure strongly depends on the aircraft, but in conclusion the airframe and fuel costs contribute approximately the same amount to the total operating cost. [21]

In reality a trade off is thus needed that takes into account the increase in manufacturing cost and the decrease in fuel costs for airliners. Balancing these interest is not straightforward because these costs are generally hard to estimate. In the next two section several approaches and drivers of composite manufacturing costs are further evaluated. [16]

4.1.1. Cost Estimation Frameworks

There exists several cost estimation frameworks that can be used to asses the cost of a component during the different stage of the design. Not every model is suitable for application at every design stage because some need more information than others. In general it can be said that the cost estimation becomes more accurate as the design progresses. Nonetheless a reasonable estimate is needed in the early design stages to determine its feasibility.

During the later stages of the design when detailed information is present a bottom-up approach can be used. Such an estimates takes into account all the subcomponents of the cost and sums it to the total. This method does only provide an accurate approximation if all the cost are included thus requiring more experience. Such an approach points out the major contributors to the price and can therefore be used to lower the total cost.

A parametric model could also be used. This model uses derivatives of the price with respect to driving design parameters to estimate the cost of a component. Such an approach allows the cost estimation to be added to a knowledge based engineering approach. As the design changes, the manufacturing costs is also updated. However, such an approach only works for manufacturing techniques that are well understood and for which the driving parameters have been identified. This is not necessarily the case for new technologies such as straight fiber variable stiffness laminates.

In early stages of the design there is often a lag of information to setup one of the previous models. A reasonable estimate can still be obtained by comparing the design to analogous parts and scaling the cost where necessary. Such an approach does only provide a gross estimate, but can be used to determine feasibility of the design in its early stages. Similar to the parametric model, a technology should be matured for this method to be applicable. [16]

4.1.2. Costs in Composite Manufacturing

Carbon fiber manufacturing is a complex process that requires experienced technicians and special tooling. Exact production techniques should be determined early in the process and these choices do directly affect the component's cost. For new technologies a bottom-up or parametric approach is most suitable because analogous designs do not exist. Some of the most important cost drivers are mentioned here. [32]

Arrangement of plies in molds accounts for approximately 40% of the manufacturing and most automation has been focussed on lowering this figure. Convex mould shapes and complex layup might require manual layup for correct manufacturing of the product. Those products do have a low initial production rate because of the human's learning rate. After several products the production ramps up because operations become second nature to the technicians. [22]

Automation has flattened this learning curve and allows for higher production rates. These machines are usually purchased for a specific parts or projects and high production volumes are needed to finance the initial purchase. [28, 42]

New material developments have also been focussed on increasing the performance of cheaper materials. The aerospace industries uses a high amount of pre-impregnated fibers because of their excellence performance. But these prepregs are expensive and should be kept at low temperatures to prevent premature curing. As a result a dense cutting pattern is desired to reduce the amount of waste. [22]

After manufacturing parts are inspected for defects and quality. Parts that do not meet the quality standards are scrapped and do not make it to the final product. The risk in engineering is the likelihood of occurrence of a defect multiplied by the impact. Large carbon composite parts that are expensive to produce and are often scrapped because of defects have a high risk. This risk is potentially high for new technologies and increases with the part's complexity. [22]

4.2. Adoption of Carbon Composites in Industry

The weight savings that composites can bring outweigh the higher manufacturing cost. However, aerospace industry is not solely driven by fuel costs and recurring manufacturing costs as discussed in the previous section. For a new technology to be accepted a thorough understanding into its behavior is needed. As a consequence adoption of carbon composites has been slow. It is only recently, with the introduction of the 787, that large structural components of the airframe are commercially made of carbon composites. This chapter sets out some milestones in the adoption of carbon composites and closes with the adoption of variable stiffness.

After introduction of carbon composite into some of its empennages, Boeing participated in the Advanced Technology Composite Aircraft Structures(ATCAS) and researched the production of fuselage section 46 using carbon composite. Figure 4.3 shows the location of section 46 in the 787. The program, funded by the government, aimed to investigate and overcome some of the design and manufacturing issues that prevented the use of carbon composites. Section 46, which is located directly aft of the wingbox, was chosen because of the simple geometry while containing all of the important structural elements such as windows and door cutouts. A design that consisted of four sections was drafted and tested part by part. [21, 32, 53]

A similar design was adopted by Airbus for the same section of the A350 XWB. Similar manufacturing technology as envisioned by the ATCAS engineers are used for the production of the four sections. Stringers are located using jigs and co-bonded to the skin during cure of the individual sections. Afterwards these sections are joined together to form the cylindrical section. [15]

Although deemed too risky during the ATCAS program, Boeing now uses a single mandrel for fuselage sections of the 787. After placement of stringers on the mandrel, an AFP machine is used to create the layup of fuselage's skin. The outer surface is formed using a floating tool during cure. Developments in mould accuracy and tooling techniques allows for these large scale structures to be cured to a high dimensional accuracy. During the nineties this was deemed impossible and large parts were avoided due to the complications that would arise in final assembly. [15, 21]

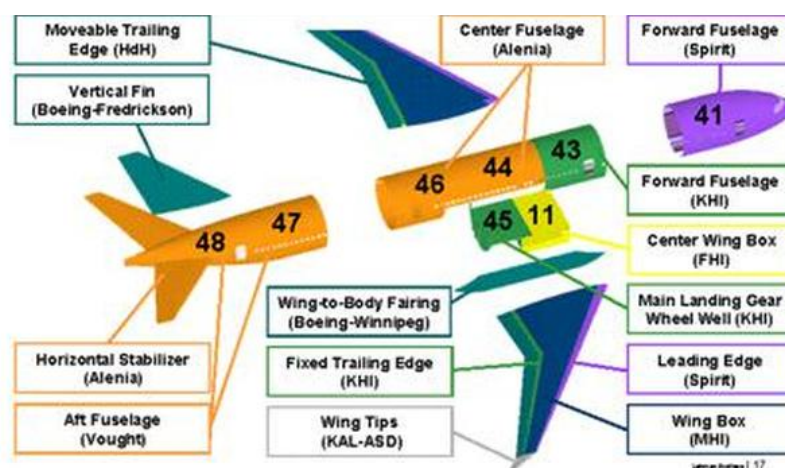


Figure 4.3: This image shows the locations and number of the sections in the modern 787 program. Boeing has always numbered the section directly aft of the wingbox 46. This particular section was the subject of ATCAS program. Image courtesy of Boeing¹.

¹<https://www.zdnet.com/article/is-787-dreamliner-novel-manufacturing-strategy-on-the-rocks/> accessed October 5, 2010

Most laminates used in aerospace have a constant layup and use additional plies at locations with structural features to increase its strength. Tailoring of fibers has not been used to its fullest extent and as a consequence the savings predicted in research have not been achieved.

But, if the figures presented in the previous chapters could be turned into reality, significant cost reductions can be accomplished. Some papers suggest the annual savings of 1kg weight reduction to be as high as US\$5000,- for short haul and US\$17000,- for long haul flights. [11] A broader accepted number relating DOC to weight savings ($\Delta C/\Delta W$) is around US\$500/kg. [23, 54] The significant difference between these numbers is partly explained by the rise in fuel prices between the report's dates. However, the price is not projected to fall; especially with the current focus on renewable fuels. Consequently, weight savings become more important than ever. [9]

In conclusion, technologies that save more cost via weight reduction than they add via manufacturing are beneficial to the industry. The cost break down in the first section shows that the airframe and fuel cost have an equal share in the direct operating cost for an airline. When average weight-to-cost numbers of US\$500/kg are used, a 1kg mass reduction of the airframe may come at a US\$500 manufacturing cost increase. Any cost increase below that number, reduces the cost of air travel. And, not to be forgotten, lightens the burden on the environment by reducing emissions.

Research for a Higher Adoption of SFVS Laminates

Although the potential of Straight Fiber Variable stiffness has been proven in experiments[12], their design methods are still mature. An integrated method to find the optimal stiffness distribution and generate a SFVS laminate that approaches the optimum does not exist. Efficient stiffness optimization, stacking sequence retrieval, and local blending algorithms do exist separately[34, 50, 52] but their integration has not been researched properly. Moreover, current algorithms sometimes do not converge to locally blended laminate and the global interpretation is not yet automated. These shortcomings prevent the application of SFVS laminates in larger structures.

5.1. Objective


This research is centered around SFVS and tries to create an understanding of their possible application in large aerospace structure. This understanding is the basis for a higher adoption in the industry and aims to create a suitable method for the application. To this end, an end-to-end framework is proposed that fully automates the design and interpretation of SFVS laminates. Then this framework can be used to assess the potential of these laminates and identify areas of interest. The formal objective is:

The objective of this research is to investigate the structural performance potential of straight fiber variable stiffness laminates in large aerospace structures by developing a framework to incorporate these laminates at a large scale.

5.2. Questions

The research is guided by a coupling of questions that aim to describe how previous research has used simple shapes to optimize real structural elements. Related to this is the modelling of complex failure phenomena and the usage of FEA in optimization. Related to the current state-of-the-art of SFVS laminates two questions are formulated to improve the design methods and cost quantification. The resulting set of research questions is:

1. *Why are the application to plates and cylinders representative for the use in large aerospace structures?*
2. *How can finite element analysis be used to design a structure with variable stiffness such that the buckling load is maximized?*
3. *How should the existing blending algorithm be adapted such that it can be used to design SFVS laminates for large aerospace structure?*
4. *How can the Pareto frontier, that both considers the buckling performances and manufacturing complexity, be generated?*



A Blended Design Framework

The result of this thesis is an integrated design framework that produces laminate designs that utilize straight fiber variable stiffness to vary the stiffness between regions. The resulting laminate structure can transfer loads via the designed paths and thereby increase the structural potential. The designs are manufacturable using automated manufacturing techniques such as Automated Fiber Placement or Automated Tape Laying. But, in contrast to fiber steered designs, can also be manufactured using manual pick and placement of composite patches.

Throughout this part of the report specific nomenclature is used to dictate specific parts of the design. These are introduced below:

Design	The design includes all the specifics of the product that is designed. This includes both the stacking sequence at every place in the laminate but also the order in which the carbon composite plies should be layed down. These together provide all the information that is necessary for production of the laminate.
Region	Part of the laminate that has a constant layup and therefore a constant stiffness. Multiple regions together make up the total design and the difference in stiffness between regions provides the laminate with variable stiffness.
Layer	A region consists of multiple individual layers through the thickness of the laminate. A layer is a single level in the laminate bounded by the bounds of the region. Layers have a single fiber orientation and thereby form the basic entity of a SFVS laminate.
Patch	Multiple layers are connected together in a Patch. A patch covers multiple regions and is the basic manufacturing entity. In practice a patch is put down in a single pick and place operation or put down at once by an AFP or ATL machine. Patches describe how a layers connect to layers in other regions and thereby determine how edges are blended.
Interface	An interface is the transition between two adjoint regions. At the interface layers end or can continue to another layer in the adjoint region.
Connection	At an interface; two layers that belong to the same patch have a connection. Connections determine how layers continue from region to region and which stopped edges are covered.

This part of the report will introduce the framework developed. In chapter 6 the finite element modelling is discussed in terms of the eigenvalue problem and variable stiffness modeling technique. Then, in chapter 7 the local optimization procedure together with the stacking sequence retrieval is explained. In chapter 8 and chapter 9 the retrieval of a local and global manufacturable laminate are described respectively. A method to qualitatively asses the cost of a design is introduced in chapter 10.

Buckling Prediction using Commercial Finite Element Analysis

Buckling of a plate in uniaxial compression is a problem that has been solved many times. For simple constant stiffness composite plate the buckling load can be accurately approximated using the following formula[22]:

$$N_0(m, n) = \frac{\pi^2 [D_{11}m^4 + 2(D_{12} + 2D_{66})m^2n^2 + D_{22}n^4(AR)^4]}{a^2m^2} \quad (6.1)$$

where D_{ij} are the stiffness components of the D matrix, a is length of the plate in the direction of loading and AR is the aspect ratio. Here the designer needs to chose the buckling mode, depict by m and n which are the number of halfwaves in the direction of and perpendicular to the load respectively, such that the minimum buckling is retrieved. Buckling of cylinders is also strongly related with the flexural stiffness of the shell.

Optimizations in research[20, 24] use custom written codes to quickly predict the structural performance of a set of design variables. However, reliance on a custom written code for the buckling prediction prohibits usage of this framework in other design cases. For every design case a new code has to be created that solves the governing eigenvalue problem associated with determining the bifurcation point. The bifurcation point is associated with buckling load and provides a reasonable estimate at which load the post-buckling regime is entered. The scope of this thesis is limited to design for buckling load and the post-buckling regime is neglected.

To overcome the reliance on custom codes and allow optimization of complex structures such as the cylinder in bending, an optimizer that simulates structures using commercially available FEA is needed. Here Abaqus is used to perform the eigenvalue analysis of the structure using the buckling step. However, the output does not have to be limited to eigenvalue analysis. A similar approach as the one presented here can also be used to optimize stiffness or strength of laminates.

This chapter provides a detailed explanation how Abaqus is used to predict buckling of SFVS laminates. First, the buckling analysis that Abaqus performs is analyzed in section 6.1. In section 6.2 the element stiffness formulation and the modelling techniques to obtain variable stiffness are discussed.

6.1. Buckling Analysis

To verify that the framework presented here creates viable and optimal designs several buckling benchmarks are used. Buckling benchmarks are excellent because they require global stress redistribution and benefit from large stiffness changes in the structure. Buckling analysis is limited to finding the bifurcation load using the Abaqus buckle routine. The post buckling regime is neglected in the design and failure of the plate is assumed after loading it beyond the bifurcation load. Failure means that the plate does no longer comply with the design constraints and here it assumed that loading a structure into its post-buckling regime is undesirable even though it can still take load.

6.1.1. Eigenvalue Problem

A static perturbation analysis in Abaqus finds the load increment for which the tangent stiffness matrix becomes singular. The tangent stiffness matrix (K^{NM}) is constructed from two parts: The stiffness matrix from the base state (K_0^{NM}) and the geometric stiffness matrix (K_δ^{NM}). This second matrix represents the change in the tangent stiffness matrix as a result of the loading pattern. The eigenvalue problem is then presented as:

$$(K_0^{NM} + \lambda_i K_\delta^{NM}) v_i^M = 0 \quad (6.2)$$

where λ_i and v_i^M are the eigenvalue and eigenvector respectively of each eigenmode i . [1]

This method provides the theoretical buckling load of a structure and does not take into account defects or imperfections. Often the real buckling load of structure is lower than the one approximated with this method due to material imperfections and differences in the thickness throughout the plate. For this work these imperfections are neglected and the bifurcation load is used as the buckling load of the structure.

6.1.2. Preconditioning

Abaqus' eigenvalue solver first finds the lowest eigenvalue, which is in the case of the plate the buckling load of interest. In cases where an negative perturbation can also cause buckling this becomes problematic. For example a constant stiffness cylinder in bending buckles at the top and at the bottom for the same but opposite moment. A variable stiffness cylinder buckles at a lower load in one direction than in the other direction. Such a constraint cannot be imposed in Abaqus directly, because the internal eigenvalue solver will compute the first eigenvalue with the lowest magnitude and does not take into account the sign.

Instead, the base state of the structure can be changed to cause buckling to happen earlier in one direction than the other. The linearity of the buckle solver allows the base state to be preconditioned with a preload, $P^N = \mu Q^N$, and then loaded using the load pattern Q^N . Preloading alters the base state and causes a lower buckling load in that direction which aids in eigenvalue extraction. The real buckling load is then obtained through the sum of the preload factor and the eigenvalue: $(\mu + \lambda_i)Q^N$. This is especially useful to extract the buckling load of a cylinders that are optimized for bending in one direction. Theses cylinders generally have a lower buckling load in the other direction. [1]

6.2. Stiffness Modelling

Abaqus accepts different types of element properties to formulate the tangent and geometric stiffness matrices needed for the eigenvalue extraction. Throughout the optimization an efficient and adequate element formulation based on lamination parameters is needed. But when the stacking sequence is known the additional information about the layup can be used.

6.2.1. Element Setup

All models used in the optimization are constructed using S4 elements which gave 6 degrees of freedom at every node. These shell elements are easy to program while also providing reasonable results. In a direct comparison to continuum shell elements SC8R, the buckling load is less than 1% lower. This result is comparable to results obtained by others. [8, 12] Moreover, S4R elements, which are based on S4, result in a more accurate buckling load prediction for lower number of elements than continuum shell elements. [2] A more detailed analysis of the effect of mesh convergence on the optimization is performed in Appendix E. Care was taken to use a mesh with the same granularity as those used in literature[5, 20] for good comparison.

The stiffness of a composite shell element can be described in several ways: A laminate layup with material properties can be defined in Abaqus; Or the stiffness of the element can be described directly; It is preferred to use the first definition because Abaqus will then take care of solving for the transverse stiffness of the laminate. But this method can only be used when the stacking sequence is known which is not the case in during initial stiffness optimization. As will be explained when introducing the optimization, lamination parameters are used here. They are not easily translated to a stacking sequence and therefore a general stiffness definition is necessary in this phase.

When using general stiffness elements in Abaqus the stiffness matrix should be provided as input. The general stiffness matrix is equivalent to the already introduced, see subsection 2.3.2, ABD matrix.

Although care should be taken of the coordinate systems in which they are derived, the FEA solver will directly work with these values.

The elements transverse shear stiffness needs additional consideration because the standard prediction of Abaqus overestimates the actual values. Instead a formulation based on homogeneous shells with a linear elastic orthotropic material is advised by the manual:[1]

$$\begin{aligned} K_{11}^{ts} &= \frac{5}{6} G_{13} h, \\ K_{22}^{ts} &= \frac{5}{6} G_{23} h, \\ K_{12}^{ts} &= 0 \end{aligned} \quad (6.3)$$

where K_{ij}^{ts} are the individual elements of the transverse shear stiffness matrix, G_{ij} are the shear moduli of the laminate and h the laminate thickness. These laminate shear moduli are equivalent to the compliance entries C_{44} and C_{55} in the compliance matrix for general laminates:

$$C_{ij} = \begin{bmatrix} C_{11} & C_{12} & C_{13} & 0 & 0 & C_{16} \\ C_{12} & C_{22} & C_{23} & 0 & 0 & C_{26} \\ C_{13} & C_{23} & C_{33} & 0 & 0 & C_{36} \\ 0 & 0 & 0 & C_{44} & C_{45} & 0 \\ 0 & 0 & 0 & C_{45} & C_{55} & 0 \\ C_{16} & C_{26} & C_{36} & 0 & 0 & C_{66} \end{bmatrix} \quad (6.4)$$

such that $G_{13} = C_{44}$ and $G_{23} = C_{55}$ for orthotropic materials. It is later shown that indeed K_{12}^{ts} can be neglected for the laminates in consideration.

By using fourth order tensor rotations the shear stiffness of a laminate can be obtained as a function of its individual layer's orientations, $\theta(z)$, as

$$\begin{aligned} C_{44}^L &= \int_{-\frac{h}{2}}^{\frac{h}{2}} \left(\frac{1 + \cos 2\theta(z)}{2} C_{44}^P + \frac{1 - \cos 2\theta(z)}{2} C_{55}^P \right) dz \\ C_{55}^L &= \int_{-\frac{h}{2}}^{\frac{h}{2}} \left(\frac{1 + \cos 2\theta(z)}{2} C_{55}^P + \frac{1 - \cos 2\theta(z)}{2} C_{44}^P \right) dz \\ C_{45}^L &= \int_{-\frac{h}{2}}^{\frac{h}{2}} \left(\frac{\sin 2\theta(z)}{2} (C_{55}^P - C_{44}^P) \right) dz \end{aligned} \quad (6.5)$$

where superscript P denotes properties associated with the ply and L properties associated with the laminate.

Collecting terms and introducing two new material invariants based on the shear stiffness results in

$$\begin{aligned} U_6 &= C_{44}^P + C_{55}^P \\ U_7 &= C_{55}^P - C_{44}^P \end{aligned} \quad (6.6)$$

allows the laminate shear stiffness to be written as a function of the first two membrane lamination parameters, V_1 and V_2 :

$$\begin{aligned} C_{44} &= U_6 + V_1 U_7 \\ C_{55} &= U_6 - V_1 U_7 \\ C_{45} &= V_2 U_7 \end{aligned} \quad (6.7)$$

Since all laminates in this thesis are considered balanced, C_{45} is neglected. C_{44} and C_{55} are directly used in Equation 6.3 to compute the equivalent shell's transverse stiffness.

6.2.2. Variable Stiffness

Variable stiffness is incorporated in the FE Model by changing the element stiffness accordingly. It depends on the level of detail known during the optimization which modelling techniques is used. During

the initial optimization only lamination parameters are known, therefore general stiffness elements are used.

A variable stiffness design consist of regions in which the lamination parameters and layup is considered constant. Care is taken to locate elements in such a way that elements strictly lay within one region only, thereby avoiding the need to interpolate stiffness between elements. As a result elements can be assigned to a single element set representing all elements in one region. Unique shell sections are assigned to the element set, representing the properties of that specific region.

This approach prevents the need for stiffness interpolation as used by IJsselmuiden et al.[20] and Van Campen et al.[49], which better represents the true structural behavior. However, it results in relation between the number of regions and the mesh size. As a consequence designs with a different number regions are harder to compare because they might have a different number of elements. A sufficiently small mesh size is chosen using a mesh convergence to overcome this problem.

To partly overcome this problem, an upper mesh size is used in the analysis of the plate. During model setup, a mesh size is selected that is near the threshold but is a multiple of the number of design regions. For example, when 9 regions are used and the upper mesh limit is 60, a total of 54 elements in both directions are used.

This method makes the comparison between different designs more fair, as the number of elements does not drastically change from model to model. The alternative is to use a constant number of elements per region, but this would have a large influence on the total number of elements and therefore on the simulation time as well.

Optimal Design of a Variable Stiffness Laminates for Maximum Buckling Load

In a review conducted by Lozano et al. in 2016, it was concluded that many researchers use multi-step methods to design composite structures. [27] The framework developed in this thesis also separates the optimization of the local stiffness and generation of a manufacturable laminate into multiple steps. The four steps of this framework each have their own purpose and their own set of constraints.

This chapter will first provide an overview of the these four steps in section 7.1. Then, the local optimization process which makes use of the lamination parameter hyperspace is detailed in section 7.2. Lastly, section 7.3 introduces the conversion from lamination parameters to stacking sequences.

7.1. Multi-Step Method

Optimization in multiple steps allows the usage of the best optimization method and parameterization for each subproblem. Here an approach is used that first tries to find an optimal design and then reduce that to a blended solution at the cost of structural performance. The end result is a producible design that complies with the relaxed guidelines and has an increased structural performance over a conventional laminate consisting of $[\pm 45]_s$. An overview of the framework is presented in Figure 7.1.

After a parametric model of the problem in Abaqus is made, the local stiffness optimizer uses lamination parameters to minimize the inverse buckling load. Minimization of the inverse buckling load renders the problem convex[20] and allows the use of an efficient gradient descent method. A surrogate model is used to provide buckling load predictions and gradient information thereby reducing the number of finite element analyses. The initial dataset for the surrogate model is created using Latin Hypercube sampling (LHS) and is further refined based on the optimization state. The optimization terminates after a specified number of iterations have not yielded improvement or a maximum number is reached. The particulars are further detailed in section 7.2

Next, for every region multiple stacking sequences are retrieved that all closely match the LPs of the previous step. A newly developed genetic algorithm (GA) based on the work of Van Campen et al.[47] is used for every region individually to generate multiple guides per region. The best guides of every region are joined to the best design. To aid in blending, alternative designs are generated by swapping some guides for a lower fitness. Structural performance is traded for blendability in subsequent steps. The GA and alternative designs are further discussed in section 7.3.

A locally blended laminate design is obtained with a cellular automaton (CA) based on the work of Van den Oord. [50] The CA has been extended with cells consisting of multiple regions to increase the butted edge removal performance and a new minimum cutting size strategy to preserve individual patches. A new relaxed blending guideline detection algorithm is introduced to assure compliance with the relaxed blending guideline. Also mathematical formulations are introduced to aid in describing the algorithms working. The cellular automaton (CA) is described in detail in chapter 8.

To create a design that complies with all the blending rules and is producible with current manufacturing technologies a newly developed global interpretation algorithm is used. This algorithm uses

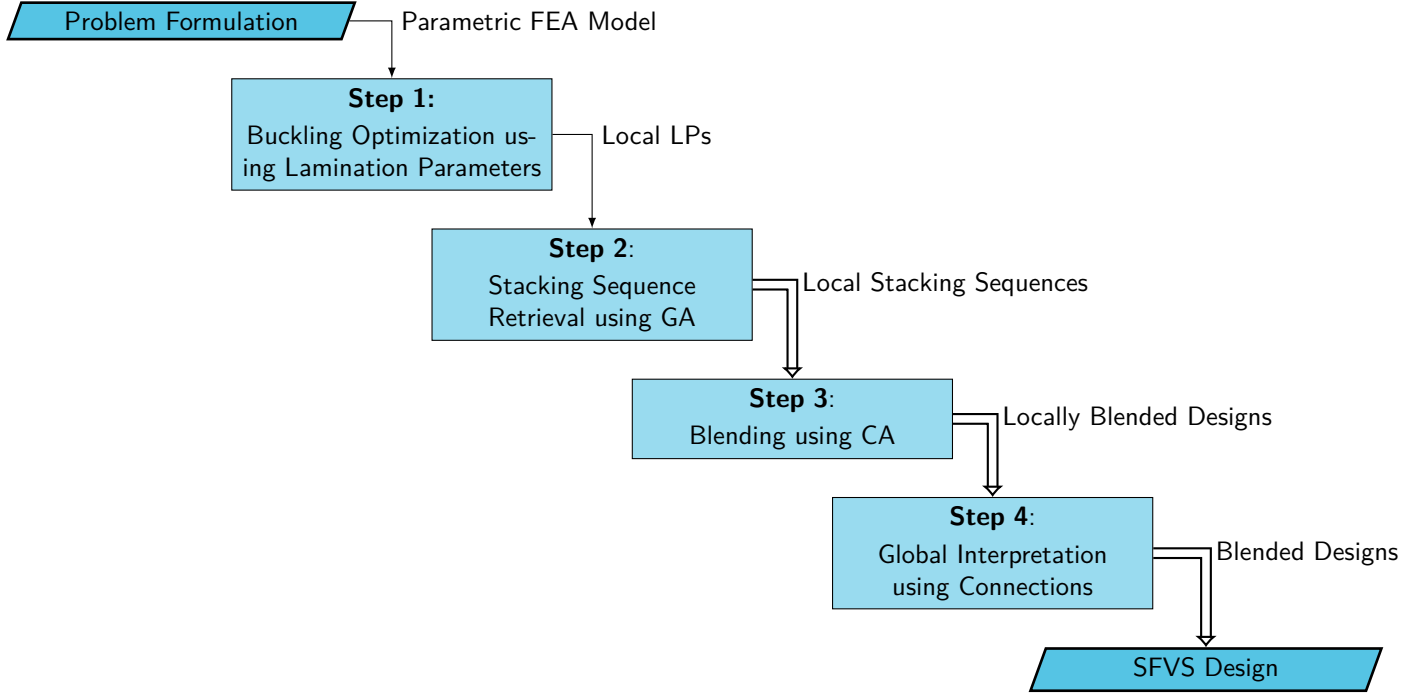


Figure 7.1: Overview of the multi-step optimization process as used by the SFVS design framework.

connections between layers to formulate patches that can be stacked in order. Interweaving of patches is circumvented by requiring a monotonically increasing patch assignment in every region. The output consists of a patch identification and orientation for every layer of every region. The patch identification serves to identify which layers are connected and as the stacking order. The algorithm is discussed in chapter 9

The output of the framework is a globally blended design that incorporates variable stiffness through varying stacking sequences. The design specifies a stacking sequence per region that complies with the relaxed blending guidelines at every interfaces and guarantees that every edge can either continue into the next region or can be covered. Additionally, the patch identification guarantees manufacturability using modern technologies.

7.2. Sampling Assisted Local Stiffness Design

The use of LPs in local stiffness optimization makes the design variables continuous. By limiting the problem to symmetric and balanced laminates only four lamination parameter (LP) have to be designed for every region of the laminate. The parameters are physically constrained to the feasible design space discussed early in subsection 2.3.1 which can be forced using linear inequalities through a convex hull approximation. The problem becomes a convex minimization problem when the inverse of the buckling load is taken as the minimization objective. The mathematical formulation of the stiffness optimization is then given by:

$$\begin{aligned}
 \min_{\mathbf{x}} f(\mathbf{x}) &= \min_{\mathbf{x}} \frac{1}{L(\mathbf{x})} \\
 \text{s.t.} \\
 g(\mathbf{x}) &\leq \mathbf{0} = A\mathbf{x} \leq B
 \end{aligned} \tag{7.1}$$

where \mathbf{x} is a vector containing four lamination parameters ($V_{1,j}$, $V_{3,j}$, $W_{1,j}$, and $W_{3,j}$) per region j , $L(\cdot)$ computes the buckling load for a given design vector, and A and B are the constants for the linear inequalities to constrain the LP to the convex hull.

Availability of the following derivate formulation allows the use of a gradient solver:

$$\begin{aligned}\frac{df}{d\mathbf{x}} &= \frac{-1}{L(\mathbf{x})^2} \frac{dL}{d\mathbf{x}} \\ \frac{dg}{d\mathbf{x}} &= A\end{aligned}\quad (7.2)$$

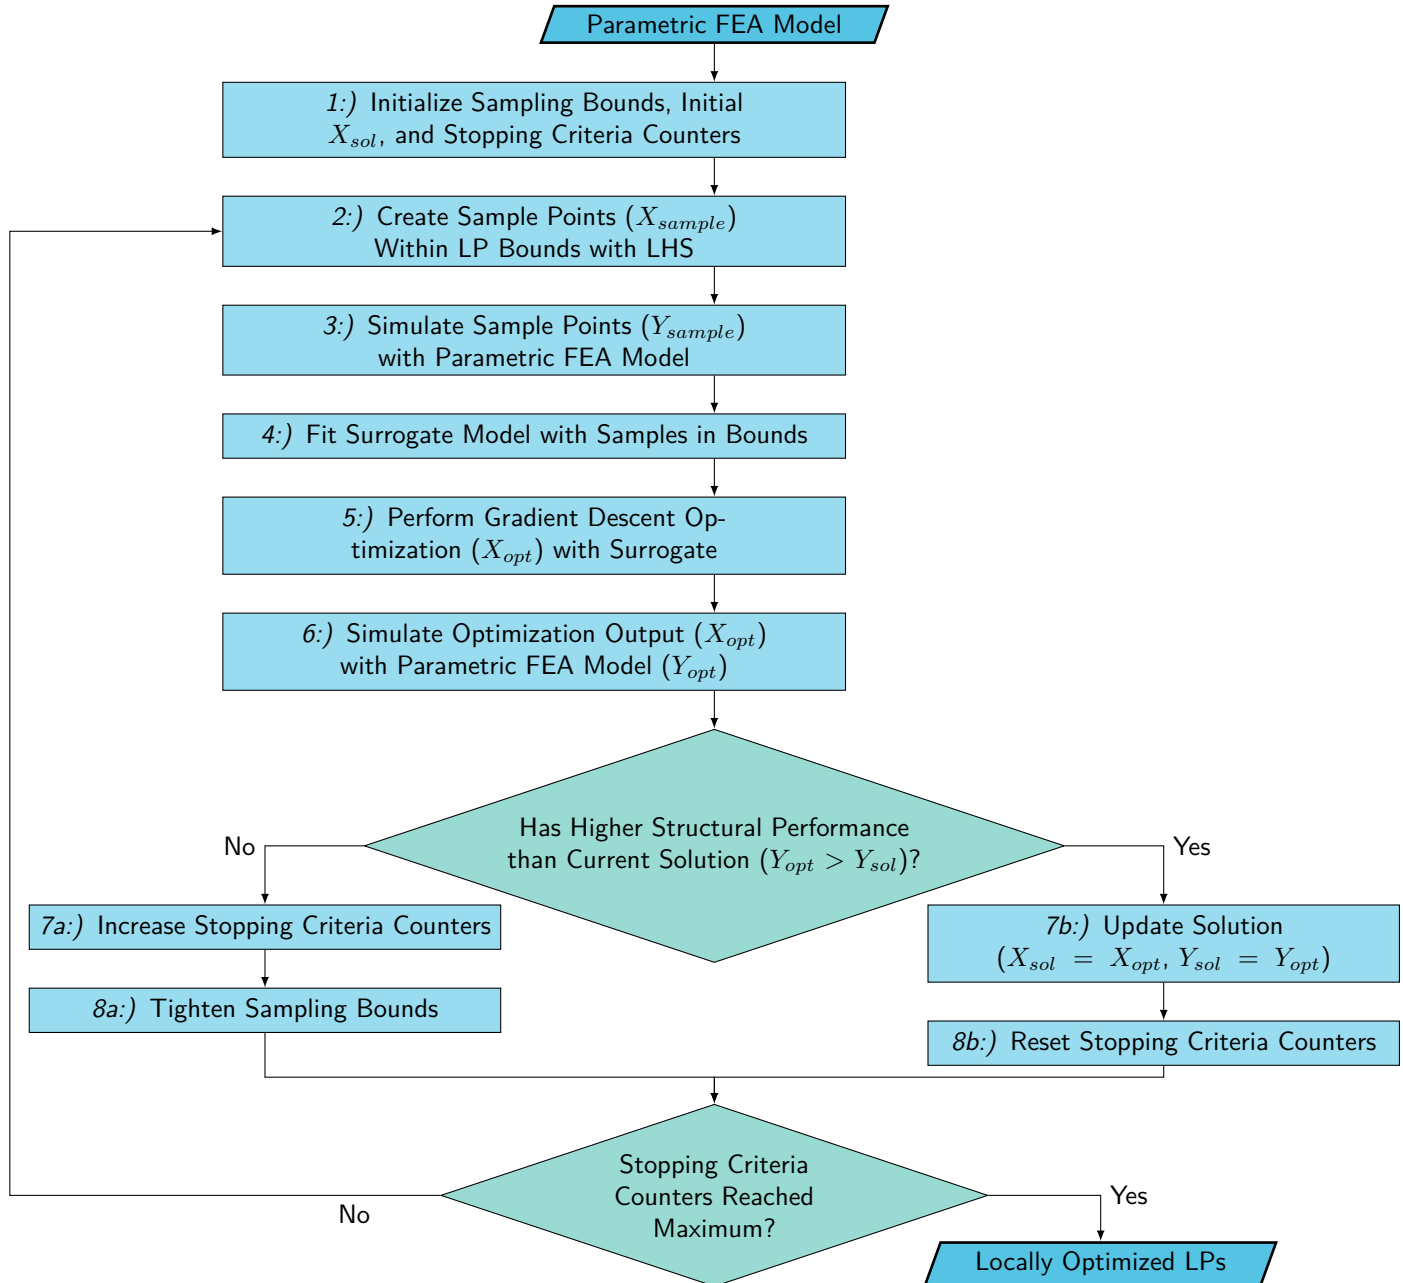


Figure 7.2: Overview of the Sampling Assisted Local Stiffness Optimization Process via a Parametric FEA Model

Evaluation of the buckling performance is computationally expensive and a surrogate modelling technique is used to predict $L(\cdot)$ and $\frac{dL}{d\mathbf{x}}$ during the optimization. For accuracy a large data set is necessary that has to be made using costly finite element analysis. To limit the number of data points needed a progressive detailing scheme is used similar to the one of Rouhi et al.. [34]. The optimization process is visualized in Figure 7.2.

The process starts (step 1) with initializing stopping criteria and setting the initial solution to zero. Sample points that cover the entire design space are drawn (step 2) and simulated using the parametric Abaqus model to form the initial dataset (step 3). After fitting of the surrogate with the dataset (step 4), a gradient descent optimization utilizes the efficient value and derivative predictions to estimate a new optimal design located in the feasible design space (step 5). Design evaluation in FEA (step 6) leads to two possible paths: 1) The new design is better than the current solution and becomes the new solution (steps 7b and 8b) or 2) the new design is not better which leads to the old design to be kept and the sampling bounds to be tightened around the current solution (steps 7a and 8a).

Convergence of the optimization process is tracked by counting the total number of iterations and the number of subsequent iteration misses. Every time an iteration does not yield an improved design is counted as a missed iteration. When one of the two counters exceeds a threshold the optimization process is terminated and the processes is said to be converged. When the counters have not exceeded their maxima a new iteration is started. The new iteration yields the same processes as previous except that new sample points are appended to the already existing dataset.

In the following subsection parts of the algorithm are detailed. First the sampling and surrogate modelling techniques are covered, then the internal gradient descent optimization, and lastly the stopping criteria and bound tightening.

7.2.1. Design of Experiments

The first step of every iteration is to determine the location of the new sample points, from here on called $X_{samples}$. These points are then simulated using the parametric Abaqus model to find the corresponding $Y_{samples}$ buckling performance values. Together they form the dataset that can be used to fit the surrogate discussed in the next section. Singular forms X_{sample} and Y_{sample} are used to indicate a single sample point.

Because the dataset forms the bases of the optimization process it is important that a technique is chosen that covers the design space equally. [30] A commonly used technique is Latin Hypercube sampling (LHS) which is visualized in Figure 7.3. This technique aims to take an unbiased sample that covers every corner of the design space. [41]



Figure 7.3: 4 by 4 square divided such that the markings comply with latin hypercube sampling: Every row and column holds a marking only once. Image taken from Sheikholeslami and Razavi[41].

Generating a Latin Hypercube sampling (LHS) sample is a simple procedure that is best explained through the example of Figure 7.3. The square represents a 2D design space where the X and Y-axes form the dimensions. In the presented case a sample of four points is generated by dividing every axis in four equal segments to make a regular grid. Four cells in the grid are selected such that only one cell per column and row is selected. There cannot be two cells in the same column or row. The center of each cell then acts as the chosen sample points. Cells with the same marking in Figure 7.3 are sampled together. In multiple dimensions every segment is selected only once for a sample point.

Such a strategy requires a constant interval for each of the dimensions that is independent of the other dimension's value. This is not necessarily the case when working with lamination parameter

(LP) because of the coupled feasible space. Some researches have generated sample points using fiber angles[34] but that is unpractical as it would require constant conversion to stacking sequence to account for sampling bounds which are computed in LP space and further discussed at the end of this section.

Here an update Latin Hypercube sampling (LHS) method is used that preserves some of the desired behavior while working with discrete points. Every region in the variable stiffness laminate acts as one dimension on the latin hypercube. For every region, i , several possible combinations of lamination parameters, $X_{\text{subsample},i}$, are created to act as a segment along the dimension. Sample points are then generated through combining sub samples ($X_{\text{subsample},i}$) in a latin fashion such that every LP combination of a region is used only once. This is similar to how every column and every row is sampled only once in Figure 7.3.

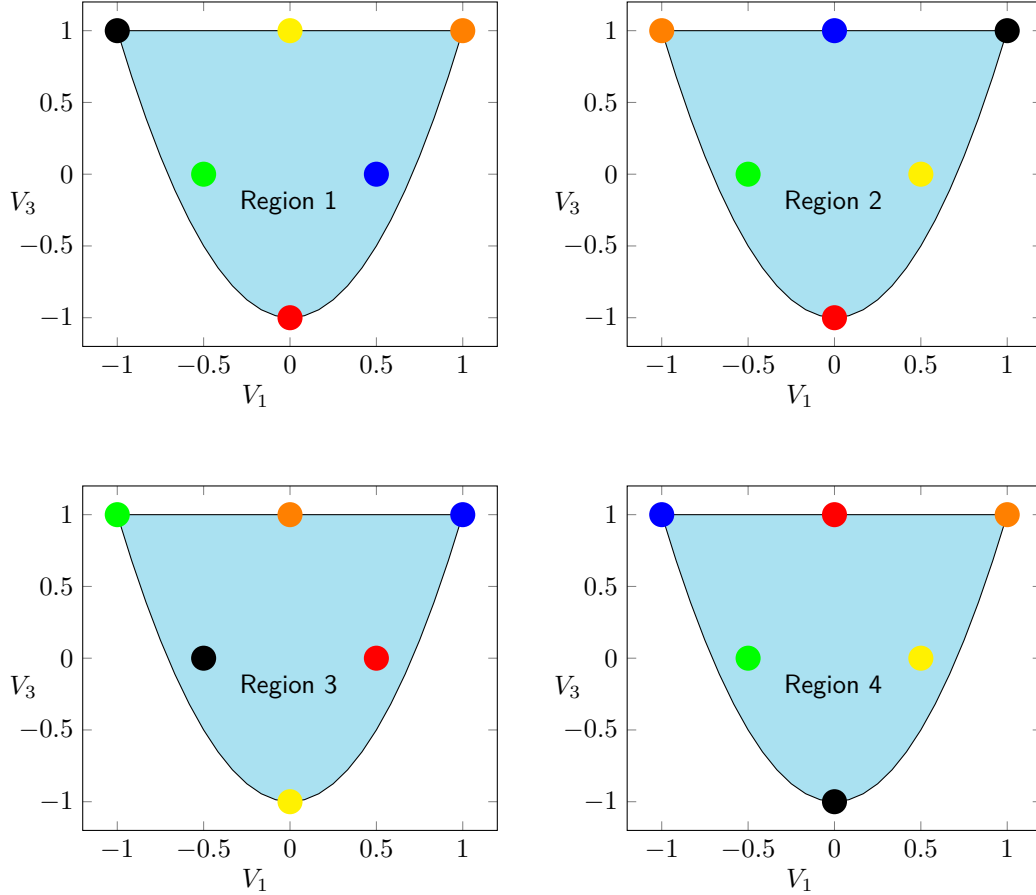


Figure 7.4: Example of 6 sampling points for 4 regions according the LHS based sampling technique. Dots with the same color belong to the same sample point and together form the sample point as $\{V_{1,1}, V_{3,1}, V_{1,2}, V_{3,2}, V_{1,3}, V_{3,3}, V_{1,4}, V_{3,4}\}$.

Figure 7.4 shows an example where 6 sample points are generated for 4 regions. Here dots in every region are the sub samples $X_{\text{subsample},i}$ for region i . Dots with the same color are combined into one sample. Every color is present in every region and is present only once to preserve the span of the design space. In this example the full red sample point is:

$$\begin{aligned}
 X_{\text{sample}} = & \\
 \text{Region 1}(X_{\text{subsample},1}) & : V_1 = 0, \quad V_3 = -1 \\
 \text{Region 2}(X_{\text{subsample},2}) & : V_1 = 0, \quad V_3 = -1 \\
 \text{Region 3}(X_{\text{subsample},3}) & : V_1 = 0.5, \quad V_3 = 0 \\
 \text{Region 4}(X_{\text{subsample},4}) & : V_1 = 0, \quad V_3 = 1
 \end{aligned}$$

It should be noted that in the real optimization problem four lamination parameters per region are

computed using this technology and this is merely a visualization of the 4D sampling space.

The sub samples of every region are created by converting a regular grid of balanced symmetric laminates in terms of fiber orientations to lamination parameters. For example, the points in Figure 7.4 use two plies and fiber orientations $\{0, \pm 45, 90\}$ to create laminates $[90]_s$, $[(90)_2 / \pm 45]_s$, $[(90)_s / (0)_2]_s$, $[0]_s$, $[\pm 45]_s$, and $[(0)_2, \pm 45]_s$.

To increase the quality of the dataset based on the optimization's state a global sampling band S is used. The sampling band is initialized to 1 at the start of the optimization process and reduced every missed iteration with

$$S_{next} = \frac{S}{\sqrt[8]{2}} \quad (7.3)$$

such that it is halved every 8 missed iterations. A minimum sample band of 0.1 is maintained throughout the optimization.

After creation of a grid with possible LP combinations, only those that are located in the hypercube of size $2S$ centered around the current solution X_{sol} are selected for the latin combination algorithm. This guarantees that every full sample point, X_{sample} has a maximum difference of S per dimension to the current solution X_{sol} :

$$|X_{sample,i} - X_{sol,i}| \leq S \quad (7.4)$$

The sampling points generated are appended to the dataset, represented by $X_{dataset}$ and $Y_{dataset}$. Although every sample $X_{samples}$ on its own is nicely distributed over the design space, this cannot be said over the union of subsequent samples. [41] As a result, the final dataset is more detailed in some areas than others.

7.2.2. Surrogate Modeling

Based on a study by Nik a simple Kriging model is selected. [30] Simple Kriging models closely resemble Gaussian Regressor models which use kernels to fit data and can compute a covariance for predictions. The uncertainty can be used to judge the prediction and can be used to tune the hyperparameters of the model.

The prediction of a surrogate model is composed of the regressor and the stochastic part. The first part is similar to a normal regressor model, but the second part takes into account the quality of the model at the local point. The standard deviation of the prediction in locations where the dataset is less dense is higher compared to more dense areas.

The regressor and stochastic part are coupled into a kernel that can be chosen by the surrogate architect. Here a Radial Basis Function similar to the one used by Nik is chosen as the kernel. During fitting the width of the RBF is tuned by maximizing the log-marginal-likelihood of the surrogate on the given data. The actual surrogate model used is implemented using a 'GaussianProcessRegressor' developed and maintained as part of the 'SciKit' module¹.

The dataset that is used for the fit is a filtered subset of the complete dataset. Only those samples that are located in the current sampling band are selected for the fit according to Equation 7.4. The selected data points are then scaled linearly such that the maximum and minimum of the data is zero and one for every dimension. Gaussian kernels work best in the range between 0-1. [33]

7.2.3. Constraint Gradient Descent Optimization

The central part of the optimization process is the internal constrained gradient descent optimization. Equation 7.1 is directly minimized where the buckling load value is provided by the surrogate model. Performance derivative $\frac{dL}{dx}$ are computed using a forward difference scheme with step 10^{-5} .

Optimization is performed using sequential quadratic programming via the 'minimize' functions of 'SciPy'² and the 'SLSQP' method that implements the work of Kraft[26]. Inequality constraints are

¹https://scikit-learn.org/stable/modules/gaussian_process.html, Last accessed on May 16, 2021

²<https://docs.scipy.org/doc/scipy/reference/optimize.minimize-slsqp.html#optimize-minimize-slsqp>, Last accessed May 16, 2021

implemented using an active-set method which takes benefit from the linearization and convex hull approximation of the lamination parameter space. [26]

The starting point for the optimizer is the current best solution X_{sol} and the objective is to find a new optimum design in the design space approximated by the surrogate. The output is a set of lamination parameters X_{opt} that represents the highest structural performance according to the surrogate.

Before advancing to this new solution which entails setting X_{sol} to X_{opt} , an Abaqus simulation is ran using this design. The output is a new buckling load that is expected to be higher than the current solution. If this is the case the iteration was successful and the solution is advanced. If this is not the case the iteration counts as a miss and the new solution is disregarded.

7.3. Stacking Sequence Retrieval

In the second step of the local optimization a stacking sequence is designed that closely matches the lamination parameters. When it is assumed that the local lamination parameters are known for all the regions i , here called $LP_{local,i}$, the stacking sequence retrieval process can be formulated as an optimization problem:

$$\min_{\theta_{i,j}} \left(LP_{local,i} - LP(\theta_{i,j}) \right)^2 \quad (7.5)$$

where $\theta_{i,j}$ denotes the angle of ply j in guide i , and $LP(\cdot)$ is a function that converts the stacking sequence, $\theta_{i,j}$, to lamination parameters per guide i .

The ply orientation is chosen from an integer set of plies, which makes the optimization an integer problem. It has been shown in literature that finding an exact match for the lamination parameters is impossible except for some special cases. [35]. Here an evolutionary approach is used to find stacking sequences that closely match the guide lamination parameters. This is explained in the next subsection.

It might be beneficiary for the blending process to start from a different combination of guides. Because several stacking sequences can approximate the LP to a similar degree, several combinations are passed to the blending step. These combinations are generated by selecting other options for guides as explained in the last subsection.

7.3.1. Genetic Algorithm

The genetic algorithm used is relatively simple multi-chromosomal genetic algorithm (GA) as also presented by Van Campen et al.[47] and originally described by Adams et al.[4]. The fitness of each individual is directly determined by the objective formulated in Equation 7.5, where those individuals with a lower fitness have a higher chance of survival. The algorithm is stopped when it passes 300 iterations or when the fitness drops below $1e-3$.

In short the GA evolves a parent population with individuals to a new child population via a crossover operator. Individuals with a higher fitness have a higher chance to be selected for the formation of the new child population. The crossover operator combines the guides of the parents to two new guides. A random position in the guide is chosen where the crossover between the parent guides occurs. The part before the crossover point belongs to parent 1 and the part after belongs to parent 2.

Then several genetic operators such as gene swap, permute, and mutate are applied to create new imperfections. The gene swap operator reverse a part of the guide at random and effectively inverses the stacking sequence of a part of the laminate. Permute swaps to locations in the gene meaning that two patches in the laminate are swapped from location. Lastly the mutate operator acts on a single patch and changes it to a random other orientation.

The child populations and 5 elite individuals form the parent parent population of the next iteration. Elite individuals have the lowest fitness of the current population. By carrying them to next evolution, the best solution is not lost due to crossover. All parameters are summarized in Table 7.1.

Table 7.1: Parameters used for the Genetic Algorithm.

Parameter	Value
Max number of iterations	300
Fitness threshold	10^{-3}
Number of individuals	300
Number of elite individuals	5
Probability of mutation	10%
Probability of permutation	10%
Probability of swap	10%
Probability of crossover	99%

For a large number of regions a single GA can struggle to converge properly. [47] The afore mentioned implements tried to optimize the buckling load directly instead of matching lamination parameter (LP). Such a problem is cannot be separated. Since here there are no dependencies between guides, the objective can be separated into a subproblem for each guide:

$$\begin{aligned}
 & \min_{\theta_{0,j}} (LP_{local,0} - LP(\theta_{0,j}))^2 \\
 & \min_{\theta_{1,j}} (LP_{local,1} - LP(\theta_{1,j}))^2 \\
 & \quad \vdots \\
 & \min_{\theta_{N,j}} (LP_{local,N} - LP(\theta_{N,j}))^2
 \end{aligned} \tag{7.6}$$

for a total of N guides. Parallelization is expected to give a significant performance and convergence improvement for large laminates.

7.3.2. Alternative Designs

The blending algorithm is deterministic and results may depend strongly on the input. As a consequence it is worthwhile to try different combination of stacking sequences to obtain more results and acquire the best design. Moreover, multiple blended results are produced that can be compared on manufacturing cost and structural performance. A pareto front of solutions can be generated.

New designs can be generated by randomly combining the best guides of every region. At this stage the design with the highest fitness is the design where for every region the guide with the highest fitness is selected. More designs are generated by changing the stacking sequence of some regions in the designs for guides with a slightly lower fitness.

To generate K unique designs for a laminate with N unique regions there are

$$M = \sqrt[N]{K} \tag{7.7}$$

unique guides per region needed, where M is rounded up to nearest integer. K combinations are selected where every combination has the same likelihood to be selected. As mentioned, the combination that contains the guides with the highest fitness for every region is always selected first.

7.3.3. Lamination Parameter Sensitivity

The stacking sequence retrieval rounds the continuous local stiffness to a discrete set of orientations that aims to match the lamination parameters. The result is a stacking sequence that is slightly higher or lower than the LP requested. In the case of the cylinder a large deviation in the buckling performance was observed depending on the bound chosen.

Lamination Parameter sensitivity information of the cylinder is obtained by applying a forward finite difference scheme. Similar to the gradient that is used in the optimizer, a small step of $h = 10^{-5}$ is used to sample the gradient from the surrogate model. Such a gradient provides information whether the load will increase or decrease when a specific bound is chosen.

A small penalty is added to the fitness function to lower the fitness on the more interesting side. The amount of skew and direction depends on the load derivative ($\frac{dL}{dLP}$) and is implemented per LP. The new fitness function becomes:

$$\left(LP_{local,i} - LP(\theta_{0,j}) \right)^2 - \kappa \left(LP(\theta_{0,j}) - LP_{local,i} \right) \frac{dL^*}{dLP} \quad (7.8)$$

where κ is weight value of the gradient term, and $\frac{dL^*}{dLP}$ the normalized derivative such that the maximum term is 1 kNm. Experimental tuning of the hyperparameter showed stable results for $\kappa = 0.05$.

Local Blending Using a Cellular Automaton

To create a manufacturable and cohesive design blending guidelines should be followed. The framework makes use of the relaxed blending guidelines described by van Campen[47]. A cellular automaton is used to transform the stacking sequences to a blended design via a set of rules. These rules transform a single cell based on its direct neighborhood in every iteration. After several iterations the algorithm converges to a blended design. The cellular automaton is based on the work of Van den Oord and consists of three stages that are visualized in Figure 8.1. [50]

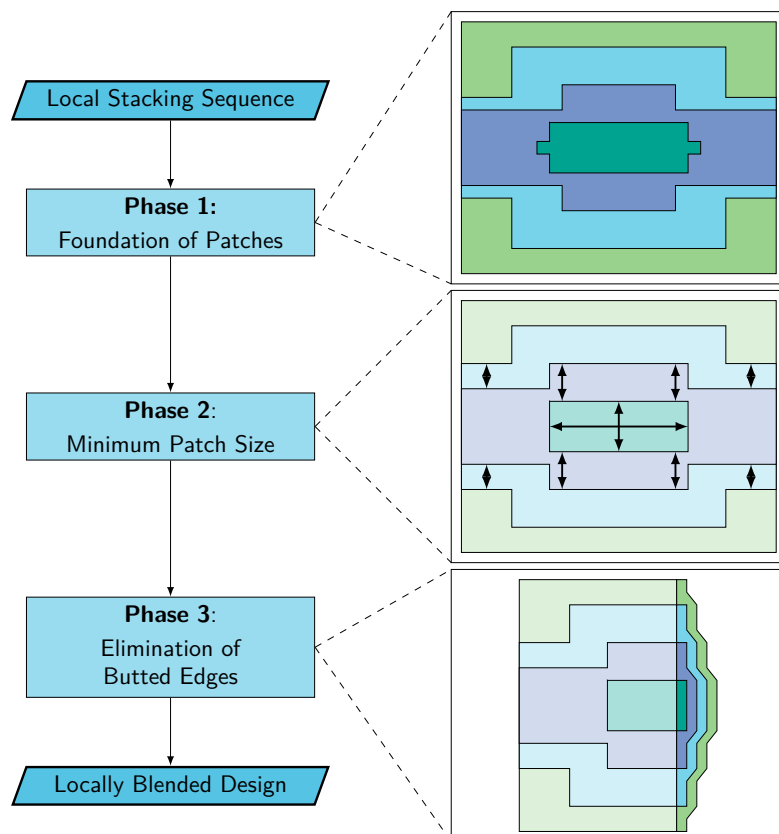


Figure 8.1: Overview of the three stages of the Cellular Automaton and visualization of the task performed during the stage. The first stage forms initial patches; the second makes the patches compliant with minimum cutting constraints; and the last removes butted edges to comply with blending guidelines.

The first stage forms patches by rotating layers toward each other such that they join. Since the output of the genetic algorithm can be relatively irregular small changes are preferred to create the preliminary patches. The rotating process is converged when a rotation would not yield improvement and would break the current patch. These first patch form an entity that can be further transformed.

Then, the minimum cutting size constraints are imposed. Automated machines and humans need a minimum patch size to work. Since the CA is capable of working with large designs, the detail level that can be reached is smaller than humans or machines can handle. As a consequence patches should have a minimum size. This stage was developed as part of the thesis, but not used due to problems that were caused by the last stage. Therefore the work is not considered here and has been added in Appendix B.

Lastly, the laminate is made compliant with the relaxed blending guideline by elimination of the butted edges. This process looks at every cell and makes a change that minimizes the number of butted edges present at the internal and exterior interfaces of that cell. The adoption of a multi-scale CA with multiple cells per region improved the convergence for designs with a high number of regions.

This chapter will discuss the first and third stage of the algorithm in more detail in section 8.2 and 8.3. The second stage is added as Appendix B because errors were introduced in the third step again. Therefore, although improvements with regard to the old method were made, it was not used to obtain a result. But first the setup of CA including the cell composition, neighbor layout and evolution schemes are discussed.

8.1. CA Setup

To simplify the problem and reduce complexity, only balanced symmetric laminates are considered. It is assumed that for every positive ply, a negative ply is located directly below. Furthermore, these plies stick together and follow the same path through the entire laminate. Every patch in the cellular automaton, as well as for the global interpretation, describes both the positive and negative ply. The fiber orientations of a patch are abstracted using a geno, $\tilde{\theta}$, between 1 and 7 to represent 0° to 90° in discrete steps of 14 degrees. These patches are mirrored on the symmetry plane of the laminate. The representation is summarized in Table 8.1.

Table 8.1: Patch representation number used to describe the fiber orientation.

Patch' Fiber Orientations	Representation($\tilde{\theta}$)
$[0_2]$	1
$[\pm 15]$	2
$[\pm 30]$	3
$[\pm 45]$	4
$[\pm 60]$	5
$[\pm 75]$	6
$[90_2]$	7

Cells are the basic entity of a cellular automaton that evolve according to an evolution scheme. To obtain a new state, every cell considers its neighborhood and uses a set of predefined rules to evolve. These rules should be designed such that a global optimum is reached. The layout of a cell, its neighborhood, and the evolution scheme affect the convergence and quality of the solution. Here first the cell layout and then the evolution scheme are discussed.

8.1.1. Cell and Neighborhood

Every cell is given a subset of all the information to compute its next state. The set of rules used to evolve the cell locally is chosen such that the global state shall converge to the optimum. It might be tempting to use a large neighborhood but this increases the complexity of the rules and thus is not feasible.

To increase the likelihood of convergence several scales can be used. Whereas the original CA called for a single region per cell, here multiple regions can evolve in the same cell. This allows the CA to solve more complex situation that were not possible before. The effect is most notable in the

butted edge removal process. Figure 8.2 shows the three different cell arrangements in cyan and the neighborhood of the cells in a shade of cyan.




Figure 8.2: Multi-scale cells that considers several regions and evolves them as one. Larger cells take more computational time to evolve but allows more complex situations to be solved.

The neighborhood of a cell determines how much information is given to the cell to evolve. The state of neighboring cells and the own state is passed to a deterministic function to produce a new state for the cell. The size of the neighborhood determines the amount of extra information that is given to a cell.

To keep the algorithm simple only cells that share an interface belong to the neighborhood of that cell. For the different cell configuration introduced above, the neighborhood is shaded in Figure 8.2. Depending on the scale, the neighborhood consists of 4 or 6 cells.

8.1.2. Evolution Scheme

The evolution scheme determines which cells evolve at the same time. Ideally, all the cells are evolved at the same to make the evolution process homogeneous. However, in practice this leads to two neighbors constantly adapting to each other and thereby flipping state indefinitely. To prevent this from happening two different evolution schemes have been developed. They are based on the following principles respectively:

1. *Cells evolving at the same time shall not be in each other neighborhood*
2. *Cells evolving at the same time shall not share neighbors*

The first scheme is equivalent to a checkerboard and is visualized in Figure 8.3. This scheme is suitable when a change in the cell's state only affects its own compliance or that of the interface with a neighbor. A good example for its application is the removal of butted edges. When the cell is updated only butted edges at the interface of cell can be removed. Other interfaces of that neighbor are not affected and therefore cells on the other side of the neighbor can evolve simultaneously. The exact algorithm is discussed later in section 8.3.

The second scheme, drawn in Figure 8.4, is useful when an evolution of the current cell also influences the compliance of an adjacent cell. For example, when a cell joins a neighbor this can make both cells compliant with minimum cutting constraints. Therefore, if neighbors on the left and on the right of a cell both change the result might be a patch that is wider than strictly necessary. This can be generalized as follows: The information the cell uses to decide how to evolve is also changing as a result of other changes in the same iteration. In the example the compliance of the neighbor has already




Figure 8.3: First order evolution scheme where cells with the same color evolve during the same iteration; Cells do not have each other as neighbor.




Figure 8.4: Second order evolution scheme where cells with the same color evolve during the same iteration; Cells do not share a neighbor.

been fixed by another cell. The purpose of this second scheme is to further decouple the information a cell uses from other cells that evolve at the same time. Similar to the first scheme, the starting order influences the end result and there is no good method to determine which cells to start with.

8.2. Foundation of Patches

At the start of the CA no notion of patches is present. Layers may have uncorrelated orientation that changes from region to region. Firstly, layers with similar orientation should be grouped and unified to form preliminary patches. This process is called foundation of patches and has been modified from the algorithm developed by Van den Oord[50]. The description here uses mathematical formulations for each of the steps to improve its performance using tensor algebra software packages.

During every evolution a cell changes the orientation of one layer towards that of its neighbor to eventually join that patch. Cells always consist of one region and evolution happens using the first order evolution, the checkerboard pattern. During an evolution only one layer of a region is allowed to change one orientation step. This forces gradual change throughout the entire laminate.

To determine which layer to change and to which orientation, a deviation distance ($DD_{i,j}$) is computed for every layer i and every neighbor j of the cell:

$$DD_{i,j} = \tilde{\theta}_{neighbor,i,j} - \tilde{\theta}_{region,i} \quad (8.1)$$

where $\theta_{neighbor,i,j}$ is the angle of the j th neighbor at level i , and $\theta_{region,i}$ the angle of the cell at level i . The deviation distances describes both the difference with the neighbor and in which direction to rotate, via the magnitude and sign respectively.

The deviation distance is grouped by the magnitude, by the level, and by the orientation direction. The occurrence is counted to form the deviation scheme (DS). The deviation scheme- summarize how many neighbors are located at a specific change:

$$DS_{i,k,l} = \sum_{j=1}^J \left[|DD_{i,j}| = k \quad \wedge \quad sgn(DD_{i,j}) = l \right] \quad (8.2)$$

where k denotes the magnitude range from 0 to 6, l the sign of the deviation distance(+ or -), and J the number of neighbors. The $sgn(\cdot)$ extracts the sign of the number and $[\cdot]$ converts a statement to 1 and 0¹. This score indicates the number of neighbors with a similar orientation if the layer at level i

¹This formerly known as the Iverson Bracket: <https://www.jsoftware.com/papers/APL.htm>, accessed September 10, 2020

would change k steps in direction l . For example, if $DS_{2,2,+} = 2$ then after rotating layer 2 by two steps in the positive direction, it would join with two neighbors.

It might be tempting to rotate towards the highest DS score, but that might sometimes cause the solution to diverge. For example, if the result of the positive rotation $DS_{i,k,+}$ is equal to that of the negative rotation $DS_{i,k,-}$ then there is no net gain in rotating towards either of these directions. This is caught by the neighbor count NC that scores the number of neighbors that would improve if level i is modified with k in the positive direction. When no improvement is made the $NC = 0$. This number is obtained through subtraction of deviation scheme:

$$NC_{i,k} = DS_{i,k,+} - DS_{i,k,-} \quad (8.3)$$

The neighbor count only gives the number of neighbors that is affected by a specific modification. To select the modification that improves the global solution the following set of rules is applied until one modification remains:

1. Modifications with a NC score of zero are not considered because those would not improve the situation.
2. The modification with the lowest k are selected first. A lower k indicates that less change is needed to improve the situation and thereby large changes are postponed. The end result is less change of the global solution.
3. If there are multiple modification at equidistance the one with the absolute highest score is chosen because the highest number of neighbors will benefit from that change.
4. If there are multiple modifications at equidistance with the same score, the one closest to the mid-plane is chosen. Modifications closer to the mid-plane have a smaller impact on the D matrix while the effect on the in-plane stiffness is similar.

After the modification is selected, the sign of the selected neighbor count marks the orientation direction. A single step is taken into that direction. For example, if $NC_{4,4} = -2$ is selected then layer 4 is rotated one orientation in the negative direction. Two neighbors that are 4 steps away are now only 3 steps away.

If all the neighbor counts are zero, the cell has converged. At this points there is no modification possible that would improve the cohesion further without sacrifice. Therefore no change is made in this cell. Ultimately when all cells have converged the CA moves to the next stage.

8.3. Elimination of Butted Edges

Butted edges occur at the interface between two regions and depend on the cell state. Since the cell's state influences all the interfaces at the same time this becomes a 3D problem where all the layers of the cell and neighbors should be taken into account together. As a consequence finding a good set of rules to transform the cell becomes impractical.

Instead, inspired by the work of Van den Oord[50], the new state of a cell is formulated as an optimization problem that tries to find the modification that minimizes the number of butted edges in the cell. A modification consists of a change of a single layer of every region in the cell into any other orientation. For example, if a cell consists of two regions then one layer in each of the two regions is allowed to change. This is different from the referenced work where only one region was considered. As a result this new implementation can eliminate more butted edges. [12]

The objective of the optimization is to minimize the number of butted edges present in the cell (this includes interfaces within the cell for multi region cells). Guide based butted edge detection algorithms, such as the one developed by van Campen, require a switch guide. [47] The absences of these guides in the CA algorithm does prevent the use of those algorithms. The scenario based algorithm used in the original CA was not general enough to detect incorrect covering of a butted edge. This was especially problematic when the two butted edges were at different levels in the laminates.

A new mathematical formulation based on a connectivity matrix was drafted. This formulation is set out in Appendix A and is able to identify blended edges. The number of butted edges is then determined by counting those that are not blended.

Since the optimization is an integer problem and the algorithm is not differentiable, an enumeration method is used to select the best modification. The following rules are followed to select the best modification:

1. The modification that minimizes the number of butted edges is chosen.
2. If multiple modifications result in the same number of butted edges, the modification that changes the least number of regions is chosen.
3. If multiple modifications are left, those that have the smallest absolute change are chosen. For example, if one modification changes two regions by 2 and 3 orientation steps and the other does so by 1 and 5 steps, then the first modification is chosen. The absolute sum of these modification is 5 and 6 respectively.
4. If still multiple modifications are left, those that make a modification closest to the mid-plane are preferred. In the case of multiple regions the absolute sum is used. If an option modifies level 2 and 3 that is preferred over an option that modifies level 1 and 5.

Similar to the foundation of patches phase, a cell has converged if none of the modification would decrease the number of butted edges present in the cell.

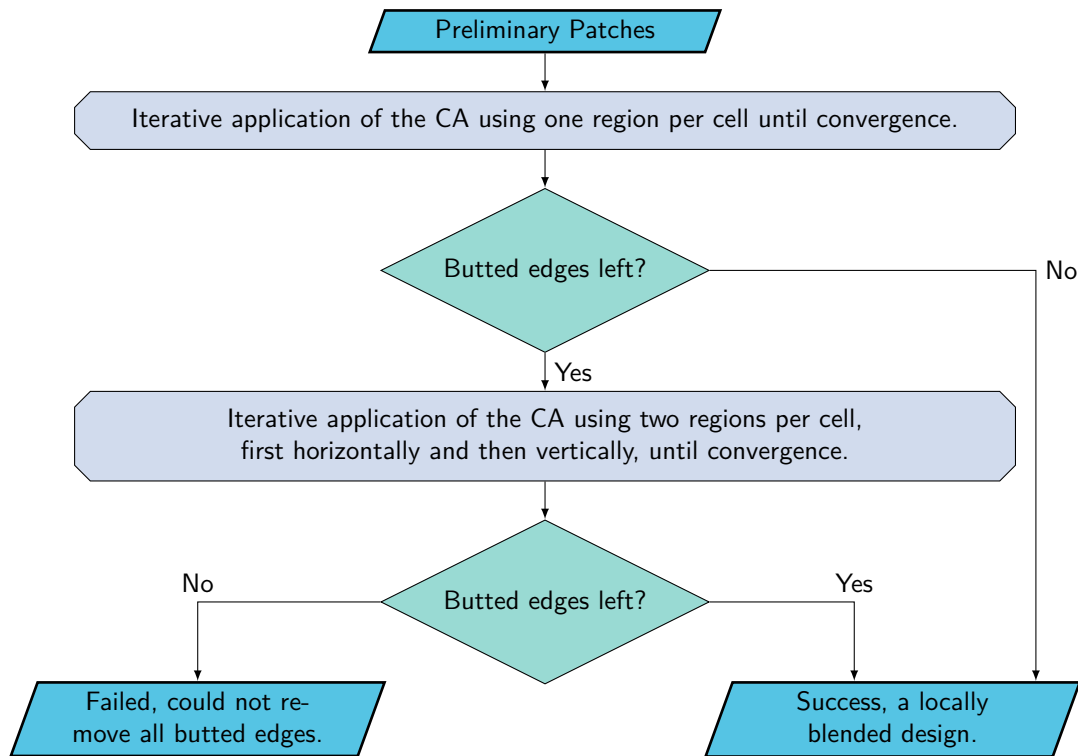


Figure 8.5: Strategy of application for scales involving multiple regions per cell. The computational cost rises exponentially with the number of regions per cell and therefore the number of higher order evolutions are minimized.

The elimination of butted edges phase first uses the simple scale where every cell contains only one region. If there are butted edges left after this scale, the second scale that considers two regions per cell is also used. This process is visualized in Figure 8.5.

The first scale that utilizes only one region per cell is run until convergence first to keep the computational cost down. The number of possible modifications is directly related to the computational cost because an enumeration method is used. For a single cell there are only 7 modifications per level possible if the representation in Table 8.1 is used. But the number of modifications grows exponentially as more regions are considered per cell:

$$\text{Number of Modifications} = (7 \cdot N)^K \quad (8.4)$$

where N is the number of levels in a region and K the number of regions per cell.

So even while after the second level their might still be butted edges unaccounted for, removing those in a higher scale cell layout is impractical. The number of modifications in cell containing three regions is close to 200,000 and is too high to enumerate. Therefore, as shown in Figure 8.5, designs are dropped from consideration further if they contain butted edges after the second scale.

Designs that make it past the third phase of the CA are presumed locally blended. Locally blended indicates that every stopped edge can be blended individually, but not necessary in the global solution. In some cases a layer should continue into two other layers of the adjacent region to blend all the layers. This leads to a conflict that should be solved at a global scale.

Also a locally blended design could require a patch interpretation where plies are interweaved. Manufacturing of interweaved plies is not possible with current automated processes and requires skilled technicians. Another step is needed to create a global interpretation of the laminate.

9

Global Patch Interpretation

A locally blended laminate cannot be produced directly, an interpretation is needed to extract the patches. A patch describes the shape, orientation, and location of a ply in the laminate. Ultimately a robot or human should place these patches on a mold to build up the laminate. The process of extracting the patch information from the locally blended design is called global interpretation (GI).

In previous work this process was done by hand and there did not exist a method to do it automatically. However, the global interpretation of large laminates is impractical by hand and hence an automated algorithm is developed. The algorithm tries to formulate the patches through connecting layers and proposes fixes for butted edges that could otherwise not be removed.

The objective of the global interpretation is to find a manufacturable laminate, which is defined in section 9.1. The algorithm is introduced in section 9.2 and then discussed in two distinct parts: The optimization of connected patches and the removal of butted edges, these are discussed in section 9.3 and 9.4 respectively. Lastly, the convergence through the ironing effect is reviewed in section 9.5.

9.1. Manufacturable Laminates

The framework shall generate fully blended laminate designs that are manufacturable using nowadays technologies. To this end all laminates generated should comply with relaxed general blending guideline and may not contain butted edges. Furthermore, after a patch has been placed in the mould it should not be necessary to move it.

A patch identification is used to connect layers together and determine the final shape of the patch. The patch identification also determines the order in which the patches are assembled. The number represents the order in which patches should be stacked where lower ones are assembled first. The number zero is reserved for an empty spot in the laminate. For example, the laminate visualized in Figure 9.1, is identified by the stacking table shown in Table 9.1. The laminate is build up from the top of the pictures to the symmetry plane. Likewise, the stacking table is read from the top down to identify the next ply to stack.

If a lower number would be followed by a higher number that ply had to be placed below the already stacked one. This would require the already stacked ply to be lifted to give space for the new one, so called weaving of plies. Current automated manufacturing robots cannot weave plies, hence for the table to be manufactured, each column should be strictly ascending.

The complete state of the laminate is now defined by two variables for each layer: its orientation and its ply assignment. Together they describe the ply orientation and the stacking order respectively. The patch assignment should assure that there are no butted edges left. This can be checked with the same algorithm as used in the elimination of the butted edges and is described in Appendix A




Figure 9.1: A blended design consists of several patches, constant in orientation, that span multiple region on different levels.

Table 9.1: The assigned patch ids for the laminate in Figure 9.1. Layers with the same number correspond to the same patch. Patches are stacked from the top of the table, where the mould is, downwards. A manufacturable laminate has a monotonically increasing patch assignment in every column.

	I	II	III	IV
Level 1	1	0	0	0
Level 2	2	1	0	0
Level 3	3	2	1	3
Level 4	4	3	3	5

It should be noted that this chapter visualizes laminates in 2D because the representation of a 3D laminate is nontrivial. Visualizing a 2D order inconsistency is impossible using the flat figures here and only occurs in 2D continuous laminates where plies wrap around to the other side of the laminate. One such example is a laminate around a cylinder. An order inconsistency in 3D would be visualized by a crossing in a 2D visualization.

9.2. Algorithm Overview

The global interpretation tries to connect layers with the goal of finding a set of connections that produces no butted edges and results in a manufacturable laminate. The input is a locally blended design where for every layer in a region the orientations are defined. A dual approach is used where a connection search identifies the problematic area in the laminate and a fixing algorithm proposes new designs that aim to resolve the conflicts. After solving all of the problems in the laminate a blended configuration is obtained.

An overview of the algorithm is visualized in Figure 9.2

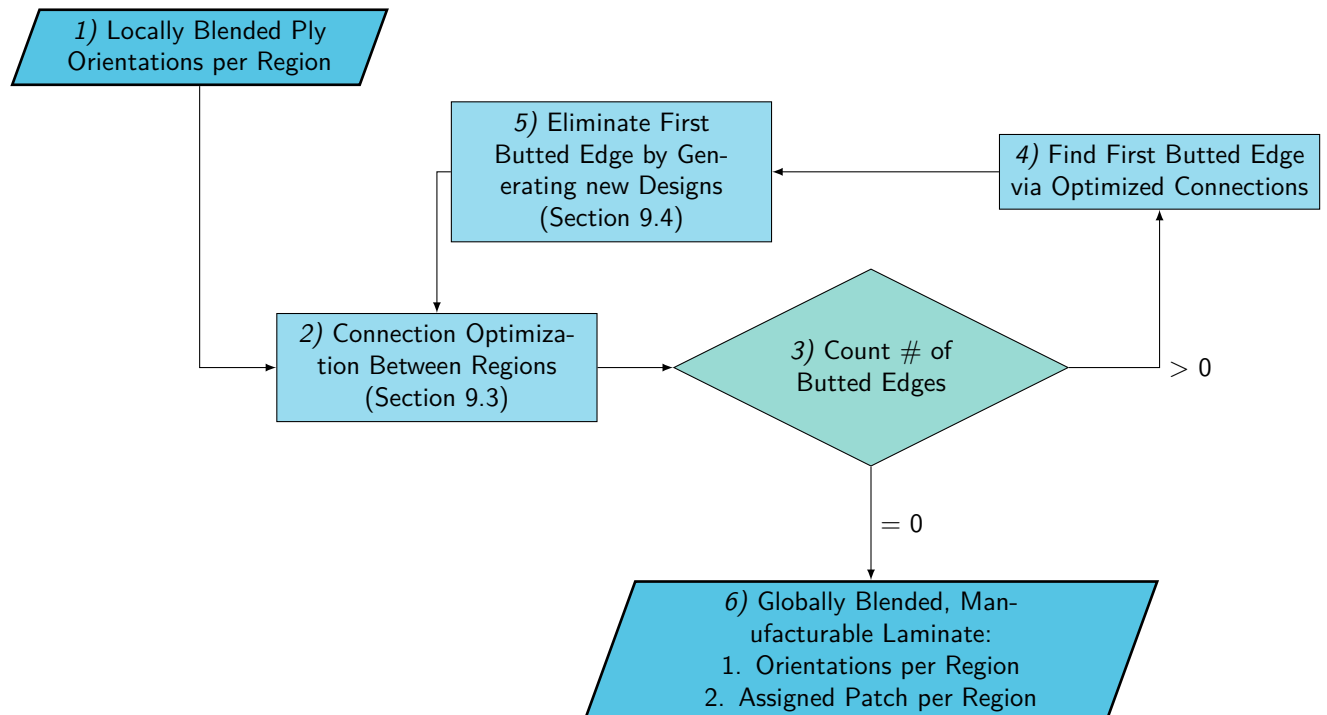


Figure 9.2: Simplified overview of the global interpretation process.

Here the different process blocks perform the following tasks:

1. The input to the algorithm is a locally blended laminate where each stopped edge can be blended individually by covering connections. Nevertheless, there is no guarantee that all stopped edges at an interface can be blended simultaneously or that the resulting patches can be stacked monotonically.
2. The patches of a design are optimized using a novel connection search that minimizes the number of butted edges and pushes them to the mid-plane. The search first identifies all the connections that are mandatory to make and then uses a depth first search to resolve the remaining connections. A connection is allowed when the orientation of the connected layers is the same and the connection does not violate patch ordering constraints. The output of this process is patch assignment for the laminate and a list with butted edges that could not be removed. The search is described in section 9.3.
3. The list of butted edges left is counted to determine the number of butted edges left in the current design.
4. If there are butted edges left the first butted edge is selected for repair.
5. The layers and neighboring regions of the selected edge are modified by changing their orientation according to repair strategies. Multiple repair strategies lead to multiple new designs to be reviewed. Repair strategies include continuing the ply in the neighboring region and covering it from above and below. The strategies aim to remove the butted edge, but there is no guarantee that it is removed globally. These strategies are discussed in section 9.4. New designs are queued for reexamination by the connection search. It could happen that a strategy was not able to fix a butted edges as will be further discussed in section 9.5.
6. If the connections search produces a result with no butted edges left a globally blended solution is found. The global result contains a fiber orientation per layer and patch identification. Patch assignments are designed such that it complies with blending guidelines and is monotonically increasing in every region. This last fact allows the patches to be manufactured using modern techniques.

9.3. Connection Search

The connection search forms the basis of the global interpretation by determining the layout of the patches. Instead of directly assigning every layer in a region to a patch, it designs the relations between layers. The patch assignment is then derived by tracing the relations.

Since the connection search lays at the basis of the global interpretation, its objective is to find a connection design that resolves all the butted edges. Such a solution might not exist at the start and therefore the secondary objective is to find a solution where the first butted edge is as close to the mid-plane as possible. The first butted edges is the closest to the outside of the laminate because a laminate is assembled from outside in.

First of all, bringing the butted edge to the mid-plane is beneficiary because the fixing strategies then operate in a region of the laminate closer to the mid-plane and hence have less impact on the flexural stiffness. Secondly, the butted edges are pushed closer to the mid-plane in every iteration, effectively ironing out the defects. This ironing effect is important for convergence and is further discussed in section 9.5.

9.3.1. Nodes and Connections

To discuss how the algorithm finds the optimum relations between the layers, an introduction into the nodes representation and derived views is needed. By representing layers in a laminate as nodes and layer relations as edges, graph theory can be utilized to find patch assignments and stacking order compliance.

Every layer in a region is represented using a node in the graph where the identifier first increments inside the region before considering the adjacent region. Every node tracks the layer it belongs to and holds the same orientation information. The nodes for the laminate of Figure 9.1 are shown in Figure 9.3 as circles where the background color represents an imaginary orientation.

The stacking order of every region is represented using directed edges between nodes in the graph.




Figure 9.3: Layers in regions from Figure 9.1 are represented as nodes in a graph. Undirected edges (thick black lines) represent connections between layers and mark the formation of patches. Directed edges (cyan arrows) show the required stacking order per region. The direction is from top to the mid-plane (bottom of the picture), because outer plies are stacked first.

The edge originates by the node that is stacked first and ends at the node that is stacked later. For manufacturable laminate every next layer should be stacked after the previous, going from the outside to the mid-plane of the laminate. This is represented in the graph by directed edges going from the outer nodes to the ones below them for every region individually. Directed edges are shown with cyan arrows in Figure 9.3.

Connections in the laminate are made using an undirected edge in the graph, visualized using a black line in Figure 9.3. These undirected edges connect the layers that together form a patch. The connection search designs these edges through the connection matrix (CM) with a 1 for edges that do exist and a 0 otherwise. The row and column represent the start and end of the edge, such that $CM_{ij} = 1$ represents an edge from node i to node j . The CM is symmetric by virtue of the undirected edges.

9.3.2. Grouping of Nodes in a Patch

Nodes are grouped in a patch through their undirected edges. The patch assignment PA assigns every node to a patch through a unique number for that patch. This patch assignment can be used to examine which edges are blended and which should be resolved via the butted algorithm set out in Appendix A.

The PA is constructed from the connection matrix through a depth first search that traverses all the edges of a node it visits. The algorithm is shown in Figure 9.4.

The algorithm first selects a starting node, which can be any of the nodes, and assigns it to the first patch. It then traverses through all of the connected nodes and assigns them to the same patch if they were not assigned yet. When all of the nodes that can be reached have been assigned, it continues to another unassigned node and assigns it to a new patch. Again all the connected nodes are traversed and the algorithm continues until all the nodes have been assigned. The result is shown in Figure 9.5

Depth first search algorithms are $O(n)$ which means that their computational time increases linearly with the number of nodes. Still, shortcuts can be taken during the connection search to further decrease the computation needed. When the connection matrix is updated and a new edge is created only the updated patch should be computed. Therefore, instead of traversing all the nodes, only the nodes that are involved in the connection and those that can be reached have to be updated. This is accomplished by selecting the nodes involved in the connection as the 'starting node' for the algorithm. Similarly, when the connection is removed the algorithm is started once from each of the nodes involved in the old connection. As a result only a limited number of nodes is updated.

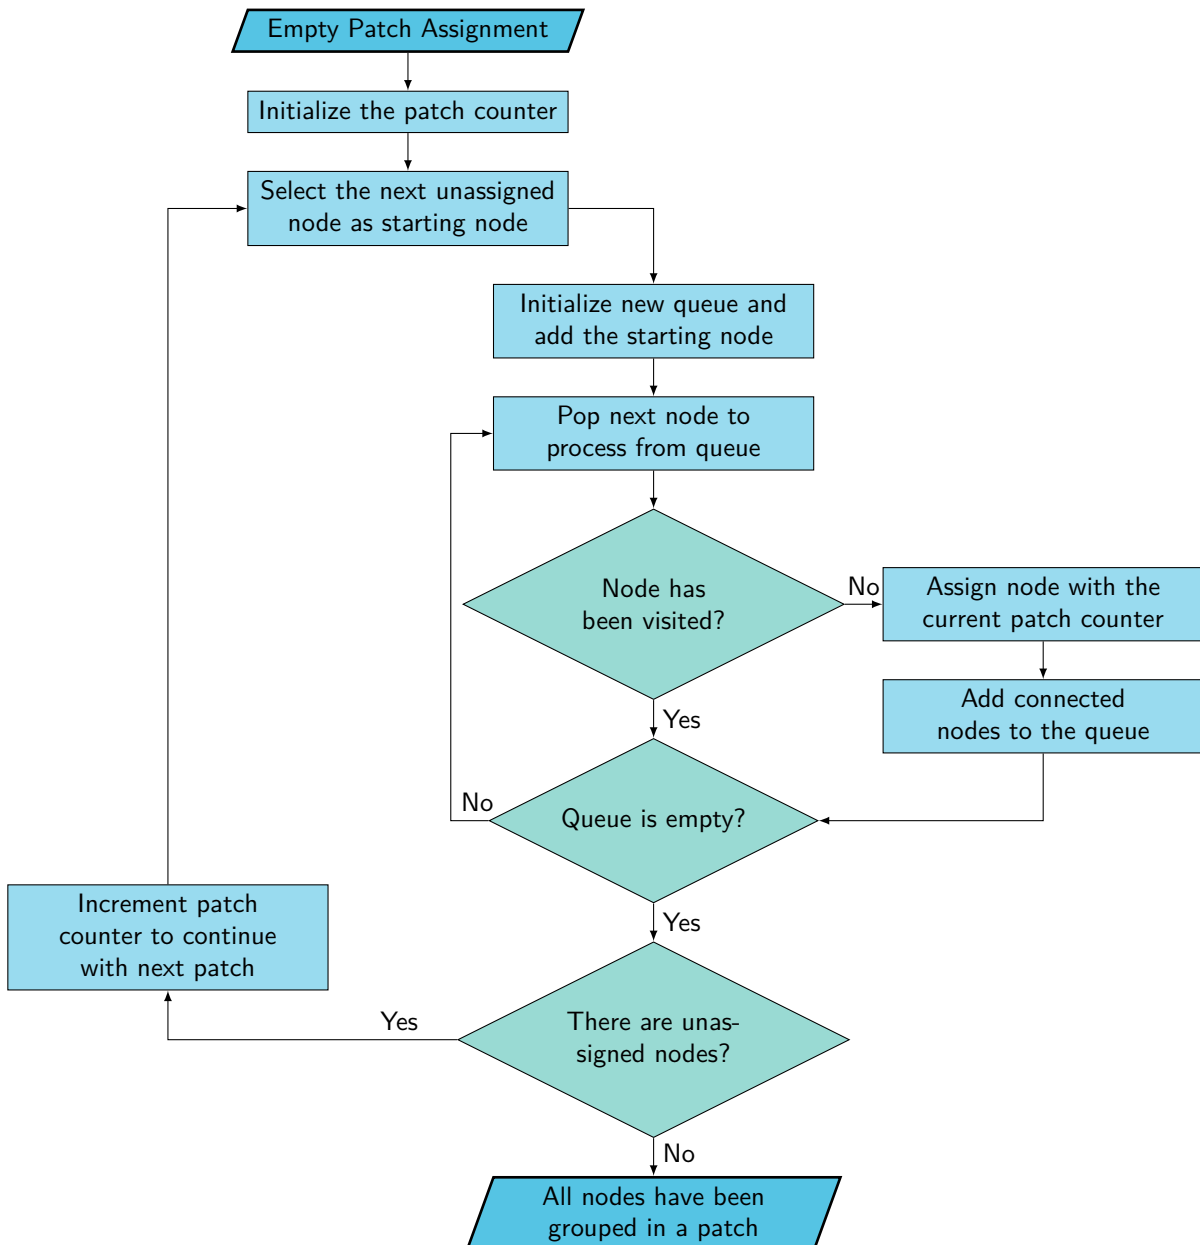


Figure 9.4: Depth First Search of the patch assignment. All the connected nodes are traversed and assigned to the same patch until all the nodes are visited at least once.

9.3.3. Stacking Order of Patches

Directed edges from node to node are not enough to detect inconsistencies in stacking order. Instead directed edges of all the nodes in a patch should be grouped together to get a patch order. To do so, all the edges ending and originating from and at a any of the nodes in a patch are combined in a new graph as shown in Figure 9.5.

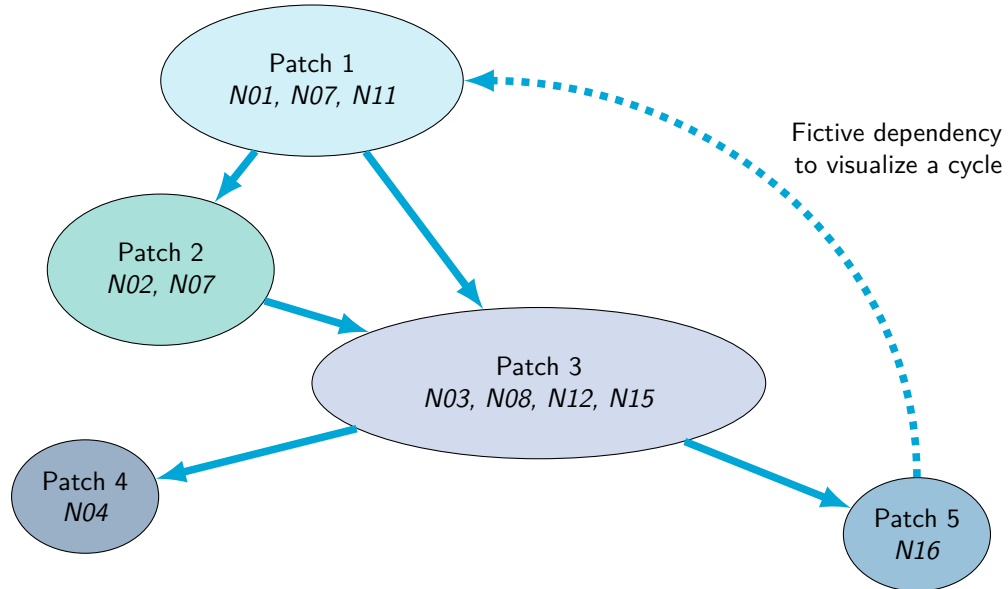


Figure 9.5: Directed edges are grouped by the patch assignment to create an order dependency graph for the example shown in Figure 9.3. The nodes that are grouped in a patch are added in italic to the patch. Directed edges originate at a patch that is stacked before a the patch it points to. A fictive cycle for illustrative purposes is created by making patch 5 depend on patch 1.

The directed edges in the graph shown in Figure 9.5 define the stacking dependency of the laminate. The origin of a edge has to be stacked before the destination. However, if a series of connections is traversed and it continues through a patch that it has already been passed through, a cycle exists. A cycle in the order dependency means that a patch can only be placed after a patch that ultimately depends on the initial patch again. Interweaving is not possible using current manufacturing method and leads to a laminate that cannot be easily manufactured.

One such an example of a cycle is created if the fictive connections is considered; The cycle originates at patch 1 to patch 3 and patch 5, back to patch 1 again. A second cycle similar cycle exists that also contains patch 2.

Cycle detection is well studied problem and can be solved in $O(n)$ complexity. However, the problem at hand tries to avoid these cycles altogether and therefore their occurrence should be detected before the new connection has been made. Instead of detecting a cycle, a matrix is constructed that prevents cycles from being created. This order matrix (OM) describes the relative order of the patches. $OM_{i,j} = 1$ indicates that patch i depends on j and should be stacked later; 0 yields no dependency. The OM for the graph in Figure 9.5 without the fictive connection is

	$p1$	$p2$	$p3$	$p4$	$p5$
$p1$	0	0	0	0	0
$p2$	1	0	0	0	0
$p3$	1	1	0	0	0
$p4$	1	1	1	0	0
$p5$	1	1	1	0	0

and with the fictive connection:

	<i>p1</i>	<i>p2</i>	<i>p3</i>	<i>p4</i>	<i>p5</i>
<i>p1</i>	1	1	1	0	1
<i>p2</i>	1	1	1	0	1
<i>p3</i>	1	1	1	0	1
<i>p4</i>	1	1	1	1	1
<i>p5</i>	1	1	1	0	1

A cycle in such a matrix is then represented by 1 on the main diagonal (highlighted in the second matrix).

To verify whether a new connection between nodes, here after referred to as *a* and *b*, is allowed, simple assertions involving the *OM* are enough. If the patch assigned to nodes *a* and *b* is $A = PA_a$ and $B = PA_b$ respectively, the following two statements suffice:

1. There may not exist a relative order dependency between the patches:

$$OM_{AB} = 0 \quad \wedge \quad OM_{BA} = 0$$

The existence of such a relation implies order dependence some where else in the laminate and could cause the patch to be stacked on top of itself. This is also referred to as an Escher staircase situation, named after the continuous, impossible staircase in Escher's drawings.

2. Neither of the patches should be stacked after a patch that the other should be stacked before of:

$$OM_{Ai} \cdot OM_{iB} == \mathbf{0} \quad \wedge \quad OM_{Bi} \cdot OM_{iA} == \mathbf{0}, \quad i \in \{1, 2, \dots, N\}$$

Here the intrinsic property that if *A* is above *B*, then *B* is below *A* is used through the transpose of the order matrix.

These checks combined overcome the need for back tracing to determine if the resulting patch would violate order constraints. Additionally, a statement is required to verify that the two nodes contain the same fiber orientation.

Construction of the order matrix happens through an iterative algorithm that traverses patches. The flow diagram is shown in Figure 9.6. First a patch located at the top of the laminate is selected as a starting point. For every patch processed the nodes and patches that are located directly above are determined by reversing the direct edges from node to node shown in Figure 9.3.

In order to continue processing of the selected patch each of the patches above should already been processed. This prevents patches from being processed too early when the complete situation above the current analyzed patch is not determined yet. If any of the patches has not been processed they are added to the queue and processing of the current patch is postponed by also adding it to the end of the queue.

If all of the patches that lay atop are processed, all of the patches that the current patch depends on are aggregated. Aggregation is done by combining the dependencies of the patches above and adding the patches directly above themselves. The result are all the patches that lay directly above plus all the patches that are located above those. This dependency is set in the order matrix and the next patch is popped from the queue. The process is complete when all patches have been processed.

Similar to the patch assignment process, only patches that have actually changed need updating. After setting or removing a connection between nodes, only the patches that can be reached from those nodes need to be updated. This reduces the amount of iterations needed to formulate the new order matrix.

9.3.4. Depth First Search

The connection search optimizes the connections between layers to push the first butted edges as far as possible to the mid-plane of the laminate, while accounting for order constraints. It uses the

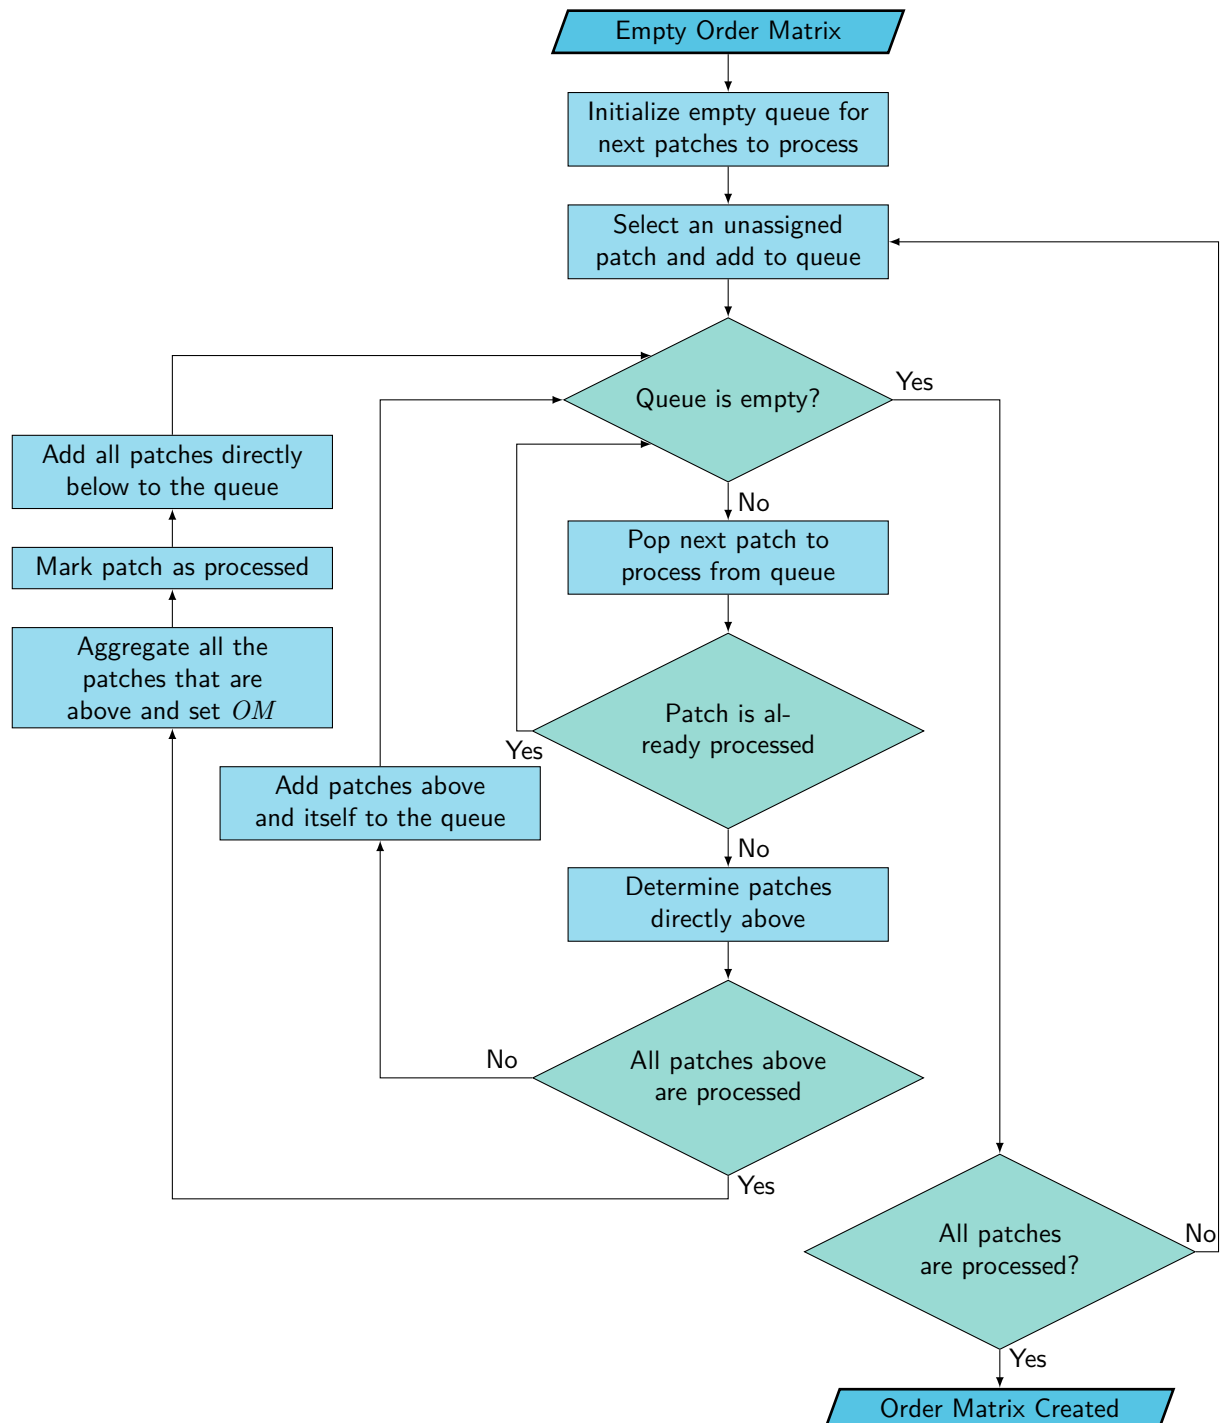


Figure 9.6: Order matrix construction is an iterative process that traverse through the patches which are fully defined and queues those that still depend on others.

previously described algorithms to update the patch assignment and order matrix continuously after applying or removing a connection. A flow diagram of the algorithm is represented in Figure 9.7.

The search starts with the initialization of any empty connection matrix and construction of the patch assignment and order matrix. At this time all the edges in the laminate are still butted because no connections have been made. Butted edges are detected based on the patch assignment through the algorithm in Appendix A.

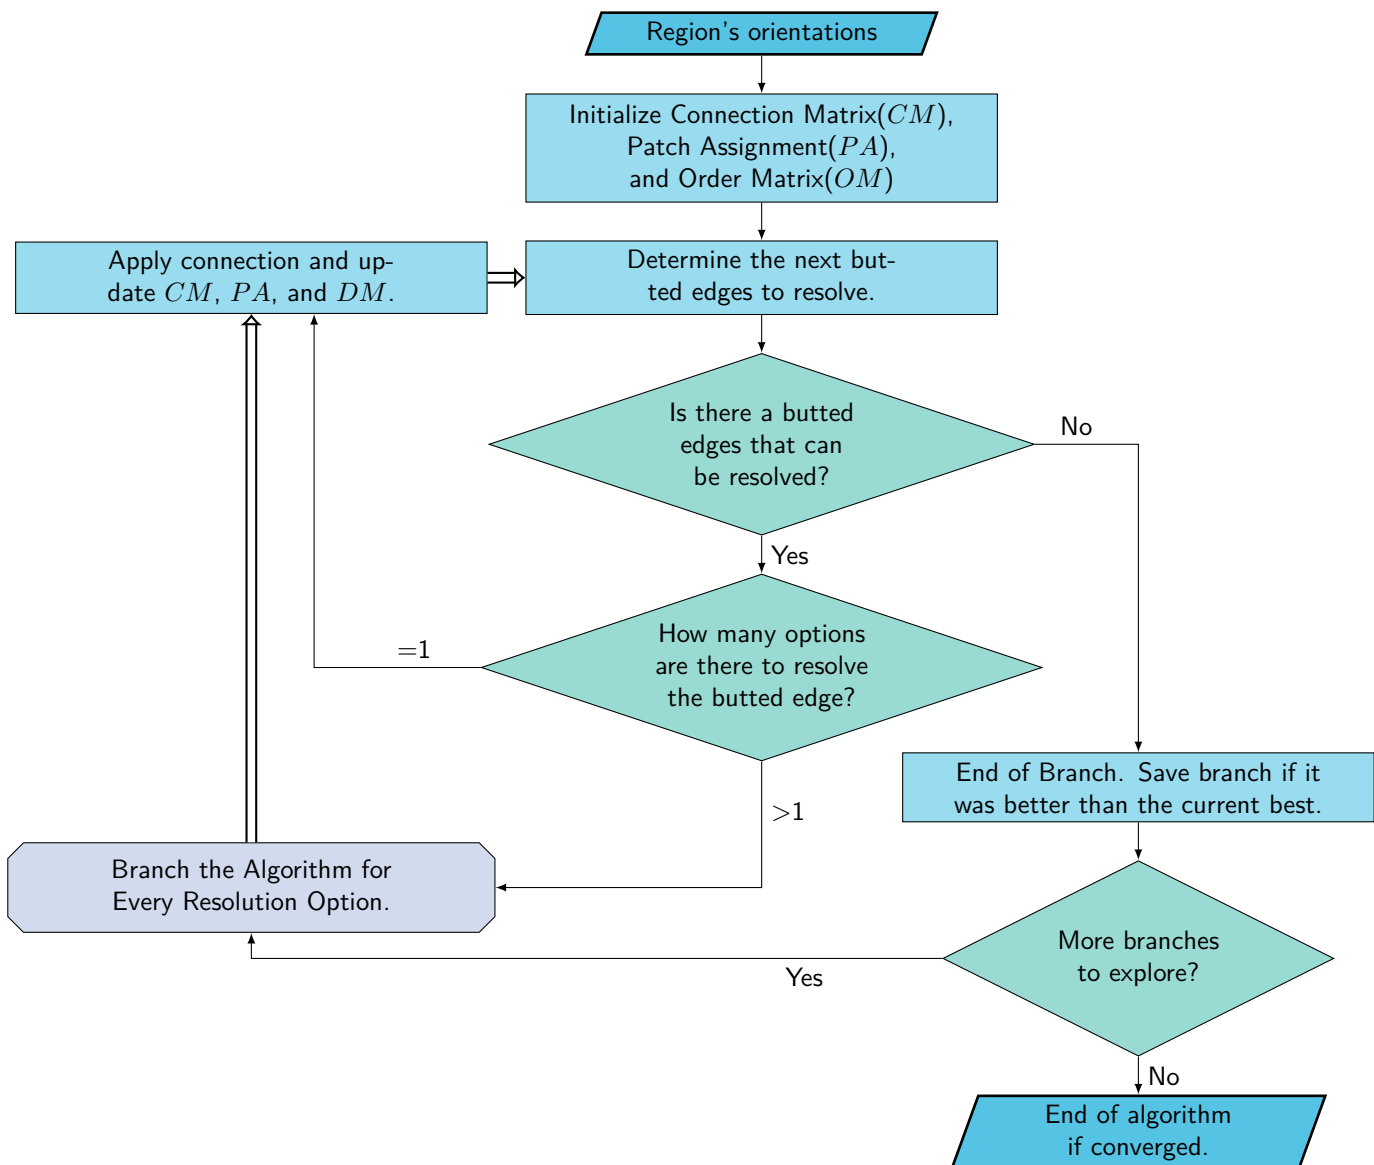


Figure 9.7: Flow diagram of the connection search algorithm. Butted edges are removed one by one by making connections between layers.

Every iteration the algorithm searches for a butted edge that can be removed and selects the one with the least number of options. If there is only one option to remove a butted edge that connection is called a mandatory connection. After checking if this connection is valid in terms of orientation similarity between the nodes and order consistency, it is applied and the next iteration is started.

In the case that there exists no feasible connections to resolve a butted edge, the butted edge is marked as definitely butted. Although the edge is locally blended, connecting the two nodes with the same orientation might not be possible in the global solution because it is already connected to another node or because the result would violate the stacking order constraints.

After all the mandatory connections have been made, several options exist to resolve the next

butted edge. At this stage it is unclear which of these options leads to the best solution and hence they are all explored. The algorithm selects the first option and saves the others for later exploration. After application of the first option it continues to the next iteration and searches for new mandatory connections.

After all the butted edges are removed or when there are no more options to remove those still present, the search has come to the end of a branch. At the end of a branch a solution is obtained with a location for the first butted edge. If this location is higher, meaning closer to the mid-plane, than the current best then this branch has yielded an improved result. After a branch has been explored, the algorithm continues from the last branching point in the next branch.

In the case the first butted edge is definitely worse than the current best, a branch is terminated early. Continuation could never yield a better result and is considered a waste of computational effort. Similarly, to restrict the computational effort the number of branches is limited to 10,000 after the first solution has been found.

9.4. Fixing Strategy

Butted edges detected by the connection search need to be resolved to achieve a global solution. Although the algorithm finds all the butted edges, only one butted edge is resolved at a time. This makes it possible to check whether the proposed fix has worked. It is not trivial to compute at forehand the effect of a change and therefore the proposed fixes are added to a queue for evaluation in the connection search.

The fixing strategies aim to resolve the first butted edge by altering the surrounding layer's orientations. These strategies do not incorporate any patch data and only act in the 2D plane surrounding the butted edge. 3D effects of the strategies are not taken into account because they are complex to compute.

Two strategies exist that are visualized in Figure 9.8 and explained below. Each of these strategies is applied to the detected butted edges which leads to multiple branches that each consider a new design. Branches are terminated when no valid solution is found using the connection search.

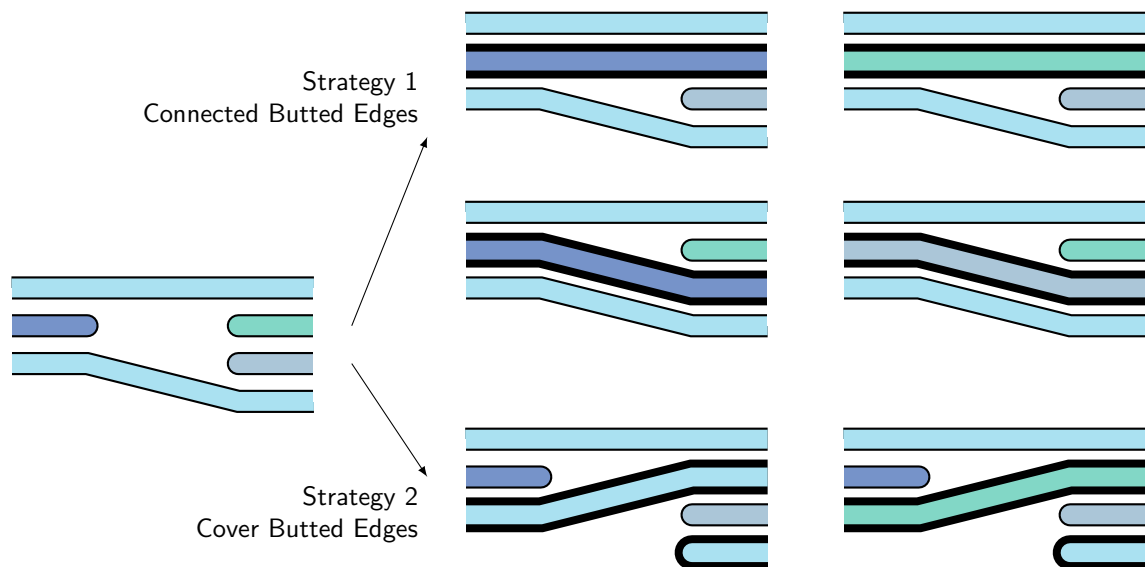


Figure 9.8: Two fixing strategies result in several new designs that cause the global interpretation algorithm to branch. The first strategy tries to combine butted edges and the second tries to cover them.

Butted edges are removed by giving the layers that are butted an equal orientation. This allows a continuation of the butted edges. In the case that one edge of one region is butted against two of the neighbor region, multiple possible resolution options exist. Each of these options is added to the queue for evaluation which causes the global interpretation to branch.

Instead of directly connecting the butted edges, the second strategy connects one of the butted edges to a layer below of the other. This covers the stopped edges of the other region and removes

the butted edges. The layer is already covered from above since all the layers above the first butted edges adhere to the blending guidelines by definition of the first butted edge.

The application of multiple fixing strategies leads to branching of the global interpretation. Several new designs are proposed and added to a processing queue for evaluation by the connection search. At this stage it is uncertain which strategy actually resolves the conflict or has introduced other defects.

9.5. Convergence

The fixing strategy creates numerous new designs to fix a butted edge and queues them for evaluation in the connection search. New designs only pass evaluation if they actually removed the butted edge and pushed the first butted edge closer to the mid-plane. If the butted edge remains present or if a new butted edge was introduced further from the mid-plane, the branch is terminated and the proposed fix is ignored.

Through termination of the design branches that do not improve the first butted edge, the algorithm convergences by pushing all the butted edges to the mid-plane. This ironing effect also minimizes the effect on the buckling performance because the change to the region's orientations are made at the mid-plane of the laminate. These have less impact on the flexural stiffness than those closer to the outside.

10

Qualitative Cost Comparison

The global interpretation algorithm will come up with several designs. These designs differ in load and simply picking the one with the highest structural performance is not always sufficient. Often an engineer has to take into account multiple criteria and trade cost for structural performance. The highest performing designs are generally impossible or too costly to manufacture and the engineer has to trade performance for a more cost effective design. This is preferably done by a measure of the complexity of the design which indicates the manufacturing cost including risk, labour, and material[22].

It is hard to quantify the manufacturing cost of a design because this depends on the production volume, the material, and the technique used. However, by quantifying those properties of design that have a negative qualitative influence on the cost, a trade-off between solutions of the global interpretation can be made. [42] Such an approach based on several metrics resembles a parametric model where drivers predicted the influence on the price. Here three drivers are discussed in terms of their qualitative influence on the cost and a proposed way to compute those metrics.

10.1. Number of Patches

The first measure for the cost of a design is the number of individual patches. A design with a higher number of patches takes more time to manufacture because either more patches have to be placed by hand or an automated machine needs more movements. Both consequences drive up the manufacturing cost as well as the manufacturing risk. Computation of the metric is straight forward and involves counting the individual patches.

10.2. Patch Shape Efficiency

There is no trivial way to measure the efficiency of a shape because this depends on the purpose of the shape. However, a common used metric in natural sciences is the area to volume ratio. The optimum in 3D, a sphere, states how a certain volume can be distributed most optimal such that the surface area is smallest. Several natural phenomena obey this shapes, including weightless water droplets for example. It can be seen as a measure for the optimal shape to fit a certain volume in.

A 2D equivalent is the ratio between the circumference and enclosed area of a shape. When two shapes have the same area, the one with the large circumference will be more irregular. An irregular shape is more expensive to cut because the cut lengths are larger. Moreover, complex shapes are harder to pack when creating a cutting pattern and cause more carbon fiber to be wasted. Similarly, automated fiber placement using tows require more cuts and movements as well. An irregular shape is undesirable and has higher associated costs than simple shapes.

To compare designs that use different region layouts a metric is needed that is independent of the region shape. Additionally, it should not change when a patch is scaled since these have the same complexity according to the definition above. The following expression, based on the aspect ratio of




Figure 10.1: Theoretical patch shapes that can be designed by the algorithm, ranging from simple on the left to more complex on the right. Patches are normalized such that regions have dimensions of 1 by 1.

an airplane's wing, is used for the normalized aspect ratio(NAR):

$$\text{NAR} = \frac{(s^*/4)^2}{A^*} \quad (10.1)$$

where A^* is the normalized shape area such that every region has area 1 and s^* is the normalized circumference such that all edges of a region have length 1. Essentially the area and circumference of a shape are measured after the shape is projected on an unit grid. A factor of four is chosen such that a patch of a single region will have $\text{NAR} = 1$

The normalized aspect ratio is calculated for the shapes in Figure 10.1 and tabulated in Table 10.1 together with the complexity that is introduced in the next section:

Table 10.1: The NAR for various shapes drawn in Figure 10.1 to show the property of the efficiency factor.

	A	B	C	D	E
s^*	4.0	8.0	6.0	8.0	12.0
A^*	1.0	4.0	2.0	3.0	5.0
NAR	1.00	1.00	1.13	1.33	1.80
Complexity, C	4	4	4	6	12

The table shows that the normalized aspect ratio does not change when a shape is scaled as is the case for shape *A* and *B*. It also shows that the factor goes up for irregular shapes, like *D* and *E*. It should be noted that although the equation for the factor closely resembles the aspect ratio equation the factor is not. Clearly, the aspect ratio of shape *D* is 2 which is different from the NAR.

10.3. Patch Complexity

Another way to measure the irregularity of patches involves counting the number of corners. A shape with a high number is considered irregular and more costly. Machines need to stop and change direction adding to the cost of manufacturing. An inefficient shape is caused by a high number of corners.

For a grid based design with closed shapes two styles of corners can be considered; interior and exterior corners. Interior corners close in on the shape making it smaller and have an angle of 90° . The rectangle presented in shape *A* of Figure 10.1 has four interior corners. Exterior corners are the opposite and open up a shape by having an angle of 270° .

Detection of interior corner is a trivial process after the edges of a patch are known. In every region two adjoint edges count as a single corners. These can then be summed to find the number of interior corners in a patch. Since the sum of the corner's angles in a shape should be equal to $180(c - 2)$, where c is the total number of corners, there should be one exterior corner for every interior corner above four. This provides an easy relation to determine the total number of corners, C , in a patch based only on the number of interior corners c_{interior} :

$$C = 4 + 2(c_{\text{interior}} - 4) \quad (10.2)$$

Equation 10.2 and 10.1 can be used during post processing to determine the complexity and irregularity of the patches used in the design.

A Blended Outcome

In the previous part a complex framework largely based on other research was described. This method is supposed to be the basis for framework that can be used to successfully optimize large scale aerospace structure that incorporate variable stiffness through straight fiber variable stiffness laminates.

To apply a new method in practice, careful validation should be performed. Several benchmark problems were introduced in the literature study that provide a baseline for the several stage of the algorithms. These baselines can be used to validate the output of the framework. Additionally, results can be obtained via different methods described in research. The validation of results against those of others or with results obtained through different method is valuable feedback on the performance of the framework. Such validations should be performed before adoption of the framework in other study cases. Validations are performed in chapter 11.

After validation of the methodology the other two benchmarks, the clamped plate and cylinder in bending, are optimized. The clamped plate was originally design using SFVS laminates, but violated the relaxed blending guideline at some interfaces. It achieved a buckling load of 132 kN using 5x5 regions. The cylinder in bending buckles in the compressive zone and was optimized using fiber steering. That design reach an improvement of over 27% when material failure was neglected. The framework is put to the test with these benchmarks and the results are summarized in chapter 12.

Framework Validation

Proper validation is needed to check convergence of the individual algorithms. Benchmarks from literature were selected to compare the results obtained using this framework before new designs for the cylinder and the clamped plate were generated. Validation of the tools used to create is needed to ensure the quality of the results.

This chapter first discusses the local optimization techniques used. The stiffness modeling using general shell elements and the convergence of the local optimization method are discussed in section 11.1. Then, in section 11.2, the differences between a single and a parallel application of the genetic algorithm to obtain the stacking sequences is discussed.

11.1. Local Stiffness using Lamination Parameters

The first step in the optimization framework optimizes the local stiffness distribution to maximize the buckling performance. It interfaces directly with Abaqus to optimize the lamination parameters per region. The general element stiffness modeling technique is validated first and then the lamination parameter results are compared to those obtained by IJsselmuiden et al.[20] for the simple supported benchmark problem.

11.1.1. Comparison of Modelling Techniques

Two different modelling techniques are used in conjunction to design the plate. It is expected that, after applying the correction for the transverse shear stiffness, the modelling techniques should give equivalent results. Five layups with equal laminate thickness have been simulated using the direct stacking sequence and the lamination parameters approach.

For each stacking sequence several plates with a varying thickness over size ratio are modelled to see the effects of the element slenderness on the modelling techniques. Element slenderness is here defined as the average in-plane size over the thickness. The influence of the transverse shear correction is expected to be more severe in thicker elements. Since the correction is idealized, a difference in buckling load is expected for small slenderness ratios.

The stacking sequences and material properties are provided in Table 11.1 and Table 11.2 respectively.

Table 11.1: Layups considered to verify general stiffness modelling technique.

	Stacking Sequence
Layup 1	$[(0/90)_3]_s$
Layup 2	$[(90/0)_3]_s$
Layup 3	$[(45/-45)_3]_s$
Layup 4	$[(90)_2/(45/-45)_2]_s$
Layup 5	$[(0)_2/(45/-45)_2]_s$

Table 11.2: Material Properties of AS4/8552.

Property	Symbol	Value
Longitudinal Modulus	E_1	141GPa
Transverse Modulus	E_2	9.03GPa
Poisson Ratio	ν	0.32
Shear Modulus 1-2	G_{12}	4.27GPa
Shear Modulus 1-3	G_{13}	4.7GPa
Shear Modulus 2-3	G_{23}	2.9GPa

A mesh convergence study of all plate resulted in a stable solution for 60 elements in the both direction regardless of the size. Absolute buckling loads shown in Figure 11.1 are for an element slenderness ratio of 1.33 which is considered an extreme case. The bar chart also incorporates the lowest theoretical buckling load that is obtained using Equation 6.1. The analytical approximation of the buckling load




Figure 11.1: Comparison of element definition techniques for prediction of the buckling load of a plate in uniaxial compression.

overestimates the buckling for every stacking sequence. Especially the laminates that contain a large number of 45 degrees layer are over estimated. The third and fourth layup that both contain a large amount of fibers in that direction, are over estimated by approximately 15%. This large difference could be explained by the transverse shear effects that are not properly take into account in Equation 6.1.

But, more importantly, it is remarkable how good the general stiffness method agrees with Abaqus' native composite modeling technique. Only a small difference in the buckling load is visible for layup 3 and 5. A more detailed examination of the difference is made in Figure 11.2 where the relative difference to the buckling load estimate using Abaqus' native approach is plotted for increasing slenderness. The result further substantiated that for thicker cross section the transverse stiffness is important to con-




Figure 11.2: Influence of element slenderness, a/h , on the accuracy of the buckling load prediction for several modelling techniques.

sider. The correction factor used here shows a good agreement with the composite modeling technique

natively used in Abaqus. For a slenderness ratio higher than 12 the difference in buckling load is less than 0.05% percent. That is considered negligible, especially considering that this modeling technique is only used during the first design phase. But also for slenderness ratios as low as 5, the agreement is more than satisfactory for the use in early design phases at difference of merely 0.25%.

Similar to the results in Figure 11.1, the 45 degrees dominated layups show a large error for low slenderness ratios. Still the general shell approach is a viable option for these type of laminates, as long as care is taken that elements are thin.

11.1.2. Optimization

The local stiffness optimization is validated by comparing the results obtained for the first benchmark, a simple supported plate in uniaxial compression, with the results obtained by IJsselmuiden et al.. [20] Five individual optimization have been performed using the framework with 18 samples per iteration. The optimization processed was terminated after 300 iterations or when the number of missed subsequent iterations passed the 30. In all five cases the last criteria was the actual stopping criteria. The buckling loads reached are plotted in Figure 11.3 against the theoretical loads obtained by IJsselmuiden et al.[20]. The designs are fed through the same Abaqus model and added to the figure to correct for differences in their custom FEA implementation. All the loads are normalized according to

$$\tilde{N}_{cr} = \frac{N_{cr} a^2}{E_{11} \cdot h^3} \quad (11.1)$$

where a is the width of the plate, h is the thickness, and E_{11} the stiffness of the material in the principal direction.




Figure 11.3: Comparison of the optimization results for the simple supported plate between the frameworks' optimizer and the results reported by IJsselmuiden. [20] These results are for a square division of the number of regions, such that a $m \times m$ layout is obtained.

Optimized designs generated using the framework increase in performance when more design freedom is provided. A higher number of regions leads to a better structural performance, initially. For 20 by 20 regions a lower buckling load is found by the optimizer. Literature reports an increase in buckling performance when the number of regions is increased, where the $m = 20$ plate reaches a normalized buckling load of 3.93. More striking, the lamination parameters reported in literature provide a higher structural performance in the same FEA model than those optimized by the framework. Closer inspection of the differences in these results is performed later.

The reference results were optimized using a custom FEA code and here converted to a model in Abaqus. The literature results for the 10 by 10 plate are 5% lower than those simulated using the

FEA model of the framework. Results reported for the 20 by 20 plate are 10% higher than those simulated here. The differences could be explained by the more advanced elements used in Abaqus or a difference in the implementation of the boundary conditions.

Since the objective was to optimize the same problem as described in literature, a similar buckling performance was expected. Plates with a low number of regions show a similar structural performance as those found in literature, but for a higher number of regions the optimizer can likely not negotiate all of the design variables. This is supported by the lower structural performance that is obtained for these designs, unlike those reported in literature. The difference in the optimized lamination parameters is shown in Figure 11.4 using a heat map.




Figure 11.4: Difference of the four designed Lamination Parameters with the design of IJsselmuiden. [20] Red means that the optimization of this framework resulted in a lamination parameter to high, while blue means to low.

The heat map shows a low difference for the V_1 and W_1 lamination parameters, with a maximum difference of 0.4 for both. Coherence with literature for V_3 and W_3 is considerably lower and some region differ nearly 0.8. Both V_3 and W_3 are especially underestimated at the top and bottom of the laminate, and overestimated in the center. This difference could be in part explained by a slight difference in modelling technique of the boundary conditions. It could also be that the stiffness optimization was terminated too early because to many design variables were changed. When many variables are used, a less optimal choice for one variable can reduce the buckling performance of the entire plate and undo also favorable changes.

An overestimation of W_3 in the center reduces the D_{66} term in the ABD matrix due to the negative sign in Γ_3 in Equation 2.7. The reduced D_{66} strongly influences the local buckling performance and could explain the reduction in buckling load. Secondly, underestimation of V_3 , and in part also W_3 , at the top and bottom of the laminate reduces the local A_{11} stiffness terms. Hence, the sides are less stiff in the load direction and they take up less load than the solution presented in literature. This results in lower buckling performance for the entire plate because the center takes more load and is activated in buckling earlier.

A comparison per region is made by computing the normalized l2 error of every region and creating a histogram of the errors. This diagram is used to quantify the significance of the difference between the results and is shown in Figure 11.5.

The histogram shows that the 60% of the samples is within 2.5% of the literature optimized result and 90% of the regions are within 10%. The highest outliers have a normalized difference of 20%. Still the difference in the optimized load is roughly 10% which emphasize how sensitive the buckling load is to small change in stiffness distribution. Although the histogram shows how close the results are, it fails to account for the relative importance of these difference.

11.2. Stacking Sequence Retrieval

As part of the framework a parallel genetic algorithm (GA) was introduced to both speed up the retrieval and improve the fitness of the stacking sequences. These two claims are validated by comparing the retrieval performance for plates with up to 25 unique regions in the first subsection.




Figure 11.5: Histogram of the normalized L2 error per region shows that 50% percent is within 0.025 from the optimum.

The GA was allowed a maximum of 100 iterations using 100 individuals per population; the other parameters can be found in Table 7.1. All of the runs were terminated because the maximum number of iterations was reached. At termination the parallel GA were converged because there was no improvement for the last 20 evolutions, while minor improvements were still being made in the single GA when more than 9 guides were considered. The L2 error of the lamination speed, retrieval time, and load reproduction of the parallel and single GA are measured for two different runs of each plate layout and plotted in Figure 11.6.




Figure 11.6: Convergence of the stacking sequence retrieval process for the single GA and the parallel GA.

Still, the parallel GA took less time to run for all of the runs and for both methods a linear time dependency is shown in the right most graph. However, the first two graphs show that the parallel method can better approximate the given LPs for a high number of guides. The shaded area around the single GA shows that there is a large uncertainty in the load reproduction. That same uncertainty is also visible in the L2 error although less severe. It is observed that for a design with low number of guides the single GA performs similar to the parallel.

The trends seen in Figure 11.6 point out that the parallel GA has a good performance and scales linearly in time with the number of guides. The results of the single GA are less conclusive because the optimization did not fully converge in the limited iteration. Nevertheless, when both computational time and accuracy are considered, the parallel application of the GA performs better.

12

Results

The developed and validated framework is used to obtain several design results. Based on the plate manufactured by Goma[12] new automated designs are generated that incorporate more unique regions. Also the effect of a higher number of regions perpendicular to the load is investigated for this type of plate. Then results of a more complex structure, a cylinder in bending, are discussed. This cylinder is representative for a hull under bending and relates to the results obtained by Blom et al.. [5] All three design problems also consider the complexity of the designs through the patch metrics.

Secondly, the framework is evaluated as a whole. The framework is designed to find a solution that closely resembles the buckling performance of the local stiffness design, while being manufacturable and complying with blending guidelines. Subsequent steps in the framework make the design manufacturable but degrade the structural performance of design. The output of the framework is a manufacturable laminate of which an example is shown in Figure 12.1 for the 5x5 layout of the plate.




Figure 12.1: Cross section view perpendicular to the load direction at multiple y locations of the optimal 5x5 plate design. This illustrates an output of the framework where all the plies have been designed such that the laminate is manufacturable.

Lastly, the weight and cost perspective is discussed. This provides insight in the advantages of SFVS laminates and creates a method trade design complexity for additional weight savings.

12.1. Clamped Plate with Uniform Region Layout

The original benchmark problem introduced in subsection 2.4.2 was designed using 5x5 regions. It was expected that an increase in sections would yield improved structural performance but could not be designed by hand. The new framework allows exploration plates with more regions. Designs have been generated up to 15 by 15 regions using a mesh size close to 60 elements¹ in each direction. Similar to the validation, the optimization uses 18 samples per iteration and stops after 30 missed iterations. The stopping criteria of 300 iterations was never reached. The importance of the sample settings is discussed in Appendix D. The structural performance is measured after each design stage and visualized in Figure 12.2.




Figure 12.2: The maximum buckling load is plotted for every design stage for a changing number of regions for a uniform plate.

The best design is a 5x5 plate, similar to the one manufactured in 2019 that is marked with the star, with a structural performance of 118 kN. That is 37% higher than the CS design with 86 kN and 12% lower than the plate manufactured by Gomaa². Global designs for the 7x7 and 9x9 plate reach 116 kN and 114 kN respectively. Design generation with more regions were terminate after 60 hours because their global interpretation did not converge.

The same convergence problem encountered in the validation chapter shows up in the lamination parameter (LP) designs. The load increases up to $m = 7$ and then decreases as the number of regions is increased further. The stacking sequence retrieval process works as expected and the buckling performance is only reduced by 5% on average.

After application of the cellular automaton the structural performance of the 5x5 plate has been reduced by 30% from the original. However, this reduction is much lower for the 15 by 15 region plate and is also less for the 7x7 and 9x9 plates. Also, the locally blended design does not seem affected by the convergence problems for these low number of regions. Blended designs seem constrained by the cohesion that is required between the regions and therefore designs with more region can reach higher loads.

Through the application of the patch metrics, the complexity of the designs can be assessed. The num-

¹A mesh convergence showed no change above 4 elements per region or more than 25 elements in both directions. A total mesh of 60 elements in both direction was chosen because that allows 4 elements per region for the 15 by 15 plate

²The CA used to manufacture that plate was not able to remove butted edges at different levels and hence that plate did not comply with the relaxed blending guideline.

ber of patches, mean normalized aspect ratio, and mean number of corners have been computed for several designs generated by the global interpretation. These results are scattered against the buckling load in Figure 12.3, 12.4, and 12.5 respectively. The pareto optima are enlarged and connected to mark the pareto front.




Figure 12.3: The plot describes how higher number of patches influences the structural performance of different region layouts. The maximum number of patches is shown by the vertical line. This is equivalent to the number of regions.




Figure 12.4: The complexity of design is described by the mean normalized aspect ratio and is traded with the structural performance.




Figure 12.5: The complexity of design is described by the average number of corners in a patch and is traded with the structural performance.

The increase in number of regions causes an increase in the number of patches and average number of corners per patch. The mean normalized aspect ratio remains constant at 1.20 for all of the designs. The mean number of corners for the best design is between 4.5 and 5 corners per patch. Four corners per patch is equivalent to a rectangular patch and is considered easy to manufacture.

On closer inspection, an increase in the complexity and normalized aspect ratio causes an increase in the buckling load. This describes the tradeoff between manufacturing complexity and buckling performance that was expected at forehand. This trend is also visible for the number of patches which suggests that to obtain a higher buckling performance more patches are needed for the same region layout.

12.2. Clamped Plate with Varying Cross Axis Regions

It has been reported that increasing the number of regions perpendicular to the load leads to load redistribution and a higher buckling load. [20] Since the number of design regions seems limited by the global interpretation algorithm, a restructuring of the region layout to better distribute the number of unique regions can result in higher structural performance. This is investigated by designing several plates using a varying number of regions perpendicular to the loading direction. The design loads for these 5xm designs are shown in Figure 12.6.

The global interpreted design of the 5x7 region plate outperforms the 5x5 region by 5 kN; The buckling load is 123 kN. The buckling performance of the 5x9 region is considerably lower at 115 kN. Those loads are an improvement of 42% and 37% improvement over the CS design respectively. The 5x3 plate, on the other hand, only has a 5% load improvement in the global interpreted design whilst the local stiffness design managed a 63% improvement. It is a known issue that the preliminary patch foundation step cannot cope with designs containing only 3 regions in one direction, thus explaining the large reduction in performance gain.

Also, the local stiffness convergence issues had a big influence on the performance of the 5x9 plate. The 5x7 plate and 5x9 plate managed a top buckling performance of 173 kN and 146 kN respectively, whilst the 7x7 and 9x9 reached 186 kN and 177 kN respectively. Additional samples and/or runs might be necessary to further increase the yield of the 5x9 plate.




Figure 12.6: The maximum design load for each design stage as a function of the number of design region perpendicular to the loaded direction.




Figure 12.7: The number of patches present in the plate is traded with the buckling performance through the plotted pareto distribution. The pareto front is shown by connecting the individual pareto optima.

The pareto distribution considering the number of patches and buckling load is shown in Figure 12.7. In general, the increase in the number of patches is less than was observed for the uniform plates. This could be because the total number of regions is less, but could also be explained by an improved layout of the patches. The exact dimensions and placement of regions is different between these designs and a finer grid provides the framework with more freedom to start and stop a patch.

Although the actual load of the 5x9 plate is not as high as that of the 5x5, the pareto front is more favorable for a low number of patches. A load increase of 20% is achieved with just 2 additional patches when the 5x9 region configuration is used. An equivalent design using the 5x5 region layout would have used close to 20 patches, almost twice the amount. Also the designs with the highest load use approximately the same number of patches. The higher number of regions might thus also provide the freedom to improve the location of patches. The 5x7 plate uses more patches in virtually every design and achieves the best design with 29 patches. The design with lowest amount of patches is 24 and still outperforms the 5x5 plate with a buckling load of 121 kN

The mean normalized aspect ratio and mean complexity of the patches are not substantially different from the results obtained for the uniform plate. These distribution plots have been added to Appendix C for reference.

12.3. Cylinder in Bending

The framework was also used to optimize a cylinder in bending and increase its buckling performance. Layouts with 1 region axially and 8, 18, 28, 38 circumferentially were created to allow redistribution of load. Axial redistribution was not considered because redistribution in that direction is not likely to increase the buckling performance. [5] The optimization was performed using 75 elements axially and 152, 162, 168, and 152 elements circumferentially for the four layouts respectively. These meshes are similar to the ones used by Blom et al.(71, 170)[5] and should allow for a good comparison to her case. Additionally, these meshes are a good trade off between optimization time and convergence, see Appendix E. The buckling load was predicted after every design step and the results are shown in Figure 12.8.




Figure 12.8: The maximum design load for the design stage of a cylinder. Convergence is more difficult hence the result are more erratic than for the other design results.

The 1x18 cylinder has the best global interpreted design at 817 kNm, 8 kNm higher than results currently

reported in literature for fiber steered designs and 30% higher than the baseline design. The highest local stiffness load is also obtained for the same region layout at 830 kNm. Similarly, the 1x28 and 1x38 cylinders have final buckling loads of 816 kNm and 810 kNm respectively. The 1x8 region layout has a buckling performance of just 725 kNm, although that is still an improvement of 15% over the baseline design.

The local stiffness design of the 1x18 cylinder has the highest load of 830 kNm, while the 1x28 and 1x38 designs have approximately the same load, 825 kNm, 829 kNm. The fact that a similar load is reached suggests that the optimization did converge and it is likely that a finer region layout does not increase in the buckling performance any higher. The 1x18 region layout provides enough design freedom to soften the bottom of the cylinder enough and increase the buckling performance of the cylinder to the maximum achievable.

It is remarkable that for each of the region layouts the unblended load, which is the load of the design that matches the optimal stiffness closest, is lower than the blended. It was expected that introduction of the blending guidelines using the CA resulted in a lower load. This is the case for plates and is inline with normal optimization theory where more constraints generally lower the objective. However, here the buckling performance of the cylinder seems to benefit from cohesion introduced by cellular automaton. The match of the stacking sequence can be quantified by computing the L2 error from the optimal lamination parameters. This is shown for the cylinder with region layout 1x38 in Figure 12.9. The buckling load of the stacking sequences is 30 kNm lower than the optimal design.




Figure 12.9: L2 error in the lamination parameters between the optimal design and the design after stacking sequence retrieval.

The stacking sequence does not match the desired stiffness in the lower side of the cylinder. This area is activated in buckling and a high difference here explains the discrepancy in the observed buckling load. It could be that the optimal lamination parameters for these regions are in a part of the lamination parameter design space that does not contain many stacking sequences. As a result there does not exist a set of fiber orientations that closely matches the requested lamination parameters. But this requires more investigation. Similar observations can be made for the other cylinder designs.

Figure 12.10 shows the pareto front of the number of patches versus the buckling performance for each of the cylinders. The pareto front shows that a small difference in the number of patches does impact the load. A slightly lower number of patches reduces the buckling load significantly. The 1x38 cylinder requires just four patch more to go from 734 kNm to 809 kNm at 40 and 44 patches respectively. The optimal design of the 1x18 cylinder uses only 22 patches. The optimal design of the 1x8 cylinder uses only 13 patches, 7 more than the baseline, and increases the load by 15%.




Figure 12.10: The number of patches is traded with the buckling performance of the cylinder through the plotted pareto distribution. Pareto optima are marked and connected.

Pareto distributions of the other patch metrics have not been generated because in the case of the cylinder all the patches have a rectangular shape.

12.4. Framework Performance

It is clear that the framework's performance strongly influences the design results. Whereas the results of the local optimization were validated in the validation chapter against known values, here the global interpretation algorithm's performance is investigated. The lack of reference algorithms prevents actual validation and is unknown if better designs might be possible.

The algorithm takes multiple steps to resolve several butted globally. However, each applied fix generally brings the design further from the optimum. Results have been obtained for plates with a maximum of 81 regions because design with more regions did converge within 60 hours. The changes introduced to obtain a final interpreted design are traced and shown as tree branches in Figure 12.11.

Designs grow from the bottom left and are first fed through the CA algorithm. Then the GI algorithm takes over and gradually changes the design one step at a time, these steps are shown on the y-axis. The normalized change used in the graph's x-axis is the average difference with respect to the unblended geno:

$$\text{Normalized Change} = \frac{|\tilde{\theta}_{unblended} - \tilde{\theta}_{interpreted}|}{N \cdot K}$$

where N and K are the number of regions and layers respectively.

The inset graphs shows the relation of the normalized change to the normalized load. The normalized load is the global interpreted load divided by the local optimized load that makes use of the lamination parameters. The normalized load is a direct measure of the structural degradation caused by the framework. The inset graphs shows a good agreement between the normalized change and the load degradation. For each design a relation between the normalized change and the load is observed. This is supported by the coefficient of determination and the randomness of the sample points around the linear fitted line.

The tree diagram shows that more steps are needed for designs that contain a high number of regions. On average solutions for the 9x9 plate require 35 steps in the GI, while the 5x5 only requires 10 steps.




Figure 12.11: Branches that lead to a solution in the global interpretation search, where the x axis shows the difference and y shows the iteration increments. The inset graph shows that there is a strong correlation between the load and the normalized change plotted on the x axis of the parent graph.

Also the 7x7 plate requires 22 steps. Each step increases the distance from the local optimum further. From the inset graph it can be determined that the increase in normalized change caused by the step results in a decrease in the structural performance.

Interestingly the 5x7 plate requires less normalized change to come to a globally blended results than the 5x5. Moreover, the 5x9 plates requires a similar amount of adjustments as the 5x5 plate. Seemingly less errors were present at the start of the global interpretation and a solution could be found more quickly. There is no clear relation between the input to the blending stages and the blendability of the design.

Through back tracking of the solution it can be seen that many global interpreted designs originate from the same base design. The 9x9 plate does not branch until after 27 steps. Branches that were tried earlier might not have led to a global interpreted design. This is not the case for the 7x7 plate where multiple branches can be seen after 5 steps already, similarly for the 5x7 design.

Each of these designs started with 16 different combinations of stacking sequences. Only the 5x5 and 5x7 plate have multiple branches directly aft the CA stage that lead to solutions. Either the CA produces locally blended designs that could not be successfully globally interpreted, or several input seeds have been blended locally to the same design. In the latter case the branches cannot be distinguished from each other.

The tree of Figure 12.11 suggests a relation between the problem size and the normalized change needed to obtain a blended design. This is plotted in Figure 12.12. The problem size is quantified using the number of potential butted edges in the region layout.




Figure 12.12: The normalized change is plotted as function of the problem size to investigate a relation between problem size and blending.

Although both the 7x7 and 9x9 plate show that more change is needed when the number of interfaces increases, this is not the case for the 5x7 and 5x9 plate. Moreover, the highest mean normalized change measured for the 5x5 plate is higher than that of the lowest 7x7 plate. The change needed to blend the 7x7 and 9x9 plate is also similar while there are nearly 500 edges more in the 9x9 plate.

The lowest change needed is seen for the 5x7 plate at just 0.9 indicating that every layer has changed on average less than one orientation step. The highest change is measured for the 9x9 plate where on average every region changed 1.6 orientation steps. That is equivalent to approximately 25 degrees in fiber orientations and could explain the large drop structural performance.

12.5. Weight Savings

So far, the discussion has been centered around increasing the structural performance. However, it is much more interesting to see how such an improvement relates to the weight that can be saved. Here a link is made between the structural improvement and possible weight savings by scaling the thickness of the laminates to match the baseline buckling performance. The reduction in thickness directly relates to the weight that can be saved.

According to Equation 6.1, the buckling performance of plates is strongly related to the flexural stiffness of the laminate. The thickness of the laminate relates linear to weight, but has a third order relation with the flexural stiffness. This leads to the following relation:

$$\frac{t_{eq}}{t_{cs}} = \sqrt[3]{\frac{L_{cs}}{L_{opt}}} \quad (12.1)$$

with $\frac{t_{eq}}{t_{cs}}$ being the thickness ratio between the panel with a buckling load equivalent to the baseline panel and the constant stiffness panel's thickness, L_{cs} the constant stiffness panel's load, and L_{opt} the achieved optimal load. Here, it is assumed that the relative structural improvement remains the same when the thickness is decreased and that the thickness is continuously scalable. In reality a design with fewer and/or thinner plies has to be made and there is no guarantee that the same improvement can be accomplished. However, it does indicate the possible savings.

In this thesis 42% structural improvement was reached which yields a thickness decreases of 11.3%. Such a thickness decreases also leads to a weight saving of roughly 11.3%. A pareto distribution that contains all the plate's region layouts and their possible savings is shown in Figure 12.13.




Figure 12.13: Pareto distribution of the weight savings and the number of additional patches used in the design.

Both the 7x7 and 9x9 region layouts do not contribute to the weight savings because their design load was not higher than designs with an equivalent number of patches. The pareto optima are connected to form the frontier that starts at the constant stiffness design and continues via the 5x9 and 5x5 designs to the 5x7 plates. The plate seems to benefit from a higher number of regions at a lower number of additional plies. That is similar to what is seen in the buckling performance distributions of Figure 12.3 and 12.7.

The pareto distribution for the cylinder has been added to Appendix C. A 7% weight savings can be achieved in the cylinder using the 1x18 region layout.

12.6. Cost

Lastly, the weight savings should be seen in relation to the increase in manufacturing cost. The cost is difficult to estimate because exact manufacturing details are not available. Hence, cost estimates of \$1, \$2, \$5, \$10, and \$20 dollar per additional patch are used. The cost for a specific weight saving is then obtained by multiplying the number of additional patches with these numbers and is presented in Figure 12.14. Straight lines can be drawn to mark the break even cost for airlines based on the reduction in the Direct Operating Cost (DOC) that is caused by the weight saving. Three lines for \$500/kg, \$1000/kg, and \$2000/kg are drawn and the favored zone, to the bottom right, is marked with an arrow. On this side of the break even line the Direct Operating Cost (DOC) for airlines will actually decrease. A similar distribution for the cylinder is added to Appendix C.




Figure 12.14: Influence on the Direct Operating Cost caused by the weight savings and increase in manufacturing cost due to the additional patches. Break even lines mark the line of equivalent DOC where the reduction in the fuel cost equals the increase in manufacturing cost.

The addition of just a few patches allows significant weight decreases of more than 0.1kg. Adding a few patches is not expensive and thus results in a DOC decrease regardless of the cost-to-weight ratio. Depending on the combination of cost-to-weight and cost per patch, a cross-over can be determined where further weight savings result in a higher DOC because of the additional manufacturing cost. These points are however not yet reached when a cost-to-weight ratio of \$2000/kg is used.

The annotations mark the developments in manufacturing research and fuel price prospects. An increase in fuel prices favors more weight savings because additional manufacturing cost are cancelled out. Also, weight savings will become cheaper due to advances in manufacturing technology. The demand for more weight savings requires further development of the algorithms, but will eventually result in more sustainable air travel.

13

Conclusion

Straight Fiber Variable Stiffness(SFVS) laminates have the potential to significantly reduce the weight of large structures. Their usage is currently limited through their inherent design complexity caused by the large design space provided by modern guidelines. A novel framework is developed that allows application of these laminates in larger structures. Through the use of lamination parameters an optimal local stiffness can be obtained and used as a starting point for blending. A novel blending algorithm converges large scale designs to locally blended solutions and small designs to globally interpreted and manufacturable laminates. The framework allows for end-to-end optimization and brings the adoption of these laminates in large scale structures closer. Here conclusions on the performance of the framework and the produced designs are discussed

The Framework

Before the framework was put to use it was validated using the optimal lamination parameters provided by IJsselmuiden et al.. [20] A simply supported plate in uniaxial compression was optimized using four lamination parameters for each region through a gradient descent optimization that uses surrogate modelling and finite element analysis. The obtained results show good agreement with the optima reported in literature for small scale designs. Differences in buckling load of 5% or less are observed. However, when 100 unique design regions are considered the optimization was not able to converge to the same optimum as reported in literature. A difference of 10% in buckling performance is observed while 90% of the design regions have a difference of less than 10% from the optimal lamination parameters. It is concluded that although the optimizer can improve the buckling performance it has convergence issues when a large number of variables is used.

During the next step of the optimization method, a genetic algorithm based on the work of van Van Campen[52] is used to convert the lamination parameters to stacking sequences. Parallelization of the genetic algorithm separates the optimization into individual problems for each regions and yields a five fold speed increase for designs considering 25 regions. Also the convergence of more than 9 regions is found to be better with this new approach. Derivative information of the surrogate model can be integrated in the fitness function to aid in the convergence. To sum up, separation of the stacking sequence retrieval process allows the method to approximate the solution closer in less time.

Lastly, a major contribution to the framework is the automated global interpretation of the results. This method uses connections to model the cohesion between regions and uses graph theory to constraint manufacturing order. This solver requires exponentially more computation when the number of regions grows. Thus, plates with more than 9x9 regions were not able to converge within 60 hours.

SFVS Laminates

A total of three different designs problems were optimized using the developed framework. Two clamped plates in uniaxial compression with an uniform varying and nonuniform varying region layout and a cylinder in bending were optimized for optimal buckling performance. They are discussed in order.

The design of a plate with a uniform 5x5 region layout shows the highest buckling performance at 118 kN of compressive force. Although this is 10% lower than the predicted buckling load of 132 kN for the

plate of Goma [12], the design does not contain butted edges and is manufacturable using automated processes. Furthermore, this is still a 37% improvement over the best constant stiffness design. The expected load improvements of the 7x7 and 9x9 are not observed because both designs topped their structural performance at 116 kN and 114 kN respectively. Through the use of solution tracking it is found that the GI required 22 and 32 fixes respectively for these plates. In each of these fixes the normalized angle change increases which is directly related to the improvement degradation. The high number of fixes causes these designs to lose their potential because the locally blended configuration had a load of nearly 150 kN in the case of 9x9 plate.

Through the use of the patch metric it is concluded that plates with a higher number of regions will be more expensive to manufacture. On average the 9x9 plate has five times the number of patches of the 5x5 plate. When each of the region layout is considered individually, a slightly higher number of patches results in higher structural performance. The 5x5 plate reaches its maximum load of 118 kN at 23 patches and its minimum of 91 kN at 16 patches. Similar trends are seen for the mean normalized aspect ratio and the mean number of corners. The highest buckling load is achieved at 1.27 and 4.7 respectively. Unsurprisingly, an increase in the buckling load comes at the cost of added complexity. Interestingly, the highest buckling load of the 7x7 plate has a lower mean normalized aspect ratio of 1.16, thus suggesting that these designs contain fewer complex shaped patches at the cost of more.

The buckling load of the 5x7 region layout is higher than that of the 5x5 plate. A load of 123 kN is found which is a 4% improvement over the 5x5 plate and a 42% improvement over the constant stiffness design. This is in accordance with the expectancy that a higher number of regions perpendicular to the loading direction allows the algorithm to redistribute the load better. This is however not the case for the 5x9 plate which has buckling load of just 115 kN.

However, the 5x9 design is able to reach this load with the same number of patches (24 patches) while their were nearly twice as many design regions. This shows that the algorithm is able to distribute patches properly over a high number of regions. Also, at just 10 patches a buckling load of 100 kN is achieved. Thus by adding two patches on top of the constant stiffness (CS) design, the buckling load is increased with 20%. Adding two patches is only a small increase in manufacturing complexity and cost.

Thirdly, a cylinder representative of a fuselage is optimized using the framework. Literature reports a 27% increase over the constant stiffness to a buckling moment of 809 kNm. A lamination parameter design that is created as part of this thesis and that uses only 10 unique regions reaches a buckling load of 830 kNm. After blending and global interpretation, the improvement dropped to 817(30% improvement) kNm which is still 8 kNm more than the best design reported in literature. Also the 1x28 and 1x38 region layout produced manufacturable designs that have a similar or better performance than those reported in literature. The cylinder with 5 unique regions(1x8) only achieved a buckling load of 725 kNm(15%).

Blending of these cylinders illustrates how cohesion between regions can also increase the buckling load. The buckling load of each of the blended designs is higher than that of the design directly after retrieval of the stacking sequence. In the case of the cylinder with 20 unique regions the load is 30 kNm lower in the second step and then rises by 23 kNm in third step. These discrepancies can be attributed to the high errors between the desired and retrieved lamination parameters in the lower side of the cylinder.

Weight savings of more than 11% were achieved using the 5x7 plate. Similarly, up to 7% of the cylinder's weight can be saved when straight fiber variable stiffness laminates are used. Depending on the exact cost of adding an additional patch, the generated design reduce the direct operating cost for airlines for nearly every cost-to-weight ratio. Fuel price prospects and manufacturing advances will open up more opportunities to save weight.

In conclusion, the framework developed can be used to optimize several designs using SFVS laminates in a standardize way. Weight savings of more than 11% for plates in compression are shown and more gain is still expected for finer grids. The laminates designed are fully manufacturable using current technologies and have the potential to decrease the direct operating cost for airlines and reduce emissions.

Discussion and Recommendations

During this thesis a framework was developed that automates the design process of straight fiber variable stiffness (SFVS) laminates. Variable stiffness designs were generated that increase the buckling performance with a maximum of 42% depending on the problem. Although this shows the potential of the framework and SFVS laminates as a whole, more work is needed for the adoption of these laminates. Here, potential improvements to the frameworks are discussed first and then additional research opportunities into SFVS laminates to accommodate a faster adoption are proposed.

Framework

A start has been made on a framework for the end-to-end design of SFVS laminates but more work is needed. Shortcomings are observed in both the convergence of the local stiffness optimization and in the blending algorithms.

The optimizer does not properly convergence to the local optimum when a high number of design variables is used. This could be accounted to the sampling method that does not take into account previous samples, or to the parameterization technique. Progressive latin hypercube sampling can be introduced to ensure good coverage of the design space during progressive sampling. [41] Moreover, the surrogate model's uncertainty can be used to indicate areas of interest and to selectively sample design spaces that are likely to yield a higher result. To this end bayesian optimization techniques might be implemented.

Experiments with up sampling of the design as a starting point require more research. Such methods allow parameterization of the design using less design variables and would thus increase the likelihood of convergence. Depending on the exact parameterization, this could also introduce cohesion between regions. For example, control points and linear interpolation of the stiffness results in a more cohesive design. Also the impact of variable thickness on the buckling performance can be investigated. Variable thickness allows for finer tailoring of the stiffness and might aid in identifying areas where weight can be saved.

Other parts of the algorithm can be improved as well. Global interpretation of designs with a large number of sections is still problematic. They do not properly converge and do require a large amount of change from the locally blended configuration. It is likely beneficiary to invest in improvements of the CA such that interfaces can be blended as a whole. Currently, the CA does sometimes require that a layer of a region connects to two other layers in a neighboring region. This can be circumvented when a different butted edge detection algorithm is used that takes into account the whole interface instead of just one edge.

The stacking sequence retrieval step causes a large degradation of the buckling performance in the case of the cylinder. Large differences in the stiffness between the lamination parameter optimum and retrieved stacking sequences are observed. More investigation is needed to determine how the cohesion introduced by blending increases the buckling load. Additionally, variable regions layouts might be used such that the lower side of the cylinder is made up of one large region. The derivative of the optimal lamination parameters might be used to determine the region density in a specific area of

the design. Those areas with a higher derivative in the lamination parameters might benefit from more regions. A higher number of regions provides more freedom to the blending algorithm and allows a better match with the desired lamination parameters without increasing the computational cost.

SFVS Laminates

After this research, the framework can be used to design other structures that could benefit from straight fiber variable stiffness. For example, similar to the work of Çelebi et al.[7], variable stiffness might be used to redirect loads and accommodate cutouts in cylinders or plates. Also other loading patterns that can benefit from load redistribution might be investigated using this framework. To this end, the current methodology might be used in the design of an actual wingbox or fuselage section to estimate the real savings. These should also include multiple load cases.

However, an increased understanding of these laminates is needed before they are permitted in industry. The fatigue performance of these laminates is still largely unknown. Especially, the presence of ply drops might provide path ways for cracks to grow faster. Although a sudden change in the laminate, caused by ply drop, might also stop cracks from propagating. Such research should be executed at the laminate level and might then be extrapolated to a whole structure. A thorough understanding of the fatigue behavior is critical for their wide spread adoption.

More research is needed into the production processes of these laminates. The cost drivers for these laminates are largely unknown and a good trade-off cannot be made yet. Cost predictions are limited to per patch cost because the influence of patch shape on the cost could not be estimate. Such a cost estimation study should take into account the specifics of the production method and should also account for the risks. Moreover, such a study might identify which developments in manufacturing technology are needed to drive down the cost further. For example, it is a known issue that plates created by this method are metastable and deform after cure. A thermal stress model can be created to predict the curing behavior of these laminates.

Straight fiber variable stiffness laminates have the potential to increase the structural performance of common structures such as plates and cylinders. These structural improvements ultimately lead to weight reduction, fuel consumption reductions, and more affordable air travel. However, currently more research on the industrial and scientific level is needed for these laminates to be adopted.


Bibliography

- [1] Abaqus. *Analysis User's Guide*. Dassault Systemes, Providence, RI, USA, 2016.
- [2] Abaqus. *User Subroutines Reference Guide*. Dassault Systemes, 2016.
- [3] Mostafa Abdalla, Zafer Gürdal, and Christos Kassapoglu. Formulation of Composite Laminate Robustness Constraint in Lamination Parameters Space. In *50th AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*, Palm Springs, California, May 2009. American Institute of Aeronautics and Astronautics. ISBN 978-1-60086-975-4. doi: 10.2514/6.2009-2478.
- [4] David B. Adams, Layne T. Watson, Zafer Gürdal, and Christine M. Anderson-Cook. Genetic algorithm optimization and blending of composite laminates by locally reducing laminate thickness. *Advances in Engineering Software*, 35(1):35–43, January 2004. ISSN 09659978. doi: 10.1016/j.advengsoft.2003.09.001.
- [5] Adriana W. Blom, Patrick B. Stickler, and Zafer Gürdal. Optimization of a composite cylinder under bending by tailoring stiffness properties in circumferential direction. *Composites Part B: Engineering*, 41(2):157–165, March 2010. ISSN 13598368. doi: 10.1016/j.compositesb.2009.10.004.
- [6] A.W Blom. *Structural Performance of Fiber-Placed, Variable-Stiffness Composite Conical and Cylindrical Shells*. PhD Thesis, Technology University of Delft, Delft, 2010.
- [7] Mansur Çelebi, Zafer Gürdal, Brian Tatting, Agnes Blom-Schieber, Mostafa Rassaian, Steven Wanthal, and Halit Süleyman Türkmen. Bending of Composite Cylindrical Shells with Circular Cutouts: Buckling and Failure Analysis. *Journal of Aircraft*, 56(4):1551–1564, July 2019. ISSN 1533-3868. doi: 10.2514/1.C035246.
- [8] David J Eberlein. *Composite Cylindrical Shell Buckling: Simulation & Experimental Correlation*. MSc. Thesis, Technology University of Delft, Delft, October 2019.
- [9] Bill Gates. *How to Avoid a Climate Disaster*. Allen Lane, Dublin, 2001.
- [10] Hossein Ghiasi, Kazem Fayazbakhsh, Damiano Pasini, and Larry Lessard. Optimum stacking sequence design of composite materials Part II: Variable stiffness design. *Composite Structures*, 93(1):1–13, December 2010. ISSN 0263-8223. doi: 10.1016/j.compstruct.2010.06.001.
- [11] Muhaned Gilani and Durmuş Sinan Körpe. Airline Weight Reduction to Minimize Direct Operating Cost. page 12, 2019.
- [12] M Goma. *Laminate Blending Demonstrator*. MSc, Delft University of Technology, Delft, 2018.
- [13] Zafer Gürdal, Raphael T. Haftka, and Prabhat Hajela. *Design and Optimization of Laminated Composite Materials*. Wiley, New York, 1999. ISBN 978-0-471-25276-4.
- [14] V.B. Hammer, M.P. Bendsøe, R. Lipton, and P. Pedersen. Parametrization in laminate design for optimal compliance. *International Journal of Solids and Structures*, 34(4):415–434, February 1997. ISSN 00207683. doi: 10.1016/S0020-7683(96)00023-6.
- [15] Alan Hiken. The Evolution of the Composite Fuselage: A Manufacturing Perspective. In George Dekoulis, editor, *Aerospace Engineering*. IntechOpen, November 2019. ISBN 978-1-83962-784-2 978-1-83962-786-6. doi: 10.5772/intechopen.82353.
- [16] Ch. Hueber, K. Horejsi, and R. Schledjewski. Review of cost estimation: Methods and models for aerospace composite manufacturing. *Advanced Manufacturing: Polymer & Composites Science*, 2(1):1–13, January 2016. ISSN 2055-0340, 2055-0359. doi: 10.1080/20550340.2016.1154642.
- [17] M.W. Hyer and H.H. Lee. The Use of curvilinear Fiber Format to Improve Buckling Resistance of Composite Plates with Central Circular Holes. *Composite Structures*, 18:239–261, 1991.
- [18] Samuel T. IJsselmuiden, Mostafa M. Abdalla, and Zafer Gürdal. Implementation of Strength-Based Failure Criteria in the Lamination Parameter Design Space. *AIAA Journal*, 46(7):1826–1834, July 2008. ISSN 0001-1452, 1533-385X. doi: 10.2514/1.35565.
- [19] Samuel T. IJsselmuiden, Mostafa M. Abdalla, Omprakash Seresta, and Zafer Gürdal. Multi-step blended stacking sequence design of panel assemblies with buckling constraints. *Composites*

- Part B: Engineering*, 40(4):329–336, June 2009. ISSN 13598368. doi: 10.1016/j.compositesb.2008.12.002.
- [20] Samuel T. IJsselmuiden, Mostafa M. Abdalla, and Zafer Gurdal. Optimization of Variable-Stiffness Panels for Maximum Buckling Load Using Lamination Parameters. *AIAA Journal*, 48(1):134–143, January 2010. ISSN 0001-1452, 1533-385X. doi: 10.2514/1.42490.
 - [21] L. B. Ilcewicz, P. J. Smith, C. T. Hanson, T. H. Walker, S. L. Metschan, G. E. Mabson, K. S. Wilden, B. W. Flynn, D. B. Scholz, D. R. Polland, H. G. Fredrikson, J. T. Olson, and B. F. Backman. Advanced Technology Composite Fuselage - Program Overview. Technical report, The Boeing Company, Seattle, 1997.
 - [22] Christos Kassapoglou. *Design and Analysis of Composite Structures*. John Wiley & Sons Ltd., United Kingdom, second edition, 2013.
 - [23] Markus Kaufmann, Dan Zenkert, and Per Wennhage. Integrated cost/weight optimization of aircraft structures. *Structural and Multidisciplinary Optimization*, 41(2):325–334, March 2010. ISSN 1615-147X, 1615-1488. doi: 10.1007/s00158-009-0413-1.
 - [24] A. Khani, M.M. Abdalla, and Z. Gürdal. Circumferential stiffness tailoring of general cross section cylinders for maximum buckling load with strength constraints. *Composite Structures*, 94(9):2851–2860, September 2012. ISSN 02638223. doi: 10.1016/j.compstruct.2012.04.018.
 - [25] Byung Chul Kim, Kevin Potter, and Paul M. Weaver. Continuous tow shearing for manufacturing variable angle tow composites. *Composites Part A: Applied Science and Manufacturing*, 43(8):1347–1356, August 2012. ISSN 1359835X. doi: 10.1016/j.compositesa.2012.02.024.
 - [26] Dieter Kraft. A Software Package For Sequential Quadratic Programming. Technical report, Institut für Dynamik der Flugsysteme, OberpfaffenHofen, July 1988.
 - [27] Gustavo Gonzalez Lozano, Ashutosh Tiwari, Christopher Turner, and Simon Astwood. A review on design for manufacture of variable stiffness composite laminates. *Proceedings of the Institution of Mechanical Engineers, Part B: Journal of Engineering Manufacture*, 230(6):981–992, June 2016. ISSN 0954-4054, 2041-2975. doi: 10.1177/0954405415600012.
 - [28] Dirk H.-J.A. Lukaszewicz, Carwyn Ward, and Kevin D. Potter. The engineering aspects of automated prepreg layup: History, present and future. *Composites Part B: Engineering*, 43(3):997–1009, April 2012. ISSN 13598368. doi: 10.1016/j.compositesb.2011.12.003.
 - [29] Y. Narita. Layerwise optimization for the maximum fundamental frequency of laminated composite plates. *Journal of Sound and Vibration*, 263(5):1005–1016, June 2003. ISSN 0022460X. doi: 10.1016/S0022-460X(03)00270-0.
 - [30] Mahdi Arian Nik. A comparative study of metamodeling methods for the design optimization of variable stiffness composites. *Composite Structures*, page 8, 2014.
 - [31] Marco Picchi Scardaoni, Marco Montemurro, Enrico Panettieri, and Anita Catapano. New blending constraints and a stack-recovery strategy for the multi-scale design of composite laminates. *Structural and Multidisciplinary Optimization*, September 2020. ISSN 1615-147X, 1615-1488. doi: 10.1007/s00158-020-02725-x.
 - [32] D. R. Polland, S. R. Finn, K. H. Griess, J. L. Hafenrichter, C. T. Hanson, L. B. Ilcewicz, S. L. Metschan, D. B. Scholz, and P. J. Smith. Global Cost and Weight Evaluation of Fuselage Side Panel Design Concepts. Technical report, The Boeing Company, Seattle, 1997.
 - [33] Carl Edward Rasmussen and Christopher K. I. Williams. *Gaussian Processes for Machine Learning*. Adaptive Computation and Machine Learning. MIT Press, Cambridge, Mass, 2006. ISBN 978-0-262-18253-9.
 - [34] Mohammad Rouhi, Hossein Ghayoor, Suong V. Hoa, and Mehdi Hojjati. Computational efficiency and accuracy of multi-step design optimization method for variable stiffness composite structures. *Thin-Walled Structures*, 113:136–143, April 2017. ISSN 02638231. doi: 10.1016/j.tws.2017.01.019.
 - [35] Mohammad Rouhi, Hossein Ghayoor, Jeffrey Fortin-Simpson, Tom T. Zacchia, Suong V. Hoa, and Mehdi Hojjati. Design, manufacturing, and testing of a variable stiffness composite cylinder. *Composite Structures*, 184:146–152, January 2018. ISSN 02638223. doi: 10.1016/j.compstruct.2017.09.090.
 - [36] Anthony Sabido, Luis Bahamonde, Ramy Harik, and Michel J.L. van Tooren. Maturity assessment of the laminate variable stiffness design process. *Composite Structures*, 160:804–812, January 2017. ISSN 02638223. doi: 10.1016/j.compstruct.2016.10.081.

- [37] Marco Picchi Scardaoni and Marco Montemurro. A general global-local modelling framework for the deterministic optimisation of composite structures. *Structural and Multidisciplinary Optimization*, 62(4):1927–1949, October 2020. ISSN 1615-147X, 1615-1488. doi: 10.1007/s00158-020-02586-4.
- [38] Shahriar Setoodeh, Mostafa Abdallah, Zafer Gurdal, and Brian Tatting. Design of Variable-Stiffness Composite Laminates for Maximum In-Plane Stiffness Using Lamination Parameters. In *46th AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics and Materials Conference*, Austin, Texas, April 2005. American Institute of Aeronautics and Astronautics. ISBN 978-1-62410-065-9. doi: 10.2514/6.2005-2083.
- [39] Shahriar Setoodeh, Mostafa M. Abdalla, and Zafer Gürdal. Design of variable–stiffness laminates using lamination parameters. *Composites Part B: Engineering*, 37(4-5):301–309, June 2006. ISSN 13598368. doi: 10.1016/j.compositesb.2005.12.001.
- [40] Shahriar Setoodeh, Adriana Blom, Mostafa Abdalla, and Zafer Gürdal. Generating Curvilinear Fiber Paths from Lamination Parameters Distribution. In *47th AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference & 14th AIAA/ASME/AHS Adaptive Structures Conference & 7th*, Newport, Rhode Island, May 2006. American Institute of Aeronautics and Astronautics. ISBN 978-1-62410-040-6. doi: 10.2514/6.2006-1875.
- [41] Razi Sheikholeslami and Saman Razavi. Progressive Latin Hypercube Sampling: An efficient approach for robust sampling-based analysis of environmental models. *Environmental Modelling & Software*, 93:109–126, July 2017. ISSN 13648152. doi: 10.1016/j.envsoft.2017.03.010.
- [42] B. A. R. Soares, E. Henriques, I. Ribeiro, and M. Freitas. Cost analysis of alternative automated technologies for composite parts production. *International Journal of Production Research*, 57(6):1797–1810, March 2019. ISSN 0020-7543, 1366-588X. doi: 10.1080/00207543.2018.1508903.
- [43] Grant Soremekun and Zafer Gu. Stacking sequence blending of multiple composite laminates using genetic algorithms. *Composite Structures*, page 10, 2002.
- [44] Hiroshi Suemasu, Yuichiro Aoki, Sunao Sugimoto, and Toshiya Nakamura. Effect of gap on strengths of automated fiber placement manufactured laminates. *Composite Structures*, 263: 113677, May 2021. ISSN 02638223. doi: 10.1016/j.compstruct.2021.113677.
- [45] B.F. Tatting. *Analysis and Design of Variable Stiffness Composite Cylinders*. PhD Thesis, Virginia Tech, Blacksburg, Virginia, 1998.
- [46] Julien Van Campen and Zafer Gürdal. Retrieving Variable Stiffness Laminates from Lamination Parameters Distribution. In *50th AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*, Palm Springs, California, May 2009. American Institute of Aeronautics and Astronautics. ISBN 978-1-60086-975-4. doi: 10.2514/6.2009-2183.
- [47] Julien Van Campen, Omprakash Seresta, Mostafa Abdalla, and Zafer Gürdal. General Blending Definitions for Stacking Sequence Design of Composite Laminate Structures. In *49th AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference & 16th AIAA/ASME/AHS Adaptive Structures Conference & 10th*, Schaumburg, IL, April 2008. American Institute of Aeronautics and Astronautics. ISBN 978-1-60086-993-8. doi: 10.2514/6.2008-1798.
- [48] Julien Van Campen, Christos Kassapoglou, and Zafer Gurdal. Design of Fiber-steered Variable-stiffness Laminates Based on a Given Lamination Parameters Distribution. In *52nd AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics and Materials Conference*, Denver, Colorado, April 2011. American Institute of Aeronautics and Astronautics. ISBN 978-1-60086-951-8. doi: 10.2514/6.2011-1894.
- [49] Julien M.J.F. Van Campen, Christos Kassapoglou, and Zafer Gürdal. Generating realistic laminate fiber angle distributions for optimal variable stiffness laminates. *Composites Part B: Engineering*, 43(2):354–360, March 2012. ISSN 13598368. doi: 10.1016/j.compositesb.2011.10.014.
- [50] Ellen Van den Oord. *Overcoming the Curse of Dimensionality in Composite Laminate Blending*. Msc Thesis, TU Delft, Delft, 2018.
- [51] Michael J. Van Tooren, Ifat Jahangir, and Ali Elham. Optimization of variable stiffness composite plates with cut-outs subjected to compression, tension and shear using an adjoint formulation. In *57th AIAA/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*, San Diego, California, USA, January 2016. American Institute of Aeronautics and Astronautics. ISBN 978-1-62410-392-6. doi: 10.2514/6.2016-1970.

- [52] Julien Marie Jan Ferdinand van Van Campen. *Optimum Lay-up Design of Variable Stiffness Composite Structures*. PhD Thesis, Technology University of Delft, Delft, 2011.
- [53] T. H. Walker, P. J. Minguet, B. W. Flynn, D. J. Carbery, G. D. Swanson, and L. B. Ilcewicz. *Advanced Technology Composite Fuselage-Structural Performance*. Technical report, The Boeing Company, Seattle, 1997.
- [54] K Wang, D Kelly, and S Dutton. Multi-objective optimisation of composite aerospace structures q. *Composite Structures*, page 8, 2002.



Detection of Blended Edges

Detection of butted edges is done by checking for each edge if it either continues in the next region or if the it is covered from above and below, such that there are not to stopped edges. If either of these cases exists that edge is called blended.

A.1. Detection of Blended Edges

Detection starts with the determining how layers can continue in the next via the interface connectivity(IC). The Ic can be computed based on the layer identification, either patch assignment or orientation) and provides an insight how layers can be connected across the interface. The interface connectivity is defined as:

$$C_{ij} = \begin{cases} 1 & \text{If layer } i \text{ of region A can continue into layer } j \text{ of region B} \\ 0 & \text{Otherwise} \end{cases} \quad (\text{A.1})$$

Given a geno $\theta_{A,i}$ and $\theta_{B,i}$ for region A and B respectively, the interface connectivity is computed by comparing all individual layers from A to all individual layers of region B as:

$$C_{ij} = [\theta_{A,i} == \theta_{B,j}] \quad (\text{A.2})$$

The formulation based on the patch id is comparable to the Equation A.2 but replaces the geno with the patch identification.

A layer i is blended if one of the following conditions is met:


$$\begin{aligned} 1. \quad & \exists j \quad \text{s.t.} \quad C_{ij} == 1, \\ & \quad \quad \quad j \in \{0, \dots, N-1\} \\ 2. \quad & \exists j, k, l \quad \text{s.t.} \quad C_{jl} == 1 \wedge C_{k(l+1)} == 1, \\ & \quad \quad \quad j \in \{0, \dots, i-1\}, \\ & \quad \quad \quad k \in \{i+1, \dots, N-1\}, \\ & \quad \quad \quad l \in \{0, \dots, N-2\} \end{aligned} \quad (\text{A.3})$$

where N is the number of layers.

Here the first case corresponds to a continuation of the layer in the next region. The second case searches for a set of connections that cover the edge from above and below.

As presented, Equation A.3 will mark inner and outer blended edge as butted. This can be circumvented by adding a artificial layer with identification 0 to the outsides of the layer, effectively introducing a straight connection above and below of the laminate. The equivalent IC matrix then becomes:

$$\begin{bmatrix} 1 & 0 & 0 \\ 0 & C_{i,j} & 0 \\ 0 & 0 & 1 \end{bmatrix}$$



Minimum Cutting Size

A new method for to comply with minimum cutting constraints at a local level was developed. The method was however not used because later stages in the algorithm would violate these constraints again and no practical solution for a combined algorithm was found. Hence this development has been added as an appendix.

B.1. Minimum Manufacturing Length

Since directional requirements add complexity to the already complex problem, the constraint is generalized to any direction. This allows for the formulating the constraint to be satisfied when both the size in x and y at any point in a patch are bigger than the minimum cutting length. Mathematically this constraint can be written as

$$\begin{aligned}\min(Size_x(x, y)) &\geq 100mm \\ \min(Size_y(x, y)) &\geq 100mm\end{aligned}\tag{B.1}$$

where $Size_x(x, y)$ and $Size_y(x, y)$ is a function over the domain of the laminate which gives the size of that part in x and y respectively.

The global constraint is satisfied when it is also satisfied for every local cell, or

$$\begin{aligned}Size_{x,i} &\geq 100mm \\ Size_{y,i} &\geq 100mm\end{aligned}\tag{B.2}$$

where $Size_{x,i}$ and $Size_{y,i}$ are the local path size for cell i . This allows a cellular automaton to evolve based on the patch size of a single cell.

Local size, as shown in Figure B.1, gives an indication about the size of the patch as observed from that cell. In and of itself it does not guarantee manufacturability of the path, but it does when it is satisfied for all cells in the patch. For example the patch highlighted in Figure B.1 is completely manufacturable in the center portion, but the top left corner only has an height of cell. That is too small to be manufactured albeit the other cells being manufacturable.

B.2. Neighbor length

For evolution purpose it is also valuable to understand the surrounding patch sizes and the composition of the own. To local patch size can be computed by the sum of the strings of similar cells attached to each side plus one. In other words the local size in x is the sum of the number of similar cells to the left, the number of similar cells to the right, and 1 (the own cell).

The length of the string of cells attached to one side is called the neighbor length (NL). For a simple problem it indicates a single length of the patch in that direction, whereas for complex problems the minimum length in that direction is used. The minimum length is most valuable because it indicates whether the minimum cut length is respected in that direction. This more complex is especially valuable




Figure B.1: The local size of a patch is the observed distance to the patch's sides.

where one side of a cell is split into multiple neighbors, which is the case in for example the horseshoe pattern.

The neighbor length for a simple plate is visualized in Figure B.2. To the side of each cell a single is shown which indicates the number of neighbors that has the same orientation as that cell. For example th




Figure B.2: The neighbor length (NL) of each cell is based on the length of the string of similar neighbors attached to that side and gives a measure of the remaining patch length in that direction.

Computation of the neighbor length is a straightforward iterative process which follows the following rule: If the neighbor is equal to the analyzed cell, the neighbor length at that side is equal to opposite NL of that neighbor + 1. Or in mathematical terms

$$\begin{aligned}
 NL_{west,i} &= 1 + NL_{west,i:west} \\
 NL_{north,i} &= 1 + NL_{north,i:north} \\
 NL_{east,i} &= 1 + NL_{east,i:east} \\
 NL_{south,i} &= 1 + NL_{south,i:south}
 \end{aligned} \tag{B.3}$$

where $NL_{direction,i}$ indicates the neighbor length of the i th cell on the left, top, right, and bottom side

respectively, and $NL_{direction,i:direction}$ indicates the neighbor length on side $direction$ of the neighbor located at side $direction$.

Equation B.3 is used repeatedly until the neighbor length is settled for all the cells. The iterative process is bound because ultimately every patch has a bound and after enough iterations the NL has crossed the entire patch. For example a patch of two cells wide has the neighbor length of cells settled after just two iterations.

The local patch size can now easily be computed as a function of the neighbor length of the current cell by

$$\begin{aligned} Size_x &= 1 + NL_{west,i} + NL_{east,i} \\ Size_y &= 1 + NL_{north,i} + NL_{south,i} \end{aligned} \quad (B.4)$$

B.3. Local Minimum Cut Length Compliance Procedure

During the second step in the blending procedure, patches should be formed that satisfy the minimum cut length constraints. Whereas in the previous step neighbors were not directly influenced, they now are because changing the orientation of the current cells allows the constraints of the neighboring cells to be improved or degraded. Furthermore, some local situation require the change of several neighbors to comply with the minimum patch size.

Even more so, in some cases the global change is less when a neighbor changes angle instead of the cell that is actually too small. The best example for this case is the grid shown Figure B.1. If every cell should only look at improving its own manufacturability the entire grid would become green. This is easiest envisioned when looking at the center green cells; They will not change because they all satisfy the minimum cut length of two. Therefore all neighbors can only comply with the constraints by also orientating in the greens' direction.

Because by changing the ply angle of the current cell, the constraints of a neighbor can also improved, its condition is also taken into account. To this end, the update procedure uses a second order evolution to ensure constraints of neighbors are not improved simultaneously by cells evolving at same time. This provides stability to the algorithm and prevents jumping between solutions because of cells alternating their solution.

During the update procedure every layer of the analyzed cell is update to the orientation that is deemed most fit. To determine the most fit orientation, all possible orientation are given a score, where the lowest score is deemed most fit for that layer.

Scoring of the layer is primarily focused on reducing the number of violated constraints. Hence the total number of constraints satisfied and those that can be influenced are computed. Here a split is made between constraints that matter and can be influenced, and once that do not matter or cannot be influenced by the analysed cell.

The latter states that constraints that cannot be improved or degraded by the current cell are not taken into account. Neighbors that are attached to the east and west sides cannot have their size in y changed by the analysed cell, because that is on the cross axis. Therefore the analysed cell cannot improve any minimum cut length constraint in that direction and the whole constraint is neglected.

The importance of the neighbor is a more complex phenomenon and indicates the likelihood of that cell evolving regardless of the orientation of the analysed cell. For example a neighbor to the west that is too small in the y direction is likely to change regardless of the configuration chosen for the analysed cell. The analysed cell cannot improve nor degrade the state of the neighbor and therefore the constraints for that neighbor or not taken into consideration. If the neighbor remains in that configuration, the neighborhood of that neighbor shall eventually become stable. This will ultimately result in the compliance of the neighbors' constraints which in term makes the neighbor important during further iterations.

Every iteration the score for each orientation possibility is computed according to the following rules:

- Every important influenceable minimum patch size constraint not satisfied is penalized by one.
- The orientation difference from the current state angle is penalized by 0.1, making it less important than the number of constraints.

- For every orientation already present in one of the four neighbors 0.01 is subtracted. This ensures that orientations that are more common in the neighborhood are preferred when the previous two constraints did not yield a conclusion.
- A penalty of 0.1 is added for every time the orientation was also considered during a previous orientation. This ensures that when two orientation are at tie break when considering their improvement, the algorithm will ultimately decide to try an angle that is further away with the assumption that the neighborhood will help out. Consequently it pushes the global algorithm to search for a solution that is not stuck in a local optimum.

One could argue that using the string length makes this an exponential problem, but a shortcut can be taken to overcome this. All the rules used in evolution process only look at the neighbor size being less than the minimum length, see Equation B.2. Since the maximum size is ultimately one more than the biggest neighbor length, ref Equation B.4, the computation of the neighbor length does not have to be more precise than just above the minimum cut length. This means that the maximum neighbor length that has to be computed is equal to the minimum cut length minus one. In other words the computation stated in Equation B.3 only needs as many iterations needed to guarantee compliance with the minimum patch size.

B.4. Results

Figure B.3 shows preliminary before application of the minimum cut compliance stage of the cellular automaton (CA). Several patches are only one region in length or width and are in this example too small to be manufactured. Minimum cut constraints prescribe a minimum size of 2 regions for both direction in this example.




Figure B.3: Patch design before the minimum cut compliance procedure. Several patches have local dimension of only one region what is deemed too small to manufacture.

After application of the cellular automaton's minimum cut compliance stage all the patch comply with the minimum cutting size. The result is shown in Figure B.4. Through the local size the algorithm has determined that every region now complies with minimum cutting size constraints in both directions.

However, during the butted removal process regions are processed individually and minimum cutting constraints are neglected. This allows the CA to find a solution that is conform the blending guidelines




Figure B.4: Patch design directly after application of the minimum cut fase of the CA. All patches have a local patch size greater or equal than 2 adjacent regions.

but violates the minimum cutting constraints in many regions. Figure B.5 shows the design after butted edge removal; every region that violates the minimum cutting size constraints is marked with a star.

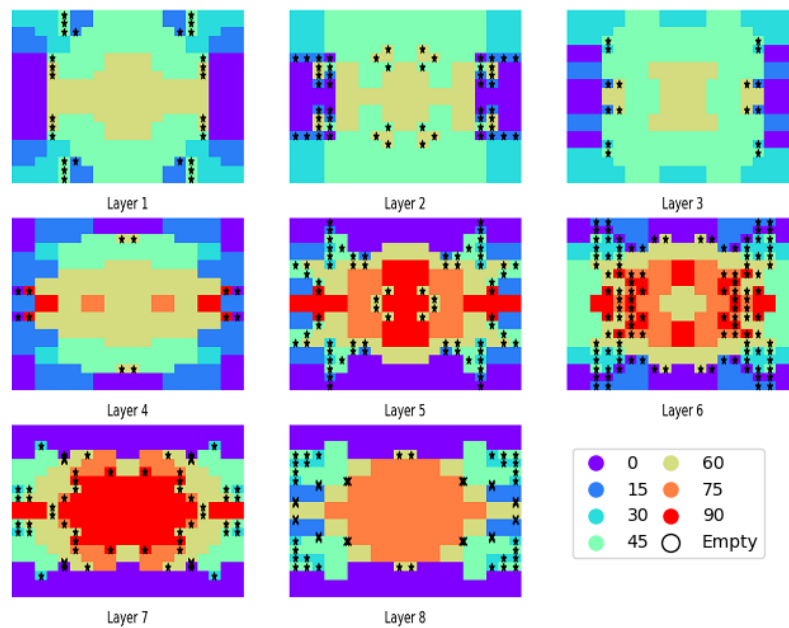



Figure B.5: Patch design directly after elimination of the butted. All region marked with a star do not comply with the minimum cut length of 2 regions. The elimination fase has undone most of the work of the minimum cutting constraints fase.

An integrated approach is needed to produce designs that comply with the blending guidelines and minimum cutting constraints. Additionally, more research is needed to create a suitable method that also tracks cutting size over multiple levels and ensures compliance after global interpretation.



Additional Design Results

C.1. Patch Metrics of the 5xM Plate

These patch metric show a similar story as the patch metrics of the MxM plate. Average aspect ratio is 1.2 and remains constant when the number of regions is increased.




Figure C.1: Pareto distribution considering load and mean normalized aspect ratio




Figure C.2: Pareto distribution considering load and shape complexity

C.2. Results for the Mx5 Plate

Optimizations for a varying number of regions along the loading direction were also performed. As reported in literature, increasing the number of regions in this direction did not yield significant improvement. This is most notable by the absence of improvement between 5x5 and 7x5 sections.




Figure C.3: Buckling performance after each design stage




Figure C.4: Pareto distribution of patches versus buckling load

C.3. Global Interpretation Tree of Cylinder Results

A tree of the global interpretation search of the cylinder has also been created. It was not added to the report because the global interpretation does not involve solving order problems and is only a 2D problem.

C.4. Cost and Weight of cylinders

Figure C.6 and C.7 contain the weight and cost analysis for the cylinders.




Figure C.5: Branches that lead to a solution in the global interpretation search, where the x axis shows the difference and y shows the iteration increments. The inset graph shows that in the case of the cylinder the change is almost unrelated to the change in load.




Figure C.6: Weight savings that can be achieved using SFVS laminates in cylinders.




Figure C.7: Effect of SFVS laminates on the direct operating cost for airlines.

D

Determination of the Number of Samples

Different sampling quantities were tried to investigate the influence on the optimization convergence of the 5x5 plate. The initial number of samples was changed from 36 to 108 in increments of 18, and the resample count ranged from 9 till 36 in steps of 9. The results are shown in Figure D.1

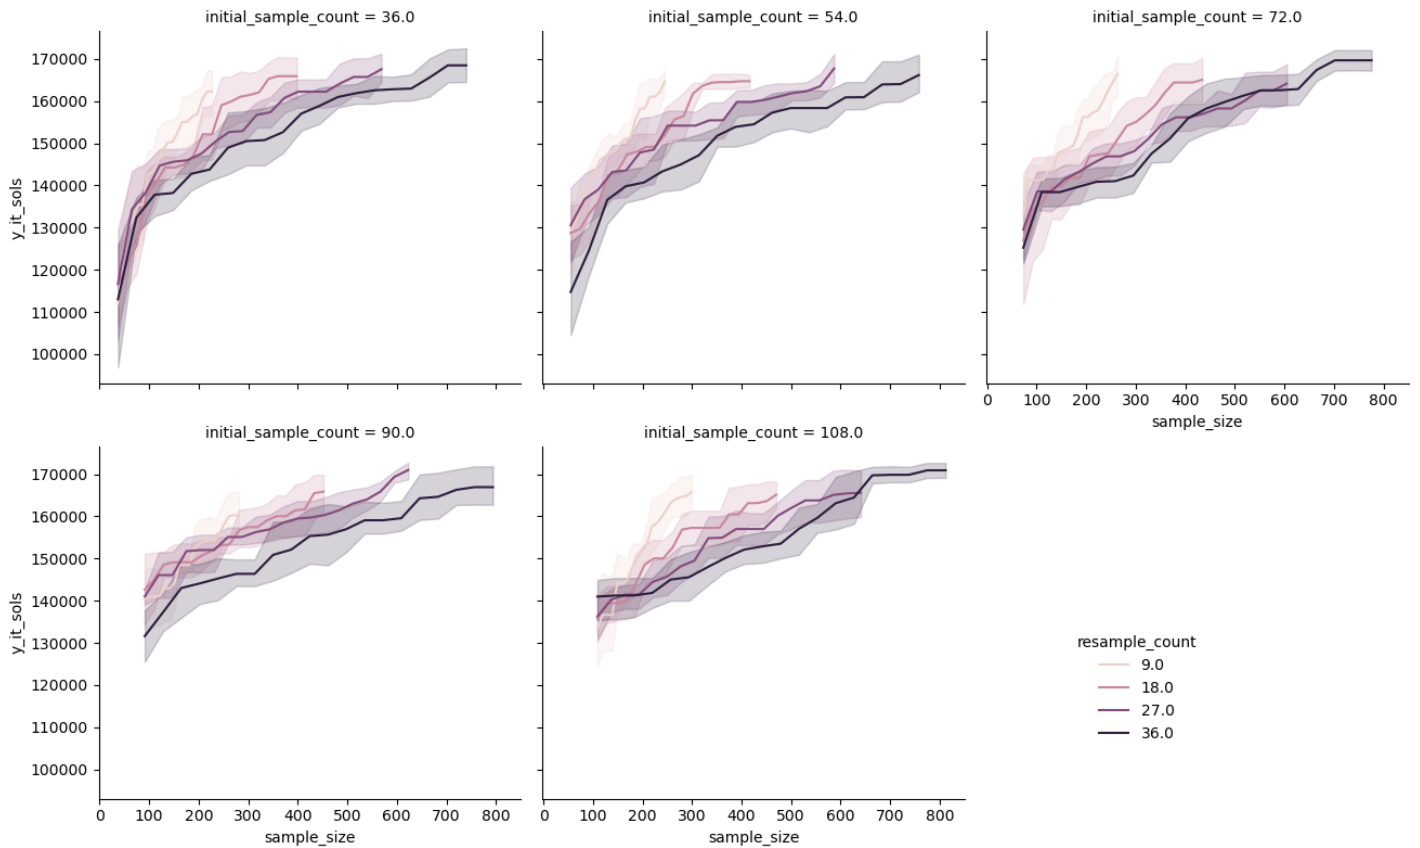



Figure D.1: Plots for different resampling size and different initial sample counts. The number of samples simulated using FEA is shown on the x-axis, and the optimized load is shown on the y-axis.

A lower number of samples per iterations results in convergence to a buckling load at a lower total number of samples. This effect can be seen in each of the initial sample count settings. Moreover, the initial sample setting count only seems to affect the initial convergence. Subsequent iterations benefit from more resampling and closing of the sample bands.



Cylinder Optimization Convergence

The optimization of the smallest cylinder has been performed several times using different mesh sizes to study the effect of the mesh on the obtained local stiffness optimum. The finest mesh size was used to simulate the end result of each optimization again and judge the quality of the optimization. The results are shown in Figure E.1, where the refined load yields the simulation of the end result using the finest mesh. The optimization prediction was the load output of the optimizer. The shaded area shows the spread in the results.




Figure E.1: Convergence study on the optimization of the cylinder with 8 regions of which 5 unique.

For a finer mesh the solutions do converge faster. The spread for the small mesh is over 60 kNm. And also at 5000 elements, the setting used for all of the cylinders reported, the spread is still 50 kNm. It can be concluded that a finer mesh might lead to a better convergence of the optimization because the finer mesh has a spread of roughly 10 kNm. This does come at a time an increased computation cost. Optimization using the largest mesh takes more than three hours, where the finer ones take approximately 1. Parallel execution of the FE analysis might improve this time. In the end all optimization were performed with roughly 5000 elements because the time is significantly less and good results can still be obtained.