

Toolpath Generation for Fused Filament Fabrication of Functionally Graded Materials

Kuipers, T.

DOI

[10.4233/uuid:7b2ddf0f-3e0f-49d8-b425-b6df0b540e32](https://doi.org/10.4233/uuid:7b2ddf0f-3e0f-49d8-b425-b6df0b540e32)

Publication date

2022

Document Version

Final published version

Citation (APA)

Kuipers, T. (2022). *Toolpath Generation for Fused Filament Fabrication of Functionally Graded Materials*. [Dissertation (TU Delft), Delft University of Technology]. <https://doi.org/10.4233/uuid:7b2ddf0f-3e0f-49d8-b425-b6df0b540e32>

Important note

To cite this publication, please use the final published version (if applicable). Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights. We will remove access to the work immediately and investigate your claim.

Toolpath Generation
for
Fused Filament Fabrication
of
Functionally Graded Materials

Dissertation

for the purpose of obtaining the degree of doctor
at Delft University of Technology
by the authority of the Rector Magnificus prof. dr. ir. T.H.J.J. van der Hagen
chair of the Board for Doctorates
to be defended publicly on
Thursday 14 April 2022 at 12:30 o'clock

by

Tim KUIPERS

Master of Science in Computer Science,
Utrecht University, The Netherlands
born in Eindhoven, The Netherlands

This dissertation has been approved by the promotor.

Composition of the doctoral committee:

Rector Magnificus,	chairperson
Prof. dr. C.C.L. Wang,	Delft University of Technology, promotor
Dr. J. Wu,	Delft University of Technology, copromotor

Independent members:

Prof. dr. ir. K.M.B. Jansen,	Delft University of Technology
Prof. dr. S.C. Pont,	Delft University of Technology
Prof. dr. I. Gibson,	University of Twente
Dr. S. Lefebvre,	Université de Lorraine
Dr. K. Masania,	Delft University of Technology
Prof. dr. P. Vink,	Delft University of Technology, reserve member

Ultimaker

This research was funded by Ultimaker B.V.

Keywords: Functionally graded materials, Mechanical metamaterials, Cellular materials, Lattice structures, Toolpath generation, Additive manufacturing, Fused Deposition Modeling, Fused Filament Fabrication, Material Extrusion

Printing: Ridderprint, www.ridderprint.nl

ISBN 978-94-6458-132-4

Copyright © 2022 by T. Kuipers. All rights reserved. No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by means, without prior written permission of the author.

An electronic version of this dissertation is available at <http://repository.tudelft.nl/>.

TABLE OF CONTENTS

Glossary	v
Summary	vii
Samenvatting	xi
I Introduction	1
I.1 Functionally graded materials	2
I.2 Fused Filament Fabrication Technology	3
I.3 Research objective	7
I.4 Dissertation organization	7
II Spatially graded grayscale tone	11
II.1 Introduction	12
II.2 Background	14
II.3 Method	17
II.4 Results	25
II.5 Discussion	27
II.6 Conclusion and future work	31
III Spatially graded compliance	35
III.1 Introduction	36
III.2 Related work	38
III.3 Overview	39
III.4 CrossFill	41
III.5 Adaptive subdivision	48
III.6 Toolpath generation	53
III.7 Results	55
III.8 Conclusion and future work	63
IV Enabling spatially graded thickness	69
IV.1 Introduction	70
IV.2 Related Work	72
IV.3 Method	73
IV.4 Beading schemes	84
IV.5 Fabrication	88
IV.6 Results and discussion	92
IV.7 Conclusion	97
IV.A Edge discretization	102
IV.B Data set	103
IV.C Accuracy calculation	104
V Enabling dual-material structures	105
V.1 Introduction	106
V.2 Method	110
V.3 Results	118
V.4 Discussion	125

V.5	Conclusion	129
VI	Applications	135
VI.1	Spatially graded roughness	136
VI.2	Spatially graded density in 2D	138
VI.3	Spatially graded thickness	140
VI.4	Spatially graded dual material proportions	143
VII	Conclusion	147
VII.1	Research objective realization	148
VII.2	Implications	148
VII.3	Future outlook	155
	Acknowledgments	157
	Curriculum Vitæ	159
	Publications	161

GLOSSARY

ABBREVIATIONS

3D	Three-dimensional
AM	Additive Manufacturing
CAD	Computer-Aided Design
FDM	Fused Deposition Modeling
FFF	Fused Filament Fabrication
PLA	Polylactic Acid
PP	Polypropylene
TPU	Thermoplastic Polyurethane
TPMS	Triply-Periodic Minimal Surface
FGM	Functionally Graded Material

DEFINITIONS

3D Printing See *Additive Manufacturing*

Additive Manufacturing (AM) A method for manufacturing three-dimensional objects by successively adding material

Build plate The bottom plate onto which the first layer is fastened

Compliance The inverse of stiffness

Feeder The mechanism by which the filament is fed to the nozzle

Filament A thread of thermoplastic material used as stock for Fused Deposition Modeling.

Functionally Graded Material (FGM) A substrate with a *Spatially graded material property*

Fused Deposition Modeling (FDM) An additive manufacturing technology based on the extrusion of a thermoplastic filament

Fused Filament Fabrication (FFF) See *Fused Deposition Modeling*

Gcode A common file format which encodes the printing process: toolpaths, nozzle temperature, fan speed, etc.

Lattice See *Microstructure*

Macro-scale The order of magnitude of the printed product – typically in the order of decimeters

Meso-scale An order of magnitude in between the *micro-scale* and the *macro-scale* – typically in the order of centimeters

Micro-scale The order of magnitude of the printing resolution – typically in the order of millimeters

Microstructure A (semi-)repeating geometry at a scale in between the nozzle size and the manufactured product

Nozzle The tip of the print head from which the thermoplastic filament is extruded

Nozzle size The diameter of the hole in the tip of the nozzle

Print bed See *Build plate*

Retraction A backwards move of the filament in order to prevent the nozzle from leaking when traveling in between extrusion lines

Spatially graded material property A visual or mechanical material property which varies from location to location in the manufactured product

Toolpaths The paths traversed by the print head and the amount of material extruded along those paths

Triply-Periodic Minimal Surface A microstructure consisting of a smooth surface with locally minimal surface area which connects to instances of itself in the 3 directions of 3D space.

SUMMARY

The products in our day to day lives have different requirements in different regions of the product. One way of dealing with such spatially varying requirements is to divide the product into multiple parts and assign each part a different material. For example, the handle of a drill is often fitted with a rubbery material which gives it more grip when holding it. This gives the designer a choice of which regions exactly to use that material. Designers can consider the large variety of hands and ways of holding a drill. Equipped with an average distribution of gripping pressure throughout the handle, they could identify the regions with a pressure higher than some cutoff value and assign the gripping material to those regions. This means that the material properties vary abruptly over the surface; the adjacent regions on the one side of the cutoff value provide needlessly much grip, while on the other side the regions provide too little grip. The final distribution of material properties is *segmented* and therefore does not follow the *gradual* distribution of requirements optimally. In this workflow the designer is forced to reduce the continuous gripping pressure information into a binary material choice.

But what if we could manufacture products with a gradient in their material properties? To answer this question we consider the material and the manufacturing technique.

A Functionally Graded Material (FGM) is any substrate with material properties made to vary from region to region. FGMs find application for example in personalized footwear, implants, tires and airplane wings. They can improve a product's performance by optimizing the spatial gradation of material properties throughout the product. Rather than a homogeneous block of material, FGMs consist of a fine-scale geometry of one or more base materials. The material properties of an FGM can be governed by controlling the shape of that fine-scale structure.

Fused Filament Fabrication (FFF) is an additive manufacturing technique which can produce complex geometry cheaply. Thermoplastic material is heated and extruded out of a nozzle to deposit extrusion lines. These extrusion lines accumulate to form layers, which are added on top of each other to form the final product. A 3D model is converted into toolpaths for the 3D printer, which describe the geometry of the extrusion lines the nozzle should traverse. Because of the physics involved in such machines, there are several manufacturing constraints to which the print job must adhere, such as (i) the maximum overhang angle to prevent printing in mid air, (ii) (semi-)continuous extrusion to prevent print defects at the ends of extrusion lines, (iii) integer thickness geometry ($N * \text{linewidth}$) to prevent overlapping extrusion lines, and (iv) chemically compatible materials to prevent a multi-material print job from disassembling during the manufacturing process.

These manufacturing constraints hamper the development of the FFF production of FGMs. If FFF technology could be used to manufacture FGMs, then designers and researchers with access to such additive manufacturing systems could design more complex products with improved performance. Conversely, if we were able to manufacture FGMs by using FFF technology, then FGMs can be produced more cheaply and with a wider

range of materials, unlocking a wider range of applications. However, the above manufacturing constraints mean that developing FGMs for FFF is a complex task.

How can we generate toolpaths for fused filament fabrication 3D printers which enable functionally graded materials? In this dissertation I present the results of four research cycles, which answer this research question from different perspectives. The first two research cycles propose FGMs which are generated by making the fine-scale geometry adhere to the manufacturing constraints. The last two research cycles propose toolpath generation algorithms which eliminate manufacturing constraints, making existing FGMs available for FFF. Moreover, the research cycles differ in that they relate to the spatial gradation of different material properties: surface properties such as color or volumetric properties such as stiffness. Furthermore, two of the research cycles relate to combining different materials, while the other two are aimed at properties of single material parts.

The first research cycle proposes an FGM which controls surface properties of the product, such as the (grayscale) color and the surface roughness. These properties can be modulated by careful manipulation of the deposition lines which follow the layer contours. The perceived grayscale tone can be controlled by adjusting the contours of layers when those layers are printed with a black and a white material alternatingly. Surface roughness can be controlled by adjusting the contour-following lines of a single material print. Such properties can be used for visual purposes such as aesthetics and presenting information. Surface roughness can furthermore be used for its effect on gripping.

The second research cycle proposes a single material FGM called 'CrossFill'. CrossFill is a foam structure with spatially graded density. The structure adheres to the overhang constraint and consists of constant-width strictly continuous extrusion lines. Such foams can be printed with flexible material to allow for a spatially graded flexibility, which is interesting for applications such as a bike saddle. A 2D variant of this foam structure can be used for creating grayscale images, which can be laser engraved, plotted or 3D printed on textile.

The third research cycle proposes a toolpath generation framework which generates toolpaths to fill layer contours with a varying thickness. Varying thickness geometry is filled using a varying number of toolpaths with varying extrusion width. I also present a control mechanism to achieve such extrusion width variation. As such the manufacturing constraint of integer thickness is alleviated. Hence a class of FGMs becomes FFF manufacturable, which spatially varies the stiffness by modulating the thickness of a regular repeating cell geometry. Furthermore, I show how a spatially graded thickness can be used to control the translucency of a thin shelled object, in order to manufacture custom lamp shades.

The last research cycle proposes a mechanism for interlocking two materials using a repeating lattice structure at the interface between the two bodies. By securing the bonding between two bodies mechanically, the chemical compatibility constraint is alleviated. The lattice facilitates continuous extrusion and can be applied for interfaces in any direction. The lattice can be used to define multi-scale FGMs, where the interlocking lattice ensures the bonding at a fine scale, while the meso-scale geometry of the two materials is used to determine the spatial gradation in material properties. Furthermore, the interlocking lattice can be made to fill the volume, so that material properties can be spatially graded by controlling the lattice geometry directly.

With FGMs we can design products where different regions are optimized for different requirements. This dissertation presents several new FGMs for FFF and unlocks FFF

printing of a wide range of existing FGMs. Unlocking FGMs for FFF technology makes these advanced materials more accessible in terms of manufacturing costs and material choice versatility. I present algorithmic frameworks which automatically determine the fine-scale structure of 3D prints based on a spatial requirement specification. The relations between the fine-scale structure and its emerging material properties are verified experimentally using commercially available FFF systems. The FGM generation frameworks have the potential to revolutionize the product development process and improve product performance. Furthermore, the presented methods are beneficial to FFF technology in general, making it compatible with a wider range of 3D models and unlocking new applications. This dissertation establishes the groundwork for moving towards a future in which FFF is used to produce everyday products exhibiting spatially graded properties with improved performance.

SAMENVATTING

De producten in ons dagelijks leven hebben verschillende vereisten in verschillende delen van het product. Een van de manieren om met zulke ruimtelijk gevarieerde vereisten om te gaan is door het product op te breken in meerdere onderdelen en elk onderdeel een verschillend materiaal toe te wijzen. Zo wordt bijvoorbeeld het handvat van een boor vaak uitgerust met een rubberachtig materiaal, wat het meer grip geeft om vast te houden. Dit geeft de ontwerper een keuze over welke gebieden precies uitgerust moeten worden met dat materiaal. Ontwerpers kunnen rekening houden met de grote verscheidenheid aan handen en manieren om een boor vast te houden. Gegeven een gemiddelde grijpdrukverdeling over het handvat, kunnen ze een gebied afbakenen met een grijpdruk hoger dan een gegeven grenswaarde en dat gebied uitrusten met het grijpmateriaal. Dit betekent dat de materiaaleigenschappen abrupt veranderen over het oppervlak; de aangelegen gebieden aan de ene kant van de grenswaarde geven overbodig veel grip, terwijl aan de andere kant de gebieden te weinig grip leveren. De uiteindelijk verdeling van materiaaleigenschappen is *gesegmenteerd* en volgt daarom de *graduele* verdeling van vereisten niet optimaal. In deze werkwijze wordt de ontwerper gedwongen om continue gegradeerde grijpdrukgegevens te reduceren tot een binaire materiaalkeuze.

Maar wat als we producten zouden kunnen fabriceren die een ruimtelijke gradiënt in hun materiaaleigenschappen hebben? Om deze vraag te beantwoorden beschouwen we het materiaal en het fabricageproces.

Een Functioneel Gegradeerd Materiaal (FGM) is welk substraat dan ook met aangebrachte materiaaleigenschappen die van plek tot plek verschillen. FGMs hebben toepassingsmogelijkheden in bijvoorbeeld gepersonaliseerd schoeisel, implantaten, banden en vliegtuigvleugels. Ze kunnen de prestatie van een product verbeteren, door de ruimtelijke gradatie van materiaaleigenschappen te optimaliseren. In plaats van een homogeen blok materiaal, hebben FGMs een fijnmazige structuur van een of meerdere basismaterialen. De materiaaleigenschappen van een FGM kunnen worden gestuurd door de vorm van die fijnmazige structuur te manipuleren.

Filament Fusie Fabricage (FFF) is een additieve fabricagetechniek die goedkoop complexe vormen kan produceren. Thermoplastisch materiaal wordt verhit en ge-extrudeerd uit een mondstuk om extrusielijnen te deponeren. Deze extrusielijnen accumuleren tot lagen, die op elkaar gestapeld worden om het uiteindelijke product te vormen. Een 3D model wordt omgezet in printbanen (toolpaths) voor de 3D printer, die de geometrie beschrijven van de paden die het mondstuk moet afleggen. Vanwege de fysica die bij zulke apparaten komt kijken zijn er enkele fabricagerestricties waar de printopdracht aan moet voldoen, zoals (i) de maximale hellingshoek om te voorkomen dat er midden in de lucht geprint wordt, (ii) (semi-)continue extrusie om defecten aan de uiteindes van printbanen te voorkomen, (iii) gehele waarde dikte geometrie ($N * \text{lijnbreedte}$) om overlappende printbanen te voorkomen, en (iv) chemische compatibiliteit om te voorkomen dat een printopdracht met meerdere materialen uit elkaar valt tijdens het fabricageproces.

Deze fabricagerestricties hinderen de ontwikkeling van het met FFF produceren van

FGMs. Als FFF technologie gebruikt kan worden om FGMs mee te fabriceren, dan kunnen ontwerpers en onderzoekers, die toegang hebben tot een dergelijk additief fabricagesysteem, complexere producten vormgeven met verbeterde prestaties. Omgekeerd, als het mogelijk zou zijn om FGMs te fabriceren met FFF-technologie, dan kunnen FGMs goedkoper geproduceerd worden en met een breder scala aan materialen, waardoor een grotere verscheidenheid aan toepassingen mogelijk wordt. Nochtans betekenen de bovenstaande fabricagerestricties dat het ontwikkelen van FGMs voor FFF een complexe taak is.

Hoe kunnen we printbanen voor filament fusie fabricage 3D printers genereren die functioneel gegradeerde materialen mogelijk maken? In dit proefschrift presenteer ik vier onderzoekscycli, die deze onderzoeksvraag vanuit verschillende perspectieven benaderen. De eerste twee onderzoekscycli opperen FGMs waarbij de kleinschalige geometrie zodanig wordt aangepast dat aan de fabricagerestricties voldaan wordt. De laatste twee onderzoekscycli opperen algoritmes voor het genereren van printbanen die fabricagerestrictie elimineren, zodat bestaande FGMs beschikbaar komen voor FFF. Daarnaast verschillen de onderzoekscycli in termen van de materiaaleigenschappen die ruimtelijk gevarieerd worden: oppervlakte-eigenschappen zoals kleur of inhoudelijke eigenschappen zoals stijfheid. Bovendien zijn twee van de onderzoekscycli gerelateerd aan het combineren van verschillende materialen, terwijl de andere twee zich richten op de eigenschappen van structuren van een enkel materiaal.

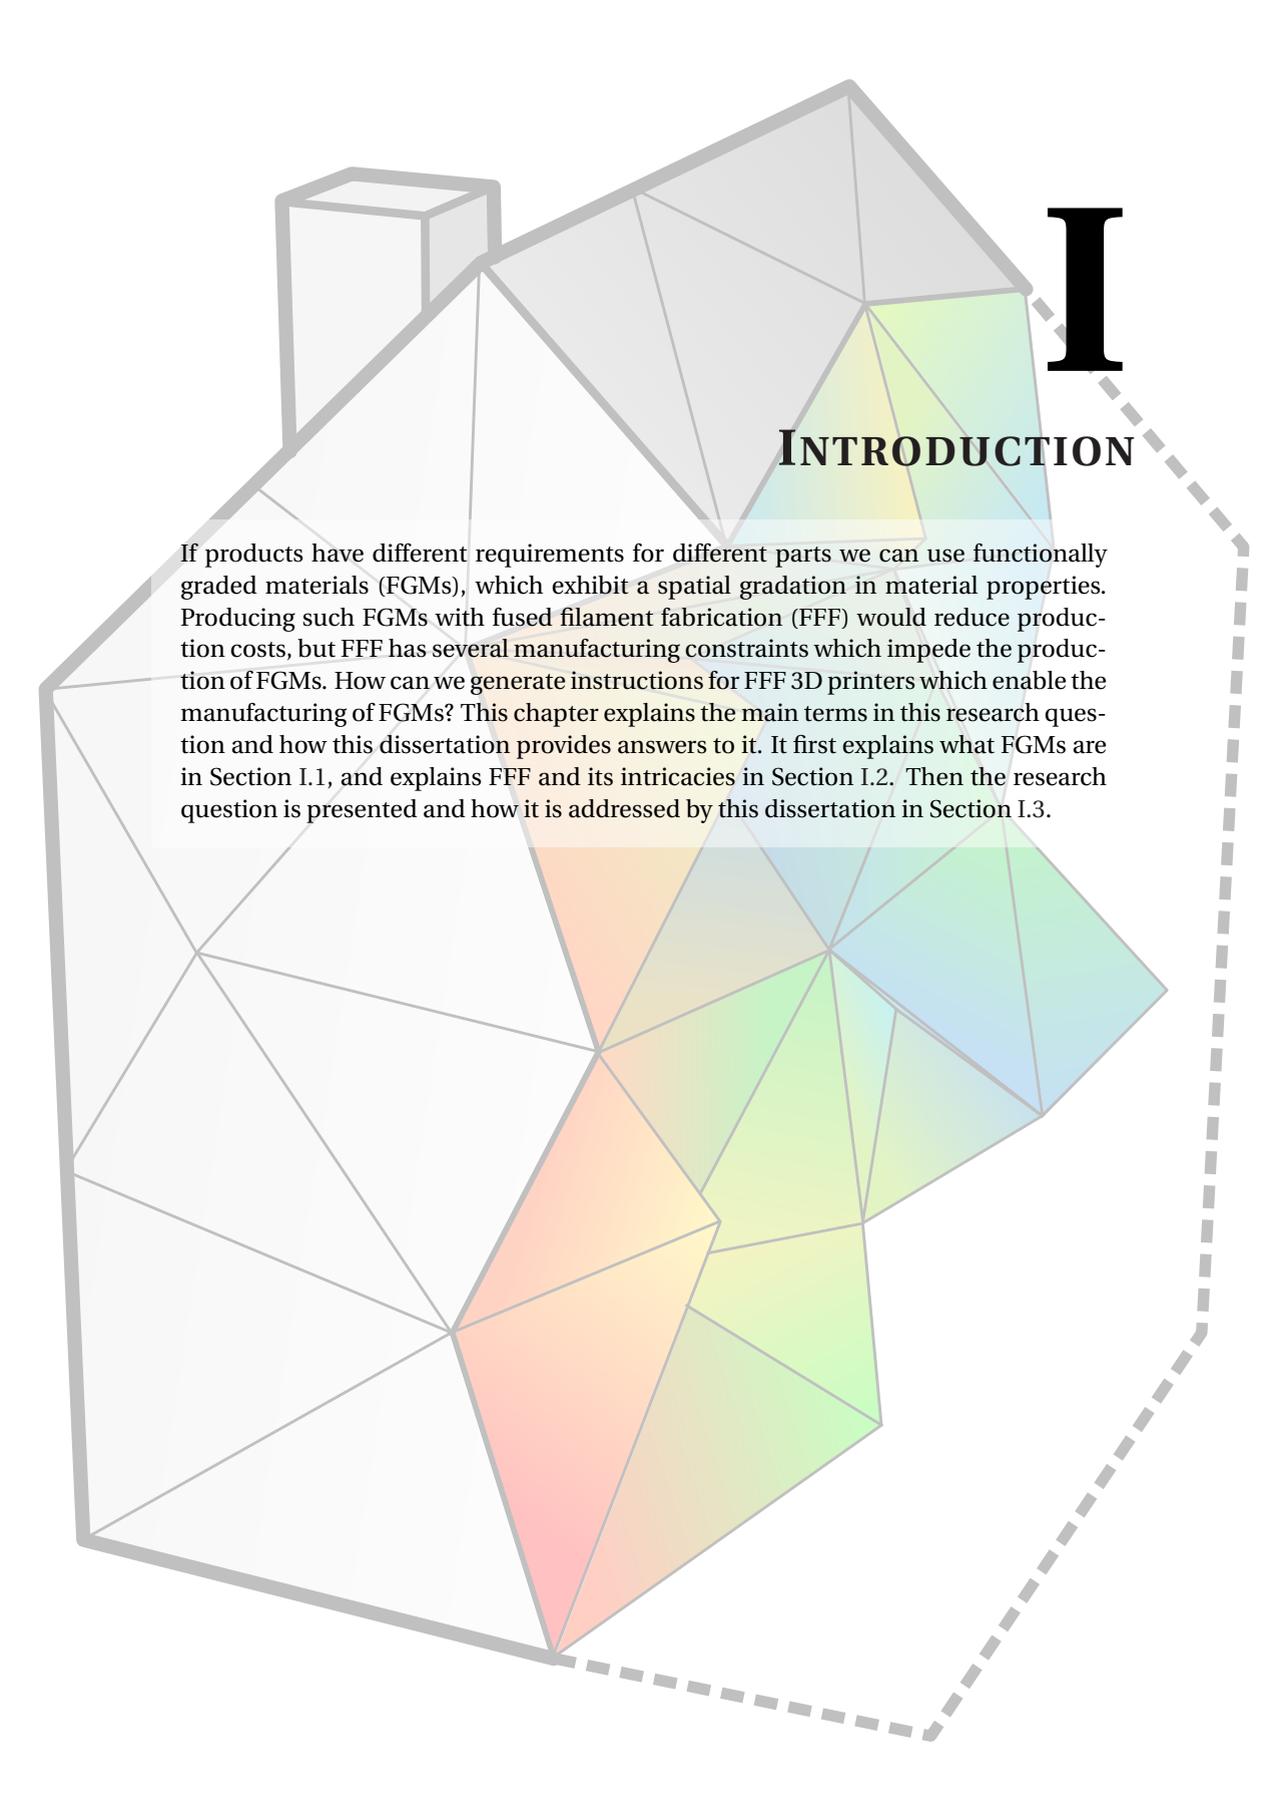
De eerste onderzoekscyclus stelt een FGM voor die oppervlakte-eigenschappen van het product sturen, zoals grijswaarde en ruwheid. Deze eigenschappen kunnen worden gestuurd door weloverwogen manipulatie van de extrusielijnen in de buitenste contouren van elke laag. De waargenomen grijswaarde kan gemoduleerd worden door de contour aan te passen van de lagen, wanneer die lagen afwisselen met een zwart en wit materiaal geprint worden. Oppervlakteruwheid kan gestuurd worden door aanpassing van de extrusielijnen die de contour volgen van een object dat met een enkel materiaal geprint wordt. Zulke eigenschappen kunnen worden gebruikt voor visuele doeleinden zoals esthetiek en het presenteren van informatie. Oppervlakteruwheid kan bovendien gebruikt worden voor grip-oppervlaktes.

De tweede onderzoekscyclus presenteert een FGM dat geprint kan worden met een enkel materiaal, genaamd 'CrossFill'. CrossFill is een schuimstructuur met een ruimtelijke dichtheidsgradiënt. De structuur voldoet aan de hellingshoekrestrictie en bestaat uit een strikt continue extrusielijn met constante breedte. Zulke schuimen kunnen geprint worden met flexibele materialen, om een ruimtelijk gevarieerde flexibiliteit te vertonen, wat interessant is voor toepassingen zoals bijvoorbeeld een fietszadel. Een 2D variant van dit schuim kan bovendien gebruikt worden om grijswaarde afbeeldingen mee te genereren, die vervolgens kunnen worden lasergesneden, geplot of geprint op textiel.

De derde onderzoekscyclus legt een raamwerk voor om printbanen te berekenen die laagcontouren met variërende dikte opvult. Geometrie met variërende dikte wordt opgevuld met een variërend aantal printbanen met variërende extrusiebreedte. Daarnaast presenteer ik een stuurmechanisme om zulke extrusiebreedtevariatie te verwezenlijken. Als zodanig wordt de fabricagerestrictie van gehele waarde diktes opgeheven. Dit betekent dat een klasse van FGMs hiermee beschikbaar komt voor FFF, waarbij de stijfheid ruimtelijk gevarieerd wordt door de dikte te moduleren van een regelmatig herhalende cel vorm. Bovendien laat ik zien hoe een ruimtelijk gegradeerde dikte gebruikt kan worden om de doorschijnendheid van dun gehulde objecten te sturen, om bijzondere lampenkappen te fabriceren.

De laatste onderzoekscyclus presenteert een mechanisme om twee materialen met elkaar te verbinden door middel van een regelmatig raster rondom het raakvlak tussen de twee lichamen. Door de binding tussen twee lichamen mechanisch te verwezenlijken wordt de fabricagerestrictie van chemische compatibiliteit opgeheven. Het raster faciliteert continue extrusie en kan toegepast worden op raakvlakken in elke richting. Het raster kan worden gebruikt om multischaal FGMs te definiëren, waarbij het verbindingsraster de hechting waarborgt op een fijne schaal, terwijl de mesoschaal geometrie van de twee materialen gebruikt wordt om de ruimtelijke gradatie in materiaaleigenschappen te bewerkstelligen. Bovendien kan het vergrendelraster gebruikt worden om het gehele volume op te vullen, zodat materiaaleigenschappen ruimtelijk kunnen worden gegradieerd door direct de rastergeometrie aan te passen.

Met FGMs kunnen we producten ontwerpen waarbij verschillende delen geoptimaliseerd zijn voor verschillende vereisten. Dit proefschrift presenteert verscheidene nieuwe FGMs voor FFF en faciliteert het FFF-printen van een breed aanbod aan bestaande FGMs. Door FGMs te ontgrendelen voor FFF-technologie worden deze geavanceerde materialen toegankelijker in termen van productiekosten en de veelzijdigheid aan materiaalkeuze. Ik presenteer algoritmische raamwerken die automatisch bepalen wat de kleinschalige opbouw van een 3D print moet zijn gebaseerd op een ruimtelijk gegradieerde vereistespecificatie. De wisselwerking tussen de fijnmazige structuren en hun emergente materiaaleigenschappen zijn experimenteel gecontroleerd met behulp van commercieel beschikbare FFF-systemen. De voorgestelde FGM-omgeving heeft de mogelijkheid om een revolutie teweeg te brengen in het productontwikkelingsproces en de productprestaties te verbeteren. Ten slotte bevorderen de gepresenteerde technieken de FFF-technologie in het algemeen, waardoor FFF een breder scala aan 3D modellen ondersteunt en nieuwe toepassingen mogelijk worden gemaakt. Dit proefschrift legt de basis voor een toekomst waarin FFF wordt gebruikt om alledaagse producten te produceren die ruimtelijk gegradieerde eigenschappen vertonen met verbeterde prestaties.



I

INTRODUCTION

If products have different requirements for different parts we can use functionally graded materials (FGMs), which exhibit a spatial gradation in material properties. Producing such FGMs with fused filament fabrication (FFF) would reduce production costs, but FFF has several manufacturing constraints which impede the production of FGMs. How can we generate instructions for FFF 3D printers which enable the manufacturing of FGMs? This chapter explains the main terms in this research question and how this dissertation provides answers to it. It first explains what FGMs are in Section I.1, and explains FFF and its intricacies in Section I.2. Then the research question is presented and how it is addressed by this dissertation in Section I.3.

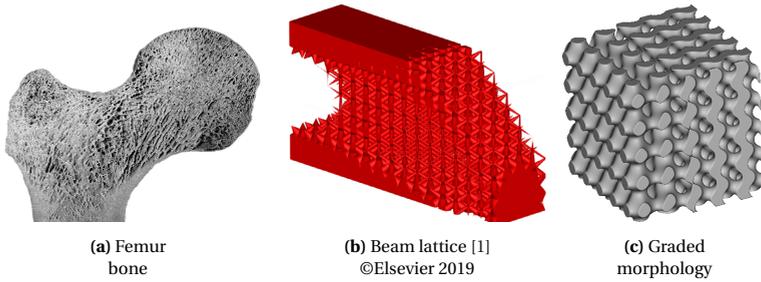


Fig. 1.1. Examples of functionally graded materials

I.1 FUNCTIONALLY GRADED MATERIALS

Functionally graded materials (FGMs) are substrates with a spatial gradient in material properties, such as stiffness, color, Poisson's ratio, etc. Material properties of FGMs vary from location to location (the *macro-scale*) through adjustments of the *meso-scale* geometry, which can be 3D printed at a *micro-scale* resolution. As such FGMs have a wide application area; they are currently utilized by aerospace, automotive, biomedical and defense. FGMs have been used to create optimized designs of wings, tires, implants, personalized footwear and artistic head phones. They are also referred to as spatially graded properties or (mechanical) metamaterials.

FGMs appear extensively throughout nature; for example bone structures exhibit a spatially graded stiffness. See Fig. 1.1a. Through evolution bones have been optimized to withstand wear and impact forces while using minimal material. FGMs could be optimized so that different segments of a product adhere to different constraints, while minimizing weight, production cost, manufacturing time, etc. Moreover, FGMs can be tweaked to exhibit properties not found in the base material(s); for example *auxetic* materials have a negative Poisson's ratio, meaning that when they are squeezed in one direction they contract in the others rather than bulging out.

FGMs can be classified into two categories: graded surface properties and graded volumetric properties. While color is a surface property, stiffness is considered as a volumetric property. However, surface properties can also be considered volumetric properties confined to a thin shell of the total volume. State-of-the-art color reproduction techniques consider how light penetrates the volume and how it scatters before it departs the object, thereby blurring the dichotomy between surface and volumetric properties.

Functionally graded surface properties can be achieved using various mechanisms. For example, dithering is typically used in inkjet 3D printing to control the relative densities of droplets in the base colors of the inks. Lithophanes are thin sheets with varying thickness to control the translucency, so that an image appears when light is shone through. Similarly, a glossy coating of varying thickness can be used to achieve spatially graded glossiness [2].

Mechanisms to control volumetric properties are researched extensively. Volumetric FGMs typically consist of semi-repeating cells, for which the geometry is spatially varied in order to achieve the material property gradient. Such a pattern of similar cells defines a *microstructure* (or more aptly a *mesostructure*). Microstructures range from regular tilings to stochastic structures: on one end we find regular honeycombs such as the cubic lattice, in the middle are fractal structures which are semi-repeating, and on the other end

we find 3D Voronoi diagrams. Another way to categorize microstructures is according to the dimensionality of the submanifold of the cell geometry: 3D voluminous cells such as solid triply periodic minimal surfaces (TPMS), 2D sheetlike cells such as 3D Voronoi diagrams, 1D edge-based cells such as beam lattices, and 0D point-based cells such as voxel-based techniques. Their material properties can be controlled in various ways; for instance: 3D exploring the effect of (topological and morphological) geometry changes [3], 2D level-set methods to control the sheet thickness [4], 1D beam thickness control [5], and 0D ordered dithering to control the relative placement of voxels [6].

Volumetric FGMs which do not use microstructures are also discussed in literature. A large volume of literature is dedicated to producing FGMs using analog techniques not based on additive manufacturing (AM). Foams are an example of mechanical metamaterials which typically consists of a base material and air; some analog processes can be exploited to manufacture spatially graded foams. Metals, ceramics and/or thermoplastics can be fused together more robustly by introducing a micron-scale gradient between the materials [7].¹ For additive manufacturing there are techniques to mix two materials together during the production process, such as a mixing nozzle [8].

For an extensive overview of techniques for FGMs I refer the reader to state-of-the-art review papers [9, 10, 11, 12, 13, 7].

I.2 FUSED FILAMENT FABRICATION TECHNOLOGY

FGMs are commonly produced by additive manufacturing (AM) a.k.a. 3D printing. The flexibility of additive manufacturing systems allows for the production of complex geometry at little to no extra cost, and the lead time is low compared to traditional manufacturing processes. In recent decades, the popularity of additive manufacturing has exploded. Advances in precision, reliability, material range and the expiration of crucial patents have played a key role in this development. Not only are 3D printing systems used for rapid prototyping, but increasingly also for small quantity production of end products, personalized products, tooling, jigs and fixtures. 3D printing is prevalent in aerospace, defense, healthcare, automotive, education and research.

This dissertation focuses on a type of AM called Fused Filament Fabrication (FFF), a.k.a. Fused Deposition Modeling (FDM). FFF operates by melting thermoplastic material and depositing it in layers which are stacked on top of each other to produce the end part. As the process of extruding material is relatively low-tech, FFF allows for many different types of material: thermoplastics as well as other pliable materials and pastes. Moreover, it means that 3D printing using FFF is comparatively cheap. The worldwide adoption of FFF has skyrocketed in recent decades. For these reasons, this dissertation focuses on FFF to manufacture FGMs.

However, the adoption of FFF for manufacturing FGMs is modest in comparison to other AM technologies. In order to explain the limitations of FFF, Section I.2.1 describes the hardware of a typical FFF system. From the hardware configuration we derive the manufacturing constraints in Section I.2.2. These constraints have to be taken into account when generating the 3D printer instructions, which are covered in Section I.2.3.



Fig. 1.2. Ultimaker S3 hardware overview

1.2.1. HARDWARE

The limitations of FFF technology derive from the hardware of such systems. In order to explain the mechanisms of FFF systems, I use 3D printers from Ultimaker as an example; see Fig. 1.2

In order to extrude material, the feeder forwards filament toward the print head, where the material gets melted and protrudes from the nozzle. The speed at which material is ejected from the nozzle depends on factors such as the material viscosity, the temperature and the amount of pressure delivered by the feeder. The material gets flattened from a cylindrical shape by squeezing it in between the nozzle and the previous layer, to which it then fuses because of the residual heat. The thus extruded line is also called a ‘bead’ or a ‘trace’. The first lines are extruded directly onto the build plate. When the first layer is finished the print head moves up to print the next layer on top of it. In order to travel in between extrusion lines without extruding, the filament is retracted to reduce the pressure in the system and to prevent molten material from oozing out; then it performs an unretract in order to resume the extrusion process.

The latest models from Ultimaker are dual extrusion machines; the print head has two nozzles from which two separate materials can be extruded. The second material is often used to print a sacrificial support structure, which can be easily removed after printing. However, the second material can also be used to manufacture dual material products where different parts of a product have different (proportion of) materials.

1.2.2. MANUFACTURING CONSTRAINTS

From such a hardware setup we can derive several limitations typical to FFF systems. See Fig. 1.3. In order to print reliably, the following manufacturing constraints should be met – both when printing FGMs and when printing regular 3D models.

Resolution Any printed geometry cannot be smaller than the hole in the nozzle (a.k.a. the *nozzle size*), which is typically 0.4 mm. However, the accuracy of the positioning system is an order of magnitude smaller than the nozzle size, so the accuracy of the outer dimensions is far better. In the Z direction the resolution of a print is determined by the

¹In the context of metals the term ‘microstructure’ is commonly used to refer to some micron-scale phenomena, whereas in the field of AM the term commonly refers to geometry at the millimeter- or centimeter-scale.

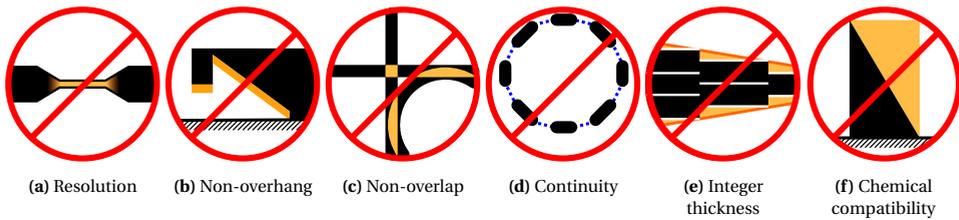


Fig. 1.3. Manufacturing constraints for the FFF process.

thickness of the layers, which is typically in the range of 0.05 mm to 0.2 mm.

Support In order to squeeze the extrudate onto the previous layer, there should be material present underneath the extruded line. Overhanging geometry is considered to be self-supporting up to a certain incline. If the 3D model has a too steep overhang, the extruded lines will not be supported enough by the layer below, and an auxiliary support structure should be printed below the overhang.

Non-overlap Material cannot be extruded twice in the same location. When the printer is instructed to do so, the pressure in the system will increase up to a point where the hardware will fail. However, for small amounts of such over-extrusion, the pressure is released in consecutive extrusion lines, where a small blob defect can appear.

Continuity Printed lines should be relatively long without intermediate interruptions. In order to stop extrusion a retraction is performed, which leaves an imprint on the filament, causing future feeding to be less accurate. When too many retractions occur on the same portion of filament, it gets stuck ('grinding') and the printing process fails. Moreover, because of inaccuracies in the extrusion control mechanism, the exact shape of the starts and ends of extruded lines are difficult to control.

Integer thickness Thin geometry should have a thickness equal to an integer multiple of the nozzle size. Although this constraint is often overseen, it is a consequence of the resolution constraint and the non-overlap constraint.

Chemical compatibility When printing with two different materials, the adhesion or fusion between the materials can be negligibly small. When the boundary between two regions of such chemically incompatible materials is a simple smooth interface, the model can easily fall apart after or during printing. In the latter case a print failure would ensue. Multi-material designs should therefore only be assigned combinations of materials which are mutually compatible.

I.2.3. TOOLPATH GENERATION

These constraints need to be taken into account when printing a 3D model. The 3D model is loaded into a toolpath generation software package (a.k.a. a *licer*) such as Cura, which computes the 3D toolpaths and other instructions the 3D printer should follow. The tool-

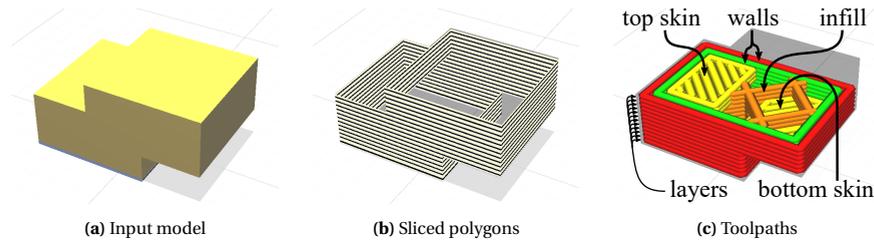


Fig. I.4. Toolpath generation for FFF

path generation algorithms consist of several stages, each of which observes the aforementioned manufacturing constraints. See Fig. I.4.

Slicing The 3D model is sliced by a stack of horizontal planes into polygons which describe the areas to be filled on each layer. Often the layer thickness is homogeneous throughout the model, because varying the layer thickness accurately requires varying the feeder pressure, temperature and/or movement speed, which requires an often prohibitively sophisticated model of how these settings interact.

Walls From the polygons of each layer several extrusion passes are determined which follow the contours, called the *walls*. These are commonly computed by performing offsets of an odd multiple of half the nozzle size from the outline. These extrusion lines follow the contours of the model because the positioning accuracy is better than the extrusion resolution. Moreover, contour lines adhere to the continuity constraint given that the layer polygons are generally relatively large.

Top/bottom skin Parallel extrusion lines are generated for the regions inside the walls which are close to the tops or bottoms of the 3D model. Such straight lines can be printed quickly since the positioning system does not need to decelerate for any corners along those lines. Several layers of top/bottom skin will be printed in order to ensure that the final product is strong.

Infill The remaining regions are filled using some sparse *infill* structure, which uses less time and material to print, but which is just dense enough for the topmost top skin not to exhibit defects. The infill density and geometry can be adjusted depending on the use case for the 3D model. Depending on the use case for the 3D model the infill density and geometry can be adjusted.

These steps are just a tiny part of the total number of algorithms employed by state-of-the-art toolpath generation packages such as Cura. Cura has a vast number of settings, which have to be carefully tuned and verified in order to have a setting profile which can print a wide range of 3D models accurately. For a more extensive review of the toolpath generation process I refer the reader to a review of the state-of-the-art [14].

I.3 RESEARCH OBJECTIVE

While FFF allows for considerably more design complexity than traditional fabrication methods, the manufacturing constraints listed in Section I.2.2 make FFF more restricted than other additive manufacturing techniques. Research into FGMs is frequently focused on additive manufacturing methods other than FFF. Dithering approaches work well for AM processes like jetting, because at each position a droplet of a different material can be ejected. Beam lattices are often manufactured using powder bed fusion techniques, because the powder itself can act as support for overhanging beams. However, such FGMs violate the continuity constraint of FFF, because the layers would consist of many small islands of material. In fact, many of the currently proposed FGM techniques violate the above manufacturing constraints and therefore cannot reliably be printed using FFF.

Therefore, this dissertation aims to answer the following question: **How can we generate toolpaths for fused filament fabrication 3D printers which enable functionally graded materials?** The main research question involves answering the following sub-questions:

- What material properties of FFF products can be controlled?
- How can we adjust the printing process such that these properties are graded spatially?
- How can these FGMs be realized such that they adhere to manufacturing constraints?

The last question above can be advanced by working toward either of two objectives:

- Generating FGMs which adhere to FFF manufacturing constraints.
- Generating FFF toolpaths in such a fashion as to alleviate manufacturing constraints which impede FGMs.

This dissertation covers several techniques for spatially grading various material properties, using FFF. The proposed techniques contribute to these two objectives in different proportions; some are concerned more with generating FGMs, others are concerned more with toolpath generation, but all of the presented techniques contribute to both these objectives in some respect.

Because materials have many properties which can be of interest, the field of FGMs is vast; current research has barely scratched the surface of this subject. The research is therefore limited to unmodified commercially available FFF hardware systems. Furthermore, this research is constrained to toolpath generation techniques which take the manufacturing constraints and fine-level workings of FFF into account. Additionally, our research is focused on toolpath geometry; other interesting ways to control material properties locally such as temperature control and fan speed control are beyond the boundaries of this dissertation. Nevertheless, the scope of this research is so wide that the presented work leaves a large body of potential research for future endeavors.

I.4 DISSERTATION ORGANIZATION

The upcoming chapters of this dissertation present several toolpath generation techniques which unlock the spatial gradation of several material properties. The four

Toolpath generation for fused filament fabrication of functionally graded materials

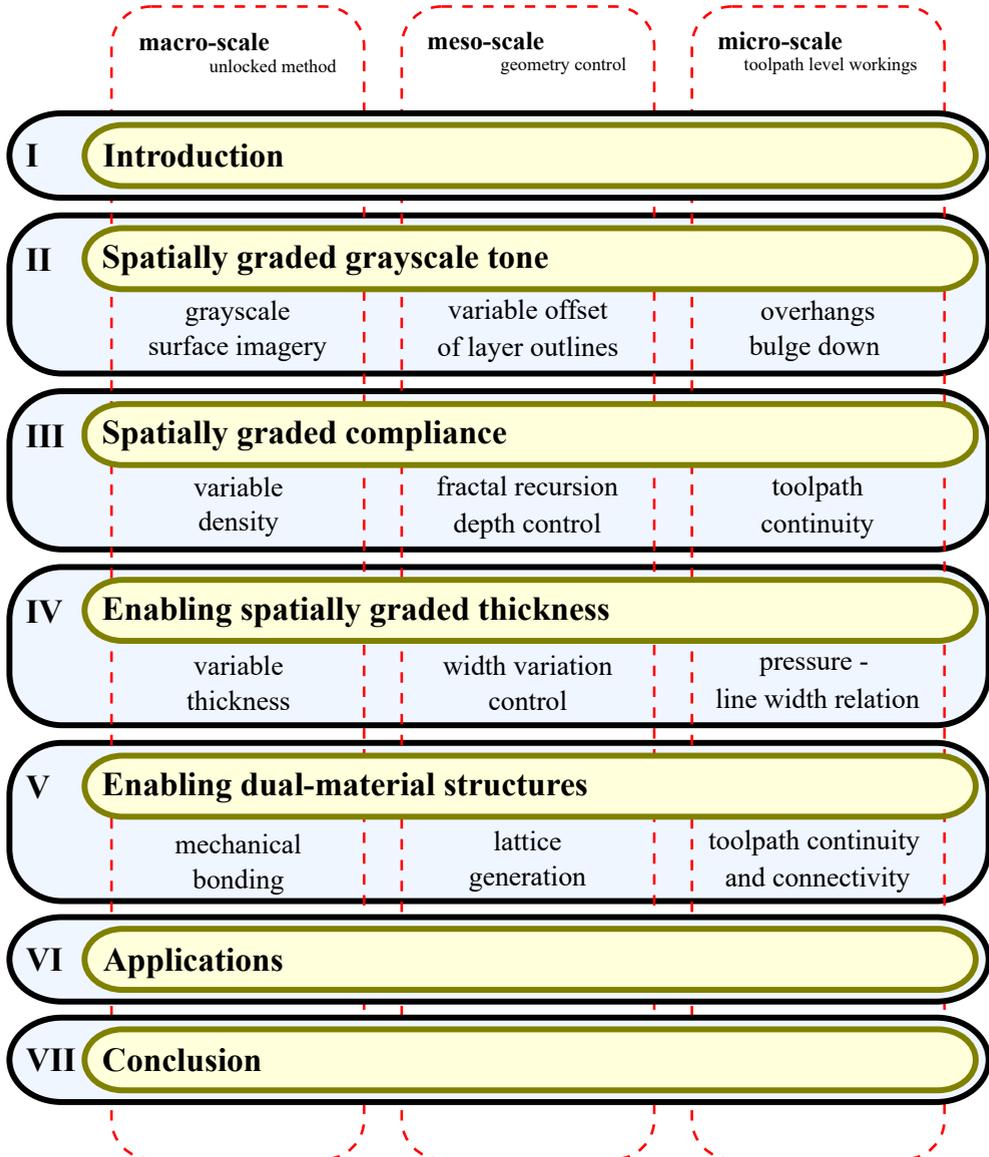


Fig. I.5. Dissertation overview and common themes throughout the chapters.

research cycles are covered by Chapters II to V and they are extended in Chapter VI. Three themes are common to the main research cycles: a) the micro-scale physical mechanisms of FFF on a toolpath level b) the meso-scale toolpath geometry c) the macro-scale material properties. See Fig. I.5.

Chapter II introduces a method to vary the material color. The technique can be used to produce grayscale imagery on the surface of 3D prints using a black and a white filament on off-the-shelf dual material FFF systems. By varying the geometry of the outlines of each layer, we can control the geometry of overhanging printed lines, which then determines the relative amount of black or white material visible at a given location on the surface.

Chapter III proposes a technique to vary the compliance (or conversely the stiffness) of a print. It presents a framework for generating a fractal which, when sliced, produces strictly continuous toolpaths on each layer, which is an important requirement specifically when printing with flexible materials. By using a different recursion depth of the fractal at different locations, we control the density and thereby the stiffness of the structure from region to region. The structure can thereby be adjusted to mimic different types of foam in different locations.

Chapter IV presents a framework for varying the thickness of thin geometry. A control mechanism to reliably manufacture varying width toolpaths is proposed, and an algorithm to generate varying width toolpaths to fill thin geometry which has a varying thickness. This research alleviates the integer thickness manufacturing constraint. It thereby unlocks several types of FGM which control material properties through the thickness of surfaces, such as triply periodic minimal surfaces. By spatially altering the thickness of the surfaces, the product obtains spatially graded stiffness.

Chapter V investigates a microstructure for bonding incompatible materials together. The manufacturing of multi-material FFF products is hampered by the fact that certain combinations of material do not adhere to each other. A method is proposed for interlocking two parts of a different material by introducing a microstructure at the interface between the two regions. The thickness of the microstructure geometry is then optimized for the best strength. Alleviating the chemical compatibility manufacturing constraint unlocks a wide range of multi-material FGMs.

Chapter VI discusses some novel applications which are unlocked by the presented techniques. The technologies developed in Chapters II to V have implications not only for their intended goal, but also unlock purposes which fall beyond the goal of the individual chapters.

Chapter VII concludes this dissertation. It reflects on the research question and anticipate future research.

REFERENCES

- [1] Hongming Zong et al. "VCUT level set method for topology optimization of functionally graded cellular structures". In: *Computer Methods in Applied Mechanics and Engineering* 354 (2019), pp. 487–505. ISSN: 00457825. DOI: 10.1016/j.cma.2019.05.029.
- [2] Willemijn Elkhuisen et al. "Gloss, color, and topography scanning for reproducing a Painting's appearance using 3D printing". In: *Journal on Computing and Cultural Heritage* 12.4 (2019). ISSN: 15564711. DOI: 10.1145/3317949. arXiv: 1910.10836.

- [3] Christian Schumacher et al. “Microstructures to control elasticity in 3D printing”. In: *ACM Transactions on Graphics* 34.4 (July 2015), 136:1–136:13. ISSN: 07300301. DOI: 10.1145/2766926.
- [4] Simon R.G. Bates, Ian R. Farrow, and Richard S. Trask. “Compressive behaviour of 3D printed thermoplastic polyurethane honeycombs with graded densities”. In: *Materials & Design* 162 (2018), pp. 130–142. ISSN: 02641275. DOI: 10.1016/j.matdes.2018.11.019.
- [5] Sing Ying Choy et al. “Compressive properties of functionally graded lattice structures manufactured by selective laser melting”. In: *Materials and Design* 131.May (2017), pp. 112–120. ISSN: 0264-1275. DOI: 10.1016/j.matdes.2017.06.006.
- [6] Qun Lou and Peter Stucki. “Fundamentals of 3D Halftoning”. In: *Electronic Publishing, Artistic Imaging, and Digital Typography*. Springer Berlin Heidelberg, 1998, pp. 224–239.
- [7] Vasavi Boggarapu et al. “State of the art in functionally graded materials”. In: *Composite Structures* 262.January (2021), p. 113596. ISSN: 02638223. DOI: 10.1016/j.compstruct.2021.113596.
- [8] Haichuan Song et al. “Colored Fused Filament Fabrication”. In: *ACM Transactions on Graphics* 38.5 (2019), pp. 1–11. ISSN: 07300301. DOI: 10.1145/3183793. arXiv: 1709.09689.
- [9] Francesco Tamburrino, Serena Graziosi, and Monica Bordegoni. “The design process of additive manufactured Meso-Scale Lattice Structures: a review”. In: *Journal of Computing and Information Science in Engineering* 18.4 (July 2018). ISSN: 1530-9827. DOI: 10.1115/1.4040131.
- [10] Yan Li et al. “A Review on Functionally Graded Materials and Structures via Additive Manufacturing: From Multi-Scale Design to Versatile Functional Properties”. In: *Advanced Materials Technologies* 5.6 (2020). ISSN: 2365709X. DOI: 10.1002/admt.201900981.
- [11] Dalia Mahmoud and Mohamed Elbestawi. “Lattice Structures and Functionally Graded Materials Applications in Additive Manufacturing of Orthopedic Implants: A Review”. In: *Journal of Manufacturing and Materials Processing* 1.2 (2017), p. 13. ISSN: 2504-4494. DOI: 10.3390/jmmp1020013.
- [12] G. Loh et al. “An overview of functionally graded additive manufacturing”. In: *Additive Manufacturing* 23 (2018). ISSN: 22148604. DOI: 10.1016/j.addma.2018.06.023.
- [13] Bassiouny Saleh et al. “30 Years of functionally graded materials: An overview of manufacturing methods, Applications and Future Challenges”. In: *Composites Part B: Engineering* 201.January (2020), p. 108376. ISSN: 13598368. DOI: 10.1016/j.compositesb.2020.108376.
- [14] Marco Livesu et al. “From 3D models to 3D prints: an overview of the processing pipeline”. In: *Computer Graphics Forum* 36.2 (May 2017), pp. 537–564. ISSN: 14678659. DOI: 10.1111/cgf.13147. arXiv: 1705.03811.

II

SPATIALLY GRADED GRAYSCALE TONE

In the previous chapter we introduced the idea of functionally graded material properties and the challenges and opportunities we face when aiming to manufacture them using fused deposition modeling (a.k.a. fused filament fabrication). One type of FGM pertains to surface properties, such as color. By spatially varying the color we can create objects with imagery on the surface.

This chapter proposes a technique which uses an off-the-shelf dual material FDM printer to create grayscale imagery on 3D models. By altering the geometry of the outlines of layers printed alternately with black and white filament, we can modulate the amount of black or white material visible in each location. Viewed from a distance these ratios of visible black and white blend into an intermediate grayscale tone. We present a framework which automatically translates the colors of textured 3D models into variable offsets, and an algorithm which automatically applies such variable offsets to the layer outline polygons.

Sections II.1 and II.2 introduce the functionally graded surface color technique and provide its background in FDM, and in coloring techniques in general. The grayscale toning principle is explained in Section II.3 and the required algorithmic framework is laid out. Sections II.4 and II.5 present the results of this method and discuss these, which leads us to the conclusion in Section II.6.

This chapter is published as T. Kuipers, W. Elkhuisen, J. Verlinden, and E. Doubrovski. “Hatching for 3D prints: Line-based halftoning for dual extrusion fused deposition modeling”. In: *Computers & Graphics* 74 (2018), pp. 23–32. ISSN: 00978493. DOI: 10.1016/j.cag.2018.04.006; some minor corrections have been made.

ABSTRACT

This work presents a halftoning technique to manufacture 3D objects with the appearance of continuous grayscale imagery for Fused Deposition Modeling (FDM) printers. While droplet-based dithering is a common halftoning technique, this is not applicable to FDM printing, since FDM builds up objects by extruding material in semi-continuous paths. The line-based halftoning principle called 'hatching' is applied to the line patterns naturally occurring in FDM prints, which are built up in a layer-by-layer fashion. The proposed halftoning technique is not limited by the challenges existing techniques face; existing FDM coloring techniques greatly influence the surface geometry and deteriorate with surface slopes deviating from vertical or greatly influence the basic parameters of the printing process and thereby the structural properties of the resulting product. Furthermore, the proposed technique has little effect on printing time. Experiments on a dual nozzle FDM printer show promising results. Future work is required to calibrate the perceived tone.

II.1 INTRODUCTION

The ability to apply color to 3D printed parts is relevant for both prototyping and manufacturing. Possible applications include reproduction of color-scanned 3D objects and fabrication of products with logos and labeling. Color can also be used as a design feature or to visualize geometric information such as the results of finite element analyses. See Fig. II.1.

At present, 3D printing in color is available for a variety of Additive Manufacturing (AM) systems that are predominantly based on ink-jet technology. Techniques for printing color using only Fused Deposition Modeling (FDM) are sparse and suffer from a low resolution or have radical impact on the printing process.

This paper presents a novel technique for fabricating 3D grayscale objects using the FDM 3D printing method. It uses a principle which is based on modulating the visible width of printed lines of two alternating colors to produce the appearance of continuous tone gradients. Creating the perception of continuous tones by generating small patterns of discrete colors is termed *halftoning*.

Implementing color variation with high frequency details using FDM is a challenge. FDM builds up objects by extruding material in semi-continuous paths, which makes it impossible to apply droplet-based halftoning principles that are commonplace in existing color 3D printers.

A promising technique to fabricate continuous tone objects using FDM has been presented by Reiner et al. [2]. However, since that technique inherently produces textures at a relatively low sample rate, it does not allow the fabrication of high frequency details. Furthermore, the technique does not allow fabrication of textures on horizontal surfaces and degrades for diagonal surfaces with a slope approaching horizontal.

Addressing these issues, we propose a novel halftoning technique for dual extrusion FDM systems. The proposed halftoning technique is based on hatching, an established 2D halftoning principle based on lines rather than dots. The implementation of the technique described in this paper is open source and can be found at github.com/Ultimaker/CuraEngine [3].

This paper is an extension of the techniques proposed by a conference paper [4]. The technique proposed there changes the geometry of alternating black and white layers to



Fig. II.1. 3D prints obtained by applying hatching on a 14 cm 3D portrait, a 15 cm artistic figurine, a full size soda can with textual information, and the result of a stress analysis performed on a 16 cm connecting rod of a piston engine.

modulate the perceived grayscale tone when viewed either from directly above or when viewed straight from the side.

Rather than presenting different hatching techniques for two viewing angles, this paper presents a unified hatching technique for viewing surfaces from any angle and in particular a viewing angle locally perpendicular to the surface. The phenomenon that an overhanging line occludes the previous layer, also known as sagging, is exploited for an edge case of that general hatching technique. This paper provides a model of the sagging behavior which is used to derive the proportion of visible white to black filament from any viewing angle. Experimental data of sagging is collected and analysed in order to grasp the limitations of our model.

II.2 BACKGROUND

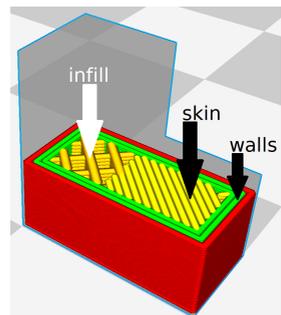
II.2.1. PROCESS PLANNING FOR FDM

The following section will briefly explain the basics of process planning for FDM, a.k.a. *slicing*. Some elementary concepts and processes are explained required for understanding the presented hatching technique. The terminology employed and the techniques described apply to the open source slicing software called Cura [3].

FDM generally builds up 3D prints in a layer-by-layer fashion. One of the first stages in slicing is generating the outlines of each layer. The outlines are the boundaries of the regions which are to be filled with material. Line segments are generated by intersecting each triangle of the input mesh with horizontal planes at heights corresponding to each layer. All line segments of a layer are then stitched into polygons which form the outlines of that layer.

Because starting and stopping extrusion of filament causes blemishes, the outlines of a layer are achieved by following the contours of the layer: the walls. Several consecutive walls are printed next to each other. The outer walls are generated by applying an inward offset of half the line width to the outlines. Successive walls are then generated by applying offsets to previous walls. These walls define printed lines which follow the contours of the object.

The remaining area within the innermost wall is split into infill and skin. By applying boolean operations on the leftover region with the outlines of layers above and below we calculate the areas which are close to the top and bottom of the model boundary surface: the skin. By applying a difference operation we can then determine the infill areas from the skin areas and the region left over from the walls. The skins are several layers thick and they are densely filled with a pattern of parallel lines.



II.2.2. COMMERCIAL COLOR 3D PRINTERS

The first commercial full color 3D printing systems date back to 1993 [5]. These systems use ink-jet technology to apply colored binder onto white powder [6]. Consecutive layers of bound powder form the final 3D model. Instead of jetting a binder onto a substrate, Mcor developed a process in which conventional ink is jetted onto sheets of paper, which are then cut and stacked [7]. Stratasys uses ink-jet technology to print the building material itself. Their recent system incorporates six heads, each able to print a colored ma-

terial [8]. More recently, HP Inc. introduced a printing technology in which liquid agents are jetted onto powder in order to alter the powder's fusing behavior. According to the company, these agents may also include color in the future [9].

II.2.3. 3D HALFTONING

Because printers work with a limited number of base colors, specific strategies need to be applied to make full color prints. In 2D printing, this is usually done through different halftoning techniques.

2D ink-jet technologies apply a halftoning principle called dithering. In dithering, the distance between printed colored dots is varied to create perceived variations of colors. While halftoning for 2D printing industry is well developed, halftoning for 3D printers is still an active field of research. The first mention of halftoning in 3D printing is not focused on color reproduction but on material density variation for stereolithography [10]. Techniques for 3D color dithering have been presented for binder jetting printers [11]. Vidim, Wang, and Ragan-kelley presented a programmable pipeline for multi-material 3D printing. In their pipeline, dithering is applied for both visual and mechanical properties [12]. With the aim to produce full-color prints using material jetting technology, Brunton, Arikan, and Urban presented a halftone technique that takes into account the translucency of the printed material.

The different commercial systems discussed above all use ink-jet technology, and the color halftoning techniques considered are based on the ability of ink-jet to deposit discrete droplets of color in the micrometer range. This allows the production of high frequency details. However, FDM lacks the ability to deposit discrete features in this size range, since it builds up objects by extruding semi-continuous lines of material. Therefore, the halftoning techniques discussed above cannot be directly adapted to FDM. To create high frequency details using FDM, the development of novel halftoning techniques is needed.

II.2.4. COLOR FDM

Limited by the semi-continuous material extrusion of FDM, techniques have been proposed that aim to reproduce continuous tones in FDM prints. These can be categorized into two tactics. 1. Mixing material and color prior to extrusion. 2. Applying halftone principles.

Corbett's FDM grayscale printing technique performs continuous color mixing by implementing a mixing nozzle in the print head [14]. In this setup, multiple feeders are connected to a single nozzle. Multiple materials are fed into a volume of the nozzle where they are molten and mixed. The main limitation of this technique is that color changes require the full volume inside the mixing nozzle to be flushed; the horizontal resolution is limited by the length a printed line needs to be to flush the nozzle.

Reiner et al. have shown that a dot-based halftoning principle can be applied to dual nozzle FDM systems in order to produce two-tone texture mapped 3D prints [2]. Their technique involves applying sine patterns to the outer contour of each layer, with alternating each layer between black and white filament and shifting the sine patterns by half the wave length every two layers. A texture-based amplitude modulation is then applied in order to make the peaks of the one filament protrude more than the other, which results in a shift in perceived color toward the former filament. The main challenge of this technique is to align the phase of the sine pattern across consecutive layers with outlines of

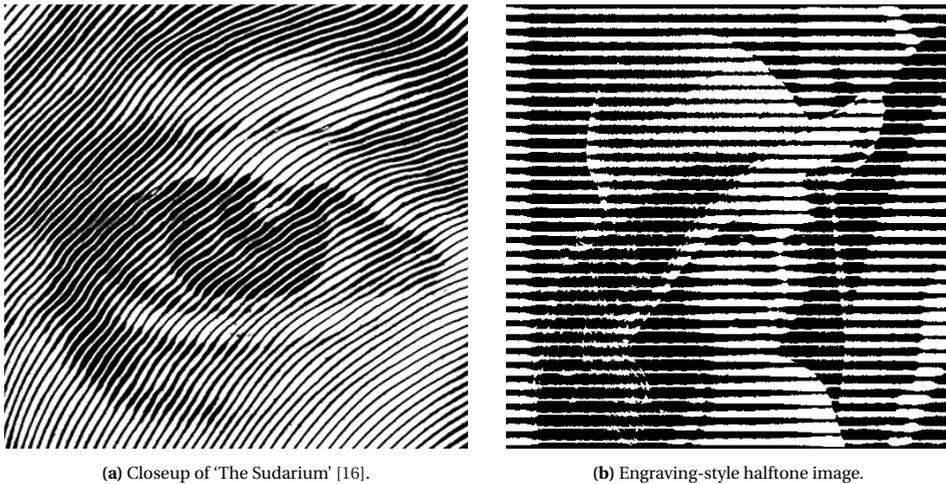


Fig. II.2. Examples of 2D hatching.

a different geometry. The horizontal resolution of the produced textures is limited by the sine's wavelength, which in turn is limited by the width of the extruded lines. The application of the sine pattern therefore results in loss of high frequency details in the geometry. Moreover, this technique works best for vertical surfaces and performs significantly less on surfaces with lower slopes.

More recently Song and Lefebvre have demonstrated a continuous tone imagery technique for FDM printing [15]. Though they make use of a mixing nozzle to alleviate several calibration issues, the fundament of their technique is inherently applicable to any multi-extrusion system which supports two or more base materials. Each layer of the 3D print consists of several sublayers of slightly translucent material with different color. By varying the thickness of these sublayers the perceived color can take any color within the gamut spanned by the colors of the base materials. The technique is limited by the smallest sub-layer thickness which can reliably be achieved on an FDM printer.

While the technique is able to produce stunning results, relying on variable sublayer thickness can cause several problems. Because the layer thickness is one of the most elementary properties of the printing process, many parameters may depend on it: the overhang angle for determining where to place support structures, the required amount of infill, the optimal movement speeds during extrusion, the optimal cooling speed of the fans, etc. While it is reasonable to find a local optimum in print settings given a static layer thickness, it is hard to find out how all settings relate to each other in order to support varying layer thickness with optimal print settings. A suboptimal combination of these settings might lead to structural problems such as a decreased tensile strength.

Another problem with varying layer thickness is that the nozzle may collide with higher regions of a layer when printing lower regions. Most commercial nozzles have a flat horizontal ring around the hole from which material protrudes. This flat part helps to fuse the printed line together with neighboring lines of the current layer. When trying to print a line with a low thickness adjacent to a line with a high thickness, the nozzle collides with the already printed line, which can cause the print to fail.

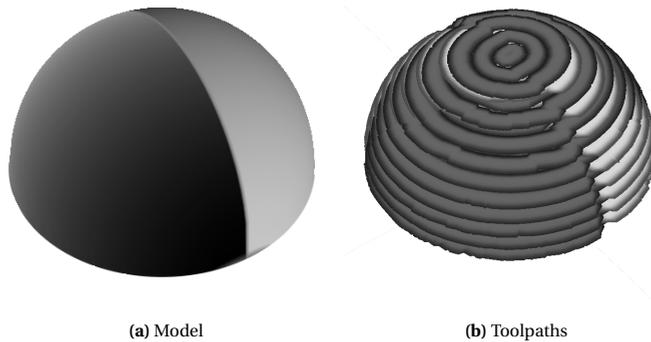


Fig. II.3. Example of a 4.5 mm textured half dome and a visualization of the corresponding toolpaths.

II.2.5. HATCHING

Hatching is a halftoning principle which has lines as elements, rather than dots, which are commonplace in halftoning techniques. Hatching dates back to 17th century engraving techniques [17]. An example of this can be seen in Fig. II.2a. Variations in the perceived tone are achieved by varying the local ratio between the width of black lines and the width of the white background surrounding it [18]. In more recent developments, image processing algorithms have been proposed which convert grayscale images into black and white engraving-style halftone images [19]. Their results resemble Fig. II.2b, which was produced using the linear Newsprint filter from the GNU Image Manipulation Program. The linear characteristic of hatching makes it a particularly suitable technique for halftoning in FDM.

II.3 METHOD

The presented halftoning technique leverages the discrete nature of FDM. Layers are printed with alternating filament. All even layers are printed in black while all odd layers are printed in white. The region of a layer which is visible from the outside resembles a line. Changing the perceived widths of these lines controls the observed grayscale tone. This is done by changing the outlines of a layer. See Fig. II.3.

At places where the model should be darker the outlines of the black layers are expanded and the outlines of the white layers are contracted, so that more black filament will be visible on the surface. Because only the outlines are changed while the patterns with which these outlines are filled up remain the same, the structural properties of 3D prints remain unaltered.

At places on the model where the texture tone is moderate or where the surface slope is high, achieving the desired ratio of perceived material simply follows from the viewing angle. For extremal texture tones or near vertical surfaces a more elaborate approach is needed, which is based on sagging. In this section we first illustrate the simple hatching method after which the method is expanded to also make use of sagging.

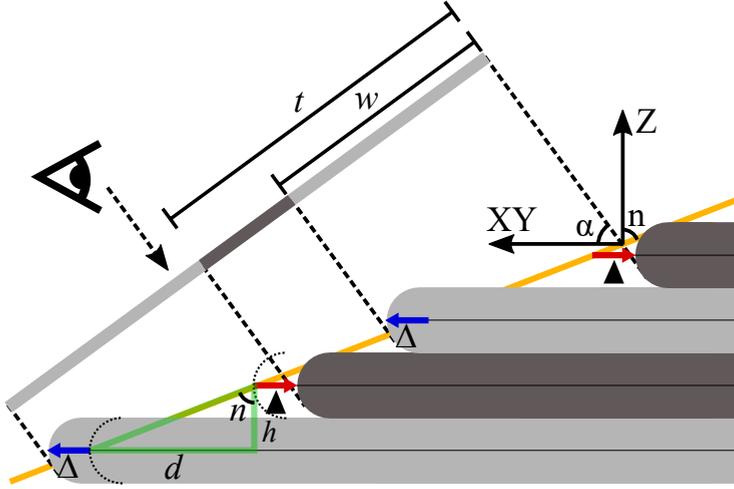


Fig. II.4. Vertical cross-section of a print showing the basic hatching technique. The cross-section is perpendicular to the local surface and to the build plate. The model surface is indicated in orange. The applied offsets in red and blue cause the amount of visible dark filament to be less than the amount of visible light filament.

II.3.1. HATCHING

Because objects are built up by discrete planar layers, sloped surfaces are discretized; this effect is known as the stair-stepping effect. By applying varying offsets to the outline polygons of each layer we can modulate the local widths of these stair steps. A local grayscale tone arises from the proportion of white filament visible w.r.t. black filament at that region. In places where the model needs to be light according to the surface texture color information of the 3D model, the white layers are offset outward and the black layers are offset inward.

The amplitudes of the variable offset should depend on the basic characteristics of the stair-stepping effect. In regions where the surface is nearly horizontal the width of the stair steps is large and so those regions require a larger offset to change the ratio between visible white and black filament.

The amplitudes of the offset should also depend on the viewing angle. When viewing an object from a higher altitude, the sides of the steps have less impact than when viewing the stair steps from a lower altitude.

The ratio r of visible white to the total of visible white and black can then be calculated by projecting the horizontal component (the width of the stair step) and vertical component (the layer height) onto the viewing plane. The perceived ratio r and the corresponding offset Δ are therefore related by the following formula:

$$\begin{aligned} \Delta &= -\blacktriangle \\ r &= \frac{w}{t} = \frac{(d + 2\Delta) \sin(\alpha) + h \cos(\alpha)}{2d \sin(\alpha) + 2h \cos(\alpha)} \\ d &= h \tan(n) \end{aligned} \quad (\text{II.1})$$

, where the ratio r is a value between 0 and 1, \blacktriangle is the offset applied to black layers and Δ to white, w is the projection of visible white filament onto the viewing plane, t is

the projection of two full stair steps into the viewing plane, $0 \leq n \leq 0.5\pi$ is the absolute angle between the surface normal vector \vec{n} and its projection downward on the horizontal plane, α is the absolute viewing angle, h is the layer thickness and d is the horizontal stair step distance corresponding to the given surface slope. See Fig. II.4.

From the texture image we can derive the required ratio of visible white filament r . Disregarding lighting conditions, the amount of visible white filament at a given location should be proportional to the luminance of the corresponding texture image coordinates. (Section II.3.4 covers mapping outline locations to texture UV coordinates.) The luminance values are obtained by applying a gamma expansion to the Luma component of the pixel: $r = (0.2126R + 0.7152G + 0.0722B)^{1/2.2}$ [20]. Using Eq. (II.1), we can then obtain a formula for the required offset at a given location on the model:

$$\Delta = h \left(\frac{1}{2} + r \right) \frac{\sin(n) \sin(\alpha) + \cos(n) \cos(\alpha)}{\cos(n) \sin(\alpha)} \quad (\text{II.2})$$

Performing trigonometric functions on a computer is a relatively expensive task. Luckily, given that $0 \leq n \leq 0.5\pi$, they can be derived from the x , y and z component of the vector: $\sin(n) = \text{abs } \vec{n}_z$ and $\cos(n) = \sqrt{\vec{n}_x^2 + \vec{n}_y^2}$.

The formula above calculates the offset required to optimize the perceived luminance value *w.r.t. a certain viewing angle*. For other viewing angles the perceived luminance value will be off. Choosing a viewing angle could be a user input which depends on the model and its function.

An alternative is to optimize the perceived tone *w.r.t. a viewing angle perpendicular to the surface*. In such a case the normal angle n aligns with the viewing angle α and the above formulae simplify to:

$$r = \frac{1}{2} + \Delta \frac{\sin(n) \cos(n)}{h} \quad (\text{II.3})$$

$$\Delta = h \frac{\frac{1}{2} + r}{\vec{n}_z \sqrt{\vec{n}_x^2 + \vec{n}_y^2}} \quad (\text{II.4})$$

It should be noted that these formulae are independent of the geometry of a cross-section of a layer. As long as the cross section of a white layer follows the same geometry as a black layer, it does not matter whether the side of a layer is straight, a semicircle or something in between.

II.3.2. SAGGING

However, when a layer is offset by an amount such that it extends beyond the previous layer, i.e. when $2 \text{abs } \Delta > d$, the geometry of the cross-section of a layer will change. The phenomenon of an overhanging layer occluding the layer below is known as *sagging*. Figure II.5 shows cross section microscopy photos of a print with various different overhang distances. Reiner et al. proposed that this effect could be used to perform halftoning in FDM: “subtle geometric offsetting between layers enables the control of variation of tone due to occlusion and gravity.” This section explores ways in which to accurately use sagging for hatching in FDM.

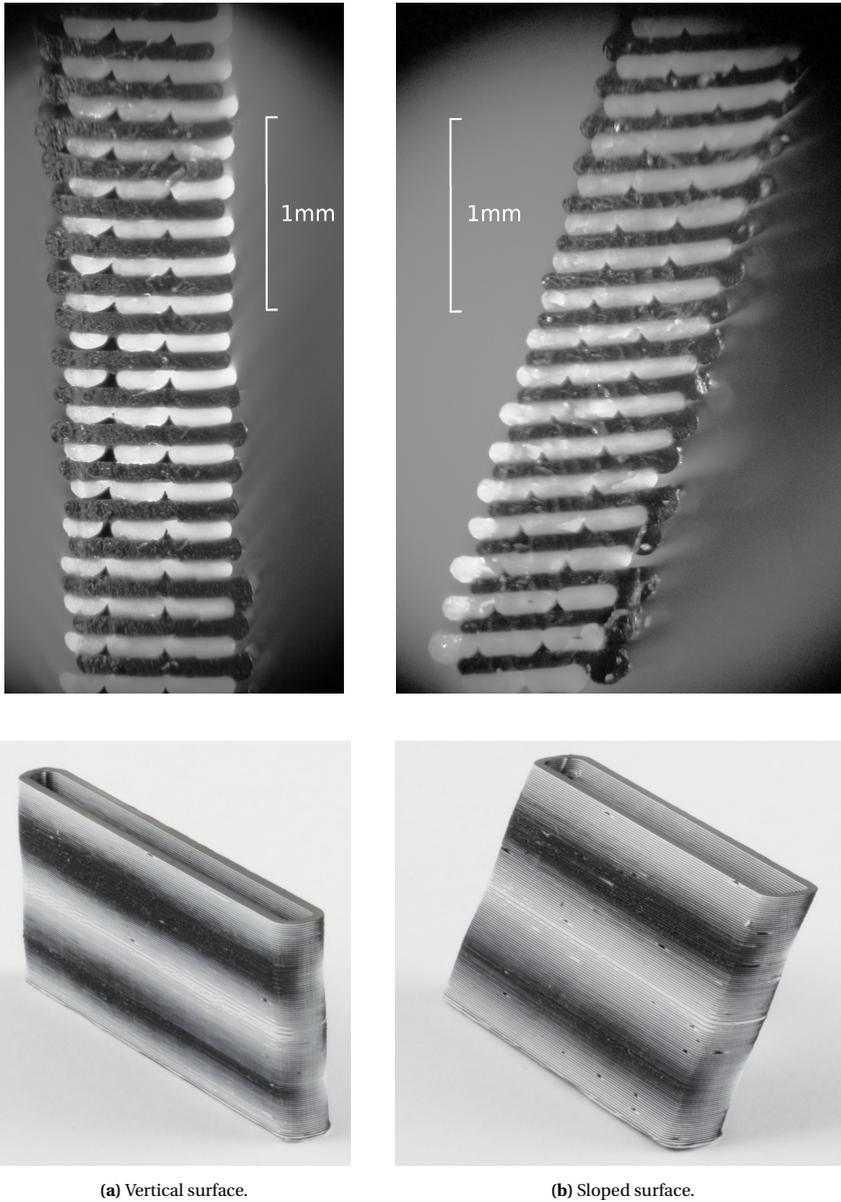


Fig. II.5. Cross-section photos of sagging made with the Olympus SZ61 microscope. The specimen have various overhang distances while the layer thickness is 0.1 mm. The grayscale levels have been adjusted to increase the contrast in the dark and white regions at the cost of contrast of in between tones, so that details in both black and white filament are clearly visible.

Though gravity could influence sagging, it is conceivably a phenomenon caused by the uneven distribution of back pressure due to there being material of the previous layer under only some locations beneath the nozzle. When printing a line on top of a previous layer, the existing layer creates back pressure which forces the material from the round shape of the nozzle to a flatter line. This back pressure drops where there is no layer below, so given the same amount of pressure in the nozzle, more material is extruded there. Given that there is no reason for that material to stay in the plane of that layer it follows that it will start to occlude the layer below.

MODELING SAGGING

In order to produce the right grayscale tone we need to know how the occlusion relates to the amount of overhang. The proportion of white to black filament visible from a given angle is influenced by the amount of occlusion due to sagging. A model which predicts the amount of visible filament from any viewing angle forms the basis of a grayscale tone calibration.

A geometrical model of the cross-section of a sagged line can provide a mathematical formula for calculating the amount of visible black and white filament from any viewing angle. We can model the cross-section of a printed line as a rectangle and two circles for which the center and the radius depend on the amount of overhang. See Fig. II.6. For simplicity we assume that the black and white filament exhibit the same sagging behavior. The amount of occlusion in our model follows the following formula:

$$f(o) = (o - \delta x) \sin(\alpha) + \delta y \cos(\alpha) + \delta r \quad (\text{II.5})$$

for positive angles α , where o is the overhang distance by which the top line is extended beyond the bottom line, f is the amount of occlusion as viewed from α as a function of the amount of overhang and δx , δy and δr are the change in position and radius of the circle which is used to describe the side of the sagged line.

Supplementing Eq. (II.5) to Eq. (II.1), the perceived luminance in case of sagging when viewed from positive α becomes as follows:

$$r_{\rightarrow} = \frac{(d + 2\Delta_{\rightarrow}) \sin(\alpha) + h \cos(\alpha) + s}{2d \sin(\alpha) + 2h \cos(\alpha)} \quad (\text{II.6})$$

$$s = f(\max(0, 2\Delta - 2d)) - f(\max(0, 2\Delta - 2d)) \quad (\text{II.7})$$

If we specify that the circle of the sagging layer touches the layer below then one can easily see that $(o - \delta x)^2 + (h - \delta y)^2 = (h + \delta r)^2$. See Fig. II.6b. When we further specify that

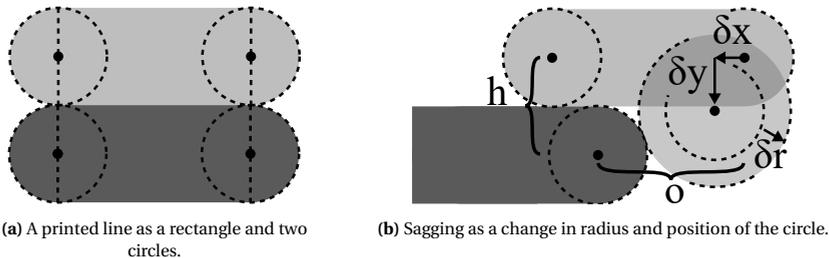


Fig. II.6. Cross-sections of the geometrical model of a printed line.

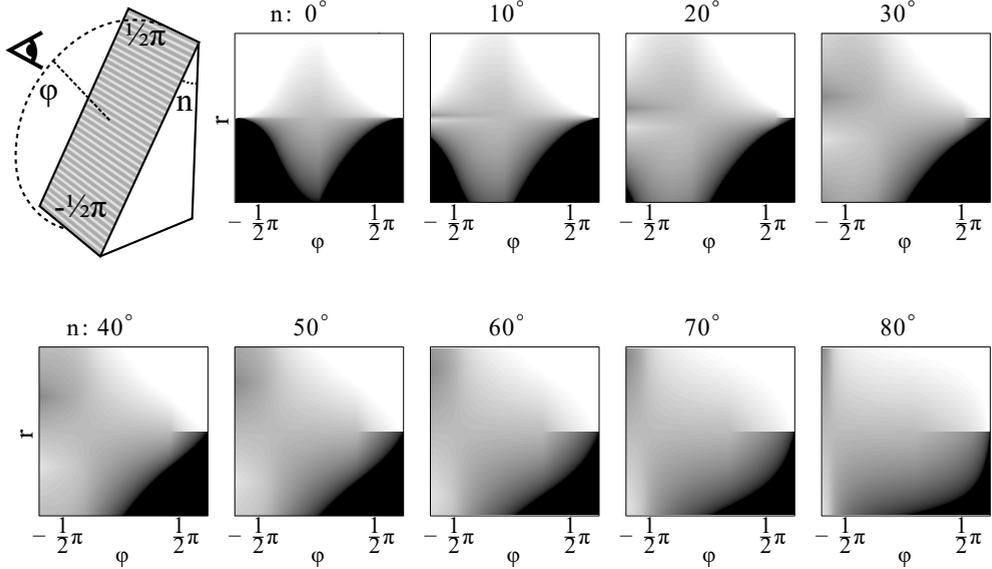


Fig. II.7. Calculated luminance according to the geometrical model optimized for perpendicular viewing under various viewing angles and various ratios of visible white to total filament r . ϕ is the viewing angle w.r.t. the surface: $\phi = \alpha - n$. The image on the top left shows how ϕ and n relate to the print.

the top of the circle coincides with the top of the layer $\delta y = \delta r$ and that the circle touches the previous layer, we can use the Pythagorean theorem to derive that $\delta r = (o - \delta x)^2 / 4h$. For simplicity we model the relation between the circle center receding and the overhang as a linear function going through the origin: $\delta x = Cx \cdot o$. Using the Pythagorean theorem we derive that $Cx = 1 - \sqrt{2}h/w$, where w is the overhang distance at which a layer fully occludes the previous layer when viewed from the side. From preliminary experiments we concluded that at $w = 0.2$ mm a layer of $h = 0.1$ mm would occlude the previous layer fully, which leads to $Cx \approx 0.2929$.

We can then use the quadratic formula to derive the offset required for a viewing angle perpendicular to the surface from the equations above, which results in Algorithm 1:

Note that this algorithm only accurately describes our geometrical model for positive surface slopes, and also positive viewing angles because it is optimized for perpendicular viewing. When relating the perceived proportion of white filament to the offset for viewing angles below zero one should take great care in distinguishing between different occlusion scenarios. When the sagging layer is farther from the viewing plane than the layer below, the layer below occludes the bottom of the sagging layer.

Figure II.7 plots the luminance as function of the viewing angle, surface angle and input grayscale tone. It should be noted that the middle of these graphs – corresponding to a perpendicular viewing angle – coincides with the full gradient from black to white for each model angle. Higher viewing angles result in higher contrast because of the parallax effect. Extreme negative angles result in grayscale tones at the opposite side of the spectrum compared to the input tone; this is due to the fact that sagging causes the top of the layer to be more round, which reduces the amount of visible filament from the sagged layer.

Algorithm 1 Calculating offset distance

```

1: procedure GETOFFSET
2:    $n \leftarrow \text{get}N()$  ▷ Normal vector of the model surface
3:    $L \leftarrow \text{get}L()$  ▷ Luminance of the texture image
4:    $r \leftarrow L^{1/\gamma}$  ▷ Gamma expansion
5:    $\sin\_n \leftarrow \text{abs } n.z$ 
6:    $\cos\_n \leftarrow \sqrt{n.x^2 + n.y^2}$ 
7:    $\tan\_n \leftarrow \sin\_n / \cos\_n$ 
8:    $\text{diag} \leftarrow h / \sin\_n$ 
9:    $d \leftarrow \text{diag} * \cos\_n$ 
10:   $\text{dir} \leftarrow 1$ 
11:  if  $r > 0.5$  then
12:     $r \leftarrow 1 - r$ 
13:     $\text{dir} \leftarrow -\text{dir}$ 
14:   $\Delta \leftarrow ((0.5 - r) / \cos\_n) \text{diag}$ 
15:  if  $2 \text{abs} \Delta \leq d$  then
16:    return  $\text{dir} * \Delta$ 
17:   $Cx \leftarrow \text{get}Cx()$  ▷ Ratio of receding to overhang
18:   $C \leftarrow 1 - 2Cx + Cx^2$ 
19:   $a \leftarrow -0.5 / (h * \text{diag}) C(1 + \sin\_n)$ 
20:   $b \leftarrow 0.5 / \text{diag} (C * \tan\_n(1 + \sin\_n) + 2\cos\_n(Cx - 1))$ 
21:   $c \leftarrow 0.5 - 0.5\cos\_n(C/4 * \tan\_n(1 + \sin\_n) - Cx * \cos\_n) - r$ 
22:   $\text{det} \leftarrow \max(0, b^2 - 4a * c)$ 
23:   $\Delta \leftarrow (-b - \sqrt{\text{det}}) / (2a)$ 
24:  return  $\text{dir} * \Delta$ 

```

Negative surface slopes pose a plethora of problems. Because there are viewing angles at which the sagging reduces the amount of visible filament, there can be combinations of surface normal angle and input tone which have multiple offsets as solution. If the algorithm to compute the offset from the normal and the input luminance is not stable, the perceived ratio of black to white filament can vary wildly for texture colors which only vary mildly.

Note that there is a maximum overhang up to which our sagging model makes sense. If the overhang distance is larger than the line width, no more sagging can occur. Moreover, at a surface slope of about -45° the surface requires a support structure, which inevitably also affects the sagging behavior.

II.3.3. HORIZONTAL HATCHING

The presented hatching technique makes use of the stair-stepping effect and sagging, but these do not occur on horizontal surfaces. Therefore, horizontal top and bottom surfaces of 3D meshes are instead hatched by modulating the widths of the lines used to print the top/bottom skin. The most commonly used pattern to fill in these areas is a regular grid of equidistant straight lines. Grayscale gradients are achieved by modulating the width of these lines. When printing a black layer the line widths are modulated so that parts of the adjacent white layer become visible and vice versa. A visual representation of this is presented in Fig. II.8b. The resulting halftoning images resemble the image in Fig. II.2b.

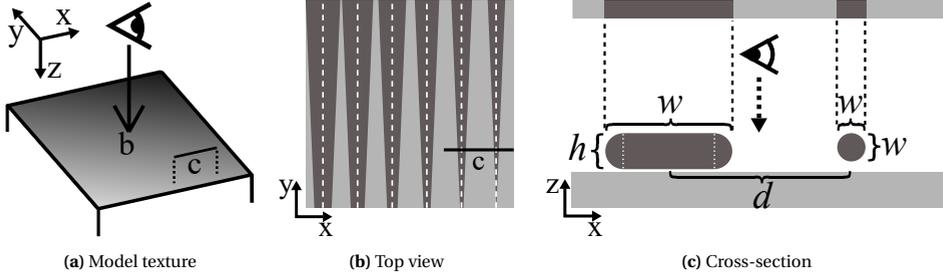


Fig. II.8. Horizontal hatching of a linear gradient. The 'c' marks where the cross section of figure (c) is located. The dashed lines in (b) show the travel paths of the center of the nozzle while printing the black lines. (c) shows the model used for achieving lines of a given width.

The texture image is sampled at regular intervals along the skin lines and the width w of each line segment between two consecutive sample points is determined by the line distance d and the average texture luminance L , which is in the range $(0, 1)$: $w = Ld$.

Varying the width of lines is achieved by increasing and decreasing the amount of deposited material. For Bowden style FDM printers it takes relatively long to change the amount of material departing the nozzle per second – a.k.a. the *flow*. Therefore we propose keeping the flow constant and varying the movement speed v of the print head rather than varying the flow at a constant movement speed: $v = c/A$, where A is the area of a cross section of the printed line and c is an empirically determined constant flow. We model the cross section as a rectangle with semicircular sides: $A = \pi(0.5h)^2 + h(w-h)$, where h is the layer thickness. See Fig. II.8c. For lines narrower than the layer thickness we model the cross section as a circle: if $w < h$ then $A = \pi(0.5w)^2$.

II.3.4. IMPLEMENTATION

This section describes the implementation of the variable offsets required for vertical and diagonal hatching.

TEXTURE MAPPING

Before the variable offsets are applied, the model textures are mapped onto the outlines of each layer. This is done in the phase where the surface mesh is sliced into outlines; for each face which produces a line segment of an outline, the corresponding texture line segment is recorded. Each location on the polygons of the outlines of a layer can then be mapped to a UV coordinate of the texture image. The grayscale value at the point in the image is then used to calculate by which distance the outline is displaced at the location.

VARIABLE OFFSET

The displacement given by the equations above is used to generate a variable offset polygon from the existing outlines of a layer. Both vertices and intermediate points on line segments are displaced in the outward direction perpendicular to the polygon. Line segments are offset by displacing points sampled at a regular interval, as can be seen in Fig. II.9a.

Given that vertices in the outline generally belong to an edge of the 3D model connecting *two* faces, which can have different texture image locations at that vertex, applying a variable offset should take two displacement values into account. The offset Δ_B applied to

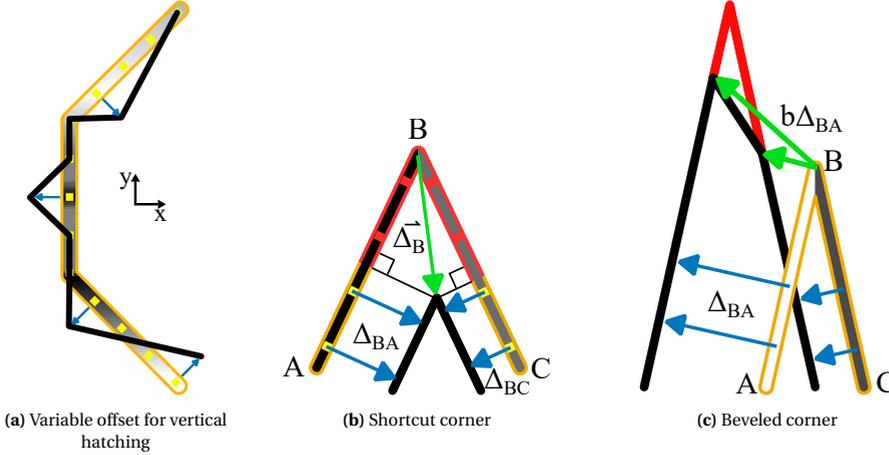


Fig. II.9. Visual explanation of variable offset applied to parts of a polygon. The original polygon is shown in orange, the resulting offset polygon in black, the sampling points in yellow and the offsets in blue. The grayscale values determine the amplitude of the offset. The green offsets show how variable offsetting of corners should be handled. The red items are omitted from the end result.

the vertex at the corner between two line segments BA and BC is given by

$$\vec{\Delta}_B = \frac{\Delta_{BA} \text{abs } \vec{B}\vec{A}\vec{B}\vec{C} + \Delta_{BC} \text{abs } \vec{B}\vec{C}\vec{B}\vec{A}}{\det[\vec{B}\vec{A}^T \vec{B}\vec{C}^T]} \quad (\text{II.8})$$

, where Δ_{BA} is the offset at B which follows from the texture coordinates at B on the mesh face which AB is located on and likewise for Δ_{BC} .

For inward offsets like the one shown in Fig. II.9b, the offset corner bypasses part of the connected line segments. Projecting the offset $\vec{\Delta}_B$ onto the line segments BA and BC gives the distance by which we disregard the sampling points along those line segments. The disregarded sampling points on the edges are not displaced; they are omitted from the variably offset polygon.

When using the above formula, sharp corners which take an outward displacement into account, could result in corners which are displaced by a distance greatly exceeding either displacement value. The resulting corner is therefore beveled when the displacement $\vec{\Delta}_B$ is larger than the bevel distance of both line segments. The bevel distance of a line segment $\vec{B}\vec{A}$ is given by $b\Delta_{BA}$, where $b > 1.0$ is a constant ratio which determines how much of the corner is beveled off. In Fig. II.9c a bevel ratio $b = 1.1$ was used.

Where line segments in the original outlines are close to each other and the offsets applied are larger than the distance between the line segments, the variably offset polygon would contain self-intersections. The self-intersecting parts of the polygon should be removed by applying a polygon clipping operation which uses a filling rule based on a positive winding number. For further reading, see [21].

II.4 RESULTS

Experiments were performed on an Ultimaker 3 machine, using black and white polylactic acid (PLA) – resp. Ultimaker 9014 and Ultimaker 9016. The print cores used in this setup



Fig. II.10. A 35×35 mm printed surface for various surface angles n optimized for a perpendicular viewing angle. Figure (b) shows a reproduction of the method by Reiner et al. on a vertical surface without tone calibration. Figure (l) shows the horizontal hatching technique. Image courtesy of ©Playboy Magazine 1972.

have a nozzle size of 0.4 mm and a line width setting of 0.35 mm was used. A default layer thickness of 0.1 mm was used. Different head movement speeds were used throughout the printing process; most notably, the outer walls were printed with a speed of 15 mm/s. In order to follow the outline as accurately as possible, the outer wall was printed before the inner walls. We used a sampling distance of 0.1 mm. Furthermore we applied a static offset of 0.1 mm to all outlines in order to prevent self-intersections due to negative offsets on either side of a thin outline polygon.

In an initial testing phase of sagging behavior multiple overhang distances o were tested at 0.01 mm intervals. After these experiments it seemed that applying an offset of 0.2 mm yields satisfactory results: 0.2 mm was the minimal tested offset at which black pixels in the texture appeared fully black when viewed from the side and vice versa for white. The sagging ratio was therefore estimated at $w/h = 2.0$, which was then used along with the geometrical sagging model to determine the offsets from texture color information.

The standard Lena test image (Fig. II.10a) has been printed on surfaces of different slope using our hatching model optimized for perpendicular viewing. The results can be seen in Fig. II.10.

Table II.1: Comparison of slice times of the total process planning, processing time spent on performing the variable offset alone and print times between normal FDM printing and our hatching technique on the prints presented in Fig. II.1.

model	normal		offset	hatching	
	slice	print		slice	print
Figurine	158 s	32 h37 min	32 s	300 s	36 h8 min
Portrait	69 s	23 h30 min	9 s	88 s	24 h20 min
Can	34 s	16 h28 min	6 s	47 s	17 h9 min
Rod	16 s	6 h35 min	4 s	21 s	7 h29 min

Tests on horizontal hatching culminated in a result which can be seen in Fig. II.10l, which shows the top of a test print containing a single wall and skin lines. The horizontal hatching technique was performed on the top black layer with a fully dense white layer below. The texture was sampled at 0.4 mm intervals. A line distance of 0.7 mm was used, a reference speed of 25 mm/s and a reference line width of 0.35 mm to produce the constant flow of $0.875 \text{ mm}^3/\text{s}$.

Figure II.1 shows hatched objects with geometric and organic shapes and with low and high frequency texture detail. The time it takes to print these objects using the 3D hatching techniques was up to 15% longer when compared to printing the same models with a single extruder. See Table II.1. This was unforeseen, provided that switching extruder on the Ultimaker 3 typically takes up less than 1% of the print time. We postulate that the difference in printing time is caused by the irregularity introduced by the variable offset. This irregularity causes the print head movement speed to be limited by its acceleration settings.

Because our hatching algorithms change the outlines by offsetting points sampled each 0.1 mm, the number of points in the outlines increases drastically. This further impacts a lot of the other process planning algorithms down the line. We can see in Table II.1 that applying hatching to the rhino figurine only added 20 % extra process planning time, while that caused an added slice time of 70 % down the line.

II.5 DISCUSSION

II.5.1. TEXTURE RESOLUTION

Figure II.10 shows the texture resolution obtained with the 3D hatching technique, compared to the input image (Fig. II.10a) and the technique of [2] (Fig. II.10b). Figure II.10c show a vertical print hatched for horizontal viewing. The vertical resolution depends on the layer thickness, which was 0.1 mm for all prints. The results obtained using vertical halftoning have a higher horizontal resolution than results obtained using the method by Reiner et al.: Fig. II.10b. Because vertical hatching is not limited by the wave length, the resolution of the resulting print is not limited by the line width, which is related to the physical size of the hole in the nozzle.

Figure II.10b was created using the technique from [2]. A wave length of 1.2 mm was used; a minimum amplitude of 0.375 mm and a maximum amplitude of 0.75 mm. Because the object is a simple cube for which each slice has the same square cross-section, projection threshold and relaxation do not play a role. No tone calibration has been performed; the amplitude depends linearly on the luminance of the texture input. The implementa-

tion is included in Cura [3].

Figure II.10c to (k) show result obtained using hatching optimized for perpendicular viewing on surfaces with various surface slopes n . The vertical resolution decreases with higher surface angles n , because the stair-step width gets larger. However, the grayscale tones remain roughly the same across the different surface slopes — notwithstanding the lack of tone calibration and the presence of visual artifacts.

Horizontal hatching shows a relatively low resolution. This is because the resolution is determined predominantly by the nozzle size, which was high in comparison to the layer thickness in our tests. Because the black lines sit on top of a white layer, the perceived tone tends toward black when viewing the surface from a lower altitude. Close examination of Fig. II.10l reveals that at places where the luminance value of the texture is high, black lines appear to be printed thinner, rather than narrower. It seems that these thinner lines only partially block light from the previous layer, resulting in a grayscale value close to the texture luminance even though their width is larger than intended.

It should be noted that all of the results in Fig. II.10 seem darker than the input image. This might be explained by several factors which affect the perceived luminance. Several such factors are mentioned in the remainder of this section.

II.5.2. DIMENSIONAL ACCURACY AND SLICING

Because the presented halftoning technique works by altering the geometry of a layer it affects the dimensional accuracy of the print. The size of the offset is always below $0.5d + w$, i.e. it is smaller than the stair step width aside for the extra offset required for sagging. This means that if it were not for sagging, the dimensional accuracy of the resulting object would be at least as good as the dimensional accuracy of a print job performed with twice the layer thickness. The maximum distance from the surface mesh to the surface of the printed product incurred by the stair-stepping effect is half the layer thickness. For our tests with a layer thickness of $h = 0.1$ mm the range in which our technique employs the sagging effect is $0.5w = 0.1$ mm. Since for perpendicular viewing the sagging is used when the stair step are narrow and less so when the stair steps are wide, the maximum dimensional error incurred is 0.1 mm. For near horizontal surfaces this inaccuracy incurred drops to 0.05 mm on top of the existing inaccuracy due to the stair-stepping effect.

The proposed hatching technique incurs minimal changes to the patterns by which conventional slicing software applications generate semi-continuous lines to build up a 3D printed object. The layer thickness remains unaltered and the layers are still built up by the walls, skin and infill. Keeping the elementary properties of FDM unaltered means that the structural properties of 3D prints are minimally influenced.

II.5.3. VISUAL ARTIFACTS

Variations in print parameters and inaccuracies in the printing system lead to various (visual) artifacts. These artifacts are particularly pronounced in hatching, due to the high contrast between consecutive layers. These artifacts cannot be well described with existing quality metrics for FDM, like (global) geometric accuracy and/or surface roughness (i.e. [22, 23]), both which do not capture the effect of the geometric variation on the perceived gray scale. Even though the sagging of vertical surfaces is employed within the range of 0.2 mm, we found the human eye can easily detect these inaccuracies, especially in regions with a uniform gray scale.

Figure II.11 shows the artifacts that could be distinguished in our experimental results.

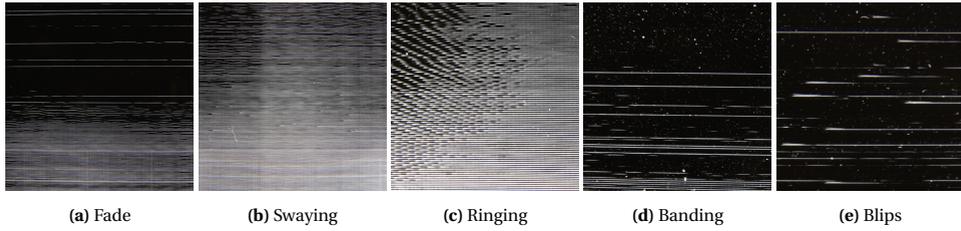


Fig. II.11. Various artifacts in the FDM printing process.

Note that there are potentially more types of defects, caused by different faulty elements in the FDM printing system and other variations in print parameters [24]. Vertical waves of lighter and darker colors, as can be seen in Fig. II.11a, might be caused by inaccuracies in the bed positioning or by temperature gradients due to the heated build plate. Slightly bent axes could account for low frequency horizontal wave patterns as can be seen in Fig. II.11b. Jerks of the print head can cause ringing, which results in a high frequency horizontal wave pattern in the perceived luminance of vertical surfaces (see Fig. II.11c). Banding – high frequency, vertical variations (see Fig. II.11d), can be caused by various (combined) effects, that stem from inaccuracies in the printing system and/or variation in printing parameters. Small discontinuities in the surface, shown in Fig. II.11e are caused by excess material oozing from the unused nozzle onto the model, while printing with the other nozzle. Inaccurate Z offsets between the two nozzles can lead to one of both filaments sagging more than the other, which can be seen in Fig. II.5, where the black layer is more compressed than the white layer. Visual artifacts in horizontal hatching (used in Fig. II.10l) may result from reduced layer bonding when printing thinner lines. Most of the inaccuracies described above, might be greatly reduced by using a single mixing nozzle, rather than an FDM system with two separate nozzles.

While these inaccuracies also affect the technique proposed by Reiner et al., the resulting grayscale tone of their technique is less sensitive to these. While the modulation of the amplitude of sine waves to the outer wall is limited to 0.75 mm our hatching technique is limited to $1/2w = 0.1$ mm. This means that the technique presented in this paper is affected 7.5 times as much by such inaccuracies.

One important observation from the sagging tests is that an overhang of 0.2 mm already causes a full layer of vertical occlusion. According to the geometrical sagging model this means that the area of the cross-section is increased by $\pi h^2 - \pi(0.5h)^2 - Cx.w.h \approx 0.018\text{mm}^2$. This means that sagging causes filament bleeding; extra filament is extruded, reducing the pressure in the nozzle, which impacts sagging in the subsequent region as well as influencing other aspects of the FDM process. This effect is not accounted for by our algorithm.

Other visual artifacts in hatching stem from the fact that a grayscale value based variable offset is applied on a layer-by-layer basis. The perceived grayscale value at a place in a given layer depends on the offsets applied at that place in the layer below and the layer above. Visual artifacts occur where the texture tones or face angles in an area wildly differ between layers. The implementation of vertical and diagonal hatching assume consecutive layers have the same layer thickness. Where this assumption is violated visual defects occur as well.

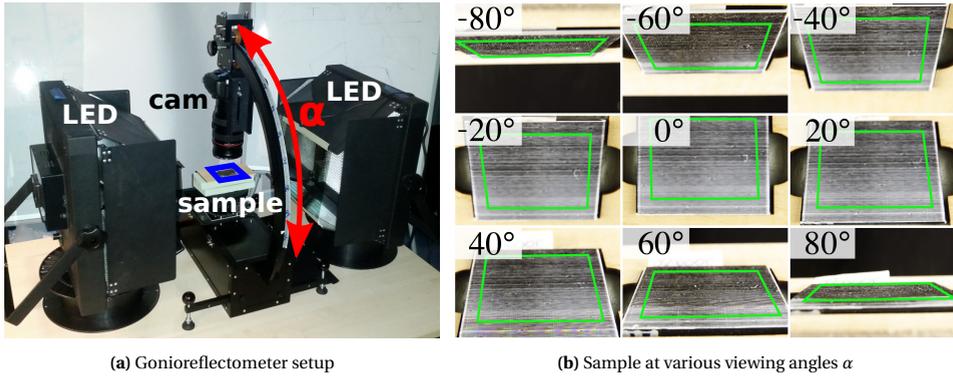


Fig. II.12. Measurement of printed samples

II.5.4. VIEWING ANGLE DEPENDENCY

Similar to the technique of Reiner et al., the perceived tone is dependent on the viewing angle. However, as the surface produced with 3D hatching is less irregular than with the dithering technique of Reiner et al., the perceived tone depends less on the azimuth of the viewing angle. In other words, when shifting the viewing angle sideways, the perceived grayscale values show little change, as you are looking along the print layers. The largest effect on the perceived grayscale is to be expected in the elevation viewing angle (ϕ in Fig. II.7). The relative offset between black and white consecutive layers and the sagging effect both lead to more occlusion with more extreme viewing angles.

In the following part the predicted ratios of visible filament according to our sagging model are compared to printed results for the case corresponding to $n = 0^\circ$ and $\phi = \alpha$ (see Fig. II.7). 18 vertical walls of 40×50 mm were printed, with a relative offset ranging from -0.2 mm to 0.2 mm. Note that the artifacts described above, were also present in these printed samples to varying degrees, which influence the measurement result.

In order to capture these prints under various angles, a Canon 5DS camera was installed in a hatching/gonioreflectometer setup. Samples were illuminated using two LED array light sources, placed at either side, in a cross-polarized setup, to avoid specular reflections as much as possible (see Fig. II.12a) [25]. The white balance of the camera was set using a completely white PLA print, and the exposure and aperture was set to avoid under- and over-exposure. Images were captured slightly out of focus to avoid moiré sampling effects, which are likely to occur due to the high frequency line patterns of the hatching technique.

Samples were captured at viewing angles α ranging between -80° and 80° with 10° increments. A masked region of 35×40 mm was sampled in all images. Note that foreshortening occurs in the images due to the measurement angle (see Fig. II.12b). The luminance of the sampled area was recovered by converting the RGB image to grayscale following the Luma standard [20]. The mean luminance was calculated for every sample, applying the gamma compression of 2.2 afterwards. The measured intensities are normalized relative to the minimum and maximum measured intensities.

Figure II.13b show the measured intensities for different overhangs from different viewing positions. Because of various visual artifacts and lighting conditions, comparing these results to Fig. II.13a does not yield the required insights to verify or falsify our

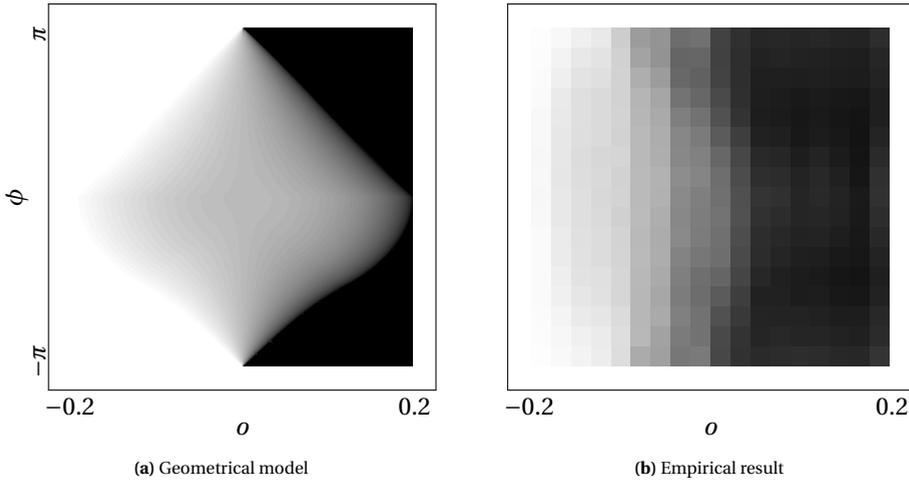


Fig. II.13. Luminance value for different amounts of overhang o and viewing angles ϕ as predicted by our model (a) and as obtained from experimental results (b).

geometrical model of sagging. However, looking at the differences between the two graphs can shed some light on the influence of the lighting conditions at play.

These results do not show a 50-50 gray line for all viewing angles when there is no offset applied. While a geometrical model should always satisfy such a condition, in reality less light is reflected to more extreme viewing angles than to a viewing angle perpendicular to the surface.

The fact that the graph in Fig. II.13b is not rotationally symmetric around $o = 0$ can also not be explained by a geometrical model of sagging which is equally applied to both filaments. Deviation from this symmetry could stem from multiple sources. In the microscope image (Fig. II.5) the black filament appears to sag in a different way from white filament; translucency and subsurface scattering might account for white filament appearing darker when enclosed by black filament; shadows might play a role and changing the shape of the outlines can contribute to different amounts of light being refracted in a specific direction. Understanding the influence of these factors is needed to calibrate the appearance of full color FDM prints.

II.6 CONCLUSION AND FUTURE WORK

We presented hatching halftoning techniques for 3D prints from dual extrusion FDM printing systems. We show how to perform offsets of varying amplitude on polygonal outlines and specifically how to determine the amplitude of offsets required for achieving the right grayscale tones on the surface of 3D prints. We investigate the phenomenon called *sagging* and propose a model for it, which serves as basis of a thorough tone calibration, which is future work. The results demonstrate the ability to manufacture objects with the appearance of full grayscale textures, while the techniques have little effect on printing time (maximum 20% more than monochrome). Because the way in which layers are built up is unaltered, structural properties of the resulting product are unaffected. The applicability of our hatching technique stands irrespective of the

geometry and surface angle. While it obtains a higher resolution than the technique presented by Reiner et al., it is also more sensitive to inaccuracies in the mechanical FDM printing system. This technique is applicable to any dual extrusion FDM system, although using a mixing nozzle can alleviate a lot of accuracy problems otherwise affecting the perceived grayscale tones.

One possible course of future research would be focused on improving on the unreliability caused by sagging. A dithering technique might be adopted for near horizontal surfaces and surfaces with extremal grayscale tones. The technique presented by Reiner et al. could serve this purpose once a way to gradually shift from their technique to the hatching technique for more horizontal surfaces. The ability to apply the dithering technique gradually for more extreme combinations surface angle and texture color is required in order for the algorithm to be stable; introducing a hard cutoff boundary between hatching and dithering introduces problems for surfaces around that cutoff boundary with slight variations in texture tone or surface angle.

A thorough tone calibration is needed to reproduce grayscale tones correctly. In order to get a grip on the different factors which influence the perceived tone, different factors could be investigated, for example material translucency, subsurface scattering, shadows induced by the offsetting and various lighting conditions. Together with a verified geometrical model of sagging, the tone calibration results can then be explained. The formula by which the print speed is computed from the line width in horizontal hatching also provides an opportunity for tone calibration. Given that thin lines appear translucent, additional tone calibration might impact the required amount of material per line segment, which affects the line width modulation.

One could adopt a similar technique for FDM printers which have more than two extruders; if it is capable of printing with cyan, magenta, yellow, black and white filament, hatching could be used to produce full color prints. When different or less colors are available, a mapping between color spaces should be performed to make the print appear as close to the textured mesh as possible.

A different line of research could be devoted to determining the optimal viewing angle from the model geometry and the texture image. A more natural viewing angle could be obtained from the surface normal vector of the nearest point on the convex hull of the full geometry. Where the texture image has more high frequency detail the optimal viewing angle could be determined more by the local surface angle of the model. A deeper investigation of how humans look at 3D objects is needed to determine what the optimal viewing angle is at any point on the surface.

This paper presents techniques to print heterogeneous surface colors on FDM printing systems. Future endeavors could be devoted to enabling other heterogeneous surface properties for FDM, such as surface roughness and specularly. Going beyond just the surface of the mesh it could be researched how heterogeneous volumetric properties could be achieved in FDM. What process planning techniques can be used to satisfy a heterogeneous infill density specification?

ACKNOWLEDGEMENTS

We would like to thank Leo Haslam (Blockade figurine), Beerend Groot (tin can) and COM-SOL (connecting rod) for permission to use their models. We would also like to thank Jun Wu and Charlie Wang for suggestions and improvements.

Funding: This work was supported by Ultimaker.

REFERENCES

- [1] T. Kuipers et al. “Hatching for 3D prints: Line-based halftoning for dual extrusion fused deposition modeling”. In: *Computers & Graphics* 74 (2018), pp. 23–32. ISSN: 00978493. DOI: 10.1016/j.cag.2018.04.006.
- [2] Tim Reiner et al. “Dual-color mixing for fused deposition modeling printers”. In: *Computer Graphics Forum*. Vol. 33. 2014, pp. 479–486.
- [3] Ultimaker. *Ultimaker Cura 4.2.1 software*. Version 4.2.1. Aug. 31, 2019. DOI: 10.5281/zenodo.1216744.
- [4] Tim Kuipers, Eugeni Doubrovski, and Jouke Verlinden. “3D hatching: linear halftoning for dual extrusion fused deposition modeling”. In: *Proceedings of the 1st Annual ACM Symposium on Computational Fabrication*. ACM. 2017, p. 2. DOI: 10.1145/3083157.3083163.
- [5] Z Corporation. *Z Corporation 3D Printing Technology – fast, affordable and uniquely versatile*. 2005. URL: https://www.ucy.ac.cy/arch/documents/3d%7B%5C_%7DPrinter%7B%5C_%7DLab/3D%7B%5C_%7DPrinting%7B%5C_%7DTechnology.pdf.
- [6] Mihaela Iliescu et al. “Z Printing Rapid Prototyping Technique and SolidWorks Simulation—Major Tools in New Product Design”. In: *Proceedings of the 2nd WSEAS International Conference on Sensors, and Signals and Visualization, Imaging and Simulation and Materials Science*. World Scientific, Engineering Academy, and Society (WSEAS). 2009, pp. 148–153.
- [7] Mcor. “How Paper-based 3D Printing Works: The Technology and Advantages”. In: (2013). (accessed April, 2017). URL: http://mcor technologies.com/wp-content/uploads/2013/04/MCOR-WP-19032013-EU%7B%5C_%7Dlow.pdf.
- [8] Stratasys. *Stratasys J750*. 2016. URL: http://usglobalimages.stratasys.com/Main/Files/Machine%7B%5C_%7DSpec%7B%5C_%7DSheets/PSS%7B%5C_%7DPJ%7B%5C_%7DStratasysJ750%7B%5C_%7D0217a%7B%5C_%7DWeb.pdf (visited on 02/01/2018).
- [9] HP. *HP Multi Jet Fusion technology*. 2014. URL: h41367.www4.hp.com/campaigns/ga/3dprinting/4AA5-5472ENW.pdf (visited on 07/04/2017).
- [10] Qun Lou and Peter Stucki. “Fundamentals of 3D Halftoning”. In: *Electronic Publishing, Artistic Imaging, and Digital Typography*. Springer Berlin Heidelberg, 1998, pp. 224–239.
- [11] W Cho et al. “A dithering algorithm for local composition control with three-dimensional printing”. In: *Computer-Aided Design* 35.9 (2003), pp. 851–867. ISSN: 00104485.
- [12] Kiril Vidim, Szu-Po Wang, and Jonathan Ragan-kelley. “OpenFab: A Programmable Pipeline for Multi-Material Fabrication”. In: *ACM Transactions on Graphics* 32.4 (2013), pp. 1–11. ISSN: 0730-0301.

- [13] A Brunton, C A Arikan, and P Urban. "Pushing the limits of 3d color printing: Error diffusion with translucent materials." In: *ACM Transactions on Graphics (TOG)* 35.1 (2015), p. 4.
- [14] James Corbett. "Reprap colour mixing project". In: *Final Year MEng Project, Department of Mechanical Engineering, Faculty of Engineering and Design, University of Bath, Bath* (2012).
- [15] Haichuan Song and Sylvain Lefebvre. "Colored fused filament fabrication". In: *ACM Transactions on Graphics* 38.5 (2017), pp. 1–11. ISSN: 07300301. DOI: 10 . 1145 / 3183793. arXiv: 1709 . 09689. unpublished manuscript.
- [16] Claude Mellan. *The Sudarium*. 1649. URL: <http://hdl.handle.net/10934/RM0001.collect.152641>.
- [17] Irving Lavin. "Claude Mellan's Holy Face". In: (2004).
- [18] Emil Praun et al. "Real-time hatching". In: *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*. ACM. 2001, p. 581.
- [19] Akio Yamamoto. *Producing engraving-style halftone images*. US Patent 7,126,723. 2006.
- [20] Charles Poynton and Garrett Johnson. "Color science and color appearance models for CG, HDTV, and D-CINEMA". In: *ACM SIGGRAPH 2004 Course Notes*. ACM. 2004, p. 3.
- [21] Bala R Vatti. "A generic solution to polygon clipping". In: *Communications of the ACM* 35.7 (1992), pp. 56–63. ISSN: 15577317. DOI: 10 . 1145/129902 . 129906.
- [22] Antonio Armillotta and Marco Cavallaro. "Edge quality in Fused Deposition Modeling: I. Definition and analysis". In: *Rapid Prototyping Journal* 10.1 (2017), pp. 5–6. ISSN: 1355-2546. DOI: 10 . 1108/RPJ-02-2016-0020.
- [23] Min Kyung Kim, In Hwan Lee, and Ho-Chan Kim. "Effect of fabrication parameters on surface roughness of FDM parts". In: *International Journal of Precision Engineering and Manufacturing* 19.1 (Jan. 2018), pp. 137–142. ISSN: 2005-4602. DOI: 10 . 1007/s12541-018-0016-0.
- [24] Brian N. Turner, Robert Strong, and Scott A. Gold. "A review of melt extrusion additive manufacturing processes: I. Process design and modeling". In: *Rapid Prototyping Journal* 20.3 (2014), pp. 192–204. ISSN: 1355-2546. DOI: 10 . 1108/RPJ-01-2013-0012.
- [25] Eugene Hecht. "Chapter 8: Polarization". In: *Optics*. 4th Edition. San Francisco: Addison-Wesley, 2001. Chap. 8. ISBN: 978-0805385663.

III

SPATIALLY GRADED COMPLIANCE

In the previous chapter we presented a method to manufacture products with a spatially graded grayscale tone on the surface. We now shift our attention to the second category of FGM: volumetric properties. This chapter proposes an FGM with spatially graded compliance, called *CrossFill*. By increasing the density of the structure in a region, that region becomes more stiff – i.e. less compliant. Because compliance is the primary focus, we aim to print such structure using a flexible material, such as thermoplastic polyurethane (TPU). Because TPU is flexible, the extrusion process control is difficult to master, which makes retracting more problematic than in other materials. We therefore propose a method which produces strictly continuous constant width toolpaths. We propose a novel space-filling surface fractal which when sliced produces long continuous polygons. By adjusting the recursion depth of the fractal locally, we can generate a structure which has a spatially graded compliance.

Sections III.1 and III.2 introduce the method and discuss related work. Section III.3 gives an overview of the technique to generate spatially graded density structures. Section III.4 introduces the fractal used for *CrossFill* and Section III.5 explains how we can use adaptive subdivision levels of the fractal to adhere to a given density specification. Section III.6 explains how we can extract continuous extrusion toolpaths from the generated fractal structure. Section III.7 presents experimental results and discusses them, which leads us to the conclusion in Section III.8.

This chapter is published as Tim Kuipers, Jun Wu, and Charlie C.L. Wang. "CrossFill: Foam Structures with Graded Density for Continuous Material Extrusion". In: *Computer-Aided Design* 114 (Sept. 2019), pp. 37–50. ISSN: 00104485. DOI: 10.1016/j.cad.2019.05.003; some minor corrections have been made.

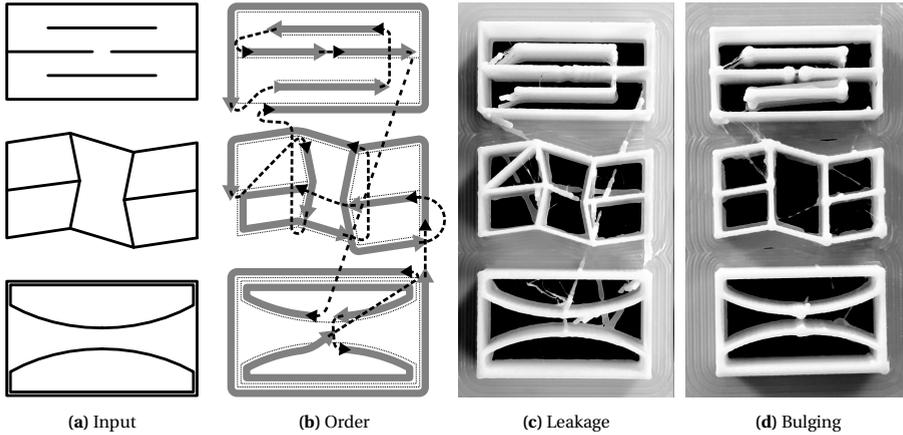


Fig. III.1. Discontinuity in extrusion paths causes defects. (TPU, Ultimaker 3) End points (top), T-junctions (middle) and isolated polygons (bottom) all introduce discontinuity in the extrusion process and rapid travel moves are required. (c) When the extrusion motor stops, the material still leaks out during rapid travel moves. (d) Retraction reduces leakage, but instead introduces bulging.

ABSTRACT

The fabrication flexibility of 3D printing has sparked a lot of interest in designing structures with spatially graded material properties. In this paper, we propose a new type of density graded structure that is particularly designed for 3D printing systems based on filament extrusion. In order to ensure high-quality fabrication results, extrusion-based 3D printing requires not only that the structures are self-supporting, but also that extrusion toolpaths are continuous and free of self-overlap. The structure proposed in this paper, called *CrossFill*, complies with these requirements. In particular, *CrossFill* is a self-supporting foam structure, for which each layer is fabricated by a single, continuous and overlap-free path of material extrusion. Our method for generating *CrossFill* is based on a space-filling surface that employs spatially varying subdivision levels. Dithering of the subdivision levels is performed to accurately reproduce a prescribed density distribution. We demonstrate the effectiveness of *CrossFill* on a number of experimental tests and applications.

III.1 INTRODUCTION

3D printing enables the fabrication of complex structures with unprecedented geometric detail. This creates the opportunity to realize 3D shapes with complex internal structures. Physical properties of these *infill structures* are determined by their geometry and the constitutive material by which they are made. Even with a single constitutive material, 3D printing allows to achieve graded physical properties (e.g. density and stiffness) by spatially varying the geometry of infill structures. This enables *functionally graded materials* (FGM) at a manufacturable scale. Precise realization of graded physical properties can lead to many applications, such as customized insoles, comfort cushioning and medical phantoms.

Fused deposition modeling (FDM) is one of the most widely used 3D printing processes as it has a comparatively low running cost and supports a wide variety of materials. FDM

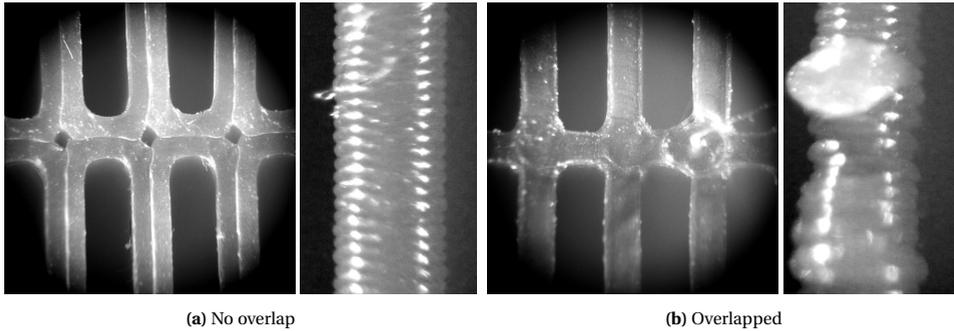


Fig. III.2. Microscopic photos of top and side views of printing results with a 0.38 mm wide extrusion path: (a) without versus (b) with overlapping by 0.36 mm respectively. Overlapping extrusion paths exhibit over-extrusion of material at the overlapping region, which results in unwanted blobs on the surface of the print.

systems work by extruding melted streams of material from a moving nozzle to form a quickly solidified path. However, there is limited study on using FDM to reliably fabricate FGM. This task is challenging since the complicated geometry of a functionally graded infill structure is difficult to meet the different manufacturing constraints required by FDM to ensure printing quality:

- *Overhanging geometry:* If a geometric feature is not properly supported by lower layers, it is said to be overhanging. While overhanging geometry can be printed using support structures, the complexity of infill structures inside a 3D model does not allow an easy removal of support structures. Therefore, infill structures are expected to be self-supporting.
- *Discontinuous material extrusion:* In extrusion-based fabrication, frequent restarting and stopping extrusion creates defects. Simply stopping the extrusion motor will lead to material leakage. One common way to prevent that is to retract the filament backward a bit before starting a rapid travel move, but that in turn introduces bulges where the extrusion paths start and end. See Fig. III.1. In order to fabricate infill structures reliably, it is desired that each layer of the structures is fabricated by continuous extrusion along a single toolpath without any interruption.
- *Overlapping extrusion paths:* Since material is extruded along a toolpath typically with a constant width, there is an excess of material in a layer when extrusion paths are too close to each other. The overlap of extrusion paths causes blobs and wider lines, which make it difficult to control the density of infill structures in corresponding regions (see Fig. III.2). It is preferred to solve this problem intrinsically by generating infill structures without overlapping toolpaths.

In order to use FDM for fabricating foam structures with graded density inside a given model, a method needs to be developed for generating infill structures according to a user-specified density distribution, which should also avoid the above manufacturing problems.

In this paper, we propose a novel type of foam structure that can achieve the aforementioned objective. Specifically, we develop a space-filling surface, called *CrossFill*, an

FDM printable foam structure as infill for a 3D model. Each layer of CrossFill is a space-filling curve that can be continuously extruded along a single overlap-free toolpath. The space-filling surface consists of surface patches which are embedded in prism-shaped cells, which can be adaptively subdivided to match the user-specified density distribution. The adaptive subdivision level results in graded mechanical properties throughout the foam structure. Our method consists of a step to determine a lower bound for the subdivision levels at each location and a dithering step to refine the local average densities, so that we can generate CrossFill that closely matches the required density distribution. A simple and effective algorithm is developed to merge a space-filling curve of CrossFill of a layer into the closed polygonal areas sliced from the input model. Physical printing tests have been conducted to verify the performance of the CrossFill structures.

Our approach provides three technical contributions:

- A novel self-supporting space-filling surface which supports spatially graded density;
- A scheme for refining the structure to match a prescribed density distribution;
- An algorithm for merging the toolpath of an infill structure with the input model's boundary so as to retain continuity.

III.2 RELATED WORK

For an overview of techniques involved in 3D printing let us refer to two survey articles [2, 3]. In this section we review the design of microstructures and in particular for extrusion-based 3D printing, as well as the use of space-filling curves and surfaces.

III.2.1. STRUCTURES WITH GRADED PROPERTIES

A variety of graded structures have been proposed in recent years, including lattice structures with varying thickness (e.g. [4, 5, 6, 7, 8]), triply periodic minimal surfaces (e.g. [9, 10]), free-form microstructures [11], microstructures for expressive deformation (e.g. [12, 13, 14]), bone-like microstructures (e.g. [15, 16]) and microstructures optimized by inverse homogenization (e.g. [17, 18]).

Most of these complex structures are fabricated with powder-based 3D printing systems such as selective laser sintering (SLS) or with stereolithography (SLA). Density gradation is typically achieved by varying the thickness of geometric primitives. However, reliably fabricating microstructures with varying thickness is challenging for extrusion-based 3D printing. When printing beads narrower than the nozzle size, it is difficult to predict at which location exactly the bead will end up; furthermore, when the geometry is wider than the nozzle size the toolpath generation needs to switch from a single bead into several in a controlled manner. Furthermore, lattice structures are sliced into small disconnected components for each layer which violates the continuous extrusion constraint (e.g. [4, 13, 19]).

III.2.2. STRUCTURES FOR EXTRUSION-BASED PRINTING

3D shapes fabricated by FDM typically comprise uniform infill structures. Recently Martínez et al. proposed a method to generate infill structures with spatially graded

density by printing the cell membranes of 3D Voronoi diagrams [20]. The cells center locations are randomly sampled from a 3D user-specified probability distribution in order to create the spatially graded infill. Overhang constraints are satisfied by carefully constructing a distance measure which forms the basis of defining the cell membranes.

The generated structures are limited by the following factors. The extrusion paths are not continuous; the amount of retractions reduces reliability and increases print time. The density is controlled indirectly through the density of the cell centers. The actual relation between the two remains unclear. At high densities the method is likely to generate overlapping extrusion paths, leading to over-extrusion, which causes defects in the print. These problems are well resolved by our approach.

Wu et al. proposed using a subdivision grid of slanted cubes called *rhombuses* for extrusion-based 3D printing [21] and proposed optimizing the subdivision structure for stiffness [22]. The sides of the rhombuses can then be printed using a single bead. However, the T-junctions require retractions, which are problematic especially for flexible materials. Our method makes use of a subdivision grid as well, but generates a continuous toolpath. Moreover, compared to the rigid rhombic structures, the CrossFill structure is more compliant and acts like a foam, which is beneficial for several applications such as an insole.

In this paper we assert the density distribution is prescribed by the user, and present a method to reliably reproduce the distribution using extrusion-based printing. The graded density can be specified by the user or by an optimization process [15, 23].

III.2.3. SPACE-FILLING CURVES AND SURFACES

Our method makes use of space-filling surfaces, which are analogous to space-filling curves in 2D (e.g. Hilbert curve [24], Sierpiński curve [25, 26]). Using space-filling curves with varying degrees of subdivision level has been explored for purposes other than 3D printing, such as robotic exploration tasks [27], finite element analysis [28, 29], paths for CNC milling [30]. In the context of 3D printing, Kumar, Pandithevan, and Ambatti combined several square based space-filling curves (e.g. Hilbert curve) to generate porous infill structures with spatially graded density and semi-continuous extrusion [31]. However, this method allows for spatially varying density only in the horizontal plane – not in the vertical direction.

While there is much literature on extending such space-filling curves to *polylines* in 3D, a space-filling curve in 2D can also be extended into a space-filling *surface* in 3D. Space-filling surfaces are first defined by Ahmed and Bokhari [32]. Although the space-filling surfaces are continuous, any layer-wise cross section is still discontinuous. In this paper we propose a new type of space-filling surfaces, which provides cross sections that are continuous.

III.3 OVERVIEW

CrossFill is a space-filling surface that is constructed using subdivision rules on prism-shaped cells. Each cell contains a patch of the surface, which is sliced into a line segment on each layer to be a segment of the extrusion toolpath. Since the toolpath will be fabricated with a *constant* width, the size of a cell determines the regional fraction of solid material (hereafter referred to as ‘density’). By adaptively applying the subdivision rules to the prism cells, we create a subdivision structure of cells with a density distribution that

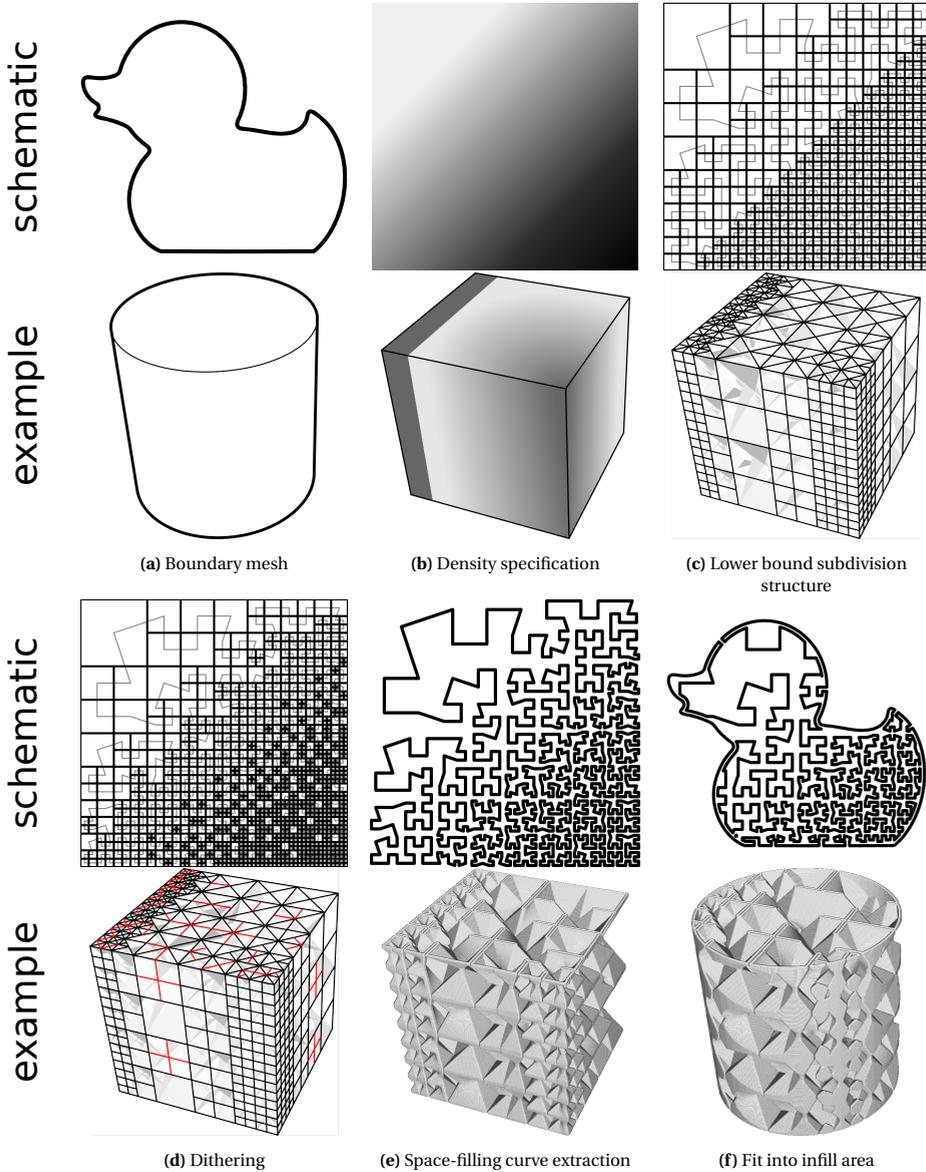


Fig. III.3. Schematic overview of our method. The top rows show a 2D analogue of our method for clear visualization. The prism-shaped cells in the bottom row are visualized as semi-opaque solids to keep the visualization uncluttered. Red lines in the bottom row highlight the local subdivisions performed in the dithering phase. Note that, the shell of the 3D model is not displayed in (f) for the illustration purpose.

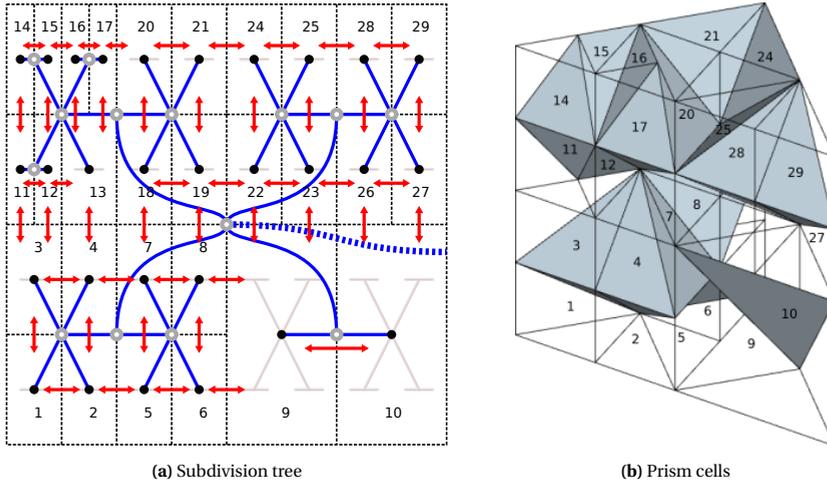


Fig. III.4. The hybrid tree and graph data-structure employed for CrossFill, where the blue links indicate the subdivision relationship between prism cells and the black dots represent the leaf-nodes on the tree each of which embeds a patch of the space-filling surface. The red arrows denote the connectivity graph for easily traveling from one cell to its neighbors.

closely matches a user-specified input. Continuity of the space-filling surface across adjacent cells with different subdivision levels – both horizontally and vertically – is ensured by the subdivision rules and by post-processing of the surface patches in neighboring cells.

Figure III.3 provides an overview of our method using a simple 3D example (bottom) and a 2D schematic illustration (top). From a user-specified 3D density field (Fig. III.3b), we first increase the subdivision levels everywhere until one further subdivision would result in an average cell density higher than the average requested density in that region (Fig. III.3c). The resulting subdivision structure forms the *lower bound* of the final subdivision levels. In order to closely match the input density distribution, we develop a dithering method in which we alternate the subdivision level between the lower bound and one level deeper (Fig. III.3d). Once the subdivision structure is finalized, we slice it into space-filling curves for the toolpaths on each layer (Fig. III.3e). In this step, we adjust the surface patches in the cells such that overlapping toolpaths are prevented. Lastly, the space-filling curves are trimmed to the infill area and connected to the shell of the input 3D models to form the final toolpaths preserving continuous extrusion (Fig. III.3f).

III.4 CROSSFILL

CrossFill is an infill structure which consists of a space-filling surface. CrossFill is so named because the toolpath of this structure resembles crosses (see Fig. III.9).

III.4.1. INITIALIZATION AND DATA-STRUCTURE

The space-filling surface of CrossFill is generated using a subdivision scheme that starts from a cube which encompasses the input 3D model with side lengths $2^i w$ for integer values of i and a constant extrusion width w . This cube is divided into four prisms by splitting along the two diagonals of the horizontal faces. Starting from these four prisms, the prisms are adaptively subdivided into smaller ones. In lieu of the commonly used 1:8

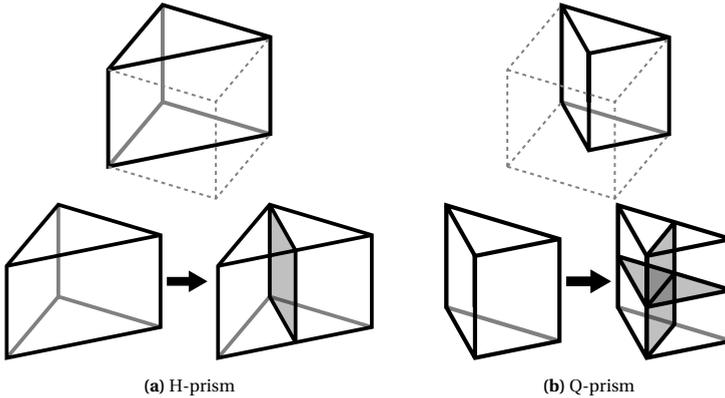


Fig. III.5. The basic types of prisms in CrossFill: the shape of an H-prism can be obtained by cutting a cube in *half* and the shape of a Q-prism by cutting a cube in *quarters*. A H-prism can be subdivided into two Q-prisms and a Q-prism is always subdivided into four H-prisms in our system.

subdivision of cubic cells, we make use of 1:2 and 1:4 subdivision of prism-shaped cells. This allows more granularity, which is beneficial for matching the requested density distribution. Details on the types of prism-shaped cells and subdivision rules will be presented in Section III.4.2 and III.4.3, respectively.

The hierarchy and connectivity of cells are encoded by a combination of a tree and a graph. The *subdivision tree* connects a cell to its subdivided constituent cells. The root node of the tree corresponds to the starting cube, while other nodes correspond to prism cells. The leaf nodes of the tree constitute the current subdivision structure. The *connectivity graph* stores the connectivity information among the leaf cells. Two neighboring leaf cells are linked if their surface patches are connected through the space-filling surface; however, two cells which share a face which is not crossed by the space-filling surface are not linked. The links store the relative spatial location of the neighboring cell – up, down, left or right along the space-filling surface. See Fig. III.4 for an illustration. This connectivity graph facilitates efficient traveling between neighboring cells. The tree and graph are updated each time a new subdivision is applied to a leaf node.

III.4.2. TYPES OF CELL

The construction of CrossFill depends on a subdivision tiling consisting of prism-shaped cells. As shown in Fig. III.5, two types of prism are distinguished in our subdivision system:

H-prism is constructed by vertically cutting a cube in *half* along a diagonal of the horizontal faces.

Q-prism is generated by splitting a cube into *quarters* along both diagonals of the horizontal faces.

The subdivision tiling is generated by subdividing an H-prism into two Q-prisms and subdividing a Q-prism into four H-prisms (see the bottom row of Fig. III.5).

Each prism cell encompasses a triangular patch of the space-filling surface. The prism cells are categorized according to which side faces of the prism are crossed by the triangle surface patch (see the horizontal cross sections visualized in Fig. III.7a):

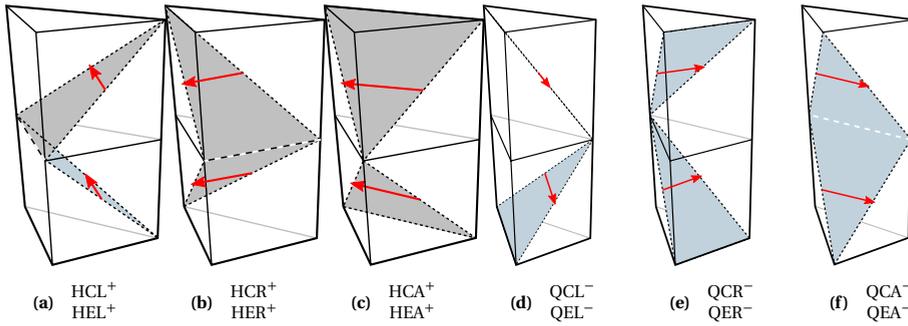


Fig. III.6. The types of prism cells by considering the embedded space-filling surfaces.

A-route the surface is spanned across the faces connected to the two catheti of the right triangle at the base of the prism.

L-route the surface crosses between the face connected to the hypotenuse and the left cathetus.

R-route the surface crosses between the face connected to the hypotenuse and the right cathetus.

In order to keep track of the spatial ordering between cells when subdividing and when updating the connectivity graph, we introduce the horizontal direction of traversal (see Fig. III.6):

+ **direction** travel from left to right, from left to hypotenuse or from hypotenuse to right.

– **direction** travel from right to left, from right to hypotenuse or from hypotenuse to left.

In order to fully characterize a cell, we also distinguish between the two fashions in which the 3D surface patch is embedded in the prism (compare the bottom and top in Fig. III.6):

E-embedding when vertically exploring a cell from bottom to top by horizontal cross sections, the embedded surface is said to be *expanding* if it is moving from the right of the traveling direction to the left.

C-embedding the embedded surface is *contracting* otherwise. (Note that the embedding is not defined in terms of whether the triangle surface patch points up or down.)

In total we consider 12 different types of prism cell; see Fig. III.6. It can be easily verified that the overhanging angle of embedded surfaces in all types of prism cell is always less than 45° – i.e., the *self-supporting* constraint is satisfied.

III.4.3. SUBDIVISION RULES

We now define the subdivision rules on the prism cells, which depend on the type of prism (H- or Q-prism), the route type (A-, L- or R-route), the direction of traversal (+ or –) and

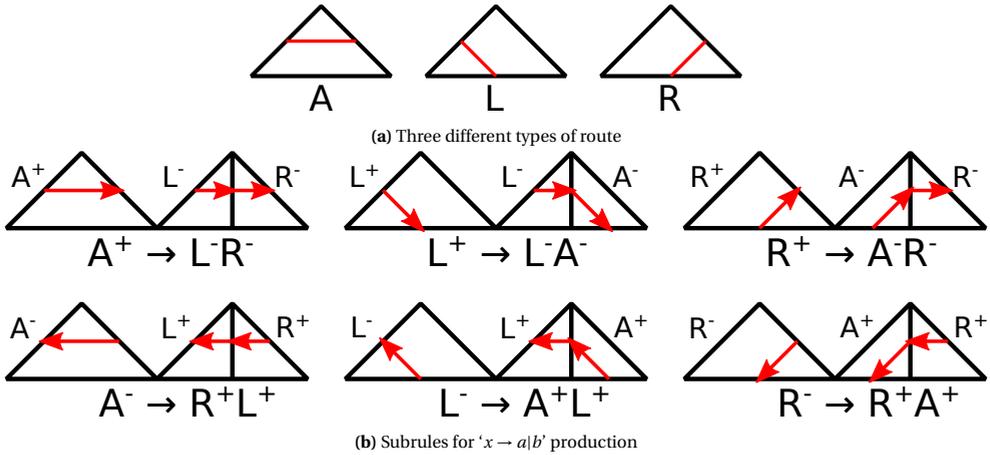


Fig. III.7. Types of surface patch distinguished by the route (which sides of the cell the segment crosses) and the direction. When a cell is subdivided, the type of route x is substituted with different routes in the two newly constructed cells as a and b respectively.

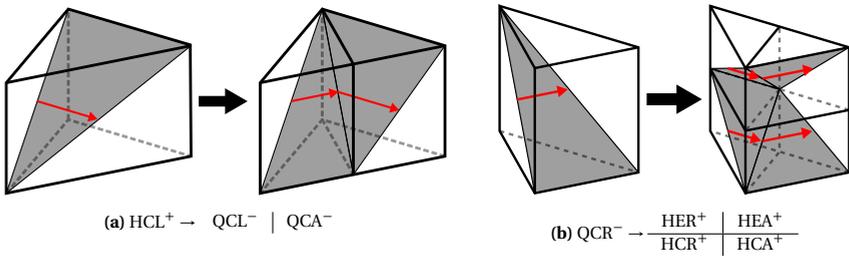


Fig. III.8. Examples of combined subdivision rules.

the type of embedding (C- or E-embedding):

$$\begin{array}{l}
 HCLx \rightarrow QCa \mid QCb \quad HEx \rightarrow QEa \mid QEb \\
 QCx \rightarrow \frac{HEa \mid HEb}{HCa \mid HCb} \quad QEx \rightarrow \frac{HCa \mid HCb}{HEa \mid HEb}
 \end{array} \tag{III.1}$$

where the pattern ' $x \rightarrow a|b$ ' can be filled with any of the subrules for route type and traveling direction as given in Fig. III.7b. The spatial ordering on the right hand side of the rule indicates the spatial ordering of the cells: on top of each other and horizontally next to each other when following the route along its direction. Two examples of filled subdivision rules described by Eq. (III.1) are illustrated in Fig. III.8.

We initialize the CrossFill fractal with four QCA^- cells such that the surface patches form a pyramid. Together with the subdivision rules this forms a system closely related to an L-system.

Relation to 2D L-systems As illustrated in Fig. III.9a and (b), a subdivision tiling of right triangles and an embedded 2D space-filling curve can be obtained when intersecting a uniformly subdivided CrossFill structure with a horizontal plane at an altitude of half

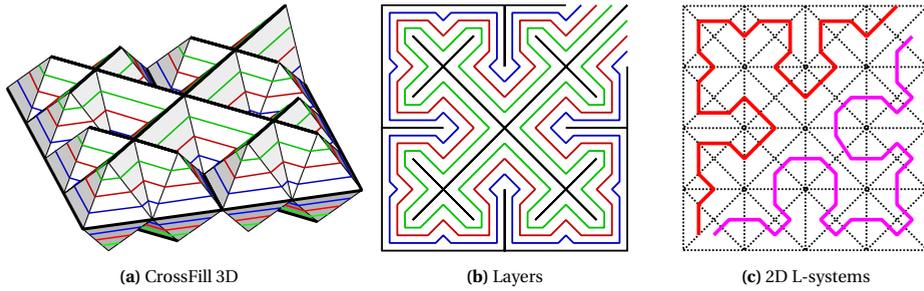


Fig. III.9. Intersecting the uniformly subdivided CrossFill (a) by horizontal planes results in space-filling curves (b). The red space-filling curve can be obtained by applying the subrules of Fig. III.8b. The Sierpiński curve is visualized in purple in (c).

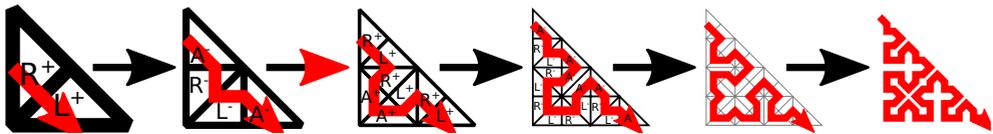


Fig. III.10. Repeated application of the subdivision rules defined in Fig. III.7b. Denser and denser space-filling curves can be generated from the subdivided space tiling.

the prism’s height. The rules used above with integrated traversal information can help generate a 2D space-filling curve similar to the Sierpiński curve (see the purple curve in Fig. III.9c). Whereas the Sierpiński curve is generated by connecting the centers of the triangle cells, the CrossFill pattern is generated by connecting the vertices located on the edges which are crossed by the curve. Compared to the 1:4 subdivision of square cells (e.g. for constructing the Hilbert curve), the 1:2 subdivision of triangles allows more granularity. Similarly, our 3D subdivision rules provide more granularity than a cubic 1:8 subdivision. As illustrated in Fig. III.10, repeatedly denser space-filling curves can be obtained when reapplying the subdivision rules defined in Fig. III.7b.

III.4.4. COMPATIBILITY AND CONTINUITY

The surface patches embedded in a subdivision structure with uniform subdivision level form a continuous space-filling surface. With adaptive subdivision levels, neighboring cells can have different sizes. The boundary of the surface patch embedded in a big cell does not always match with the boundary of surface patches in neighboring linked smaller cells (e.g. prisms 10 and 27 in Fig. III.4).

To solve this problem in general is difficult; however, it is solvable when neighboring linked cells have only one level difference in the subdivision tree, because there is only a limited set of configurations. The continuity can be enforced on such a structure with heterogeneous subdivision level in three steps.

Step 1: Horizontal continuity enforcement When horizontally neighboring cells have a different subdivision level their embedded surface patches may not match at the sides where the cells meet. For example the edge **a** in Fig. III.11 does not align with **b** and **c**. In such a case the edge of the lower subdivision level cell is transformed such that it matches the higher subdivision level cells and its surface is converted into a ruled surface

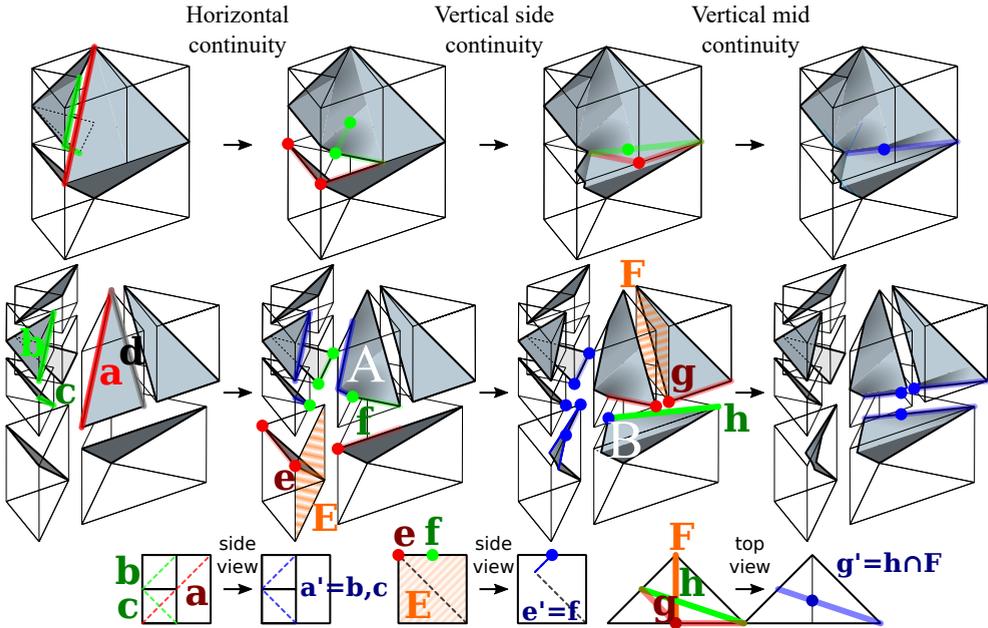


Fig. III.11. Three steps for adjusting the space-filling surface to enforce the continuity between neighboring linked cells. Adjustments are made to edges on the sides of the surface patches and the patches are thereby transformed into ruled surfaces. Red elements (vertices and edges) are to be changed into the corresponding blue elements in order to match the green elements in this figure. (top) Assembled view. (middle) Exploded view. (bottom) Closeups of the interface between two cells.

\mathcal{A} : on each layer we connect the one edge (\mathbf{d}) of the surface patch with either of the two edges (\mathbf{b} or \mathbf{c}) of the smaller surface patches using a horizontally straight line segment. See Fig. III.11 and Fig. III.12a.

Step 2: Vertical side continuity enforcement When vertically neighboring cells have a different subdivision level the edges of their surfaces may not end in the same location at the interface where the cells meet. For example the vertex \mathbf{e} does not coincide with \mathbf{f} in Fig. III.11. In such a case the edge(s) of the lower subdivision level cell are transformed: part of the edge is flipped horizontally in the plane \mathbf{E} along the cell side where the edge resides. This adjustment is also performed if the discontinuity was introduced because of the horizontal continuity enforcement, as is illustrated by surface patch \mathcal{B} in the middle of Fig. III.11.

Step 3: Vertical mid continuity enforcement When vertically neighboring cells have different subdivision level and have ruled surfaces the horizontal edges may not align on the horizontal side where the cells meet. For example vertex \mathbf{g} does not lie on the edge \mathbf{h} in Fig. III.11. In such a case the vertices of the higher subdivision level cells are adjusted to lie on the intersection between edge \mathbf{h} and the side \mathbf{F} where the two higher subdivision cells meet horizontally. Similar to vertical side continuity enforcement, we flip part of the edge to which the adjusted vertex belongs and introduce ruled surfaces.

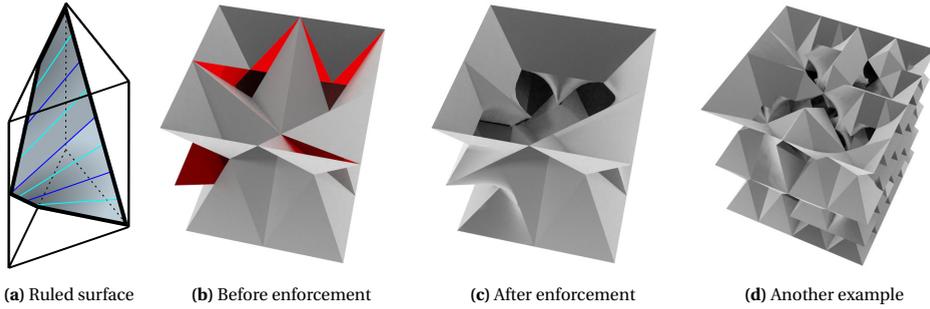


Fig. III.12. The impact of continuity enforcement on CrossFill structures. The enforcement causes surface patches to become ruled surfaces. Blue lines show the isolines which are straight at each Z height. Red patches show discontinuities in the space filling surface.

The space-filling surface in linked cells with subdivision level differences can be effectively enforced to be continuous by this approach. Examples of CrossFill structures with continuity enforcement are shown in Fig. III.12. For more extensive examples we refer to the video in the supplementary material. Note that in the implementation we do not actually construct the surface; instead, we compute the vertex locations of the edge segments of the surface patches, slice those at the height of a printing layer and connect the resulting locations using straight line segments.

Self-supporting It should be noted that the curved surfaces introduced by enforcing the continuity of neighboring linked prism cells will not violate the self-supporting property of the space-filling surface. With the help of the carefully designed continuity enforcement algorithm, we generate surfaces that have overhang $\leq 45^\circ$ in most places. There is only one exceptional case for the side enforcement in an H-prism, where the overhang can be increased to $\tan^{-1} \sqrt{2} \approx 55^\circ$. However, geometry overhanging with an angle of 55° is not a problem for most FDM 3D printers, so the self-supporting constraint is not violated.

Density Sudden jumps in density are hard to be realized on our infill structure. For example, a density distribution which is 10% in the bottom half and 80% in the top half is not easily realized while satisfying the overhang constraint. Our space-filling surface requires some distance to change from the low to the very high density along the vertical direction. Also, the surface patch with enforced continuity would be considerably different from the original triangular surface patches, which might have a large influence on the physical properties associated with a given density. The situation is controlled by imposing the constraint that cells linked to each other only allow to differ by a single subdivision level at most.

The distance required to change from a low density to a high density depends heavily on the size of the cell associated with the lower density. For a simple square subdivision grid the distance between the side of a cell with subdivision level n and height h to a cell with subdivision level m is minimally $h \cdot \left(\frac{1}{2^1} + \frac{1}{2^2} + \dots + \frac{1}{2^{m-n}} \right)$, which converges to h for $m \rightarrow \infty$. For our prism based subdivision approach the distance converges to $2h$ vertically and $0.97h$ horizontally. However, depending on the positioning of the most dense cell w.r.t. the grid of the least dense cell, the required distance can increase to $4h$ vertically and

$2.75h$ horizontally. The horizontal distance is measured along the space filling surface in terms of the average length of segments in the 2D L-system. This means that two cells which are spatially next to each other can have a large difference in density so long as the space filling surface takes a large detour between the two.

III.5 ADAPTIVE SUBDIVISION

This section presents our approach for generating a subdivision structure with subdivision levels which closely matches the requested density distribution. Our approach consists of two steps: the first decides a lower bound subdivision level at each location and the second fine-tunes the local density distribution by dithering between the lower bound and a higher bound. Before presenting this approach, we introduce the methods for evaluating the target and the current density in a cell.

III.5.1. TARGET CELL DENSITY

Common ways of specifying a density distribution are as a scalar field defined on a tetrahedral mesh or a voxel model or as a procedural function. In order to make our program compatible with several commercial software packages, we construct a voxel model from a sequence of grayscale image files. The relationship between the grayscale value and the physical density is specified by the user. The target density ρ_T of a cell \mathcal{P} is computed as the average density of voxels $\{v_i\}$ covered by \mathcal{P} :

$$\rho_T(\mathcal{P}) = \frac{\sum_{v_i} \text{Vol}(v_i \cap \mathcal{P}) \rho(v_i)}{\sum_{v_i} \text{Vol}(v_i \cap \mathcal{P})} \quad (\text{III.2})$$

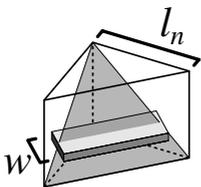
where the function $\text{Vol}(v_i \cap \mathcal{P})$ computes the volume of the common region of a voxel v_i and a cell \mathcal{P} . We define the *target mass* M_T as the requested amount of volume to be filled with solid material in a cell:

$$M_T(\mathcal{P}) = \rho_T(\mathcal{P}) \text{Vol}(\mathcal{P}). \quad (\text{III.3})$$

Because the size of starting cube of CrossFill is a power of 2 times the extrusion width w , the fractal can start with a volume which is considerably larger than the input model and its density distribution. For a cell lying completely outside the voxel set, we use the density of its nearest voxel as its density $\rho_T(\cdot)$.

III.5.2. CURRENT CELL DENSITY

For a cell \mathcal{P}_n located at level n in the subdivision tree, we calculate the current amount of material in the surface patch according to the size of \mathcal{P} and the type of route (L, R or A). By considering the configurations of embedded triangles in a prism cell (Fig. III.6), we calculate the current density estimate ρ_C of \mathcal{P}_n and the corresponding current mass M_C as follows:



$$\rho_C(\mathcal{P}_n) = \frac{w}{l_n} \cdot \begin{cases} \sqrt{2} & \text{for A} \\ 1 & \text{for L and R} \end{cases} \quad (\text{III.4})$$

$$l_n = l_{\text{init}} 2^{-n/2}$$

$$M_C(\mathcal{P}) = \rho_C(\mathcal{P}_n) \text{Vol}(\mathcal{P}_n)$$

Algorithm 2 Lower bound subdivision structure generation

```

function LOWERBOUNDSUBD(cell  $\mathcal{P}$ )
  if  $\mathcal{P}$  is a leaf-node then
    Compute  $M_C$ ,  $M_N$  and  $M_T$ ;
    if  $M_N < M_T$  then                                      $\triangleright \mathcal{P}$  needs to be subdivided
      SUBDIVIDE( $\mathcal{P}$ );
    for all  $c \in \mathcal{P}$ .children do
      LOWERBOUNDSUBD( $c$ );
  function SUBDIVIDE(cell  $\mathcal{P}$ )
    for all  $c \in \mathcal{P}$ .links do                                $\triangleright$  For linked neighbors of  $\mathcal{P}$ 
      if  $c$ .depth <  $\mathcal{P}$ .depth then                          $\triangleright$  For level constraint
        SUBDIVIDE( $c$ );
    Subdivide  $\mathcal{P}$  according to the rules;
    Update the corresponding links of neighbors;

```

where w is the constant width of material extrusion, l_n is the length of the cathetus at the top of the prism and l_{init} is the side length of the starting cube. Note that w is the horizontal width which differs from the thickness in the direction normal to the surface. Note also that the density, i.e., the fraction of solid material, is independent of the height of the prism, the embedding of the surface patch and the direction; it is determined by the average of horizontal segment length and the extrusion width w , which are unaffected by these factors. In this analysis, for the sake of simplicity, we neglect the effect that linked cells influence the density of a given cell due to the continuity enforcement. We provide a method for compensating for the errors induced by continuity enforcement in Section III.7.1.

From Eq. (III.4) and the subdivision rules in Fig. III.7 we can derive the increment in mass when performing a subdivision, which is used to supervise the generation of an adaptive subdivision structure. The new mass the children cells would have after a subdivision M_N is as follows:

$$M_N(\mathcal{P}) = M_C(\mathcal{P}) \cdot \begin{cases} 1 & \text{for A} \\ 1 + \frac{1}{2}\sqrt{2} & \text{for L and R} \end{cases} \quad (\text{III.5})$$

The ratio of A-, L- and R-route cells in a CrossFill structure with uniform subdivision level quickly converges to $1/3$ after a depth of only 5 subdivisions. If the current density of A-route cells is 100 % then the other $2/3$ of all cells is still at a lower density, so the maximum attainable density in the whole structure is $1/3 + 1/3\sqrt{2} \approx 80\%$.

From the above we can derive that the average density increment factor is $\sqrt{2}$. This provides more granularity than other fractal structures, e.g. the Hilbert curve, which has a density increment factor of 2 for each subdivision. One step of the Hilbert curve quadruples the amount of cells, while it only doubles the distance between connected cells, so the total length of the curve is doubled, which means that the curve with a constant line width comes to cover double the area.

III.5.3. LOWER BOUND SUBDIVISION LEVELS

Given the specified density distribution, we adaptively refine the subdivision structure such that the current density of each cell approaches, but does not exceed, the average

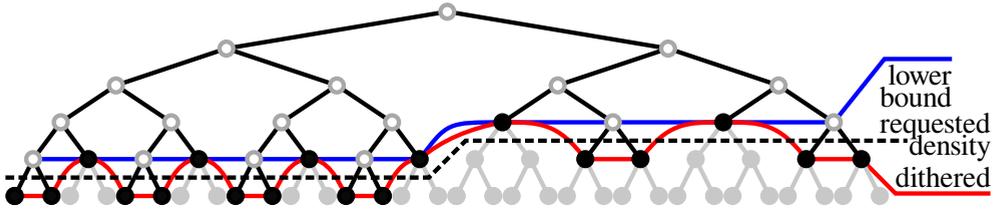


Fig. III.13. Schematic overview of dithering from a subdivision tree with lower bound subdivision levels.

target density (i.e., Eq. (III.2)). This is achieved by a top-down pass on the subdivision tree. In order to accomplish that we subdivide a cell \mathcal{P} if it satisfies the following condition: $M_N(\mathcal{P}) < M_T(\mathcal{P})$. To restrict the subdivision level difference between linked cells to be at most one, before subdividing a candidate cell \mathcal{P} that satisfies the condition, we subdivide its linked cells with a shallower subdivision level first.

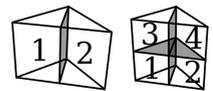
The pseudo-code of our algorithm is presented in Algorithm 2. By calling the function `LOWERBOUNDSSUBD(·)` on the root of the tree, the subdivision level is decided in each location. This constitutes the lower bound subdivision levels. The following subsection presents a dithering approach to further reduce the approximation error.

III.5.4. DITHERING

Because the input is a continuous density distribution while the output only admits a limited set of subdivision levels, choosing a subdivision level of a CrossFill cell always induces a discretization error a.k.a. *quantization error*. The idea is to *diffuse* this quantization error to linked cells in the neighborhood so as to influence the chosen subdivision level there. This causes the subdivision levels to oscillate between the subdivision levels closest to the target density (see Fig. III.13). This is akin to the widely employed *dithering* technique in multimedia processing.

We define the quantization error as the difference between target mass and current mass. Diffusing this error to dither the subdivision level leads to a CrossFill structure with densities better matching the target distribution regionally. See Fig. III.14 for an example on a simple square subdivision tiling. The lower bound subdivision structure exhibits strong banding artifacts; these artifacts are eliminated by dithering.

Dithering order When processing cells for dithering we consider the leaf nodes in a sequence analogous to the *Morton order* [33]: see figure in the right margin. We traverse the tree in depth-first order and at each non-leaf cell we recurse the children in the following order:



first the bottom left, then the bottom right and, if it is a Q-prism, then the top left and the top right. Here ‘right’ refers to linked cells along the direction of the horizontal traversal of the surface patches, which is clockwise around the space-filling polygon of each layer. Quantization error of the current cell \mathcal{P} is then redistributed to those cells in the neighborhood of \mathcal{P} which have not yet been processed in dithering.

Neighborhood Dithering along a manifold with a lower dimension makes the implementation simpler while retaining the advantages of dithering. Velho and Gomes [34] have shown that propagating quantization error along the directions of a space-filling curve (as

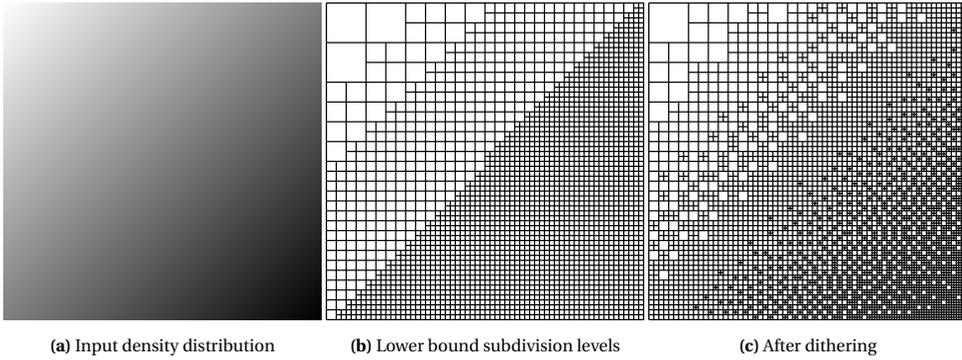


Fig. III.14. A square subdivision tiling fitted to an input density distribution image with a diagonal gradient from 0 % to 100 % density. The dithering step eliminates banding artifacts from the lower bound subdivision structure and creates a smooth density distribution.

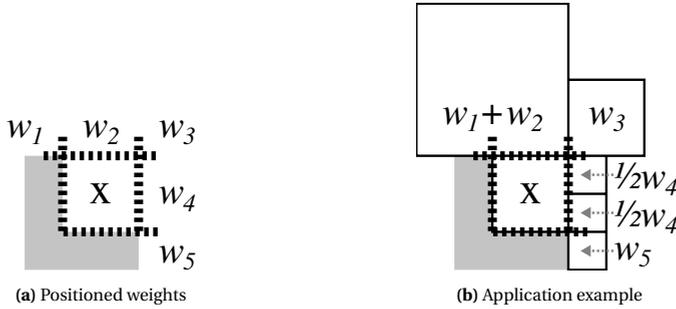


Fig. III.15. Quantization error propagation weights in various scenarios. Connected cells may cover an area at multiple positions relative to the checking cell marked by 'x'. The horizontal dimension depicted here is along the traversal of the space-filling surface (i.e., two-manifold). The gray area indicates locations which have already been processed.

1D manifold) for 2D images can yield a halftoning technique with appealing properties. Similarly, rather than considering all geometrically neighboring cells of a prism cell \mathcal{P} , we only consider the cells which are neighboring \mathcal{P} in the connectivity graph. As such we only disperse quantization error along the 2D manifold of the space-filling surface within 3D space. In order to visualize the dithering process effectively, we ‘unfold’ the space-filling surface and consider the resulting 2D topology (see Fig. III.4a and the supplementary video).

Weights The amount of the quantization error distributed to each (yet unprocessed) cell in the neighborhood depends on its relative position with respect to the current cell \mathcal{P} being checked. Connected cells to the left and bottom have always already been processed because of the Morton order. For a subdivision structure with uniform subdivision, the configuration is therefore as shown in Fig. III.15a. The diagonal cells are obtained by accessing the links of directly linked neighbors. There are various error diffusion schemes with different weights and different configurations (e.g. [35, 36, 37]). A comprehensive study about their performance is beyond the scope of this paper. We employ a simple error diffusion scheme with the following weights: $w_1 = w_3 = 1.0$, $w_2 = w_4 = 2.0$ and

Algorithm 3 Dithering

```

function DITHER(cell  $\mathcal{P}$ )
  if  $\mathcal{P}$  is NOT a leaf-node then
    for all  $c \in \mathcal{P}$ .children do                                 $\triangleright$  using the Morton order
      DITHER( $c$ );
  else
    Compute  $M_C$ ,  $M_N$  and  $M_T$ ;
    Define  $M = M_C$ ;
    if  $\frac{1}{2}(M_C + M_N) + \mathcal{P}.M_{\mathcal{E}} < M_T$  then
      Subdivide  $\mathcal{P}$  according to the rules;
      Update the corresponding links to neighbors;
      Update mass as  $M = M_N$ ;
     $\mathcal{E} = M + \mathcal{P}.M_{\mathcal{E}} - M_T$ ;
    Re-distribute  $\mathcal{E}$  to  $\mathcal{P}$ 's unprocessed neighbors;
     $\triangleright$  The  $M_{\mathcal{E}}$  of cells in the neighborhood is updated.

```

$w_5 = 0.0$. Because not all positions relative to the checking cell are always occupied by (unprocessed) cells, the weights are normalized to compute the quantization error to be diffused to unprocessed linked neighbors: $\hat{w}_i = w_i / \sum_{c \in \mathcal{P}.\text{neighborhood}} c.w$

Because linked cells can have a size different from \mathcal{P} , the weights need to be adjusted to account for the change of configuration. A linked cell can occupy more space than \mathcal{P} or only a portion of the space occupied by \mathcal{P} . In the former case multiple weights are added together, while in the latter case the weight is split equally. The diagonal positions always retain the same weight. An illustration of such a configuration can be found in Fig. III.15b.

Algorithm The dithering algorithm decides on the final subdivision level by choosing between the lower bound subdivision level and a higher bound of one subdivision level deeper by comparing the quantization error of those two subdivision levels. The quantization error \mathcal{E} of a prism cell \mathcal{P} before subdivision (\mathcal{E}_C) and after subdivision (\mathcal{E}_N) are calculated by

$$\begin{aligned}
 \mathcal{E}_C(\mathcal{P}) &= M_C(\mathcal{P}) + M_{\mathcal{E}}(\mathcal{P}) - M_T(\mathcal{P}) \\
 \mathcal{E}_N(\mathcal{P}) &= M_N(\mathcal{P}) + M_{\mathcal{E}}(\mathcal{P}) - M_T(\mathcal{P})
 \end{aligned}
 \tag{III.6}$$

where $M_{\mathcal{E}}(\mathcal{P})$ is the quantization error diffused to \mathcal{P} from already processed cells. \mathcal{P} is subdivided if the absolute value of $\mathcal{E}_N(\mathcal{P})$ is smaller than the absolute value of $\mathcal{E}_C(\mathcal{P})$. An exception to this rule is if \mathcal{P} has linked neighbors with a lower subdivision level in order to comply with the constraint that linked cell can only differ by a single subdivision level. After the subdivision decision the corresponding quantization error is diffused to the unprocessed linked neighbors according to the weighing scheme described above. Pseudo-code of this dithering approach is given in Algorithm 3. The diffused error $\mathcal{P}.M_{\mathcal{E}}$ is initialized as *zero* on each cell \mathcal{P} before calling DITHER(\cdot) on the root of the subdivision tree. Figure III.14c shows the result on a 2D example.

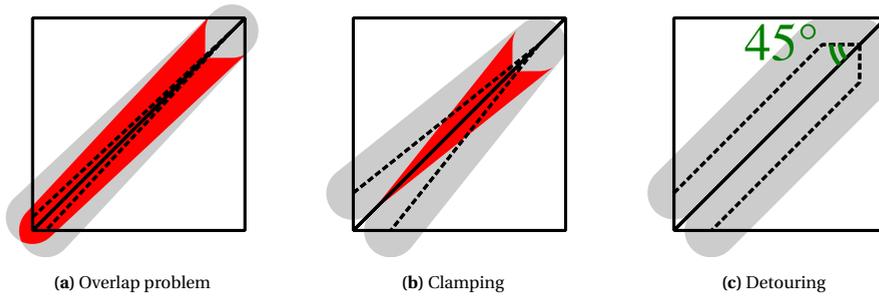


Fig. III.16. Dealing with overlapping paths in neighboring prisms (the red color indicates overlapping regions) – the figures illustrate a top-view of two prisms crossed by extrusion paths (dashed black lines) which initially lie close to each other (i.e., overlap occurs in (a)). By clamping the endpoints to a position that is at least $w/2$ the extrusion width away from the other prism (b) and introducing an additional 45° turn on the line segments (c), the overlap can be avoided.

III.6 TOOLPATH GENERATION

We now have a method for creating an infill structure with spatially graded density according to a user-specified density distribution. The infill structure is defined in a cubic region and can be fabricated by continuous material extrusion. In this section, we will first explain how to effectively slice the structure into a continuous 2D polygonal curve for each layer. After that, we will fit the 2D polygonal curve of a layer into the region of an input 3D model.

III.6.1. SLICING

The first step toward generating the toolpath of a layer for 3D printing is to generate the space-filling curve which lies on the intersection between the surface of CrossFill and the horizontal plane at the height z of the printing layer. As mentioned above, the space-filling surface only exists conceptually in our implementation. We directly generate the space-filling curves from the type of prism cells and their linkage in the connectivity graph. Given a height z , we first find the sequence of cells covering this height in the subdivision structure. The cell which is closest to the last point on the toolpath of the previous layer, is chosen as the first cell for exploring the horizontally linked cells. The whole sequence of cells can be traced out by following the links in the connectivity graph which are pointing to the right, i.e., by following the cells along the horizontal traversal direction (see Fig. III.6). When appending a cell to the sequence we take the right linked cell which intersects with the z height. After employing the continuity enforcement rules from Section III.4.4 we are left with the edges of the surface patches, which are sliced at z to serve as the vertices of the space-filling curve of that layer.

III.6.2. OVERLAP PREVENTION

These space-filling curves can have overlap near the boundary of a prism when a surface patch segment is too close to the neighboring prism (see Fig. III.16a for an example). This occurs when the distance between the two segments is less than the horizontal width w of extrusion toolpaths. In order to prevent path overlap, we *clamp* the endpoints of sliced line segments to a position that is $w/2$ away from the neighboring prism (see Fig. III.16b).

However, only applying the clamping step cannot completely avoid overlap in situa-

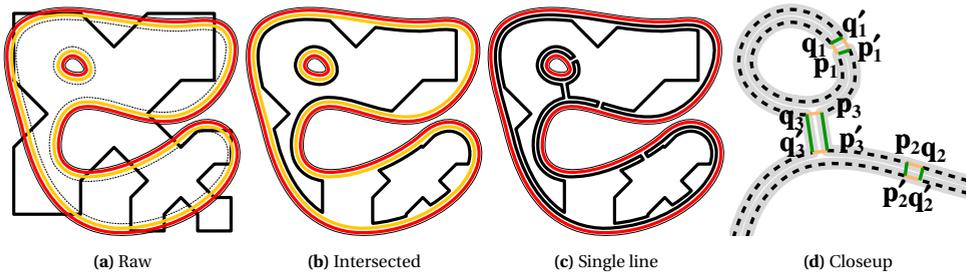


Fig. III.17. Inserting the cross-filling curve into to an infill area, where the cross-filling curve is displayed in black color, the outer wall is in red color and the second wall is shown in orange color. (a) The original cross-filling curve. (b) The intersected space-filling structure. (c) The finally connected toolpaths of multiple walls and the infill structure. (d) A zoom-view of (c).

tions where the toolpath makes a sharp turn. When the *turning angle between the segment and the side of the cell* is sharper than 45° , we introduce an additional vertex at a position with distance $w \cdot 1/2\sqrt{2}$ away from the vertex on the side of the cell – called *detouring*. See Fig. III.16c. It should be noted that in a uniformly subdivided CrossFill structure, the turning angle on a space-filling curve generated by slicing is either 45° or 90° (see Fig. III.9b for an example), which means that detouring is not needed in such a context.

Detouring does not violate the overhang constraint. The introduced vertex is supported either by a detoured vertex below or a segment with a turning angle just above 45° . Moreover, detouring does not reduce any area covered by the extrusion path, so a detoured layer still supports the layer above. The increased density caused by detouring can be compensated for using a method introduced in Section III.7.1.

III.6.3. CONVERSION INTO INFILL STRUCTURE

Now that we have a non-overlapping space-filling curve as toolpath for fabricating the CrossFill structure of each layer, We limit it to the interior area of an input 3D model while retaining the continuity of the toolpaths. Specifically, we first intersect the space-filling curve with the infill area shrunk by $w/2$ (see Fig. III.17c). This operation makes the trimmed space-filling curve connected to itself via the perimeters of that layer.

Performing an intersection between the space-filling curve and the infill area could result in a single polygon. However, there are cases in which additional polygons are generated:

1. the infill area splits the space-filling curve into multiple parts (in the bottom right of Fig. III.17b);
2. the infill area contains a polygon which does not touch the space-filling curve (in the top left of Fig. III.17b).

We tackle both these problems by connecting all polygons to the innermost perimeter of the shell of the print and to each other afterwards. The problem to be solved here for generating continuous toolpaths is different from the method of Zhao et al. [38], in which spiraling toolpaths are generated to completely fill a given region. By contrast, we are tackling the problem of connecting multiple polygons into a single polygonal curve.

Two polygons are connected into a single polygon by building a bridge between them as follows. First, two points \mathbf{p} and \mathbf{q} with $\|\mathbf{p} - \mathbf{q}\| = w$ on a polygon are considered. We

consider the point \mathbf{p}' on the other polygons closest to \mathbf{p} and find a point \mathbf{q}' on the other polygon such that \mathbf{qq}' is parallel to \mathbf{pp}' . We search for such pairs of points until we find a bridge for which the length is at most $3/2w$. New line segments \mathbf{pp}' and \mathbf{qq}' are then added and the line segments \mathbf{pq} and $\mathbf{p'q'}$ are removed. Examples for building bridges can be found in Fig. III.17d. Repeatedly building bridges between polygons can connect all into a single polygon. In order to minimize the number of sub-optimal bridges, we start by connecting the smallest polygons to suitable neighbors and work our way outward.

Our method has several advantages. When there are many possible candidate locations for building a bridge, we can select the 'optimal' position according to various criteria. For example, we can promote bridges at regions with low curvature in order to minimize the influence on the extruded amount of material. Another option is to build bridges that are closer to interior regions so as to minimize the visual surface impact.

Connecting polygons which have distance more than $2w$ introduces new line segments hanging in the air. Strictly this conflicts with the overhang constraint; however, this often is not a problem for FDM printing. Such distances are rarely long, and these lines of bridges do not have to support any material above. In practice, the extruded beads of FDM can stay in mid-air because of the high viscosity of the melted plastic.

III.7 RESULTS

III.7.1. EXPERIMENTS

Experiments were performed on an Intel Core i7-7500U CPU @ 2.70 GHz using a single core and 16.3 GB memory. We have printed test structures on several Ultimaker 3 machines, loaded with white Ultimaker TPU 95A in AA 0.4 mm Print Cores. The basic print path settings were taken from the default Cura 4.0 profile of 0.1 mm layer thickness for this setup. Most notably the setting for Infill Line Width was $w = 0.38$ mm and the Speed was 25 mm/s. In order to enable the connect polygons functionality in some of our tests, we set the Extra Skin Wall Count to zero and instead set the Extra Infill Wall Count to one, so that we can enable the setting Connect Infill Polygons. The Infill Line Distance setting was set to 0.76 mm meaning that the smallest possible prism has sides of length $l_{\max} = 2w\sqrt{2}$, which corresponds to a density of 40 % in order to save computation time.

For various tests we have used a test cube with side lengths of 48.64 mm, which is $2^7 w$, so that the starting cube of the subdivision tree matches the 3D model. In order to isolate the infill structure, the settings Top/Bottom Thickness and the Wall Thickness have been set to 0 mm.

For some of the models, the requested infill densities are too low to support the dense top skin. In order to overcome this problem, we enforced a minimal subdivision level for prism cells which overlap with top skin after the dithering stage, by iteratively calling SUB-DIVIDE from Algorithm 2. By modifying the subdivision structure in this way, we guarantee a required percentage of infill for supporting top skin regions.

SIMPLIFIED DENSITY MEASURE COMPENSATION

The method we propose uses a simplified density estimate based on cell size and route type. Because this does not take into account effects from continuity enforcement (Section III.4.4) or from clamping and detouring (Section III.6.2) the output density is different from the simplified density estimates. Once we evaluate the discrepancy, we can compensate for it.

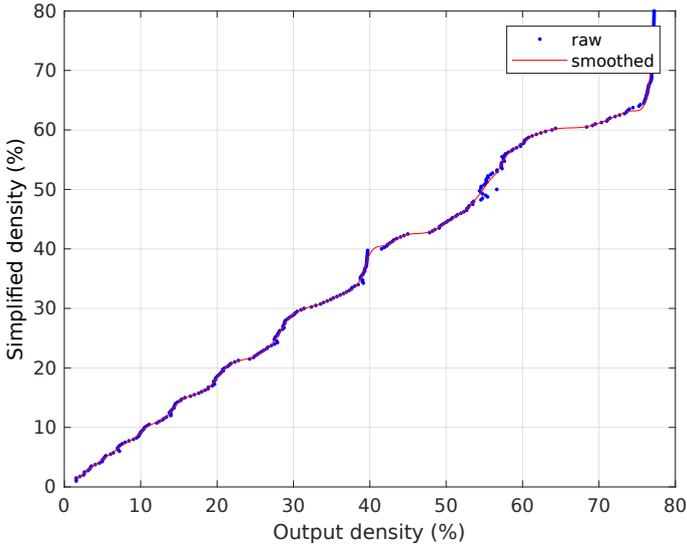


Fig. III.18. Compensating for inaccuracies of the simplified density measure. For several simplified density values the actual total amount of volume is recorded in blue, so that we can map required output densities to the corresponding simplified input densities using a smoothing spline fitted to that data in red.

We have generated test cubes with a homogeneous simplified density (Eq. (III.4)) within the range of 1% to 80% and analyzed the total extruded volume compared to the total volume of the cube. These results are shown in blue in Fig. III.18. We fit a MATLAB smoothing spline to the data with a smoothing parameter of 0.75. The resulting curve is then used to compensate for the disparity between the simplified density estimates and the actual densities by mapping the density requirements to the corresponding simplified density estimates prior to applying our algorithms.

ACCURACY

The accuracy of a functionally graded material is inherently related to a viewing resolution. When viewing any print at a resolution close to the printer resolution, the density is either 100% or 0% regardless of the user specified density at each location, which means the accuracy at that resolution is low. In order to evaluate the accuracy of our functionally graded material, we evaluate the average local error at a range of resolutions. For each resolution we divide the specification and the generated infill structure into smaller cubes and compute the local error as the absolute difference between the average specification density throughout that cube and the average realized density throughout that cube. The accuracy is then given by the average local error across all subcubes for that resolution (see Fig. III.19). Note that at kernel sizes of powers of two the subcubes align with the prism-shaped cells, which lowers the error measure.

We define several density specifications for our test cube on which we evaluate the accuracy. a) Homogeneous at 20%. b) Homogeneous at 40%. c) Gradient: a smooth linear gradient from 10% in one bottom corner to 40% in the diametrically opposite corner. d) Contrast plane: half of the cube is 10% infill density and the other half is 40%. The

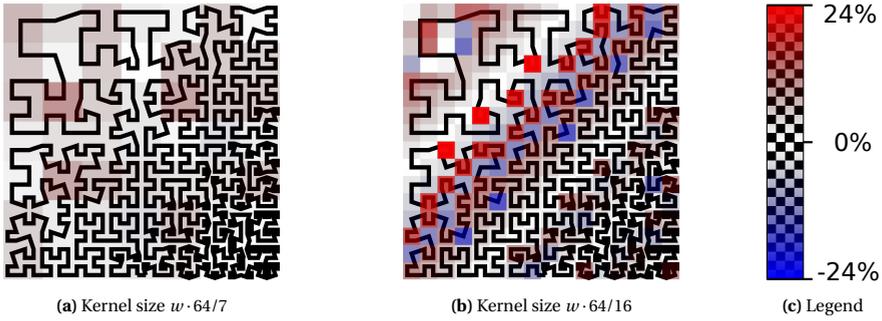


Fig. III.19. Example of local errors on a 2D space-filling curve generated from a diagonal gradient from 10% to 80% density. The errors are translated to opacity and overlaid with the space-filling curve. w is the line width. When analyzing at lower resolution the errors are higher. The overall error is positive because the diagonal line segments introduced at locations where consecutive cells are a different subdivision level have a higher density than the simplified density estimate used.

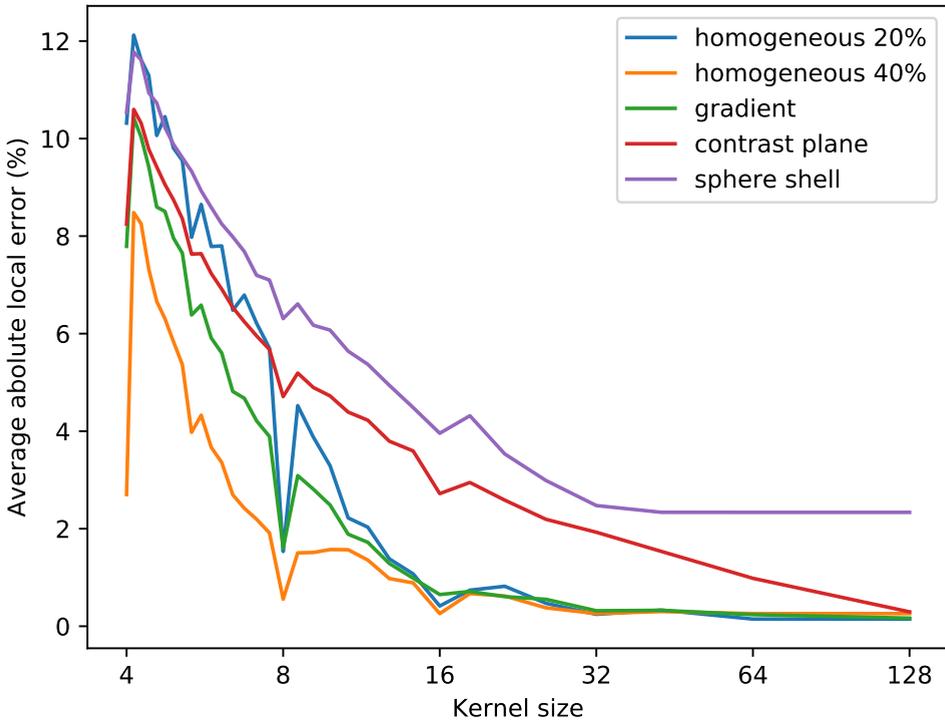


Fig. III.20. Local average error for a range of kernel sizes on several test specifications. Kernel size is in multiples of the line width, while the average local error is measured in terms of infill density percentage. When viewing the structures at lower resolutions, the accuracy is higher. High-frequency specifications such as the sphere shell perform worse.

Table III.1: Example model settings.

	white	black	top	spec size (px)	phys. size (mm)
Sole	5 %	40 %	0 %	1456 × 564 × 1	155 × 58 × 14
Bunny	10 %	80 %	20 %	45 × 35 × 2785	129 × 100 × 126
Phantom	0 %	100 %	20 %	417 × 412 × 146	123 × 121 × 87
Saddle	40 %	10 %	0 %	60 × 44 × 47	250 × 188 × 63

Table III.2: Computation time in seconds.

	Sole	Bunny	Phantom	Saddle
Lower bound	6.2	11.2	4.3	150.9
Dithering	0.5	2.3	0.4	10.5
Extract polygon	0.4	15.7	3.4	12.8
Limit polygon	0.3	11.3	3.8	12.0
Reconnecting	4.6	98.9	65.4	135
Total gcode	14	182	105	346

plane which separates these two halves makes an angle of 22.5° with the X axis in the horizontal plane and has an overhang angle of 45° . e) Sphere shell: a sphere with a radius of half the cube side length and a shell thickness of $1/7$ the size of the side lengths of the cube. The density of the shell is 40 % and the density inside and outside is 10 %. All these are specifications consisting of 512 images of 512×512 px. The accuracy results are shown in Fig. III.20.

COMPUTATION TIME

In order to evaluate the running time of our algorithms, we consider four of the application models discussed in Section III.7.3. We consider the test models and corresponding settings displayed in Table III.1. The computation times are shown in Table III.2.

ELASTIC BEHAVIOR

Because the generated structures are similar to foams, which are often used in a compressive context, it would be interesting to find out their compressive behavior. Because the non-linear material properties at high strain values when compressing a foam are difficult to capture when using a finite elements method, we have performed actual physical tests instead.

We have printed samples with homogeneous subdivision level, and we have printed samples using dithering to approximate several homogeneous density specifications with simplified densities between 10 % and 30 %. Compressions were performed in both the vertical and the horizontal direction. Because our structures are rotationally symmetric around the Z axis we only need to test 2 of the 3 dimensions. In total 78 prints were made in 42 test configurations (some configurations were tested multiple times): 2 testing directions, 4 homogeneous subdivision levels (10.1 %, 14.0 %, 20.1 %, 28.5 %) and 17 heterogeneous subdivision levels in the same range.

We have performed compression tests on the Instron 3366, fitted with compression plates. See Fig. III.21a. Compressions were performed at a speed of 0.5 mm/s up to a maximum force of 2 kN after which the sample was decompressed.

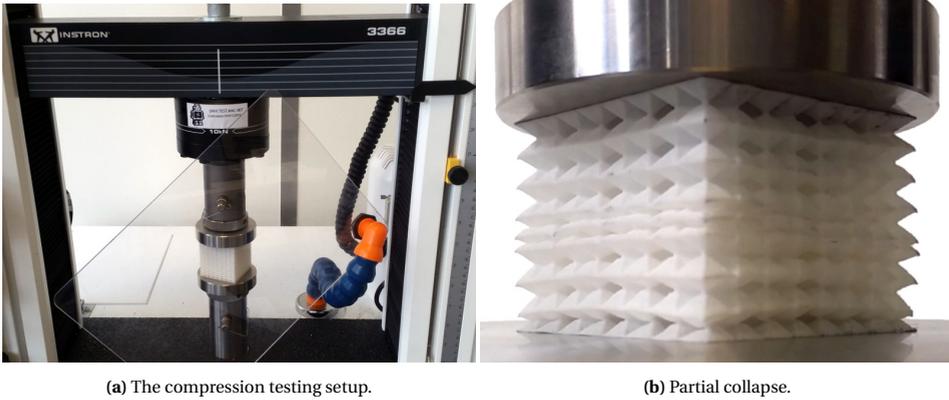


Fig. III.21. Compression testing. Under stress some cells in the structure start to collapse. Cells tend to collapse in groups on the same heights.

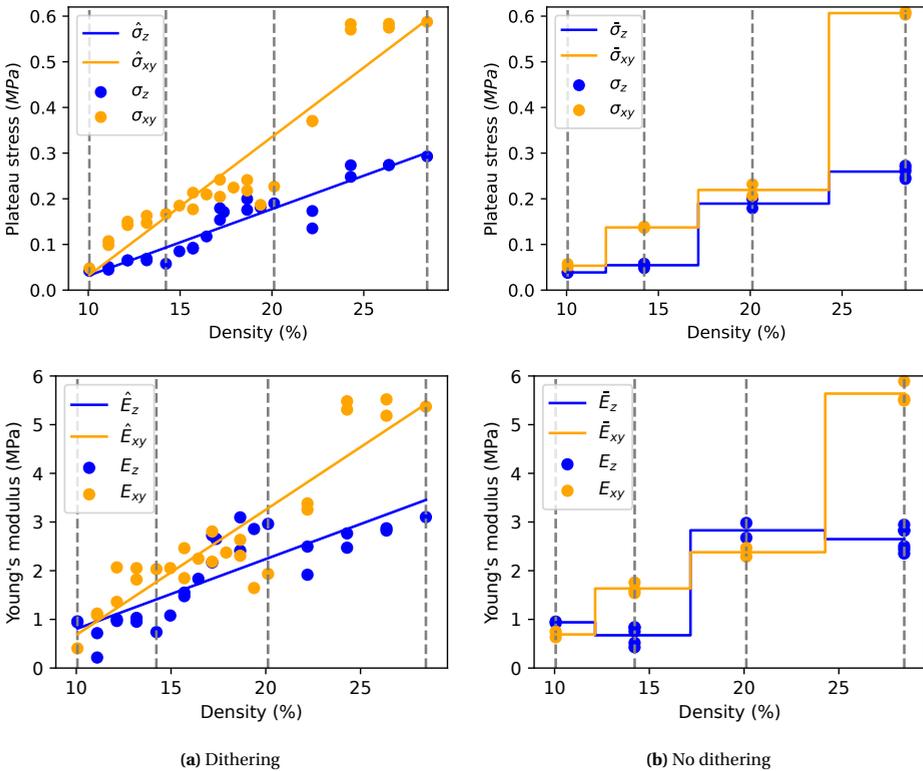


Fig. III.22. Plotting Young's modulus and plateau stress for different densities. Lines represent linear regression results over dithered data and average values of not dithered data. Dithering provides material properties in between the discrete densities. However, they are not monotonically increasing with the density, which makes mapping required material properties to corresponding infill densities non-trivial.

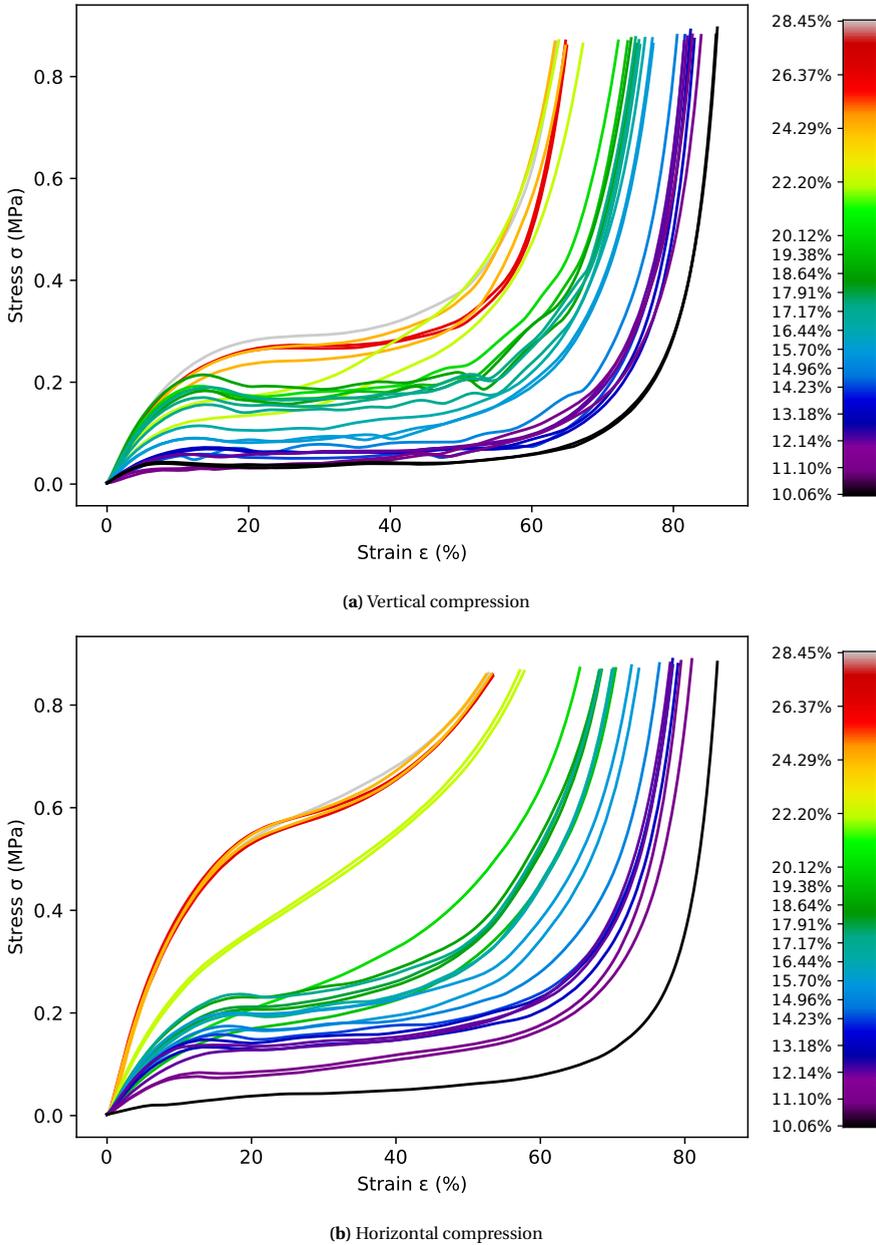


Fig. III.23. Stress – strain – density graphs of the results of our compression experiments. The structure is more compliant in the vertical direction than in the horizontal direction. From these results we can conclude that dithering can provide a spectrum of material properties; higher densities start collapsing at higher strain values and their collapse trajectory is shorter.

The stress-strain results are plotted in Fig. III.23. One interesting observation from the data is that the stress-strain graphs are roughly horizontal for a long range of strain values. This is caused by the structure collapsing and folding in on itself. Such plateaus are typical of foams [39, 40], and they are important design variables for common applications of foams [41]. Figure III.21b shows how such collapse can be localized to only a particular Z range in the case of vertical compression.

The constant stress along that range is a fundamental characteristic of such a plateau. We estimate it by taking the average stress along the part of the stress-strain graph which has a local tangent modulus below 0.4 MPa. The resulting values have been plotted alongside the Young's modulus values in Fig. III.22b.

III.7.2. DISCUSSION

ACCURACY

At a kernel size of $16w$ (6.08 mm) the average absolute local error is low for smooth input density distributions. Because of the constraint that neighboring linked cells can only differ by a single subdivision level the two distributions with sharp contrast edges or high frequency detail score considerably worse. See Fig. III.20, 'contrast plane' and 'sphere shell'. The relative error decreases as the resolution increases. Therefore, depending on the application the user might decide that at a specific resolution the accuracy is good enough.

COMPUTATION TIME

Our algorithms take up over 80 % of the total gcode generation time; there is room for improvement. It should be noted that we not optimized the code for loading the density specification data into the subdivision tree or the code for connecting polygons, which currently consume the majority of the processing times. Polygon extraction and polygon limiting times depend highly on the number of layers, which is relatively small for the shoe sole example.

ELASTIC BEHAVIOR

It should be noted that the first local maximum in the stress-strain graphs of Fig. III.23 are not the yield points. Most of the deformation applied to the structure is elastic deformation. We observed the material to creep back to its original shape at ever decreasing rates. After 24 hours a strain of merely 0.8 % remained after an initial strain above 50 %, meaning that the plastic deformation is negligible at stresses with magnitudes as high as in our tests (2 kN).

For larger densities the strain at which the structure is fully collapsed is lower. When we compress a sample with 40 % density by 60 % then the stress-strain graph should exhibit a local tangent equal to the base properties of the material used. In our test results we see that the tangent tends toward the Young's modulus of TPU: 26 MPa [42]. This agrees with literature on the densification of foams [39].

In most aspects the structures are more compliant in the vertical direction than in the horizontal. The plateau heights in Fig. III.23a are lower than the corresponding ones in (b) and the strain regions of the plateaus are wider as well. Of course some aspect of this anisotropy is caused by the layer-wise buildup of FDM, but a larger part of the difference is most likely caused by the geometric structure. The structure can collapse in the Z direction easily because of the alternating E- and C-embeddings.

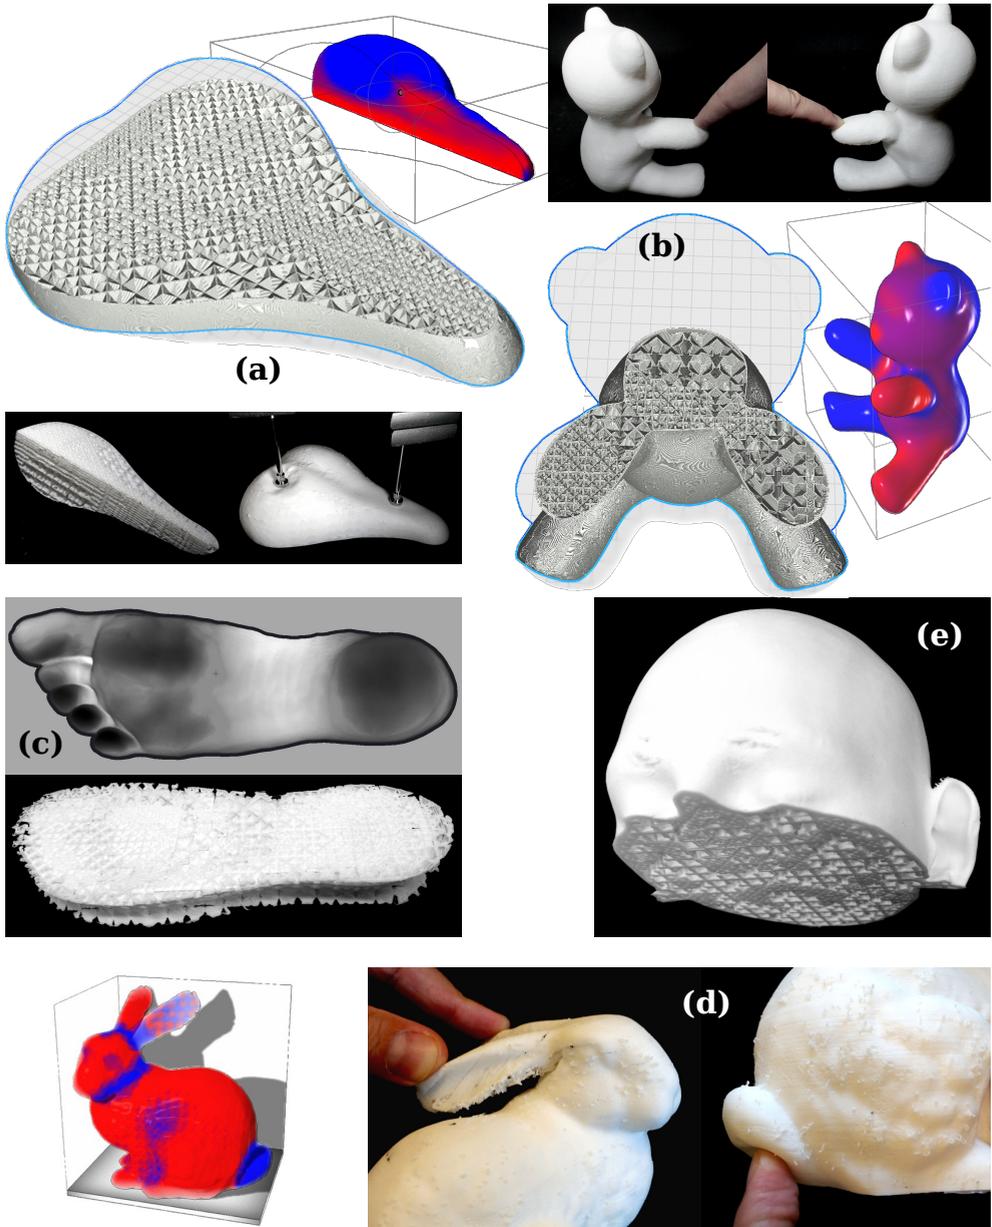


Fig. III.24. Various examples of applications of CrossFill. (a) A bicycle saddle with a density specification. A weight of 33 N is added on various locations to show the different response of different density infill. (b) A teddy bear with a density specification. (c) A shoe sole with densities based on a pressure map of a foot. (d) The Stanford bunny painted with a density specification. (e) A medical phantom with an example density distribution for calibrating an MRI scanning procedure.

The vertical Young's modulus E_z is lower than the horizontal Young's modulus E_{xy} for most densities. For densities around 10 % and 20 %, that relationship is reversed, though. That can be explained by the fact that around those densities the subdivision structure contains mostly Q-prism cells which are filled with CrossFill surface patches with a more vertical slope than those for H-prisms. More vertical elements increase the stiffness in the vertical direction.

The prisms only subdivide vertically every two iterations: only the Q-prisms do. The H-prisms of 20.1 % are subdivided horizontally compared to 14.0 %, while the prisms of 14.0 % and 28.5 % are also subdivided vertically compared to 10.1 % and 20.1 % respectively. This irregularity can explain the nonmonotonicity we see in the Young's modulus with respect to infill densities in Fig. III.22b: horizontal subdivision decreases the slope of the surface patches in the structure, which decreases the thickness of these patches, which in turn decreases the overall stiffness.

Because the deeper subdivision level cells determine the ruled surface in horizontal continuity enforcement, surface patches of a dithered subdivision structure are predominantly determined by the higher density cells. We therefore expect to see large changes after 10 % and after 20 %. We can see large jumps in the vertical Young's modulus at those places in Fig. III.22b. Also the plateau heights for horizontal compressions show large jumps.

When comparing the dithered results in Fig. III.22 to the ones without dithering we can conclude that dithering indeed provides more granular control of the overall material properties of the manufactured part. However, because of the nonmonotonicity in these results it is not trivial to define a process for the designer to choose which infill density is needed at a location in a design.

III.7.3. APPLICATIONS

The structures generated by our method support variable compliance, which can control the deformation of an object under certain loads. We have designed a density specification for the Stanford bunny model using Autodesk Monolith [43], as shown in Fig. III.24d.

Because the CrossFill structures behave much like foams, it could be attractive for "cushioning, packaging and energy absorption" [39]. Foams with variable stiffness could be used for a personalized bike saddle (see Fig. III.24a) or for personalized shoe soles (see Fig. III.24c).

A boundary mesh of a teddy bear along with a density distribution show how different densities lead to different bending behavior (see Fig. III.24b). The difference in density distribution in the two arms causes the arms to deform in a different way under the same load.

Because our method is principally density-based, it could also be useful in situations where the relevant material properties are volumetric; CrossFill could prove useful for imaging *phantoms* - objects which can be used in the medical field to evaluate magnetic resonance imaging (MRI) scan procedures. We have printed an example of what could serve as a phantom in Fig. III.24e.

III.8 CONCLUSION AND FUTURE WORK

In this paper, we have introduced a new infill structure, CrossFill, which can provide spatially graded density to match a user-specified density distribution. CrossFill is carefully

designed so that it is self-supporting and can be fabricated from a single, continuous and overlap-free toolpath on each layer. Algorithms for generating the lower bound subdivision levels and dithering the subdivision levels have been developed to accurately match the prescribed density distribution. To use CrossFill as infill structures of a given 3D model, we have presented an algorithm to connect the toolpaths of CrossFill and the toolpaths for the given model's shell into a single continuous extrusion path. The performance of CrossFill has been verified on a variety of experimental tests and applications.

The study of experimental tests shows that CrossFill acts very much like a foam although future work needs to be conducted to further explore the mapping between density and other material properties. Another line of research is to further enhance the dithering technique, e.g. changing the weighing scheme of error diffusion.

REFERENCES

- [1] Tim Kuipers, Jun Wu, and Charlie C.L. Wang. "CrossFill: Foam Structures with Graded Density for Continuous Material Extrusion". In: *Computer-Aided Design* 114 (Sept. 2019), pp. 37–50. ISSN: 00104485. DOI: 10.1016/j.cad.2019.05.003.
- [2] Marco Attene et al. "Design, Representations, and Processing for Additive Manufacturing". In: *Synthesis Lectures on Visual Computing: Computer Graphics, Animation, Computational Photography, and Imaging* 10.2 (June 2018), pp. 1–146. ISSN: 2469-4215. DOI: 10.2200/S00847ED1V01Y201804VCP031.
- [3] Marco Livesu et al. "From 3D models to 3D prints: an overview of the processing pipeline". In: *Computer Graphics Forum* 36.2 (May 2017), pp. 537–564. ISSN: 14678659. DOI: 10.1111/cgf.13147. arXiv: 1705.03811.
- [4] Jonàs Martínez, Jérémie Dumas, and Sylvain Lefebvre. "Procedural voronoi foams for additive manufacturing". In: *ACM Transactions on Graphics* 35.4 (2016), pp. 1–12. ISSN: 07300301. DOI: 10.1145/2897824.2925922.
- [5] A. O. Aremu et al. "A voxel-based method of constructing and skinning conformal and functionally graded lattice structures suitable for additive manufacturing". In: *Additive Manufacturing* 13 (2017), pp. 1–13. ISSN: 22148604. DOI: 10.1016/j.addma.2016.10.006.
- [6] Simon R.G. Bates, Ian R. Farrow, and Richard S. Trask. "Compressive behaviour of 3D printed thermoplastic polyurethane honeycombs with graded densities". In: *Materials & Design* 162 (2018), pp. 130–142. ISSN: 02641275. DOI: 10.1016/j.matdes.2018.11.019.
- [7] Sing Ying Choy et al. "Compressive properties of functionally graded lattice structures manufactured by selective laser melting". In: *Materials and Design* 131 (2017), pp. 112–120. ISSN: 0264-1275. DOI: 10.1016/j.matdes.2017.06.006.
- [8] Sakkadech Limmahakhun et al. "Stiffness and strength tailoring of cobalt chromium graded cellular structures for stress-shielding reduction". In: *Materials and Design* 114 (2017), pp. 633–641. ISSN: 18734197. DOI: 10.1016/j.matdes.2016.11.090.
- [9] Dawei Li et al. "Optimal design and modeling of gyroid-based functionally graded cellular structures for additive manufacturing". In: *Computer-Aided Design* 104 (July 2018), pp. 87–99. ISSN: 00104485. DOI: 10.1016/j.cad.2018.06.003.

- [10] Diab W. Abueidda et al. “Mechanical properties of 3D printed polymeric cellular materials with triply periodic minimal surface architectures”. In: *Materials & Design* 122 (2017), pp. 255–267. ISSN: 0264-1275. DOI: 10.1016/j.matdes.2017.03.018.
- [11] Fady Massarwi et al. “Hierarchical, random and bifurcation tiling with heterogeneity in micro-structures construction via functional composition”. In: *Computer-Aided Design* 102 (2018), pp. 148–159. ISSN: 00104485. DOI: 10.1016/j.cad.2018.04.017.
- [12] Christian Schumacher et al. “Microstructures to control elasticity in 3D printing”. In: *ACM Transactions on Graphics* 34.4 (July 2015), 136:1–136:13. ISSN: 07300301. DOI: 10.1145/2766926.
- [13] Julian Panetta et al. “Elastic textures for additive fabrication”. In: *ACM Transactions on Graphics* 34.4 (July 2015), 135:1–135:12. ISSN: 07300301. DOI: 10.1145/2766937.
- [14] Corentin Coulais et al. “Combinatorial design of textured mechanical metamaterials”. In: *Nature* 535.7613 (2016), p. 529.
- [15] Xingchen Liu and Vadim Shapiro. “Random heterogeneous materials via texture synthesis”. In: *Computational Materials Science* 99 (2015), pp. 177–189. ISSN: 0927-0256. DOI: <https://doi.org/10.1016/j.commatsci.2014.12.017>.
- [16] Jun Wu et al. “Infill Optimization for Additive Manufacturing – Approaching Bone-like Porous Structures”. In: *IEEE Transactions on Visualization and Computer Graphics* 24.2 (Feb. 2018), pp. 1127–1140. ISSN: 1077-2626. DOI: 10.1109/TVCG.2017.2655523.
- [17] Erik Andreassen, Boyan S. Lazarov, and Ole Sigmund. “Design of manufacturable 3D extremal elastic microstructure”. In: *Mechanics of Materials* 69.1 (2014), pp. 1–10. ISSN: 0167-6636. DOI: 10.1016/j.mechmat.2013.09.018.
- [18] Eric Garner et al. “Compatibility in microstructural optimization for additive manufacturing”. In: *Additive Manufacturing* 26 (2019), pp. 65–75. ISSN: 2214-8604. DOI: 10.1016/j.addma.2018.12.007.
- [19] Ole Sigmund. “Tailoring materials with prescribed elastic properties”. In: *Mechanics of Materials* 20.4 (1995), pp. 351–368. ISSN: 01676636. DOI: 10.1016/0167-6636(94)00069-7.
- [20] Jonàs Martínez et al. “Polyhedral Voronoi diagrams for additive manufacturing”. In: *ACM Transactions on Graphics*. Proceedings of SIGGRAPH 2018 37.4 (Aug. 2018), p. 15. DOI: 10.1145/3197517.3201343.
- [21] Jun Wu et al. “Self-supporting rhombic infill structures for additive manufacturing”. In: *Computer-aided Design* 80 (Nov. 2016), pp. 32–42. ISSN: 00104485. DOI: 10.1016/j.cad.2016.07.006.
- [22] Jun Wu. “Continuous optimization of adaptive quadtree structures”. In: *Computer-Aided Design* 102 (2018), pp. 72–82. ISSN: 00104485. DOI: 10.1016/j.cad.2018.04.008. arXiv: 1712.04523.
- [23] Jonàs Martínez et al. “Orthotropic k-nearest foams for additive manufacturing”. In: *ACM Transactions on Graphics* 36.4 (July 2017), pp. 1–12. ISSN: 07300301. DOI: 10.1145/3072959.3073638.

- [24] David Hilbert. “Ueber die stetige Abbildung einer Line auf ein Flächenstück”. In: *Mathematische Annalen* 38.3 (1891), pp. 459–460.
- [25] Waław Franciszek Sierpiński. “Sur une nouvelle courbe continue qui remplit toute une aire plane”. In: *Bulletin de l'Académie des Sciences de Cracovie - Série A* (1912), pp. 463–478.
- [26] George Polya. “Über eine Peanosche kurve”. In: *Bull. Acad. Sci. Cracovie (Sci. math. et nat. Série A)* (1913), pp. 305–313.
- [27] Siddharth H. Nair, Arpita Sinha, and Leena Vachhani. “Hilbert’s Space-filling Curve for Regions with Holes”. In: *arXiv preprint arXiv:1709.02938* (2017). arXiv: 1709 . 02938. URL: <http://arxiv.org/abs/1709.02938>.
- [28] Sanderson Gonzaga et al. “Sierpinski-like Space-filling Curve for Total Ordering of Adaptive Triangular Discretized Domains”. In: February 2015 (2015).
- [29] M. Bader et al. “Memory efficient adaptive mesh generation and implementation of multigrid algorithms using Sierpinski curves”. In: *International Journal of Computational Science and Engineering* 4.1 (2008), p. 12. ISSN: 1742-7185, 1742-7193. DOI: 10.1504/IJCSE.2008.021108.
- [30] J. G. Griffiths. “Toolpath based on Hilbert’s curve”. In: *Computer-Aided Design* 26.11 (1994), pp. 839–844. ISSN: 00104485. DOI: 10.1016/0010-4485(94)90098-1.
- [31] Gurunathan Saravana Kumar, Ponnusamy Pandithevan, and Appa Rao Ambatti. “Fractal raster tool paths for layered manufacturing of porous objects”. In: *Virtual and Physical Prototyping* 4.2 (June 2009), pp. 91–104. ISSN: 17452759. DOI: 10.1080/17452750802688215.
- [32] Masood Ahmed and Shahid Bokhari. “Mapping with space filling surfaces”. In: *IEEE Transactions on Parallel and Distributed Systems* 18.9 (2007), pp. 1258–1269. ISSN: 10459219. DOI: 10.1109/TPDS.2007.1049.
- [33] Guy M Morton. “A computer oriented geodetic data base and a new technique in file sequencing”. In: (Mar. 1966), p. 302.
- [34] Luiz Velho and Jonas de Miranda Gomes. “Digital halftoning with space filling curves”. In: *ACM SIGGRAPH Computer Graphics* 25.4 (July 1991), pp. 81–90. ISSN: 00978930. DOI: 10.1145/127719.122727.
- [35] Robert W Floyd. “An adaptive algorithm for spatial gray-scale”. In: *Proc. Soc. Inf. Disp.* Vol. 17. 1976, pp. 75–77.
- [36] J. F. Jarvis, C. N. Judice, and W. H. Ninke. “A survey of techniques for the display of continuous tone pictures on bilevel displays”. In: *Computer Graphics and Image Processing* 5.1 (1976), pp. 13–40. ISSN: 0146664X. DOI: 10.1016/S0146-664X(76)80003-2.
- [37] Peter Stucki. *MECCA - A Multiple-Error Correcting Computation Algorithm for Bilevel Image Hardcopy Reproduction*. Tech. rep. Zurich: IBM Research Lab, 1981.
- [38] Haisen Zhao et al. “Connected fermat spirals for layered fabrication”. In: *ACM Transactions on Graphics* 35.4 (2016), pp. 1–10. ISSN: 07300301. DOI: 10.1145/2897824.2925958.

- [39] M. F. Ashby. “The properties of foams and lattices”. In: *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 364.1838 (2006), pp. 15–30. ISSN: 1364503X. DOI: 10.1098/rsta.2005.1678.
- [40] M. Avalle, G. Belingardi, and R. Montanini. “Characterization of polymeric structural foams under compressive impact loading by means of energy-absorption diagram”. In: *International Journal of Impact Engineering* 25.5 (2001), pp. 455–472. ISSN: 0734743X. DOI: 10.1016/S0734-743X(00)00060-9.
- [41] N. J. Mills et al. “Polymer foams for personal protection: Cushions, shoes and helmets”. In: *Composites Science and Technology* 63.16 (2003), pp. 2389–2400. ISSN: 02663538. DOI: 10.1016/S0266-3538(03)00272-0.
- [42] Ultimaker. *Technical data sheet TPU 95A*. Tech. rep. 2018, pp. 1–3. URL: <https://ultimaker.com/download/74605/UM180821%20TDS%20TPU%2095A%20RB%20V10.pdf>.
- [43] Andrew O. Payne (Autodesk) Michalatos, Panagiotis (Harvard Graduate School of Design). *Autodesk Monolith*. 2018. URL: <http://www.monolith.zone/>.

IV

ENABLING SPATIALLY GRADED THICKNESS

In the previous chapter we presented a method to generate FGMs with spatially varying stiffness, consisting of strictly continuous constant width toolpaths. While strict continuity and constant width help to manage the extrusion process, they lead to very stringent manufacturing constraints. If toolpaths are not allowed to overlap and have a constant width then we can only manufacture geometry with a thickness of an integer multiple of the extrusion width. Such an integer thickness constraint impedes a class of FGMs which alter the mechanical properties of microstructures by spatially grading their thickness.

In this chapter we propose a method for generating varying width toolpaths to fill geometry of varying thickness and a control mechanism to achieve varying width extrusion lines. By looking at the radial distance of the outline polygons of a layer we can decide on the width of each of the extrusion lines across the width of the polygons. The toolpath widths can then be chosen such that they stay close to the nozzle size which limits the impact of extrusion inaccuracy. As such we eliminate the integer thickness manufacturing constraint and thereby unlock the class of FGMs which modulates thickness in order to achieve a spatial gradation.

Sections IV.1 and IV.2 provide an introduction and related work to the presented technique. Section IV.3.2 presents the general framework for controlling the number and width of toolpaths to fill geometry of an arbitrary thickness and Section IV.4 proposes several schemes which determine the number and width in different ways. Section IV.5 pertains to the fabrication of variable width toolpaths: a control mechanism to achieve it and its results. Section IV.6 evaluates the framework and the various schemes and discusses its limitations and applications, in order to arrive at the conclusion presented in Section IV.7.

This chapter is published as Tim Kuipers, Eugeni L. Doubrovski, Jun Wu, and Charlie C.L. Wang. "A Framework for Adaptive Width Control of Dense Contour-Parallel Toolpaths in Fused Deposition Modeling". In: *Computer-Aided Design* 128 (Nov. 2020), p. 102907. ISSN: 00104485. DOI: 10.1016/j.cad.2020.102907; some minor corrections have been made.

The method described in this chapter is patent pending [2], but open source available.

ABSTRACT

3D printing techniques such as Fused Deposition Modeling (FDM) have enabled the fabrication of complex geometry quickly and cheaply. High stiffness parts are produced by filling the 2D polygons of consecutive layers with contour-parallel extrusion toolpaths. Uniform width toolpaths consisting of inward offsets from the outline polygons produce over- and underfill regions in the center of the shape, which are especially detrimental to the mechanical performance of thin parts. In order to fill shapes with arbitrary diameter densely the toolpaths require adaptive width. Existing approaches for generating toolpaths with adaptive width result in a large variation in widths, which for some hardware systems is difficult to realize accurately. In this paper we present a framework which supports multiple schemes to generate toolpaths with adaptive width, by employing a function to decide the number of beads and their widths. Furthermore, we propose a novel scheme which reduces extreme bead widths, while limiting the number of altered toolpaths. We statistically validate the effectiveness of our framework and this novel scheme on a data set of representative 3D models, and physically validate it by developing a technique, called *back pressure compensation*, for off-the-shelf FDM systems to effectively realize adaptive width.

IV.1 INTRODUCTION

3D printing enables the fabrication of complex geometry under few design constraints compared to conventional fabrication techniques. Recent developments have seen a rapid growth in both the use and capabilities of desktop 3D printing systems. Fused Deposition Modeling (FDM) is one of the most common 3D printing techniques. It is widely used because of the versatility in the types of plastic which can be used and the relatively low running costs. FDM printers are used, for example, in showcasing scale models of buildings, casings for electronics, prototypes for blow molded parts, jigs and fixtures. Recent research addressed manufacturing complex volumetric structures such as microstructures [3, 4, 5] and topology optimized structures [6, 7, 8]. Many of these applications involve 3D models with detailed features within the order of magnitude of the nozzle size, which restrains the field of the process planning algorithms.

FDM printers extrude semi-continuous beads of molten plastic through a nozzle, which moves along a planned toolpath within each layer of a 3D object. A common strategy is to extrude along a number of parallel toolpaths which follow the shape of the contour of the layer and fill up the remaining area using parallel straight toolpaths. Contour-parallel toolpaths fit to the layer outlines more accurately, because the resolution of the positioning system is an order of magnitude smaller than the size of the hole in the nozzle. This paper is concerned with the generation of such contour-parallel toolpaths and addresses several issues which commonly occur in 3D models with narrow geometry.

The simple technique for generating the dense contour-parallel toolpaths of a layer consists of performing uniform inward offsets with the size of the nozzle from the outline shape. However, for geometrical features which are not an exact multiple of the nozzle size this method produces areas where an extrusion bead is placed twice: *overflow* areas; and areas which are not filled at all: *underfill* areas. See Fig. IV.1a. Overfills cause a buildup of pressure in the mechanical extrusion system, which can result in bulges or even a full print failure. Underfills, on the other hand, can cause a drastic decrease in the part stiffness or even for small features not to be printed at all. These problems are exacerbated for models

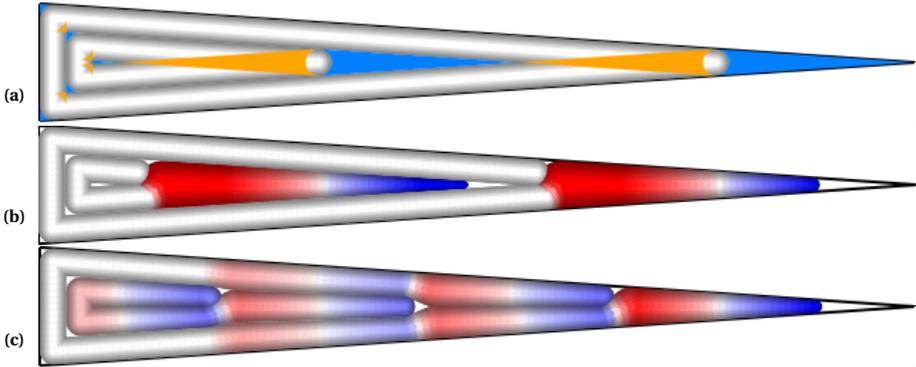


Fig. IV.1. Illustration of different toolpaths for a shape showcasing a range of shape radii (black). These results can be read as a graph with feature size on the horizontal axis and its corresponding beading along the vertical axis. (a) Toolpaths using uniform offsetting results in large overfill (orange) and underfill (azure). (b) Toolpaths with adaptive width [11] where beads that are wider or narrower than the nozzle size are indicated in red and blue, respectively. (c) Our approach minimizes over- and underfill with less extreme widths.

with layer outlines with small features, because the over- and underfill areas are relatively large compared to the those features.

One promising direction to avoid over- and underfills is to employ toolpaths with adaptive width. Ding et al. developed a toolpath strategy for wire and arc additive manufacturing which produces a width variation typically lower than a factor of 3, but is far greater for some parts [9, 10]. However, the range of bead widths manufacturable by FDM systems is limited. A nozzle of $w = 0.4$ mm will typically start to cause fluttered extrusion around lines narrower than 0.3 mm and lines will start to bulge upward if they are wider than the flat part of the nozzle, which is typically 1.0 mm.

The current state of the art of contour-parallel toolpath generation developed by Jin, Du, and He employs a strategy which alters the widths of the centermost beads within a range of widths $[0.25w, 1.8w]$ [11], which is similar to the strategy employed by the open source industry standard software package Ultimaker Cura [12]. See Fig. IV.1b. Still, controlling the extrusion width through movement speed changes or through volumetric flow control (e.g. linear advance) yields diminishing accuracy for deposition widths deviating more from the nozzle size, since process parameters such as nozzle temperature are optimized for beads with the nozzle size. Moreover, reducing the variation in width is beneficial for limiting the variation in mechanical properties of the resulting product, meaning it conforms better to a simulation which employs a homogeneity assumption. We therefore reduce the bead width range by distributing the workload from the centermost bead over neighboring beads.

Our contributions are as follows:

- A geometric framework allowing various adaptive bead width control schemes used to generate contour-parallel toolpaths which minimize under- and overfill.
- A specific beading scheme, which reduces the variation in the extrusion widths to within $[0.75w, 1.5w]$.
- A back pressure compensation approach to accurately realize adaptive bead width on Bowden style hardware systems.

IV.2 RELATED WORK

Toolpath generation consists in generating a path in the a planar contour, representing the intersection of a plane and a 3D solid object. The nozzle is then instructed to move along the path while extruding material. Sites along the toolpaths are assigned several properties such as movement speed, but for this paper we will focus on the assigned width of the extruded bead. Toolpath generation is an integral part of process planning for 3D printing. For an overview of the processing pipeline, we refer to the survey by Livesu et al. [13]. For reducing printing time and material cost, sparse infill structures such as triangular and hexagonal patterns have been used to approximate the interior of 3D shapes. In this paper, we focus on the generation of toolpaths for dense regions, such as the boundary shell of 3D shapes. This is sometimes called 'dense infill' [13].

The toolpath has a direct influence on the printing time, material cost, and mechanical properties of the printed object [14, 15]. FDM calls for toolpaths with several desirable properties. First, the extrusion path should be as continuous as possible. A discontinuous path requires to stop and restart material extrusion. For certain materials, such as TPU, this could lead to printing defects or even print failure [16]. Second, the toolpath is preferred to be smooth. Sharp turns require to reduce the movement speed of the nozzle, and so this prolongs the printing process. Third, the extruded path should cover the region of the contour without underfilling. Such underfill negatively influences the mechanical performance of the parts. Fourth, the extrusion paths should not overlap with one another. Such overfill causes a pressure build up in the mechanical system, which leads to over-extrusion in later paths and in extreme cases cause print failure [16]. Han et al. presented a method for analyzing the under- and overfill of a vertical cross-section [17]. Our method is primarily concerned with minimizing under- and overfill within horizontal cross-sections.

Two basic strategies for dense toolpath generation are the direction-parallel strategy and the contour-parallel strategy. Direction-parallel (or zig-zag) toolpaths fill an arbitrarily shaped contour with a set of parallel, equally spaced line segments. These parallel segments are linked together at one of their extremities, to avoid discontinuous extrusion. Contour-parallel toolpaths typically consists of a set of equally spaced offsets from the slice boundary outline polygons. Steuben, Iliopoulos, and Michopoulos presented a method for generating sparse infill toolpaths based on the isocontours of surface plots of some variable generated on each 2D contour [18]. In order to increase the continuity of contour-parallel toolpaths, a strategy to connect dense toolpaths into spirals was introduced by Zhao et al. [19] and later extended to also connect a mixture of dense and sparse toolpaths together [16]. Jin et al. discusses several approaches for connecting direction-parallel and contour-parallel toolpaths into continuous paths [20]. Spiral toolpaths have also been applied to (CNC) machining [21, 22]. One of the problems with contour-parallel toolpaths is that it tends to leave gaps between the toolpaths (see Fig. IV.1a). This is due to the fact that the diameter of the part is not exact multiple of the (constant) deposition width in those regions. To avoid problems with such gaps, hybrid approaches that combine direction and contour-parallel are often used [23, 24]. Close to the slice boundary, there are several contour-parallel curves, while the interior is filled using a zig-zag pattern. For complex shapes, the entire cross-section could be decomposed into a set of patches, and for each of them the basic strategies can be applied [25, 20]. Alternative toolpath patterns, seen also in CNC machining, include space-filling curves [26, 27, 28].

Reducing under- and over-filling can be accurately achieved by making use of adaptive

deposition width. Adaptive width can be used to locally match the nonuniform space between adjacent paths, and thus to ensure a better filling of the area. Kao and Prinz propose smooth adaptive toolpaths based on the medial axis skeleton of the boundary contour [29]. Their approach handles simple geometry where there are no branches in the medial axis. An extension was proposed by Ding et al. to handle complex shapes [9]. However, this extension inherits a problem in the original method: from any point in the skeleton to the boundary, the number of toolpaths is constant. Depending on the size of small and large features in the layer outlines, this strategy can require a range of bead widths beyond the capabilities of the manufacturing system. Jin, Du, and He proposed a strategy of adding toolpaths with varying width along the center edges of the skeleton, while leaving other paths unchanged [11]. The resulting beads have widths within the wide range of $[0.25w, 1.8w]$ (see Fig. IV.1b). In this paper we propose a novel scheme to distribute the width alterations throughout a region around the center, and thus limit the occurrence of extreme variation in width (see Fig. IV.1c).

Under- and over-filling issues have a large proportional impact for thin geometric features. Jin, He, and Du proposed a sparse wavy path pattern for thin-walled parts [30]. Besides under- and over-filling, there are a few other robustness issues in toolpath generation for thin geometric features. Moesen et al. proposed a method to reliably manufacture thin geometric features using laser-based additive manufacturing techniques [31]. » achange!! Behandish, Mirzendehtel, and Nelaturi presented a method to characterize local- topological discrepancies due to material under- and over-deposition, and used this information to modify the as-manufactured outcomes [32].

For ease of reference we have included a legend showing the terms employed in this manuscript in Fig. IV.4. These terms will be further explained as they first appear throughout this paper.

IV.3 METHOD

IV.3.1. OVERVIEW

Given arbitrarily shaped polygons which represent the 2D outline of a layer of a 3D model, our method generates extrusion toolpaths with varying width, i.e. a set of polylines where each site consists of a location and an extrusion width; in between the sites we linearly interpolate the position and extrusion width.

Our method starts with computing a graph which represents the topology of the input polygon: its *skeleton*. Our skeleton is based on the medial axis transform (MAT), a strategy that has been commonly used for generating contour-parallel toolpaths [33]. We visualize the skeleton as the union of cones (UoC) (Section IV.3.2), by raising each point in the domain to a height that equals the shortest distance from the point to the polygon boundary (Fig. IV.2c). Contour-parallel toolpaths with uniform width can be interpreted as the intersection of the union of cones with a set of horizontal planes at equally spaced heights (Figs. IV.2a to IV.2e).

As depicted in Fig. IV.2f, our method first identifies edges and nodes of the skeleton in the center of the polygon, which correspond to ridges and peaks in the UoC (Section IV.3.3). The heights \hat{b} at these elements are then quantized to an integer number of beads \hat{b} . To ensure a smooth toolpath between regions with quantized integer heights that differ, we add new nodes in the skeleton with quantized heights and interpolate the heights \hat{b} in between (Section IV.3.4). The union of cones corresponding to the smoothed

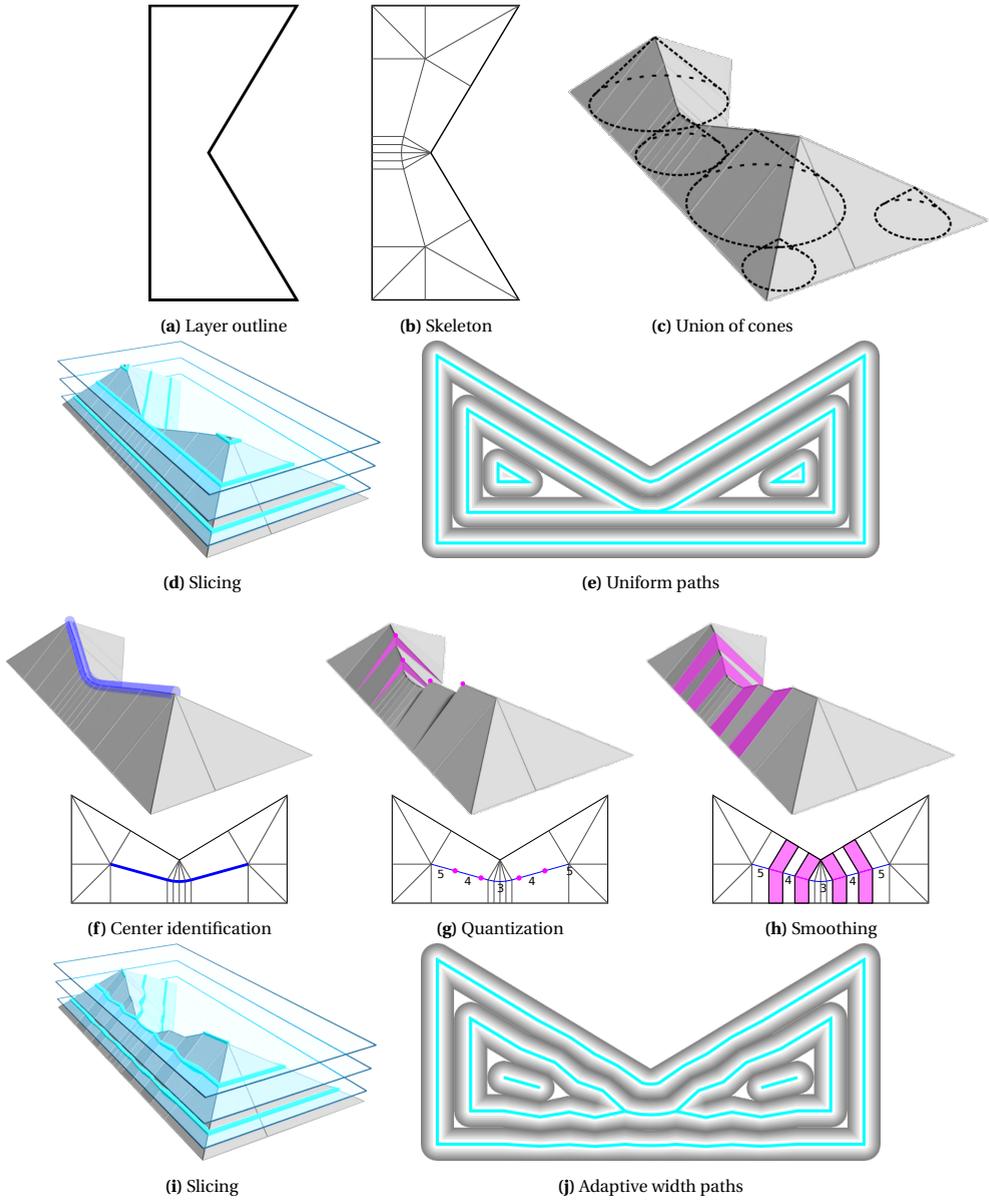


Fig. IV.2. The first row illustrates the generation of uniform paths (e) by interpreting the path as the intersection between horizontal planes and the union of cones (c), which is an alternative visualization of the skeleton (b). The second row depicts the stages with both 2D and 3D visualizations for generating paths with adaptive width (j). Central elements in the skeleton are first identified (blue in (f)). The heights are then quantized in terms of number of beads (the integer values in (g)), and smoothed (h).

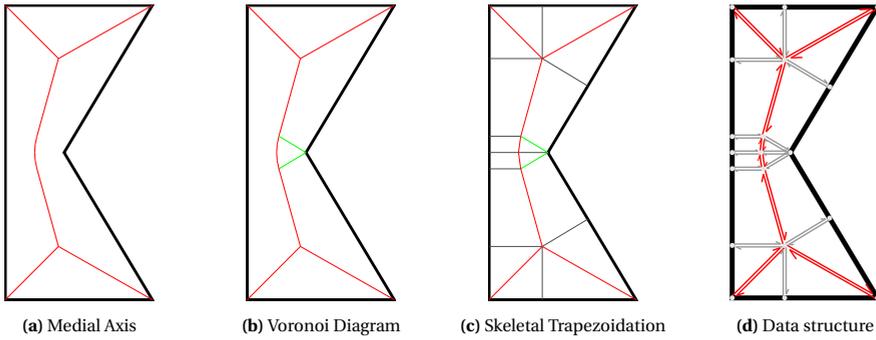


Fig. IV.3. Skeletonization of an outline shape (black). Relation between the medial axis (red), the limited Voronoi Diagram (red and green) and the Skeletal Trapezoidation (red, green and gray): $\text{MAT} \subset \text{Limited VD} \subset \text{ST}$. (d) The skeleton is represented using a half-edge data-structure.

skeleton is then sliced at regular intervals to obtain toolpaths with varying spacing, which translates into varying width (Section IV.3.6). The video in the supplementary material provides an example animation of this approach.

This section explains how we generate toolpaths using our framework with uniform bead widths and evenly distributed locations between the center of the polygon and the outline. In Section IV.4 we describe how to apply the framework to different beading schemes and we show several such beading schemes.

IV.3.2. UNION OF CONES

The union of cones (UoC) is derived from a common skeletonization of the polygonal outline shape: the medial axis. By assigning each node in the skeleton a height equal to its shortest distance to the outline we obtain the shape of the UoC. Starting from the medial axis we further decompose the shape into simple fragments, so that the domain contains only quads and triangles. This decomposition constitutes an approximation of the UoC.

Medial axis transform The medial axis is a representation commonly used to analyze a shape. It is defined as the set of positions where the inscribed circle meets the boundary in at least two locations [34, 35]. The resulting skeleton consists of straight edges and parabolic edges. An example is illustrated in Fig. IV.3a. We call the set of points on the outline polygon P closest to a skeletal point v its *support*:

$$\text{sup}(v) = \underset{x \in P}{\text{argmin}} |x - v|. \quad (\text{IV.1})$$

The shortest distance for a point on the skeleton is called its feature radius, $R(v)$. The medial axis along with the feature radius values along the skeleton form a complete shape descriptor, known as the medial axis transform (MAT).

By vertically raising the center of an inscribed circle to a height that equals the center's feature radius, a cone is formed. The union of all such cones forms a 3D solid volume. The medial axis can thus also be interpreted as ribs of the surface of the union of cones [34].

Skeletal trapezoidation Starting from the medial axis we decompose the input polygon into a set of quads and triangles, so that we can perform the slicing stage on simple shapes.

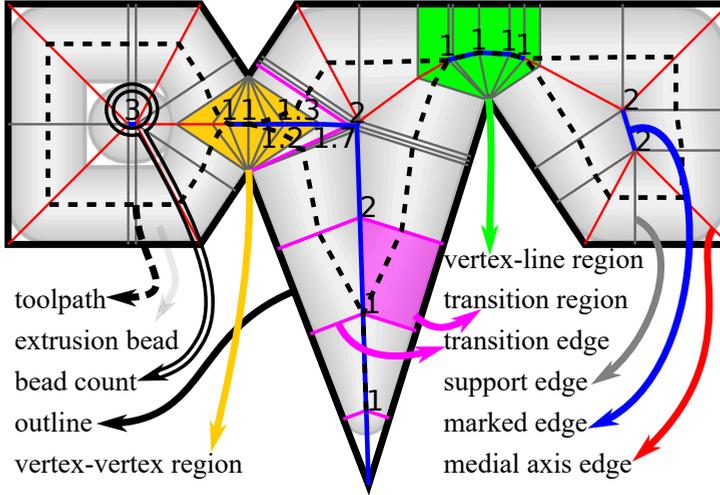


Fig. IV.4. Illustrative explanation of terms and color coding that are consistently used in this paper.

We employ a shape decomposition similar to the one proposed by Ding et al. [9]. The basic idea is to add edges connecting each node v on the medial axis to each of its support points $\text{sup}(v)$. The resulting skeleton decomposes the outline shape into trapezoids and triangles. Considering the fact that the concept of trapezoidation conventionally allows for the degenerate case where a trapezoid resolves into a triangle [36, 37], we call this shape decomposition the *Skeletal Trapezoidation (ST)*.

The edges generated by the MAT are classified into three types: 1. line-line edge – straight edge generated from two line segments in the outline polygons, 2. vertex-line edge – parabolic edge resulting from an outline vertex and a line segment in the outline, and 3. vertex-vertex edge – straight edge resulting from two outline vertices. The vertex-line and vertex-vertex edges are discretized into pieces with a length up to 0.2 mm, which gives an approximation error of only ± 0.01 mm. This allows to approximate the feature radius between two discretized nodes v_0 and v_1 by linear interpolation. Again we connect the newly inserted nodes to their support, which results in vertex-line regions and vertex-vertex regions such as depicted in Fig. IV.4.

Approximation of union of cones The skeletal trapezoidation (ST) provides a means to visualize the union of cones (UoC) approximated by a 3D surface mesh composed of quadrilateral and triangular patches. We assign each node in ST a (real number) height value measured in terms of beads, referred to as the *bead count* b . We define the bead count as the number of beads to fit along the *diameter* of the inscribed circle centered at node v , i.e. $2R(v)$, by

$$\tilde{b}_v = 2R(v) / w^* \quad (\text{IV.2})$$

where w^* is the nozzle size. We divide the diameter rather than the radius as this allows to deal with an odd number of beads while using integer logic. Note that although the overview of the method was described geometrically in terms of the UoC, the actual tool-path generation relies on the two-dimensional ST; the use of the bead count as a height value is only a visualization aid.

Implementation The medial axis of a polygonal shape is a subset of the Voronoi Diagram generated from the line segments and vertices of the shape [35]. The edges of the Voronoi diagram that fall outside of the outline shape are irrelevant for our purpose and are thus discarded. Note that besides the full medial axis, the Voronoi diagram also contains edges connecting to concave vertices in the outline shape (see Fig. IV.3b). These extra edges are a subset of the edges connecting a node to its support, so we keep them in. From the Voronoi diagram we add nodes to discretize parabolic edges and edges formed by two concave outline vertices, and then connect all nodes to their supports, forming a skeletal trapezoidation. We then assign each node the bead count values using Eq. (IV.2). We compute the Voronoi diagram using the Boost C++ libraries [38], which implements the algorithm proposed by Fortune [39]. A half-edge data-structure is used to represent the Voronoi diagram (Fig. IV.3d).

IV.3.3. CENTER CLASSIFICATION

In order to prevent over- and underfill from occurring in central regions, parts of the ST are marked as being central.

Our framework will decide on a beading at all the marked nodes in ‘the center’ and apply the beading outward to the unmarked nodes (Section IV.3.5).

A node in the ST is marked as as central if its feature radius is larger than that of all its neighboring nodes, i.e. a local maxima. An edge and its two nodes are also marked as being central if it is significant according to a significance measure.

Significance measure We make use of the *bisector angle* as an indicator of significance which is commonly used in shape analysis. The bisector angle α is the interior angle $\angle p_0 l p_1 \leq 180^\circ$, between any location l on an edge of the ST and its two supporting points $\{p_0, p_1\} = \text{sup}(l)$ [40]. An edge is significant if the bisector angle on any location on the edge exceeds a prescribed α_{\max} . As illustrated in Fig. IV.5a, for a polygon with a pointy wedge area of an angle β , we have $\alpha = 180^\circ - \beta$. This corresponds to overfill areas and underfill areas the size of $1/4(w^*)^2 (\tan(\alpha/2) - \alpha/2)$ when filled using the simple technique of uniform bead width w^* . A too large α_{\max} may leave a lot of under-/overfill, while a too small value may introduce toolpaths to fill in negligibly small underfills. We therefore set $\alpha_{\max} = 135^\circ$. Although significance measures are commonly used as a heuristic for finding the parts of a skeleton which are in some sense relevant [40, 41], we use the bisector angle as an *exact* indicator of the amount of overfill and underfill in the uniform toolpaths of constant width.

To avoid evaluating the bisector angle at any location on all edges, we devise an efficient measure which operates only on the two nodes of an edge. Because all locations along a line-line edge have the same bisector angle we can evaluate whether the edge is significant by checking whether

$$|R(v_1) - R(v_0)|/|v_1 - v_0| > \cos(\alpha_{\max}/2) \quad (\text{IV.3})$$

(see Fig. IV.5b). This ratio has a clear geometrical interpretation as the slope of the ridge in the UoC surface. For vertex-line edges and vertex-vertex edges only a portion of the edge is significant. We therefore introduce nodes at the boundaries of the significant portion during the discretization of such edges (see Appendix IV.A). The significance of all edges can then accurately be evaluated using Eq. (IV.3).

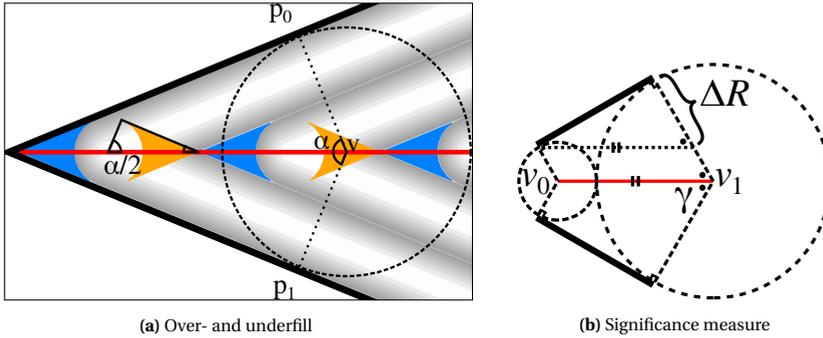


Fig. IV.5. Properties of the significance measure along a skeletal edge (red) generated from two polygon lines (black) using the properties of inscribed circles (gray) and their radii (dashed). (a) The size of overfill (orange) and underfill areas (azure) for the uniform toolpathing technique can be calculated from the bisector angle. (b) The significance measure can be simplified using $\alpha = 2\gamma = 2\cos^{-1}\Delta R/|v_1 - v_0|$.

Marking filtering After initializing the marking at all edges and nodes, we filter out high frequency changes in the marking in order to ensure that the generated toolpath is smooth. The filtering is performed by additionally marking some unmarked elements, rather than the opposite since unmarking central regions reintroduces large over- and underfill areas. From each marked node v_0 with an upward unmarked edge attached we walk along the upward edges; if the total length traversed until we reach another marked node v_1 is shorter than some filter distance $d_{\max}^{\text{unmarked}}$, we mark all edges encountered as being central. We use $d_{\max}^{\text{unmarked}} = w^*$ in order to filter out high frequency oscillations in the order of magnitude of the nozzle size, while keeping close to the significance measure.

IV.3.4. CENTRAL HEIGHT ADJUSTMENT

After the central regions have been identified, their heights are quantized. First, the initial bead count \bar{b} is quantized into an integer bead count \hat{b} at the marked nodes using a quantization operator q , then the locations along the edges where q makes a jump from one bead count n to another $n+1$ are identified and then ramps are introduced to smoothly transition from n to $n+1$ using fractional bead counts \hat{b} along the smooth transition.

Quantization We define a quantization operator q to map a feature diameter ($d = 2R(v)$) to a bead count: $q: \mathbb{R} \rightarrow \mathbb{N}$. Because our quantization scheme should round to the nearest integer multiple of the nozzle size, we have $q(d) = \lfloor d/w^* + 1/2 \rfloor$. Alternative quantization schemes are discussed in Section IV.4. By applying q to the heights of central nodes we quantize the bead count:

$$\bar{b}_v = q(2R(v)) = \lfloor \bar{b}_v + 1/2 \rfloor \quad (\text{IV.4})$$

Transition anchors For a marked edge which connects nodes v_0 and v_1 with $\bar{b}_{v_0} \leq n < \bar{b}_{v_1}$, we determine the *transition anchor locations* at which the bead count transitions from n to $n+1$. To this end, we introduce the function

$$q^{-1}(n) := \arg\max_d q(d) = n, \quad (\text{IV.5})$$

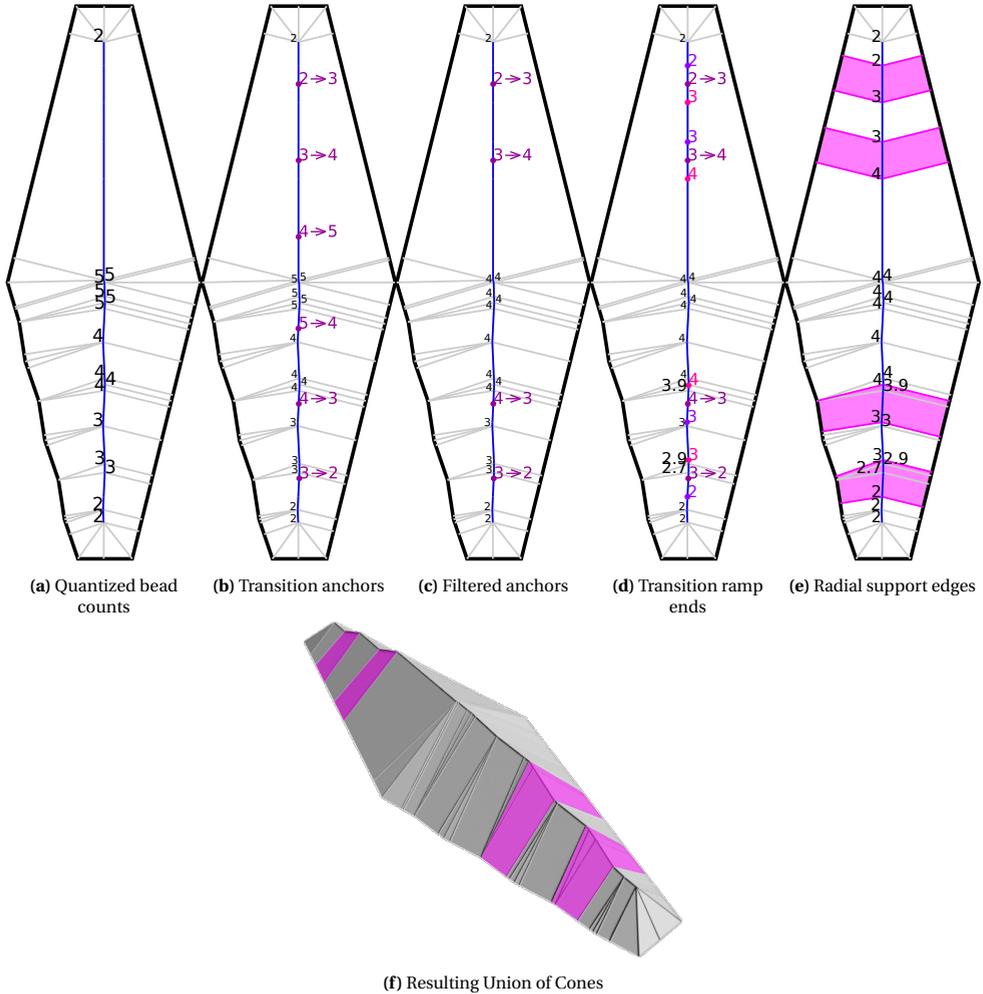


Fig. IV.6. Applying bead counts and transitioning on a shape showing the difference between a simple ST (top) and a mirrored version with small perturbations in the outline (bottom). Outline in black, central edges marked in blue, radial edges in grey. (a) First we initialize the bead counts (black) in the marked edges (blue). (b) We then extract the anchor locations (purple) where the bead count transitions. (c) We then filter out regions which exhibit frequent transition. (d) We then calculate the end locations (magenta, pink) of the transitions and modify the bead count at nodes in between to fractional values. (e) Finally we introduce nodes at the ends and introduce radial edges (purple) as per the trapezoidation constraint. The symmetry in the result shows that transitioning is robust against small perturbations in the outline shape.

which gives the feature diameter d at which the bead count q transitions from n to $n + 1$. The location of the anchor v_x is then computed by inversely interpolating $R(v_x) = q^{-1}(n)$, i.e.

$$v_x = v_0 + (v_1 - v_0) \frac{q^{-1}(n) - R(v_0)}{R(v_1) - R(v_0)}. \quad (\text{IV.6})$$

An illustration of the anchors is shown in Fig. IV.6b.

We perform a filtering step to prevent frequently changing the bead count back and forth within a short distance. For two consecutive anchors which transition to opposite directions, if the distance between them is smaller than some limit $d_{\max}^{\text{transition}}$, the bead counts at all nodes in-between are set to the surrounding bead counts, and consequently these anchors are removed (See Fig. IV.6c). A value of $d_{\max}^{\text{transition}} = 1$ mm seems to produce satisfactory results. This means that for some small regions we generate toolpaths with bead widths outside the typical range.

Smooth transitions A sharp transition from n to $n + 1$ beads at an anchor location creates sharp turns in the toolpath (see Fig. IV.7 top). We introduce a transition length $t(n)$ to ensure a smooth transition (see Fig. IV.7). The length of the transition is set to $t(n) = w^*$ and it is centered at the anchor, i.e. the distance from the lower end v_0 to the anchor position v_x is set to $t_0(n) \equiv \Delta(v_0 v_x) = t(n) (q^{-1}(n) / w^* - n)$, where Δ is the total distance along the edges between two nodes. The transition length $t(n)$ ensures that the center beads do not overlap with the innermost transitioning beads, while keeping the amount of underfill low and the toolpath smooth. The transition anchor position $t_0(n)$ ensures that the transitions never overlap with each other or with locations where all beads have the preferred width w^* .

We discard any transition anchor which is too close to the end of a chain of marked edges for the smoothed transition to fully fit within the marked region. In order to make the transition ramps robust against small perturbations in the outline shape which cause extra (support) edges in the skeleton, we modify the nodes v_x which are between the two ends v_0 and v_1 of the transition by (re-)assigning them a fractional bead count \hat{b} which is linearly interpolated between the two ends of the transition (see Fig. IV.6e):

$$\hat{b}_{v_x} = n + \Delta(v_0 v_x) / \Delta(v_0 v_1) \quad (\text{IV.7})$$

Note that although the ST is not stable against noise in the boundary shape, the distance field itself is, so by designing our algorithms such that they are stable against changes in the topology of the skeleton our method is stable against small perturbations in the outline. Finally we update the ST by adding support edges at the transition ends. As shown in Fig. IV.6f, the marked regions in the UoC mesh have become horizontal at integer multiples of $1/2w^*$ for long stretches with ramps in between.

IV.3.5. BEADING

Now that we know how to determine the bead counts in the marked central regions the question is how the unmarked regions are handled. Determining bead count values for the unmarked nodes and interpolating linearly along the unmarked edges would mean that toolpath sites would be distributed evenly along each unmarked bone; while that would suffice for the evenly distributed beading scheme, it would not allow for more sophisticated, non-linear schemes. Instead we determine the radial distance to the boundary at

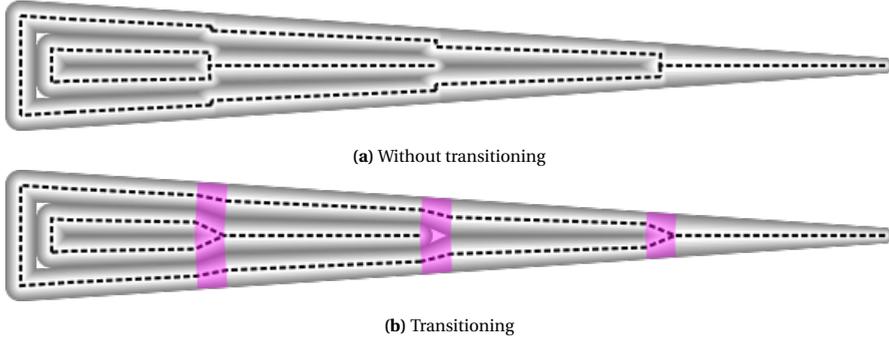


Fig. IV.7. Sharp turns around regions where the bead count changes are prevented by transition regions (highlighted in cyan).

which each bead should occur from the boundary to the center. Each central node is associated with a sequence of radial distances L which control the locations of the beads, starting from the outer bead and ending in the center. Together with a sequence of bead widths W , these form what we call a *beading* B . For our distributed beading scheme we compute the beading for a central node ν with $n = \lfloor \hat{b}_\nu \rfloor$ beads and a diameter $r = R(\nu)$ as:

$$\begin{aligned}
 B(n, r) &= (W(n, r), L(n, r)) &&= (\{w_0 \dots w_{\lfloor n/2 \rfloor - 1}\}, \{l_0 \dots l_{\lfloor n/2 \rfloor - 1}\}) \\
 w_i &= r/n &&\text{for all } i \in \mathbb{N} : i < n/2 \\
 l_i &= r/n(i + 1/2) &&\text{for all } i \in \mathbb{N} : i < n/2
 \end{aligned}$$

where w_i and l_i are the width and location of the i th bead, respectively, counting from the outline inward. Example beadings for an odd and even bead count with arbitrary widths are visualized in Fig. IV.8a.

Beading interpolation The beading is defined in terms of an integer number of beads, while we have assigned a fractional bead count to nodes within a transition region. In order to generate a beading for a node ν with $n < b_\nu^* < n + 1$ we linearly interpolate the bead widths and locations between a beading B^1 based on n and a beading B^2 based on $n + 1$ (see Fig. IV.8). Such interpolation is also used to deal with beading conflicts (see Fig. IV.9). There we also apply beading interpolation from a marked node ν_m upward along unmarked bones, and interpolate between ν_m and the beading at the top of the slope over some distance t_{beading} from the lower marked node, which we set to $t_{\text{beading}} = w^*$, so that the transition is not too swift.

Beading propagation The beading information is then broadcast throughout the ST from central regions outward, so that each unmarked node ν knows the beading of the marked node on top of the ramp on which ν is placed. We first broadcast the beading information upward from all marked nodes, so that we can then deal with beading conflicts in a downward phase. The downward phase makes sure that all nodes have a beading associated with it, so that the slicing algorithm can efficiently slice the edges leading up to a marked or unmarked node.

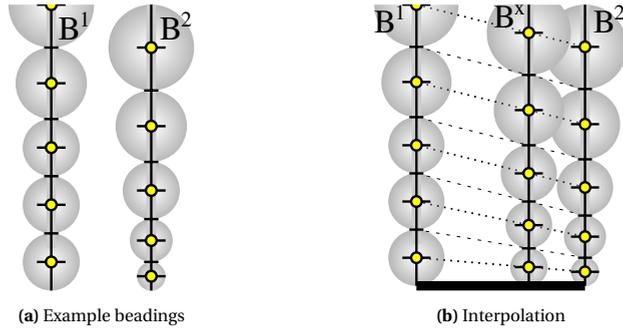


Fig. IV.8. Interpolation between two beadings B^1 and B^2 with odd and even bead count resulting in a beading B^x at $n + 2/3$. Bead indices are counted inward from the outline (thick black). Interpolation of locations in dots, interpolation of widths in dashes.

IV.3.6. TOOLPATH EXTRACTION

Once each node has been assigned a beading, we proceed to generate the toolpath sites along the edges of the ST. A site S consists of a location v a width w and an index i , which are computed for an edge $v_0 v_1$ from the beading B of the upper node v_1 :

$$S = \{v, w, i\}$$

$$v = v_1 + (v_0 - v_1) \frac{R(v_1) - l_i^B}{R(v_1) - R(v_0)}$$

$$w = w_i^B$$

for any i for which $R(v_0) < l_i^B \leq R(v_1)$. See Fig. IV.10. We store all sites of an edge in a mapping from edge to a list of sites.

We then generate extrusion segments for each trapezoid by connecting together the sites of the same index. See Fig. IV.11. If the amount of sites on both sides of the trapezoid is not the same then this trapezoid is in a transition and we leave one inner site unconnected.

Because the bead count is defined in terms of the feature diameter rather than the radius, only some of the bead count values \hat{b} in a central region coincide with a slicing height. When the bead count \hat{b} is even, the ridge is sliced as normal; the intersection between a slicing plane and the mesh surface results in a polyline on both sides of the ridge, which are connected together into a polygonal toolpath. When the bead count \hat{b} is odd, the ridge will coincide exactly with a slicing height, which results in a single polyline toolpath being generated along the middle of the feature. In that case we should prevent the algorithm from generating the center extrusion segment twice from the trapezoids on either side of that segment. We therefore use some arbitrary condition to decide which one of the two to include based on the ordering of the coordinates of v_0 and v_1 : $x_0 < x_1 \vee (x_0 = x_1 \wedge y_0 < y_1)$.

All trapezoids in the ST are assigned to separate domains, corresponding to which boundary polygon they are connected to (see Fig. IV.11a) [9]. By traversing the trapezoids per domain in order we can efficiently connect all segments into polylines. See Fig. IV.11. In a final step we connect the ends of polylines together, so that the final toolpaths contain both polygons and polylines.

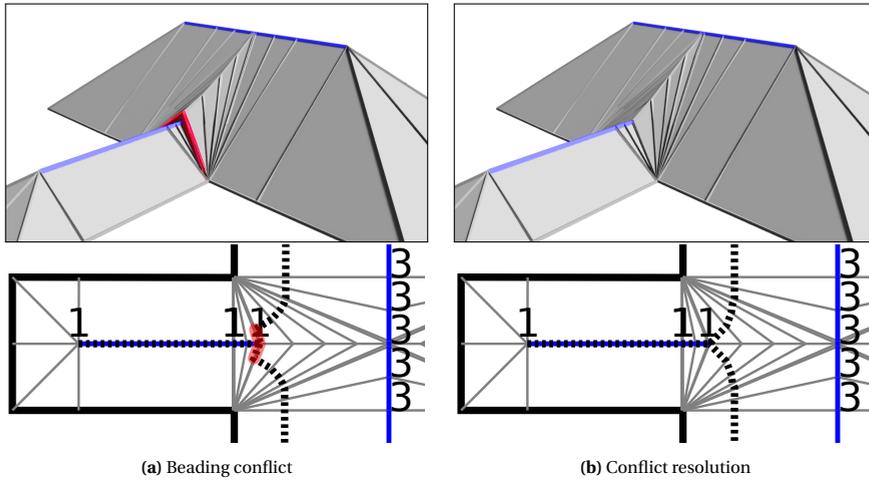


Fig. IV.9. (a) The beading propagated from above conflicts with the beading below. (b) The beading conflict is resolved by gradually interpolating between the two beadings. The ramp to the upper ridge does not line up with the lower ridge, which means that the toolpaths (dashed) resulting from the beading propagated from above does not align with the beading from the thin outline feature (highlighted in red).

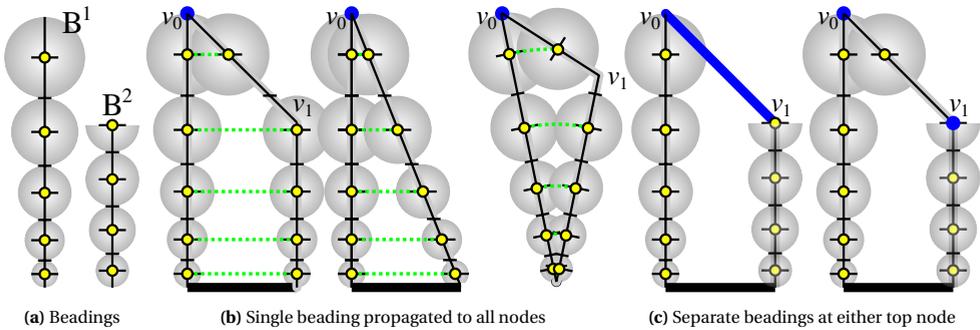


Fig. IV.10. Applying beadings to generate sites along trapezoids. (a) shows the locations l_i and widths w_i of two arbitrary different beadings. (b) shows the application of B^1 to the various types of trapezoid. (c) shows how a trapezoid with a marked edge will have two different beadings assigned, which will generate their respective sites along the support edges. No sites will be generated along marked edges. Wide black lines are outline segments, marked nodes and edges in blue, the sites in yellow and green wavefronts of equidistant radial distance at $R = l_i$.

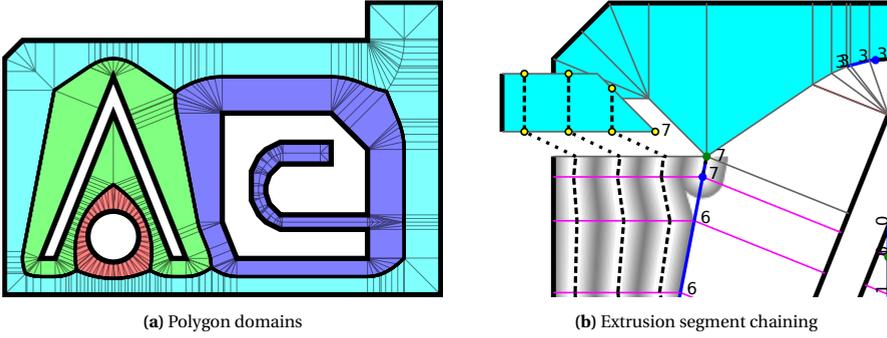


Fig. IV.11. Generating toolpaths on a part of the test outline shape by chaining together extrusion segments along each polygon domain. Each edge is assigned toolpath sites (yellow) which are connected together as shown in the singled out trapezoid. By following the trapezoids along the domain (cyan) of a single outline polygon, the extrusion segments can efficiently be connected into existing polyline toolpaths (light and dark gray).

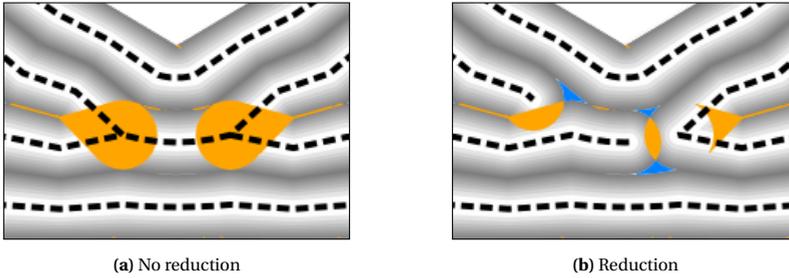


Fig. IV.12. Reducing polyline toolpaths away from intersections in order to prevent overfill. Toolpath locations in black, underfill in azure and overfill in orange.

Around the transition locations and around nodes with odd bead count and more than two marked edges attached there will be intersections in the toolpaths. Such intersections cause overfill because the nozzle passes the location multiple times. We deal with this special case by forcing a new polyline when traversing the trapezoids, and in the final polyline connection step we greedily connect the first two polylines ending in the same location and retreat all other polylines ending in that same location in order to prevent the overfill. In order to retreat a polyline which ends in a site S , we remove part of the polyline paths up to the intersection by a distance of $w^S d_{\max}^{\text{intersection}}$. We set $d_{\max}^{\text{intersection}} = 75\%$ in order to slightly favor overfilling over underfilling. This ratio effectively deals with the balance between overfill and underfill generated at that location after the retreat has been applied. See Fig. IV.12.

IV.4 BEADING SCHEMES

A critical component in toolpath generation is how to distribute the beads over the feature radius. While the framework presented in the previous section takes evenly distributed beads as an example, it allows to apply different beading schemes to configure the bead distribution to cater for specific requirements from the application, 3D printer or material.

Definition IV.4.1. A beading scheme is defined by the quantization operator q and

the beading operator $B: \{q(d), B(n, r)\}$. The beading function $B(n, r)$ consists of $(W(n, r), L(n, r))$, which provides sequences of n bead widths and of n distances from the outline to fill up a radial distance r .

For the smoothness and continuity of toolpaths we require that W_n is monotonic and continuous at each bead index n for constant bead count c : $0 \leq \frac{\partial W(c, r)_n}{\partial r} \leq 1$. We further ensure that beads do not overlap, that beads are extruded from the center of where they end up and that odd bead counts produce a single polyline toolpath exactly in the center by determining the bead locations from the widths:

$$L(n, r)_i = \begin{cases} -\frac{1}{2}W(n, r)_i + \sum_{j=0}^i W(n, r)_j & \text{if } i < \frac{1}{2}(n-1) \\ r & \text{if } i = \frac{1}{2}(n-1) \end{cases}$$

We introduce several beading schemes which determine the bead count and their widths in various ways. We can emulate a variety of toolpath generation methods from related literature by defining new beading schemes. We also introduce new beading schemes which produce toolpaths with less extreme widths compared to techniques from existing literature.

Uniform beading scheme We can define a beading scheme which emulates the uniform width offsetting technique by disabling the marking of edges, so that we never employ transitioning. We can simply set $\alpha_{\max} = 180^\circ$ and supply a simple beading scheme given by Table IV.1a.

Outer bead We can emulate the method from Moesen et al. by carefully choosing how the beading scheme functions deal with the outermost bead. Also we turn off the reduction of toolpaths near 3-way intersections $d_{\max}^{\text{intersection}} = 0\%$, so that the polygonal toolpaths emulate the remaining area to be filled by another path planning technique similar to their technique. We do not need transitioning, so we also set $t(n) = 0$. See Table IV.1b.

Constant bead count We can emulate the method from Ding et al. by dividing the feature radius over the widths of a constant number of beads. Additionally in order to emulate their definition of “branches” we mark all ST edges ($\alpha_{\max} = 0^\circ$) and we unmark the outer edges connected to the outline shape in a separate algorithm. Note that this deviation from the proposed framework violates the robustness against small perturbations in the outline polygon, since this marking depends on the topology of the graph of the ST. See Table IV.1c.

Centered We can emulate the method from Jin, Du, and He by transcribing how they deviate from the uniform width toolpaths. We therefore base the beading scheme on the bead count $q^-(d)$ defined by the uniform beading scheme. Jin, Du, and He replace two beads from the uniform toolpaths by a single one when the distance between the center of those beads falls short of $d_{\min} = 0.8w^*$, which gives us $w_{\max} = d_{\min} + 2 \cdot \frac{1}{2}w^* = 1.8w^*$. Conversely, they place an extra bead when the distance exceeds $d_{\max} = 1.25w^*$ [11], which gives us $w_{\min} = d_{\max} - 2 \cdot \frac{1}{2}w^* = 0.25w^*$ [11, p. 72]. We emulate the rounded polygonal path rerouting they define by supplying a transition length $t(n) = \frac{1}{2}w^*$ which results in a discretized version of their rounded polygon segment. See Table IV.1e.

Table IV.1: Beading schemes.

<p>(a) Uniform scheme</p> $q^-(d) = 2 \left\lfloor \frac{d}{2w^*} + \frac{1}{2} \right\rfloor$ $W(n, r)_i = w^* \text{ for all } i$	<p>(b) Outer bead</p> $q(d) = \begin{cases} 1 & \text{if } d < w^* \\ 2 & \text{otherwise} \end{cases}$ $W(n, r)_i = \begin{cases} 2r & \text{if } n = 1 \\ w^* & \text{otherwise} \end{cases}$
<p>(c) Constant bead count</p> $q(d) = C$ $W(n, r)_i = 2r/n \text{ for all } i$	<p>(d) Evenly distributed</p> $q(d) = \left\lfloor \frac{d}{w^*} + \frac{1}{2} \right\rfloor$ $W(n, r)_i = 2r/n \text{ for all } i$
<p>(e) Centered</p> $q(d) = q^-(d) + \begin{cases} -1 & \text{if } q^-(d)w^* - d > w^* - d_{\max} \\ 1 & \text{if } q^-(d)w^* - d < w^* - d_{\min} \\ 0 & \text{otherwise} \end{cases}$ $W(n, r)_i = \begin{cases} 2r - (n-1)w^* & \text{if } i = \frac{1}{2}(n-1) \\ w^* & \text{otherwise} \end{cases}$	
<p>(f) Inward distributed</p> $q(d) = \left\lfloor \frac{d}{w^*} + \frac{1}{2} \right\rfloor$ $W(n, r)_i = w^* + E(n, r) \frac{\omega(n, r)_i}{\sum_{j=0}^{n-1} \omega(n, r)_j} \text{ for all } i$ $E(n, r) = 2r - nw^*$ $\omega(n, r)_i = \max(0, 1 - N^{-2}(i - (n-1)/2)^2)$	

Evenly distributed By taking the advantages of the above two schemes we can define a beading scheme which constitutes a novel toolpathing technique. We can evenly divide the local feature diameter over the widths of all beads, but choose a local bead count better matching the local feature size. We determine the local bead count by dividing the diameter by the preferred bead width and rounding to the nearest integer. This reduces the demands on the system and deviation from mechanical properties caused by beads with extreme deviations from the preferred width. See Table IV.1d.

Inward distributed scheme The evenly distributed scheme can be conceptualized as calculating the total discrepancy E between the actual feature diameter d and the total preferred width nw^* , dividing the total discrepancy by the number of beads and setting the width of each bead to $w^* + E/n$. However, depending on the application we might want a different distribution of widths. We therefore supply a beading scheme which supports an arbitrary distribution of the discrepancy. The distribution is determined by some weighing function $\omega(n, r)$, which defines the portion of the discrepancy to distribute to each bead. See Table IV.1f. For example, we can choose an ω which distributes the discrepancy over the innermost N beads, and distribute most of it to the inner beads. See Fig. IV.13. That way we limit the region of impact of the distributed scheme to a central region and have the preferred bead width w^* in regions farther away. This limits the impact of transitioning regions so that transitions keep the toolpaths smooth farther away from the central regions.

Widening meta-scheme Complementary to any of these schemes we can enforce a minimum feature size and minimum bead width in our framework. Regions where the model is narrower than some r_{\min} can be printed with a bead width w_{\min} larger than the model

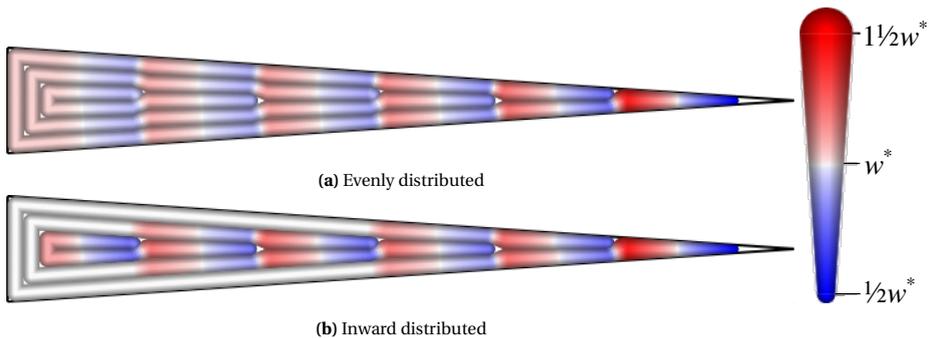


Fig. IV.13. Closeup of toolpaths generated with the distributed and inward ($N = 1.5$) beading schemes for a large wedge shape. Colors represent bead widths.

thickness. See Figs. IV.14a and IV.14b. We can simply override

$$q'(d) = \begin{cases} 0 & \text{if } 0 \leq d < 2r_{\min} \\ 1 & \text{if } 2r_{\min} \leq d < w^* \\ q(d) & \text{otherwise} \end{cases}$$

$$W'(n, r)_0 = \begin{cases} \max(w_{\min}, 2r) & \text{if } 2r < w^* \\ W(n, r)_0 & \text{otherwise} \end{cases}$$

Shell meta-scheme The industry standard of FDM is to generate only a limited contour-parallel perimeters and to fill the remainder using a direction-parallel strategy. We therefore provide a meta-scheme to generate adaptive bead width toolpaths only in narrow regions and generate the limited number of perimeters M using the preferred width in regions which are wide enough. We also take care not to leave gaps which are too small to be filled using the direction-parallel strategy:

$$q'(d) = \min(M, q(d))$$

$$W'(n, r)_i = \begin{cases} W(n, Mw^*)_i & \text{if } 2r > q^{-1}(M) \\ W(n, r)_i & \text{otherwise} \end{cases}$$

These meta-schemes introduce non-linearities in the quantization function. Because the beading is only evaluated at nodes in the skeleton, we need to make sure that there are nodes at the locations along the skeleton where the non-linearities happen. We therefore insert extra nodes along with their ribs at locations ν with a radial distance $R(\nu) = r_{\min}$ for widening and at $R(\nu) \in \{Mw^*, q^{-1}(M), q^{-1}(M) + 1/2w^*\}$ for the transition from narrow shell to unconstrained shell. Combining all meta-schemes functionality we can generate results such as depicted in Fig. IV.14c.

IV.5 FABRICATION

In order to accurately manufacture adaptive width toolpaths using an off-the-shelf 3D printing system, we need a model which relates the required width to process parameters such as movement speed and filament extrusion speed. A different approach might be appropriate depending on whether the filament feeder is mounted directly on the print head (a.k.a. *direct drive*) or the filament fed from the back of the printer to the print head via a *Bowden tube*. Because Bowden style 3D printing systems have the filament feeder relatively far away from the nozzle, changing the internal pressure in the system requires a large amount of filament movement, which requires a prohibitive amount of time.

IV.5.1. BACK PRESSURE COMPENSATION

Because changing the internal pressure is difficult in our setup, we keep the internal pressure constant, and vary the movement speed instead. One approach would be to keep the filament inflow f (in mm^3/s) constant by varying movement speed accordingly [42]. However, that does not result in the intended filament outflow variation – see Fig. IV.15a. We conjecture that the filament outflow is related to the total pressure in the system, which depends not only on the amount of filament in between the feeder wheel and the nozzle (which we keep constant), but also depends on the back pressure that the previous

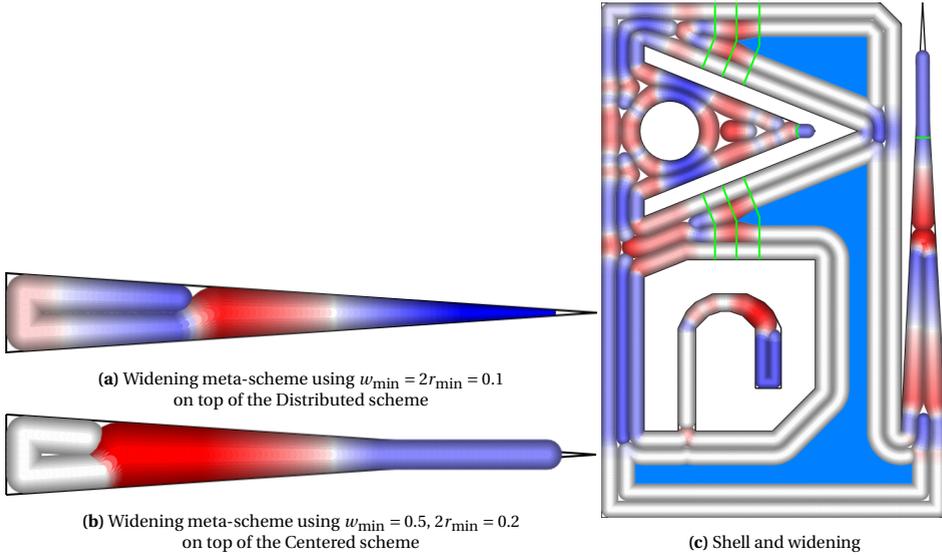


Fig. IV.14. Toolpaths using the widening and shell meta-schemes. (a) and (b) show widening. (c) show toolpaths generated with the inward distributed strategy ($N = 1.5$) in conjunction with the shell meta-scheme ($M = 4$) and the same widening as in (b). Widening and shell require extra edges (green) at key locations in the skeleton. The azure area is to be filled using some direction-parallel toolpaths.

layer exerts on the filament protruding from the nozzle. The amount of back pressure is most likely monotonically related to the requested line width. We compensate for the back pressure using a simple linear model:

$$v(w) = \frac{f(w)}{hw} \quad (\text{IV.8})$$

$$f(w) = f_0 - k(w/w_0 - 1) \quad (\text{IV.9})$$

where $v(w)$ is the movement speed as a function of requested bead width w , $f(w)$ is the filament outflow, f_0 is a constant reference flow, w_0 is a constant reference bead width and k is the amount of back pressure compensation.

Our back pressure compensation method effectively changes the speed to realize adaptive width, but this approach is limited, since the movement speed is constrained by acceleration considerations near bends in the toolpath [43]. Moreover, as the layer height is decreased the back pressure becomes larger compared to the internal pressure, which might cause the back pressure compensation method to demand prohibitively slow movement speeds. Furthermore, the shape and filling of the previous layer might influence the amount of back pressure. Accurate flow control can be further enhanced by using a direct drive hardware system and by employing *pressure advance algorithms* which dynamically change the internal pressure [44]. Conversely such a setup might benefit from some form of back pressure compensation as well.

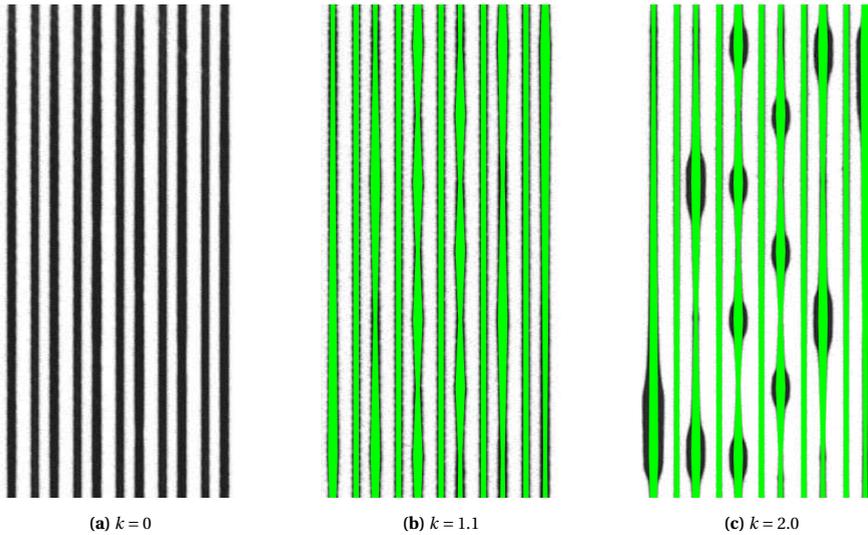


Fig. IV.15. Print results (black) of the varying width test on top of a dense white raft. Target widths in green. (a) Simple flow equalization without back pressure compensation results in nearly constant bead widths. (b) A value of $k = 1.1$ seems to produce good results.

IV.5.2. PRINT RESULTS

Using increments of 0.1 we established that using a factor of $k = 1.1$ yields satisfactory bead width variation for our setup where we use $f_0 = v_0 w_0 h$ with $v_0 = 30 \text{ mm/s}$, $w_0 = 0.4 \text{ mm}$ and $h = 0.1 \text{ mm}$. See Fig. IV.15b. The fact that the printed lines are wider than intended is compensated for using a flow reduction to 90%. Test prints were performed on an unmodified Ultimaker S5 system, with a standard 0.4 mm nozzle and PLA filament. The printing order is determined greedily by choosing the closest point of a polygonal extrusion path, or the closest of either end point in case of an open polyline extrusion path. Because the machine instructions file format *G-code* does not natively support adaptive width beads, we discretize adaptive width extrusions into 0.2 mm long segments of the average width. The print results can be viewed in Fig. IV.16.

In Fig. IV.16b the underfill problem of the naive uniform offset approach is most prevalent for the Ultimaker word mark, which negatively impacts the visual quality and the stiffness of the part. Moreover, in the case of the spatially graded honeycomb, there are several fully disconnected hexagons, which means the object falls apart when picked up. The honeycomb print is also missing all parts which are slightly more thin than the preferred bead width w^* . Figure IV.16c still shows some underfill, but considerably less than the uniform approach. These prints also exhibit dark regions where the translucency of the layer is less because the bead is higher. This can be explained by inaccuracies in the back pressure compensation method, which arise for bead widths which deviate from the preferred width by a large amount. Figure IV.16d diminishes the underfill nearly completely and the visual quality of these prints is more homogenous than those of the other methods. Moreover, the absence of dark regions signifies that our proposed method is more robust against inaccuracies in the deposition system. However, both the centered and inward distributed approach introduce transitions to a different bead count in the word ‘Delft’, which reduces the dimensional accuracy on the outline around those locations.

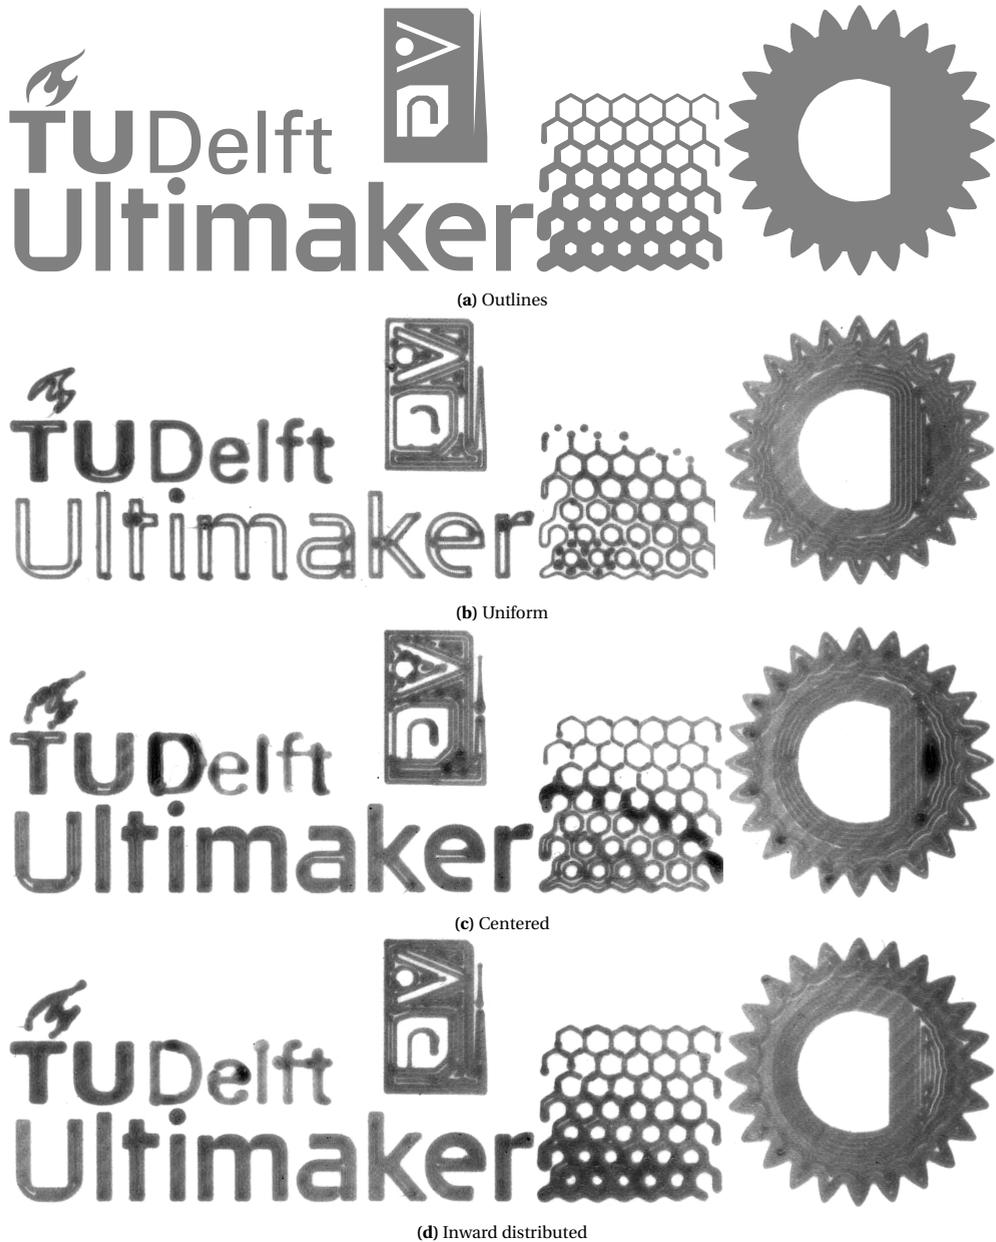


Fig. IV.16. Test shapes printed using the uniform scheme, centered scheme and the inward distributed scheme. The uniform technique produces distinct underfill areas. The centered scheme shows some defects due to inaccurate control of extreme deposition widths. The inward distributed scheme produces the least defects.

IV.6 RESULTS AND DISCUSSION

We evaluate the proposed framework and the various beading schemes on a set of different types of 3D models, ranging over various applications and various types of geometry. The data set is described in Appendix IV.B. We sliced all models in the data set and selected 300 random slices for analysis. Toolpaths of these 300 outline shapes are generated using the uniform technique as implemented by Clipper [45] – a state-of-the-art polygon offset library, and by our framework using four beading schemes, i.e. the constant bead count scheme with a bead count of $C = 4$, the centered, the evenly distributed, and the inward distributed beading scheme using $N = 2$, all with a preferred bead width of $w^* = 0.5$ mm and using the widening meta-scheme to enforce a minimum printed feature size of $w_{\min} = 2r_{\min} = 0.3$ mm. The tests were performed on a desktop PC equipped with an Intel Core i7-7500U CPU @ 2.70 GHz (a single core is used) and 16.3 GB memory. We report on the total statistics summed over the whole data set, because averaging would be biased.

IV.6.1. COMPUTATIONAL RESULTS

ACCURACY

We first evaluate the accuracy of different beading schemes in terms of the relative amount of the overfill and underfill. We construct the over- and underfill area by comparing the shapes covered by each extrusion move with each other and with the total shape of the boundary polygons. (For implementation details see Appendix IV.C.) This results in polygonal shapes such as visualized in the top half of Fig. IV.17: there are orange shapes where the beads overlap and azure shapes in the voids in between the beads. We compare the total area in mm^2 of these overfill and underfill shapes to the total area of the boundary for each sample in the data set and report the average percentages in Fig. IV.18a. The inward distributed scheme has a calculated overfill of 0.30 % and an underfill of 0.24 %. This is lower compared to the uniform scheme, which results in 1.63 % overfill and 1.62 % underfill in the data set.

UNIFORMITY

We visualize the bead widths resulting from the different schemes in the bottom of Fig. IV.17. We binned the toolpaths into width bins at 0.01 mm increments and determine the total toolpath length pertaining to each bin. From these statistics we calculate the mean and standard deviation and report them in Fig. IV.18b. We found that the mean width of the inward and evenly distributed schemes is close to the preferred bead width of 0.5 mm, while their standard deviation is lower than for the centered and constant bead count scheme. These results show that, while causing less overfill and underfill, inwards distributed and evenly distributed schemes deviate less or less often from the preferred bead width compared to the other schemes.

PRINT TIME

The total time it takes to print a part is influenced not only by our back pressure compensation scheme, but also by the geometry of the toolpaths. In order to separate these effects we report on the total print time when using back pressure compensation and when using a constant (maximum) movement speed in Fig. IV.18c. We estimate print times using a simulation of the Marlin firmware using the default movement settings of the setup described in Section IV.5.2. While the idealized print time is predominantly determined by

the total toolpath length, the print time using back pressure compensation is predominantly determined by the occurrence of wide beads, because they have a reduced the flow in mm^3/s . Because of acceleration constraints imposed by the hardware the maximum movement speed is not reached near sharp corners. We therefore also report on the angles of the bends in the toolpaths in Fig. IV.18d. Furthermore, the print time is negatively affected by discontinuities in the extrusion process. Between extrusions the printer needs to stop extrusion, travel to the next extrusion path and restart the extrusion process, which may introduce defects and incurs extra print time. For closed polygonal toolpaths we can start anywhere within the path, so we can optimize the starting location so as to minimize the travel time. We therefore report both on the open and closed path count in Fig. IV.18e.

COMPUTATIONAL PERFORMANCE

Fig. IV.18f plots the computation time against the vertex count of the layer for the full data set, comparing the uniform technique implemented using Clipper [45] to our framework with the inward distributed scheme. For polygonal shapes with as many as 10^4 vertices, the computation for both approaches is less than 1 second, with our method being approximately five times that of the uniform technique. These results could be improved upon by utilizing the locality inherent in our algorithms for parallelization on the GPU.

The computational complexity is limited by the generation of the Voronoi Diagram, which is $O(n \log n)$, where n is the number of vertices in the input shape. The other steps in our framework have a complexity of $O(m)$, where m is the number of elements in the ST. Therefore, the total running time of our algorithm is $O(n \log n)$. Results in Fig. IV.18f confirm that both our framework and the uniform technique have an expected running time of approximately $5 \times 10^{-6} n \log n$ seconds.

IV.6.2. COMPARISON OF BEADING SCHEMES

We can see from Fig. IV.17a(top) and IV.18a that the uniform technique causes a lot of overfills and underfills: on average 1.6% of the total target area is covered by underfill and likewise for overfill. To our knowledge, the uniform beading scheme, as well as the outer beading scheme, is of little use to FDM printers.

The constant bead count scheme effectively deals with underfills, but generates orders of magnitude more overfills compared to the other schemes. Also, the scheme comes at the cost of greatly varying bead widths and an average bead width that is not close to the preferred bead width. Note that most overfill areas occur near regions of alternating bead width. While the scheme results in short toolpaths, as indicated by the idealized print time, it also results in a wide range of bead widths, which cause the back pressure compensation print time to be very large. See Fig. IV.18. For an input outline shape which contains both very small and very large features, the constant bead count scheme produces bead widths which can fall outside of the range of manufacturable bead widths. Moreover the centrality marking is not robust against small perturbations in the outline; adding a small chamfer in a corner causes the unmarked ST to be very small at that location, which results in tiny bead widths. See top right of Fig. IV.17b.

In Fig. IV.17c we can see that the centered beading scheme effectively deals with overfill and produces desired bead widths in all locations, except for the extrusion paths in the center, where the bead widths range between $0.25w^*$ and $1.8w^*$. However, it does produce some narrow underfill regions. Compared to the uniform technique the centered technique increases the (open) path count, but considerably reduces over- and underfill

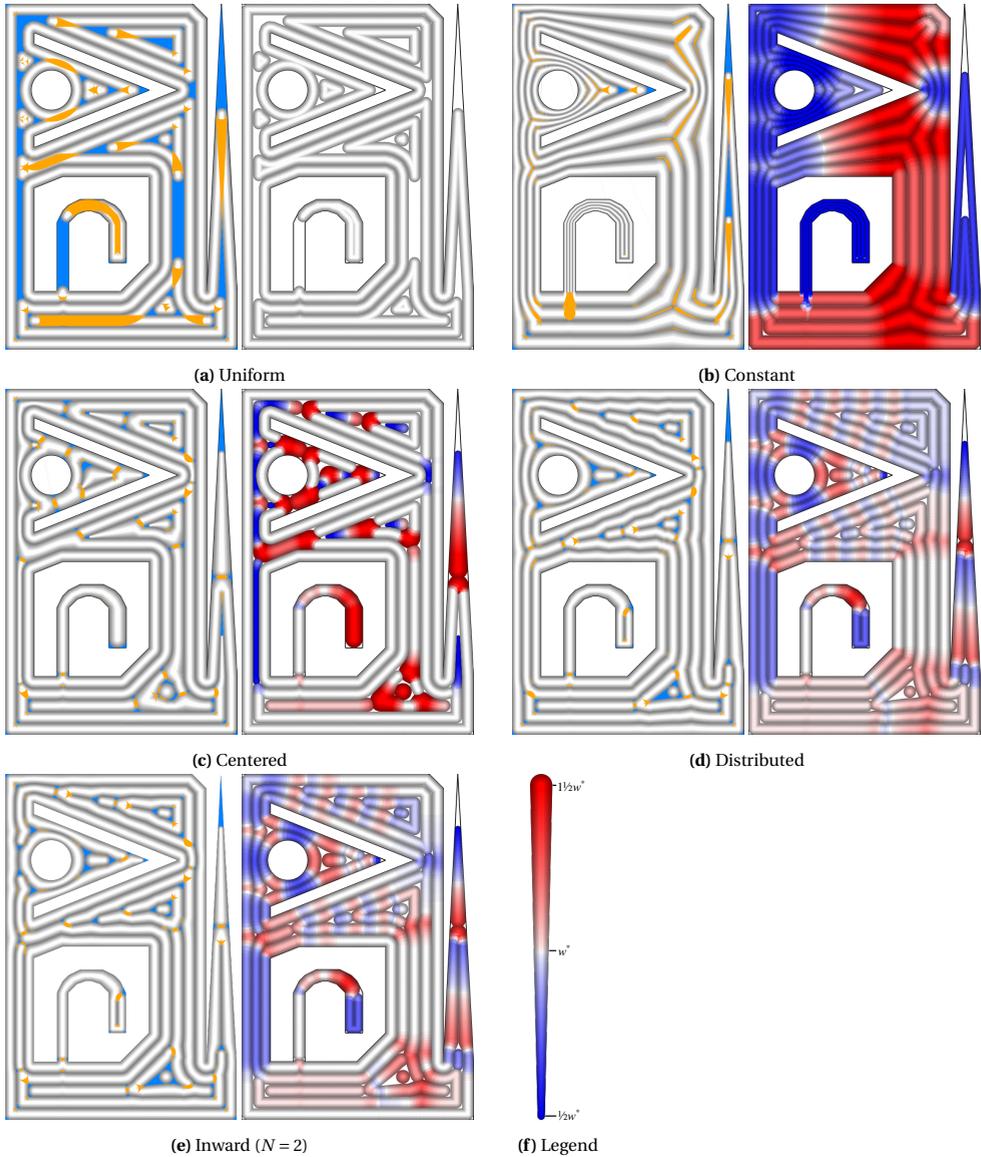


Fig. IV.17. Visualization of the overfills and underfills (top) and the widths (bottom) for various beading schemes. Extrusion beads in gray tones, overfill in orange, underfill in azure, narrow beads in blue and wide beads in red. In order to distinguish clearly from the Distributed scheme the Inward is limited to $N = 2$.

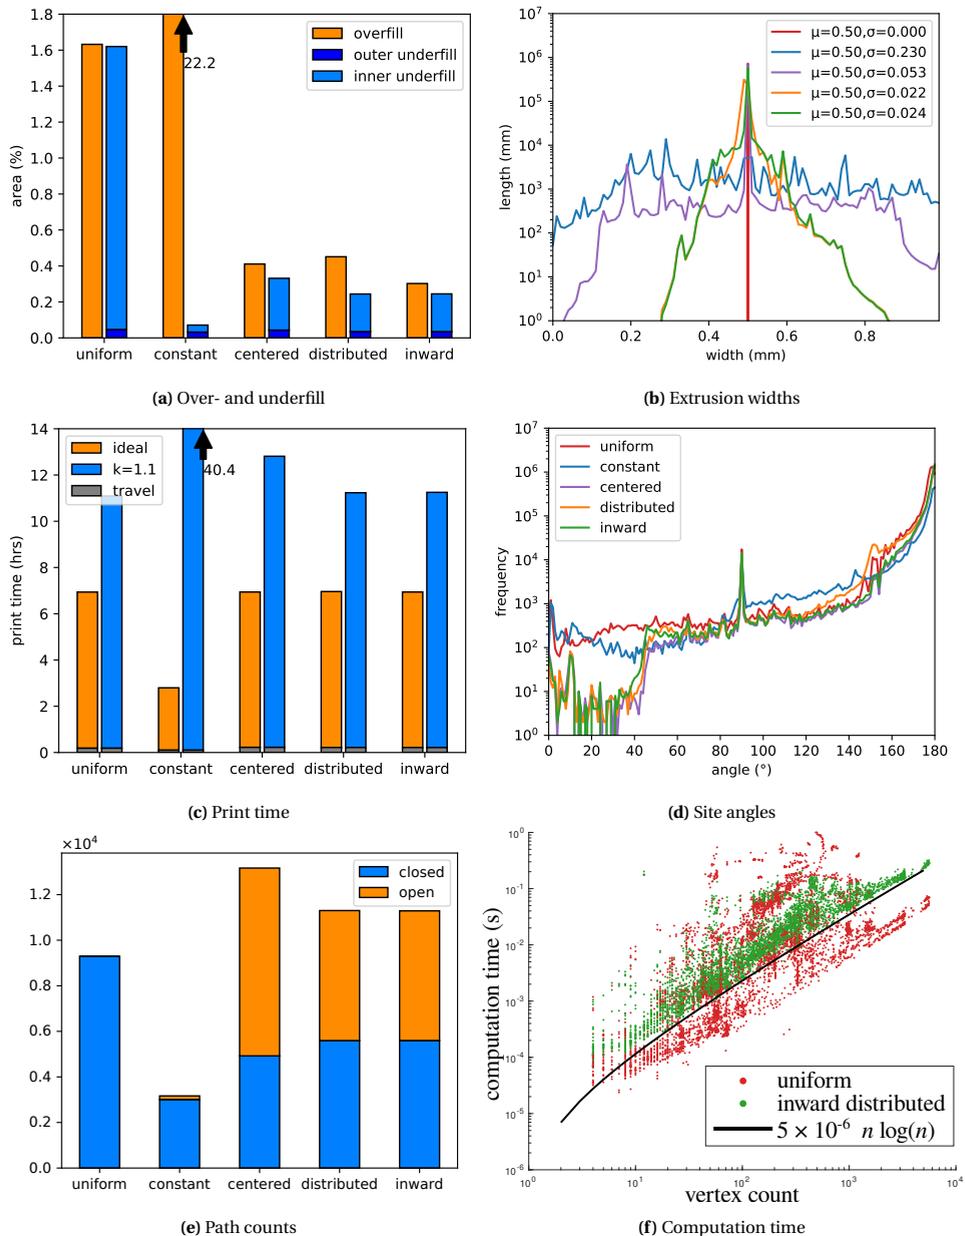


Fig. IV.18. Statistical analysis of the toolpaths from applying the uniform width technique and various beading schemes using our framework to a data set of 300 slices. Note the use of a logarithmic scale in the bottom graphs on the Y-axes and for (f) on the X-axes as well.

and decimates the number of toolpath angles below 45° . See Fig. IV.18.

However, according to Fig. IV.18b the centered scheme exhibits a wider range of bead widths than the distributed schemes: the standard deviation of the bead widths in the centered scheme is approximately $53\ \mu\text{m}$, while that of the distributed schemes is approximately $23\ \mu\text{m}$.¹ Moreover, because the quantization operator rounds to the nearest number of beads, in the worst case where we switch from a single to two beads the widths switch from $0.75w^*$ to $1.5w^*$, which is a considerably smaller range than in the centered scheme. We therefore conclude that the distributed schemes exhibit a lower bead width variation and lower (open) path count compared to the centered scheme.

Figures IV.13, IV.17d and IV.17e show that in the inward distributed scheme the outer toolpaths have the preferred width more often than in the evenly distributed scheme, which means that the outline accuracy of the inward distributed beading is less affected by inaccuracies in the adaptive width control system. Furthermore, we find that compared to evenly distributed, the inward distributed scheme produces less corners with angles above 130° and less overfill, because the area of influence that bead count transitions have is limited in the inward distributed scheme. Thus the inward distributed scheme prevents over- and underfill, generates smooth toolpaths with more homogeneous width and affects smaller more centered parts of the print than the other schemes, while incurring little to no extra print time.

IV.6.3. LIMITATIONS

Because the performance of the various toolpathing techniques depends on the geometry of a model, they have ramifications for the practice of design for additive manufacturing. Because the naive method produces under- or overfill for parts of an outline with a constant diameter $d \neq 2iw^*$ it is best practice to design a model such that horizontal cross-sections have a feature diameter of an even integer multiple i of the bead width. For the centered scheme and for the current state of the art one should only avoid parts for which $(2i + 1.8)w^* < d < (2i + 0.25)w^*$ in order to avoid underfill. For the distributed schemes however, there is no diameter at which the framework produces under- or overfill for a part with a constant diameter d . The design consideration therefore reduces to limiting the diameter of your parts to be within the range $[w_{\min}, \infty)$, where w_{\min} is a configurable parameter when using the widening meta-scheme.

The default limit bisector angle $\alpha_{\max} = 135^\circ$ ensures that we do not employ transitioning in shallow wedge regions, which would result in a lot of short odd single bead poly-lines, which would break up the semi-continuous nature of polygonal extrusion paths; $\alpha_{\max} = 135^\circ$ corresponds to $w^* / \cos 1/2\alpha_{\max} \approx 0.4\ \text{mm}$ long segments and under-/overfill areas of $1/4(w^*)^2 (\tan(\alpha/2) - \alpha/2) \approx 0.05\ \text{mm}^2$. However, future work might be aimed at reducing under-/overfill in regions with a low bisector angle without the introduction of short single polyline extrusion segments. If the over-/underfill problem is also solved for non-significant regions we might be able to increase α_{\max} and reduce the discontinuity introduced by short extrusion segments.

Another limitation of our method is that in a location v with locally maximal $R(v) \approx (i + 1/2)w^*$ the odd bead count will result in a single polyline extrusion segment consisting of only a single point. This can be viewed in the bottom right of Fig. IV.17e for example. In

¹Although the standard deviation σ of the inward distributed scheme is slightly higher than that of the evenly distributed scheme, the mean absolute deviation is lower (i.e. $9\ \mu\text{m}$ versus $11\ \mu\text{m}$), because its distribution is more peaked.

order to print such a dot, we make it into a $10\ \mu\text{m}$ long extrusion segment, with an altered width such that the total volume remains correct. A more principled way of dealing with such a situation remains future work.

Finally it should be noted that although our framework can accurately emulate the constant bead count approach by Ding et al., its emulation of the centered approach by Jin, Du, and He is imperfect. The transitions resulting from our framework introduce sharper corners and there is more width variation in those corners. Whereas the width of the connecting segment in the approach by Jin, Du, and He is the preferred width w^* , the bead widths closer to the center resulting from our framework will be twice the local radius, which is larger than w^* . However, this inflated bead width variation is expected to have an insignificant impact on the measured bead width variation.

IV.6.4. APPLICATIONS

Toolpaths with varying width is particularly meaningful for narrow parts, since there the negative effect of under- and overfill is more pronounced than in wide parts. In extreme cases, thin features will not be filled at all. Therefore, our framework, while working for wide parts as well, shows most of its potential for objects which contain thin parts.

Figure IV.19 collectively shows the application of the proposed inward distributed scheme for various types of 3D model, including both thin parts (architectural models, casings, embossed text, gears and microstructures) and wide parts (Fig. IV.19b) and organic shapes (Fig. IV.19c).

For architectural models and casings, preventing over- and underfill is expected to make them stronger. For embossed text, preventing underfill reduces the various holes in the top surfaces, which is detrimental to the visual quality of those top surfaces. For gears and similar mechanical parts that are designed with finite element analysis, the less variation in extrusion widths is closer to the assumptions under fast analysis (e.g. using homogenization [46]).

Of particular interest are microstructures that could be uniquely fabricated by 3D printing. For example, topology optimized bone-like structures [47] contain filaments of varying thickness that follow a varying stress direction (Fig. IV.19g). An angled Gyroid structure with uniform thickness also results in outline shapes with varying width (Fig. IV.19h). These structures are accurately densely filled using our framework. Another class of microstructures consists of parameterized patterns with varying thickness to achieve functional gradation. Figure IV.19i shows the contour-parallel toolpaths with varying width of a hexagonal grid neatly switches between different bead counts over the volume, preventing the jagged moves a direction-parallel toolpaths would create for such a case [3].

IV.7 CONCLUSION

In this paper we have introduced a framework for computing contour-parallel toolpaths employing adaptive bead width in order to minimize underfill and overfill areas. We introduced beading schemes which improve on the state of the art, and we have introduced a back pressure compensation method for accurate fabrication of adaptive width.

Our framework is flexible, demonstrated by the several beading schemes which emulate existing techniques. The computation times of our framework are on par with the state-of-the-art library for performing offsets of non-adaptive bead width. Our framework

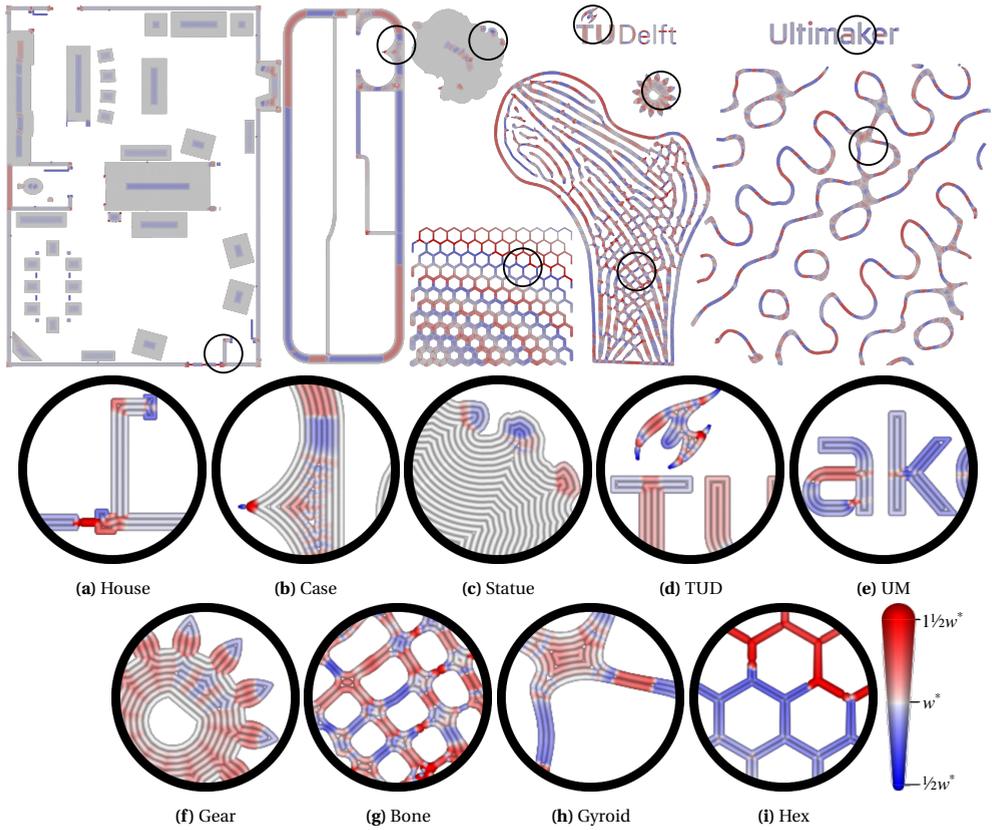


Fig. IV.19. Visualization of the widths for the output toolpaths of the inward distributed beading scheme ($N = 3$) applied to various example application objects. From left to right and top to bottom: a house, a case for electronics, a statue, two common logos, a gear, a topologically optimized bone structure, a tilted homogeneous gyroid structure and a heterogeneous thickness hexagonal grid.

is stable: small local changes in the outline shape cause only small changes in the toolpath.

Compared to the state of the art, the inward distributed beading scheme reduces the amount of beads with a width deviating extremely from the preferred bead width by changing the width of several beads near the center instead of only the center-most bead. It is therefore expected to limit the impact of varying the bead width in terms of production accuracy and homogeneity of material properties, which in turn is helpful to efficiently simulate an FDM manufactured part.

The proposed beading scheme greatly improves the process planning for parts with thin contours, which often occur for example in architectural models, prototypes for casings or microstructures. Meanwhile it leaves most of the toolpaths the same as the uniform width technique in large features, meaning that existing studies which relate process parameters with mechanical properties of the print are still applicable. Compared to the naive approach of constant width toolpaths our beading scheme is expected to improve the stiffness, dimensional accuracy and visual qualities of the manufactured model. It is expected that as distributed beading schemes are implemented in commercial software

packages and bead width variation control become commonplace, the practice of design for additive manufacturing can disregard some of the nozzle size considerations.

The presented framework is open source available at github.com/Ultimaker/libArachne

REFERENCES

- [1] Tim Kuipers et al. “A Framework for Adaptive Width Control of Dense Contour-Parallel Toolpaths in Fused Deposition Modeling”. In: *Computer-Aided Design* 128 (Nov. 2020), p. 102907. ISSN: 00104485. DOI: 10.1016/j.cad.2020.102907.
- [2] Tim Kuipers. “Method of determining toolpaths for an infill structure for a digital 3d model”. Pat. req. Ultimaker B.V. July 15, 2021. URL: <https://patentscope.wipo.int/search/en/detail.jsf?docId=W02021139928>.
- [3] Simon R.G. Bates, Ian R. Farrow, and Richard S. Trask. “Compressive behaviour of 3D printed thermoplastic polyurethane honeycombs with graded densities”. In: *Materials & Design* 162 (2018), pp. 130–142. ISSN: 02641275. DOI: 10.1016/j.matdes.2018.11.019.
- [4] Oraib Al-Ketan, Reza Rowshan, and Rashid K. Abu Al-Rub. “Topology-mechanical property relationship of 3D printed strut, skeletal, and sheet based periodic metallic cellular materials”. In: *Additive Manufacturing* 19 (Feb. 2018), pp. 167–183. ISSN: 22148604. DOI: 10.1016/j.addma.2017.12.006.
- [5] I Maskery et al. “Insights into the mechanical properties of several triply periodic minimal surface lattice structures made by polymer additive manufacturing”. In: *Polymer* 152 (2018), pp. 62–71. ISSN: 0032-3861. DOI: 10.1016/j.polymer.2017.11.049.
- [6] Tomás Zegard and Glaucio H. Paulino. “Bridging topology optimization and additive manufacturing”. In: *Structural and Multidisciplinary Optimization* 53.1 (Jan. 2016), pp. 175–192. ISSN: 1615-1488. DOI: 10.1007/s00158-015-1274-4.
- [7] Jun Wu, Weiming Wang, and Xifeng Gao. “Design and Optimization of Conforming Lattice Structures”. In: *IEEE Transactions on Visualization and Computer Graphics* (2019), pp. 1–14. arXiv: 1905.02902. URL: <http://arxiv.org/abs/1905.02902>.
- [8] Lin Cheng, Jiayi Bai, and Albert C. To. “Functionally graded lattice structure topology optimization for the design of additive manufactured components with stress constraints”. In: *Computer Methods in Applied Mechanics and Engineering* 344 (2019), pp. 334–359. ISSN: 00457825. DOI: 10.1016/j.cma.2018.10.010.
- [9] Donghong Ding et al. “Adaptive path planning for wire-feed additive manufacturing using medial axis transformation”. In: *Journal of Cleaner Production* 133 (2016), pp. 942–952. ISSN: 09596526. DOI: 10.1016/j.jclepro.2016.06.036.
- [10] Yi Xiong et al. “Process planning for adaptive contour parallel toolpath in additive manufacturing with variable bead width”. In: *The International Journal of Advanced Manufacturing Technology* (June 2019). ISSN: 14333015. DOI: 10.1007/s00170-019-03954-1.
- [11] Yuan Jin, Jianke Du, and Yong He. “Optimization of process planning for reducing material consumption in additive manufacturing”. In: *Journal of Manufacturing Systems* 44 (2017), pp. 65–78. ISSN: 02786125. DOI: 10.1016/j.jmsy.2017.05.003.

- [12] Ultimaker. *Ultimaker Cura 4.2.1 software*. Version 4.2.1. Aug. 31, 2019. DOI: 10.5281/zenodo.1216744.
- [13] Marco Livesu et al. "From 3D models to 3D prints: an overview of the processing pipeline". In: *Computer Graphics Forum* 36.2 (May 2017), pp. 537–564. ISSN: 14678659. DOI: 10.1111/cgf.13147. arXiv: 1705.03811.
- [14] Brian N. Turner, Robert Strong, and Scott A. Gold. "A review of melt extrusion additive manufacturing processes: I. Process design and modeling". In: *Rapid Prototyping Journal* 20.3 (2014), pp. 192–204. ISSN: 1355-2546. DOI: 10.1108/RPJ-01-2013-0012.
- [15] Sung Hoon Ahn et al. "Anisotropic material properties of fused deposition modeling ABS". In: *Rapid Prototyping Journal* 8.4 (2002), pp. 248–257. ISSN: 13552546. DOI: 10.1108/13552540210441166.
- [16] Tim Kuipers, Jun Wu, and Charlie C.L. Wang. "CrossFill: Foam Structures with Graded Density for Continuous Material Extrusion". In: *Computer-Aided Design* 114 (Sept. 2019), pp. 37–50. ISSN: 00104485. DOI: 10.1016/j.cad.2019.05.003.
- [17] Wenbiao Han et al. "Tool Path-Based Deposition Planning in Fused Deposition Processes". In: *Journal of Manufacturing Science and Engineering* 124.2 (Apr. 2002), pp. 462–472. DOI: 10.1115/1.1455026. eprint: https://asmedigitalcollection.asme.org/manufacturingscience/article-pdf/124/2/462/4691900/462_1.pdf.
- [18] John C. Steuben, Athanasios P. Iliopoulos, and John G. Michopoulos. "Implicit slicing for functionally tailored additive manufacturing". In: *Computer-Aided Design* 77 (Aug. 2016), pp. 107–119. ISSN: 00104485. DOI: 10.1016/j.cad.2016.04.003.
- [19] Haisen Zhao et al. "Connected fermat spirals for layered fabrication". In: *ACM Transactions on Graphics* 35.4 (2016), pp. 1–10. ISSN: 07300301. DOI: 10.1145/2897824.2925958.
- [20] Yuan Jin et al. "A non-retraction path planning approach for extrusion-based additive manufacturing". In: *Robotics and Computer-Integrated Manufacturing* 48. December 2016 (2017), pp. 132–144. ISSN: 07365845. DOI: 10.1016/j.rcim.2017.03.008.
- [21] Martin Held and Christian Spielberger. "A smooth spiral tool path for high speed machining of 2D pockets". In: *Computer-Aided Design* 41.7 (2009), pp. 539–550. ISSN: 00104485. DOI: 10.1016/j.cad.2009.04.002.
- [22] Nuodi Huang, Roby Lynn, and Thomas Kurfess. "Aggressive Spiral Toolpaths for Pocket Machining Based on Medial Axis Transformation". In: *Journal of Manufacturing Science and Engineering* 139.5 (2017), p. 051011. ISSN: 1087-1357. DOI: 10.1115/1.4035720.
- [23] Sara McMains et al. "Layered manufacturing of thin-walled parts". In: *ASME Design Engineering Technical Conference, Baltimore, Maryland*. Citeseer. 2000.
- [24] Yu An Jin, Yong He, and Jian Zhong Fu. "An Adaptive Tool Path Generation for Fused Deposition Modeling". In: *Advanced Materials Research*. Vol. 819. Trans Tech Publ. Sept. 2013, pp. 7–12. DOI: 10.4028/www.scientific.net/amr.819.7.

- [25] Donghong Ding et al. "A tool-path generation strategy for wire and arc additive manufacturing". In: *International Journal of Advanced Manufacturing Technology* 73.1-4 (2014), pp. 173–183. ISSN: 14333015. DOI: 10.1007/s00170-014-5808-5.
- [26] Jordan J Cox et al. "Space-filling curves in tool-path applications". In: *Computer-Aided Design* 26.3 (1994), pp. 215–224. ISSN: 0010-4485. DOI: [https://doi.org/10.1016/0010-4485\(94\)90044-2](https://doi.org/10.1016/0010-4485(94)90044-2).
- [27] J. G. Griffiths. "Toolpath based on Hilbert's curve". In: *Computer-Aided Design* 26.11 (1994), pp. 839–844. ISSN: 00104485. DOI: 10.1016/0010-4485(94)90098-1.
- [28] Saquib Shaikh et al. "Hilbert curve based toolpath for FDM process". In: *CAD/CAM, Robotics and Factories of the Future*. Ed. by CS Mandal, DK and Syan. Lecture Notes in Mechanical Engineering. Springer, 2016, pp. 751–759. DOI: 10.1007/978-81-322-2740-3_72.
- [29] Ju-hsien Kao and Fritz B Prinz. "Optimal motion planning for deposition in layered manufacturing". In: *Proceedings of DETC*. Vol. 98. Citeseer, 1998, pp. 13–16.
- [30] Yuan Jin, Yong He, and Jianke Du. "A novel path planning methodology for extrusion-based additive manufacturing of thin-walled parts". In: *International Journal of Computer Integrated Manufacturing* 30.12 (2017), pp. 1301–1315. ISSN: 13623052. DOI: 10.1080/0951192X.2017.1307526.
- [31] Maarten Moesen et al. "Robust beam compensation for laser-based additive manufacturing". In: *Computer-Aided Design* 43.8 (2011), pp. 876–888. ISSN: 00104485. DOI: 10.1016/j.cad.2011.03.004.
- [32] Morad Behandish, Amir M. Mirzendehtdel, and Saigopal Nelaturi. "A Classification of Topological Discrepancies in Additive Manufacturing". In: *Computer-Aided Design* (2019). ISSN: 0010-4485. DOI: 10.1016/j.cad.2019.05.032. arXiv: 1904.13210.
- [33] Kunnayut Eiamsa-ard et al. "Toward automatic process planning of a multi-axis hybrid laser aided manufacturing system: skeleton-based offset edge generation". In: *ASME 2003 International Design Engineering Technical Conferences and Computers and Information in Engineering Conference*. American Society of Mechanical Engineers Digital Collection, 2003, pp. 227–235. DOI: 10.1115/DETC2003/DAC-48726. eprint: https://asmedigitalcollection.asme.org/IDETC-CIE/proceedings-pdf/IDETC-CIE2003/37009/227/2594325/227_1.pdf.
- [34] Harry Blum et al. "A transformation for extracting new descriptors of shape". In: *Models for the perception of speech and visual form* 19.5 (1967), pp. 362–380.
- [35] Der-Tsai T. Lee. "Medial Axis Transformation of a Planar Shape". In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* PAMI-4.4 (1982), pp. 363–369. ISSN: 01628828. DOI: 10.1109/TPAMI.1982.4767267.
- [36] B. Chazelle and J. Incerpi. "Triangulation and shape-complexity". In: *ACM Transactions on Graphics* 3.2 (1984), pp. 135–152. ISSN: 07300301. DOI: 10.1145/357337.357340.
- [37] A. Fournier and D. Y. Montuno. "Triangulating Simple Polygons and Equivalent Problems". In: *ACM Transactions on Graphics* 3.2 (1984), pp. 153–174. ISSN: 07300301. DOI: 10.1145/357337.357341.

- [38] Boris Schäling. *The boost C++ libraries*. Boris Schäling, 2011.
- [39] Steven Fortune. “A sweepline algorithm for Voronoi diagrams”. In: *Algorithmica*. SCG ’86 2.1 (Nov. 1987), p. 153. DOI: 10.1007/BF01840357.
- [40] Dominique Attali and Annick Montanvert. “Modeling noise for a better simplification of skeletons”. In: *Proceedings of 3rd IEEE International Conference on Image Processing*. Vol. 3. IEEE, 1996, pp. 13–16. ISBN: 078033258X. DOI: 10.1109/ICIP.1996.560357.
- [41] Avneesh Sud, Mark Foskey, and Dinesh Manocha. “Homotopy-Preserving Medial Axis Simplification”. In: *Proceedings of the 2005 ACM Symposium on Solid and Physical Modeling*. Cambridge, Massachusetts: Association for Computing Machinery, 2005, pp. 39–50. DOI: 10.1145/1060244.1060250.
- [42] T. Kuipers et al. “Hatching for 3D prints: Line-based halftoning for dual extrusion fused deposition modeling”. In: *Computers & Graphics* 74 (2018), pp. 23–32. ISSN: 00978493. DOI: 10.1016/j.cag.2018.04.006.
- [43] Deniz Sera Ertay, Alexander Yuen, and Yusuf Altintas. “Synchronized material deposition rate control with path velocity on fused filament fabrication machines”. In: *Additive Manufacturing* 19 (2018), pp. 205–213. ISSN: 22148604. DOI: 10.1016/j.addma.2017.05.011.
- [44] Sigmund Arntsonn Tronvoll et al. “Investigating pressure advance algorithms for filament-based melt extrusion additive manufacturing: theory, practice and simulations”. In: *Rapid Prototyping Journal* 25.5 (Jan. 2019), pp. 830–839. ISSN: 13552546. DOI: 10.1108/RPJ-10-2018-0275.
- [45] Angus Johnson. *Clipper 6.4.2 - an open source freeware library for clipping and off-setting lines and polygons*. Version 6.4.2. 2017.
- [46] Xingchen Liu and Vadim Shapiro. “Homogenization of material properties in additively manufactured structures”. In: *Computer-Aided Design* 78 (2016), pp. 71–82. DOI: <https://doi.org/10.1016/j.cad.2016.05.017>.
- [47] J. Wu et al. “Infill Optimization for Additive Manufacturing – Approaching Bone-Like Porous Structures”. In: *IEEE Transactions on Visualization and Computer Graphics* 24.2 (Feb. 2018), pp. 1127–1140. DOI: 10.1109/TVCG.2017.2655523.

IV.A EDGE DISCRETIZATION

We calculate the location l of the boundary between a significant and nonsignificant portion of an edge analytically. For example, the parabolic MAT edge generated from the outline vertex $(0, 1)$ and an outline segment aligned with the X-axis follows $y(x) = 1/2x^2$ and $R(x) = y(x)$. We can determine the significant portion $[-x_{\text{bound}}, x_{\text{bound}}]$ by evaluating $\frac{\partial R}{\partial x} > \cos(\alpha_{\text{max}}/2)$, which is $|x_{\text{bound}}| = (\tan(\alpha_{\text{max}}/2))^{-1}$. Similarly, a MAT edge generated from two vertices at $(0, 0)$ and $(0, 1)$ follows $y(x) = 1/2$ and $R(x) = \sqrt{1/4 + x^2}$. The boundaries of significance are given by $|x_{\text{bound}}| = 1/2(\tan(\alpha_{\text{max}}/2))^{-1}$. From these we can derive the locations $l = (\pm x_{\text{bound}}, y(x_{\text{bound}}))$. These specific cases can easily be transformed into all possible cases using scaling and rotation operations.

Table IV.2: 3D models used for validation

Model Name	Author
AirCasting Air Monitor Casing	HabitatMap
Air hose splitter	frizinko
Al Hamra Tower	TurnerConstructionCompany
canon NP-E3 battery cap	kosuyoung
David	Thunk3D
Deck Assembly Tool	PSomeone
Ender 3 Cable Chain	johnniewhiskey
Ergonomic Hacksaw Handle	mmOne
Gap measurement tool	ravm84
G-Clamp fully printable	johann517
Gyroid	Tim Kuipers
3D Printable Jet Engine	CATIAV5FTW
Lawn Mower Throttle Replacement	Spammington
OpenRC F1 Internal gear box mod	intoxikated
PCB Test Fixture	JMadison
Pioneer Radio Holder for Ford Focus	Perugino
Replicator Dual Fan Mount	aubenc
Atuador versão 2 * actuator version 2	Caroline Holanda
TE Pocket Operator Hard case	Salvation76
Screw sizer	Pierrolalune63
Bone-like optimized infill	Jun Wu
Two-Story Spec House	pwc-phil

IV.B DATA SET

The data set we tested on was a custom selected set of open source 3D models found on the internet which was selected to cover a broad range of different types of application and geometry. Applications range from prototypes, to fixtures and mechanical end-use parts. The geometry covers a wide range including thin filaments, smooth surfaces, organic shapes, chamfered shapes, small shapes and large shapes. The models are described in Table IV.2.

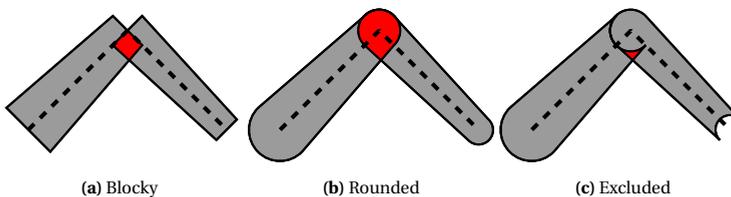


Fig. IV.20. Extruded area of two extrusion segments. Red areas signify doubly extruded areas.

IV.C ACCURACY CALCULATION

In order to estimate the overfill and underfill, we need to accurately calculate the area covered by a single extrusion path. If we would simply use an isosceles trapezoidal area, we would get overfill artifacts at corners in the toolpath (Fig. IV.20a). We therefore use a semi-circle (Fig. IV.20b) with a diameter equal to the starting width in the one end of each segment, and exclude it at the other end, because it will be included in the next segment. For polyline extrusion paths which are not closed, we also include the semi-circle of the destination location (Fig. IV.20c).

Using boolean operations we can obtain the polygonal regions for overfill and those for underfill. In order to deal with rounding errors we perform a morphological close of $5\ \mu\text{m}$, before calculating the total area in mm^2 . We also calculate regions which are covered thrice by different extrusion segments and add twice its area to the total overfill area amount.

V

ENABLING DUAL-MATERIAL STRUCTURES

In the previous chapter we presented a technique which makes it possible to reliably print geometry of arbitrary thickness, which unlocks a class of FGMs which have a spatially graded thickness. Now that a large body of FGMs is made available through the elimination of a manufacturing constraint, they fall beyond the scope of this research, which is limited to toolpath generation techniques. We therefore shift our focus to multi-material FGMs. Using a commercially available dual material FDM system we can manufacture structures which have a spatial gradation of properties in between the material properties of the base materials. However, a large subset of multi-material FGMs would violate the chemical compatibility manufacturing constraint. Two materials are said to be incompatible if their chemical makeup does not allow for a sufficient bonding between the two materials. Two bodies that do not connect to each other can easily break apart after or during printing, in which case a print failure would ensue.

This chapter proposes a microstructure which can be applied at the interface between two bodies, in order to achieve a mechanical type of bonding between two materials. By defining the geometry of the structure such that the two materials are interlocking topologically, the bodies cannot separate by deformation of either of the materials alone – either material would have to break. We then optimize the geometry such that the base materials would break at the highest applied force. Eliminating the chemical compatibility constraint allows for reliable manufacturing of multi-material FGMs.

Section V.1 introduces the technique and consider related work. Section V.2 presents two variants of the interlocking structure and analytical models to optimize them. The mechanical performance of the structures is validated using both simulation software and empirically in Section V.3. Section V.4 discusses the validity of the analytical models and the performance of various interlocking structures, which leads to the conclusions presented in Section V.5.

This chapter is published as Tim Kuipers, Renbo Su, Jun Wu, and Charlie C.L. Wang. “TTIL: Interlaced Topologically Interlocking Lattice for continuous dual-material extrusion”. In: *Additive Manufacturing* 50 (2022), p. 102495. ISSN: 2214-8604. DOI: <https://doi.org/10.1016/j.addma.2021.102495>; some minor corrections have been made.

ABSTRACT

Material Extrusion (MEX) systems with dual-material capability can unlock interesting applications where flexible and rigid materials are combined. When chemically incompatible materials are concerned the adhesion between the two might be insufficient. Therefore researchers typically rely on dovetail type interlocking geometries in order to affix two bodies mechanically. However, dovetail type interlocking introduces extrusion discontinuities and relies on the material's resistance to deformation, which is difficult to model.

We propose a simple and effective 3D lattice consisting of interlaced horizontal beams in vertically alternating directions which interlock topologically: the interlaced topologically interlocking lattice (ITIL). It ensures continuous extrusion and ensures an interlock even for highly flexible materials. We develop analytical models for optimizing the ultimate tensile strength of the ITIL lattice in two different orientations relative to the interface: straight and diagonal. The analytical models are applied to polypropylene (PP) and polylactic acid (PLA) and verified by finite elements method (FEM) simulations and physical tensile experiments. In the diagonal orientation ITIL can obtain 82% of the theoretical upper bound of 8.6MPa. ITIL seems to perform comparably to dovetail interlocking designs, while it lends itself to application to non-vertical interfaces. Optimizing the lattice for non-vertical interfaces, however, remains future work.

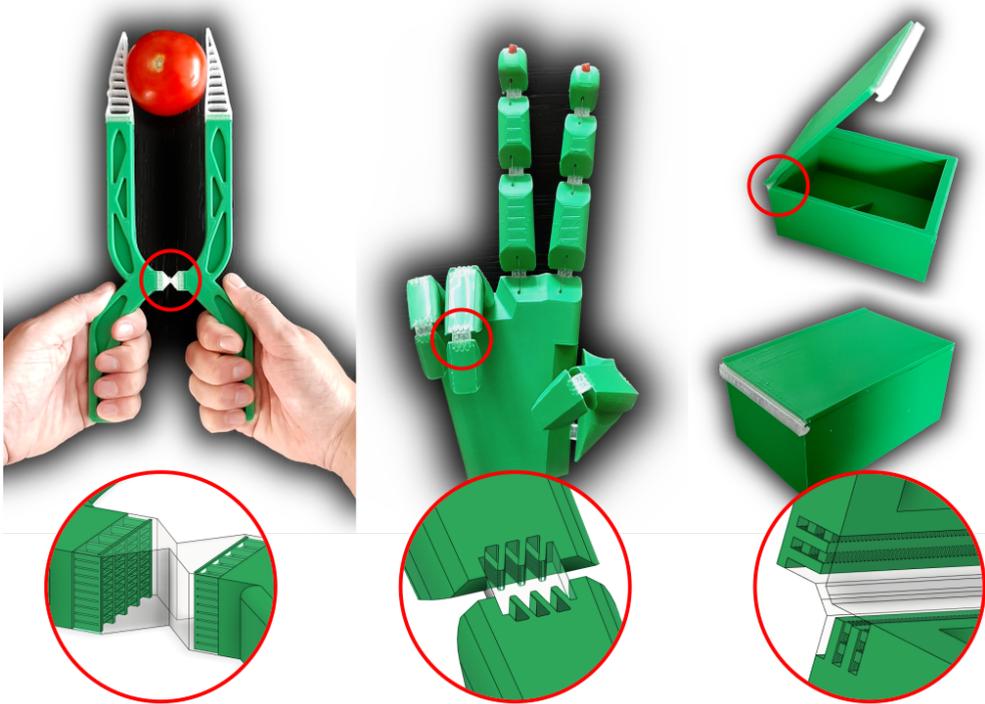
V

V.1 INTRODUCTION

Multi-material extrusion 3D printers unlock a plethora of applications through combining the unique material properties of various materials. However, depending on the combination of materials the adhesion between the materials can be excessively weak. While only a small number of all possible 3D printing material combinations may exhibit any incompatibility issues, it is often precisely those incompatible combinations where the different chemical make-up produces interesting applications. For example, polypropylene (PP) is semi-flexible and fatigue-resistant, but has a very weak chemical bond to typical rigid filaments such as polylactic acid (PLA). PP is often used for living hinges, which consist of a single part rather than moving components. Combining these materials unlocks compliant mechanism applications such as visualized in Fig. V.1. In such cases it is necessary to rely on mechanical interlocking to prevent the materials from breaking apart from each other.

Dovetail interlocking is a common strategy to affix two bodies together; one example can be found in jigsaw puzzles. The pieces interlock and stay connected because of the material stiffness. However, the interlock can be broken in plane by deformation of the pieces, and if the dovetail is not accompanied by horizontal features above and below the pieces could be disassembled by translation alone. See Fig. V.2a. Moreover, because the dovetails widen toward their tip, they cannot easily be printed with continuous extrusion toolpaths.

Therefore, we propose *topological interlocking*, which is a type of interlocking where the interlock is preserved under continuous deformations such as stretching, twisting or bending of any magnitude – that is: the bodies can only be unlocked by discontinuous deformations such as fracture. In order to achieve topological interlocking both materials have a high genus topology: the holes or tunnels in one material are filled with the other material and vice versa, similar to how the rings of chain mail are linked together. A topological interlock remains effective under any deformation of the base materials and can



V

Fig. V.1. Applications of dual material products using the ITIL lattice for interlocking Ultimaker transparent PP and Ultimaker green tough PLA material. A gripper making use of the straight ITIL variant. A prosthetic hand using the diagonal ITIL variant. A storage box utilizing the straight ITIL variant.

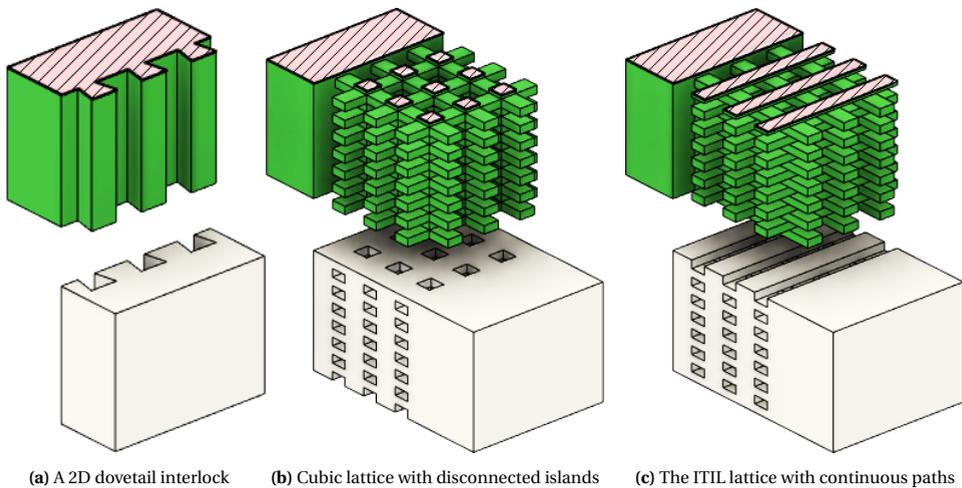


Fig. V.2. Interlocking principles; exploded view with a section cut at the top layer. The dovetail can be disassembled by translation; the cubic lattice causes highly discontinuous extrusion on some layers; ITIL solves these problems.

only be broken by failure in either material. This interlocking principle is robust especially when flexible and deformable materials are concerned.

However, most topologically interlocking geometries would introduce discontinuities in the extrusion process when sliced into layers for 3D printing, as each slice would contain disconnected islands. See Fig. V.2b. Such discontinuities can easily lead to defects, which influence the dimensional accuracy and the mechanical properties of the resulting part. We therefore have to generate interlocking geometry for which the layers consist of long continuously connected areas for both materials.

It seems impossible to generate topologically interlocking geometry with holes or tunnels while enforcing continuous extrusion for both materials; if the one material leaves a hole in a layer then filling that hole with the other material will cause it to be disconnected from the other regions of the second material. The ITIL lattice consists of long horizontal beams which ensure continuous extrusion, which alternate in orientation along the Z axis in order to connect all beams together, thereby creating linked hoops which constitute the topological interlock. Where the beams of the two orientations meet they form long vertical pillars. However, the small island cross-sections of these pillars are all located in between consecutive layers, so that they never result in disconnected islands. See Fig. V.2c.

V

V.1.1. RELATED WORK

Multi-material additive manufacturing Multi-material additive manufacturing has unlocked a plethora of applications, by making use of functionally graded material properties, tailored composite materials or multi-material designs. Several review papers cover a wide range of techniques on these topics [2, 3]. Using multiple materials to create colored surface imagery is commonly performed using MultiJet technologies, but can alternatively be performed using inkjet techniques [4] or even DLP resin printers [5]. Such techniques can even be extended to deal with translucency [6] and gloss [7]. Several color printing techniques have also been proposed for MEX [8, 9, 10].

Besides visual attributes, multi-material systems can be used to create parts with graded material properties, by generating composite structures with varying densities of soft and hard materials [11]. Fine-tuning the small-scale geometry in which such materials are deposited can give a more sophisticated control over the induced material properties [12, 13], and adjusting the small-scale geometry throughout the product on a meso scale can increase the performance of the product even more [14].

Some MEX systems for multi-material extrusion have been proposed which operate by extruding multiple materials out of a single nozzle, e.g. by routing multiple filaments into a single mixing nozzle [15] or by creating a single strand of multi-material filament [16, 17], but such systems can exhibit hardware issues when the different materials require vastly different processing parameters. Therefore the more upscale multi-material extrusion systems use a separate nozzle for each material [18].

Lattices Lattices such as beam lattices, triply periodic surfaces and ordered dithering have widely been studied for their mechanical properties. Several review papers provide a comprehensive overview of lattices, which are also known as ‘microstructures’, ‘meso-scale structures’, ‘cellular materials’, etc. [19, 20, 21]. Single material lattices can constitute auxetic materials, or light weight structures with tailored material properties, through fine adjustments of the lattice geometry. Multi-material lattices inherently exhibit topological interlocking, which makes them good candidates for interlocking [22]. However,

optimizing lattices for adhesion between incompatible materials while adhering to MEX manufacturing constraints has been studied scarcely.

Adhesion Important factors for adhesion between polymers are entanglement and dissipation [23]. The adhesion between layers produced within a body of a single material produced by MEX can be influenced by various process parameters [24, 25], as well as the geometry of the toolpaths [26, 27, 28]. The adhesion by which two bodies of different MEX printed materials stick together can be influenced by a wide spectrum of pre-treatment methods, process parameters and material properties [22, 29]. Increasing surface roughness might improve the adhesion between materials [30, 31], but this supposed benefit is contested [23]. For MEX one could try mixing the materials by overlapping their toolpaths to increase adhesion or create simple straight protrusions in order to increase the friction between the two materials [32]. However, if such protrusions bulge outward the adhesion does not merely increase because of the increased friction, but also because it would constitute a dovetail type of interlocking structure.

Interlocking An interesting type of interlocking can be found in interlocking assemblies where “locking of an individual element within the assemblage is furnished by the kinematic constraints provided by its neighbors by virtue of the element geometry and the mutual arrangements of the elements within the structure.” [33] Although this concept has also been referred to as ‘topological interlocking’ in literature, it is rather different from the type of topological interlocking proposed in this paper. This paper pertains to bodies of different materials, and the interlocking between these materials is preserved under continuous deformations, i.e., *topologically*. While interlocking assemblies lock orthogonal movement when the elements are globally constrained transversely, the interlock is nullified when no such global constraint is present. This limits the application area of interlocking assemblies.

Textiles exhibit a different type of interlocking, where individual strands are woven together into what can be viewed as a lattice of knots. Knot theory is a vast area of research within the mathematical discipline of topology. The mechanical properties of textiles are also thoroughly studied [34]. The concept of weaving can be exploited in MEX: weaving extruded strands together across several layers can improve layer adhesion within a part of a single material [35]. However, because the strands are extruded from the top, no topological knots can be achieved in MEX. It is not possible for MEX to weave strands into an interlocking textile lattice.

Literature on interlocking patterns for adhesion between incompatible materials in MEX 3D printing has focused on extended 2D interlocking dovetail type of designs, such as jigsaw shapes [36], trapezoidal sutures [37], T-shapes [38, 39] and star shapes [40] in the horizontal direction as well as in the vertical direction across several layers [28]. Topology optimization can generate complex 2D dovetail interlocking shapes, which fit to the specifics of the design locally [41]. However, the dovetail shapes are often relatively large, which limits their applicability and if they are shrunk the widening of the dovetails would introduce discontinuous toolpaths. If such interlocking designs are considered on their own without horizontal features above and below they would allow for disassembly in the vertical direction. By considering aspects of the interlock along the Z axis, the dovetail interlocking concept can be expanded into 3D interlocking structures.

The jigsaw idea can be expanded into a 3D interlocking structure by protruding not

only sideways, but also vertically, resulting in a shape resembling a tree [42]. Similarly the T-shaped interlocking design with horizontal bars could be expanded with bars in the vertical directions. However, such designs violate the semi-continuous extrusion requirement because the layers above and below the base of the T would contain separated islands of one material, which are difficult to print accurately when the two materials do not adhere to each other.

These issues can be addressed by generating a repeating lattice where the interlocks are connected together; simple straight I shaped extrusions can be linked together by cross beams in order to form a topologically interlocking structure resembling a ladder [43]. Such a topologically interlocking structure satisfies the continuity constraint only for a single material and is therefore interesting for applications where the one material is 3D printed, while the other is overmolded silicone. However, a different lattice is required if the extrusion continuity constraint is to be met by both materials.

V.1.2. CONTRIBUTIONS

We propose an interlocking lattice consisting of interlaced horizontal beams which satisfy the extrusion continuity constraint. The lattice interlocks all degrees of freedom in 3D space and as such lends itself to interfaces between two bodies of different material of arbitrary orientation. We consider the situation of a horizontally applied tensile force to bodies with a vertical interface, and optimize the lattice in two different orientations w.r.t. the interface using a simple analytical model. The optimized structures are validated using numerical simulation as well as physical experiments. Results show that the ITIL lattice on a vertical interface performs comparably to dovetail type interlocking geometries. Optimizing the lattice for interfaces of arbitrary orientation is left as future work.

V.2 METHOD

This paper considers topologically interlocking structures satisfying extrusion continuity constraints, which will be optimized for tensile strength; the geometry of the designs are optimized to yield or break at the highest horizontally applied tensile stress. The highest administrable force applied to a unit cell of any interlocking structure will be carried by both materials a and b . In the plane where tensile failure of material a occurs the total force F will be distributed over the cross-sectional area A_a ; likewise for material b . The theoretically optimal interlocking lattice would be such that the force F will be homogeneously distributed over A_a and A_b in such manner that both materials will fail at the same ultimate force: $F = \sigma_{y,a}A_a = \sigma_{y,b}A_b$. The ultimate force before material a breaks can be improved by increasing the area A_a , but this is at the cost of material b (and vice versa), since their combined area is limited to the total cross sectional area A_{total} of the lattice. The ultimate strength of any multi-material lattice is therefore limited to:

$$\begin{aligned}\sigma^* &= \frac{F}{A_{\text{total}}} = \frac{F}{A_a + A_b} = \frac{\sigma_{y,a}A_a}{A_a + A_b} \\ A_a &= A_b\sigma_{y,b}/\sigma_{y,a} \\ \sigma^* &= \frac{\sigma_{y,a}A_b\sigma_{y,b}/\sigma_{y,a}}{A_b\sigma_{y,b}/\sigma_{y,a} + A_b} = \frac{1}{1/\sigma_{y,a} + 1/\sigma_{y,b}}\end{aligned}\tag{V.1}$$

This formula for the ultimate tensile strength of interlocking lattices is a theoretical

Table V.1: Yield properties of single material MEX printed samples.

material	σ_y	σ_{yZ}	ϵ_y	ϵ_{yZ}
PLA	47 MPa	33 MPa	3.5 %	2.6 %
PP	10.5 MPa	10.6 MPa	29 %	22 %

upper bound. In practice the worst cross-sectional area A_a is in a different plane than the worst cross-sectional area A_b , so they do not add up to A_{total} . For example, in the dovetail geometry (Fig. V.2a) the highest stress of a dovetail of material a is at its base, while the highest stress in b will be at the end of the dovetail of material a . The geometry by which the interlock is secured reduces the strength compared to the theoretical upper bound of the ultimate tensile strength. The upper bound can therefore never be reached.

V.2.1. INTERLACED TOPOLOGICALLY INTERLOCKING LATTICE

The topologically interlocking structure we propose consists of horizontal beams alternating in material. On top of that we print another set of beams rotated about Z and on top beams of the first rotation again, etc. Long horizontal beams assure continuous extrusion, while the alternating direction of the beams assures the interlock. See Fig. V.2c.

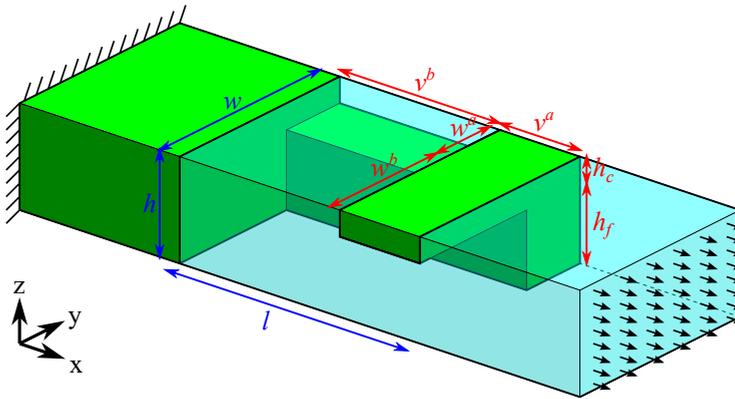
When joining bodies of different materials, some region around the interface between the two bodies should be replaced by a number of unit cells from such a lattice structure. The length of the transition region L is limited by the design of the product within which the lattice is to be employed. We consider a design constraint of $L_{\text{max}} = 12w_{\text{min},m} = 3.6\text{mm}$; this is enough to have some design freedom for the lattice, while limiting the impact on the rest of the product.

We consider a single layer of ITIL cells along the interface. Adding more cells in the direction orthogonal to the interface would make them protrude farther into the two bodies, which limits design applicability. If the two layers of cells would have the same geometry then we would not expect any gain in the ultimate strength of the interlock. Optimizing the geometry of the two layers of cells separately could lead to improved ultimate strength, but falls beyond the scope of this paper.

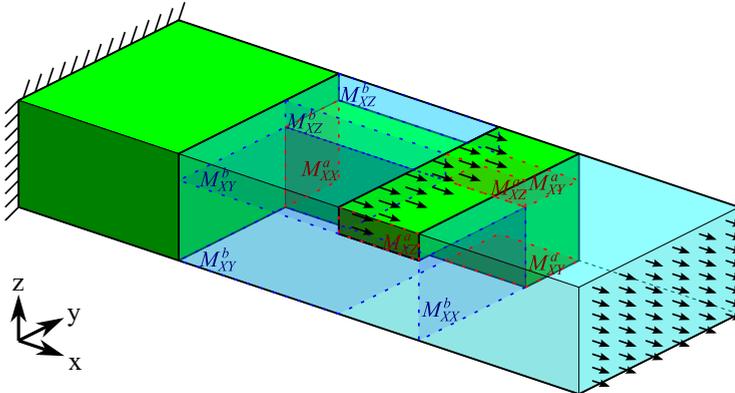
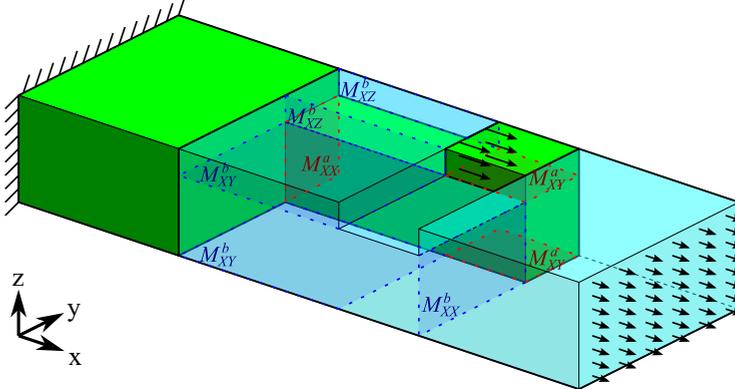
According to the reasoning above we should fill the interface with a single layer of cells, but how should those cells be oriented w.r.t. the interface? While rotating the lattice about X or Y is impossible due to the continuous extrusion constraint, we are free to rotate about Z. However, given that the tensile stress is applied in a single direction it only makes sense to consider two orientations: straight and diagonal.

For the materials chosen by this study, Ultimaker Tough Polylactic Acid (PLA) and Ultimaker PolyPropylene (PP), the theoretical upper bound comes out to be 8.6 MPa. The structure may not only be subject to tensile failure, but also to shear failure modes. Because of the layer-wise build-up employed by MEX, the tensile properties in the Z direction are different from those in the horizontal plane. See Table V.1. The structure is furthermore subject to manufacturing constraints determined by the nozzle sizes and the layer thickness. The standard layer thickness h_{min} is 0.1 mm and the minimal line width w_{min} for a 0.4 mm nozzle is 0.3 mm.

If all dimensions of the beams are set to minimal and the angle between the beams is set to 90° , the finger of the straight ITIL variant takes up a quarter of the area at the base of the cell, so we can expect the strength to be close to $1/4\sigma_{y,b} \approx 2.6\text{MPa}$. For the diagonal ITIL variant the beams take up half of the area at the base of the cell, which would be close to



(a) Design variables (in red)

(b) Failure modes for material *a* (in red) and for *b* (in blue)

(c) Partly broken, partly interlocking cell

Fig. V.3. Unit cell of the straight ITIL variant. Force is transferred from material *a* (green) to material *b* (transparent cyan) through the contact area on the cross beams. In the partly broken situation there is still interlocking, but the total force is transferred through a narrower area.

Stresses in the straight ITIL variant

$$\sigma_{XX,m} = \frac{F}{w_m h_f} \quad (V.2)$$

$$\tau_{XZ,m} = \frac{w_{\neg m}}{w} \frac{F}{2v_m h_c} \quad (V.3)$$

$$\tau_{XY,m} = \frac{w_m}{w} \frac{F}{2v_m w_m} \quad (V.4)$$

where $m \in \{a, b\}$ and $\neg a = b$ and $\neg b = a$

$1/2\sigma_{y,b} \approx 5.3$ MPa. However, these ultimate strengths can be improved upon considerably by optimizing the geometry. This section considers these two ITIL variants and analytical models to optimize them.

V.2.2. STRAIGHT ITIL VARIANT

A unit cell of the straight ITIL variant is visualized in Fig. V.3a. A cell consists of a single *finger* of height h_f protruding from the body of material a outward and part of a *cross beam* of height h_c angled at 90° . In order to print the relatively short fingers using continuous extrusion, we employ the constraints that $w_m \geq 2w_{\min,m}$ (where m is either material a or b), so that the toolpaths for the outline of each layer can go back and forth along the finger without interruption. The cross beams are long and continuous enough, so they could be printed using a single extrusion path: $v_m \geq w_{\min,m}$. We set the minimum height of the beams to twice the layer height, so as to be able to recover from manufacturing inaccuracies¹: $h_f \geq 2h_{\min}$ and $h_c \geq 2h_{\min}$.

STRAIGHT ITIL VARIANT (WHOLE)

In order to optimize the straight ITIL variant for a maximal tensile strength we consider three types of stress, related to three types of failure mode for either material m : tensile stress $\sigma_{XX,m}$ for M_{XX}^m , cross beam shear stress $\tau_{XZ,m}$ for M_{XZ}^m and Z shear stress $\tau_{XY,m}$ for M_{XY}^m . See Fig. V.3b. These three types of failure mode for either material are modeled using classical beam theory as having a homogeneous stress distribution. Because the total force F is modeled as being homogeneously distributed over the whole cross beam, the two shear stresses obtain only a portion of the total force. Furthermore, they are divided by two because the beams and pillars are fixed on both sides. See Eqs. (V.2) to (V.4).

The tensile stress of the whole cell is given by $F/(w_a+w_b)(h_f+h_c)$. Combining the above we obtain the constrained optimization problem given by Eqs. (V.5) to (V.12). The $\sqrt{3}$ in Eqs. (V.11) and (V.12) comes from the von Mises yield criterion. Although one might consider combining the stresses together using the same criterion, this does not increase the accuracy of the model, because the failure mode planes do not overlap.

Because the objective and all stresses are invariant under various scaling operations, we can choose the value of a subset of the design variables. Because of invariance to uni-

¹If the structure would consist of alternating geometry each layer, then the inaccuracy of the one layer can cause over-extrusion in the next layer, which snowballs the problem upward during printing.

Model of the straight ITIL variant

$$f : \max \frac{F}{(w_a + w_b)(h_f + h_c)} \quad (V.5)$$

subject to

$$g_{wb} : w_b \geq 2w_{\min,b} \quad (V.6)$$

$$g_{va} : v_a \geq w_{\min,a} \quad (V.7)$$

$$g_{vb} : v_b \geq w_{\min,b} \quad (V.8)$$

$$g_{hf} : h_f \geq h_{\min} \quad (V.9)$$

$$g_{tm} : \frac{F}{w_m h_f} \leq \sigma_{y,m} \quad M_{XX}^m \quad (V.10)$$

$$g_{ca} : \frac{w_b}{w} \frac{F}{2v_a h_c} \leq \frac{1}{\sqrt{3}} \sigma_{y,a} \quad M_{XZ}^a \quad (V.11)$$

$$g_{zm} : \frac{w_m}{w} \frac{F}{2v_b w_b} \leq \frac{1}{\sqrt{3}} \sigma_{yZ,m} \quad M_{XY}^b \quad (V.12)$$

where Eqs. (V.10) and (V.12) are duplicated
for both materials $m \in \{a, b\}$

form scaling we employ the design constraint; scaling up the width and force only increases the cross beam shear values $\tau_{XZ,m}$, so we employ the minimum width constraint; scaling up the height and force only increases the Z shear values $\tau_{XY,m}$, so we employ the minimum height constraint. This way we limit the design space from six to three dimensions, i.e. w_b, h_f, v_a :

$$w_a = 2w_{\min,a} \quad h_c = 2h_{\min} \quad v_b = L_{\max} - v_a \quad (V.13)$$

Using the formulae of the constrained optimization problem one can find the maximum force and thus the maximum stress of any design; we can rewrite each of the mechanical constraints Eqs. (V.10) to (V.12) to give a formula for F and the lowest value of those formulae will give the active failure mode for a given design. The resulting response surfaces are shown in Fig. V.4a. The optimum is at the intersection of four constraint surfaces, which is remarkable for a 3D space.

BROKEN CROSS BEAMS MODEL

A careful analysis of the geometry will show that if shear failure M_{XZ}^a has occurred, there is still interlocking between the two materials. If a part of the cross beams of a has sheared off, still the pillar of material a remains, which is surrounded by material b . See Fig. V.3c. Once failure mode M_{XZ}^a has occurred, still any other failure mode has to occur for the interlock to fail completely. Since the failure mode can only happen by part of the PP cross beam pushing against the part of the PLA cross beam which is being sheared off, both cross beams shear together. Because PLA breaks at a lower strain than PP (see Table V.1), we know that for these materials M_{XZ}^b never occurs unless M_{XZ}^a has occurred. We will con-

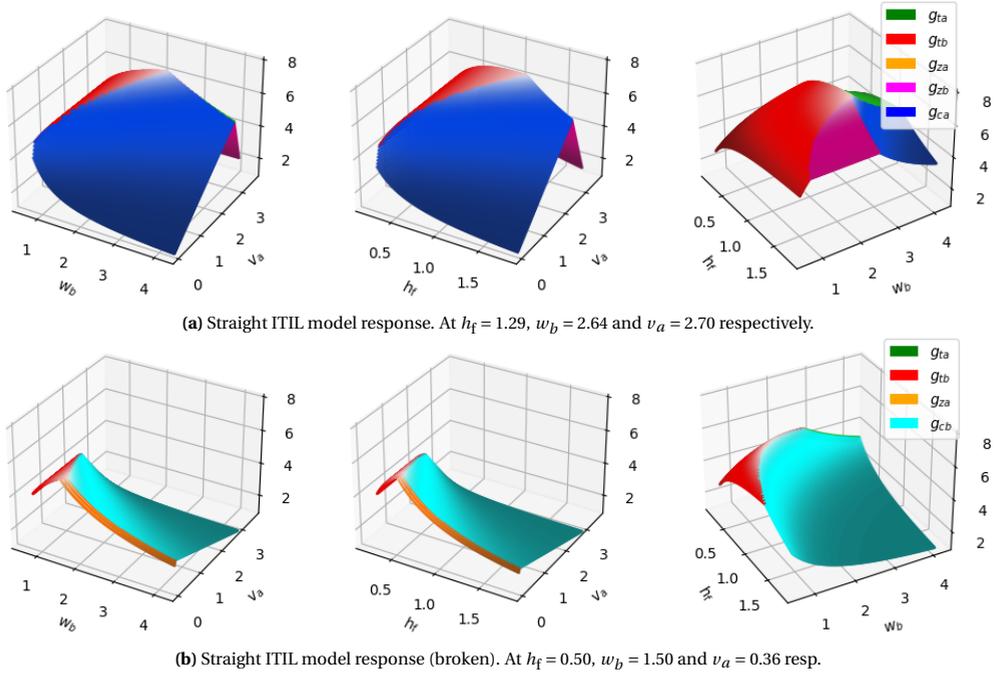


Fig. V.4. Maximum strength according to analytical models for the straight ITIL variant along three 2D slices of the 3D design space. Four constraint surfaces meet at the optimum. By taking the maximum values of these two models the maximum stress before separation can be calculated.

Straight ITIL variant - shear constraints for broken case

$$g_{cb} \dagger : \frac{F}{2(L_{\max} - v_a)2h_{\min}} \leq \frac{1}{\sqrt{3}}\sigma_{y,b} \quad M_{XZ}^b \quad (V.14)$$

$$g_{zb} \dagger : \frac{F}{2v_a w_a} \leq \frac{1}{\sqrt{3}}\sigma_{yZ,a} \quad M_{XY}^a \quad (V.15)$$

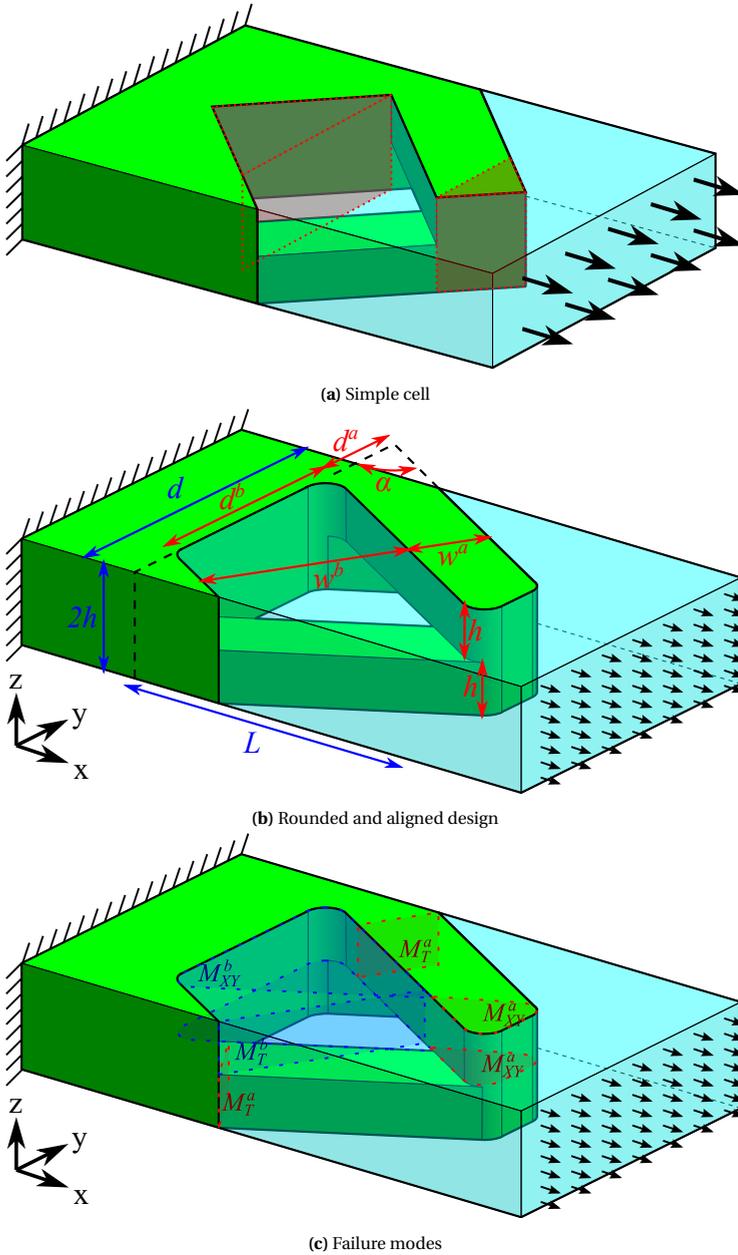


Fig. V.6. Unit cell of the diagonal ITIL variant. The M_T^m failure modes are orthogonal to the beams and influenced by both tensile and shear components of the total force.

Model for the diagonal ITIL variant

$$f : \max \frac{F}{2hd} \quad (\text{V.16})$$

subject to

$$g_{wa} : w_a \geq 2w_{\min,a} \quad (\text{V.17})$$

$$g_{wb} : w_b \geq 2w_{\min,b} \quad (\text{V.18})$$

$$g_d : L \leq L_{\max} \quad (\text{V.19})$$

$$g_{tm} : \frac{F}{2w_m h} \sqrt{\left(\frac{w}{d}\right)^2 + 3\left(\frac{w}{2M}\right)^2} \leq \sigma_{y,m} \quad M_T^m \quad (\text{V.20})$$

$$g_{zm} : \frac{F}{2A_{z,m}} \leq \frac{1}{\sqrt{3}} \sigma_{yZ,m} \quad M_{XY}^m \quad (\text{V.21})$$

where

$$d = 2Mw / \sqrt{4M^2 - w^2}$$

$$M = L - 2r$$

$$w = w_a + w_b$$

$$A_{z,m} = \frac{1}{2}\pi r^2 + r(w_m - 2r)\frac{d}{w} + dM\left(\frac{w_m}{w}\right)^2$$

where Eqs. (V.20) and (V.21) are duplicated

for both materials $m \in \{a, b\}$

V

that the rounded edges at the ends of both fingers are aligned. See Fig. V.5a.

Because of the diagonal geometry we expect that the stress distribution throughout the structure will be quite complex, and the failure will depend on the deformation of both materials during stretching. Nevertheless, we provide a simplified analytical model assuming the stress is homogeneously distributed and disregarding the influence of deformation on stress distribution. We decompose the force F into one component parallel and another orthogonal to the direction of the beams. We then use the parallel component to determine the tensile stress and the orthogonal component to determine the shear stress in the beam, which combine into a single constraint using the von Mises yield criterion in Eq. (V.20). The diagonal ITIL model is then defined by the constrained optimization problem given by Eqs. (V.16) to (V.21).

Again the stresses are all scale-invariant, so we set $w_a = 2w_{\min,a}$. However, because changes in w_b given the same L will change the angle α of the beams, we cannot assume the design constraint g_d is active.

The resulting response surface can be viewed in Fig. V.5b. Because the height of the fingers is minimal the Z shear constraints are both dominated.

V.3 RESULTS

In order to validate our analytical models we compare its predictions against both simulation results and physical tensile tests. While the physical tests constitute the final arbiter

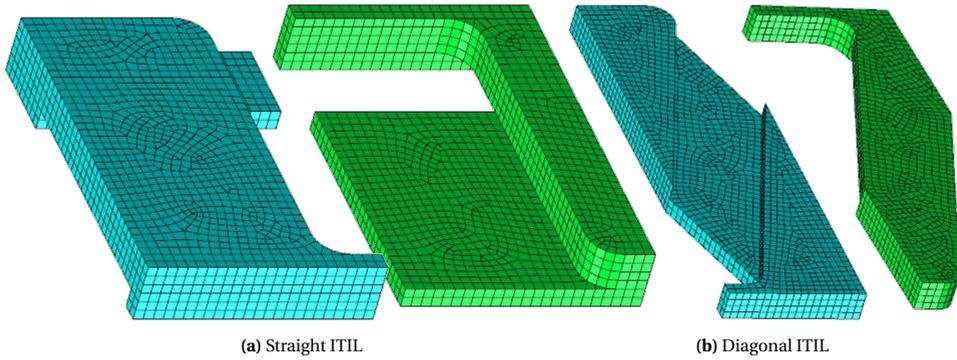


Fig. V.7. Example simulation meshes. The mesh of the diagonal ITIL variant is half a unit cell covering only one of the two diagonal fingers. Because of the rounded corners an axis aligned meshing was impossible, leading to singularities in the mesh.

on the matter, our resources for performing physical tests are limited. Simulations can easily be performed by running a script for multiple days without user interaction. Given that the simulations make use of the same material properties which were acquired from tensile tests performed by Ultimaker, the simulations can teach us about the validity of the homogeneity assumptions in the analytical models. Moreover, the physical test results are afflicted with a spread in manufacturing inaccuracies. The simulations can therefore enrich the understanding we gain from physical experiments.

V.3.1. NUMERICAL SIMULATION

In order to simulate interlocking structures with a range of design parameters we automatically generate INP files in Abaqus CAE (2020) using a script. Solving an INP file gives us the force-displacement graph, from which we can determine the ultimate tensile strength of a particular design. In order to simulate accurately we used the stress-strain curves from tensile tests on the base materials printed flat on the build platform as the plasticity in tabular form. The simulations were performed in the Abaqus/Explicit solver where the Dynamic, Explicit procedure step was used with a mass scaling factor of 10^7 , using geometric non-linearity and general contact (explicit) to disregard friction for simplicity.

The repeating nature of the interlocking patterns was captured by modeling half of the unit cell and applying symmetry constraints to the sides, top and bottom. The model was meshed using C3D8R hexahedral elements of $\pm 75\mu\text{m}$.

A grid search was used to measure the influence on the ultimate strength along each of the design variables w_b, v_a, h_c , along with the total length L . The search space was therefore 4D and 2D for the straight and diagonal ITIL model respectively. In order to estimate the optimum we fit a smooth response surface to these data points using a radial basis function (RBF) network [44], with a smoothness of $\lambda = 1$.

STRAIGHT ITIL VARIANT

In order to prevent stress concentrations and adhere to manufacturing accuracy, the vertical edges of the straight ITIL variant were rounded with $r = 0.15\text{mm}$; see Fig. V.7.

Newton's method was used to determine the optimum, starting from the best sampled point. This step only considered the dimensions w_b and v_a , because L_{max} is given and h_f

Table V.2: Optimal designs according to the hypersurface fitted to the FEM simulations for the straight ITIL variant and for the diagonal ITIL variant.

	L_{\max} (mm)	3.6	3.0	2.4	1.8
straight	σ_{\max} (MPa)	6.17	6.12	5.89	5.59
	h_f (mm)	0.8	0.8	0.7	0.6
	w_b (mm)	2.58	2.42	2.17	2.05
	ν_a (mm)	2.67	2.23	1.78	1.35
diag	σ_{\max} (MPa)	6.30	6.37	5.86	4.69
	w_b (mm)	1.21	1.19	1.18	1.04

has to be an integer multiple of h_{\min} . The resulting hypersurfaces are visualized in Fig. V.8. The obtained optima are shown in Table V.2.

We compare these results to our analytical model by adjusting the analytical model to capture the inaccurate Z strength used in the simulations: $\sigma_{yZ,m} := \sigma_{y,m}$. See Fig. V.9. We then observe that our analytical model on average predicts only 7.8% higher ultimate strength values than then the FEM simulations, with a standard deviation of 16.2%.

V

DIAGONAL ITIL VARIANT

Modeling the diagonal ITIL variant in Abaqus can be quite cumbersome, since it does not natively support periodic boundary constraints. Whereas this problem can be overcome in the straight ITIL variant because it is symmetric, the diagonal variant is only *rotationally* symmetric. While a symmetry constraint can be used on the top and bottom, the two sides of the design are mirror images of each other, but also flipped vertically.

However, since the height of the beams is relatively low compared to their width we have observed that the stresses and strains are quite similar in the top and bottom. If we model half of the diagonal ITIL cell by cutting it vertically and apply symmetry constraints to the sides, the induced error is only $\pm 10\%$ compared to simulating an interface consisting of two whole cells.

The results of these simulations are shown in Fig. V.10a. We compare the simulations against our analytical model in Fig. V.10b. The analytical model predicts only 0.4% lower ultimate strength values on average with a standard deviation of 10%.

V.3.2. PHYSICAL EXPERIMENTAL TESTS

Tensile tests were performed on an Instron 3366 Universal Testing machine at 5 mm/min. Prints were manufactured on a Ultimaker S5 systems in 5-fold with Ultimaker Green Tough PLA and Ultimaker PP using the default 0.1 mm layer thickness profile, with 100% infill and a custom brim to make sure both materials stick to the build platform. For PP we print the outer before the inner walls so as to improve the dimensional accuracy. In order to deal with the various widths of the beams we generate toolpaths from STL 3D models using the Cura Arachne Engine beta release [45], which implements a framework for generating variable line width toolpaths to fill small geometry of arbitrary dimensions [46]. The Inward Distributed and the Distributed strategy were used on PLA and PP respectively. See Fig. V.11.

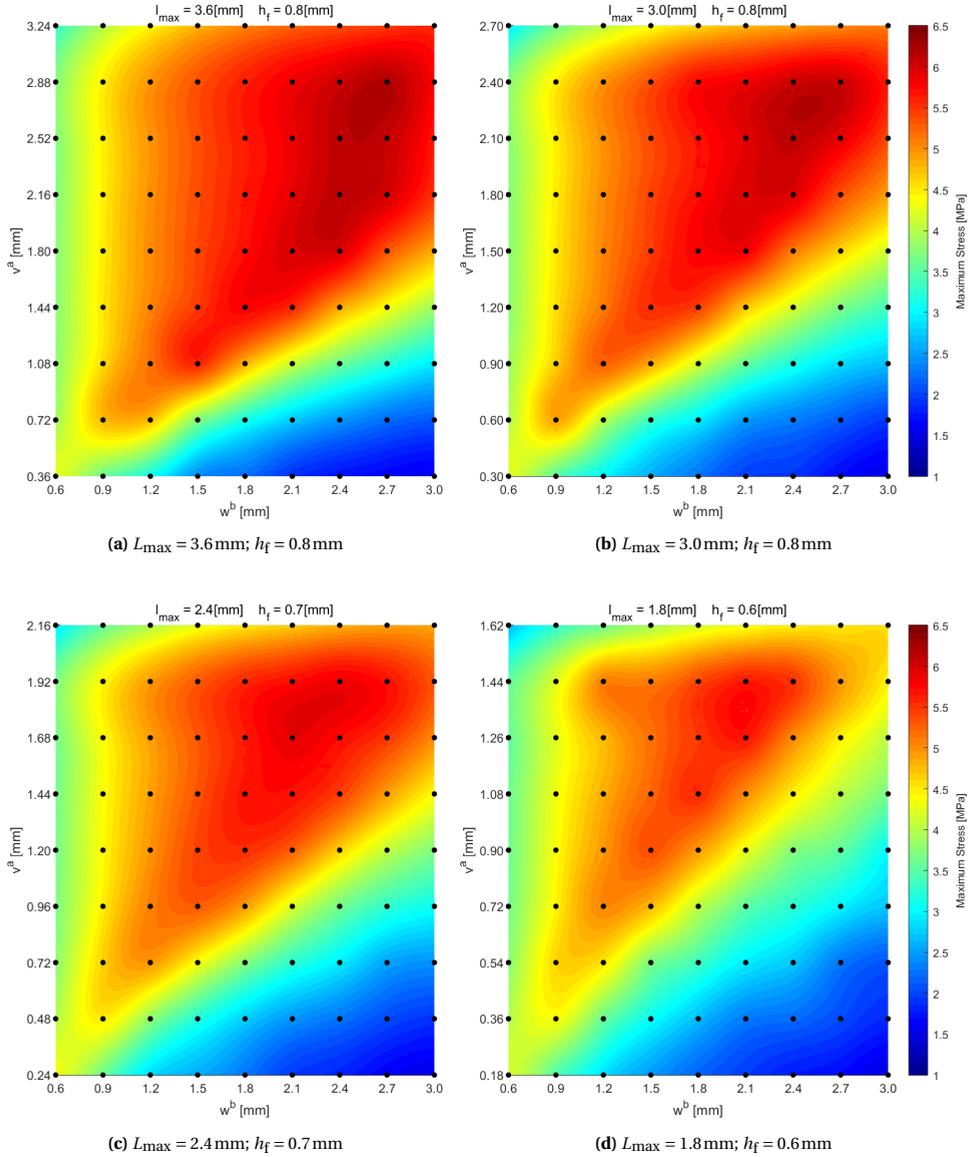


Fig. V.8. 2D slices of the 4D simulation results and fitted RBF hypersurface for the straight ITIL variant. Sampled data points in black.

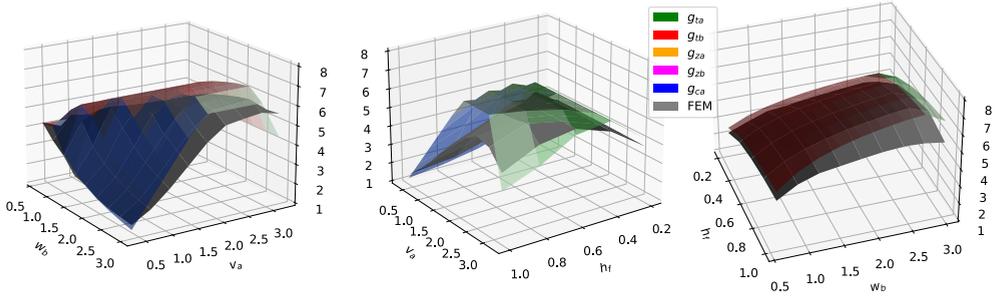


Fig. V.9. Ultimate strength according to the analytical models and the simulation results for the straight variant of the ITIL lattice. The analytical models follow roughly the same shape and same height as the simulation results. The response on three 2D slices of the 3D design space are shown, from top to bottom at $h_f = 0.8$, $w_b = 2.7$ and at $v_a = 2.88$.

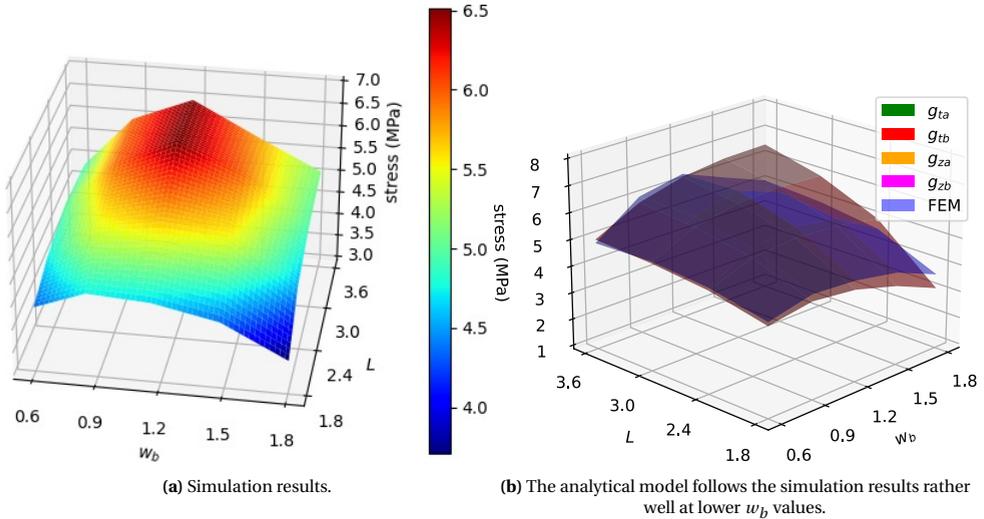


Fig. V.10. Simulation results for diagonal ITIL variant using linear interpolation between the simulation results.

MODEL PARAMETERS

Straight ITIL The straight ITIL variant suffers from the curse of dimensionality; even when setting $w_a = 0.6$ mm, $h_c = 0.1$ mm and $L = 3.6$ mm, there are still the three free design variables w_b , v_a and h_f to determine. With 5 specimens per sample point and limited resources, the total number of data points we are able to test is limited. We therefore chose to sample close to the two optima of the analytical models: whole and broken, as well as deviations from those optima in both directions along the axis of each design variable of 0.3 mm in w_b and v_a and 0.2 mm in h_f . See Fig. V.12c.

Each sample of the straight ITIL variant has 5×5 cells. Because the repetition of cells is broken at the sides of the specimen, the boundary cells are adjusted for manufacturability and stability. The specimens end with a PLA finger on both sides and in cross beams on both top and bottom. See Fig. V.12a.

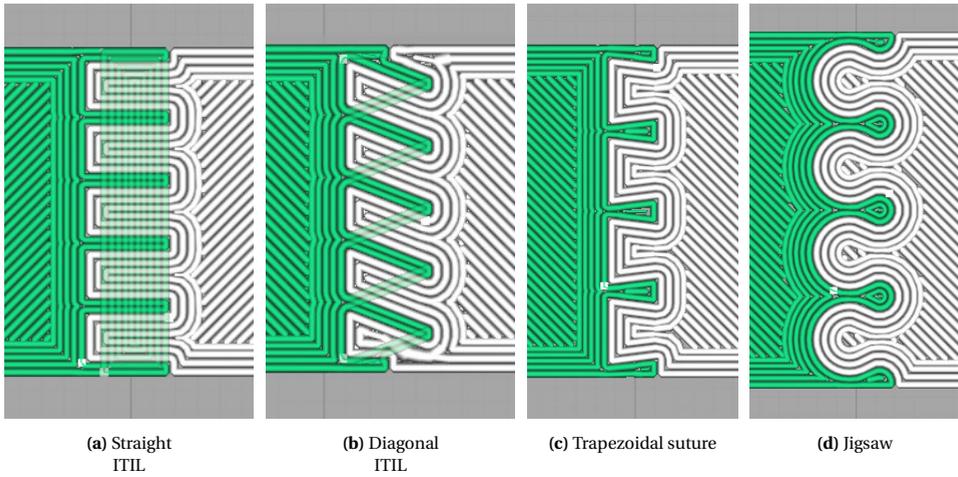


Fig. V.11. Gcodes generated with Cura Arachne engine beta. The layers of main fingers are shown and the other layers in a transparent overlay. While the straight and diagonal ITIL lattice produce continuous extrusion beads, the toolpaths for the dovetail designs include small separated segments.

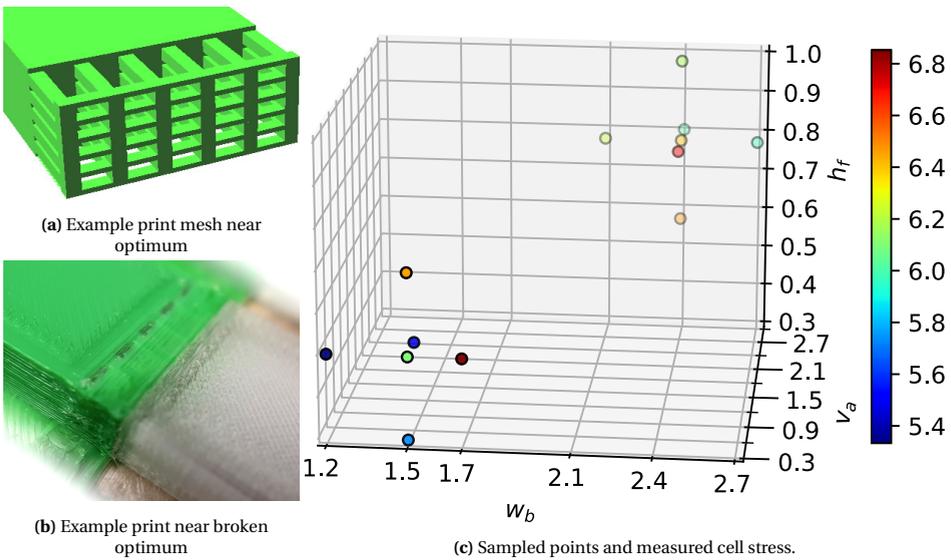


Fig. V.12. Experimental setup of straight ITIL variant. Note that for the broken model v_a is minimal, so v_a is not a valid sample.

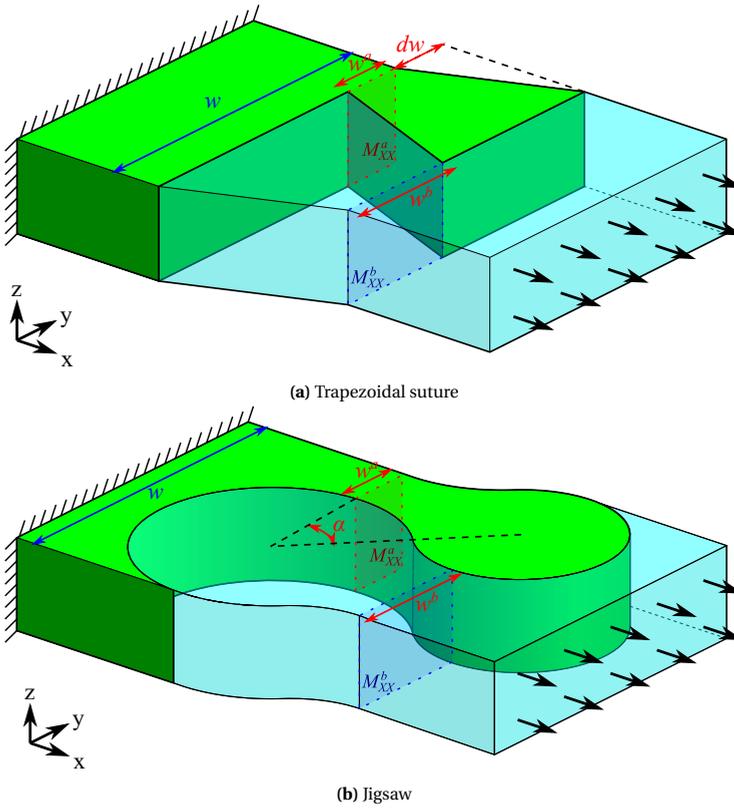


Fig. V.13. Simple 2D dovetail interlocking lattices.

Diagonal ITIL variant Each sample of the diagonal ITIL variant contains 5 cells in the horizontal direction, but 13 repetition in Z because of the low unit cell height. Extra finger beams are added to the sides of the specimen to prevent any part of the beam to be less than $2w_{\min,} = 0.6\text{ mm}$ wide. See Fig. V.11b. Because with a given $L = 3.6\text{ mm}$, $h = 0.2\text{ mm}$ and $w_a = 0.6\text{ mm}$ the remaining design space is only one-dimensional, we can simply sample various points along w_b : (0.6, 1.2, 1.8, 2.4, 3.0, 3.6).

Dovetail interlocking We compared our interlocking structures against two interlocking designs: trapezoidal sutures and jigsaw interlocking. See Fig. V.13. We used $w_a = 2w_{\min,a}$, $dw = 0.3\text{ mm}$ and $L = 2.4$ for the trapezoidal suture, and $w_a = 2w_{\min,a}$ and $\alpha = 35^\circ$ for the jigsaw interlocking design. We printed samples with both $w_b = 3w_a$ and $w_b = \sigma_{y,a}/\sigma_{y,b}w_a = 4.48w_a$. We used 6 and 4 repetition respectively and a height of 5 mm.

Note that the jigsaw interlocking structure is quite similar to the trapezoidal suture, with the addition of semicircles to the ends of the trapezoids. The total length L of the jigsaw structure is 2.96 mm and 4.05 mm for the two w_b values respectively, so the L_{\max} constraint is violated by that structure.

The boundaries of these two structures end in half a PLA lobe, because that is the stiffer material. In order to meet the minimum width constraint there, the sides of the specimen are extruded by $w_{\min,a}$. However, because the dovetails wider toward the tip, they

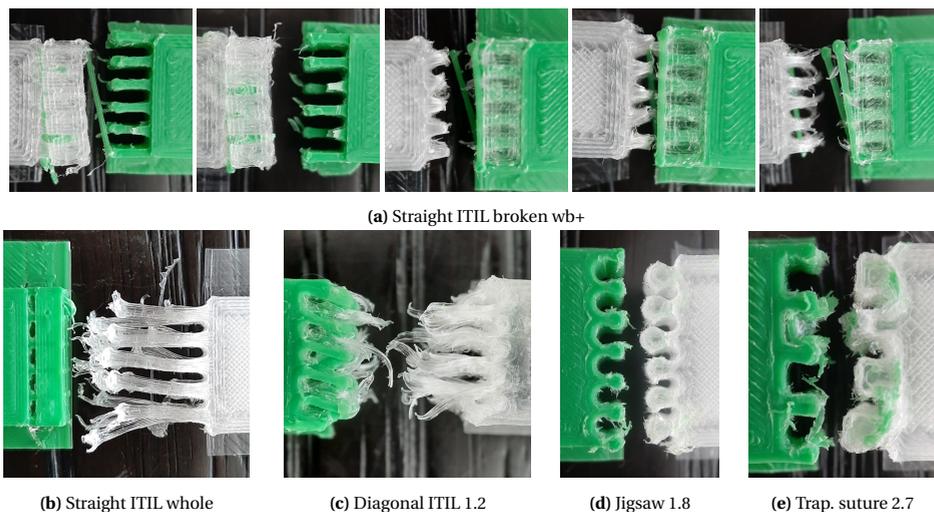


Fig. V.14. Samples after tensile tests of the best performing designs. The broken wb+ samples from the straight ITIL variant exhibit multiple failure modes, indicating that this sample was close to the intersection of several constraint surfaces.

require an extra toolpath to be filled densely, which is disconnected from the other toolpaths, thereby violating extrusion continuity constraints. See Figs. V.11c and V.11d

PHYSICAL TEST RESULTS

After tensile testing we can observe various failure modes, such as in Fig. V.14. The tensile tests performed result in force-displacement graphs, from which the ultimate tensile strength values are derived. The slope of these graphs after the optimum has been reached furthermore tells us something about the failure mode by which the sample has failed. See Fig. V.15.

Because the boundary cells deviate from the regular pattern, computing the ultimate tensile strength can be done in two ways, giving rise to two statistics. We compute the *cell stress* by dividing the maximum force of the force-displacement graphs by the number of cells and then divide it by the cross sectional area of the cell. We compute the *total stress* by dividing the force by the total cross-sectional area of the sample, including the extra geometry at the boundaries of the sample. Ideally we would compensate for manufacturing inaccuracies by using measured dimensions of the specimens, but measuring the internal geometry of the interlocking structure is practically infeasible, so we use the dimensions of the 3D mesh instead. The results from the physical experiments, along with the analytical and simulated predictions are gathered in Fig. V.16.

V.4 DISCUSSION

V.4.1. VALIDITY OF ANALYTICAL MODELS

When comparing the tensile test results to the analytical models and the simulation results (Fig. V.16), we observe that the whole straight ITIL variant is predicted rather well by the models. The analytical model tends to overestimate the strength because of homogeneous stress assumptions, while the simulations tend to underestimate the strength.

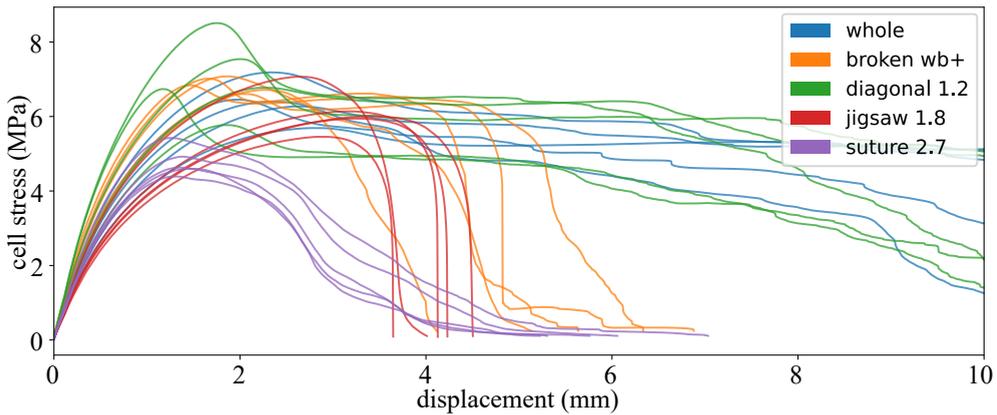


Fig. V.15. Comparison of the best performing design of each type. The diagonal ITIL design with $w_b = 1.2$ mm showed the highest maximum tensile stress. By inspecting the shape of the graph you can determine the dominant failure mode: PLA breaking (sharp drop) or PP yielding (plateau). Dovetail unlocking has either a sharp drop or a more gradual decline.

V

Although the standard deviation is too high to capture the differences in response with significance, the analytical model follows the trends of the FEM simulations rather well and the simulation results generally fall within one standard error of the physical results.

The response around the anticipated optimum of the broken cross beams model is predicted by our models considerably worse. The fact that the simulations also performed worse in that region was expected, since after cells in the mesh are broken some self-intersection may occur in the simulation. When viewing the best performing straight ITIL sample (broken w_b+ , Fig. V.14a) we observe that several failure modes occur throughout the sample (due to manufacturing inaccuracy), but shearing off of part of the cross beams is not one of them. One possible explanation is that for geometry so small as the order of the nozzle size the micro-gaps in between adjacent extruded strands impair the material properties less than in the relatively large test samples with which the material properties of MEX printed parts were determined. The empirically obtained material properties of MEX printed parts may be less accurate at smaller scales.

By comparing the analytical model for the diagonal ITIL variant to the simulations results we observe that at higher w_b values the analytical model overestimates the strength. The height of the constraint surface for g_{ta} (see Fig. V.5b) was higher than it was simulated to be, which can be related to the fact that for a higher w_b the angle α of the beams is lower, which means that the shear component of the stress was higher and that the stress might have been less uniform.

This overestimation of the analytical model is compensated by an underestimation of the simulations, which makes the analytical predictions fit well to the empirical results; see Fig. V.16. The underestimation of the simulations can be related to the type of boundary conditions which were applied, but also to an underestimation of the empirically determined material properties of 3D printed parts. Overall the diagonal ITIL model is supported by the data.

The analytical modeling of the dovetail geometries considers only homogeneous tensile stress, but disregards the fact that the interlock is broken by deformation alone. The fact that the dovetail designs do not rely on topological interlocking means that they are

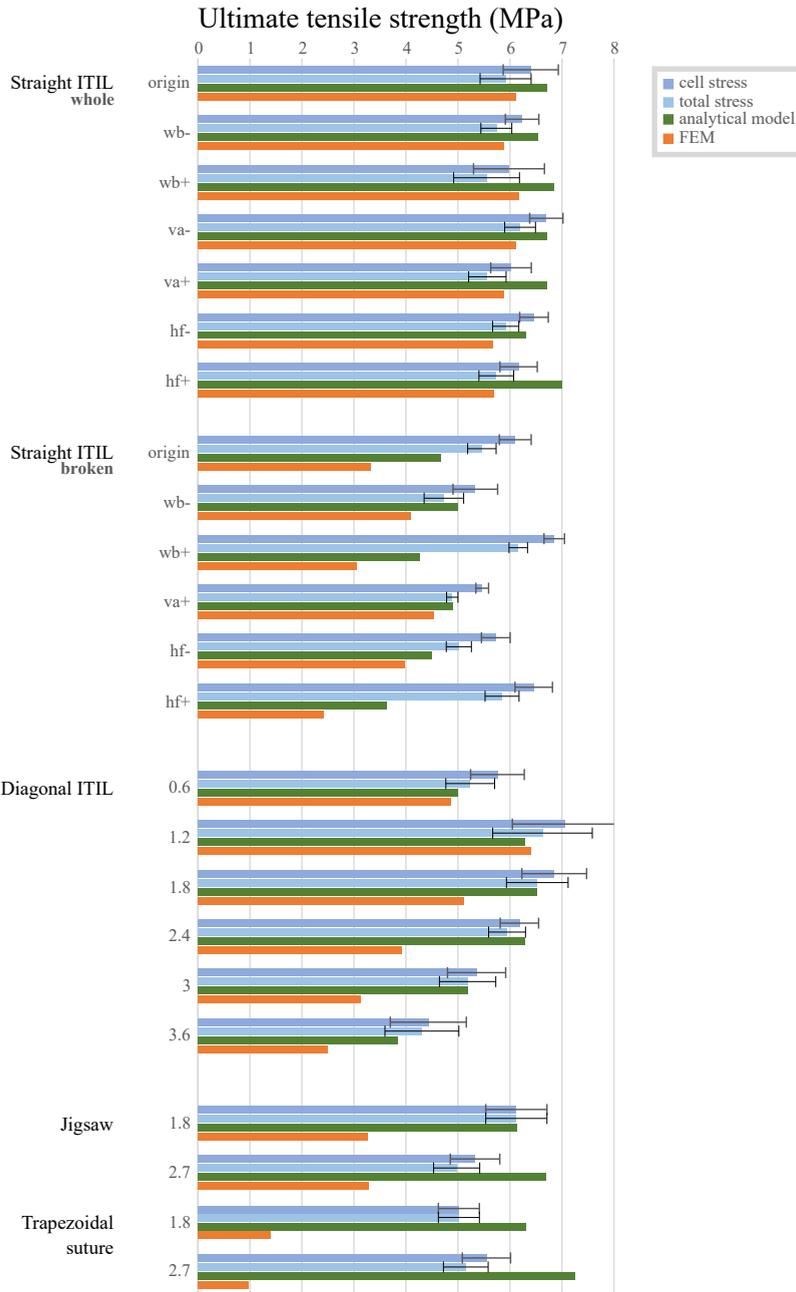


Fig. V.16. Test results compared to the predictions according to the analytical model and the RBF network fitted to the simulation results. The straight ITIL variant samples are labelled relative to the whole and broken optimum, while the rest is labelled by their w_b value.

more difficult to model.

V.4.2. COMPARISON OF DIFFERENT INTERLOCKING STRUCTURES

All interlocking designs considered can reach roughly 6 to 7 MPa, which means that ± 2 MPa of the theoretical upper bound of 8.6 MPa from Eq. (V.1) is used to secure the interlock. The diagonal ITIL design with a cell stress of 7.07 ± 1.0 MPa seems to outperform the others, but not significantly. See Fig. V.15.

The dovetail designs performed considerably worse than the other designs – even when the L_{\max} constraint was violated. Moreover, some prints of the dovetail designs showed contamination between the two materials (see Fig. V.14e), which might be caused by the toolpath discontinuities (see Fig. V.11c).

Some designs of the straight ITIL variant near where the broken model is optimal perform significantly better than the other tested straight ITIL designs; given the high dimensionality of the design space it might be the case that there is a straight ITIL design in between the whole and broken optima which outperforms the diagonal ITIL variant nonetheless.

For designs where a lower L_{\max} is allowed the diagonal ITIL variant performs worse. For $L_{\max} = 1.8$ mm the diagonal is simulated to outperform the straight ITIL variant by 5.53 MPa to 4.69 MPa. See Table V.2. The performance of the diagonal ITIL model greatly reduces for lower L_{\max} , because the angle of the beams is reduced, which causes them to be more susceptible to shear stresses.

V.4.3. LIMITATIONS

The biggest obstacle to drawing definitive conclusions about the relative strengths of the interlocking structures is the size of the standard deviation. The variation in results has a multitude of causes, related to print head position inaccuracy, filament diameter fluctuations, temperature oscillations and even the location on the build platform.² In order to mitigate these factors we have spread out the various geometries over 5 different Ultimaker S5 printers and different locations on the build platform. Although this increases the variability in the results it does increase the reliability.

Because we needed to test a large quantity of samples, we had to limit the size of the samples in terms of interface dimensions. For the 5×5 cells of the straight ITIL designs 12 out of 25 cells were boundary cells, so it is challenging to extrapolate these results to larger interfaces. Moreover, the different geometries were tested with different amounts of cells, because the cell geometries were very different. The dovetail designs were modelled without taking friction into account and the size of the interlock (dw, α) was not optimized for. This makes it hard to draw definitive conclusions comparing the different interlocking geometries.

Finally, the validity of our models turned out to be limited, partly because of homogeneity assumptions and partly because we used empirical material properties of 3D prints, rather than simulating how those properties come about. In order to get more accurate models we could emulate the contact area and polymer entanglement between neighbouring traces, fit that to the empirically obtained tensile properties and use the resulting fitted model to simulate the structures on a toolpath level.

²With different locations on the build platform the Bowden tube is bent differently, altering the pressure and thus the amount of extrusion.

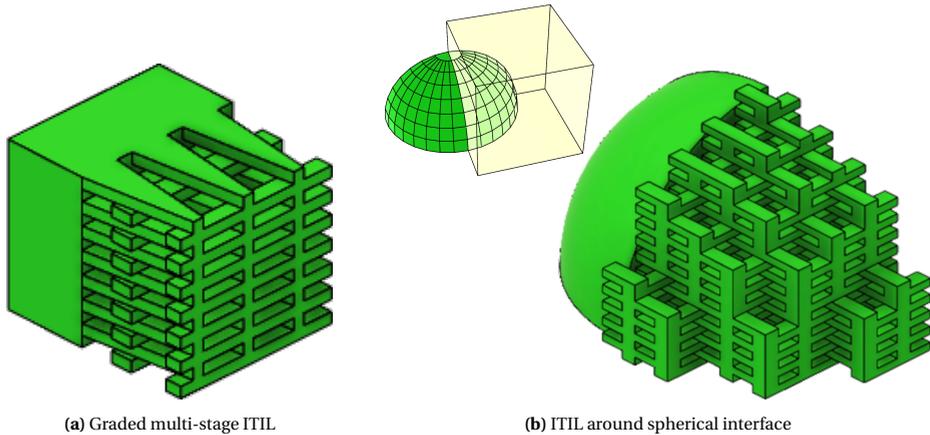


Fig. V.17. Extensions of ITIL. (Only PLA is shown.)

V.5 CONCLUSION

In this paper we have presented the ITIL lattice: a highly interconnected interlocking structure, which cannot be broken by deformation alone and satisfies extrusion continuity constraints. The lattice locks all three axes of 3D space, and because the lattice is repeating in all three dimensions it is not limited to flat vertical interfaces. The lattice was optimized for flat vertical interfaces and shown to exhibit comparable ultimate strength to dovetail type interlocking designs.

When optimizing for a horizontally applied tensile stress orthogonal to the interface, two orientations of the ITIL lattice present themselves: the straight ITIL variant and the diagonal ITIL variant. We have developed analytical models for these two and compared them against simulation results and physical tensile tests.

Based on just the average tensile strength of the physical specimens we could conclude that the diagonal ITIL variant outperforms the straight one and existing dovetail designs, but the standard deviation is too high to make that conclusion definitive. The diagonal ITIL model is more simple, has less free design variables, has no active Z shear constraints, the cells are smaller and the toolpaths are fully continuous because all beams are directly connected to the main body. All tested interlocking designs can reach between 6 and 7 MPa, which is roughly three quarters of the theoretical upper bound of 8.6 MPa, but the optimal diagonal ITIL design produces the best average tensile strength. However, according to the analytical models and simulation results the diagonal ITIL variant is outperformed by the straight variant for $L_{\max} < 2.4$ mm., so for products with a smaller margin available for the interlocking structure the straight ITIL variant is preferred.

V.5.1. FUTURE WORK

Future work might be aimed at loading scenarios different from tensile, different materials, different nozzle sizes and different layer heights. Investigating other properties such as the displacement at break, stiffness or toughness would also be interesting under compressive loads. It is specifically compelling to investigate the resilience of the interlocking structure against vertically applied loads. If a larger space around the interface between the two bodies is allowed to be altered by the structure, one could optimize for a multi-

layer interlocking structure, where the cells of the different layers have different geometry, such as visualized in Fig. V.17a. Another interesting route is to optimize for an interface which has a complex geometry and heterogeneous stress distribution. While the ITIL lattice inherently allows for complex interfaces (see Fig. V.17b), optimizing the orientation and geometry of cells for complex surfaces remains future work. In the light of sustainability and recycling one might select for specific failure modes, such as Z shear failure, which do not leave parts of the one material attached to the body of the other material after failure. Taking a broader perspective one might design an interlocking structure which is both topologically interlocking and also has dovetail features.

REFERENCES

- [1] Tim Kuipers et al. “ITIL: Interlaced Topologically Interlocking Lattice for continuous dual-material extrusion”. In: *Additive Manufacturing* 50 (2022), p. 102495. ISSN: 2214-8604. DOI: <https://doi.org/10.1016/j.addma.2021.102495>.
- [2] Mohammad Vaezi et al. “Multiple material additive manufacturing - Part 1: A review: This review paper covers a decade of research on multiple material additive manufacturing technologies which can produce complex geometry parts with different materials”. In: *Virtual and Physical Prototyping* 8.1 (2013), pp. 19–50. ISSN: 17452759. DOI: 10.1080/17452759.2013.778175.
- [3] Mohammad Rafiee, Rouhollah D. Farahani, and Daniel Therriault. “Multi-Material 3D and 4D Printing: A Survey”. In: *Advanced Science* 7.12 (2020), pp. 1–26. ISSN: 21983844. DOI: 10.1002/advs.201902307.
- [4] Emanuel M Sachs et al. *Three-dimensional printing techniques*. 1994.
- [5] Chi Zhou et al. “Development of a multi-material mask-image-projection-based stereolithography for the fabrication of digital materials”. In: *22nd Annual International Solid Freeform Fabrication Symposium - An Additive Manufacturing Conference, SFF 2011* (2011), pp. 65–80.
- [6] Alan Brunton et al. “3D printing spatially varying color and translucency”. In: *ACM Transactions on Graphics* 37.4 (2018), pp. 1–13. ISSN: 07300301. DOI: 10.1145/3197517.3201349.
- [7] Willemijn Elkhuisen et al. “Gloss, color, and topography scanning for reproducing a Painting’s appearance using 3D printing”. In: *Journal on Computing and Cultural Heritage* 12.4 (2019). ISSN: 15564711. DOI: 10.1145/3317949. arXiv: 1910.10836.
- [8] Tim Reiner et al. “Dual-color mixing for fused deposition modeling printers”. In: *Computer Graphics Forum*. Vol. 33. 2014, pp. 479–486.
- [9] Haichuan Song et al. “Colored Fused Filament Fabrication”. In: *ACM Transactions on Graphics* 38.5 (2019), pp. 1–11. ISSN: 07300301. DOI: 10.1145/3183793. arXiv: 1709.09689.
- [10] T. Kuipers et al. “Hatching for 3D prints: Line-based halftoning for dual extrusion fused deposition modeling”. In: *Computers & Graphics* 74 (2018), pp. 23–32. ISSN: 00978493. DOI: 10.1016/j.cag.2018.04.006.
- [11] W Cho et al. “A dithering algorithm for local composition control with three-dimensional printing”. In: *Computer-Aided Design* 35.9 (2003), pp. 851–867. ISSN: 00104485.

- [12] Yuen-Shan Leung et al. “Digital Material Design Using Tensor-Based Error Diffusion for Additive Manufacturing”. In: *Computer-Aided Design* 114 (2019), pp. 224–235. ISSN: 00104485. DOI: 10.1016/j.cad.2019.05.031.
- [13] Mohammad J. Mirzaali et al. “Multi-Material 3D Printing of Functionally Graded Hierarchical Soft–Hard Composites”. In: *Advanced Engineering Materials* 22.7 (2020). ISSN: 15272648. DOI: 10.1002/adem.201901142.
- [14] Bo Zhu et al. “Two-Scale Topology Optimization with Microstructures”. In: (2017). ISSN: 07300301. DOI: 10.1145/3095815. arXiv: 1706.03189.
- [15] RepRap.me. *Diamond Hotend*. 2015. URL: reprap.org/wiki/Diamond%7B%5C_%7DHotend.
- [16] Haruki Takahashi, Parinya Punpongsanon, and Jeeun Kim. “Programmable Filament: Printed Filaments for Multi-material 3D Printing”. In: *Proceedings of the 33rd Annual ACM Symposium on User Interface Software and Technology*. New York, NY, USA: ACM, 2020. ISBN: 9781450375146. URL: <https://doi.org/10.1145/3379337.3415863>.
- [17] Mosaic. *Palette*. 2015. URL: www.mosaicmfg.com/products/palette-3.
- [18] Ultimaker. *Ultimaker S5*. 2018. URL: ultimaker.com/3d-printers/ultimaker-s5.
- [19] Joseph E. Cadman et al. “On design of multi-functional microstructural materials”. In: *Journal of Materials Science* 48.1 (2013), pp. 51–66. ISSN: 00222461. DOI: 10.1007/s10853-012-6643-4.
- [20] Binbin Zhang et al. “Additive Manufacturing of Functionally Graded Material Objects: A Review”. In: *Journal of Computing and Information Science in Engineering* 18.4 (2018), p. 041002. ISSN: 1530-9827. DOI: 10.1115/1.4039683.
- [21] Francesco Tamburrino, Serena Graziosi, and Monica Bordegoni. “The design process of additive manufactured Meso-Scale Lattice Structures: a review”. In: *Journal of Computing and Information Science in Engineering* 18.4 (July 2018). ISSN: 1530-9827. DOI: 10.1115/1.4040131.
- [22] Raphael Freund et al. “Determination of influencing factors on interface strength of additively manufactured multi-material parts by material extrusion”. In: *Applied Sciences* 9.9 (2019), p. 1782.
- [23] Steven Abbott. *Adhesion science: principles and practice*. DEStech Publications, Inc, 2015. DOI: 10.1177/0340035206070163.
- [24] Xia Gao et al. “Fused filament fabrication of polymer materials: A review of inter-layer bond”. In: *Additive Manufacturing* 37 (2021), p. 101658. ISSN: 2214-8604. DOI: <https://doi.org/10.1016/j.addma.2020.101658>.
- [25] Ana Elisa Costa, Alexandre Ferreira da Silva, and Olga Sousa Carneiro. “A study on extruded filament bonding in fused filament fabrication”. In: *Rapid Prototyping Journal* 25.3 (2019), pp. 555–565. ISSN: 13552546. DOI: 10.1108/RPJ-03-2018-0062.
- [26] Chad Duty et al. “Z-Pinning approach for 3D printing mechanically isotropic materials”. In: *Additive Manufacturing* 27.March (2019), pp. 175–184. ISSN: 22148604. DOI: 10.1016/j.addma.2019.03.007.

- [27] David O. Kazmer and Austin Colon. “Injection printing: additive molding via shell material extrusion and filling”. In: *Additive Manufacturing* 36.May (2020). ISSN: 22148604. DOI: 10.1016/j.addma.2020.101469.
- [28] Mitchell Oliver Debra et al. *Joiners, Methods of Joining, and Related Systems for Additive Manufacturing*. 2020.
- [29] L.R. Lopes, A.F. Silva, and O.S. Carneiro. “Multi-material 3D printing: The relevance of materials affinity on the boundary interface performance”. In: *Additive Manufacturing* 23 (2018), pp. 45–52. ISSN: 2214-8604. DOI: <https://doi.org/10.1016/j.addma.2018.06.027>.
- [30] S Hüttenbach et al. “The interface between two strongly incompatible polymers: interfacial broadening and roughening near T_g”. In: *Langmuir* 7.11 (1991), pp. 2438–2442.
- [31] A N Gent and C-W Lin. “Model studies of the effect of surface roughness and mechanical interlocking on adhesion”. In: *The Journal of Adhesion* 32.2-3 (1990), pp. 113–125. ISSN: 15455823. DOI: 10.1080/00218469008030185.
- [32] Francesco Tamburrino, Serena Graziosi, and Monica Bordegoni. “The influence of slicing parameters on the multi-material adhesion mechanisms of FDM printed parts: an exploratory study”. In: *Virtual and Physical Prototyping* 14.4 (2019), pp. 316–332. ISSN: 17452767. DOI: 10.1080/17452759.2019.1607758.
- [33] Yuri Estrin, Vinayak R. Krishnamurthy, and Ergun Akleman. “Design of architected materials based on topological and geometrical interlocking”. In: *Journal of Materials Research and Technology* 15 (2021), pp. 1165–1178. ISSN: 2238-7854. DOI: <https://doi.org/10.1016/j.jmrt.2021.08.064>.
- [34] Y. Rahali et al. “Computation of the effective mechanical properties including non-classical moduli of 2.5D and 3D interlocks by micromechanical approaches”. In: *Composites Part B: Engineering* 98 (2016), pp. 194–212. ISSN: 1359-8368. DOI: <https://doi.org/10.1016/j.compositesb.2016.04.066>.
- [35] Yuan Yao et al. “A 3D weaving infill pattern for fused filament fabrication”. In: (2021), pp. 1–21.
- [36] I. A. Malik, M. Mirkhalaf, and F. Barthelat. “Bio-inspired “jigsaw”-like interlocking sutures: Modeling, optimization, 3D printing and testing”. In: *Journal of the Mechanics and Physics of Solids* 102 (2017), pp. 224–238. ISSN: 00225096. DOI: 10.1016/j.jmps.2017.03.003.
- [37] Yaning Li, Christine Ortiz, and Mary C. Boyce. “A generalized mechanical model for suture interfaces of arbitrary geometry”. In: *Journal of the Mechanics and Physics of Solids* 61.4 (2013), pp. 1144–1167. ISSN: 00225096. DOI: 10.1016/j.jmps.2012.10.004.
- [38] Micaela Ribeiro, Olga Sousa Carneiro, and Alexandre Ferreira da Silva. “Interface geometries in 3D multi-material prints by fused filament fabrication”. In: *Rapid Prototyping Journal* 25.1 (2019), pp. 38–46. ISSN: 13552546. DOI: 10.1108/RPJ-05-2017-0107.
- [39] Irfan Mustafa and Tsz Ho Kwok. “Development of Intertwined Infills to Improve Multi-Material Interfacial Bond Strength”. In: *Journal of Manufacturing Science and Engineering* (2021), pp. 1–10. DOI: 10.1115/1.4050490.

- [40] Wenzhi Wang et al. “Tensile behavior of bio-inspired hierarchical suture joint with uniform fractal interlocking design”. In: *Journal of the Mechanical Behavior of Biomedical Materials* 113.October 2020 (2021), p. 104137. ISSN: 18780180. DOI: 10.1016/j.jmbbm.2020.104137.
- [41] Lior Aharoni, Ido Bachelet, and Josephine V. Carstensen. “Topology optimization of rigid interlocking assemblies”. In: *Computers & Structures* 250 (2021), p. 106521. ISSN: 00457949. DOI: 10.1016/j.compstruc.2021.106521.
- [42] Regina M. Gouker et al. “Manufacturing of multi-material compliant mechanisms using multi-material molding”. In: *The international journal of advanced manufacturing technology* 30.11-12 (2006), pp. 1049–1075. ISSN: 02683768. DOI: 10.1007/s00170-005-0152-4.
- [43] Lars Rossing et al. “Bonding between silicones and thermoplastics using 3D printed mechanical interlocking”. In: *Materials and Design* 186 (2020), p. 108254. ISSN: 18734197. DOI: 10.1016/j.matdes.2019.108254.
- [44] Huong Quynh Dinh, Greg Turk, and Greg Slabaugh. “Reconstructing surfaces by volumetric regularization using radial basis functions”. In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* 24.10 (2002), pp. 1358–1371. ISSN: 01628828. DOI: 10.1109/TPAMI.2002.1039207.
- [45] Tim Kuipers. *Cura Arachne Engine beta*. 2021. DOI: 10.5281/zenodo.5106098.
- [46] Tim Kuipers et al. “A Framework for Adaptive Width Control of Dense Contour-Parallel Toolpaths in Fused Deposition Modeling”. In: *Computer-Aided Design* 128 (Nov. 2020), p. 102907. ISSN: 00104485. DOI: 10.1016/j.cad.2020.102907.

VI

APPLICATIONS

In the previous chapters I presented several techniques for toolpath generation, required for manufacturing functionally graded materials. Besides the FGMs already addressed, the proposed frameworks can also be employed for novel FGMs which were not the main focus of those chapters. This chapter discusses applications and techniques which are unlocked by the research presented in the previous chapters of this dissertation. These applications show how the previous chapters are closely linked together and they exemplify the impact of this dissertation on the field of fused filament fabrication of functionally graded materials.

Section VI.1 presents a technique to create a spatially graded surface roughness, which is based on the algorithmic framework presented in Chapter II. Section VI.2 presents a method to generate 2D images based on Chapter III and a technique to print these on textile using FFF. Section VI.3 presents a technique to imbue surfaces with a spatially graded thickness, which has two applications: spatially graded stiffness of sheet-based lattices and spatially graded light permeation for producing lithophane imagery on complex 3D models. This technique lends functionality from both Chapters II and IV. Lastly, Section VI.4 presents a technique for using the ITIL structure from Chapter V for spatially graded stiffness of dense parts.

VI.1 SPATIALLY GRADED ROUGHNESS

In Chapter II I presented a computational framework to translate a surface specification of colors into variable offsets on the layer outlines of layers with alternating color in order to reproduce the surface imagery of the specification. A similar framework can also be adopted for other surface property specifications such as roughness.

Spatially graded surface roughness has several interesting applications. It can be used to create more grip on a product in the areas where a user would hold it. It can also be used for its visual properties; by imbuing the surface with a rough texture it will appear more diffuse, which could be used to hide the layered buildup of FFF 3D printing.

An FFF product can be instilled a roughness by jittering the print head while printing the outer walls of each layer. Rather than following the original polygons of the contour of a layer, the print head can be made to follow a perturbed version of those polygons. The polygons are sampled at semi-regular intervals and those points are then offset with a random amplitude. The sample distance is chosen to be around twice the nozzle size so that the perturbations will not influence each other. Each sample step is chosen randomly between 0.8 mm and 1.2 mm so that the oscillation pattern will not align over several layers which could otherwise result in an unintended interference pattern. The amplitude of the offset at each sampled location is chosen to be within the range 0.0 mm to 0.4 mm and is determined by the amount of surface roughness required.

A possible alternative to using a surface roughness specification is to design the surface to have bumps in a modeling program. It should be noted that modeling such a surface texture would require a highly detailed model, which could require prohibitively much storage space. Moreover, the effect would not result in the same surface finish because the offsets would then be applied to the surface, rather than to the outer wall, which is at half the nozzle size inward from the surface at each layer. Jittering the movement of the print head directly gives more direct control over the jittering and thus over the final surface finish. By altering the toolpath generation pipeline for FFF 3D printing we can precisely control spatially graded surface roughness.

In order to translate the surface specification into jittered print head movement, I propose an adapted version of the computational framework presented in Chapter II. In contrast to the presented framework, the sampling of the polygons is performed at random intervals and the offset applied has a random magnitude within the prescribed amplitude. Each sampled point is offset in the normal direction by a distance d determined by the rgb value of the surface texture image at the sample point: $d = (2\text{rand}() - 1)(r + g + b)/3/255$. The resulting jittered polygons are then converted to Gcode, using the conventional toolpath processing pipeline in order to have the print head follow the jittered path.

Some inspirational examples are presented in Fig. VI.1. Figures VI.1a and VI.1b show applications where parts of the product are given more grip for a better manual adjustment, while other parts are kept smooth for dimensional accuracy or aesthetics respectively. Figure VI.1c presents an example of surface roughness being used artistically to convey the impression of different materials with different surface qualities. Figure VI.1d shows how a spatially graded roughness can be used to encode visual information on the surface. Finally, Fig. VI.1e presents an example where the two computational frameworks for surface properties have been combined for a reproduction of the authors face in both color and roughness. These applications demonstrate that an FGM for surface roughness is rather versatile.



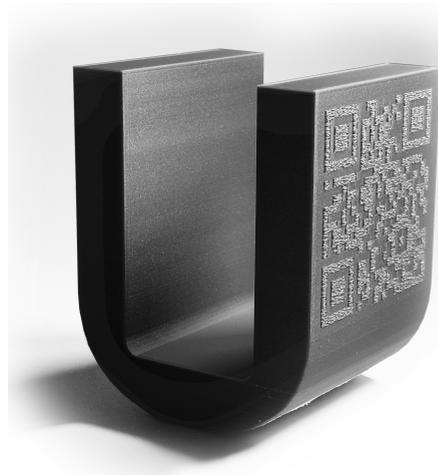
(a) Knob with grip and internal accuracy



(b) Handle with grip



(c) Artistic expression



(d) Visual information



(e) Gloss reproduction

Fig. VI.1. Applications of spatially graded surface roughness.

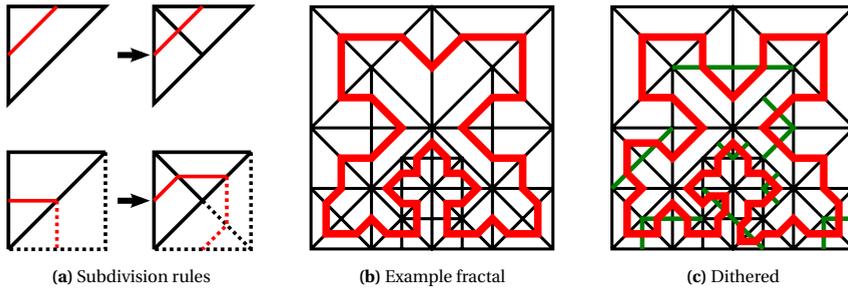


Fig. VI.2. Method summary of the 2D CrossFill variant. The area to be covered is divided into triangles (black), which are traversed by toolpaths (red). When a triangle is subdivided according to the rules in (a) the density of the traversing toolpath is increased by a discrete step. By dithering the subdivision levels such as in (c) the overall density in a region can be made to lie in between the discrete density levels, so that the brightness values of the input image are more accurately reproduced.

VI.2 SPATIALLY GRADED DENSITY IN 2D

Chapter III presented CrossFill: a 3D fractal surface which allows for a spatially graded density, while ensuring that each layer consists of a single continuous extrusion line. This section explores applications of a 2D variant of that fractal. By increasing the density of a black material on a white background in a particular region, that region will appear darker on average. The 2D CrossFill variant can as such be used to reproduce images. This section considers using FFF for producing such figures as well as inherently 2D production techniques, such as laser engraving and plotting.

The 2D variant can be derived by only considering the 2D subrules of 3D CrossFill, which are presented in Section III.4.2. Similarly, a 2D variant of the dithering approach is adopted: error is propagated linearly along the fractal in order to make local decisions on whether cells should be subdivided once more or not. In order to enforce toolpath continuity we adjust the subdivision rules for L and R, such that, when the hypotenuse edge of a cell is subdivided, its neighboring cell must also be subdivided; when subdividing, such pairs of cells are always considered jointly. The method is summarized in Fig. VI.2.

When using an input image as a density specification the 2D CrossFill fractal can be used to generate continuous toolpath figures. A standard FFF system can be used for printing such figures onto textile. The textile is spanned onto the build plate using clips. See Fig. VI.3a.

By letting the first layer permeate into the woven textile material it will form an interlocking structure analogous to the lattice presented in Chapter V; the interlock ensures that the print bonds to the textile throughout its eventual use. In order to get a good permeation, the material flow percentage and the initial layer height are tweaked for each textile separately, since each textile may have a different density and thickness. A small patch on the back may be used to find the optimal settings for each textile material. See Figs. VI.3b and VI.3c.

By printing with flexible material we ensure that the print is able to deform along with the fabric when it is worn or washed. Figure VI.4a shows several examples of FFF prints of TPU onto T-shirts. These prints have been worn and washed several times over the past years and show little to no wear.

2D fractal figures can also be fabricated using other production techniques. Fig-

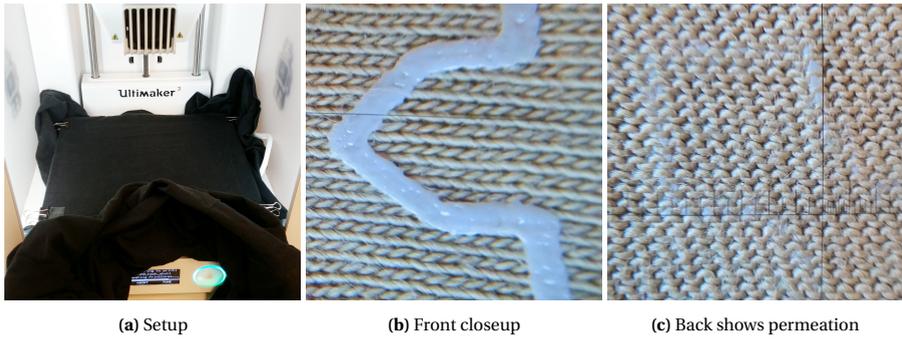
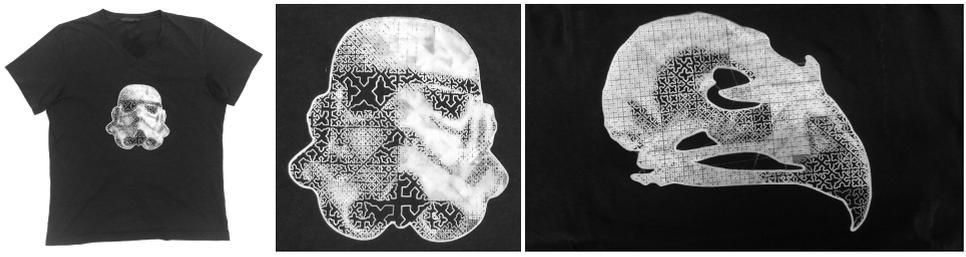
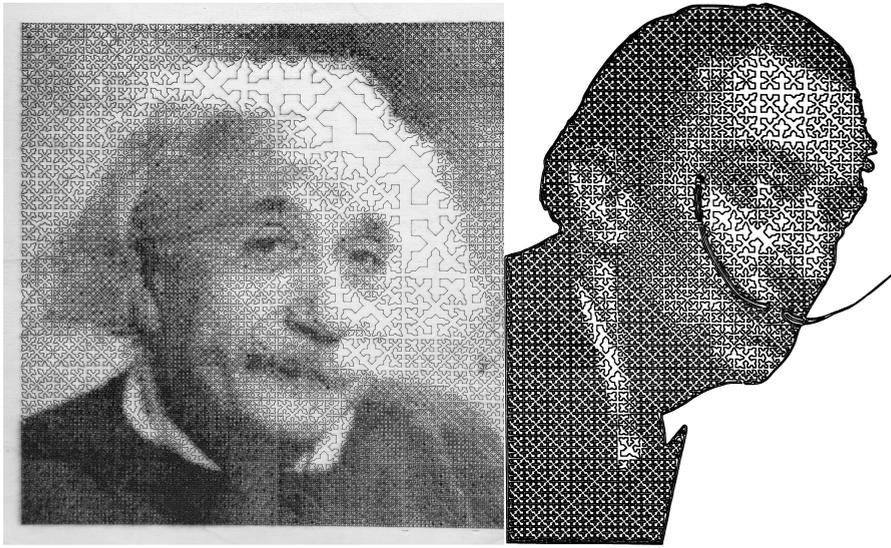


Fig. VI.3. 3D printing onto textile.



(a) 3D printed onto textile.



(b) Laser engraved in wood.

(c) Plotted on paper.

Fig. VI.4. 2D CrossFill applications.

ures VI.4b and VI.4c shows examples of images produced using a laser engraver and a plotter.

VI.3 SPATIALLY GRADED THICKNESS

In Chapter IV I introduced a method to print thin geometry of varying width using varying width extrusion lines. As mentioned in that chapter, one application of the method is in microstructures with spatially graded thickness. By making a repeating lattice thicker in some regions it will become more stiff there. This application area is presented in Section VI.3.1. Another application area is concerned with using spatially graded thickness in order to achieve spatially graded visual properties, which is presented in Section VI.3.2.

VI.3.1. SPATIALLY GRADED STIFFNESS

One commonly used microstructure is a triply periodic minimal surface: the gyroid. In contrast to the microstructure presented in Chapter III, the gyroid behaves less like a foam, and more like a spring; the further you compress the structure, the more force is required.

Strictly speaking the gyroid surface violates the overhang constraints. We can preclude any manufacturing problems by adjusting the surface patches of the repeating structure: the violating regions are replaced by faces at the overhang angle, and perfectly horizontal faces which can be bridged during manufacturing. The remaining patches are then lofted to connect with the adjusted patches. However, in practice the violation of overhang constraints in the standard gyroid is constrained to such a small area that the print defects are negligible. See Fig. VI.5.

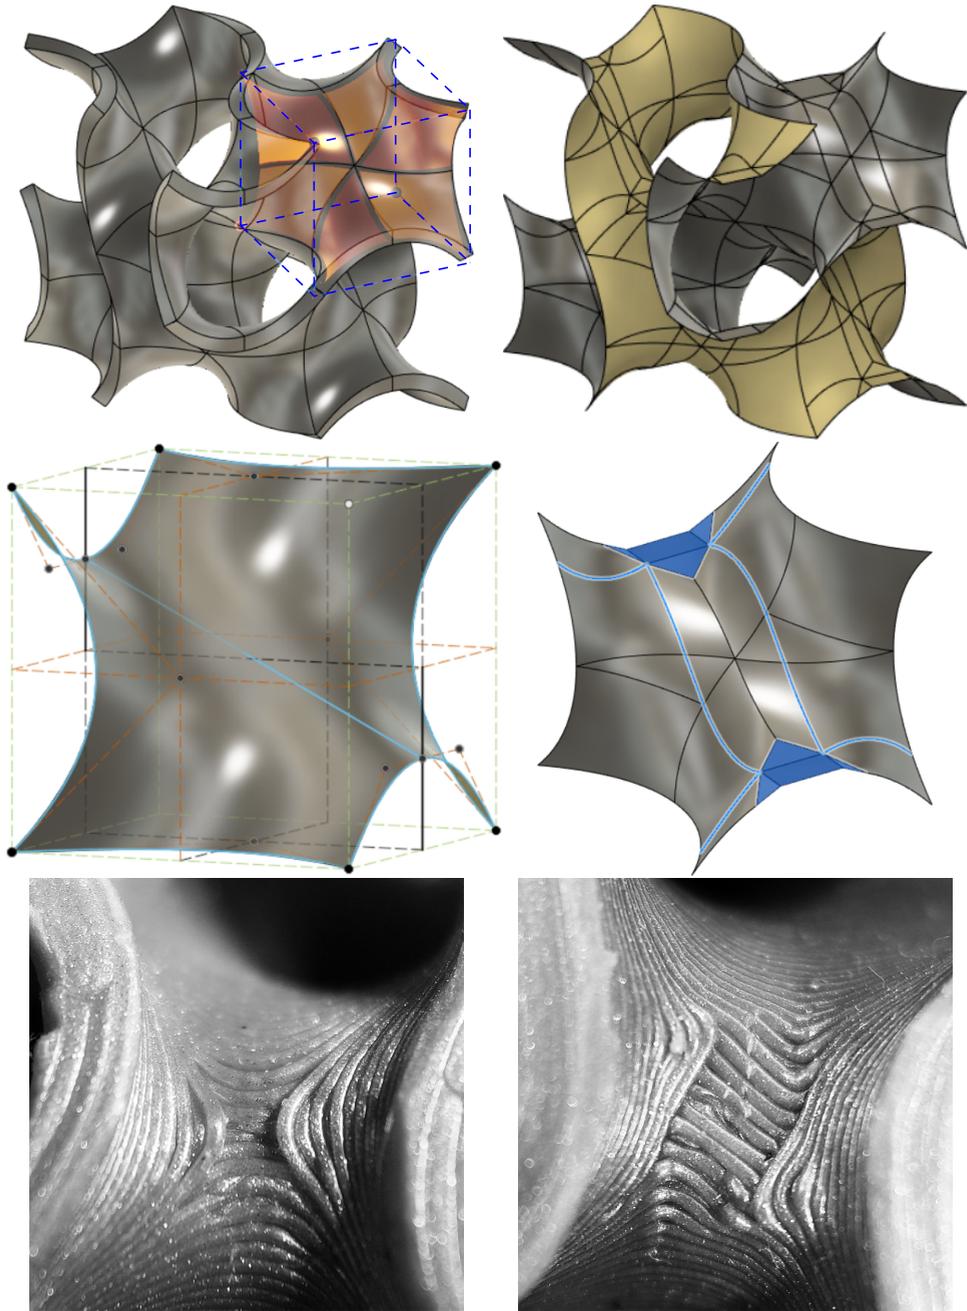
A spatially graded thickness sheet-lattice can be generated by applying a spatially graded offset to the vertices of its mesh. The thickness specification for a sheet-lattice can be supplied by surface coloring. Similar to the method presented in Chapter II, we can use a textured model and interpret the surface grayscale tones as thickness values. The surface mesh is duplicated and flipped in order to convert the surface mesh of the sheet-based lattice into a volumetric sheet-based lattice. Each vertex is then displaced along its normal direction according to the thickness dictated by the texture brightness at that location.

Using such an algorithmic framework we can generate sheet-based lattices with functionally graded thickness, such as depicted in Fig. VI.6a. These lattices can be used to manufacture personalized footwear. Figure VI.6b shows an example of a shoe sole where the thickness of the lattice is determined by pressure data of a foot standing on a pressure pad.

VI.3.2. SPATIALLY GRADED TRANSLUCENCY

A different application area of spatially graded thickness of sheet-like objects can be found in lithophanes. Lithophanes are objects which exhibit imagery when light is shone through them. By making the surface thicker in some locations we can make it appear darker there. For this effect a semi-translucent material is required, so that modulating the amount of light which is transmitted through the material can be modulated using a reasonable range of thicknesses. One such material is Ultimaker PLA pearl white, which has a transmittance of $\tau \approx 52\%$ at a reference thickness of 1 mm.

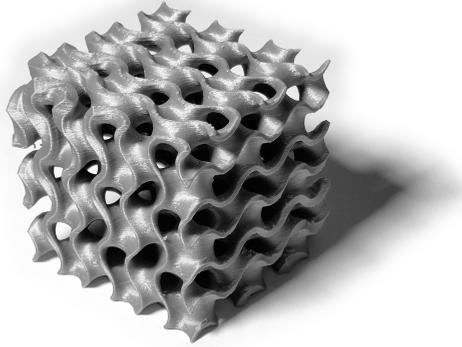
Typically lithophanes are printed flat like the image they are portraying, though several online tools are available to automatically generate lithophanes on a cylinder, cube



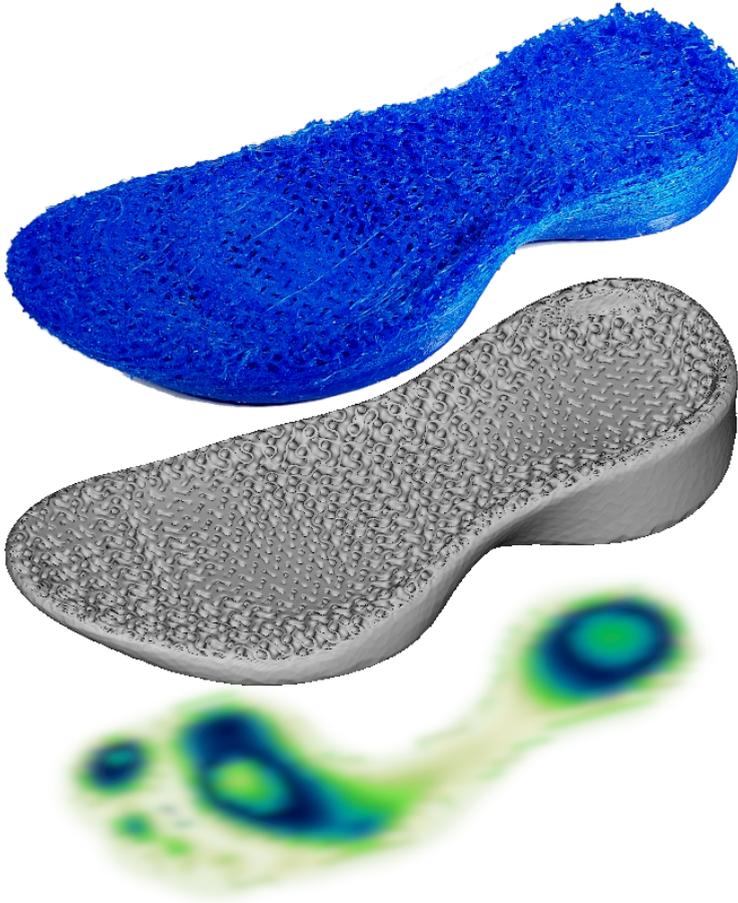
(a) Standard gyroid.

(b) Overhang adjusted gyroid.

Fig. VI.5. Repeating cell, surface patch and printed result of the standard gyroid and the overhang adjusted gyroid cell. While the standard gyroid strictly violates overhang constraints, the induced defects are often minimal.



(a) Lattice.



(b) Shoe sole.

Fig. VI.6. Sheet-based gyroid lattice with spatially graded thickness.

or sphere. The algorithmic framework for thickening mentioned in Section VI.3.1 can be extended to convert textured 3D models into 3D lithophanes. The offsets are only applied to the internal shell, so that the external geometry remains unaffected.

The length of the offset on a vertex is based on both the texture color of the vertex and the amount of light reaching the connected faces, i.e. their *illuminance*. The illuminance can be calculated in Blender Cycles by baking the diffuse lighting into a second texture map. The rgb pixel values of the two texture maps are converted into brightness values using the Y value of the CIE XYZ color space (Eq. (VI.1)), which gives us the required output luminance L_{req} and the illuminance L_{in} . Equation (VI.2) provides a model for how much light is transmitted at a vertex v , based on the local thickness $w(v)$. The required output luminance is then scaled to be within the range of possible output luminances (Eq. (VI.3)) and is used to calculate the thickness w of the shell at each vertex v in Eq. (VI.4).

$$L = 0.21r^{2.2} + 0.72g^{2.2} + 0.07b^{2.2} \quad (\text{VI.1})$$

$$L_{\text{out}}(v) = \tau^{w(v)} L_{\text{in}}(v) \quad (\text{VI.2})$$

$$L_{\text{max}} = \tau^{w_{\text{min}}} \min_v L_{\text{in}}(v)$$

$$L_{\text{min}} = \tau^{w_{\text{max}}} \max_v L_{\text{in}}(v)$$

$$\hat{L}_{\text{req}}(v) = L_{\text{min}} + L_{\text{req}}(v)(L_{\text{max}} - L_{\text{min}}) \quad (\text{VI.3})$$

$$w(v) = \log_{\tau} \frac{\hat{L}_{\text{req}}(v)}{L_{\text{in}}(v)} \quad (\text{VI.4})$$

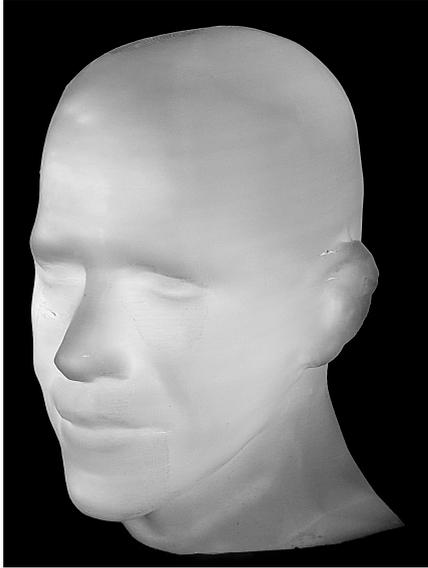
Although these formulae simplify the underlying mechanics of light, they can already be used to generate 3D lithophanes with appealing imagery. Figure VI.7 presents several lithophanes printed with Pearl white PLA and soluble PVA support material. The toolpaths for these jobs are generated using the computational framework presented in Chapter IV.

VI.4 SPATIALLY GRADED DUAL MATERIAL PROPORTIONS

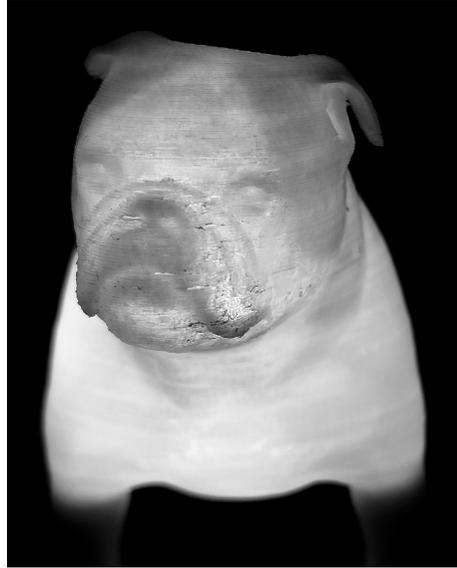
The previously presented FGMs concern either surface properties or sheetlike structures. Another class of FGMs manufacturable on FFF systems are voluminous FGMs. Rather than being made up of material and void, voluminous FGMs can be defined which consist of two materials which fill the entire volume of a product. As such voluminous FGMs have an advantage in applications where the final product needs to be watertight, needs to sink or needs to perform well in vacuum.

By using two materials the overall properties of the product can be made to lie somewhere in between the properties of the two constituent materials. By grading the material proportions we get a spatial gradient in multiple material properties simultaneously.

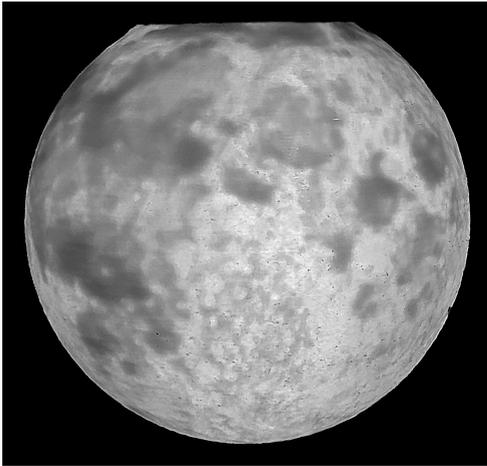
One approach to define voluminous FGMs is using a TPMS where either side of the surface is printed with a different material. TPMS like the gyroid are defined by the isosurface of a scalar field defined by a simple trigonometric function. By adjusting the isovalue, the surface will morph to something more similar to a beam lattice, which shifts the material ratio. A spatially graded material ratio can be achieved by using a spatially graded isovalue for the TPMS equation. See Figs. VI.8a to VI.8c. However, the two materials can still separate at the isosurface interface, which makes the material properties complex to simulate and hard to control.



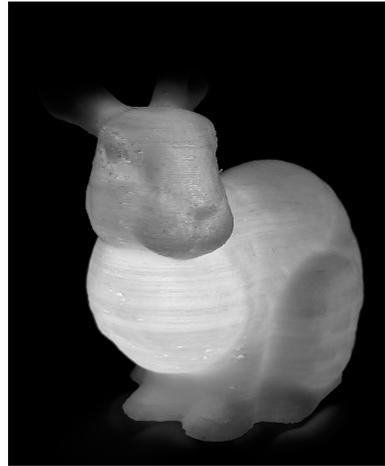
(a) Portrait.



(b) Pet.



(c) Moon.



(d) Stanford bunny.

Fig. VI.7. 3D lithophanes create surface imagery using spatially graded thickness shells which partially transmit light from an internal light source.

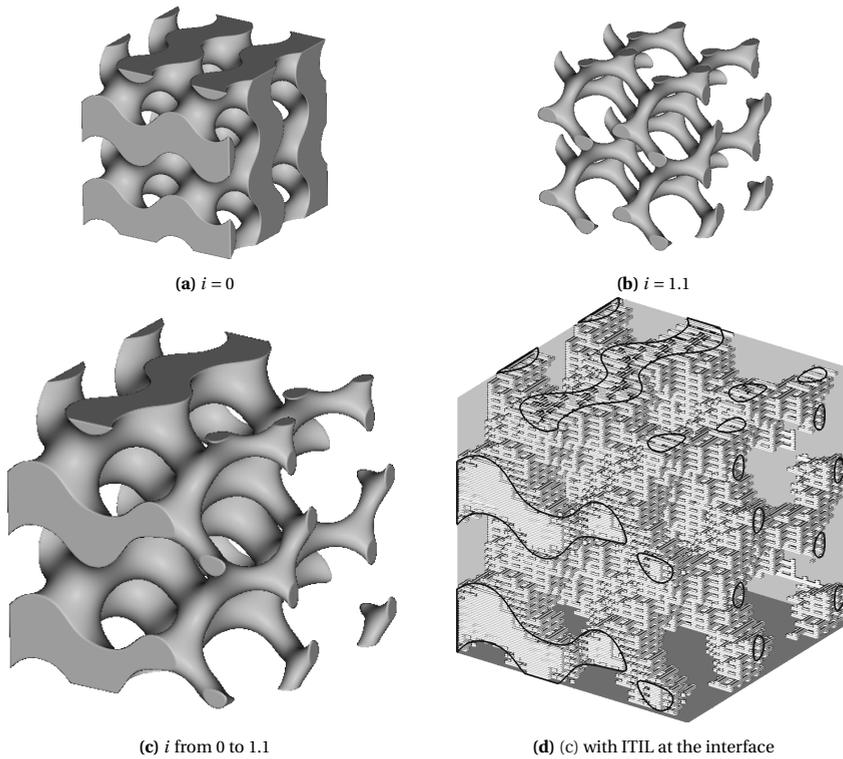


Fig. VI.8. Various isovalues i of the gyroid field, which is given by: $\sin x \cos y + \sin y \cos z + \sin z \cos x = i$

In order to counteract the separation issue, we might apply the interlocking lattice presented in Chapter V at the isosurface interface. See Fig. VI.8d. While the interlocking lattice ensures the bonding at a fine scale, the meso-scale geometry of the TPMS can be used to regulate the spatial gradation in material properties. However, the changes that the ITIL structure makes to the interface geometry can significantly affect the mechanical behavior of the TPMS. For the material properties to be less affected by the ITIL structure the TPMS would need to be relatively large, which can become problematic for products with a relatively small design.

Instead, the ITIL lattice can also directly be used to generate FGMs without introducing a meso-scale geometry. By creating a spatial gradient in the width of the beams the proportional density of the two materials can be spatially graded. Meanwhile the beams are still horizontal, which guarantees semi-continuous extrusion. The connectivity of the beams means that either material is still interconnected and eventually to the build plate, which aids the manufacturability of the FGM.

Such spatially graded lattices can be used to control the spatial gradation of stiffness in the object. This has application in the field of soft robotics. By using TPU as a flexible material and PP as a less flexible material it is possible to tune the deformation behavior of a soft body. For example, the deformation of the tail of a mechanical fish can be optimized for better propulsion, while the fins remain more stiff for better stability. See Fig. VI.9a.

The beam widths can also be modulated for surface properties such as color, when

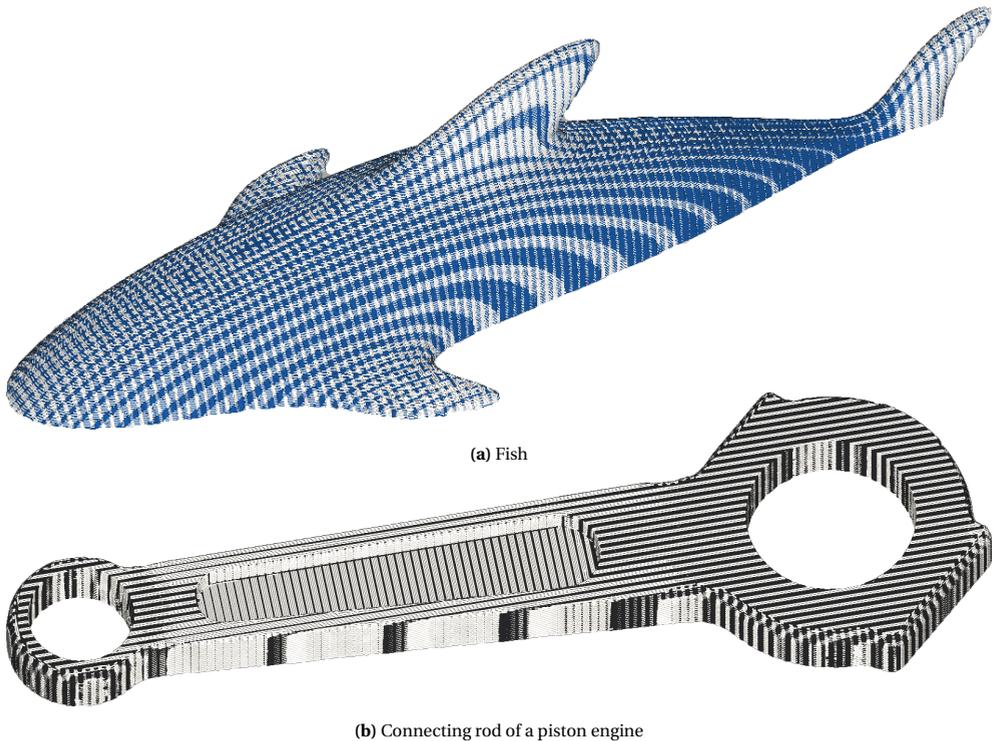
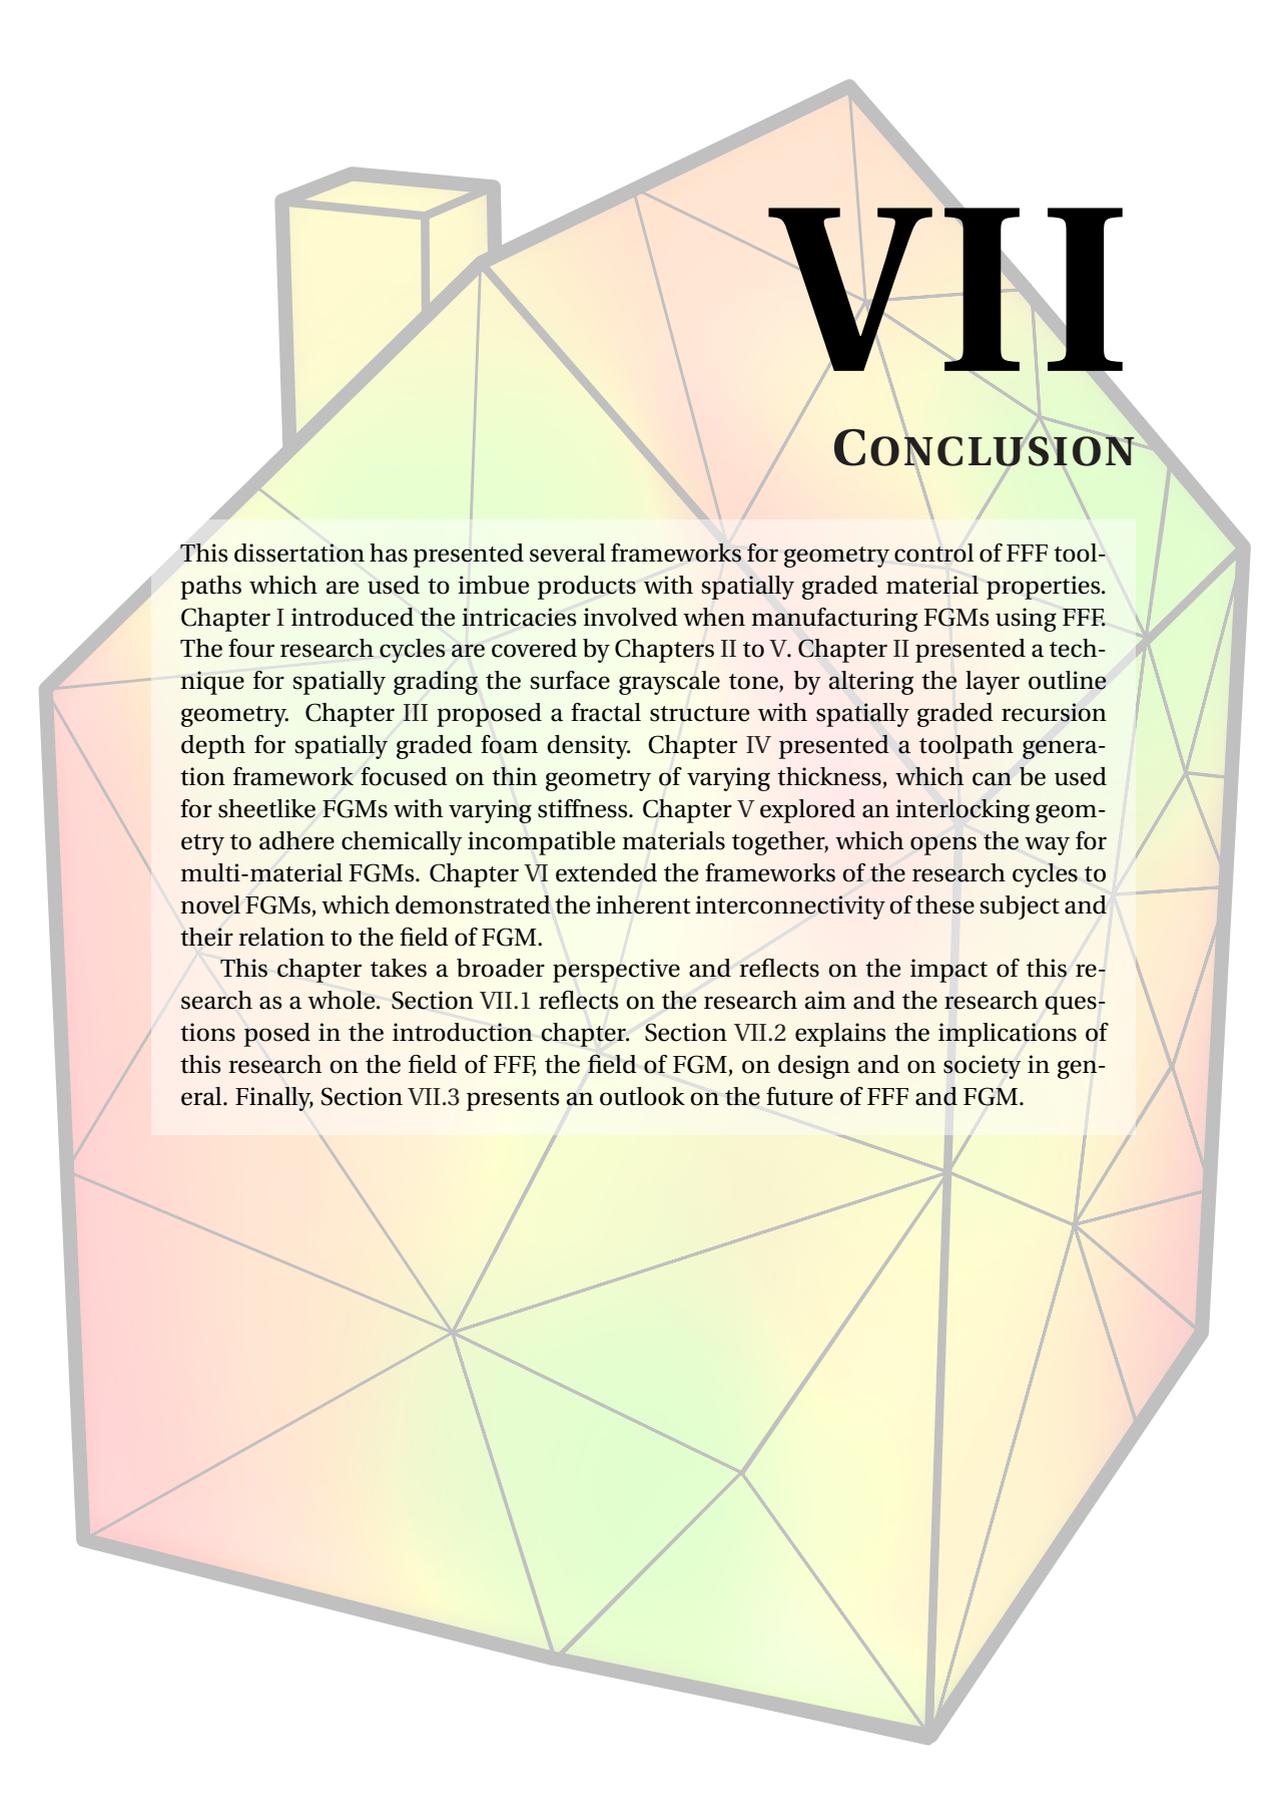


Fig. VI.9. Applications of spatially graded ITIL.

using two materials with a dark and a light color respectively. Where beams of the former material are wider than of the latter, the surface appears darker. See Fig. VI.9b.

For narrow beams, where the width is modulated in a range close to the nozzle size, the toolpath generation framework presented in Chapter IV generates a single extrusion line with varying width. In order to actualize the beam width variation the print head adjusts the speed according to the back pressure compensation method described in Section IV.5.1. The back pressure compensation technique can be seen as an extension of the flow equalization technique presented in Section II.3.3. In fact, the spatially graded ITIL itself can be seen as a type of horizontal hatching, but for both materials and reiterated over all the layers.

The commonalities between the applications presented in this chapter show that the work presented in this dissertation impacts a larger part of the FGM research area than the individual parts addressed in the previous chapters. They furthermore show how the previous chapters are linked together in ways which have not been addressed in the chapters individually. The fact that the research chapters unlock novel FGMs illustrates the impact of this dissertation on the field of FGM research.



VII

CONCLUSION

This dissertation has presented several frameworks for geometry control of FFF toolpaths which are used to imbue products with spatially graded material properties. Chapter I introduced the intricacies involved when manufacturing FGMs using FFF. The four research cycles are covered by Chapters II to V. Chapter II presented a technique for spatially grading the surface grayscale tone, by altering the layer outline geometry. Chapter III proposed a fractal structure with spatially graded recursion depth for spatially graded foam density. Chapter IV presented a toolpath generation framework focused on thin geometry of varying thickness, which can be used for sheetlike FGMs with varying stiffness. Chapter V explored an interlocking geometry to adhere chemically incompatible materials together, which opens the way for multi-material FGMs. Chapter VI extended the frameworks of the research cycles to novel FGMs, which demonstrated the inherent interconnectivity of these subject and their relation to the field of FGM.

This chapter takes a broader perspective and reflects on the impact of this research as a whole. Section VII.1 reflects on the research aim and the research questions posed in the introduction chapter. Section VII.2 explains the implications of this research on the field of FFF, the field of FGM, on design and on society in general. Finally, Section VII.3 presents an outlook on the future of FFF and FGM.

VII.1 RESEARCH OBJECTIVE REALIZATION

The main research question of this dissertation is: “How can we generate instructions for FFF 3D printers which enable the manufacturing of FGMs?” While this question is quite broad, I have limited myself to geometry control for toolpath generation of unmodified desktop FFF systems, as described in Chapter I.

This dissertation’s focus on FFF presupposes that this manufacturing technology must be taken into account when producing FGMs. While FGM is a challenging field in itself, using FFF to manufacture them poses additional challenges. In order to avoid print failures and defects, the geometry has to comply with several manufacturing constraints: resolution, non-overhang, non-overlap, continuity, integer thickness and chemical compatibility.

The research question of this dissertation has been answered in two ways: novel FGMs have been proposed which adhere to the manufacturing constraints (Chapters II, III and VI), and toolpath generation techniques have been proposed in order to alleviate some of the manufacturing constraints which inhibit FGMs (Chapters IV and V).

In Chapter I FGMs were categorized into surface properties and volumetric properties. The surface properties which were addressed in this dissertation are (grayscale) color, roughness and translucency. While this is a subset of the range of possible surface properties, it is unclear how geometric control of toolpaths can be exploited to realize other surface properties. Perhaps other types of process control can be exploited to manufacture other types of surface FGMs, but they lie beyond the scope of this research.

The volumetric FGMs were further subdivided according to the dimensionality of their elements. 0D voxel based FGMs and 1D beam lattice FGMs are not manufacturable on FFF systems, because their cross-sections at each layer would contain many isolated small islands, which violate the continuity constraint. The remainder is 2D sheetlike lattices and 3D voluminous FGMs, which are both addressed in this dissertation. Chapters III and IV relate to sheetlike lattices, while the framework of Chapter V is used in Section VI.4 to generate voluminous FGMs. There remains a vast range of sheetlike lattices and voluminous FGMs which have not been specifically addressed in this dissertation, since the number of possible geometries is innumerable.

The volumetric properties addressed in this dissertation are stiffness, density and plateau stress of foams. However, the presented FGMs have different values for other mechanical properties as well. By adjusting the recursion depth or sheet thickness of these structure, those other properties could be controlled as well. However, not all mechanical properties can be controlled by the proposed volumetric FGMs. For example, generating FGMs which have different mechanical behavior in different directions, i.e. *orthotropic* properties, are an interesting subject for further research.

VII.2 IMPLICATIONS

While the focus of this dissertation has been on geometry control and toolpath generation for FFF 3D printers to enable FGMs, the research has significance beyond this narrowly defined subject. This section discusses the implications of this research on AM, the field of FGM, and taking a broader perspective on design and on society in general.

All algorithmic frameworks discussed in this dissertation are open source available at github.com/Ultimaker and github.com/Bage10rb. The variable width toolpath generation framework of Chapter IV is also open source available, but it is patent pending: Tim

Kuipers. “Method of determining toolpaths for an infill structure for a digital 3d model”. Pat. req. Ultimaker B.V. July 15, 2021. URL: <https://patentscope.wipo.int/search/en/detail.jsf?docId=W02021139928>.

VII.2.1. IMPLICATIONS FOR ADDITIVE MANUFACTURING

This research has implications for additive AM in several ways. First, I will discuss how toolpath generation software could be extended to process FGM material specifications. Second, I will discuss the impact of this research on toolpath generation software in general. Third, I will discuss the implications of this research for FFF process control. Lastly, I will discuss the applicability of the proposed techniques on AM techniques other than FFF.

FGM SPECIFICATIONS

The algorithmic frameworks presented in Chapters II and III and Section VI.1 are modifications of the standard toolpath generation process. While these frameworks provide for the toolpath generation framework, the slicing software ecosystem would need further adjustments in order for full support of these FGMs. The spatially graded material specifications would need to be transferred to the slicing software and they would need to be displayed for the user to meaningfully interact with the software.

For surface property specifications there is a readily usable interchange format called Wavefront OBJ. OBJ files are commonly used to store surface image texture information, which can be reinterpreted as values for a certain material property. A possible extension of this format could address the conversion of an RGB pixel value to a value for a specific material property. That way the material specification is standardized so that design programs and toolpath generation software are communicating in mutually understood terms.

Volumetric material property specifications are less standardized and can come in many forms. Material specification data can be provided in a grid of voxels, which can be stored as a series of images, but many other less commonly used file formats are available as well. Another way of specifying material properties is at the vertices of a volumetric mesh consisting of tetrahedra of varying shape. Such volumetric specifications are being addressed in a novel file format called 3MF, which aims to become the new standard for 3D printing.

Besides an infrastructure for exchanging material specification data between design and toolpath generation software, it is also important to show such specifications to the user visually. While surface properties can naturally be drawn on the surface, displaying volumetric properties is less straightforward. Various design software packages tackle the problem in various ways: display the properties on the surface, on a section plane only, or at dots on a grid. However, most slicing software does not yet allow for such visualization. Toolpath generation software packages should extend their frontend functionality in order to support FGMs properly.

SLICING SOFTWARE FEATURES

Although the features of visualization and import/export have not yet been fully supported in the Ultimaker Cura slicing software package, this research has led to several other features that have been incorporated into that toolpath generation program.

The interlocking lattice presented in Chapter V is expected to be a feature of Cura 5.1. In order to deal with arbitrary complex interfaces between two bodies of a different material, a regular lattice grid is generated and the cells which overlap with the interface are used to generate the interlocking lattice only where it is needed. See Fig. V.17b. This allows designers to create multi-material parts with ease.

Cura has functionality to automatically convert image files to flat lithophanes. Although the 3D lithophane technique described in Section VI.3.2 have not been merged into Cura, the underlying light permeation model described in Eqs. (VI.1) to (VI.4) was used to reproduce those images more accurately. Likewise the surface roughness technique described in Section VI.1 is a part of Cura, but the spatial gradation is not yet available.

The CrossFill lattice presented in Chapter III is also merged to Cura and is the default infill pattern when printing with flexible materials. Printing with semi-flexible materials is now more feasible and the resulting product is less stiff than the old default infill structures. Some limited form of spatial gradation is possible when using the Cross infill pattern in Cura: there is a setting field where the user can supply a density image specification, but since it is only a single image, the specification is only 2D. Furthermore, the 2D CrossFill method used to print images on textiles presented in Section VI.2 is available to Cura users.

Chapter III also described a method to reconnect the polygons of the CrossFill pattern to the walls of a layer, which is implemented in Cura. See Fig. III.17. The Connect Infill Polygons feature can be used to ensure continuous extrusion throughout the infill when a polygonal infill pattern is selected. Better yet, any infill pattern can be converted into a polygonal infill pattern. By performing a boolean difference between the infill area and the thickened infill polylines, a polygonal infill pattern is generated, though it has twice the thickness of the original lines. A Cura user can set the Infill Line Multiplier to 2 and enable Connect Infill Polygons to ensure strict continuous extrusion for any infill structure. This enhances the reliability for printing especially flexible materials, while keeping a wide range of infill pattern choices open to the user.

Lastly the toolpath generation framework described in Chapter IV has become the new standard for generating the walls in Cura. The adoption of the framework is the reason for advancing the version number to Cura 5.0. The impact is not only limited to the accuracy and reliability with which thin geometry is printed; compared to the old framework used by Cura the new framework also produces faster prints and uses less material over all. With Cura 5.0 users can disregard the integer thickness constraint and focus on requirements of the design itself. Designers are now able to create variable thickness designs in the design software without having to consider this manufacturing constraint. Variable thickness sheet-based lattices do not need to be generated in the slicer – they can be modeled in design software.

PROCESS CONTROL

Besides features for geometry control of the toolpath generation pipeline, this research also has implications for FFF process control. Some improvements to the FFF printing process are required in order to manufacture some of the proposed FGMs.

The grayscale printing technique proposed in Chapter II required a fine calibration of the positioning system. If the calibration were off in the horizontal plane then one side of the object would come out lighter and the other darker. If the calibration were off in the Z direction then the whole product might come out too dark or too light. In order for

the product to show the correct grayscale values the offset between the nozzles had to be calibrated more precisely than the standard calibration procedure.

Another area where precision needed to be improved, is where extrusion lines of varying width are concerned. Chapter II proposes a scheme for adjusting the nozzle movement speed inversely with the line width so as to equalize the extrusion flow in mm^3/s . That flow equalization scheme was then further improved in Chapter IV to account for the extra pressure needed to squeeze out wider lines. With this adjusted process control, varying width lines could be laid down more accurately, which helps to manufacture varying thickness objects more accurately and to manufacture spatially graded grayscale tone products more accurately.

The extended flow equalization feature has therefore been adopted in Cura 5.0. Furthermore, a paper I co-authored used this process control feature to accurately manufacture varying width toolpaths generated using a wholly different toolpath generation framework than the one presented in Chapter IV: Samuel Hornus et al. “Variable-width contouring for additive manufacturing”. In: *ACM Transactions on Graphics* 39.4 (July 2020). ISSN: 0730-0301. DOI: 10.1145/3386569.3392448. This framework utilizes a wide range of extrusion widths in order to define toolpaths which exhibit a minimal amount of micro-gaps, so that we are able to manufacture more dense and therefore stronger prints. Although not related directly to FGMs, this framework is closely related to the research presented in this dissertation in terms of toolpath geometry control.

VII.2.2. APPLICATIONS TO OTHER AM TECHNOLOGIES

This dissertation was primarily focused on FFF, while other additive manufacturing technologies were not discussed extensively. Some of the manufacturing constraints are uncommon in other AM technologies, such as the continuous extrusion constraint. It is hard to imagine that the multi-material interlocking lattice from Chapter V is as relevant to manufacturing techniques different from FFF. Also the hatching technique presented in Chapter II depends heavily on the process specifics of FFF, so I do not expect this technique to find application in other AM technology. Since the goal of this research project was to enable FGMs for FFF it stands to reason that the proposed techniques apply mostly to FFF specifically.

However, we have seen several techniques which can be considered for other manufacturing processes as well. For example, the 3D lithophanes proposed in Section VI.3.2 could be manufacturing using any manufacturing process which can produce semi-translucent products. Also the 2D CrossFill figures discussed in Section VI.2 had application for technology other than FFF: laser engravers and pen plotters. In fact the fractal structure of CrossFill from Chapter III could also be produced by technology which does not require continuous extrusion or self-supporting geometry; its mechanical behavior similar to foams could prove to be valuable also for manufacturing processes.

Some of the proposed techniques impact a wider range of additive manufacturing technologies, namely those which operate using a trace; material extrusion technologies extrude material in strands which follow a trace, stereolithography uses a laser to cure a photopolymer along a trace and powder bed fusion technologies have a laser tracing a path along which material is bound to the powder below. In as much as such manufacturing technologies produce a constant width trace that produces defects when overlapped with existing traces, they are subject to the integer thickness manufacturing constraint. If a control mechanism were developed to adjust the trace width by e.g. dynamically chang-

ing the power of the laser, that constraint might be overcome. The variable width toolpath generation framework presented in Chapter IV could help these technologies to improve the quality of thin products.

VII.2.3. IMPLICATIONS FOR FGM

This dissertation has shown that FGMs are manufacturable on FFF 3D printing systems. Since FFF is relatively cheap and versatile compared to other additive manufacturing technologies, this research has made FGMs more accessible to researchers, professionals and hobbyists. FFF opens up a wide range of materials to manufacture FGMs at low cost.

The algorithmic framework presented in Chapter II unlocks FFF specific surface FGMs. Section VI.1 shows that the framework can easily be extended to other surface properties which can be tweaked using toolpath geometry, such as surface roughness. However, the number of surface properties which can be tweaked by FFF toolpath properties is limited. For other surface properties different additive manufacturing techniques might be more appropriate.

The toolpath generation framework presented in Chapter IV unlocks a family of FGMs; sheet-based lattices can be imbued a spatially graded thickness as described in Section VI.3.1. With the enhanced toolpath generation framework such lattices can now be printed without print defects or failures. The research therefore unlocks level-set approaches and spatially graded offset approaches for sheet-based FGMs.

For some materials continuous extrusion is still a requirement, as we have seen in Chapter III. For such materials using the toolpath generation framework of Chapter IV on a varying thickness sheet lattice results in discontinuous paths. Some part of the research field of FFF-produced FGMs will therefore still be concerned with toolpath generation rather than only with geometry control.

Likewise the overhang constraint still opens up a research gap for FFF-based FGMs. The overhang constraint played a role in determining the geometry of the cross fractal of Chapter III. However, we have seen in Section VI.3.1 that small patches of overhang lead to negligible print defects. Nonetheless, for larger patches it needs to be taken into account when generating sparse FGMs. Future research might be dedicated to generating other types of FGM which adhere to the overhang constraint.

This dissertation opens the way to multi-material FGMs manufactured using FFF. The method described in Chapter V shows how bodies of different and even chemically incompatible materials can be joined together. By securing the bonding between two bodies mechanically, the chemical compatibility manufacturing constraint is alleviated. The proposed structure satisfies the continuous extrusion constraint and can be applied to interfaces in any direction. As such it can be used to manufacture multi-scale multi-material FGMs. At the micro-scale the ITIL structure helps to satisfy manufacturing constraints, while at a meso-scale the designed geometry can be tweaked to spatially vary the material properties of the product at the macro-scale. However, multi-material FFF and the ITIL structure can still be improved upon in order to get more reliable multi-material prints with better adhesion.

VII.2.4. IMPLICATIONS FOR DESIGN

With the advance in FGMs and their adoption in the cheap and versatile manufacturing technology of FFF, the focus of designers will shift more toward the mechanical properties at various locations in the product. Designers could work more with spatial data in order

to design FGMs which are tailored to the consumer or optimized for its operation. For example, the design of medical devices, footwear, grippers and (re)production of spatial information could be assisted by the design of FGMs.

Algorithmic frameworks can be used to automatically convert mechanical property specification into FGMs; the task of the designer is then to design the FGM property specification of the final product. Surface property specifications can simply be painted onto 3D models using Blender, Photoshop, Maya, etc. A spatially graded shell thickness of products or lattices can be designed in programs like nTopology. Volumetric density distributions in the form of voxel grids can be designed in programs like Monolith.

Designing products with FGMs means that products can be manufactured using less components and fewer different materials, because a single material component can have different properties in different regions. FGMs simplify the design process by relaying the product requirements to the algorithmic frameworks which automatically determine how those requirements should be translated in the final product. The designer is therefore partially relieved of the task of dividing up a product into parts and designing how those parts should interact.

How multiple parts should interact is further simplified by the interlocking structure presented in Chapter V. Different parts composed of different materials can now be printed simultaneously, without the designer having to worry about how the parts will adhere together. The adhesion is guaranteed by the automatically generated interlocking structure at the interface between the bodies. The designer is therefore relieved of designing interlocking mechanisms themselves.

The practice of design for additive manufacturing is simplified, because the research has alleviated design constraints. The manufacturing constraint of integer thickness is resolved, which means that designers will no longer have to take that into account when optimizing a design for printability on FFF machines. The chemical compatibility manufacturing constraint is also alleviated, which means designers have a wider range of material combinations to choose from.

FGMs are commonly produced using additive manufacturing techniques, which means that the prototyping technique can be the same as the final production method. Using FFF for that prototyping technique reduces prototyping costs and lowers the bar for materialization, which can speed up the prototyping process. Furthermore, because the final product is produced using the same technique as prototyping, there is no need to adjust the design for traditional manufacturing methods. When developing a product with FGMs for FFF the total product development process is shortened.

FGMs broaden the spectrum of product properties for the designer to control. With FGMs designers have a vast array of material properties which could be controlled throughout the product. The fact that designers will have more tools at their disposal means that there are more considerations to be made. A more sophisticated judgment is required to choose which product properties to control, since FGMs add a range of material properties to the designer's toolbox.

Specifically, we can design surface color for artistic purposes or to display simulation results. We can design surface roughness for dexterous manipulation or to convey information. We can design foam density for optimized packaging, seat comfort or deformation behavior. We can design 2D images to be reproduced on textile or paper. We can design regular lattices with tailored mechanical properties and determine a spatial gradient in stiffness on top of that. We can design multi-material parts with living hinges and

soft robotic actuators.

VII.2.5. IMPLICATIONS FOR SOCIETY

This dissertation is concerned with FFF of FGMs. It aims to impact both the disciplines of FFF and FGM as a whole. Both these fields have their place on society and furthering these fields benefits society in different ways.

Additive manufacturing, and FFF specifically, can be used for localized production. Production can be distributed throughout multiple locations where a 3D printer is available, which can reduce the need for spare parts and shipping of parts across the globe. FFF therefore contributes to a reduction in shipping costs, in fuel and in material waste, so in that respect it is good for the environment.

However, FFF typically uses plastics and is not very energy-efficient. FFF requires plenty of energy to keep the nozzle and build plate heated, and moving the print head around also costs energy. Conventional manufacturing processes typically use considerably less energy per part. In that respect FFF causes energy waste and material pollution.

The balance between the positive and negative impacts of FFF on the environment has to be made up in a plethora of contexts and use-cases. Many other factors of the manufacturing system and the printing material also influence the environmental impact. Future developments in additive manufacturing technology could tip the balance toward either side. Evaluating the positive effects of additive manufacturing versus conventional manufacturing methods is a difficult quest which falls beyond the scope of the current study.

While conventional manufacturing technology is often used for mass production, FFF is often used for low volume production. Low volume production has various applications: specialized end-products, personalized products, tooling, jigs and fixtures, etc. These small quantity production applications benefit from the advances in FFF that this research provides.

Additive manufacturing is furthermore often used for prototyping. Because FFF 3D printers can manufacture arbitrary geometry at a low price, prototyping has become more accessible all over the world. This in turn boosts the development of new products and of technology in general. FFF helps to accelerate the advancement of technology, which benefits the world in a plethora of ways.

Such FFF printed products can then be reduced in weight by using FGMs. FGMs can produce lightweight structures which are relatively strong. FGMs can therefore contribute in reducing the amount of material used, which in turn reduces energy consumption and material usage. Furthermore, products with different requirements in different regions can be manufactured using a single material FGM, rather than several parts of different materials; FGMs reduce the need of connecting parts such as bolts and they improve the recyclability of the product. In that respect FGMs lessen the environmental impact of products.

FGMs have often been used in the fields of soft robotics, aerospace, healthcare and defense. Currently researchers are studying how FGMs can be used for ultra-personalized products. However, nearly all products have different requirements in different regions. In the future we might see FGMs being used in far more fields than they currently are.

VII.3 FUTURE OUTLOOK

The research areas of FGMs and FFF are vast and there is copious room for future research.

This dissertation has focused on controlling the emergent material properties of FFF printed products, but was restricted to properties which could be graded spatially. Various material properties of the final product can be controlled through factors which are difficult to vary spatially. For example, the mechanical performance of a part is greatly determined by how well an extruded line is fused to its neighboring lines, which depends greatly on the thermal properties of the material and the temperature of the extrudate and of the already printed geometry. Thermal control mechanisms can help to improve the mechanical performance of parts. Furthermore, the layer-wise buildup of FFF products causes a part to be weaker in the Z direction than in the XY plane. Non-planar approaches can help in (partially) alleviating this anisotropy. With the development of such approaches the overall mechanical performance of FFF printed parts can be enhanced which would boost the adoption of FFF in industrial applications.

FFF toolpath generation has several research areas which show potential for future research. Generating support structures with minimal material while adhering to the continuous extrusion constraint would be one direction of interesting research. Another is slicing layers with variable layer thickness or non-planar layers. Researching how toolpath geometry generation can be optimized for thermal performance in the FFF extrusion lines is largely uncharted research territory. Furthermore, the print defects at the ends of each extrusion line still warrant research into reducing the effect and placement of the ends of extrusion lines. Beside reducing the effect of the line ends future research could be dedicated to reducing the number of extrusion discontinuities. The variable width toolpath generation framework from Chapter IV could be extended in order to support continuous extrusion. Moreover, it could perhaps be combined with the variable width contouring framework of Hornus et al., so that we can reduce the number of micro-gaps, while simultaneously keeping the line width variation under control.

With technological advances, FFF can grow to overcome not only the integer thickness constraint and chemical compatibility constraint, but other constraints as well. A better deposition accuracy and process control, or novel hardware design can help to improve the printing resolution. Non-planar slicing can help to overcome the overhang constraint. And advanced toolpath generation frameworks can solve the problem of continuous extrusion toolpaths, which always adhere to the non-overlap constraint. With current research progressing this frontier we can expect the industry to follow suit in overcoming those manufacturing constraints.

The reduction of manufacturing constraints for additive manufacturing streamlines the development of FGMs and simplifies design for additive manufacturing. The research field of FGMs can focus on the geometry of the design and the interaction between geometry and mechanical properties without having to be concerned with toolpath generation or manufacturing technology. Designers will then be more concerned with defining specifications and FGM researchers will be more concerned with generating geometry which adheres to those specifications.

Functionally graded materials can be conceptualized as a subset of smart materials. Whereas FGMs often consider only the conventional material properties to be varied throughout a product, smart materials incorporate dimensions not found in conventional materials. Auxetics are a class of smart materials which control their lateral deformation when actuated axially. 4D printing comprises smart materials which deform when

exposed to other external stimuli, such as electricity or water. Such smart materials can be made to adapt to their context, which means they can function better in a world which is diverse and progressively evolving. The future might see an exponential rise in the adoption of FGMs and smart materials in general.

The active and reactive mechanisms of smart materials can be combined in smart ways to comprise logic circuits. Researchers can manufacture control systems to replace conventional ones which consume electricity. Given the global focus on climate change we can expect the concept of mechanical computation to evolve and take root in end products.

Related to FGMs is the discipline of topology optimization. Topology optimization is concerned with finding the optimal geometry and distribution of material(s) in such a manner that a specification of requirements on the physical behavior of a product is attained in some optimal way. For example, topology optimization can minimize the required amount of material needed for a product to withstand some forces, which is valuable for e.g. aerospace applications. Those requirements can be defined only on a couple of key locations in the product, whereas FGMs are concerned with reproducing a material specification which is defined everywhere. Current research is looking into combining topology optimization with FGMs, by making the topology optimization framework output a material specification, which is then used as input to generate an FGM. These two disciplines can be consolidated into a more unified discipline, which simultaneously optimizes both cell geometry and how that geometry is spatially graded. In the future such consolidated FGM topology optimization might become accessible in commercial design software.

The research of this dissertation has explored how to realize a spatial gradation for several material properties using FFF machines. However, the range of possible material properties to grade is broad, so the opportunities for future research are numerous. Mechanical properties which could be interesting to grade are material properties such as brittleness, creep, fatigue or slip. The presented research was often focused on isotropic mechanical material properties, such as stiffness; future research could be aimed at reproducing orthotropic properties, so that the FGMs behave differently when probed along different directions. So far we have focused primarily on mechanical material properties, but there is ample opportunity for developing FGMs which create a gradation in acoustic properties, electrical properties, magnetic properties, optical properties or flammability. In the future we might see FGMs being used for a wide range of applications.

This dissertation has explored a plethora of techniques to use FFF to manufacture structures for which physical properties can be imposed by controlling the topology of those structures. When manufacturing evolves to eliminate its design constraints, technology can focus on the interplay between the topology of matter and the physical properties which emerge from that topology. The physical reality we live in can be shaped to fit the multitude of ways in which objects are utilized, which can be diverse and can change over time. Engineering the interplay between topology and its emerging properties will drive advancement of technology and science.

ACKNOWLEDGMENTS

Throughout the research trajectory I have received a great deal of support and assistance from various people.

I would like to thank Jun Wu and Charlie Wang for supervising me and diligently helping me move forward. Jun's comments on my writing were conducive in lifting the scientific rigor of each of the published research papers and of the dissertation as a whole. Charlie's challenging suggestions guided me to radically change my perspective, which caused me to leap forward in my research projects. Their feedback has been invaluable and has helped me progress throughout my research trajectory.

I would like to express gratitude to Bas van Deursen and Peter Brier for supervising me in a manner which let me do my research independently, while connecting me with the right departments at Ultimaker when needed. Furthermore, I am grateful to Paul Heiden and Anouschka Wodrada for allowing me to embark on this journey.

I am thankful to Zjenja Doubrovski and Jouke Verlinden for the help they provided before I started my PhD. The paper we co-authored back then was the catalyst which eventually led to me starting this PhD trajectory. I would especially like to thank Zjenja for the insightful discussions we had throughout the trajectory.

I appreciate Johan Versteegh, Mercedes Fernandez, Siert Wijnia and Rijk van Manen for their insights into the intricacies of material extrusion and how toolpath generation can influence that process. Their bright outlook and critical analysis kept me optimistic and sharp.

I value Willemijn Elkhuisen and Rob Scharff for being researchers by my side – helping each other in our research and having fun in the lab.

I would like to give thanks to Ferry Helmich and Dikkie Baan for being good friends in times of need.

Finally, I am indebted to my parents Ines de Booij and Michel Kuipers for suggesting linguistic corrections and for supporting me throughout my life leading up to this milestone.

CURRICULUM VITAE

Tim Kuipers

15 Apr 1987 | Born in Eindhoven, the Netherlands

WORK

2014–2022 | **Software Engineer**
Ultimaker B.V.



EDUCATION

2017–2022 | **PhD at Industrial Design Engineering**
Delft University of Technology, the Netherlands
Dissertation: Toolpath generation for fused filament fabrication of functionally graded materials
Promotor: Prof.dr. Charlie C.L. Wang
Co-promotor: Dr. Jun Wu

2011–2014 | **MSc Technical Artificial Intelligence**
Utrecht University, the Netherlands
Track: AI techniques
Thesis: *Spread maximization – a novel unsupervised learning paradigm applied to convolutional neural networks*
Advisor: Prof.dr. Marco Wiering
Cum Laude

2006–2011 | **BA Philosophy**
Utrecht University, the Netherlands
Track: Analytic Philosophy
Thesis: *Grissiaans a priori als versmelting van het idealistische fundament van Bolland en het intuïtionistische raamwerk van Brouwer*
Advisor: Prof.dr. Albert Visser
Cum Laude

1999–2005 | **Voorbereidend Wetenschappelijk Onderwijs**
pleincollege Van Maerlant, Eindhoven, the Netherlands

PUBLICATIONS

- 2022 Tim Kuipers, Renbo Su, Jun Wu, and Charlie C.L. Wang. “ITIL: Interlaced Topologically Interlocking Lattice for continuous dual-material extrusion”. In: *Additive Manufacturing* 50 (2022), p. 102495. ISSN: 2214-8604. DOI: <https://doi.org/10.1016/j.addma.2021.102495>
- 2020 Tim Kuipers, Eugeni L. Doubrovski, Jun Wu, and Charlie C.L. Wang. “A Framework for Adaptive Width Control of Dense Contour-Parallel Toolpaths in Fused Deposition Modeling”. In: *Computer-Aided Design* 128 (Nov. 2020), p. 102907. ISSN: 00104485. DOI: [10.1016/j.cad.2020.102907](https://doi.org/10.1016/j.cad.2020.102907)
- 2020 Samuel Hornus, Tim Kuipers, Olivier Devillers, Monique Teillaud, Jonàs Martínez, Marc Glisse, Sylvain Lazard, and Sylvain Lefebvre. “Variable-width contouring for additive manufacturing”. In: *ACM Transactions on Graphics* 39.4 (July 2020). ISSN: 0730-0301. DOI: [10.1145/3386569.3392448](https://doi.org/10.1145/3386569.3392448)
- 2019 Tim Kuipers, Jun Wu, and Charlie C.L. Wang. “CrossFill: Foam Structures with Graded Density for Continuous Material Extrusion”. In: *Computer-Aided Design* 114 (Sept. 2019), pp. 37–50. ISSN: 00104485. DOI: [10.1016/j.cad.2019.05.003](https://doi.org/10.1016/j.cad.2019.05.003)
Best paper award (2nd place)
- 2018 T. Kuipers, W. Elkhuisen, J. Verlinden, and E. Doubrovski. “Hatching for 3D prints: Line-based halftoning for dual extrusion fused deposition modeling”. In: *Computers & Graphics* 74 (2018), pp. 23–32. ISSN: 00978493. DOI: [10.1016/j.cag.2018.04.006](https://doi.org/10.1016/j.cag.2018.04.006)
- 2017 Tim Kuipers, Eugeni Doubrovski, and Jouke Verlinden. “3D hatching: linear halftoning for dual extrusion fused deposition modeling”. In: *Proceedings of the 1st Annual ACM Symposium on Computational Fabrication*. ACM, 2017, p. 2. DOI: [10.1145/3083157.3083163](https://doi.org/10.1145/3083157.3083163)

PATENT

- 2021 Tim Kuipers. “Method of determining toolpaths for an infill structure for a digital 3d model”. Pat. req. Ultimaker B.V. July 15, 2021. URL: <https://patentscope.wipo.int/search/en/detail.jsf?docId=W02021139928>