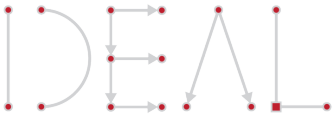
Large-Deflection Topology Optimization: An Open-Source Solver and Multi-Material 3D Printing

Elias Huber
Matriculation Number: 23-928-518

Thesis Supervisor:
Millicent Schlafly

In Partial Fulfillment of the Requirements
for the Degree of
Bachelor of Science ETH in Mechanical Engineering





Swiss Federal Institute of Technology Zurich
Department of Mechanical and Process Engineering (D-MAVT)
IDEAL – Chair of Artificial Intelligence in Engineering Design

Autumn Semester 2026

Abstract

A compliant structure that must reach large deflections is designed as a density field in a continuum, but the printer can deposit only solid material. Several transformations separate the design from the printed part, and this thesis examines three: the solver's approximation of the physics, the threshold that turns the field into a solid, and the deposition of the filament.

This thesis rebuilds the design loop and measures each transformation in its own terms. The new loop runs on Firedrake with a mixed Q1P0 element and a taped adjoint, and keeps the inherited cube, motions and optimization settings. Frozen fields were re-solved as the solid the exporter builds at three cutoffs, and VarioShore, a foaming TPU, was characterized on a Prusa XL over a nozzle-temperature ladder, in prints that combine the TPU body with a soluble support and PETG fixture plates.

The loop runs without a license and scales with ranks. On the one field that both codes solve, the stored energy differs by 6 to 36 % between the inherited Abaqus pipeline and Firedrake. Hourglass control and the integration rule are excluded as the cause. The linear model changes the design on every axis that can be read, most on the rotations, where they fall 0.30 to 0.48 below their repeat band; `tz` is confounded by its shorter stroke. Five separate single-axis designs reach 7.0 mm and 0.40 rad. At $\rho \geq 0.5$ the export cutoff moves the energy secant between 54.6 % lower and 17.9 % higher, with the sign set by the body.

The foaming TPU leaves its plateau between 215 and 220 °C, reaches 1.94× free width expansion at 240 °C, and the compensated prints stay within 0.06 mm of the 0.40 mm target, and within 0.03 mm from 230 to 240 °C. The gap is therefore narrowed to a sequence of steps with their own units, and the major number still missing is the modulus of the foamed solid. Taken together, the results describe what each step does: the cross-code gap bounds the solver, the repeat bands show when a design difference is real, the cutoff measurements give the export's effect, and the foam map sets the deposition. They are the inputs for predicting a printed part's stiffness one step at a time, and the foaming TPU is the material that would let that last step be tuned.

Acknowledgments

Thanks go to Millicent Schlafly, who supervised this thesis, met weekly, and shaped it with her feedback, suggestions and guidance. Thanks also for tolerating the many two-minute questions at her door that turned into fifteen-minute detours.

Thanks go to Federico Sánchez Trachsel, who answered many questions about the sensor, the robot, and the pipeline.

Thanks also to the other students in the group for the discussions in the student rooms, and for the feedback and questions that shaped this work.

Generative AI tools assisted with drafting, literature search and code; Appendix A.1 declares the tool, its scope and its use.

Contents

List of Figures iv

List of Tables v

Abbreviations vi

1 Introduction 1
- 1.1 Research questions . . . 2
- 1.2 Contributions . . . 2

2 Background and Related Work 3
- 2.1 Density-based topology optimization for compliant structures . . . 3
- 2.2 The Abaqus platform and its limits . . . 4
- 2.3 The hyperelastic law and near-incompressibility . . . 5
- 2.4 Foaming TPU in material extrusion . . . 6

3 Methods 8
- 3.1 The design loop . . . 8
- 3.2 The formulation and its quantities . . . 10
- 3.3 Solver and model checks . . . 11
- 3.4 From field to part . . . 12
- 3.5 VarioShore calibration . . . 14
- 3.6 Density viewer . . . 15

4 Results 17
- 4.1 Solver throughput and checks . . . 17
- 4.2 Linear and nonlinear designs . . . 18
- 4.3 Reaching the target stroke . . . 19
- 4.4 Print cutoff . . . 21
- 4.5 Foam response . . . 22

5 Discussion 24
- 5.1 The solver's bound and the model's change . . . 24
- 5.2 Reachability and feasibility . . . 25
- 5.3 The cutoff's measured step . . . 26
- 5.4 The material's operating range . . . 27
- 5.5 Limitations . . . 28

6 Conclusion 29

7 Future Work 30

References 31

A Appendix 33
- A.1 AI Use Declaration . . . 33
- A.2 Code and data . . . 34
- A.3 Optional settings . . . 38
- A.4 Solver checks . . . 39
- A.5 Foaming ladder and flow compensation . . . 40
- A.6 Print record . . . 41
- A.7 Export budget . . . 42
- A.8 Cutoff sweep . . . 43

List of Figures

1 The six motion patterns . . . 1
2 The design loop as rebuilt . . . 9
3 Calibration prints . . . 14
4 Density viewer . . . 15
5 Tangent force against stroke . . . 18
6 The six service-load geometries . . . 20
7 The print cutoff changes stored energy in both directions . . . 21
8 Relative bead width . . . 22
9 The four printed parts . . . 23

List of Tables

1 Optimization parameters . . . 10
2 Materials . . . 13
3 Compute and rank scaling . . . 17
4 Model comparison . . . 19
5 Stroke and end state . . . 19
6 AI use declaration . . . 33
7 Run matrix . . . 35
8 Cutoff evaluations . . . 36
9 Optional settings . . . 38
10 Internal energy . . . 39
11 Solver probes . . . 39
12 Foaming ladder and flow compensation . . . 40
13 Print record and settings . . . 41
14 Export budget . . . 42
15 Full cutoff sweep . . . 43

Abbreviations

Abbreviation	Expansion
DOC	Degree of constraint
DOF	Degree of freedom
MMA	Method of moving asymptotes
PETG	Polyethylene terephthalate, glycol-modified
PLA	Polylactic acid
Q1P0	Continuous linear displacement, element-constant pressure
SIMP	Solid isotropic material with penalization
SLURM	Simple Linux Utility for Resource Management, the batch scheduler on Euler
TO	Topology optimization
TPU	Thermoplastic polyurethane

1 Introduction

A compliant structure can be designed in a continuum but printed only as a mix of solid and void. The optimizer distributes density over the cells of a 40 mm cube and returns a field that is mostly zero or one, with a band of intermediate density in between. The design problem requires the top face of the cube to move freely in one or more of the six motions and to resist the rest. Those motions are the three translations, `tx`, `ty` and `tz`, in millimeters, and the three rotations, `rx`, `ry` and `rz`, in radians about an axis through the cube's center depticted in Figure 1. The target that motivates the work is a stroke of 7.0 mm and 0.40 rad at the moved face.

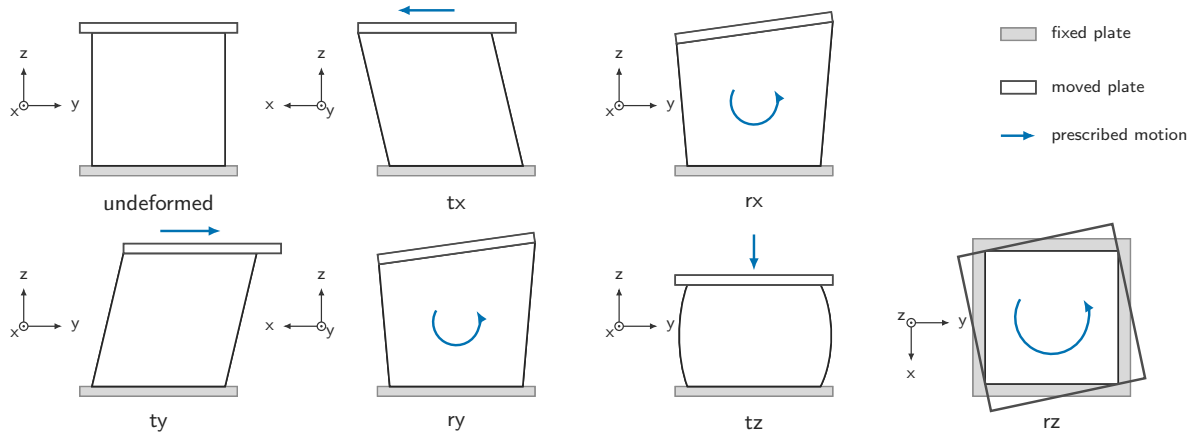


Figure 1: The six motion patterns, undeformed first. A 40 mm cube is held between a fixed lower plate (grey) and a moved upper plate; tx, ty and tz translate it along x, y and z, and rx, ry and rz rotate it about those axes; the dot marks the axis out of the page. Redrawn after Goronzi (2026).

The printer cannot deposit that intermediate density band. Therefore the exporter must decide where the solid begins and the void ends. A threshold turns the band into a mix of complete solid and complete void, and the printer lays that solid in thermoplastic polyurethane (TPU). The threshold removes the intermediate densities that graded the transition, so some sections of the printed part come out thinner or shorter than the field originally asked for.

A flexure carries its softness in its thin sections. The threshold can make them thinner or shorter. It also turns grey (between 0 and 1) density regions into full material which can make other sections stiffer than the grey field.

A printed part does not exhibit the stiffness the field asked for, and the difference cannot be assigned to a single cause. There are multiple transformations between the field the topology optimization (TO) produces and the printed part, which can move the printed stiffness away from the designed one. Three of them are investigated in this thesis and they are the solver's approximation of the physics, the threshold that turns the field into a solid, and the deposition, in which the printer builds the part from filament.

The formulation, the cube and the six motions are inherited from earlier theses in the group (Lang 2026; Trachsel 2026).

This thesis replaces the solver, the element and the sensitivities: Firedrake, Q1P0 and a taped adjoint stand in for Abaqus, the hybrid brick and the analytic adjoint, so the loop runs without a license. The same loop carries both models, and their comparison says where the inherited linear design model stops reproducing the nonlinear design at the target stroke. It adds a continuous modulus map, a measure of what the export step costs in stiffness, and an operating range for the foaming TPU the parts are printed in. The cube, the objective and the print recipe stay as they were, therefore the two pipelines share the same problem.

The bodies are printed in a foaming TPU, VarioShore. Its main advantage here is printability: a soft TPU is difficult to print, and this filament feeds and prints like a stiff one before it comes out soft. Section 2.4 sets out the mechanism, and Section 3.5 calibrates the process window. Varying stiffness through temperature (or printing speed), across several materials or as a gradient inside one part, is a route that calibration opens, and this thesis leaves it for future work.

1.1 Research questions

1. Can an open-source solver carry the design loop and speed up with more compute added?
2. Does the inherited linear design model hold over the stroke range the thesis uses?
3. Can the open-source pipeline reach large deflections without a continuation failure?
4. How much does turning the optimized density field into a printable solid change the stiffness?
5. Which process settings control the bead width and expansion of VarioShore on the Prusa XL, and what values realize a chosen bead width?

1.2 Contributions

The main contributions of this thesis are:

- A license-free design loop that scales with more compute: 1.58× at 64 ranks (parallel processes), larger than the scatter between two runs of the same rank count.
- A target of 7.0 mm and 0.40 rad, reached on five of the six axes, two of them above the 0.05 compliance cap.
- An export step whose cost in stiffness is measured: at $\rho \geq 0.5$ it moves the secant stiffness between 54.6 % lower and 17.9 % higher, and the sign follows the body.
- A measured operating range and a temperature-to-flow map for VarioShore TPU: the plateau ends between 215 and 220 °C and the map holds the compensated bead within 0.03 mm of the 0.40 mm target from 230 to 240 °C.

Appendix A carries the full records behind these numbers. It holds the run matrix, the solver checks, the foaming ladder, the print record and the export budget, and each table there is the source the body quotes.

2 Background and Related Work

2.1 Density-based topology optimization for compliant structures

A density-based topology optimization problem makes the material distribution the design variable. A scalar field ρ on the mesh cells takes the value one in solid material and a small density floor in void, and the stiffness of a cell follows from its density alone. The field can take any density between zero and one, and the optimizer moves it cell by cell under a gradient.

The interpolation between the phases decides how much stiffness a grey cell returns. The classical solid isotropic material with penalization (SIMP) form raises the density to a penalty exponent, which makes an intermediate cell inefficient and drives the optimum toward the bounds (Bendsøe et al. 2004). The SIMP penalty was not active in the inherited code. That code carried a linear map, discretized into a ladder of a thousand sections, one per density level, because Abaqus assigns one discrete material per element and takes the map only as a finite set of sections. This thesis keeps that linear two-material map, $E(\rho) = E_{\text{void}} + \rho\,(E_{\text{solid}} - E_{\text{void}})$, and evaluates it at every density, so the residual is differentiable with respect to ρ.

The map also fixes the two ends of the range. $E_{\text{void}} = 0.004\,\text{MPa}$ is a soft solid that keeps the stiffness matrix non-singular and describes a void, and $E_{\text{solid}} = 21\,\text{MPa}$ is the print-effective modulus of FiberFlex 40D (Trachsel 2026, p. 9). The void end carries a blended constitutive response: below the blend boundary of 2.8 MPa, the blend fades ν from a modeling value of 0.1 toward the solid 0.48 and scales the higher-order Yeoh terms with it, so a mostly void cell behaves as a compressible solid. That blend boundary uses the print-effective modulus of Filaflex 60A, the soft TPU of the inherited work (Trachsel 2026, p. 9), in place of a measured void property.

The design variable (density) passes through two regularizers before it becomes a material that is evaluated by finite-element analysis. A cone filter averages the density over a neighborhood and sets the smallest feature the field can express. A Heaviside projection with a ramped sharpness then drives the filtered field toward the bounds. The update is Svanberg's method of moving asymptotes (MMA) (Svanberg 1987), which carries the constraints of the problem and applies a move limit to each cycle.

The compliant structure in this and the inherited work is a flexure. One face of a 40 mm cube is held fixed and the opposite face carries a prescribed stroke in one or more of the six motions. The structure must yield in each compliant motion, which is a degree of freedom (DOF), and hold each resisting motion, a degree of constraint (DOC). A design may be compliant in more than one motion at once. The six motions are the three translations and the three rotations of the moved face, and each is solved on its own.

Classical minimum compliance minimizes the stored energy at a fixed load and would return the stiffest structure for one load path, which, at a prescribed stroke, is a solid block. The flexure problem replaces that objective with a ratio of stored energies at equal stroke, maximized over the resisting motions and capped on the compliant one (Equation 1). The cap

is what keeps the structure soft where it must move, and the objective is what keeps it stiff where it must resist.

$$\max_{\rho} \sum_{m \in \mathcal{R}} \alpha_m(\rho) \quad \text{subject to} \quad \alpha_j(\rho) \leq \bar{\alpha}, \quad j \in \mathcal{C}, \qquad \alpha_m = \frac{W_m}{W_m^0}, \tag{1}$$

where $\mathcal{C}$ and $\mathcal{R}$ are the compliant (DOF) and resisting (DOC) motion sets, with $\mathcal{R}$ the complement of $\mathcal{C}$, W_m is the strain energy stored in motion m, and W_m^0 is the energy of the solid cube at the same prescribed motion.

The normalization makes the cap dimensionless. Dividing the energy of a motion by the energy of the solid cube at the same motion makes $\alpha_m = 1$ a solid response, so $\alpha_m \leq \bar{\alpha} = 0.05$ keeps the compliant motion at most 5 % as stiff as the solid.

The cap takes over the role that a volume constraint plays in classical topology optimization. Koppen et al. (2022, p. 10) show that bounding the compliant strain energy sets how much material the design spends, so no separate volume fraction is imposed. The amount of solid therefore follows from the stiffness requirement.

This energy based formulation is inherited. Lang (2026) built the workflow on the energy formulation of Koppen et al. (2022), and Trachsel (2026) applied it to the 40 mm cube with the six motions, the cone filter, the Heaviside projection, an MMA update and the print recipe. The inherited design loop solved the problem in Abaqus with the `C3D8RH` element and a modulus ladder, and differentiated the analytic adjoint of the linear design model. Trachsel (2026) added a geometrically nonlinear evaluation of the force–displacement path; that path stopped at 2.0 to 3.0 mm for translations and 0.12 rad to 0.14 rad for rotations, against a 4 mm and 0.2 rad request. This thesis needs the larger stroke, so it keeps the nonlinear law and replaces the solver, the element and the sensitivities.

2.2 The Abaqus platform and its limits

The inherited loop runs in Abaqus/Standard, a commercial finite-element code, and it inherits the `C3D8RH` element, the material ladder and the linear adjoint of the design loop. The platform blocks the work through a license, its closed implementation and its fixed extension points.

The first limit is the license. A job on eight cores checked out 12 license tokens from a shared ETH pool that was close to full. The inherited work names license time as the binding constraint (Trachsel 2026, p. 27). Each concurrent job holds its tokens for its whole wall time, so a wider campaign of multiple concurrent runs is paid for in tokens as well as in core hours. This strictly limits the concurrent tests that can be run.

The second limit is inspection. Abaqus ships as a binary, and its source is not available for reading or modification. The element, a static stabilization and a contact algorithm enter the run as documented options (Dassault Systèmes 2016), so their effect appears in the output and their cause stays inside the code. The cross-code comparison of Section 3.3 sees this

limit: Abaqus adds an artificial energy of 1.6 to 11.05 percent of its internal energy, visible in the output and not derivable from the implementation (Section 4.1). A stopped run leaves a message file and its last converged state, and the failed iterate is not written as a field.

The third limit is adaptation. A continuous modulus map enters an Abaqus element only as the thousand-section ladder of Section 2.1, because every element carries one discrete material. The tangent of the finite-strain residual is assembled inside the solver, and the inherited sensitivity is the linear adjoint described in Section 3.1. A new element, a new sensitivity or a potential new contact law means custom tooling over a fixed interface, and the solver that drives it stays closed.

2.3 The hyperelastic law and near-incompressibility

Printed TPU is an elastomer (colorFabb 2023), and the prescribed rotation of 0.40 rad and translation of 7.0 mm move the flexure into the finite-rotation regime, where the small-rotation kinematics no longer hold (Bonet et al. 2008). Whether the printed material also leaves the linear elastic range at that strain is an assumption inherited with the model. The material is therefore modeled as an isotropic hyperelastic solid. Its strain energy is the isochoric third-order Yeoh form of Equation 2, written on the deviatoric first invariant (Yeoh 1993).

$$\Psi = C_{10}(\bar{I}_1 - 3) + C_{20}(\bar{I}_1 - 3)^2 + C_{30}(\bar{I}_1 - 3)^3, \qquad \bar{I}_1 = J^{-2/3}\,\mathrm{tr}\,\mathbf{C}, \tag{2}$$

where $\mathbf{F}$ *is the deformation gradient,* $\mathbf{C} = \mathbf{F}^{\mathsf{T}}\mathbf{F}$ *is the right Cauchy–Green tensor, and* $J = \det \mathbf{F}$ *is the volume ratio.*

The law and the coefficient ratios come from an earlier loop, where they were curve-shape choices (Trachsel 2026). No coefficient was fitted to a printed specimen in this thesis. Only C_{10} is tied to a measured modulus: $C_{10} = E/(4(1+\nu))$, with E being the 21 MPa print-effective modulus of the FiberFlex 40D that was measured in previous work. C_{20} is negative and softens the response past the initial rise, and C_{30} is positive and limits the softening at large stretch. The ratios are fixed at $C_{20}/C_{10} = -0.25$ and $C_{30}/C_{10} = 0.04$.

Poisson's ratio of the solid is fixed at $\nu = 0.48$, the threshold above which the Abaqus guide treats a solid as nearly incompressible (Dassault Systèmes 2016, §28.1.1). The value is a modeling choice. The bulk modulus follows as $\kappa = E/(3(1-2\nu))$, about 8.3 times the Young's modulus, so the material resists a change of volume far more strongly than a change of shape. The volumetric part of the energy is $\Psi_{\mathrm{vol}} = \kappa(J-1)^2/2$, a penalty on any change of volume, and the stored energy W is the isochoric part of Equation 2 plus this term, integrated over the body.

A displacement element can deform at constant volume only through the deformation modes its shape functions allow. Full integration imposes a volumetric constraint at every quadrature point, and the constraints together leave the element too few modes to meet them at constant volume. The element then answers an isochoric deformation with volumetric stiffness,

and the error survives as the bulk modulus grows; this is volumetric locking (Bonet et al. 2008, p. 256).

One volumetric constraint per element is the right number for the eight-node brick (Nagtegaal et al. 1974, p. 166). A mixed element reaches that count by testing the volume once per cell, with the pressure carried as a field that is constant on each cell.

The inherited brick is `C3D8RH`, the hybrid form of the eight-node element with one integration point. The hybrid form carries the pressure as an independent field, constant over the element, so the volume constraint is enforced once per cell, which makes it a mean-dilatation element. The cross-code comparison runs that inherited element against Firedrake, and Section 4.1 reports the result.

The element of this thesis is Q1P0, which enforces the volume constraint the same way, with a linear displacement and a piecewise-constant pressure. It integrates the displacement at full quadrature, which avoids the hourglass parameter that the reduced-integration `C3D8RH` of the inherited work required (Flanagan et al. 1981). The same fields at full integration in Abaqus are the hybrid `C3D8H`, which matches Q1P0 in interpolation, quadrature and constraint.

2.4 Foaming TPU in material extrusion

A material that is soft when printed helps reach large deflections at moderate forces, but very low shore-hardness TPUs suffer from printability issues that were also limiting in previous work (Trachsel 2026, pp. 9, 13). Foaming TPU has the properties that make it interesting here: unfoamed it is a stiff elastomer of about 95 Shore A, similar to the FiberFlex 40D of previous work, and VarioShore feeds like a hard filament while a blowing agent expands it as the melt is heated. The density and the hardness of the deposited bead fall with the heat the filament receives and therefore create a softer printed part.

The density the optimizer designs takes every value between void (0) and solid (1). Earlier work printed a hard/soft TPU pair to achieve an approximation of the value between full solid and void (Trachsel 2026). That pair bonded poorly, likely because the two formulations differ in their additives and are poorly compatible at the interface. Foaming TPU gives a continuous stiffness range that can be used as either a continuous change in density or as multiple distinct stiffness levels like the hard/soft TPU pair, with the same chemical composition and therefore reduced bonding issues. The printed bodies in this thesis take neither route: they use a single setting. The calibration printed several levels on one plate (Section 5.4), and Section 7 sets out the graded part that would use them.

The foaming behavior of these filaments is sensitive to many parameters and these differ from one physical setup to another. The main factor is the heat the filament recieves in the heating chamber of the hotend (Nofar et al. 2022). That heat depends on the nozzle temperature and on the time the material spends in the hotend. This time, in turn, is influenced by printing speed, flow rate and maximum volumetric flow. A comprehensive setup of multiple

printing parameters is therefore needed to succesfully and reliably print a material like this.

One precedent for a graded foaming part is looking into polylactic acid (PLA). Grigolato et al. (2025) characterize two foaming PLAs between 200 and 230 °C, and fit density, hardness and modulus against nozzle temperature. They then post-process the G-code to set the extrusion temperature and the extrusion ratio per move from a volumetric density map. The scheme is demonstrated on a topology-optimized bracket whose relative density sets the local temperature, and the print runs at full infill, so the deposited bead carries the gradient at its foamed density.

The method compiles a continuous density field into process parameters along a toolpath. The paper reports its own limit: because of larger thermal lag the temperature changes do not reach the nozzle immediately. The printed bracket therefore shows striations that do not follow the density map, with undesired voids and over-extruded regions (Grigolato et al. 2025, p. 17). The continuous field is realized only in part.

The scheme also does not transfer to an elastomer, because both materials are PLA and every property is measured on printed test specimens. The flow calibration in this thesis is inspired by their methods, which Section 3.5 reports.

A second precedent compiles the field into a slicer project. Wade et al. (2026) partition a heterogeneous design into a finite number of regions and write per-region settings and process states into a `.3mf` file, so toolpath planning stays with the slicer. For foaming they calibrate the wall width and the hardness of VarioShore against temperature, and apply the flow compensation as an `M221` command, the G-code instruction that scales the extrusion flow. They print a bar graded in hardness from 65A to 85A within 0.5 Shore A of its target on average.

Their fields are hand-made and not generated by an optimizer, and the graded parts are validated by hardness and by a center of mass, which leaves stiffness untested. The paper reports no print speed and no layer height for its foaming calibration. A process recipe needs the flow as well, because the same temperature gives different expansion at different flow (Nofar et al. 2022). The mechanism this thesis takes from it is that a hardness field can be realized through slicer-native process states on a VarioShore print. That route was used to calibrate the VarioShore process on the lab's Prusa XL.

An earlier filament test by CNC Kitchen saw VarioShore expand from about 200 °C and reach roughly twice its nominal width at 220 °C (Hermann 2020). The manufacturer's own flow compensation is still at nominal at 210 °C and falls to 70 % at 230 °C (colorFabb 2023).

3 Methods

A field is worth printing only once the solver that produced it is bounded, so the loop comes first: the cycle, the formulation it solves, and the checks that say how far the replacement can be trusted. The second half turns a field into a part, through the exporter, the calibration of the print and the viewer.

3.1 The design loop

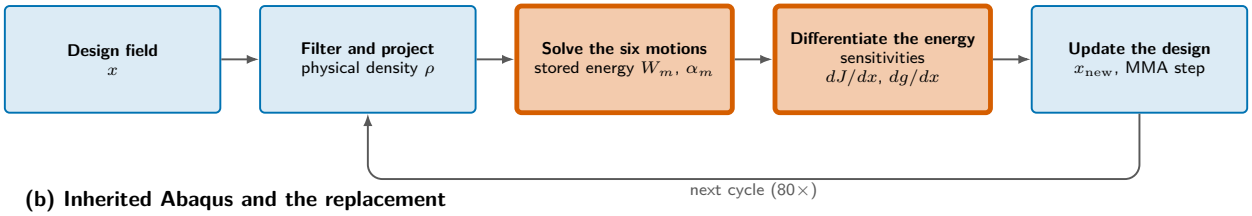
The design variable is one density per cell on a 40 mm cube, and a cycle turns that field into a stiffness distribution the printer can produce. The loop that does this was rebuilt around a different solver: Firedrake replaces Abaqus, while the cube, the six prescribed motions, the filter, the projection and the update remain inherited, and the three blocks that touch the solver are the ones this thesis changed. Figure 2 shows the sequence.

The bottom face of the cube is held fixed, and the top face carries a rigid motion: a translation, or a rotation about an axis through the cube's center. Six motions are prescribed in all: three translations and three rotations. In the six single-axis runs, one motion is the driven DOF at the full target of 7.0 mm or 0.40 rad, while the other five are its DOCs, the directions the flexure must hold, at a resisting amplitude of 2.0 mm and 0.10 rad. These six runs are the *service-load* family. A second family, the *full-load* runs, drives all six motions to the full target at once; every motion is then both driven and resisted. The family is one run per axis, six runs in all. Both families use the ceiling below.

Each motion has a ceiling, the amplitude the loop may lower when a solve attempt fails. A failed attempt gets one of two rescues, and they alternate: the design is first rolled halfway back toward the last accepted one and all six motions re-solve, and if that fails too, only the failing motion's ceiling drops one step, 0.25 mm or 0.01 rad, and the same design is retried. Four rollbacks per cycle and a bounded number of ceiling reductions per run are allowed; when both are spent the run ends on the last accepted design. A ceiling may also rise when a solve succeeds. Each motion starts at a ceiling of 2.0 mm or 0.10 rad, the resisting amplitude. In the single-axis runs, the compliant motion's ceiling climbs one step per cycle from cycle 16 toward its target, while the five DOCs stay on their initial ceiling.

The replacement this thesis undertakes touches three blocks. The first is the modulus ladder: because Abaqus limited the inherited loop to a thousand discrete material sections, a cell could take one of a thousand stiffnesses; the rebuilt loop carries the continuous map $E(\rho)$ of Section 2.1 instead, so the modulus is a field like the density and is differentiated like one. The second is the element and its solve: the inherited arm coupled `C3D8RH` to a direct factorization, which assembled and factorized the whole system, while the rebuilt loop keeps the linear brick but makes it mixed and solves the saddle-point system iteratively. The third is the gradient: the inherited sensitivity is the analytic adjoint of the linear design problem, which is not the derivative of a finite-strain residual once the kinematics leave the linear regime. The rebuilt loop records the nonlinear solve on a tape and takes the adjoint of that tape, so the gradient

(a) One optimization cycle



(b) Inherited Abaqus and the replacement

Stage	Inherited Abaqus	This thesis
Modulus	*1000-section modulus ladder*	continuous $E(\rho)$
Forward solve	*C3D8RH, direct solver*	Q1P0, block-Schur + algebraic multigrid
Gradient	*analytic adjoint of the linear problem*	automatic differentiation of the nonlinear residual
Retained	*cube, six motions, cone filter, Heaviside, MMA, print recipe*	same

Vermillion, thick outline: replaced in this thesis. Blue, thin outline: retained. *Grey italic*: inherited Abaqus block.

Figure 2: The design loop as rebuilt. (a) One of the 80 optimization cycles. (b) The three stages replaced in this thesis beside their inherited Abaqus form. W_m is a motion's stored energy, α_m its normalized value, MMA the method of moving asymptotes.

is the derivative of the residual it actually assembles.

The load schedule depends on the design, so it cannot be recorded while the design is being chosen in Firedrake. Each motion is therefore solved twice per cycle: an untaped load-ramp probe finds the load steps, walking the load from zero and halving its increment whenever a solve fails, and the taped solve replays them.

A cycle applies the same four steps. A cone filter with a three-cell radius removes grid-scale noise, and a Heaviside projection with threshold $\eta = 0.5$ and sharpness β turns the filtered field into the physical density ρ, encouraging the solver to binarise the field into solid and void. The projection is off for the first half of the 80 cycles, with $\beta = 0$, and then ramps to 16, reaching that value at 95 % of the run. The six motions are then solved, each at its own prescribed amplitude, and each returns a stored energy W_m. The energies are normalized into an objective and a set of constraints, and the taped solve yields the gradient of each with respect to ρ. One MMA step then updates the design, with a per-cell move limit of 0.2 (Table 1), and the loop runs for 80 cycles.

The objective sums the normalized energies over the resisting set $\mathcal{R}$ of Equation 1, while the compliant set $\mathcal{C}$ is excluded from that sum and stands under the cap $\bar{\alpha}$ of Equation 1 as a constraint. Two switches set the physics: the large-strain kinematics and the higher-order terms of the material law. With both on, the loop solves the finite-strain problem; with both off, it solves small-strain linear elasticity on the same modulus map and the same void blend of Section 2.1, so the two models stay comparable. The nonlinear runs have both switches on, while the linear setting, the one the inherited sensitivity was derived for, supplies the control for the comparison. Both the nonlinear and linear runs begin from the solid cube, which is feasible for every motion and leaves the loop free to remove material where none is needed.

Runs execute as one job of the SLURM scheduler on Euler, the ETH Zurich cluster. At $n = 30$ cells per cube edge, each rank is given 3 GB. A rank is one process of the message passing interface (MPI) that the solve is split across; the rank count changes only the computational cost, and Section 4.1 measures its scaling. The 32-rank point is repeated once at 6 GB/rank as a certification, a repeatability check at a different memory setting, and one $n = 40$ run measures the cost of a finer mesh.

Table 1: The settings of the service-load runs. n is the number of cells per cube edge, DOC the resisting amplitude, DOF a degree of freedom, η the projection threshold, β its sharpness, E_{solid} the solid modulus, ν Poisson's ratio. Q1P0 names the element pair: continuous linear displacement (CG1) and discontinuous constant pressure (DG0).

Parameter	Value	Note
Domain	40 mm cube	inherited
Mesh	$n = 30$ (27,000 cells)	trade-off of speed and accuracy; 1.33 mm element
Element	Q1P0 mixed	DG0 pressure, CG1 displacement
Material preset	`fibreflex40d`	$E_{\text{solid}} = 21$ MPa, $\nu = 0.48$
Projection	Heaviside, $\eta = 0.5$	β ramp $0 \to 16$ over the β window
Compliance cap	0.05 per active compliant DOF	strain-energy constraint
Loads	7.0 mm / 0.40 rad	stated target
Service DOC amplitudes	2.0 mm / 0.1 rad	service-load case
Ceiling step	0.25 mm / 0.01 rad	grid-aligned
Rescue	4 design bisections	partly in inherited loop

The released code defaults to $n = 40$; the runs reported here pin $n = 30$.

The code also exposes settings that the service-load runs leave at their default value. Table 9 in the appendix lists them and what each would change.

3.2 The formulation and its quantities

The material law is the third-order Yeoh form of Equation 2, closed by the volumetric term of Section 2.3; the bulk modulus follows from $\nu = 0.48$. The law is implemented in the Q1P0 pair throughout the loop, a continuous piecewise-linear displacement with an element-constant, discontinuous pressure equal to κ times the cell's mean dilatation. The displacement is integrated at full quadrature, so one scalar equation per cell replaces the per-point volumetric constraint, and the system is solved iteratively with an algebraic multigrid on the displacement block.

The quantities are those of Section 2.1, with two changes in the new loop. The first is the normalization: $\alpha_m = W_m / W_m^0$ uses the solid reference at the amplitude the motion actually ran at. Because the loop can change the prescribed load every cycle, that reference is tabulated on the grid of ceilings, so a motion that climbed or was lowered is normalized at its own end stroke. The second is the cutoff re-solve: the exported body is solved again on the same cube, mesh, material and motion, and K is read at the same stroke. It measures what cutting the

field into a printable solid does to stiffness, and Section 3.4 gives the process. $K = 2W/d^2$ is the energy secant at stroke d: for a translation, d is the displacement and K is a force per millimeter; for a rotation, d is the angle and K is a torque per radian. Because the stroke is the same for the two bodies compared, $\Delta K/K = \Delta\alpha/\alpha_{\text{cont}}$ identically, where α_{cont} is the value of the continuum field the export started from.

The compliance cap is a design target: a field is counted as over the cap only when its end α exceeds 0.05 by more than 1% of the cap (0.0005); a smaller excess is reported as at the cap, so the count does not turn on the last digits of the reported α. Section 4.3 reports the six exact end values.

The two models are compared by r, the Pearson correlation of the final density fields over the cells of the mesh. A value of $r = 1$ means one field is a positive affine map of the other: the pattern agrees while an offset and a scale stay free. $r = 0$ means no linear relationship. The projection drives most cells to the bounds, which keeps that freedom small in the fields compared here. In the linear model a fixed motion direction stores energy in proportion to the square of its stroke, $W \propto d^2$, and the solid reference W^0 carries the same factor, so α is independent of d: the 2 mm and 7 mm linear designs are the same field by construction. The repeat band is the correlation that two runs of one design reach; the linear runs set it at 0.9984 on `tx` and 0.4342 on `tz`, and a cross-design value is read against the band of its own model; Section 4.2 gives the nonlinear band.

The cutoff study measures the share of the energy secant that the threshold accounts for. A frozen field is one run's density field at its last cycle, held fixed; it is binarized at $\rho \geq 0.4$, 0.5 and 0.6 and exported, and the repaired and culled body is then solved again on the same cube, mesh, material and motion. The void modulus in that re-solve is not zero, because a zero modulus would make the stiffness matrix singular, so the re-solve lifts the density to $x_{\min} = 0.001$, where a void cell carries 0.025 MPa against the 0.004 MPa of the map.

The cutoff population uses three of the six service-load fields (`rx`, `ry` and `tz`), then adds a mixed-cell variant (57-01), a multi-axis run (59-01) and one legacy climb (v1.1) to complete the seven. The mixed-cell variant starts the compliant motion at full stroke while the five DOCs stay at the screening amplitude of 2.0 mm and 0.10 rad; the multi-axis run drives the three rotations together.

3.3 Solver and model checks

Because the rebuilt loop replaces the solver, its design results are only useful if the replacement itself is bounded. The forward check uses a single frozen field from the inherited Abaqus run and solves it in both codes on the same mesh, with the same material and the same prescribed motion at 2.5 mm and 0.12 rad: Abaqus with the inherited `C3D8RH` element, Firedrake with Q1P0. The stored energy of each motion and the tangent force are compared.

The energy gap is reported per motion as $g = (\texttt{ALLIE} - W)/W$, where `ALLIE` is the

Abaqus internal energy and W the Firedrake stored energy at the same stroke, so a negative g means Abaqus stores less. Because `ALLIE` includes the artificial stabilization energy `ALLAE`, g carries that term as well as the discretization difference.

The check is deliberately narrow, one field at one amplitude with no inherited optimized field re-run, so it bounds the code swap, and the model question is separate: a disagreement would pass into every later export and material result, while an agreement leaves the remaining gap to the steps this thesis measures. The gradient path is checked the same way: the taped adjoint is compared with central finite differences on the same solve, which tests the tape and its replay.

The solver-envelope probes are two synthetic density fields. The first is uniform at $\rho = 0.4$; the second is a random half-solid, half-void field. Each case freezes the field and ramps one prescribed motion from zero: the uniform field is probed on all six motions at $n = 20$ and $n = 40$, and the random field on several axes at $n = 20$ and on `tz` at $n = 40$. The measured quantity is the last stroke the solver reaches before the Newton continuation stops. These are sanity checks on the forward model and its continuation: no optimizer is involved, no design result depends on them, and the two meshes are not a convergence study. Table 11 reports the $n = 40$ strokes against the inherited 2.0 mm to 3.0 mm and 0.12 rad range; the $n = 20$ context is in its caption.

The model comparison asks a different question: whether the inherited linear design model is sufficient at the target stroke. It uses the final density fields of the 2 mm nonlinear runs, the 7 mm service-load runs and the linear design, correlated by r (Section 3.2). The 2 mm nonlinear fields come from the small-motion screen: six single-axis runs at 2.0 mm and 0.10 rad, 80 cycles, with the other settings of the service load. The linear control is one run per axis with both switches off, at the service-load amplitude. Both arms are Firedrake, so the comparison isolates the model. Four pairings separate the two effects. The 2 mm and 7 mm nonlinear designs are each set against the linear design to test the model at a fixed stroke. The other two pairings show how each model changes with stroke: the 2 mm nonlinear design against the 7 mm service-load design, and the 2 mm linear check against the 7 mm linear design.

Because the linear design is amplitude-invariant, the 2 mm linear check was run on `tx` and `tz` only. The service-load runs complete 80 cycles, but `tz` stops at 5.00 mm, short of the 7 mm target, so its 7 mm entries are a 5 mm comparison. Cross-design values are read against the repeat band of Section 3.2.

3.4 From field to part

A density field becomes a part at a threshold: the exporter binarizes at a chosen threshold, adds the minimal cells that remove diagonal-only contacts, and keeps every solid component that joins the fixed and the moved face. The threshold is chosen per purpose: 0.5 for the standard export, 0.1 or 0.2 for the inspection prints, and 0.4, 0.5 or 0.6 in the cutoff study. A component that reaches one face alone, or no face at all, is discarded, so the culled body carries

only material that spans the gap between the two faces. That body may be several compliant legs; the fixture plates bond them into a single part, and the gate accepts the arrangement. The exporter then builds the closed voxel surface, smooths it without changing its topology, and appends the fixture plates and their interfaces. The binarizing step follows the printable export that Karrer (2026) built and Trachsel (2026) used; new in this work are the removal of components that do not join both faces, the diagonal repair and the watertight surface.

Before a file is written, a printability gate tests the body against four rules. The gate repeats the connection requirement: the solid must join the fixed face to the moved face through face-sharing cells, with no component standing on one face alone. It adds three limits: no enclosed void may exceed half the domain, the solid fraction must stay below the near-solid limit of 0.95, and the fifth-percentile feature thickness must stay above the print floor. The exported surface must also close watertight. The print floor is the width below which the printer cannot lay a reliable bead, and it is the larger of twice the voxel pitch and twice the extrusion width, so at $n = 30$ it is 2.67 mm. These are the rules of the current code. Table 7 in the appendix lists every campaign behind the results, with its mesh, cycle count, job and outcome.

The print uses three materials and one simulation reference, listed with their roles in Table 2: VarioShore carries the printed flexure bodies and the wall specimens, PolyDissolve S1 the support interface, and Bambu PETG HF the fixture plates, while FiberFlex 40D supplies the simulation reference only and was not printed for this thesis. Table 13 in the appendix records the print job and the four parts it produced, with the print and filament settings that differ from stock PrusaSlicer. The run matrix notes the design, cycle and cutoff each part came from. The printer, print and filament profiles are collected in one PrusaSlicer configuration bundle in the repository, `prusa-xl-lab.ini`, which imports all 39 presets created for this thesis into PrusaSlicer 2.9.6.

Table 2: The printed materials and the simulation reference, with the role, toolhead and print setting of each. `M221` is the G-code that scales the extruder flow on top of the profile's extrusion multiplier; "per profile" means the PrusaSlicer material profile supplies the setting.

Material	Role	Toolhead	Print setting	Source
FiberFlex 40D	simulation reference only (not printed)	–	preset only (not printed)	(Fiberlogy 2026); PrusaSlicer preset
VarioShore TPU	printed material: flexure bodies and wall specimens	5	235 °C, extrusion multiplier 0.52	(colorFabb 2023); foaming ladder
PolyDissolve S1	soluble support interface	4	per profile	(Polymaker 2025)
Bambu PETG HF	fixture plates	2	per profile	(Bambu Lab 2024)

The simulation does not run on the printed material, because the printed bodies are VarioShore and its modulus was not measured; it carries 21 MPa, the print-effective modulus

that Trachsel (2026) reports for FiberFlex 40D. The solid response therefore runs from 2.8 MPa, the soft-phase modulus of the earlier pair in the inherited work and the blend boundary, to 21 MPa; below 2.8 MPa the blend fades the solid response toward the void, and the void phase itself is the separate void modulus of 0.004 MPa. That range is an assumption, and Section 5.5 sets out its basis.

3.5 VarioShore calibration

The operating range is the set of nozzle temperatures and flows that gives a chosen bead width, and it was calibrated in two steps: bead width against nozzle temperature at one volumetric flow, then flow compensation at the chosen temperature. The target was a bead of 0.40 mm, equal to the nozzle diameter.

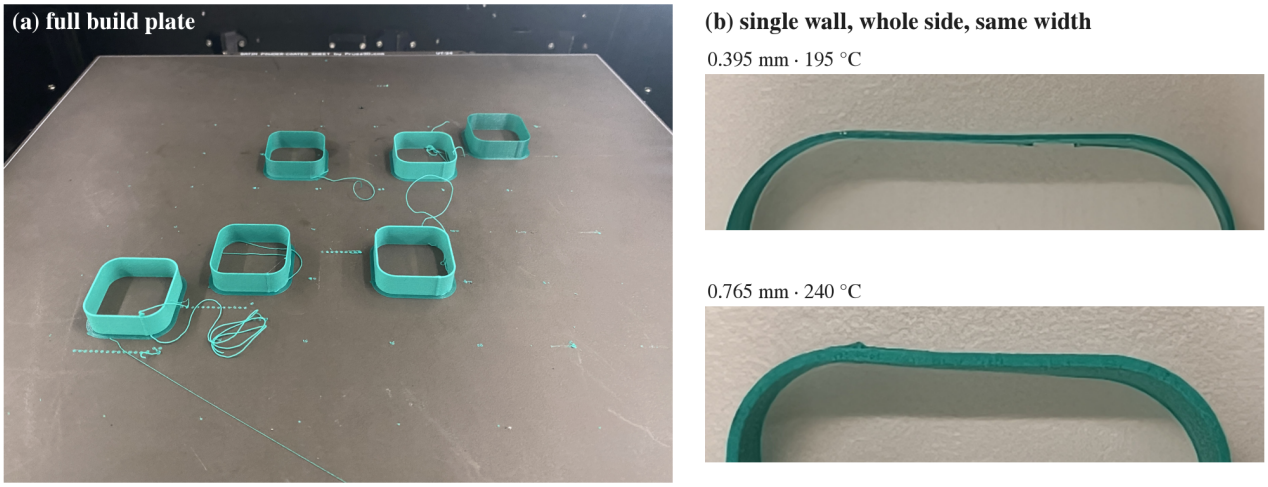


Figure 3: Calibration prints. (a) The Prusa XL build plate with the six wall specimens of the temperature ladder. (b) Whole-side views of two of them, the 0.395 mm wall at a nozzle temperature of 195 °C and the 0.765 mm wall at 240 °C; both are single perimeters printed at the same flow.

The bead width was read on a wall specimen: a 30 mm square, 10 mm tall, with a 7 mm corner radius, one perimeter wide and without infill, on a 2 mm brim. The ladder is that wall printed at a series of nozzle temperatures, and a rung is one of those temperatures.

One final specimen was printed per rung, and several rungs shared one build plate as separate objects. The nozzle waited at its setpoint, the temperature commanded for that object, before printing each one. The 0.45 mm pilot held `M221` S107, a flow of 107 %, and switched temperature per object; the 0.40 mm compensated runs carried the per-object `M221` values. The flow group was printed about four times to fix the flow at each temperature; the whole process took about 50 calibration prints. The ladder spans 195 °C to 265 °C, and the operating range is taken from 195 to 240 °C; Figure 3 shows the printed walls.

The ladder was sliced at 0.45 mm, the slicer's default line width, which sets a nominal width about 12 % wider than the nozzle; at 0.40 mm the slicer's nominal width equals the measured bead. That keeps the extrusion ratio unbiased, and it gives a smaller still reliable bead than 0.45 mm, useful for thin features.

One print speed was used throughout, 25 mm/s at the 0.2 mm layer height and the 0.4 mm nozzle, with the profile's maximum volumetric speed set to 2.3 mm^3/s. The 0.40 mm flow map then varies `M221` per rung. The filament profile's extrusion multiplier and `M221` are multiplicative: `M221` scales the profile value, so a 0.52 profile with `M221` S52 would deliver 0.27, not 0.52. The two are alternatives: use the profile value or carry the flow in `M221`, not both. The calibration held the profile at 1.00 and carried the flow in `M221`; the printed parts used the 0.52 profile and no `M221`. Widths were read with a digital caliper at four points around each specimen's mid-side, and the mean is the width quoted, with the lowest layers excluded where the slicer's Z compensation could change the section.

The non-foaming plateau is the set of rungs whose mean width lies within the caliper resolution of their mean; w_0 is that mean, and $R = \bar{w}/w_0$. The onset is the bracket where R leaves 1, with its midpoint as the onset anchor. The operating range stops where the free width expansion peaks; above that temperature the bead falls.

3.6 Density viewer

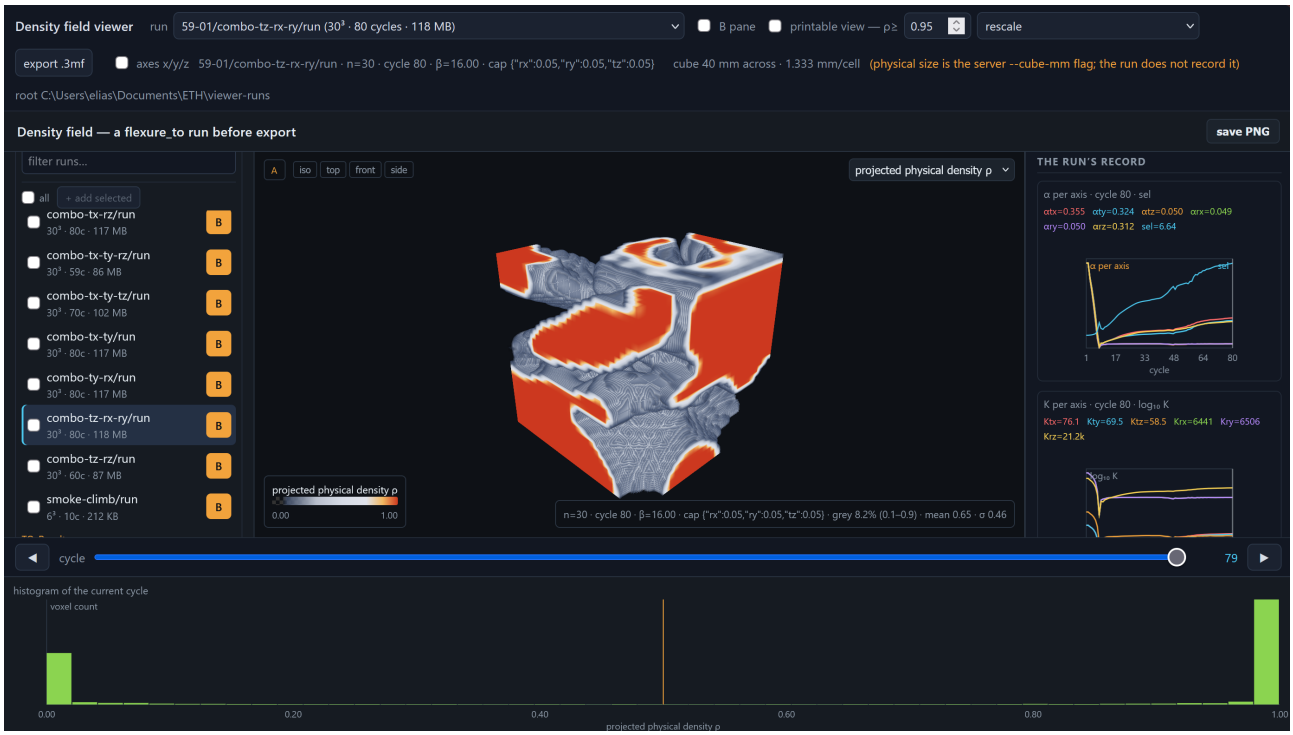


Figure 4: Density viewer. Left: the run list. Centre: the field at the selected cycle; red is solid material, grey intermediate density, and the void is transparent. Right: α and K histories, the normalized stored energy per motion and the secant stiffness. Bottom: cycle slider, density histogram. Example shown: a multi-axis run at cycle 80 (`tz` vertical translation, `rx` and `ry` rotations).

The density viewer (Figure 4) was built during this work to read a run. It takes a run directory in a browser and puts the density field and the run's own record on one page. The field is drawn as a volume, and the intermediate band carries an opaque color of its own, so the cells that a threshold deletes stay visible. Each axis can be sectioned down to a single layer. The cycle slider steps the design through the run, and the charts beside it carry the optimizer's history at the same cycle.

The record rail reads the run's files. It reports where the run ended and why, which motion reached its ceiling at which stroke, and every ceiling reduction with its cycle. A run whose snapshots were never archived still answers from that record. An evaluated run reports W and K per motion for the continuum field and for the exported cell set.

Two runs can be compared side by side, with the two cameras linked. The cycle on screen can then be exported through the same exporter and printability gate the campaign used. The viewer starts with a single command and needs only Python and NumPy, so a run can be read on a laptop that cannot run the solver. A helper copies a run tree from a cluster login host, or from another machine's viewer. It is kept in the repository for the next student and remains a work in progress; it computes no number that this thesis reports.

4 Results

The loop was run to 80 cycles on six single-axis service-load cases and the full-load family, six runs with every motion driven at once, alongside a set of frozen-field checks. The results follow the pipeline: what a cycle costs and what the two solver checks bound, how the design changes with stroke, how far the loop reaches, what the cutoff costs, and what the foam does.

4.1 Solver throughput and checks

Table 3 gives the wall time, the memory and the per-cycle cost of the six-motion loop at $n = 30$, and repeats the 32-rank point as a certification at 6 GB/rank. The loop closes without a license, and its wall time falls as ranks are added until a run has to span nodes. The 32, 64 and 128 rank counts are single-node points, the last of them on a large-memory node; every $n = 30$ row runs the same 80-cycle schedule on the same mesh, so one wall time is compared across the table.

Table 3: Compute and rank scaling for the six-motion, 80-cycle loop. n is the cells per cube edge; ranks are the MPI tasks, one per CPU core. Per cycle is the mean with the median in brackets; speed-up and efficiency are relative to the 32-rank row. The certification and $n = 40$ rows are not rank points, so they carry no speed-up.

Mesh	Ranks	Memory	Wall time	Per cycle	Speed-up	Efficiency
$n = 30$	32	3 GB/rank	**4 h 53 m**	212.6 s (median 163.8)	1.00×	100 %
$n = 30$	64	3 GB/rank	**3 h 05 m**	133.0 s	1.58×	79 %
$n = 30$	128	3 GB/rank, bigmem	**2 h 14 m**	95.9 s	2.18×	55 %
$n = 30$	256	3 GB/rank, two nodes	6 h (timeout, c51)	397.6 s	excluded	excluded
$n = 30$	**32 (certification)**	**6 GB/rank**	**5 h 31 m**	**248 s**	–	–
$n = 40$	32	8 GB/rank	**11 h 41 m**	510.4 s (median 451.0)	–	–

Wall time falls 1.58× from 32 to 64 ranks at 79 % efficiency, and 2.18× from 32 to 128 at 55 %. The baseline is the 32-rank run at 3 GB/rank, 4 h 53 m for its 80 cycles at 212.6 s per cycle and a median of 163.8 s; the certification repeats that rank count at 6 GB/rank to check repeatability and takes 5 h 31 m.

The 128-rank point runs on a different large-memory node from the 32- and 64-rank runs. The 256-rank run is the only one that spans two nodes, because a Euler node holds at most 128 cores; it stopped after six hours at cycle 51, and at 397.6 s per cycle was the slowest of the rank tests.

A single $n = 40$ run prices a higher element count: 80 cycles at 32 ranks in 11 h 41 m, at a mean 510.4 s per cycle and a median of 451.0 s, of which the forward solve takes 425.5 s and the adjoint 14.8 s.

The cross-code check is defined in Section 3.3; Table 10 gives the stored energies and Figure 5 the tangent force. The gap is −7.42 % on `tx`, −6.39 % on `ty` and −10.07 % on `rz`, and it widens to −28.38 % on `rx` and −24.32 % on `ry`; on `tz` it is −35.92 % at 98.11 % of stroke, where the `C3D8RH` arm stopped; the `C3D8H` repeat reached the full 2.5 mm. Abaqus adds an artificial energy of 1.6 % to 11.05 % of its internal energy across the six motions, inside the energy the gap compares.

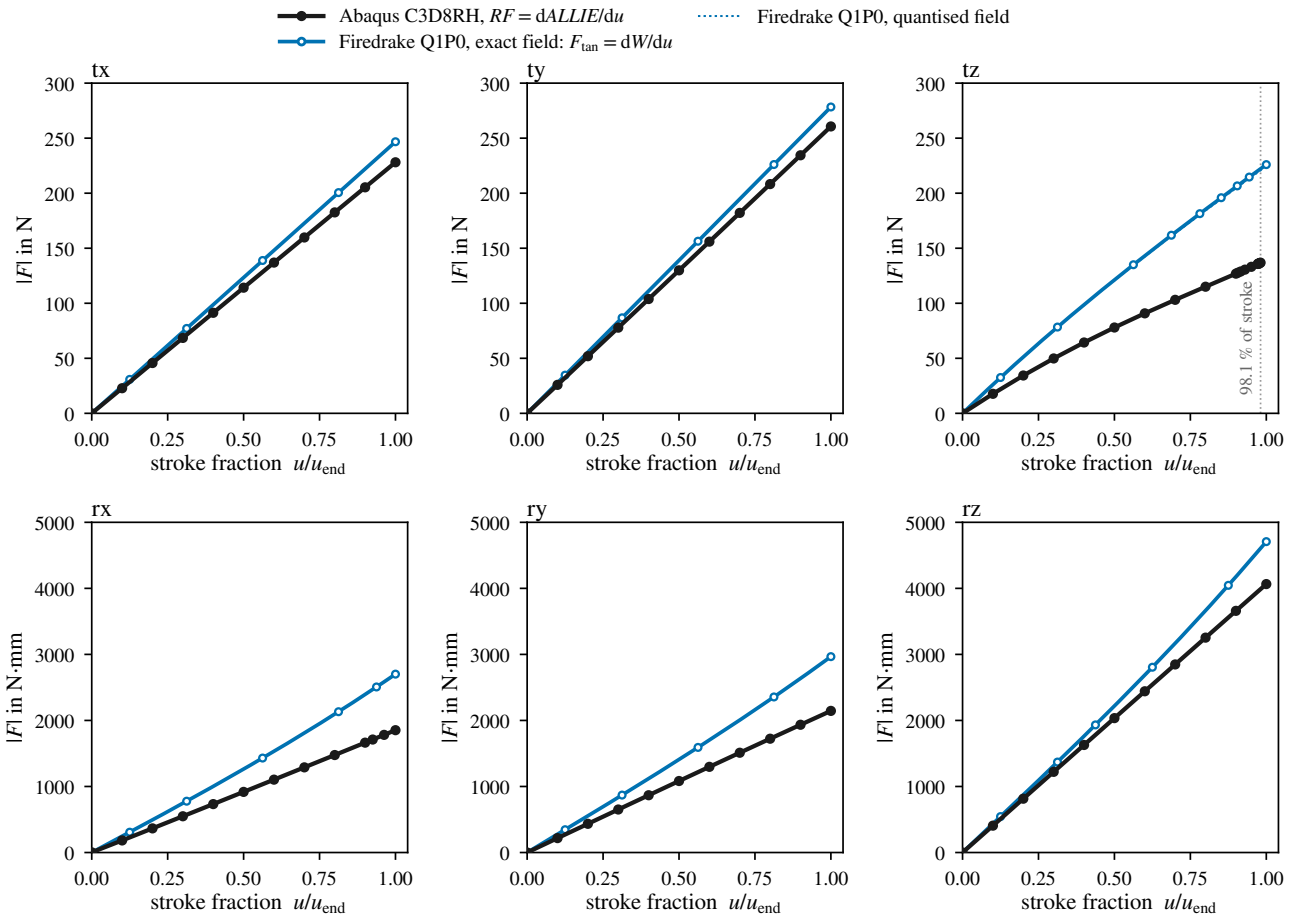


Figure 5: The cross-code comparison: tangent force against stroke on one frozen field at its maximum load, on the same mesh and material. Abaqus gives the reaction force `RF` as the derivative of `ALLIE`, its internal energy (strain plus stabilization); Firedrake gives the tangent force F_{tan} as the derivative of its stored energy W. The dashed curve repeats the Firedrake solve with the modulus from Abaqus's 1000 levels. Horizontal axis: stroke fraction u/u_{end}, u_{end} the maximum load, 2.5 mm or 0.12 rad; the vertical axis is in N for translations and N·mm for rotations.

Two checks stand outside the Abaqus comparison: the taped sensitivity and the continuation. The taped adjoint reproduces the gradient with a worst relative error of 1.1×10^{-8} and 8.3×10^{-9}, and the probe fields reach past the inherited 2.0 to 3.0 mm range, with Table 11 recording the last converged strokes.

4.2 Linear and nonlinear designs

Table 4 gives the four comparisons defined in Section 3.3. The two models stay close at the small stroke and separate at the large one, most on the rotations.

Table 4: r, the Pearson correlation of the final density fields, is 1 for identical fields. N2 is the 2 mm nonlinear design, N7 the 7 mm service-load design and L the linear design; the 7 mm arm ends at 5 mm on `tz`. The N2/N7 rows use the nonlinear repeat band (0.9985 on `tx`, 0.9903 on `rz`, 0.967–0.971 on `ry`); the linear check row uses the linear band (0.9984 on `tx`, 0.4342 on `tz`), since the amplitude drops out of the linear problem.

Comparison	`tx`	`ty`	`rx`	`ry`	`rz`	`tz`
2 mm: N2 vs L	0.988	0.985	0.938	0.809	0.695	0.274
7 mm: N7 vs L	0.937	0.940	0.659	0.693	0.507	0.299
Stroke: N2 vs N7	0.949	0.949	0.707	0.624	0.810	0.877
Linear check: 2 mm vs 7 mm	0.9980	–	–	–	–	0.3565

At 2 mm four axes lie at 0.809 or above, and `rz` at 0.695 and `tz` at 0.274 are the exceptions. At 7 mm the correlations fall into two groups: 0.937 and 0.940 on `tx` and `ty`, then 0.659, 0.693 and 0.507 on `rx`, `ry` and `rz`, and 0.299 on `tz`. A second run of the same design reaches 0.9985 on `tx` and 0.9903 on `rz` at a different rank count, and 0.9713 on `ry` at the same settings after a ceiling rescue.

The linear fields at the two amplitudes sit at the repeat-run value on `tx` ($r = 0.9980$ against 0.9984) and at the scatter level on `tz` (0.3565 against 0.4342), while the 2 mm and 7 mm nonlinear fields correlate at $r = 0.62$ to 0.95 across the stroke change. Only `tx` and `tz` were re-run at 2 mm in the linear model, so the check samples two axes.

4.3 Reaching the target stroke

Figure 6 shows the six service-load geometries in the density viewer at the cycle each was stored. Table 5 gives the stroke each service-load field reached, its end α against the compliance cap, and the cycle at which it stopped when all six motions were driven at once. All six fields pass the print gate. Reach and feasibility separate: the target is reached on five axes, and two of them end above the compliance cap.

Table 5: Stroke and end state of the six service-load runs. Target is the applied motion and Achieved the stroke the run ended on; First ceiling stop is where a run first stopped, with its cycle. α is the end-stroke compliant energy normalized to the solid cube, against the 0.05 cap. The last column gives the cycle at which the field stopped in the separate full-load campaign, with all six motions driven at once.

Axis	Target	Achieved	First ceiling stop	End α (cap 0.05)	All-six stop (cycle)
`tx`	7.00 mm	7.00 mm	–	0.0509	60
`ty`	7.00 mm	7.00 mm	–	0.0501	44
`tz`	7.00 mm	5.00 mm	5.75 mm, cycle 30	0.0514	13
`rx`	0.40 rad	0.40 rad	–	0.0497	45
`ry`	0.40 rad	0.40 rad	–	0.0696	27
`rz`	0.40 rad	0.40 rad	–	0.0499	10

`tz` climbs from 2.00 mm to its ceiling at 5.75 mm at cycle 30, takes three ceiling reductions and settles at 5.00 mm, where its stored field is evaluated; it is the only axis short of the

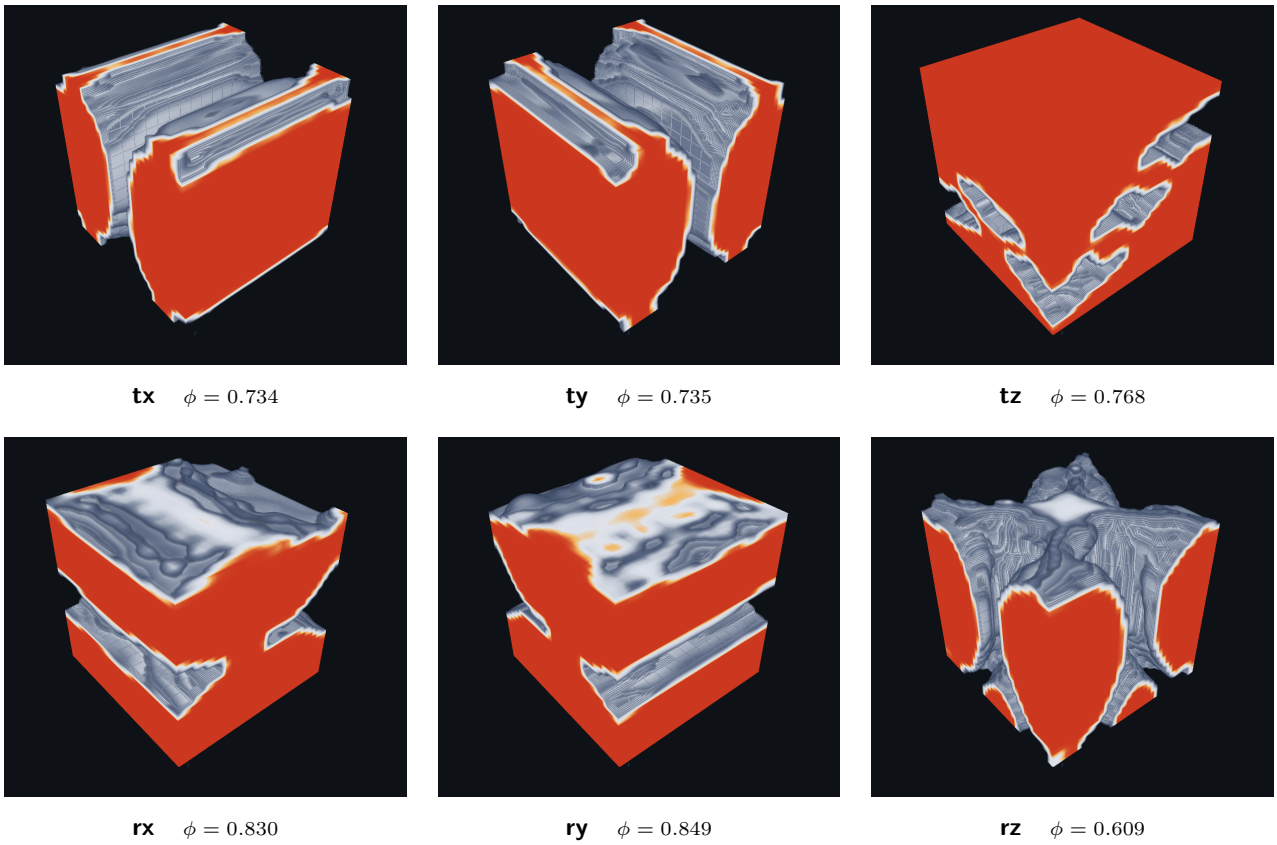


Figure 6: The six service-load geometries: full-field screenshots from the density viewer at the end of each run, labelled by the motion that run drives. Red is solid material, grey intermediate density, and the void is transparent. `tx` and `ty` each form two spanning bodies and the other four a single body; ϕ is the solid fraction at the 0.95 cutoff.

applied stroke.

The full stroke is 7.0 mm, 17.5 % of the cube edge and 0.40 rad, about 23°, both applied to the moved face. On the 1% rule of Section 3.2, two of the five full-stroke fields are over: `ry` at 0.0696 after the projection sharpening, and `tx` at 0.0509; `ty` sits at the cap within that tolerance, and `tz`, which stopped short of full stroke, is also over it.

The full-load runs drive all six motions to 7.0 mm and 0.40 rad at once and fail on every axis, between cycle 10 and cycle 60. The six failures fall into three groups: two near-void ramp stops, one mid-ramp stall and three failures at full load against a resisting motion.

`tz` and `rz` stop the ramp at 4.92 mm and 4.69 mm, at cycles 13 and 10. Both fields are near-void skeletons, both fail the connectivity rule of the gate, and the smallest cell Jacobian in the failing `tz` case is 0.015 with no inversion.

The `ry` run stops at cycle 27 with `tz`, its resisting motion, at 5.92 mm of the 7 mm ramp; the compliant `ry` motion is held at 0.40 rad.

`tx`, `ty` and `rx` reach full load and fail at cycles 60, 44 and 45, each with a stall on a resisting motion, at 0.9913 to 0.9978 of its load on `tx`, 0.9901 on `ty` and 0.9997 on `rx`.

4.4 Print cutoff

Figure 7 shows the change in energy-secant stiffness when each frozen field is replaced by the geometry the exporter produces from it at $\rho \geq 0.5$. Seven fields enter: three service-load fields (`rx`, `ry` and `tz`), two from a mixed-cell variant (`rx` and `ry`), one multi-axis run (`rx`+`ry`+`rz`) and one earlier `tz` climb. They reduce to four distinct geometries, and a separate single-axis linear field is set apart below them. Across the seven the change in stiffness spans $-54.6 \ldots +17.9\,\%$. The exported solid moves the secant stiffness in both directions, and the sign follows the body.

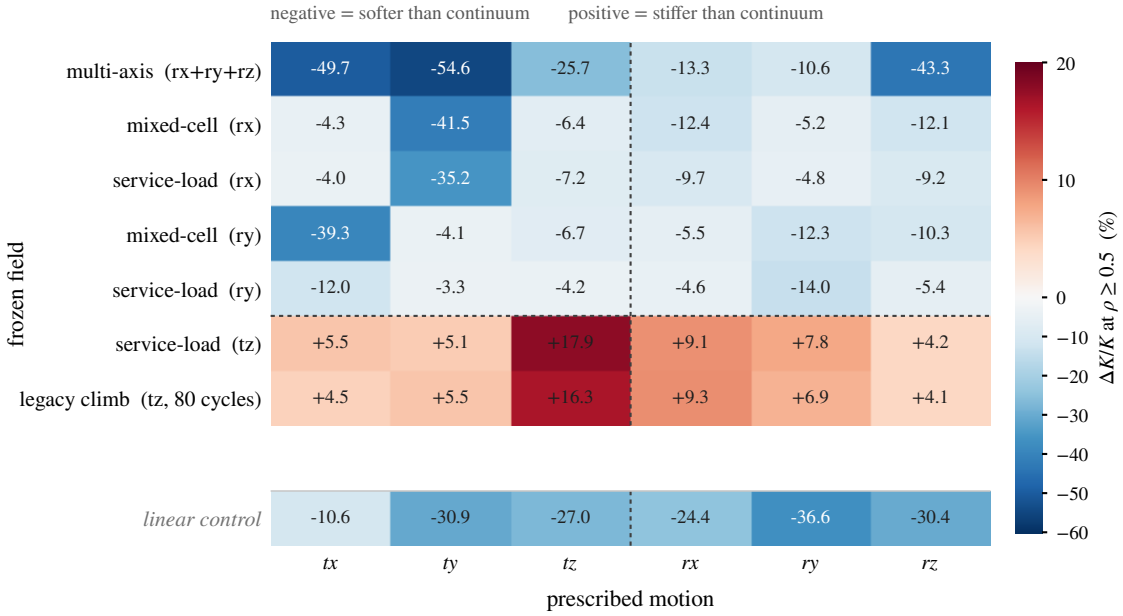


Figure 7: Change in energy-secant stiffness $K = 2W/d^2$ between a design's density field and the geometry exported from it at $\rho \geq 0.5$; the field is held fixed at the end of its run, so only the export changes. W is the stored energy at stroke d. The horizontal axis is the prescribed motion and the vertical axis the frozen field, one run's density field at its last cycle, labelled by run and motion. Each cell is the signed change $\Delta K/K$ at equal stroke: blue is a softer body, red a stiffer one. The single-axis linear field sits below the rule, set apart from the seven.

Five are softer on every solved axis, and the two `tz` fields are stiffer than the continuum field. The range widens as the cutoff rises: at $\rho \geq 0.4$ it is $-48.4 \ldots +21.7\,\%$, with four fields softer, one mixed and two stiffer, and at $\rho \geq 0.6$ it is $-61.7 \ldots +13.2\,\%$, with five softer and two stiffer. Table 15 in the appendix gives every field at all three cutoffs.

The service-load `tz` field is stiffer at all three cutoffs, and its excess falls as the cutoff rises, from 21.7 % at $\rho \geq 0.4$ through 17.9 % at 0.5 to 9.8 % at 0.6. The service-load `ry` field has one positive cell at $\rho \geq 0.4$: `rz` 0.4 %. One cell has no data: the multi-axis field at $\rho \geq 0.6$ on `rx`, where the continuum re-solve stopped short of its load, so the printed body had nothing to compare against. The sweep holds 28 evaluations: 21 for the seven fields, three for the single-axis linear field, three for the legacy ramp (33-01) and one cross-check with the density floor at zero; three of the legacy ramp's evaluations returned a report with the completion flag unset and are not plotted.

4.5 Foam response

Figure 8 shows the measured bead width relative to the unfoamed reference against nozzle temperature, on the ladder sliced at 0.45 mm. The bead holds its reference width across a plateau and widens steeply once the nozzle leaves it. The non-foaming plateau from 195 to 215 °C sets the reference w_0, and the four rungs from 195 to 210 °C lie within the caliper resolution of its mean. The bead leaves the plateau between 215 and 220 °C, and the curve carries an onset anchor at 217 °C, the midpoint of that bracket; the plotted ratio is $R = \bar{w}/w_0$.

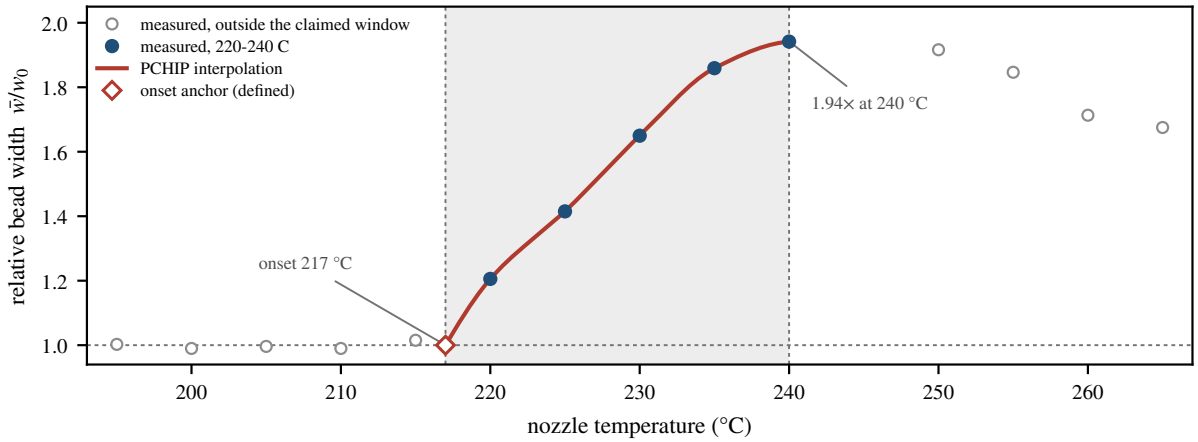
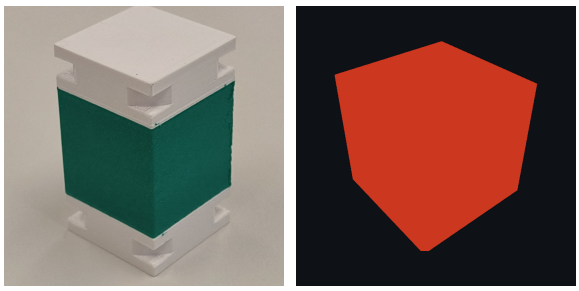


Figure 8: Relative bead width $R = \bar{w}/w_0$ on the vertical axis against nozzle temperature in °C on the horizontal axis, where $\bar{w}$ is the measured mean width and $w_0 = 0.394$ mm the 195–215 °C non-foaming plateau mean. Filled points are the 220–240 °C rungs, open points outside that window; the onset anchor at 217 °C holds $R = 1.00$ at the midpoint of the 215–220 °C bracket. The line is a piecewise cubic interpolation through the anchor and the rungs, and the 250–265 °C points are context only. One volumetric speed ($\approx 2.3\,\mathrm{mm}^3/\mathrm{s}$), $n = 1$.

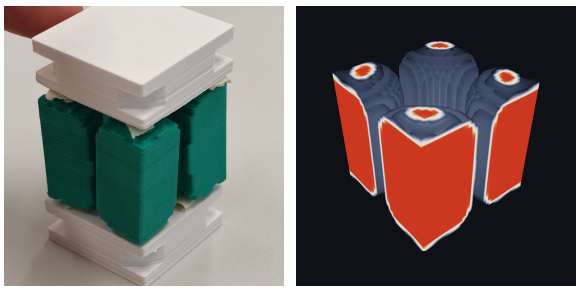
Beyond 217 °C the wall widens steeply, and the free expansion reaches 1.94× at 240 °C, where the measured width is 0.765 mm; the 250 to 265 °C rungs fall to 1.92 and 1.68 times the reference, outside the interpolation window.

The temperature-to-flow map is a step: 125 % at every rung from 195 to 215 °C, 60 % at 220 °C and 52 % at 235 and 240 °C. The proportional pass that set the map holds the compensated bead within 0.06 mm of the 0.40 mm target across the band, and within 0.03 mm from 230 to 240 °C. The final map moved the flows closer to the target but was not printed again. That margin was already small, and the four parts printed at 235 °C with an extrusion multiplier of 0.52 show no visible over- or under-extrusion, so a second ladder would have spent time and filament to shrink it further. These deviations are therefore an upper bound. Table 12 lists every rung from 195 to 265 °C, with the measured width, relative width, extrusion ratio and compensated `M221` setting.

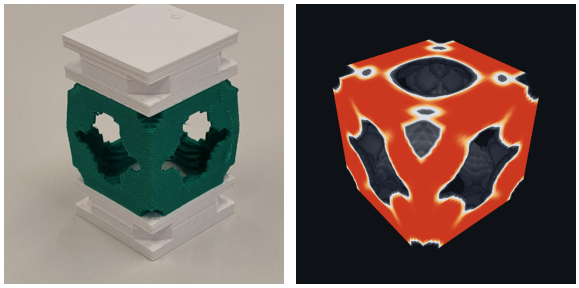
The print job of Section 3.4 also produced the four parts in Figure 9: three flexures and the solid reference, with VarioShore bodies at 235 °C and extrusion multiplier 0.52. They are from the earlier campaigns; none of the six service-load fields was printed, and none of the four was measured.



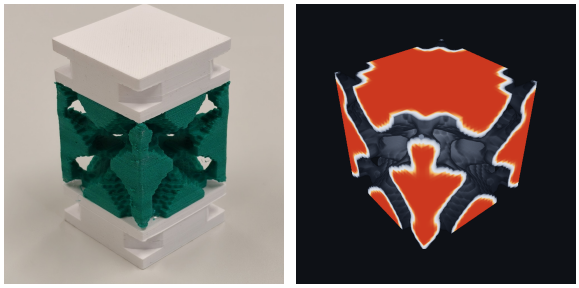
solid cycle 1, cut 0.50



txty cycle 80, cut 0.10



txtytz cycle 80, cut 0.50



tzrz cycle 61, cut 0.20

Figure 9: The four printed parts, made to demonstrate the calibrated filament and for physical tests that were planned and never carried out. VarioShore bodies at 235 °C and extrusion multiplier 0.52, each beside its projected density field at its export cut. Three are flexures; `solid` is the reference cube. The two-translation part broke at the TPU–PETG interface because too little dense TPU met the top plate, and is taped for the photograph.

5 Discussion

This section discusses the five questions of Section 1.1 in order, the first two together. Each subsection states the standard its results are judged against and what the evidence does not settle; Section 6 collects the verdicts. The five results do not share a unit (a compute time, a correlation, a stroke, a stiffness change and a temperature window), so they are taken one at a time.

5.1 The solver's bound and the model's change

This subsection answers the first two questions: the cost of the replacement is the first, and the reach of the inherited model the second. The replacement is fast, and it costs only compute: an 80-cycle loop takes 4 h 53 m on 32 ranks and 3 h 05 m on 64, with no license to check out. A 32-rank repeat at 6 GB/rank on a different node moves the per-cycle time by 17 %; the same-setting 3 GB repeat moves it by 8 %. The 1.58× is larger than both.

64 ranks is the point to choose, and 128 buys wall time when the queue is the constraint. The cost follows the mesh: 2.4 times the per-cycle time for 2.37 times the cells, which makes an $n = 40$ run a night.

The cross-code check asks how far the replacement moves the answer on the one field both codes can run. On the energies the switch moves it: Abaqus stores 6 to 36 % less than Firedrake, depending on the motion.

Neither code is checked against an exact solution, so the gap measures the difference the swap introduces. It does not identify which code is in error; it could be Firedrake, Abaqus, or both in opposite directions. At 36 % the difference is large enough to change the stiffness read from the field, which is why the element was tested next.

The element was the obvious suspect, and it was tested: repeating the Abaqus side with `C3D8H`, which matches Q1P0 in interpolation, quadrature and constraint, moves the gap by at most 1.6 points and cuts the artificial-energy share by a factor of three without moving the total. Hourglass control and integration are therefore excluded.

`C3D8H` also reached the full 2.5 mm on `tz`, where `C3D8RH` stalled at 98.11 % of stroke, so the element decides whether the continuation completes, while the energy gap persists. What remains is the formulation on one side or the other, the hybrid implementation, the material evaluation or the void treatment, and this check cannot say which.

The model question is different, and the four comparisons answer it only in part. At 2 mm the nonlinear optimum sits nearer the linear one on five axes than at 7 mm. That is what an implementation correct in the small-strain limit should do.

No repeat band was run at 2 mm, so each axis is read down its own column and the claim rests on the change from 2 mm to 7 mm. At 7 mm it moves: the correlation falls to 0.507 to 0.693 on the rotations while it stays at 0.937 and 0.940 on the translations.

Repeating a design sets the band: 0.9985 on `tx`, 0.9903 on `rz` and 0.967 to 0.971 on `ry`. The translations fall 0.06 below their band and the rotations 0.30 to 0.48, beyond the scatter of a repeated run. The linear model therefore moves the design on every axis that can be read, and it moves it most on the rotations.

The `ry` band is the loosest of the three, and `ry` at 0.693 is the closest rotation to the translations, so the claim rests on that pair. One of the three repeats behind it took a ceiling reduction partway up and still correlates at 0.971, so the band already covers a changed load history.

That is a change of optimum; it does not show that the linear model predicts the wrong force. To show the latter, the linear design would have to be re-solved at 7 mm and shown to miss the target, and that was not done.

r compares the two density fields cell by cell and weights each cell by its distance from the field mean, so the solid members carry the sum. A value below the band says the two models share less of the pattern, and the stiffness either body delivers follows from its own solve. The band rests on two rank pairs and three reruns, so it sets a scale for the comparison with no confidence level attached.

The choice of model therefore moves the design on every axis that can be read, most of all on the rotations. A designer who uses the linear model should expect a different body there. `tz` cannot be read the same way: its service-load field stopped at 5.00 mm while the linear field is at 7 mm, so its 0.299 mixes the model with a shorter stroke.

Trachsel's stop at 2.0 to 3.0 mm is the nearest reference, though the solver, element, mesh, field, objective and continuation all differ.

5.2 Reachability and feasibility

The third question is whether the pipeline reaches the target stroke, and Table 5 gives the stroke and end compliance of the six service-load fields; the first thing it establishes is that they are six single-axis designs, each with its own part. Each comes from its own optimization with one motion climbing and the other five held at 2 mm and 0.1 rad.

The resisting amplitude follows the use: a flexure moves at full stroke in its compliant direction and holds its five DOCs at a small stroke. The value sits at or below the lower bound of the range the inherited work reached, so only the driven axis goes past what earlier runs covered (Trachsel 2026). It cannot be zero, because a direction held at zero stores no energy and drops out of the objective. The five-axis count is therefore a reachability result about the pipeline.

Reachability and feasibility come apart at full stroke. `ry` and `tx` reach the target and end above the 0.05 cap, `ry` at 0.0696, so a part that meets both the stroke and the constraint would have to be optimized again under a cap it can meet. The five-axis count is a statement

about stroke; whether a design meeting both stroke and cap exists is untested.

The full-load runs drive all six motions at once and fail on every axis, between cycle 10 and cycle 60. The six failures fall into three groups: two stop on a near-void skeleton that fails the connectivity rule, one stalls short of its stroke, and three fail against a resisting motion at nearly full load. The stall is a stroke shortfall: in the two ramp groups the driven motion stops, and in the third the resisting motion does. The model has no contact, so the void phase stands in for the air gap and compacts toward very low volume at full stroke; that is an argument about the model, and whether the design can be built is untested.

`tz` is the axis that falls short in every stroke list, and it is the only motion whose normal lies along both the fixed and the moved face, so it changes the volume of the flexure. The solid cube stores 5.1 times as much energy in `tz` as in `tx` at the same 2 mm stroke, so `tz` loads the cube in bulk where the in-plane axes load it in shear. The flexures that optimize for it form members at about 45°, which Karrer reads as compression paths that stay stiff in the other motions (Karrer 2026); the geometry is observed and the reading is a hypothesis. The design landscape on that axis is the least stable as well: three 2 mm starts correlate between −0.10 and 0.29 with each other. Layer orientation does not explain it. Karrer's flat and vertical tensile specimens differ by about 18 % in modulus, this thesis measured none, and the simulation does not carry it (Karrer 2026). Q1P0 does not explain it either: the element was chosen to avoid volumetric locking. What remains is the kinematic account, with no run that isolates it.

5.3 The cutoff's measured step

The fourth question is what the threshold costs in stiffness, and Figure 7 shows the change in the energy secant when each field is replaced by the body the exporter produces from it at $\rho \geq 0.5$. The question behind it is whether the threshold softens the body, and the answer is that it moves it both ways: five fields lose stiffness on every solved axis, and the two `tz` fields gain it.

The change can be attributed term by term, and Table 14 in the appendix gives the hardening, removal, repair and floor lift terms for each field. Hardening is the stiffness change from setting the kept cells to the solid modulus, removal the change from dropping the sub-threshold cells to the printed void, and repair the change from the cells added to remove diagonal-only contacts. The floor lift is the part of removal that comes from the printed void being stiffer than the void of the continuum map; the table's first-order value reaches 26 points, and the re-solved share is 0.4 to 12 points. On the five softer fields, removing sub-threshold cells dominates on every axis, by 3 to 17 times the hardening of the cells that are kept: removal reaches −58 % while hardening adds at most 6.2 %. The removal is largest off the compliant axis in four of the five, so the body does not soften uniformly. That pattern is consistent with the lateral connections being the thinnest, where a sub-threshold skin carries part of the path; member thickness was not measured, and the reading stays a hypothesis. Culling removes no cell on any of the seven.

On the two `tz` fields the balance reverses: hardening dominates on every axis, by 4 % to 10 %, removal is near zero away from `tz`, and repair adds at most 2.6 %. The reversal fits a second reading, in which the two `tz` fields carry their compliance in a low-density phase: removing it does not soften them, and the density the kept cells gain dominates. The export therefore separates `tz` again: the two fields that gain stiffness are the two whose compliant motion is `tz`, and the balance of the terms reverses with them.

The floor lift is not physical. The re-solve lifts void to $\rho = 0.001$, so a removed cell carries 0.025 MPa, against the 0.004 MPa of the map. On the `tz` axis of the two `tz` fields it is 64 and 72 % of the entire positive number; without it the gain falls to 6.4 % and 4.5 %, so the sign holds and the size does not.

No single correction factor represents the step. The terms change sign between fields and change size within a field, and their first-order sum does not close to the measured change, because removing a cell changes the load path as well as its modulus; the closure is the re-solved chain, which is what the figure reports. What the threshold does is a property of the body, and a designer who needs the printed stiffness has to carry it as a per-body step.

5.4 The material's operating range

The fifth question is which settings set the bead, and Figure 8 gives the bead width against nozzle temperature, which describes a process as much as a material. The plateau ends between 215 and 220 °C, the free expansion peaks at 1.94× at 240 °C, and with the flow map applied the compensated bead stays within 0.06 mm of the 0.40 mm target across the band. The final map moved the flows closer to the target, and it was not printed again, so that deviation is an upper bound. The same filament therefore prints a range of widths, and the printer, the speed and the map are part of the number.

Because the map is keyed to the object and the nozzle holds its setpoint before each one, a single plate can carry several foaming levels. The ladder printed its rungs that way; the printed bodies here use one setting, so the calibration demonstrates the capability without grading a part.

Two other calibrations of this filament put the onset lower. The datasheet's own compensation is nominal at 210 °C and 70 % at 230 °C, which places its onset between those settings, and the 2020 test puts the onset at about 200 °C. Neither can be adopted as the setting for this machine, and the measured bracket sits inside the datasheet range and above the 2020 test.

The gap to the 2020 test has two candidates. The first is the higher volumetric flow used here, since the same temperature gives different expansion at different flow; the second is a change in the formulation of the filament over six years. Neither was measured; a slower 220 °C shakedown print reached a lower density, consistent with more expansion at lower flow, but the onset itself was not measured below 217 °C.

The fall above 240 °C is not explained by this measurement. For a foaming filament

in material extrusion the same fall is described as coalescence, collapse and then degradation (Grigolato et al. 2025). This work observed the fall; the mechanism is outside this measurement, and the only other calibration of this filament that runs past 240 °C reaches a plateau instead, so a width-only measurement leaves a sag artifact open (Wade et al. 2026).

The cold end of the map is more ordinary than its 125 % suggests. Under-extrusion at the cold end is documented for this filament (Wade et al. 2026), and about half of the 25-point rise is the difference between the 0.45 mm nominal and the 0.40 mm target; each print was sliced at its own width. The step from 125 to 60 % across the 5 °C between 215 and 220 °C is the 5 °C resolution of a steep onset, sampled too coarsely to resolve as a continuous transition. The map is fitted to the measured widths. A pure inversion of the free-width curve would set the flow at 64 % at 240 °C, where the measured compensation is 52 %. Part of that is the different slicing of the compensated print, part an expansion that depends on flow as well as temperature, and this thesis does not separate them. Until the modulus of the foamed solid is measured, the material step stays an operating range.

5.5 Limitations

Five limits bound what the claims above can carry.

1. **The validation is on benchmarks.** The independent check is one geometry at one mesh, and the Abaqus comparison is one frozen field at 2.5 mm and 0.12 rad. The gaps run from 6 % on `ty` to 36 % on `tz` in internal energy; hourglass control and the integration rule are excluded by the `C3D8H` repeat, and the remainder is not attributed further. The Firedrake solver should therefore be read as bounded on both counts; validation against a reference code is outside this work.
2. **The envelope check is a probe.** No optimized design was driven past its ceiling, and Table 11 in the appendix holds two frozen fields. It shows that the solver can run past the inherited range; the new limit and a mesh-convergence study are outside the probe.
3. **The resisting amplitude is a modeling choice.** The design is optimized and evaluated with the other five motions at 2 mm and 0.1 rad, so its behavior at a larger stroke in the resisting directions is untested, and full stroke is cap-infeasible on two of the five axes, `ry` by 39 % and `tx` by 1.8 %.
4. **The export is a simulation result.** The budget rests on seven fields at $\rho \geq 0.5$, and the sweep covers three cutoffs (Table 15). Four parts were printed, three of them flexures, and none was measured; no modulus-density relation exists. The stiffness change is therefore a prediction about a body that was never tested on a printed specimen.
5. **The material is an operating range.** The modulus of the foamed solid was not measured, the model carries 21 MPa from FiberFlex 40D, and the 2.8 MPa blend boundary comes from Filaflex 60A. The map holds at one volumetric speed, for this printer and this filament, was not printed again after its final flows, and the material's share of the gap is not quantified.

6 Conclusion

The gap this thesis inherited was one unexplained difference. It is now a sequence of steps, each with its own sign, unit and condition, and the useful result is that sequence: no two steps measure the same quantity, so they cannot be summed into a single correction.

The solver's step is bounded on the one benchmark both codes can run. The loop runs without a license and gains speed as ranks are added; on the frozen field the stored energy differs by 6 to 36 %, and the `C3D8H` repeat excludes hourglass control and the integration rule and leaves the residual unattributed. The modulus is a field at every density and is differentiated like one, which frees the design from the thousand discrete sections of the inherited loop. The model's step is a change of design. Each of the five readable service-load fields falls below its repeat band, the translations by 0.06 and the rotations by 0.30 to 0.48, so the linear model moves the design on all five, most on the rotations. `tz` is confounded by its shorter stroke. Past the inherited window of 2.0 to 3.0 mm in translation and 0.12 rad to 0.14 rad in rotation, the loop reaches 7.0 mm and 0.40 rad on five axes, with `tz` at 5.00 mm and two axes above the compliance cap.

The export's step is a measured percentage. At $\rho \geq 0.5$ the energy secant moves between 54.6 % lower and 17.9 % higher than the frozen field, five fields soften on every axis and the two `tz` fields stiffen, so the sign depends on the body. The material's step is an operating range. Free width expansion reaches 1.94× at 240 °C, the map holds the compensated bead within 0.03 mm of the 0.40 mm target from 230 to 240 °C, and no modulus stands behind either number. Nozzle temperature controls the expansion and the extrusion multiplier controls the bead, from 125 % on the 195 to 215 °C plateau down to 52 % at 235 to 240 °C, so a chosen width is realized by the pair at each rung. Four parts were printed, three of them flexures, and none was measured; every stiffness in this thesis is a simulation result. Measuring the modulus of the foamed solid at 235 °C, extrusion multiplier 0.52 and a 2.3 mm^3/s cap is the first experiment a next attempt should run, because it is the step that still has a range and no number. Until then the gap narrows and does not close.

7 Future Work

These follow-ups could partly close one of the limitations of Section 5.5.

1. **One measured flexure.** Three flexures and the solid reference were printed and none was measured, so the export's percentage is a prediction about bodies that were never loaded. Loading one of the three flexures at a working amplitude on the robot arm and comparing its secant with the re-solve of its own field would test the sign of the export step on a real part.
2. **A data-driven correction.** The gap cannot be attributed to a single number, and the changes vary between components. A model fitted to measured parts could capture more of the gap than one correction factor.
3. **The modulus of the printed VarioShore.** The simulation carries 21 MPa from FiberFlex 40D because the modulus of the printed VarioShore was never measured, so the operating range cannot be turned into a stiffness. One specimen printed at 235 °C, extrusion multiplier 0.52 and a 2.3 mm^3/s cap, measured on the robot arm with the partly documented protocol, would give a first data point on that modulus.
4. **Density above the onset.** The density half: printed density was not measured above the onset, so width is the only measure of the foaming level there. Weighing specimens printed at the temperature rungs would pair each density with the width already measured. Combined with the modulus measurement above, that would give a map from density to material stiffness.
5. **A graded part from the optimizer's grey band.** The optimizer designs intermediate densities that the export threshold discards, and VarioShore can realize them as several foaming levels in one print; the slicer route of Wade et al. (2026) is one way to place them, and the idea was developed here before that paper was found. This thesis calibrated the range and printed one level. The smallest test is a bar carrying two or three levels from one field, measured by width first and by stiffness once the modulus exists.
6. **A minimum interface-stock constraint.** The printability half: the gate judges the exported body, while the part that is printed is that body plus its fixture plates. The two-translation `txty` part, exported at $\rho \geq 0.10$ from a linear run, broke at the top-plate interface because too little dense TPU met the plate; the nonlinear runs also leave little dense material at the movable face. A rule in the loop that requires a minimum amount of dense material at each plate interface would make the verdict a statement about the part itself and prevent broken flexures.
7. **Contact, or a better void formulation.** The model has no contact, so the void phase stands in for the air gap and compacts toward zero volume when a resisting motion is driven to full stroke, which creates some of the problems making this loop unstable.

References

Bambu Lab (2024). *Bambu Lab PETG HF — Technical Data Sheet*. Tech. rep. URL: https://store.bblcdn.com/3a230e260a3a47c2b0db0156e07eef91.pdf.

Bendsøe, Martin P. and Ole Sigmund (2004). *Topology Optimization*. en. Berlin, Heidelberg: Springer. ISBN: 978-3-642-07698-5. DOI: 10.1007/978-3-662-05086-6. URL: http://link.springer.com/10.1007/978-3-662-05086-6 (visited on 11/02/2026).

Bonet, Javier and Richard D. Wood (Mar. 2008). *Nonlinear Continuum Mechanics for Finite Element Analysis*. 2nd. Cambridge University Press. ISBN: 978-0-521-83870-2. DOI: 10.1017/cbo9780511755446. URL: https://doi.org/10.1017/cbo9780511755446.

colorFabb (2023). *colorFabb varioShore TPU — Technical Data Sheet*. Tech. rep. URL: https://colorfabb.com/media/datasheets/tds/colorfabb/TDS_E_ColorFabb_varioShore_TPU.pdf.

Dassault Systèmes (2016). *Abaqus Analysis User's Guide*. Abaqus 2016 documentation (HTML). Sections cited: 28.1.1 "Solid (continuum) elements"; 27.1.4 "Section controls". Section titles verified against a public mirror of the 2016 HTML guides; the official copy is login-gated at Dassault Systèmes. Copyright notice in the documentation reads "Dassault Systèmes, 2015". URL: https://ceae-server.colorado.edu/v2016/books/usb/default.htm.

Fiberlogy (2026). *Fiberlogy FiberFlex 40D — Technical Data Sheet*. Tech. rep. URL: https://fiberlogy.com/app/uploads/2026/05/FIBERLOGY_FIBERFLEX40D_TDS.pdf.

Flanagan, D. P. and T. Belytschko (May 1981). "A uniform strain hexahedron and quadrilateral with orthogonal hourglass control". In: *International Journal for Numerical Methods in Engineering* 17.5, pp. 679–706. ISSN: 0029-5981. DOI: 10.1002/nme.1620170504. URL: https://doi.org/10.1002/nme.1620170504.

Goronzi, Karl (2026). "A Computational Framework for Anisotropic Stiffness Design Using Topology Optimization and Abaqus Scripting Interface". en. Bachelor's thesis. Zurich: ETH Zurich.

Grigolato, Luca and Gianpaolo Savio (Sept. 2025). "Development of functionally graded properties components via filament foaming in material extrusion". en. In: *International Journal on Interactive Design and Manufacturing (IJIDeM)*. ISSN: 1955-2505. DOI: 10.1007/s12008-025-02395-3. URL: https://doi.org/10.1007/s12008-025-02395-3 (visited on 09/17/2026).

Hermann, Stefan (Nov. 2020). *colorFabb varioShore TPU: Foaming Filament*. en. URL: https://www.cnckitchen.com/blog/testing-colorfabb-varioshore-tpu-foaming-3d-printing-filament (visited on 10/27/2026).

IDEALLab (2026). *TO fixtures (flexure_to)*. URL: https://github.com/IDEALLab/TO_fixtures.

Karrer, Kenneth (2026). "Improving a Topology Optimization Pipeline for the Design, Manufacture and Validation of 3D-Printed Flexures". en. Bachelor's thesis. Zurich: ETH Zurich.

Koppen, S., M. Langelaar, and F. Van Keulen (June 2022). "A simple and versatile topology optimization formulation for flexure synthesis". In: *Mechanism and Machine Theory* 172, p. 104743. ISSN: 0094-114X. DOI: 10.1016/j.mechmachtheory.2022.104743. URL: https://doi.org/10.1016/j.mechmachtheory.2022.104743.

Lang, Wisnu (2026). "Development of a Topology Optimization Workflow for Flexure Design". en. MA thesis. Zurich: ETH Zurich.

Nagtegaal, J.C., D.M. Parks, and J.R. Rice (Sept. 1974). "On numerically accurate finite element solutions in the fully plastic range". In: *Computer Methods in Applied Mechanics and Engineering* 4.2, pp. 153–177. ISSN: 0045-7825. DOI: 10.1016/0045-7825(74)90032-2. URL: https://doi.org/10.1016/0045-7825(74)90032-2.

Nofar, Mohammadreza et al. (Feb. 2022). "Foam 3D Printing of Thermoplastics: A Symbiosis of Additive Manufacturing and Foaming Technology". In: *Advanced Science* 9.11. ISSN: 2198-3844. DOI: 10.1002/advs.202105701. URL: https://doi.org/10.1002/advs.202105701.

Polymaker (2025). *Polymaker PolyDissolve S1 — Technical Data Sheet*. Tech. rep. URL: https://polymaker.com/wp-content/uploads/TDS_Polymaker_PolyDissolveS1_V5.5_2025-12-10_EN.pdf.

Svanberg, Krister (Feb. 1987). "The method of moving asymptotes—a new method for structural optimization". In: *International Journal for Numerical Methods in Engineering* 24.2, pp. 359–373. ISSN: 0029-5981. DOI: 10.1002/nme.1620240207. URL: https://doi.org/10.1002/nme.1620240207.

Trachsel, Federico Sanchez (2026). "Shaping the Force-Displacement Response of Topology Optimized 3D Printed Flexures". en. Bachelor's thesis. Zurich: ETH Zurich.

Wade, Charles, Devon Beck, and Robert MacCurdy (July 2026). *Functionally Grading the Slicing Process by Compiling Design Intent into Slicer Projects*. arXiv:2607.25326 [cs.GR]. DOI: 10.48550/arXiv.2607.25326. URL: http://arxiv.org/abs/2607.25326 (visited on 09/17/2026).

Yeoh, O. H. (Nov. 1993). "Some Forms of the Strain Energy Function for Rubber". In: *Rubber Chemistry and Technology* 66.5, pp. 754–771. ISSN: 1943-4804. DOI: 10.5254/1.3538343. URL: https://doi.org/10.5254/1.3538343.

A Appendix

A.1 AI Use Declaration

To ensure transparency and uphold academic integrity, Table 6 outlines the artificial intelligence (AI) tools used in the preparation of this thesis, with their use case, scope and remarks. This declaration aligns with ETH Zurich's guidelines on the responsible use of AI in scientific writing.

Table 6: The AI tools used in the preparation of this thesis, with use case, scope and remarks.

AI-Based Tool	**Use Case**	**Scope**	**Remarks**
OpenCode / OpenChamber (DeepSeek V4.1 Flash)	Drafting and editing, literature search and summarisation, LaTeX formatting, code development and debugging, analysis scripts, figure preparation	Whole repository: all chapters, code, analysis and figures	Agent sessions under the author's direction; text and results reviewed by the author; last access: 02.11.2026

A.2 Code and data

The pipeline is `flexure_to`, a Python package that solves on Firedrake and writes the watertight `.3mf` the parts are printed from. The version used here is tag `v1.5.2` of https://github.com/IDEALLab/TO_fixtures (IDEALLab 2026). The stored runs record or rebuild the settings they used, as described below; a re-run at that tag is not claimed to be bit-identical. The Abaqus pipeline this work replaced is kept at tag `abaqus-final`. The licence is MIT except the optimiser in `flexure_to/mma.py`, which is GPL-3.0, so the combined work as distributed is GPL-3.0 and removing that one file leaves an MIT work.

Three commands carry the workflow: `check` validates a run file without solving, `run` executes it inside the Firedrake container, and `export` turns a finished run directory into the `.3mf` and its printability gate. The recipe behind the certified run is `examples/climb-tz/README.md`, and the fast example exports a gate-passing part on a host that has neither Firedrake nor a compiler.

A strain-energy cap is a ratio of secant forces, so a desired force or moment at a named stroke converts into one, and back, without a solve; the force it yields holds only at that stroke. `tools/force_target.py` does both from a run's own tabulated solid reference, and reports the terminal force beside the secant one, which differ where the response is not a straight line.

The repository ships the sixteen production runs above and a reproduction bundle indexed by `research/README.md` and `research/INDEX.csv`. The bundle adds the frozen-field comparison behind Table 10 and Figure 5, the 28 printed-field evaluations and nine cutoff fields behind Tables 14 and 15 and Figure 7, the six full-load histories behind the last column of Table 5, the six service-load cycle-80 fields behind Table 4, and the recorded rank, repeat and probe tables behind Tables 3, 4 and 11. `research/reproduce.py` prints the checks the stored files support, and `research/PROVENANCE.md` names the code head and the SLURM job behind each stored run. Table 7 names every campaign, and a stored run's directory carries the campaign's identifier: `58-tx` is the `tx` run of `58-01`. The six service-load fields of Section 4.3 were produced at `a4390f5`; the gate verdicts of Table 5 at `6a7f428` and `256caa8`, which come after it.

The reproduction bundle carries the linear-control density fields and the 2 mm screen's cycle-80 fields, so the four comparisons of Table 4 recompute from the shipped fields; the `ry` entry is 0.624. The shipped fields give 0.935, 0.939 and 0.660 for the `tx`, `ty` and `rx` cells of the 7 mm row, against the printed 0.937, 0.940 and 0.659, a difference of at most 0.002. The rank-scaling run records are shipped as well. The raw probe runs are the one input that did not survive the scratch purge, so the probe table is a recorded summary, and `research/PROVENANCE.md` names that limit. The bundle also records the one place where a stored field and a printed value still differ: the compensated-bead range of Section 4.5 is read at the measured 62 % flow row while the shipped map node is 52 %.

A run says how it was made. Every run writes its resolved run file, and a run written by the current code records the same settings in `summary.json` as well, so one file names every

setting that entered the loop. The earlier runs record about a fifth of them there; for those, `tools/rebuild_run_yaml.py` rebuilds the run file, and its header names every key the run did not record, so each one can be compared against the code head named for that run in `research/PROVENANCE.md`. The distinction is deliberate: a default is what the released code uses when a run file omits a key, so the pair still describes the run, and it is not presented as the run's own record. Where a default moved between a run's head and the released tree, the rebuilt file carries the value the run used, taken from the summary where it recorded one and from the run's own history and ceilings where it did not.

The printed record is Table 13: four parts, three flexures and the solid reference, and none was measured. The photographs and sliced G-code are deliberately omitted: they belong to a companion data archive, which is not linked here. The cluster state is otherwise not shipped: beyond the cycle fields named above, the per-cycle snapshot series, the full run directories and the container homes were not retained, and `research/PROVENANCE.md` says so wherever it applies.

Table 7: The run matrix. Every simulation campaign behind a number in this thesis, one row per campaign: purpose, axes, mesh n (cells per cube edge), cycles, Euler job and outcome. A job cell carries the Euler identifiers or, where they are not repeated here, the ticket number; the four printed parts are not run rows (Table 13), and the solver probes are in Table 11.

Campaign	Purpose	Axes	n	Cycles	Job(s)	Outcome
Service load (58-01)	the service load case	six, one per run	30	80	17701364 tx, 17685221 ty, 17701370 tz, 17701372 rx, 17701373 ry, 17685218 rz	5 axes at full stroke; `tz` stall 5.75 → 5.00; all gate PASS
Full-load contrast (54b-01)	all six driven at once	six, one per run	30	80	ticket 54	mixed; `tz` reduced to 3.5, `rz` out of memory
Full load (46-01)	where the full load fails	six, one per run	30	80	ticket 46	stalls `tx` 60, `ty` 44, `rx` 45, `ry` 27, `tz` 13, `rz` 10
Small-motion screen (49b-01)	the 2 mm design family	six, one per run	30	80	17104681 tx, 17104682 ty, 17104683 tz, 17104684 rx, 17104685 ry, 17104686 rz	the model comparison's 2 mm designs
Linear design (15b-01, 04-01)	the same problem, linear	six, one per run	30	80	16304017 tz, 16393568 tx, 16393574 ty, 16393577 rx, 16393581 ry, 16393583 rz	both switches off; 7.0 mm / 0.40 rad
Band repeat (`ry`)	run-to-run scatter	`ry`, three repeats	30	80	19296799 A, 19296802 B, 19296803 C	r 0.967 to 0.971; A rescued at c61
Element check (c3d8h-01)	element or solver	six, `tz` split	40	1	19322842, 19322843	gap moves $\leq$ 1.6 pp; `tz` reaches 2.5 mm
Mixed-cell jump (57-01)	full stroke, screened DOCs	six, one per run	30	80	ticket 57	`rx`, `ry` full stroke and gate PASS; `tx`, `ty` fail the gate; `tz` 4.50 mm; `rz` 0.34 rad
Multi-axis support (59-01)	three rotations at once	`rx`+`ry`+`rz`	30	80	ticket 59	full rotations, 80 of 80

Campaign	Purpose	Axes	n	Cycles	Job(s)	Outcome
Legacy ramp (33-01)	an earlier ramp, cross-check	`tz`	30	24	`16778035`	not one of the seven fields; its three evaluations are a control in the sweep
As-printed re-solve (33-01 c24)	the exported field	`tz` design, six evaluations	30	24	eval `16977875`	energy-secant deltas
Rank scaling	cost against rank count	`tz`	30	80	`18379028` a2p 32, `18379029` 64, `18379030` 128, `18380070` 256, `18422593` g5 32	1.58×/2.18×; 256 excluded
Mesh refinement	cost at a finer mesh	`tz`	40	80	`18734745`	80 of 80; 2.40× the $n = 30$ cycle
Cross-code check (59)	the solver comparison	frozen field	40	1	`18445698` arm A, `18459968` arm B five motions, `18477228` arm B `tz`	Abaqus 6.4 % to 35.9 % below Firedrake; `tz` at 98.11 % of stroke
Release certification	the certified run, repeated	`tz`	30	80	`17833892`	5 h 31 m; α_{tz} 0.05154; code head `f5c87c8`
Foaming ladder	bead width vs temperature	printed walls	–	–	print jobs not recorded	195 °C to 265 °C, $n = 1$ per rung

The four printed parts are not run rows: `solid` is `26b-01 lin30rob-ty-det01` (job 16530158, cycle 1, cutoff 0.50); `txty` is `16b-01 combo-txty` (job 16413506, cycle 80, cutoff 0.10); `txtytz` is `16b-01 combo-trans3` (job 16413520, cycle 80, cutoff 0.50); `tzrz` is `26-01 combo-tzrz-sym` (job 16514269, cycle 61, cutoff 0.20). Table 13 lists their print preset and supports. None was measured; print date, spool and lot are unrecorded. The `txty` part broke at the TPU–PETG interface and was taped for the photo; it is not a usable print.

A *ticket* is the group's internal number for one test and the settings it ran with, the Euler jobs behind those rows were not carried into this thesis. Of the campaign rows above, the service load, the small-motion screen, the linear design and the band repeat are stored in the repository's `research/` directory, as `58-<axis>`, `49b-<axis>`, `linear-control/` and `band-repeat/`; the rank-scaling and random-start records are there as well; no other row has artifacts there.

Table 8: The cutoff evaluations. The 28 evaluations behind the full cutoff sweep (Table 15), one row per evaluation: field, campaign ID, cutoff, density floor, design cycle, mesh n, completion flag, exit code and run path. The last row is the floor-zero cross-check: it repeats the legacy ramp's $\rho \geq 0.5$ field with the density floor at zero, where the `tx` motion stalls, so the row is incomplete and does not appear as a separate row in Table 15.

field	ID	cutoff	density floor	cycle	n	complete	exit	run
legacy ramp (c24)	`33-01`	0.4	0.001	24	30	false	3	`load-ramp-tz-cap030/run`
legacy ramp (c24)	`33-01`	0.5	0.001	24	30	true	0	`load-ramp-tz-cap030/run`
legacy ramp (c24)	`33-01`	0.6	0.001	24	30	false	3	`load-ramp-tz-cap030/run`
mixed-cell (rx)	`57-01`	0.4	0.001	80	30	true	0	`dof-rx/run`
mixed-cell (rx)	`57-01`	0.5	0.001	80	30	true	0	`dof-rx/run`
mixed-cell (rx)	`57-01`	0.6	0.001	80	30	true	0	`dof-rx/run`
mixed-cell (ry)	`57-01`	0.4	0.001	80	30	true	0	`dof-ry/run`
mixed-cell (ry)	`57-01`	0.5	0.001	80	30	true	0	`dof-ry/run`
mixed-cell (ry)	`57-01`	0.6	0.001	80	30	true	0	`dof-ry/run`
service-load (rx)	`58-01`	0.4	0.001	80	30	true	0	`dof-rx/run`
service-load (rx)	`58-01`	0.5	0.001	80	30	true	0	`dof-rx/run`

field	ID	cutoff	density floor	cycle	n	complete	exit	run
service-load (rx)	`58-01`	0.6	0.001	80	30	true	0	`dof-rx/run`
service-load (ry)	`58-01`	0.4	0.001	80	30	true	0	`dof-ry/run`
service-load (ry)	`58-01`	0.5	0.001	80	30	true	0	`dof-ry/run`
service-load (ry)	`58-01`	0.6	0.001	80	30	true	0	`dof-ry/run`
service-load (tz)	`58-01`	0.4	0.001	80	30	true	0	`dof-tz/run`
service-load (tz)	`58-01`	0.5	0.001	80	30	true	0	`dof-tz/run`
service-load (tz)	`58-01`	0.6	0.001	80	30	true	0	`dof-tz/run`
multi-axis (rx+ry+rz)	`59-01`	0.4	0.001	80	30	true	0	`combo-rx-ry-rz/run`
multi-axis (rx+ry+rz)	`59-01`	0.5	0.001	80	30	true	0	`combo-rx-ry-rz/run`
multi-axis (rx+ry+rz)	`59-01`	0.6	0.001	80	30	true	0	`combo-rx-ry-rz/run`
legacy climb (tz, 80 cycles)	`v1.1`	0.4	0.001	80	30	true	0	`climb-tz/run`
legacy climb (tz, 80 cycles)	`v1.1`	0.5	0.001	80	30	true	0	`climb-tz/run`
legacy climb (tz, 80 cycles)	`v1.1`	0.6	0.001	80	30	true	0	`climb-tz/run`
single-axis linear	`v1.1`	0.4	0.001	80	30	true	0	`full-tz-80-ml005/results`
single-axis linear	`v1.1`	0.5	0.001	80	30	true	0	`full-tz-80-ml005/results`
single-axis linear	`v1.1`	0.6	0.001	80	30	true	0	`full-tz-80-ml005/results`
Floor-zero cross-check								
legacy ramp (c24)	`33-01`	0.5	0	24	30	false	3	`load-ramp-tz-cap030/run`

The run paths are the last two components: the evaluations live under `/cluster/scratch/elihuber/attempts/escape-fleet/`, except `full-tz-80-ml005`, which is under `/cluster/home/elihuber/bachelor_thesis/worktrees/v11-05/examples/`.

A.3 Optional settings

The code exposes further settings that the runs of Section 3.1 leave at their default. Table 9 gives the value of each and what it changes.

Table 9: Optional settings. Default is the value the released code applies when a run file omits the key; an empty default is an empty list or map. The last column gives the value the setting took in the runs reported here, or names where a non-default value entered a run. `optimise.initial_design` and `motion.load_ramp` are alternative modes rather than switches, and their defaults are the modes these runs use.

Setting	Default	What it changes	Used in this thesis
`optimise.symmetry`	empty	Mirror-projects the design variable onto the named planes after every MMA step; `x` and `y`, not `z`, because the top face carries the prescribed motion and the bottom face is fixed.	no
`optimise.max_vf_drop_per_cycle`	0	Caps the per-cycle drop of the mean design density; a larger step is shrunk toward the previous design, and 0 is no cap.	guard arm only
`motion.translation_load_by_motion`	empty	Overrides the prescribed translation of one DOF in mm; rotations keep `motion.rotation_load_rad`.	no
`motion.translation_load_ramp`	empty	Per-DOF load continuation: the amplitude at the first cycle and the number of cycles to the full load, held after that; empty keeps the load constant.	load-ramp route only
`optimise.initial_design`	`solid`	Start field: `solid`, `uniform`, `hollow_cube` or `random`; the last takes its seed from `optimise.initial_design_params.seed`.	`solid` in every reported run; other starts unused
`motion.load_ramp`	`adaptive`	`adaptive` probes the load steps untaped and replays them; `fixed` uses a constant increment count (`solver.load_steps`).	`adaptive` in every run; `fixed` unused

A.4 Solver checks

Two checks on frozen fields stand behind Section 4.1. Table 10 gives the stored energy of the cross-code field in Abaqus and Firedrake for the six motions, the grid behind the gaps quoted there. Table 11 gives the last stroke each of the two probe fields reached before the solver stopped, against the inherited usable range.

Table 10: One frozen 40^3 field from the Abaqus run, solved in Abaqus and Firedrake at 2.5 mm and 0.12 rad. Columns are the six motions with their prescribed amplitude in the second row. W is the Firedrake stored energy; `ALLSE`, `ALLAE` and `ALLIE` are the Abaqus strain, artificial (hourglass) and internal energies, with `ALLIE` = `ALLSE` + `ALLAE` the total that matches the reaction. B–A gives Abaqus relative to Firedrake in percent of W; `ALLAE`/`ALLIE` and `ALLSD`/`ALLIE` are the artificial and static-stabilization shares of the internal energy. `RF` = d`ALLIE`/du and $F_{\tan} = \mathrm{d}W/\mathrm{d}u$ are the tangent forces at stroke u, in N for translations and N·mm for rotations. The `tz` column is at 98.11 % of its stroke, where the Abaqus solve stopped.

	`tx`	`ty`	`rx`	`ry`	`rz`	`tz`
Amplitude	2.5 mm	2.5 mm	0.12 rad	0.12 rad	0.12 rad	2.4528 mm (98.11 %)
W (mJ)	308.511	347.551	155.084	171.949	271.680	285.656
`ALLSE` (mJ)	280.266	320.109	103.100	122.284	240.377	162.816
`ALLAE` (mJ)	5.3	5.2	8.0	7.9	4.0	20.2
`ALLIE` (mJ)	285.611	325.355	111.068	130.137	244.334	183.046
B–A `ALLSE` (%)	−9.16	−7.90	−33.52	−28.88	−11.52	−43.00
B–A `ALLIE` (%)	−7.42	−6.39	−28.38	−24.32	−10.07	−35.92
`ALLAE`/`ALLIE` (%)	1.87	1.61	7.17	6.03	1.62	11.05
`ALLSD`/`ALLIE` (%)	0.055	0.057	0.046	0.046	0.055	0.062
`RF` (N, N·mm)	228.08	260.65	1,852.45	2,144.56	4,064.87	136.88
$F_{\tan}$ (N, N·mm)	246.73	278.32	2,701.20	2,966.07	4,707.53	222.17

Table 11: The prescribed motion reached by each probe field before the solver stopped, on two synthetic 40 mm fields (uniform $\rho = 0.4$; random half-solid, half-void) and against the inherited usable range (Trachsel (2026), p. 26). n is the cells per cube edge, Target the motion requested, and the bracketed percentage the share of that target reached. At $n = 20$ the same fields reach **20.000** and **15.582** mm on `tz`, 100 and 77.9 % of a 20 mm target. Both probe rows are sanity checks on the solver.

Field	n	Target	`tz` (mm)	Other translations (mm)	Rotations (rad)	Note
uniform $\rho = 0.4$	40	**20 mm**	**16.586** (82.9 %)	`tx`, `ty` 20.000	0.600	frozen-field probe; `tz` is the short axis
random half-solid, half-void	40	**12 mm**	**11.489** (95.7 %)	not probed	not probed	worst-case synthetic contrast; no optimized load paths
inherited usable range	35	**4 mm requested**	–	**2.0 to 3.0 usable**	**0.12**	Trachsel 2026 p. 26; 4 mm requested, 2.0 to 3.0 mm usable

A.5 Foaming ladder and flow compensation

Table 12 is the full record behind Figure 8. It gives every rung from 195 to 265 °C with its measured width, relative width, expansion ratio and flow setting, including the four above the interpolation window that the figure shows as context only.

Table 12: The measured foaming ladder and its flow compensation. T is the nozzle temperature, $\bar{w}$ the mean bead width at a 0.45 mm slicer line, $R = \bar{w}/w_0$ the relative width and $E_r = 1/R$ the extrusion ratio; $w_0 = 0.394$ mm is the 195–215 °C plateau mean. `M221` is the G-code flow setting, the same value as the extrusion multiplier in the PrusaSlicer filament page. The 250–265 °C rungs lie outside the 217–240 °C interpolation window and carry no `M221`.

T (°C)	$\bar{w}$ (mm)	$R = \bar{w}/w_0$	$E_r = 1/R$	`M221` (%)
195	0.3950	1.003	0.997	125
200	0.3900	0.990	1.010	125
205	0.3925	0.996	1.004	125
210	0.3900	0.990	1.010	125
215	0.4000	1.015	0.985	125
220	0.4750	1.206	0.830	60
225	0.5575	1.415	0.707	55
230	0.6500	1.650	0.606	55
235	0.7325	1.859	0.538	52
240	0.7650	1.942	0.515	52
250	0.7550	1.916	0.522	–
255	0.7275	1.846	0.542	–
260	0.6750	1.713	0.584	–
265	0.6600	1.675	0.597	–

A.6 Print record

Table 13 is the print record behind the printed parts: the job, the four bodies, and the print and filament settings that differ from stock PrusaSlicer.

Table 13: The print record: the print job and four parts, the print and filament settings that differ from stock PrusaSlicer 2.9.6, and the settings identical in both. The four parts share the VarioShore body setup; the solid cube uses PETG supports and the three flexures print without supports. The middle column gives the thesis value, the right stock, from the print preset `0.20mm Cube v2` against `0.20mm STRUCTURAL @XLIS 0.4` and the filament `ColorFabb VarioShore Calibration 235 0.52` against `ColorFabb VarioShore TPU @XL`. Both presets ship in `print/presets/prusa-xl-lab.ini`.

Setting	This thesis	PrusaSlicer stock
Print job: Prusa XL, 0.4 mm nozzle, `XL - Lab`		
Print profile		
Line width, all but top infill	0.40 mm	0.45 mm
Top infill width	0.40 mm	0.42 mm
Infill density	100 %	15 %
Infill pattern	rectilinear	grid
Supports, automatic placement	on	off
Support style	grid	snug
Support contact distance	0 mm	0.2 mm
Support interface layers	2	3
Support and interface toolhead	4	0 (current)
Support layers synchronized	yes	no
Interlocking beams	on	off
Segmented-region maximum width	5 mm	0
Bed temperature from toolhead	2	0 (first print)
Filament profile		
Nozzle temperature	235 °C	220 °C
First-layer temperature	235 °C	220 °C
Extrusion multiplier	0.52	0.85
Maximum volumetric speed	2.3 mm^3/s	1.5 mm^3/s
Minimum fan speed	40 %	60 %
Bridge fan speed	60 %	80 %
Fan off, first layers	1	4
Full fan speed from layer	3	6
Minimum purge on wipe tower	35 mm^3	15 mm^3
Start G-code	`M900 K0; M142 S36`	`M900 K0`
Equal in both profiles		
Layer height	0.20 mm	
Perimeters, solid layers	2; 5 top, 4 bottom	
Seam, skirt and brim	aligned; none	
Retract, wipe tower	1.2 mm at 60 mm/s; 60 mm, 3 mm brim	
Ooze prevention	on, −110 °C standby; idle 70 °C	
VarioShore bed temperature	50 °C	
Filament	FLEX, 1.75 mm, 1.22 g/cm^3	

`0.20mm Cube v2 non dis` is the solid part's save, with PETG supports; the rest use `0.20mm Cube v2`.

A.7 Export budget

Table 14 gives the first-order attribution of the cutoff's stiffness change used in Section 5.3. The columns separate the three terms and the floor's share of removal: hardening is the modulus change of the cells the cutoff keeps, removal the net change from dropping the sub-threshold cells to the printed void, and repair the change from the cells added at diagonal contacts; the floor lift is the part of removal that comes from the printed void being stiffer than the continuum void. On the five softer fields removal dominates on every axis and the body softens; on the two `tz` fields hardening dominates, and in the re-solved chain the floor lift supplies most of the positive change. Each term is estimated by changing one group of cells' modulus while the strain field stays fixed at the solution before the cutoff. Any such change also moves the load path, which the estimate misses, so the terms do not add up to the measured value; the re-solved chain in Section 5.3 re-solves the field after each step and closes the sum.

Table 14: Attribution of the cutoff's stiffness change at $\rho \geq 0.5$, per field, ranges over the six axes. The first-order terms attribute modulus change and do not sum to the measured change. Floor is the void-floor lift inside the sub-threshold term; the drop to the 0.004 MPa void is the other part, so it is not the net floor contribution to the measured change.

Field	Hardening (%)	Removal (%)	Repair (%)	Floor lift (pp)
mixed-cell (`rx`)	1.4 to 4.4	−44.2 to −3.1	0	13.6
mixed-cell (`ry`)	1.4 to 4.5	−41.9 to −3.4	0	13.7
service-load (`rx`)	1.4 to 4.6	−37.9 to −2.8	0	13.8
service-load (`ry`)	1.2 to 4.9	−18.8 to −4.0	0	10.4
service-load (`tz`)	3.8 to 9.7	−5.2 to 5.7	2.6	26.4
multi-axis (`rx`+`ry`+`rz`)	2.1 to 6.2	−58.1 to −3.2	0	15.0
legacy climb (`tz`)	3.4 to 7.1	−5.7 to 5.1	2.6	26.4

A.8 Cutoff sweep

Figure 7 shows the $\rho \geq 0.5$ slice; Table 15 gives every field at all three cutoffs, and Table 8 lists the evaluations behind the sweep.

Table 15: The full cutoff sweep, one row per field and cutoff. $\Delta K/K$ is the relative change in the energy secant $K = 2W/d^2$, with W the stored energy at stroke d, in per cent, between the exported body and the continuum re-solve of the same field, frozen at the end of its run; each column is a prescribed motion, and a negative value is softer than the continuum. The linear design and the legacy ramp are flagged and excluded from the seven-field comparison.

Field	$\rho \geq$	$\Delta K/K$ (%)					
		tx	ty	tz	rx	ry	rz
Service load (58-01)							
rx	0.4	−1.9	−29.5	−4.3	−5.8	−2.6	−6.5
rx	0.5	−4.0	−35.2	−7.2	−9.7	−4.8	−9.2
rx	0.6	−7.9	−42.7	−11.9	−15.3	−7.0	−13.3
ry	0.4	−1.6	−0.5	−1.4	−1.3	−2.2	+0.4
ry	0.5	−12.0	−3.3	−4.2	−4.6	−14.0	−5.4
ry	0.6	−19.1	−5.6	−6.0	−6.8	−20.4	−10.0
tz	0.4	+7.7	+7.4	+21.7	+11.5	+9.9	+5.9
tz	0.5	+5.5	+5.1	+17.9	+9.1	+7.8	+4.2
tz	0.6	+2.1	+2.1	+9.8	+3.7	+2.9	+1.8
Mixed-cell variant (57-01)							
rx	0.4	−2.3	−34.5	−3.8	−6.6	−3.4	−9.1
rx	0.5	−4.3	−41.5	−6.4	−12.4	−5.2	−12.1
rx	0.6	−6.1	−45.0	−8.2	−16.5	−6.3	−14.3
ry	0.4	−31.9	−2.0	−3.5	−3.1	−6.4	−6.9
ry	0.5	−39.3	−4.1	−6.7	−5.5	−12.3	−10.3
ry	0.6	−46.2	−6.1	−9.1	−7.0	−16.7	−13.6
Multi-axis support (59-01)							
rx+ry+rz	0.4	−46.1	−48.4	−20.6	−7.7	−5.6	−35.6
rx+ry+rz	0.5	−49.7	−54.6	−25.7	−13.3	−10.6	−43.3
rx+ry+rz	0.6	−55.1	−61.7	−31.1	–[a]	−13.9	−53.5
Legacy climb (v1.1)							
tz	0.4	+7.4	+7.3	+20.4	+10.7	+10.0	+5.7
tz	0.5	+4.5	+5.5	+16.3	+9.3	+6.9	+4.1
tz	0.6	+3.2	+3.9	+13.2	+7.1	+3.2	+2.3
Linear design (v1.1), not pooled							
tz	0.4	+0.7	−12.6	−10.0	−13.6	−16.0	−19.9
tz	0.5	−10.6	−30.9	−27.0	−24.4	−36.6	−30.4
tz	0.6	−18.9	−37.2	−33.5	−28.9	−43.6	−33.3
Legacy ramp (33-01), not pooled							
tz	0.4	−5.6	−7.6	+0.5	+8.2	–	+4.6
tz	0.5	−35.3	−32.3	−23.3	−10.2	−12.0	−8.2
tz	0.6	−59.9	−52.3	–	–	−36.7	−24.9

[a]No data: the continuum re-solve stopped short of its load, so no ratio exists. An unmarked dash is a value the incomplete legacy reports omit: the ry case at 0.4 and the tz and rx cases at 0.6 stopped short of full load. The re-solve lifts void to $\rho = 0.001$.

Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

Declaration of originality

The signed declaration of originality is a component of every written paper or thesis authored during the course of studies. **In consultation with the supervisor**, one of the following two options must be selected:

☐ I hereby declare that I authored the work in question independently, i.e. that no one helped me to author it. Suggestions from the supervisor regarding language and content are excepted. I used no generative artificial intelligence technologies[1].

☒ I hereby declare that I authored the work in question independently. In doing so I only used the authorised aids, which included suggestions from the supervisor regarding language and content and generative artificial intelligence technologies. The use of the latter and the respective source declarations proceeded in consultation with the supervisor.

Title of paper or thesis:

Large-Deflection Topology Optimization: An Open-Source Solver and Multi-Material 3D Printing

Authored by:
If the work was compiled in a group, the names of all authors are required.

Last name(s):	**First name(s):**
Huber	Elias Gabriel

With my signature I confirm the following:
- I have adhered to the rules set out in the Citation Guidelines.
- I have documented all methods, data and processes truthfully and fully.
- I have mentioned all persons who were significant facilitators of the work.

I am aware that the work may be screened electronically for originality.

Place, date	**Signature(s)**
ZH, 02.11.2026	

If the work was compiled in a group, the names of all authors are required. Through their signatures they vouch jointly for the entire content of the written work.

[1] For further information please consult the ETH Zurich websites, e.g. https://ethz.ch/en/the-eth-zurich/education/ai-in-education.html and https://library.ethz.ch/en/researching-and-publishing/scientific-writing-at-eth-zurich.html (subject to change).