Laminated Vault: Scaffold-Free 3D Concrete Printing of Vaulted Structures via Zig-Zag Lamination

Abdallah Kamhawi[a,1,*], Mania Aghaei Meibodi[a,1]

[a]Taubman College of Architecture and Urban Planning, University of Michigan, Ann Arbor, MI 48109, USA

Abstract

Vaulted structures achieve exceptional structural efficiency by carrying loads primarily through compression, yet their construction has declined because of the specialized labor and extensive internal scaffolding required. Robotic 3D concrete printing offers a path to automated vault construction without bespoke internal scaffolding, but the fresh cementitious material cannot support itself during the cantilevering stages that arise as a vault is built from one side to the other; unsupported filaments sag and peel away from the preceding layer within seconds of deposition. This paper addresses the problem by introducing a zig-zag lamination strategy in which a secondary filament is deposited in a diagonal pattern over the primary layer, creating tension ties that anchor the growing cantilever back into stabilized material behind it. A reduced-order analytical framework is developed to design the lamination geometry, coupling vault geometry, deposition sequence, and early-age material state into closed-form admissibility checks for any candidate design. The framework was validated on a barrel-vault demonstrator: the selected lamination design sustained the complete vault with no delamination or visible distress. The scaffold-free claim refers to eliminating the internal scaffolding traditionally required to support fresh material during cantilevering construction, not eliminating the springing support inherent to any vaulted geometry.

Keywords:
3D concrete printing, scaffold-free construction, zig-zag lamination, buildability, vaulted structures

1. Introduction

1.1. The value, decline, and renaissance of vaults

By distributing loads purely through compression, vaults achieve high structural efficiency and can form long-span slabs that are difficult to realize with conventional bending systems [1–5]. Because concrete has high compressive strength but very limited tensile capacity, vaults exploit this asymmetry to create long-span structures with minimal material while reducing or eliminating the need for tensile reinforcement [6–10].

Vault construction has historically required two distinct forms of support, which are conflated in common usage but differ in structural role. The first is the *springing support*, a boundary condition at the base of the vault that transfers the compressive thrust from the shell into the foundation. Every vault, from Roman arches to Gothic ribbed vaults to Guastavino's tile vaulting and modern reinforced concrete shells, requires a springing support as a structural necessity inherent to the definition of a vault. The second is *internal scaffolding*, comprising centering frames, formwork boards, and falsework erected temporarily within the envelope of the vault to hold the fresh material in place during construction before the shell closes on itself (fig. 1). Unlike the springing support, internal scaffolding is not structurally required by the finished vault; it exists solely to carry the weight of the construction-stage material until the vault achieves its self-supporting geometry. It is the labor-intensive, waste-generating component of traditional vault construction, and it is what Guastavino's cohesive tile vaulting eliminated by cantilevering each course from the previously bonded one using fast-setting mortar. In this paper, the term *scaffold-free* refers to the elimination of internal scaffolding during the cantilevering construction stages of a printed vault. The springing support is retained as a structural boundary condition, as it must be in any vaulted structure.

*Corresponding author.
Email address: kamhawi@umich.edu (Abdallah Kamhawi)
[1]These authors contributed equally to this work.

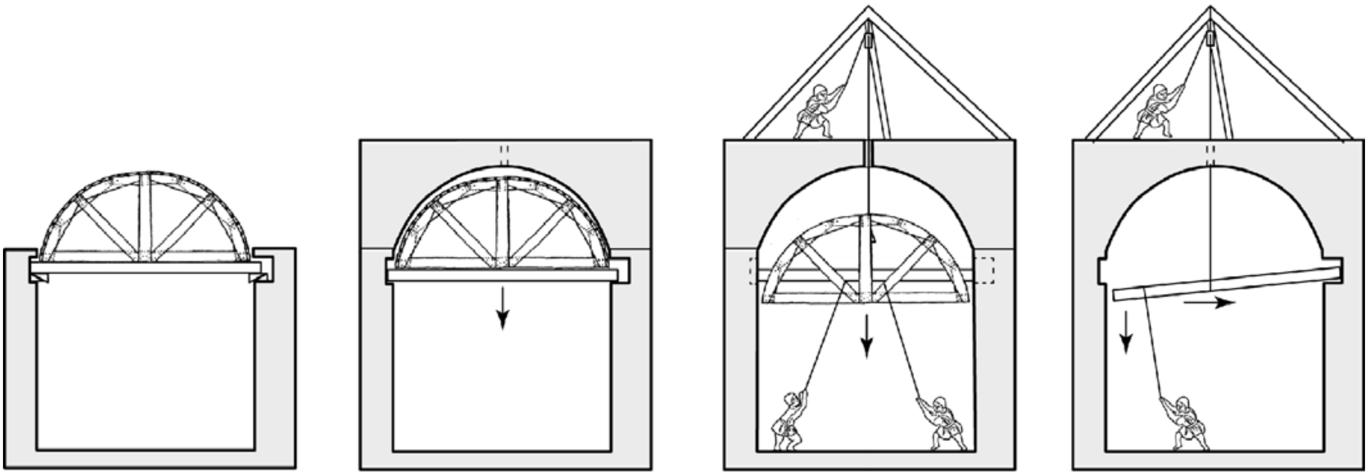

Figure 1: Traditional internal scaffolding for vault construction, illustrated for a Roman barrel vault. Centering frames, formwork boards, and timber beams are erected within the envelope of the vault to carry the weight of the fresh material during construction; the scaffolding is dismantled after the vault closes on itself. This temporary internal support is the component of vault construction that scaffold-free methods are designed to eliminate. It is distinct from the springing support at the base of the vault, which is a structural boundary condition required by any vaulted geometry [11].

Despite their structural advantages, vaults have been significantly underutilized in modern construction, owing primarily to two factors [5, 12–16]: (a) the scarcity of specialized labor with knowledge of vault construction, and (b) the need for extensive bespoke internal scaffolding and formwork that generate substantial waste and cost.

Recently, advances in computational design and robotic fabrication have renewed interest in vaulted construction [14, 17–19]. Robotic 3D concrete printing (3DCP), in particular, offers the potential to realize complex geometries through continuous, automated deposition while reducing reliance on bespoke formwork [20–25]. However, the geometric freedom of 3DCP is severely constrained by the material's rheological behavior during deposition [26–29]. Unlike thermoplastic fused deposition modeling [30] or metal directed energy deposition [31], which rely on rapid thermal phase changes to solidify material upon deposition, 3DCP depends on hydration chemistry and thixotropic structuration [26, 32–34]. The extruded cementitious paste remains in a highly deformable green state for seconds to minutes after exiting the nozzle. When fresh filament is deposited without underlying support, it acts as a soft cantilever beam subjected to its own weight, lacking the flexural and shear rigidity to resist gravitational body forces, leading to viscoplastic sagging and collapse [26, 35–38].

Current 3DCP research largely addresses this limitation through prefabricated elements that are printed in stable orientations and later assembled on-site using temporary internal scaffolding (fig. 2) [39–46]—or through geometries that remain inherently self-stabilizing during deposition [47–53]. While these approaches have significantly advanced construction with 3DCP, they are less suited to vaulted systems with complex geometries, particularly those that pass through substantial cantilevering states during construction. In such cases, the buildability problem, maintaining the stability of the fresh material throughout the print sequence, remains a primary bottleneck in scaffold-free robotic 3DCP.

Figure 2: Existing approaches to robotic 3D concrete printing of vaulted and shell structures. (a,b) Discretized, prefabricated elements assembled on temporary internal scaffolding [44, 54]. (c) Directly printed self-supporting vaults relying on inherently stable geometries [52].

A complementary body of work has developed predictive models for buildability in 3DCP, aimed at characterizing the material behavior of fresh cementitious paste during deposition. One direction is rheology-based, treating printability and structural build-up as consequences of yield stress evolution, thixotropic structuration, and hydration kinetics [26–28, 32–34, 38]. A second direction is continuum-mechanics-based, predicting the onset of construction-stage deformation and loss of stability in printed walls and shells through viscoelastic-plastic material models with progressive element activation [35–37]. These formulations answer the question of whether a given print will succeed under given material and process conditions, and they do so by modeling the printed object as a continuum whose failure is governed by bulk deformation and stability. The failure mode that governs scaffold-free printing of cantilevering vaults, however, is not bulk deformation but delamination at a specific interlayer interface, where the unsupported cantilever peels away from the stabilized material behind it. Physics-resolved prediction of this interfacial failure mode during printing is a complementary activity whose formulation differs in kind from the continuum predictions above, and it lies outside the scope of the present work.

The present paper addresses a different question. Rather than predicting how a given design will deform or fail during printing, it asks which lamination designs, specified by two integer geometric parameters, render a chosen vault geometry admissible under the governing interfacial failure mode at every stage of construction. The contribution is therefore a design tool that returns a binary admissibility decision for candidate lamination parameters, not a simulation of construction-stage mechanics. No existing method provides such a design framework for scaffold-free robotic 3DCP of vault geometries that must pass through substantial cantilevering states during the construction sequence, precisely the geometries for which the buildability bottleneck is most severe.

1.2. Contributions and objectives

This paper introduces zig-zag lamination as a scaffold-free construction strategy for robotic 3DCP of vaulted structures that pass through cantilevering construction states—the class of vault geometries for which no existing method provides a path to fabrication without internal scaffolding, centering, or temporary formwork beneath the cantilevering region. The strategy is inspired by Guastavino's cohesive tile vaulting [55] (fig. 3): a secondary filament is deposited over the primary compression-aligned layer in a diagonal alternating path, bonding the unsupported cantilever tip back to stable material and forming discrete tension ties that prevent peel failure at the root interface.

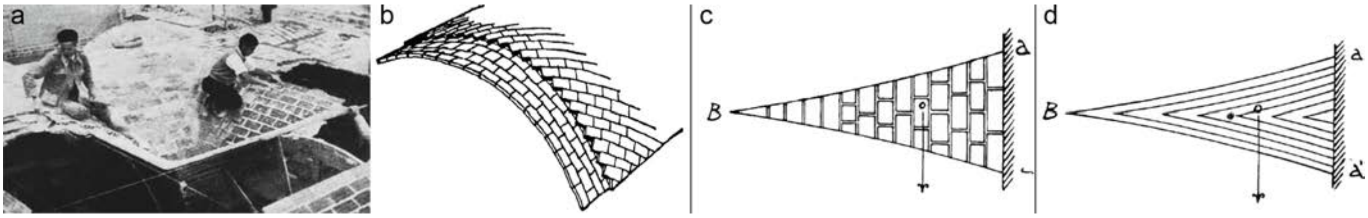

Figure 3: Cohesive tile vaulting without centering. (a) Workers lay tiles at the growing edge of a vault; each tile is cantilevered from the previously bonded course using fast-setting mortar [56]. (b) Axonometric view showing multiple laminated layers on the curved surface [57]. (c, d) Plan views of the first and second tile courses: the diagonal orientation of the second layer cross-bonds the joints of the first, providing tensile connectivity across the growing edge. The zig-zag lamination strategy developed in this paper translates this principle into robotic 3DCP through a continuous deposited filament [55].

The term *lamination* is retained because the strategy adds a secondary deposited filament over the primary layer and is conceptually related to cohesive laminated vaulting. In the construction-stage model developed here, however, the stabilizing action of this lamination is represented through discrete tie-back branches rather than through a smeared continuous layer. The diagonal branches connect the current unsupported filament to material within the stable region and mobilize local tensile action after sufficient wet-on-wet bond develops. This modeling choice does not imply that the laminated cross-section has no contribution to the stiffness or load-bearing capacity of the hardened vault; that service-stage structural role is outside the scope of the present study.

To make this strategy designable without simulation, the paper develops a reduced-order analytical framework in which any candidate lamination layout is fully specified by two integers—the restraint-node count n_r and the tie-back depth d_r—and screened against closed-form admissibility conditions that determine whether the design can survive deposition and sustain the growing vault. The framework answers two questions that must be resolved before printing: *will the cantilever survive the instant the zig-zag is deposited but not yet bonded?* and *will the vault remain stable as lamination cycles accumulate?* Both reduce to explicit inequalities that can be evaluated analytically.

The framework and the lamination strategy are validated together on a barrel-vault demonstrator. A primary-only print fails at three unsupported filaments; the selected zig-zag design sustains the complete vault to the crown with no delamination or visible distress.

2. Methods

Consider a vaulted shell structure fabricated by robotic 3D concrete printing in which internal scaffolding is eliminated from the cantilevering region and the vault must remain stable throughout the build sequence after the springing/startup boundary is established. The vault is constructed from a sequence of *primary layers*—concrete filaments deposited side by side along the shell surface, each following the local thrust line of the target geometry (fig. 6a). Each new filament bonds to the previously deposited one at a wet-on-wet interlayer interface, and adjacent filaments are separated by a center-to-center spacing p_p. Near the supports, the primary filaments rest on the springing/startup support or on previously stabilized material and are stable under their own weight. As the print advances along the curvature of the vault, however, each successive filament is carried only by its interlayer bond to the one behind it: the structure cantilevers outward from the last stable boundary, and the self-weight moment at the boundary between the stable and unsupported material—the *root interface* I_0— grows with the number of unsupported filaments. Once this moment exceeds the capacity of the still-maturing interface, the cantilever peels away at I_0 and the print fails.

A *zig-zag lamination* is introduced to prevent this failure. A secondary filament is deposited on top of the primary layer in a zig-zag path that alternates diagonally between the most recently printed unsupported filament and an earlier filament lying within the stable region (fig. 6b). In the construction-stage model developed here, the stabilizing action of the lamination is represented by discrete tie-back branches rather than by a smeared continuous layer. A zig-zag design is defined by two integers: the restraint-node count n_r and the tie-back depth d_r (fig. 6b). The printing sequence proceeds in alternating steps (fig. 4). First, one or more primary filaments are deposited along the vault thrust line, extending the unsupported cantilever by N_c increments beyond the last stable boundary (fig. 4a). A zig-zag pass is then deposited over the primary layer, producing a path of $2n_r$ diagonal branches (fig. 4b). At the instant of deposition, the zig-zag is treated as unbonded dead load on the cantilever. The model does not assume that the secondary path develops full tensile capacity immediately upon contact. Instead, the construction sequence is represented by two limiting states: a pre-bond state, in which the zig-zag contributes only dead load, and a post-bond state, in which sufficient wet-on-wet bond has developed for the diagonal branches to mobilize tension. As the interlayer bond develops, the diagonal branches engage, transferring a portion of the cantilever weight back into the stable material and reducing the net downward force that the root interface must carry (fig. 4c).

Once the transferred load sufficiently relieves the root interface, that interface is deemed stable, the root advances outward by one cycle depth $N_c\, p_p$, and the next cycle begins. The same zig-zag geometry can therefore be stabilizing or destabilizing depending on how much weight it adds before bonding versus how much load its branches can redirect after bonding. The central design question is: *which combinations of n_r and d_r keep the cantilever admissible at every instant during construction?*

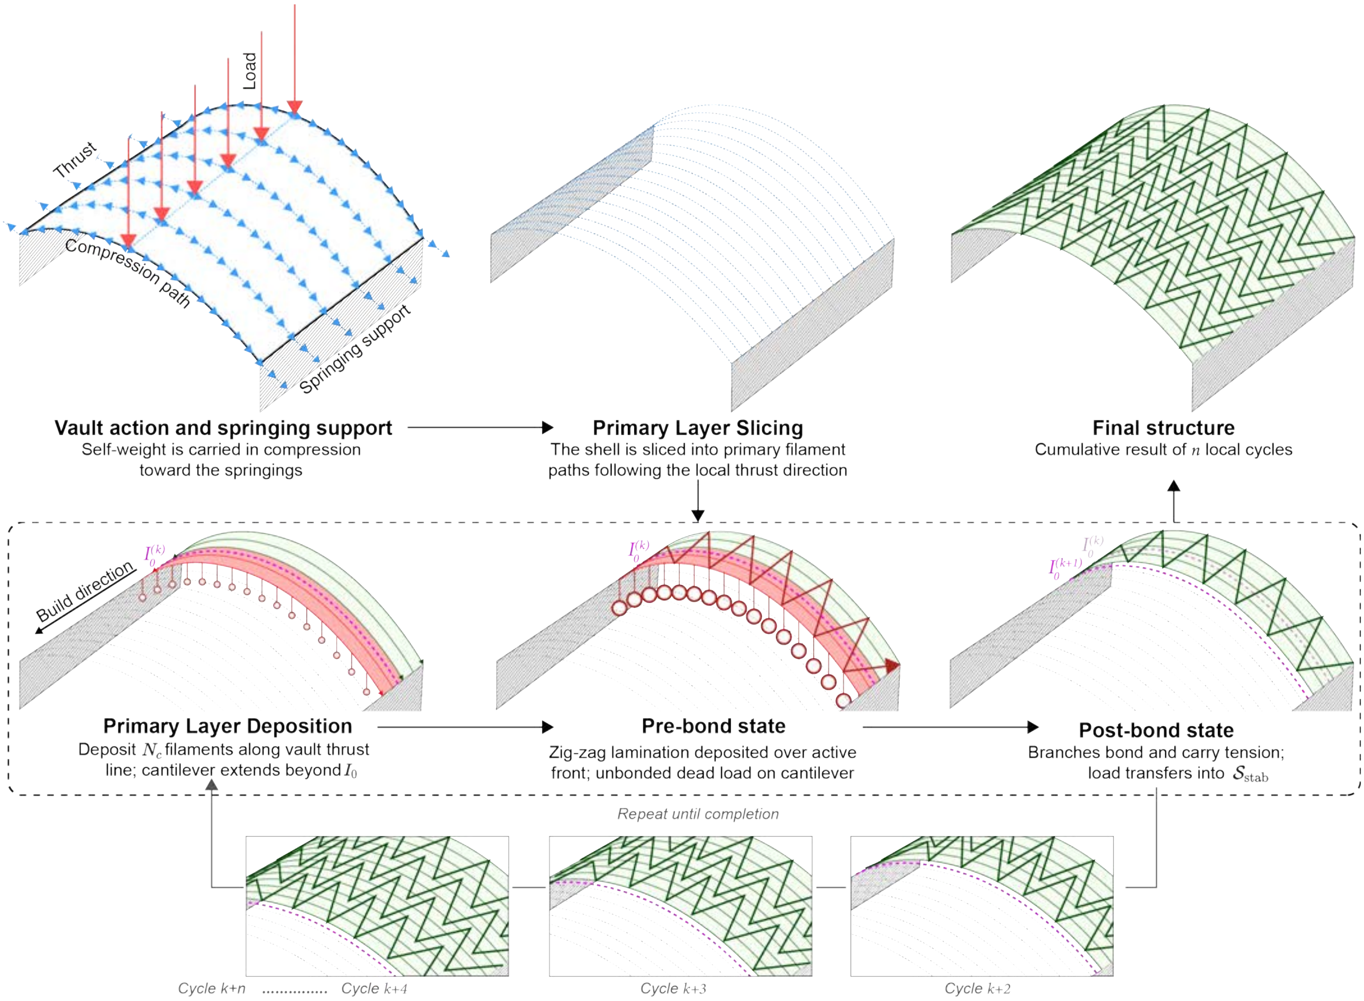

Figure 4: The zig-zag lamination acts locally and cycle by cycle. Each cycle has three states: (1) N_c primary filaments cantilever beyond the current root interface $I_0^{(k)}$; (2) a zig-zag is deposited over the active front as unbonded dead load (pre-bond); (3) the branches bond, transfer load into the stable region $\mathcal{S}_{\text{stab}}$, and the root interface advances by $N_c\, p_p$ to $I_0^{(k+1)}$. Repeating this cycle propagates the active front along the vault; the final pattern (top right) is the accumulation of these local cycles. Bright red and green mark the cycle currently being deposited or just bonded; faded green marks previously stabilized cycles now part of $\mathcal{S}_{\text{stab}}$.

This question must be answered in two distinct states: the *pre-bond state*, when the freshly deposited zig-zag is unbonded dead load on the cantilever, and the *post-bond state*, when the diagonal branches carry tension and transfer load back into the stable region. A design that survives the pre-bond state can still fail progressively if the post-bond transfer does not fully compensate the primary weight at every cycle. The framework developed below (fig. 5) resolves both questions. **Part 1** (section 2.1) develops the construction-stage geometry, establishes the primary-only demand baseline, and parameterizes the zig-zag branch mechanics. **Part 2** (section 2.2) assembles the two binding admissibility checks—deposition survivability (Check 1) and chain sustainability (Check 2)—and constructs the admissible domain $\mathcal{D}_\ell(N_c)$, from which the final design parameters can be selected.

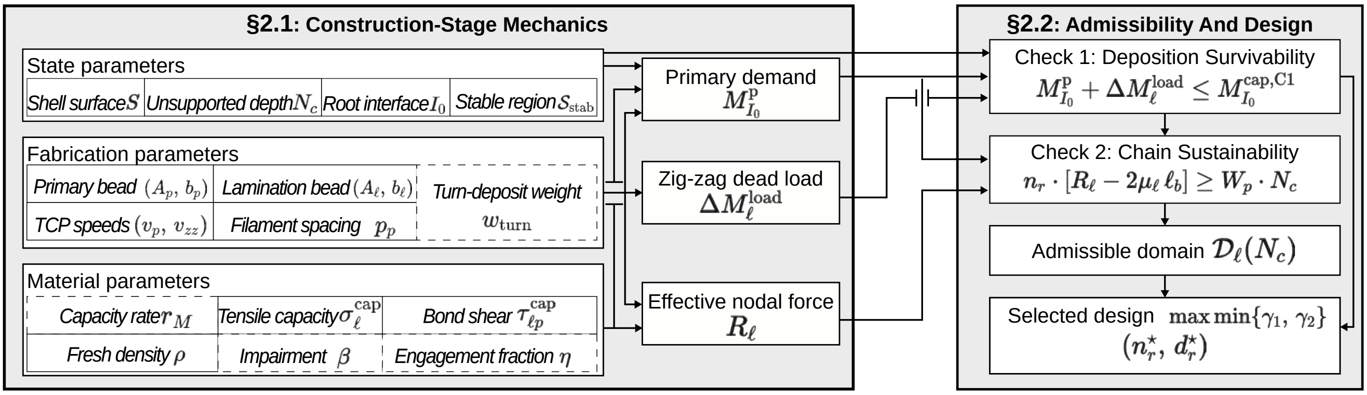

Figure 5: Overview of the lamination design framework. Construction-stage geometry, material state, and zig-zag branch mechanics define the primary demand $M_{I_0}^{p}$, the zig-zag dead load $\Delta M_\ell^{\text{load}}$, and the effective nodal force R_ℓ. Two binding checks—deposition survivability (Check 1, pre-bond moment) and chain sustainability (Check 2, post-bond force balance)—define the admissible domain $\mathcal{D}_\ell(N_c)$, from which the final design $(n_r^\star, d_r^\star)$ is selected. Dashed arrows indicate empirical inputs supplied from external characterization.

2.1. *Construction-stage geometry and governing quantities*

Given an unsupported cantilever of N_c filaments, admissibility depends on three quantities: the primary self-weight demand at I_0, the additional dead load imposed by the zig-zag, and the restoring force mobilized once its branches bond. These quantities are developed below from the construction-stage geometry on the developed shell surface. A notation key (fig. 6) and a compact symbol table (table 1) summarize the core symbols used throughout the section; secondary quantities are defined locally where they are introduced.

2.1.1. *Construction-stage geometry*

The goal of this section is to reduce the construction-stage problem to a scalar check at a single governing interface.

Consider a shell surface on which primary filaments are deposited sequentially, each of length L_p and cross-sectional area A_p, bonding to the previous filament at a wet-on-wet interlayer interface at a center-to-center spacing p_p. All geometric quantities are evaluated on the developed (unrolled) shell surface so that the layout is planar. The first several filaments rest on the springing/startup support or on previously stabilized material. As deposition advances beyond this supported zone, each new filament is carried only by its interlayer bond to the one behind it, and the structure begins to cantilever outward (fig. 6).

The interlayer bond at the boundary between the supported and unsupported material is the *root interface* I_0. The material behind I_0—the *stable region* $\mathcal{S}_{\text{stab}}$—has been hydrating longer than the cantilever and is treated as a rigid support over the timescale of a single construction stage. Ahead of I_0, the unsupported cantilever comprises N_c filaments, labeled $P_1, \ldots, P_{N_c}$ from I_0 outward. Each filament occupies a strip of width p_p whose centroid lies at transverse distance $(j - \frac{1}{2})p_p$ from I_0. Because self-weight bending demand increases monotonically toward the supported end, I_0 carries the largest peel moment and is therefore the governing location for admissibility.

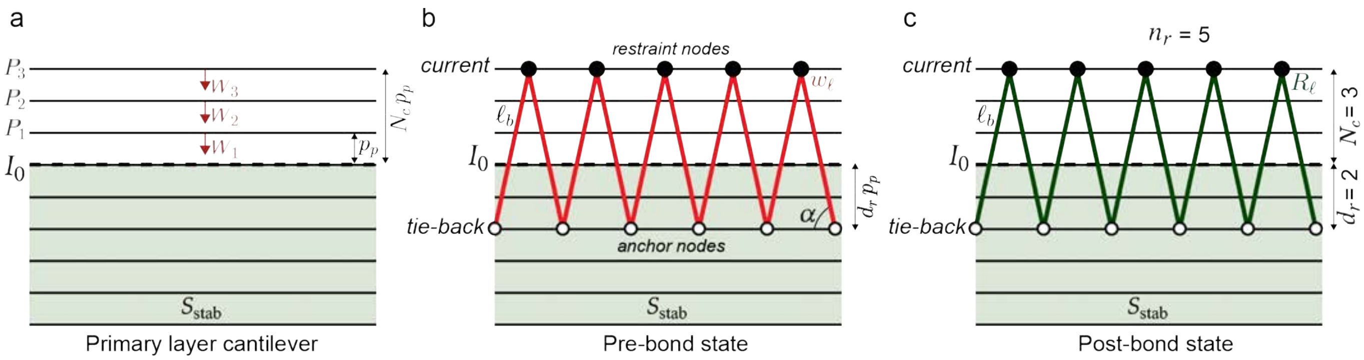

Figure 6: Geometric reference for the primary-only state and the local zig-zag lamination mechanism on the developed vault surface. **(a)** Primary-only cantilever state. The stable region S_{stab} lies below the root interface I_0, while the newly deposited primary filaments P_1, P_2, and P_3 extend beyond I_0. Their self-weights W_1, W_2, and W_3 generate the primary-only cantilever demand at the root interface. **(b)** Pre-bond zig-zag deposition. A local zig-zag lamination path is deposited over the active front and returns to anchor nodes on a tie-back filament within S_{stab}. Before sufficient wet-on-wet bond develops, this path is treated as added dead load w_ℓ on the cantilevering region. **(c)** Post-bond tie-back action. After bond engagement, the same zig-zag branches mobilize an effective nodal restoring action R_ℓ, transferring load from the current unsupported filaments back into S_{stab}. The parameters n_r, N_c, and d_r define, respectively, the restraint-node count, the number of unsupported primary increments in one cycle, and the tie-back depth measured in primary-filament spacings.

Modeling assumptions. Four idealizations reduce the construction-stage problem to a scalar check at I_0. Each addresses a distinct physical aspect of the green-state cantilever: the material-state gradient in $\mathcal{S}_{\text{stab}}$ (Assumption 1), the geometry of the cantilever and the resulting mode of interfacial failure (Assumption 2), the deflection regime over which the model is evaluated (Assumption 3), and the local curvature of the shell surface relative to the filament spacing (Assumption 4). The framework applies where these four regimes hold simultaneously; the breakdown of each is identified below together with the modeling extension required.

1. **Rigid stable region.** The material behind I_0 is treated as a rigid support, and the unsupported region is idealized as a cantilever rooted at I_0. In green-state cementitious materials, thixotropic restructuration and early hydration produce a structural build-up that grows monotonically with age over the construction window [26, 33, 36], so each interface in $\mathcal{S}_{\text{stab}}$ is older and stronger than the one ahead of it. Under uniform-spacing conditions, Check 2 ensures by construction that no residual force propagates inward to challenge interfaces other than I_0, which is therefore the governing location throughout the build. This idealization breaks down when $\mathcal{S}_{\text{stab}}$ itself contains an unresolved cantilevering substructure from a preceding cycle, so that a second interface behind I_0 may govern instead. Under those conditions the scalar check at I_0 is no longer sufficient and admissibility must be evaluated at every candidate interface along the chain.

2. **Peel-dominated loading.** The root interface is loaded primarily in opening rather than in-plane shear. Whether an interface fails in peel, shear, or a coupled mode is set by the geometry of the cantilever rather than by the orientation of the shell surface in space. The printed cantilever in this work is slender, with cantilever length large compared to the bead cross-section, so the bending stress at I_0 produced by the lever arm of the self-weight greatly exceeds the direct shear stress, and the dominant failure mode at the interface is mode-I peel opening [35, 36]. The governing stress state at I_0 is therefore represented by a single scalar bending moment M_{I_0} compared against a time-dependent capacity $M_{I_0}^{\text{cap}}(t)$. The idealization breaks down for geometries where the printed material no longer acts as a slender cantilever, such as stubby load paths, compression columns near the springing of a tall arch, or re-entrant configurations where coupled failure modes become competitive. In those regimes a coupled peel–shear criterion replaces the scalar moment check.

3. **Small deflections.** Self-weight lever arms are evaluated in the undeformed configuration. The framework predicts whether a candidate design avoids peel failure at I_0, not the dynamics of the failure event itself. Up to that failure point, the cantilever remains in the small-deflection regime, so the undeformed-geometry evaluation is sufficient throughout the regime of interest. The idealization breaks down for materials or geometries in which measurable viscoplastic sag accumulates before bond development, such as slow-setting mixes or larger bead spacings [36]. In that regime the true lever arms are larger than their undeformed values, the demand at I_0 exceeds the linear estimate, and Check 1 should be interpreted as an upper bound on the safe N_c.

4. **Developed-surface geometry.** Filament positions, lever arms, and branch lengths are evaluated on the developed shell surface. For surfaces with vanishing Gaussian curvature, the developed (unrolled) representation preserves both distances and angles exactly, so a planar evaluation of geometric quantities is rigorous; for surfaces with finite but small Gaussian curvature relative to the filament spacing, the planar evaluation remains a controlled approximation. The relevant dimensionless criterion is $K\,p_p^2 \ll 1$, where K is the local Gaussian curvature. The barrel-vault demonstrator is developable ($K = 0$ everywhere), so the planar evaluation is exact in this case. The idealization breaks down for strongly doubly curved shells, where developed-surface lever arms and branch lengths diverge from their true geodesic values. For such shells, geodesic computation on the actual surface geometry is required in place of the planar evaluation used here.

These four idealizations characterize the underlying physical model. The closed-form admissibility checks developed in section 2.2 are additionally specialized to the stronger condition of *uniform filament length and spacing*, under which a single zig-zag design (n_r, d_r) applies throughout the build and the chain analysis collapses to a single closed-form inequality (eq. (8)). For geometries that do not satisfy this condition, the underlying physics of Checks 1 and 2 is unchanged and the branch mechanics still apply at every cycle; only the chain analysis must be evaluated sequentially on a cycle-by-cycle basis, with (n_r, d_r) permitted to vary per cycle, in place of the closed-form simplification.

Table 1: Core notation used throughout the Methods. Secondary quantities are defined locally where introduced.

Symbol	**Description**	**Units**
I_0	Root interface (governing interlayer bond)	–
$\mathcal{S}_{\text{stab}}$	Stable region behind I_0	–
N_c	Number of unsupported filaments per cycle	–
$L_p,\ p_p$	Primary filament length; center-to-center spacing	m
$A_p,\ A_\ell$	Primary / lamination cross-sectional area	m^2
b_ℓ	Lamination bond width	m
ρ	Fresh concrete density	kg/m^3
W_p	Primary filament weight, $\rho A_p L_p g$	N
μ_ℓ	Lamination weight per unit length, $\rho A_\ell g$	N/m
$n_r,\ d_r$	Restraint-node count; tie-back depth	–
ℓ_b	Branch length on the developed surface	m
T_b	Branch-force capacity	N
R_ℓ	Effective nodal force at one restraint node	N
$M_{I_0}^{p}$	Primary-only root moment	N m
$\Delta M_\ell^{\text{load}}$	Total zig-zag dead-load moment	N m
r_M	Interface capacity rate	N m/s
β	Capacity-impairment factor in Check 1	–
η	Engagement fraction in post-bond transfer	–
$\gamma_1,\ \gamma_2$	Admissibility margins for Checks 1 and 2	–

2.1.2. Primary-only demand baseline

Before introducing the zig-zag, the baseline must be established: how far can the primary layer cantilever on its own before the root interface fails?

Under equal filament length and uniform spacing, every unsupported filament has the same weight $W_p = \rho A_p L_p g$, and the root moment is

$$M^{p}_{I_0}(N_c) = \tfrac{1}{2}\, W_p\, p_p\, N_c^2. \tag{1}$$

The quadratic dependence on N_c reflects the fact that each newly added filament increases both the total unsupported weight and the average lever arm of that weight about I_0.

The cementitious material at I_0 is simultaneously hydrating and gaining strength, so buildability is a race between increasing demand and increasing interface capacity. A linear capacity law is adopted:

$$M^{\text{cap}}_{I_0}(t) = r_M\, t, \tag{2}$$

where r_M is the *interface capacity rate*, in units of N m/s. This quantity is distinct from a material-level structuration rate: it is calibrated directly from the failure of a primary-only printing experiment (section 3.2) and therefore represents a structural-level rate of interfacial moment growth under the tested process conditions. The primary layer alone is admissible whenever $M^{p}_{I_0}(N_c) \le M^{\text{cap}}_{I_0}(t)$. When this condition is not satisfied, the zig-zag lamination must supply additional restraint.

2.1.3. Zig-zag geometry, branch action, and dead load

A zig-zag design is specified by two integers (fig. 6b): the restraint-node count n_r, which sets how many attachment points are distributed along the filament length, and the tie-back depth d_r, which sets how far back into the stable region each branch reaches. These two parameters jointly control the branch geometry and create the central design trade-off: increasing restraint density or tie-back depth can increase the available restoring action, but it also increases the material placed on the cantilever and therefore the dead-load penalty (fig. 7).

The zig-zag connects the current filament to a tie-back filament lying $d_r\, p_p$ behind the root within $\mathcal{S}_{\text{stab}}$. On the tie-back filament, $n_r + 1$ anchor points are evenly spaced along the length; on the current filament, n_r restraint nodes are placed at the midpoints between consecutive anchors. The path alternates between anchor and restraint node, producing $2n_r$ diagonal branches. Each branch spans a transverse distance $d_r\, p_p$ and a longitudinal distance $L_p/(2n_r)$, so its developed-surface length is

$$\ell_b(n_r, d_r) = \sqrt{(d_r\, p_p)^2 + \left(\frac{L_p}{2n_r}\right)^2}. \tag{3}$$

Its transverse force fraction is set by the corresponding branch angle α, for which $\sin\alpha = d_r p_p/\ell_b$.

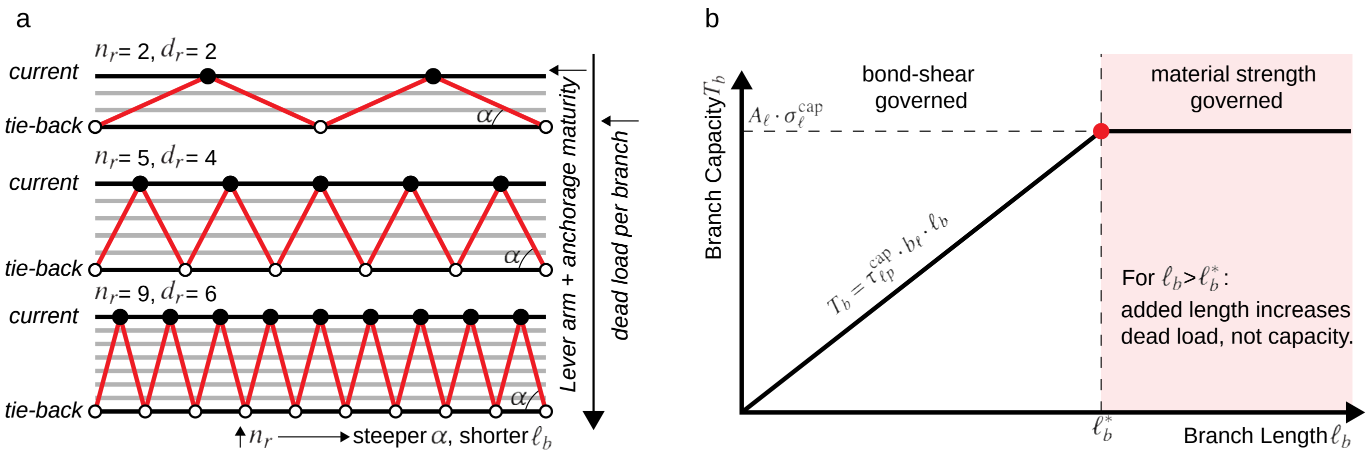

Figure 7: Zig-zag design trade-offs. **(a)** Effect of the restraint-node count n_r and tie-back depth d_r on the zig-zag layout, illustrated for three designs: $(2, 2)$, $(4, 4)$, and $(8, 6)$. Filled circles mark restraint nodes on the current filament; open circles mark anchor points on the tie-back filament. Increasing n_r shortens each branch and steepens its inclination α, improving the transverse force component but reducing bonded contact length per branch and increasing dead load per unit length of the layout. **(b)** Branch-force capacity regimes: bond-shear-governed and material-strength-governed, separated by the crossover branch length $\ell_b^\star$. Beyond $\ell_b^\star$, additional branch length contributes only dead load without gain in capacity.

Once bonded, each branch is loaded in tension and can fail either by rupture of the lamination filament or by shear transfer failure along the zig-zag/primary interface. The governing branch-force capacity is therefore

$$T_b(n_r, d_r, t) = \min\{A_\ell\, \sigma^{\text{cap}}_\ell(t),\ \tau^{\text{cap}}_{\ell p}(t)\, b_\ell\, \ell_b(n_r, d_r)\}. \tag{4}$$

The transition between the two regimes occurs at the crossover branch length $\ell_b^\star(t) = A_\ell \sigma^{\text{cap}}_\ell(t)/(\tau^{\text{cap}}_{\ell p}(t) b_\ell)$. Beyond $\ell_b^\star$, making a branch longer adds dead load without increasing the force it can carry.

The branch-force capacity is not the same as the restoring force delivered to the cantilever. Two reductions intervene. First, only the transverse component of branch tension opposes the peel moment at I_0. Second, incomplete contact, the hydration-age gradient in the anchoring zone, and substrate compliance prevent the full nominal branch force from being mobilized. These effects are captured by an empirical *engagement fraction* $\eta(t) \in (0, 1]$. The resulting *effective nodal force* at one restraint node is

$$R_\ell(n_r, d_r, t) = \eta(t)\ 2\,T_b(n_r, d_r, t)\ \sin\alpha(n_r, d_r). \tag{5}$$

Before this restoring mechanism becomes active, the freshly deposited zig-zag acts as dead load on the cantilever. Two contributions are included: the unsupported portion of the branch material itself and the excess mass deposited at each direction reversal due to extrusion continuing while the nozzle decelerates, reverses, and reaccelerates. With $\mu_\ell = \rho A_\ell g$ denoting lamination weight per unit deposited length and w_{turn} the measured excess weight per reversal, the total zig-zag dead-load moment is

$$\Delta M_\ell^{\text{load}}(n_r, d_r) = \underbrace{\mu_\ell\, n_r\, \ell_b(n_r, d_r)\, \frac{N_c^2 p_p}{d_r}}_{\text{branch dead load}} + \underbrace{2\, n_r\, w_{\text{turn}}\, N_c\, p_p}_{\text{turn-deposit load}}. \tag{6}$$

The first term decreases in effective leverage as d_r increases because a greater fraction of each branch lies on the stable side of I_0; the second grows linearly with n_r and is independent of d_r.

Across the construction stage, load transfer between the zig-zag and the stable region is mediated by the interlayer bond at the anchor nodes. Along the bonded contact, each branch is mechanically continuous with the underlying primary filament beneath it, so any load state in the branch, whether tensile from cantilever pull-back or compressive from interaction with a neighboring branch, is shared with the primary filament through this bond. Any construction-stage contribution of compressive load redistribution between adjacent branches would therefore have to be mobilized through this same interlayer bond, rather than through an independently modeled compression path along the branches. In the present framework, this transfer is represented by the bond-shear term in the branch-capacity expression (eq. (4)), which limits the force that can be transmitted from the zig-zag into the underlying primary layer.

At this point, all quantities needed for the admissibility checks are defined: $M^p_{I_0}$ (eq. (1)), $\Delta M_\ell^{\text{load}}$ (eq. (6)), and R_ℓ (eq. (5)). These can be evaluated for any candidate design (n_r, d_r) with $n_r \geq 1$ and $d_r \geq N_c$ (the latter ensures that the tie-back filament lies entirely within the stable region).

2.2. Admissibility checks and design selection

After a lamination cycle completes and its N_c filaments join the stable region, the print continues outward from the new I_0. Each completed cycle becomes one link in a growing structural chain (fig. 8a), and each link must carry not only its own weight but also the unresolved effects transmitted to it by all subsequent links. Admissibility therefore requires passing two checks: the cantilever must survive the pre-bond state when the zig-zag is unbonded dead load (Check 1), and the chain must remain sustainable after the branches bond and begin transferring load (Check 2).

In the most general case, filament length, weight, spacing, and residual force may vary from one cycle to the next, so the build must be evaluated sequentially on a cycle-by-cycle basis. The closed-form screening framework developed here is specialized to the stronger condition of equal filament length and uniform filament spacing. This condition holds for the barrel-vault demonstrator treated in section 3, where all primary filaments trace the same circular arc. Under that specialization, a single zig-zag design (n_r, d_r) is used throughout the build, the same per-cycle balance applies at every completed link, and the admissibility conditions reduce to the two checks below.

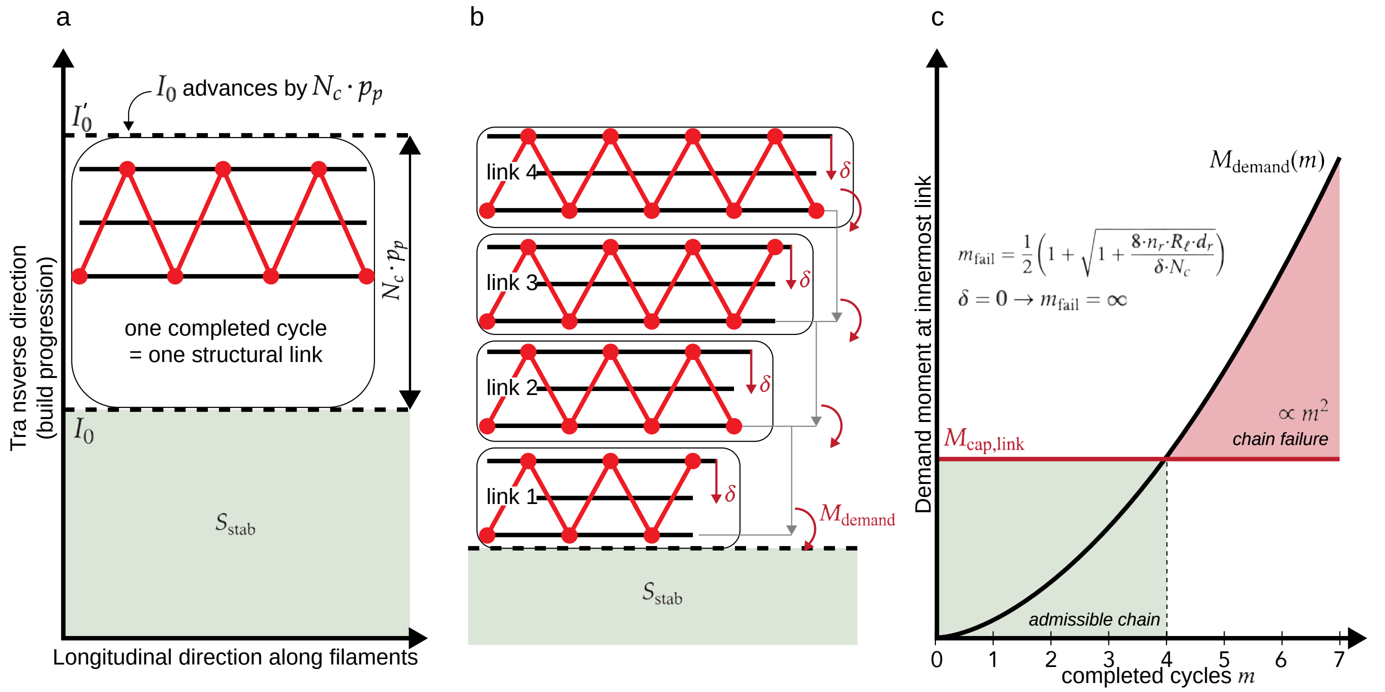

Figure 8: Progressive chain mechanics under uniform filament spacing. **(a)** Single completed lamination cycle: the N_c primary filaments and their zig-zag overlay form one link; the root interface I_0 advances by $N_c\, p_p$. **(b)** Chain of completed cycles: any unresolved downward force from each cycle is transmitted inward and acts at progressively larger lever arms. **(c)** Interpretation of Check 2: indefinite progressive construction requires each completed cycle to be self-sufficient in force balance so that no cumulative deficit propagates down the chain.

2.2.1. *Deposition survivability (Check 1)*

At the pre-bond state, the freshly deposited zig-zag is dead load on the N_c unsupported primary filaments: the tension-tie mechanism is not yet active because the secondary path has not developed sufficient interlayer bond to mobilize tensile force. The root interface must therefore survive the combined weight of the primary cantilever and the unbonded zig-zag:

$$M^{p}_{I_0}(N_c) + \Delta M^{\mathrm{load}}_{\ell}(n_r, d_r) \le M^{\mathrm{cap,C1}}_{I_0}. \qquad \text{(Check 1)}$$

The right-hand side is the interface capacity at the end of the zig-zag deposition pass:

$$M^{\mathrm{cap,C1}}_{I_0} = r_M\left(\frac{N_c\, L_p}{v_p} + \beta\,\frac{2n_r\,\ell_b}{v_{\mathrm{zz}}}\right). \qquad (7)$$

The first term is the primary deposition time and is credited directly using the capacity rate r_M calibrated from the primary-only benchmark (section 3.2). The second term is the zig-zag deposition time. Because the root interface is already loaded while the zig-zag is being deposited, only a fraction of this time is credited toward capacity growth. This fraction is represented by $\beta \in [0, 1]$, where $\beta = 1$ credits the full zig-zag deposition time and $\beta = 0$ credits none of it. With r_M and w_{turn} known independently, β is the sole empirical parameter in Check 1 and is calibrated from zig-zag identification experiments (section 3.3). No intentional dwell time is introduced between the primary pass and the zig-zag pass. The printing sequence is therefore continuous, and the root interface is loaded while it continues to mature. The parameter β represents this loaded maturation interval in a reduced-order form: it credits only the fraction of the zig-zag deposition time that is effective in increasing the root-interface capacity under simultaneous loading. Physically, $\beta < 1$ reflects the disruption of early-age bond development by the simultaneous deposition and loading of the zig-zag pass: colloidal flocculation and early hydrate-bridge formation at the wet-on-wet interface are sensitive to shear and sustained load, so the capacity gain rate under these conditions is lower than the quiescent rate r_M identified from the primary-only benchmark. Thus, Check 1 does not require a time-resolved bond-kinetics law; the net effect of bond development during the loaded zig-zag pass is identified directly from the zig-zag calibration experiments.

Because the zig-zag deposition time depends on n_r, the capacity in eq. (7) is itself design-dependent. This creates the central pre-bond trade-off: increasing n_r adds more restraint nodes, but it also increases both the deposited dead load and the elapsed time required to place it. Check 1 identifies the point beyond which added zig-zag material becomes destabilizing rather than beneficial.

2.2.2. *Chain sustainability (Check 2)*

Once the zig-zag has bonded sufficiently to the substrate, its branches carry tension and transfer load from the cantilever back into the stable region. This post-bond state is not interpreted as instantaneous full bond strength. The mobilized tie-back action is reduced by the engagement fraction η, introduced in eq. (5). This factor accounts for finite bond maturation over the construction cycle, incomplete contact, substrate compliance, and compatibility between the zig-zag path and the primary filaments.

Check 2 therefore evaluates the force balance using the portion of the nominal branch capacity that is actually mobilized over the construction time scale of one cycle.

Each completed cycle must also support the self-weight of the branches meeting at its restraint nodes. Under uniform filament spacing, indefinite progressive construction is possible only if every completed cycle is *self-sufficient* in force balance: the total post-bond restoring action must at least match the primary weight that the cycle introduces. This requirement gives

$$n_r[R_\ell(n_r, d_r, t) - 2\,\mu_\ell\,\ell_b(n_r, d_r)] \geq W_p\,N_c. \tag{Check 2}$$

When Check 2 is not satisfied, each cycle leaves a residual downward force $\delta = W_p N_c - n_r[R_\ell - 2\mu_\ell \ell_b] > 0$ that must be carried by the chain behind it. The unresolved force from cycle k acts at distance $(k-1)\,N_c\,p_p$ from the innermost link, so the cumulative moment grows quadratically with the number of completed cycles m. Equating this demand to the moment capacity of the innermost link, $n_r\,R_\ell\,d_r\,p_p$, gives a closed-form expression for the number of cycles the chain can sustain:

$$m_{\text{fail}} = \frac{1 + \sqrt{1 + \dfrac{8\,n_r\,R_\ell\,d_r}{\delta\,N_c}}}{2}, \qquad D_{\max} = m_{\text{fail}}\,N_c\,p_p. \tag{8}$$

In the limit $\delta \to 0^+$, $m_{\text{fail}} \to \infty$, recovering the indefinite growth condition of Check 2. When Check 2 is satisfied ($\delta \leq 0$), no persistent deficit propagates and the chain can in principle advance indefinitely under the assumed uniform-spacing conditions. For geometries that do not satisfy those conditions, the residual force must be accumulated sequentially on a cycle-by-cycle basis rather than through a single closed-form screening criterion.

2.2.3. Admissible domain and design selection

A design (n_r, d_r) with $n_r \geq 1$ and $d_r \geq N_c$ is admitted for indefinite progressive construction if it passes both Check 1 (eq. (Check 1)) and Check 2 (eq. (Check 2)). The set of all such designs defines the admissible domain $\mathcal{D}_\ell(N_c)$.

Not all admissible designs are equally safe. Two normalized margins quantify the reserve in each check:

$$\gamma_1 = \frac{M_{I_0}^{\text{cap,C1}} - (M_{I_0}^{p} + \Delta M_\ell^{\text{load}})}{M_{I_0}^{\text{cap,C1}}}, \tag{9}$$

$$\gamma_2 = \frac{n_r(R_\ell - 2\mu_\ell \ell_b) - W_p N_c}{W_p N_c}. \tag{10}$$

Both margins are zero at the admissibility boundary and increase toward the interior of $\mathcal{D}_\ell(N_c)$. The most robust design is the one whose weaker check has the largest reserve:

$$(n_r^\star, d_r^\star) = \underset{(n_r,d_r)\in\mathcal{D}_\ell(N_c)}{\arg\max}\ \min\{\gamma_1, \gamma_2\}. \tag{11}$$

Additional constraints—such as limits on deposited length, fabrication time, or material overhead—may be imposed as secondary selection criteria when required by the application.

3. Experimental validation: barrel-vault demonstrator

This section applies the design framework of section 2 to the 3D concrete printing of a barrel vault. The goal is to select a zig-zag design that keeps the cantilever admissible at every cycle, and to validate that selection experimentally.

The framework requires three empirical inputs that cannot be determined from geometry or material tests alone: the interface capacity rate r_M (eq. (2)), the capacity-impairment factor β (eq. (7)), and the engagement fraction η (eq. (5)). Because these three parameters enter different parts of the admissibility formulation, they are identified sequentially rather than jointly. The experimental program is organized accordingly (table 2): Experiment A calibrates r_M from a primary-only print; Experiments C1–C3 calibrate β from zig-zag prints that probe the Check 1 boundary; and the admissible domain is then constructed analytically with η left as a free parameter. Finally, Experiment B validates the selected design and provides an *a posteriori* lower bound on η.

Table 2: Experimental roadmap. Each experiment maps to a stage of the design framework (section 2). The domain construction between Experiments C and B is an analytical step, not a separate experiment.

Stage	Experiment	Design	Purpose	Framework reference
1	A	Primary only	Calibrate r_M	Primary baseline (section 2.1.2)
2	C1–C3	Various (n_r, d_r)	Calibrate β; cross-validate	Check 1 (section 2.2.1)
3	*(analytical)*	—	Construct $\mathcal{D}_\ell$; select $(n_r^\star, d_r^\star)$	Domain & selection (section 2.2.3)
4	B	Selected $(7, 4)$	Validate; bound η a posteriori	Check 2 (section 2.2.2)

3.1. Geometry, fabrication system, and input parameters

The demonstrator is a single-span barrel vault with span $S = 1200$ mm and rise $h_v = 300$ mm, printed on a custom steel support frame (fig. 9). The vault cross-section is a circular arc; from the span and rise, the radius is $R = S^2/(8h_v) + h_v/2 = 750$ mm, and each primary filament traces an arc of length $L_p = 2R \arcsin(S/2R) \approx 1391$ mm. Adjacent primary filaments are spaced at $p_p = 14$ mm center-to-center.

The steel frame provides the springing/startup boundary condition of the barrel-vault demonstrator. The frame-supported startup zone contained $N_{\text{start}} = 10$ primary filaments, corresponding to a supported startup depth of $D_{\text{start}} = N_{\text{start}} p_p = 140$ mm. These initial filaments were deposited directly on the frame and were supported by bearing contact with the steel frame. No internal centering, falsework, or formwork was placed beneath the cantilevering region beyond the startup zone.

The initial root interface $I_0^{(0)}$ is defined as the first interlayer boundary immediately outside the frame-supported startup zone. It is therefore the first interface at which the newly deposited primary filament is carried only by its wet-on-wet bond to the previously printed material, rather than by direct support from the frame. All reported unsupported depths D are measured from $I_0^{(0)}$. At the beginning of printing, the stable region S_{stab} consists of the frame-supported startup zone. After each successful lamination cycle, the progressive root interface I_0 (section 2.2) advances by $N_c p_p$, and the newly stabilized printed material is added to S_{stab}.

Because the barrel vault has a constant cross-section, all primary filaments have identical length L_p, weight W_p, and spacing p_p, satisfying the uniform-filament-spacing condition. The closed-form chain analysis (eq. (8)) therefore applies directly with a single design (n_r, d_r) for all cycles. Beyond the startup zone, each new primary filament is unsupported from below and is stabilized only by the interlayer bond to the previous filament and, after deposition of the secondary path, by the zig-zag tie-back mechanism. The cantilever mechanics of section 2.1.1 therefore govern the printed region beyond $I_0^{(0)}$.

All experiments used the same fabrication platform (fig. 9): a KUKA KR120 R2700 HA robot [58], a MAI MULTIMIX-3D pump [59] at 1.4 L min^{-1}, an accelerator dosing unit, and an inline two-component mixing head. The base material was Sikacrete-752 3D with Sika Sigunit L-5601 AF accelerator (4%) and Sika Plastiment retarder [60] (130 mL/100kg). Primary filaments were deposited at $v_p = 60$ mm s^{-1}; the zig-zag TCP speed was reduced to $v_{zz} = 50$ mm s^{-1} to accommodate turn decelerations. No intentional waiting period was introduced between primary and zig-zag deposition.

For the selected design $(n_r, d_r) = (7, 4)$, one primary pass required $t_p = L_p/v_p \approx 23.2$ s, and the zig-zag pass required $t_{zz} = 2n_r \ell_b / v_{zz} \approx 32.0$ s. The full primary–zig-zag cycle time was therefore approximately 55.2 s. The successful validation print demonstrates that, under this continuous-printing schedule, the zig-zag branches developed sufficient engagement to satisfy the post-bond force-balance condition for the selected design. No changes in material, dosage, or platform were introduced between experiments; the only variable across the test series is the zig-zag geometry.

The turn-deposit weight w_{turn} was measured from flat-surface deposition tests conducted at the same process settings used during vault printing. The zig-zag path was deposited on a flat substrate, the reversal points were isolated, and the excess mass at each reversal was weighed. The measured value is $w_{\text{turn}} = 0.66$ N per reversal. This is a fabrication input determined before the calibration experiments; it does not enter the identification of β or η.

Table 3 collects the full set of inputs required by the formulation. Bead dimensions and cross-sectional areas were measured from deposited test filaments, and the fresh density from deposited samples. The filament tensile capacity and the interfacial bond shear strength were determined from early-age characterization tests on the same mix at ages representative of the construction window (1 min to 3 min). From these material properties, the crossover branch length $\ell_b^\star = A_\ell \sigma_\ell^{\text{cap}} / (\tau_{\ell p}^{\text{cap}} b_\ell)$ evaluates to 64 mm. Any branch longer than this value is material-strength governed: its capacity is capped at $T_b = A_\ell \sigma_\ell^{\text{cap}} = 6.40$ N regardless of how much additional bonded length is available. For the barrel-vault geometry, all candidates with $d_r \geq 5$ produce branches longer than $\ell_b^\star$ for every n_r and therefore fall in the material-strength-governed regime. At shallower tie-back depths the transition depends on n_r: for example, at $d_r = 3$ designs with $n_r \lesssim 14$ produce branches exceeding $\ell_b^\star$, while denser designs produce shorter branches whose capacity is limited by the available bond area.

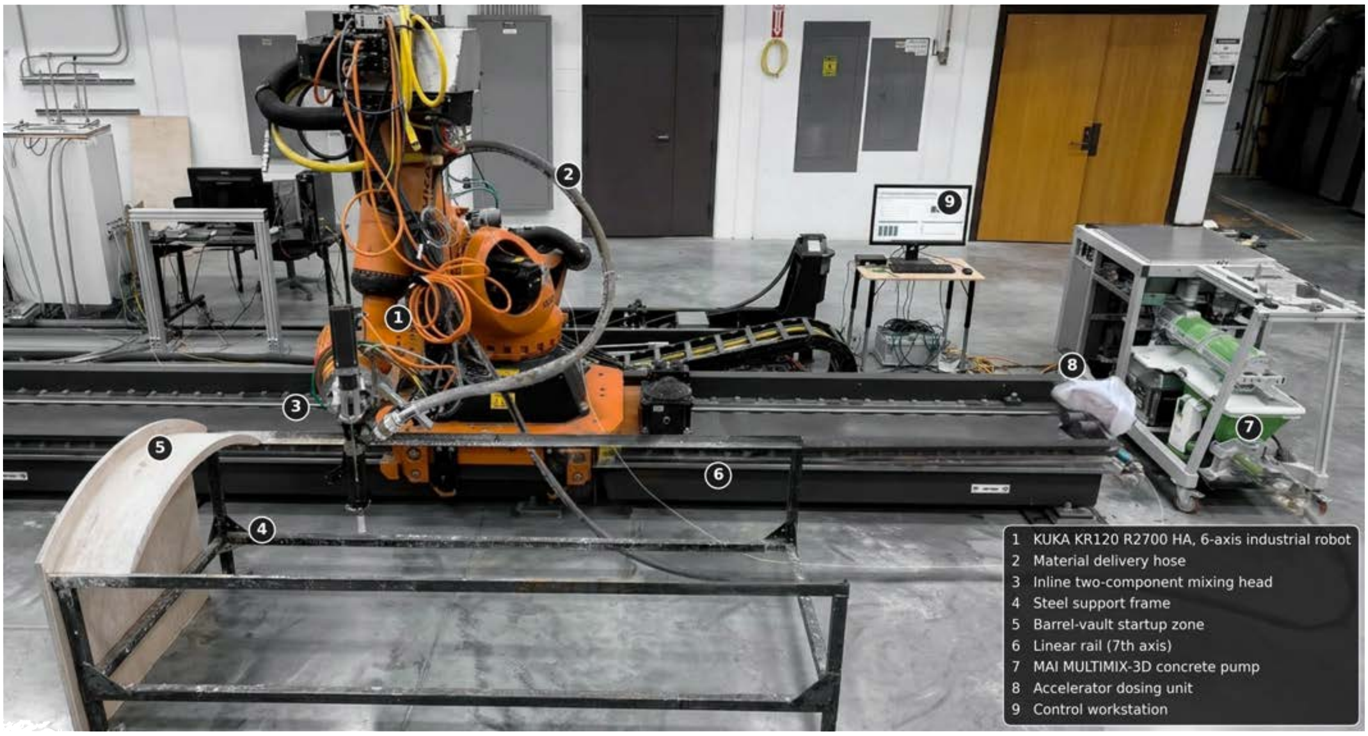

Figure 9: Barrel-vault demonstrator geometry and fabrication setup. The steel support frame (4) provides the springing/startup boundary condition and defines the frame-supported startup zone (5), which contains the first ten primary filaments over a supported depth of 140 mm. The initial root interface $I_0^{(0)}$ is located at the edge of this startup zone; beyond this boundary, the primary filaments extend as a cantilevering printed region without internal centering or formwork. Printing is performed by a KUKA KR120 R2700 HA 6-axis industrial robot (1) mounted on a linear rail serving as a 7th axis (6). Material is delivered through a hose (2) to an inline two-component mixing head (3) at the end effector. The material supply consists of a MAI MULTIMIX-3D concrete pump (7) and an accelerator dosing unit (8), controlled from a dedicated workstation (9). Span $S = 1200$ mm, rise $h_v = 300$ mm, radius $R = 750$ mm, primary arc $L_p \approx 1391$ mm.

Table 3: Input parameters and derived quantities for the barrel-vault case study. Geometric and process inputs are measured or prescribed; material inputs are from early-age characterization tests on the same mix; the capacity-impairment factor is calibrated from the zig-zag experiments (section 3.3) and the engagement fraction bounded from the validation print (section 3.5); derived quantities follow from the formulation of section 2.

Parameter	**Symbol**	**Value**	**Source**
Geometric and process inputs			
Span / rise	$S,\ h_v$	1200, 300 mm	measured
Radius (from S, h_v)	R	750 mm	computed
Primary spacing	p_p	14 mm	prescribed
Primary bead ($b_p \times h_p$; A_p)		30×10 mm; 300 mm^2	measured
Lamination bead ($b_\ell \times h_\ell$; A_ℓ)		20×8 mm; 160 mm^2	measured
TCP speed (primary / zig-zag)	$v_p,\ v_{zz}$	60, 50 mm s^{-1}	prescribed
Turn-deposit weight	w_{turn}	0.66 N	measured (flat-surface test)
Material inputs			
Fresh density	ρ	2100 kg m^{-3}	measured
Admissible axial stress	σ_ℓ^{cap}	40 kPa	characterized
Interfacial bond traction	$\tau_{\ell p}^{\text{cap}}$	5 kPa	characterized
Calibrated / bounded parameters			
Interface capacity rate	r_M	7.79×10^{-3} N m s^{-1}	calibrated (section 3.2)
Capacity-impairment factor	β	0.30	calibrated (section 3.3)
Engagement fraction	η	≥ 0.32	bounded (section 3.5)
Derived quantities			
Filament arc length	L_p	1391 mm	$2R \arcsin(S/2R)$
Primary filament weight	W_p	8.60 N	$\rho A_p L_p g$
Lamination weight per length	μ_ℓ	3.30 N m^{-1}	$\rho A_\ell g$
Crossover branch length	$\ell_b^\star$	64 mm	$A_\ell \sigma_\ell^{\text{cap}} / (\tau_{\ell p}^{\text{cap}} b_\ell)$

The three calibrated inputs—r_M, β, and η—are system-specific and encode the combined effect of mix design, accelerator type and dosage, pump flow rate, nozzle geometry, TCP speeds, and ambient conditions at the time of printing. They are not portable across fabrication setups and must be re-identified whenever any of these factors changes materially. The framework itself, however, is independent of their numerical values: Checks 1 and 2 (eqs. (Check 1) and (Check 2)), the branch mechanics (eqs. (3) to (5)), the chain analysis (eq. (8)), and the selection criterion (eq. (11)) are derived from geometry and force balance alone and apply across the class of systems defined by the structural assumptions of **??**. Applying the framework to a new system within this class therefore requires rerunning the four-print identification sequence used here: one primary-only print to fix r_M (section 3.2), two or three zig-zag prints at increasing density to bound β (section 3.3), and one validation print to confirm Check 1 and bound η from below (section 3.5), but not any modification to the underlying theory. Within a fixed fabrication system, the calibrated parameters are reusable across different vault designs that share the same primary and lamination bead geometries and deposition speeds. This transferability through recalibration is bounded by the structural assumptions stated in section 2.1.1: it accommodates variation across systems for which those assumptions hold, but does not extend the framework to systems whose underlying physics violates one or more of these assumptions, which would require new modeling rather than recalibration alone.

3.2. *Primary-only benchmark (Experiment A)*

Before introducing any zig-zag, the vault was printed with the primary layer alone to establish how far the cantilever can extend before the root interface fails. Without lamination, I_0 remains fixed at $I_0^{(0)}$ and the unsupported depth grows with each deposited filament. The cantilevering portion failed at $D = 42\,\mathrm{mm}$ beyond $I_0^{(0)}$, corresponding to $N_c^{\mathrm{crit}} = 3$ filaments (fig. 10).

Figure 10: Experiment A: primary-only benchmark. The vault is printed without lamination; the cantilevering filaments beyond $I_0^{(0)}$ delaminate at $D = 42\,\mathrm{mm}$ ($N_c^{\mathrm{crit}} = 3$ filaments), peeling away from the root interface under their own weight.

This single failure point is sufficient to calibrate the interface capacity rate r_M (eq. (2)). At the moment of failure, the root moment is $M^p_{I_0}(3) = \frac{1}{2}\,W_p\,p_p \cdot 3^2 = 0.542\,\mathrm{N\,m}$ (eq. (1)), and the elapsed time since the formation of I_0 is $t = 3\,L_p/v_p = 69.5\,\mathrm{s}$. Because the interface failed at exactly this moment, the capacity must have equaled the demand:

$$r_M = M^p_{I_0}/t = 0.542/69.5 = 7.79 \times 10^{-3}\,\mathrm{N\,m\,s^{-1}}.$$

This value is valid under the quasi-static loading conditions of primary-only deposition and over the timescale of the barrel-vault construction window.

3.3. *Calibration of β (Experiments C1–C3)*

With r_M calibrated and w_{turn} measured, every term in Check 1 is known except β. The demand does not depend on β; the capacity grows monotonically with it (eq. (7)). The task is to find the value of β that correctly separates designs that survive from designs that fail.

The parameter β is not a property of the zig-zag design—it is a property of the material and process. It captures how effectively the interface at I_0 gains strength while being actively loaded during zig-zag deposition, compared to the quasi-static conditions under which r_M was calibrated. All experiments use the same material, pump, robot, and zig-zag travel speed v_{zz}. A design with more restraint nodes takes longer to deposit, but the time term $t_{zz} = 2n_r\ell_b/v_{zz}$ already accounts for that. β scales the credited fraction of that time, and the fraction is the same because the deposition conditions are the same.

Three designs of increasing density were printed at $N_c = 1$ and continued until failure (table 4 and fig. 11). All three failed by peel at the root interface—the failure mode governed by Check 1 (section 2.2.1).

Table 4: Zig-zag calibration experiments ($N_c = 1$). Depths measured from $I_0^{(0)}$.

Test	n_r	d_r	α	ℓ_b [mm]	D_{obs} [mm]	**Outcome**
C1	10	3	31.1°	81.2	≈ 56	peel failure after 3–4 cycles
C2	12	4	44.0°	80.6	≈ 56	peel failure after 3–4 cycles
C3	35	2	54.6°	34.3	≈ 14	peel failure at first cycle

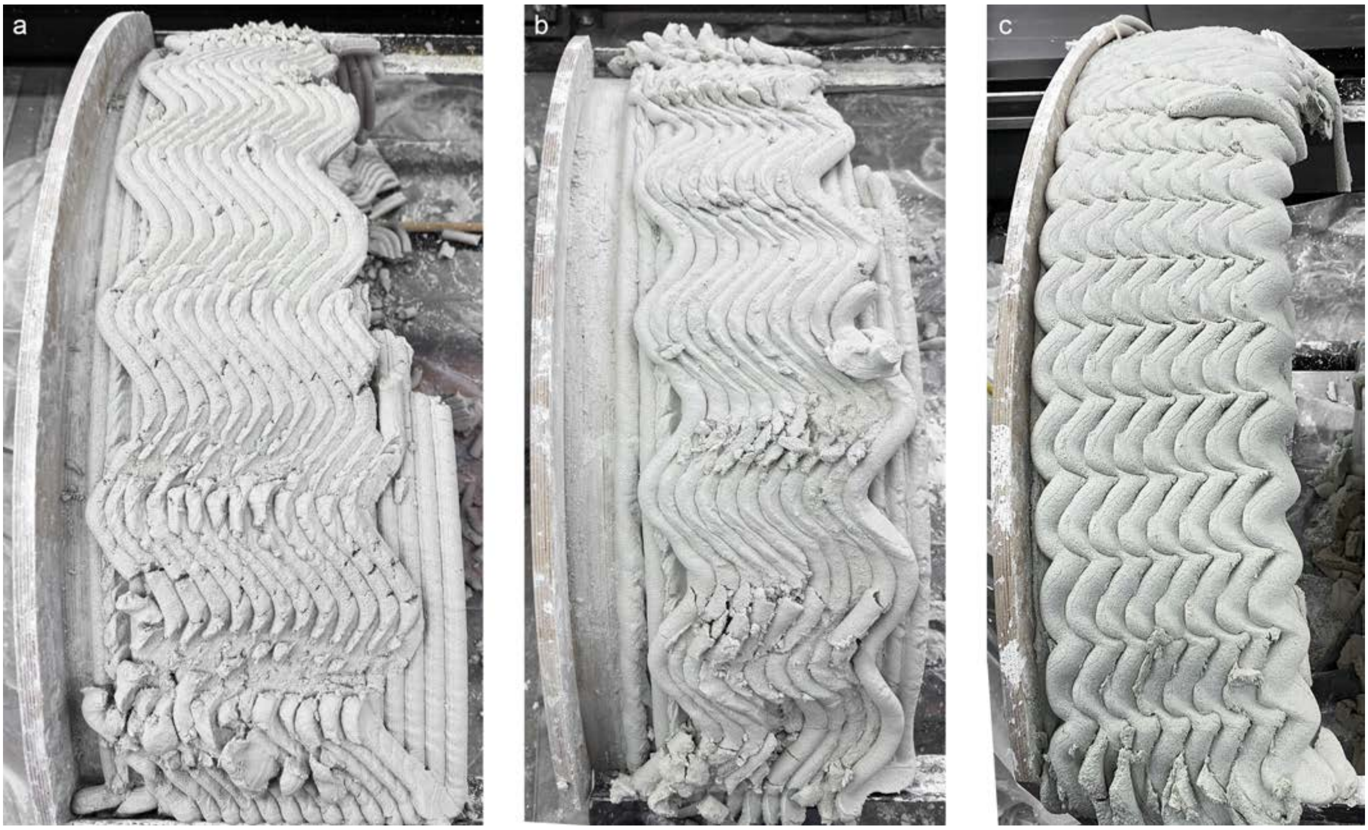

Figure 11: Calibration prints (Experiments C1–C3): photographs and failure locations. **(a)** C1: $(10, 3)$, failure at $D \approx 56$ mm. **(b)** C2: $(12, 4)$, failure at $D \approx 56$ mm. **(c)** C3: $(35, 2)$, failure at $D \approx 14$ mm. Dashed lines indicate $I_0^{(0)}$; arrows mark failure locations.

Each failed experiment provides an upper bound on β. For a given design, setting demand equal to capacity and solving for β gives the critical value $\beta^\star$ at which that design transitions from inadmissible to admissible:

$$\beta^\star = \frac{\dfrac{M^p_{I_0} + \Delta M^{\text{load}}_\ell}{r_M} - t_p}{t_{zz}}. \tag{12}$$

Since the experiment failed, the real β must lie below $\beta^\star$. The three experiments give three bounds: $\beta \leq 0.97$ (C3), $\beta \leq 0.37$ (C2), and $\beta \leq 0.30$ (C1). These thresholds differ because each design has a different demand, not because β itself changes between experiments. C3 has such extreme demand that even crediting 97% of the zig-zag time would barely make it admissible; C1, with moderate demand, becomes admissible once 30% is credited. The binding constraint is $\beta \leq 0.30$ from C1 (fig. 12). Any β at or below this value automatically classifies all three experiments as inadmissible.

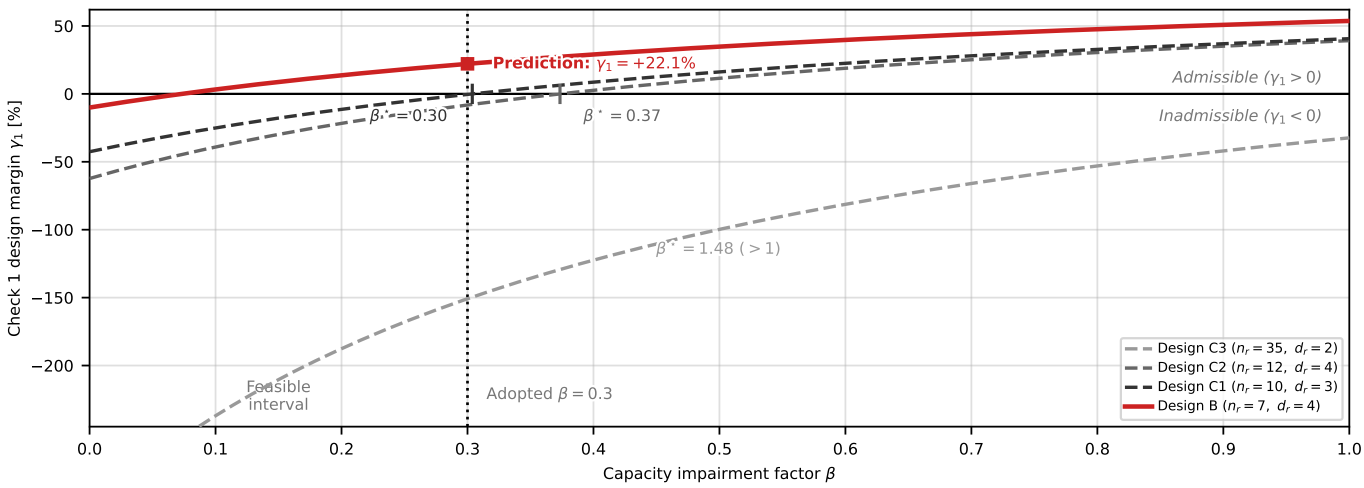

Figure 12: Identification of β. The Check 1 margin γ_1 (eq. (9)) is plotted against β for each design. Failed experiments (C1–C3, solid) require $\gamma_1 < 0$; each curve's zero-crossing is $\beta^\star$. The feasible interval (shaded) is bounded from above by C1 ($\beta^\star = 0.30$). The validation design B (dashed) is a blind prediction: $\gamma_1 = +22\%$ at the adopted β.

The adopted value is $\beta = 0.30$—the upper bound of the feasible interval and the least conservative value consistent with all observed failures. Table 5 confirms that this single constant correctly classifies all four designs across $n_r \in [7, 35]$ and $d_r \in [2, 4]$. The severity ordering ($|\gamma_1|$: C3 ≫ C2 > C1) matches the observed failure progression: C3 failed immediately, C1 and C2 survived several early cycles—consistent with a transient startup-zone capacity surplus that delays but does not prevent steady-state failure. The validation design (7, 4) is predicted admissible with $\gamma_1 = +22.1\%$; this forward prediction is tested in section 3.5.

Table 5: Check 1 evaluation at $\beta = 0.30$, $N_c = 1$, $w_{\text{turn}} = 0.66\,\text{N}$.

	(n_r, d_r)	$\Delta M_\ell^{\text{branch}}$ [N m]	$\Delta M_\ell^{\text{turn}}$ [N m]	Demand [N m]	Capacity [N m]	γ_1
C1	(10, 3)	0.013	0.185	0.258	0.257	−0.4%
C2	(12, 4)	0.011	0.222	0.293	0.271	−8.2%
C3	(35, 2)	0.028	0.647	0.735	0.293	−151%
B	(7, 4)	0.009	0.129	0.199	0.255	+22.1%

3.4. Domain construction and sensitivity analysis

With r_M and β calibrated, the Check 1 boundary is fully determined: for any design (n_r, d_r), the pre-bond demand and capacity can be computed and compared. This defines the upper boundary of the admissible domain—the maximum zig-zag density beyond which the pre-bond dead load overwhelms the interface capacity. This boundary does not depend on η.

The lower boundary comes from Check 2 (eq. (Check 2)): designs with too few restraint nodes cannot compensate the primary weight, and the chain fails progressively. This boundary depends on the engagement fraction η, which controls how much of the nominal branch capacity is actually mobilized (eq. (5)). Unlike β, η cannot be identified from the calibration experiments: all three failed Check 1, so none of them probed the Check 2 boundary. The domain is therefore constructed with η as a free parameter, and the sensitivity of the result is examined directly.

Figure 13 overlays the admissible domain at three values of η: 0.50, 0.75, and 1.0. The upper boundary (Check 1, dashed red) is identical in all three cases—it depends only on β and the measured inputs. The lower boundary (Check 2) shifts: at lower η, more restraint nodes are needed to compensate the primary weight, and the boundary moves to the right; at higher η, sparser designs become admissible, and the boundary moves to the left. The selected design (7, 4) lies well within the admissible region for all three values. The design selection is governed entirely by Check 1; the Check 2 boundary excludes only very sparse designs at low n_r and does not influence the choice.

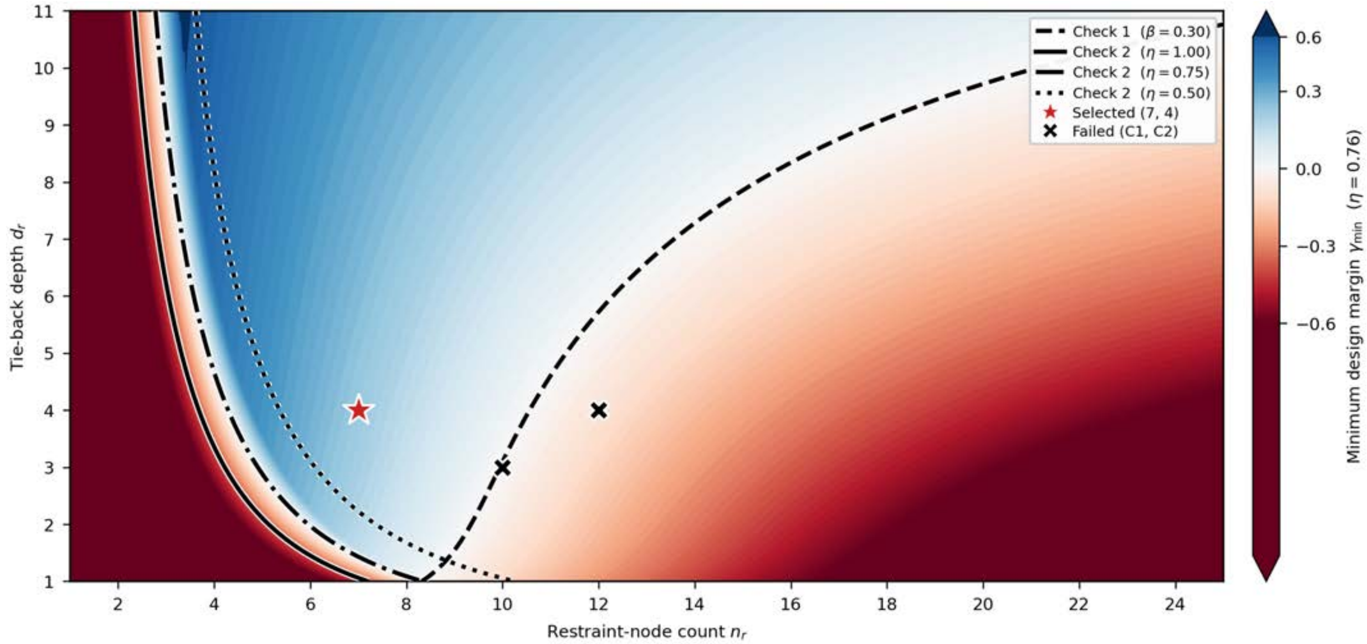

Figure 13: Sensitivity of the admissible domain to the engagement fraction η at $N_c = 1, \beta = 0.30$, $w_{\text{turn}} = 0.66$ N. The Check 1 boundary (dashed red) is identical for all η. The Check 2 boundary is shown at $\eta = 0.50$, 0.75, and 1.0; increasing η shifts it to the left, admitting sparser designs. The selected design $(7, 4)$ ($\star$) lies inside the admissible region for all three values. Calibration tests C1 and C2 ($\times$) fall outside the domain on the high-n_r side.

For design $(7, 4)$ to pass Check 2, η must exceed the critical value at which $\gamma_2 = 0$. Evaluating eq. (Check 2) at $(7, 4)$ gives $\eta_{\text{crit}} = 0.32$. This is a necessary condition verified by the validation print (section 3.5).

Applying the selection criterion (eq. (11)) at $N_c = 1$ yields $(n_r^\star, d_r^\star) = (7, 4)$ regardless of the value of η within [0.32, 1.0]. Table 6 reports the admissibility margins at two representative values.

Table 6: Selected design $(7, 4)$ at $N_c = 1$. The Check 1 margin is independent of η; the Check 2 margin is shown at two illustrative values within the admissible range $\eta \in [\eta_{\text{crit}}, 1.0] = [0.32, 1.0]$ used in fig. 13.

N_c	$n_r^\star$	$d_r^\star$	$\alpha^\star$	γ_1	γ_2 (η=0.50)	γ_2 (η=0.75)
1	7	4	29.4°	+22.1%	+94.7%	+223%

Figure 14 and table 7 together make the two-stage role of the zig-zag explicit, on a single force scale and with the supporting moment-level accounting. Before bonding, the zig-zag is destabilizing: it deposits branch material weighing $2n_r\mu_\ell\ell_b = 5.27$ N and turn deposits weighing $2n_r w_{\text{turn}} = 9.24$ N onto a cantilever that already carries 8.60 N of primary self-weight, raising the pre-bond root moment from 0.060 N m for the primary filament alone to 0.199 N m after the branch and turn-deposit dead loads are included. Check 1 verifies that this temporary penalty remains below the available root-interface capacity of 0.255 N m, credited over the $\beta\, t_{zz}$ deposition interval. After bonding, the same geometry becomes a force amplifier. At $\eta = 0.50$, each restraint node provides an effective transverse force of $R_\ell = 3.14$ N; subtracting the branch self-weight $2\mu_\ell\ell_b = 0.75$ N that each node also carries, the seven engaged nodes together provide a net restoring action of 16.7 N, roughly 1.95 times the primary weight introduced in one cycle (8.60 N). The per-cycle surplus of 8.1 N is what allows the root interface to advance and the build to repeat indefinitely under Check 2. Figure 14 places these three quantities on a single force scale so the cost-versus-benefit of the lamination can be read at a glance.

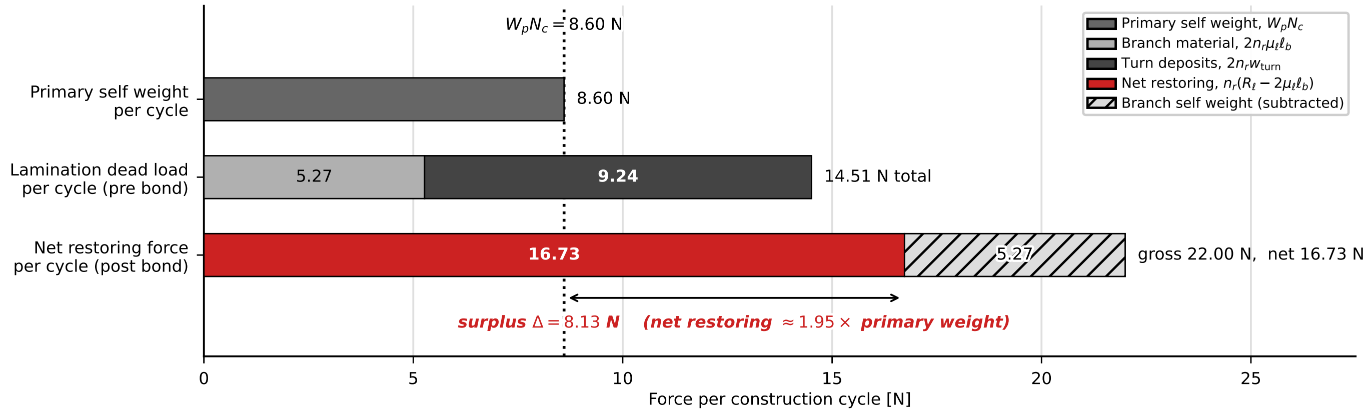

Figure 14: Per-cycle force balance for the selected design $(n_r, d_r) = (7, 4)$ at $N_c = 1$ and illustrative engagement fraction $\eta = 0.50$, expressed in newtons. **Top bar:** primary self-weight introduced per cycle, $W_pN_c = 8.60$ N, the load that each cycle must support. **Middle bar:** total lamination dead load deposited per cycle, decomposed into branch material $2n_r\mu_\ell\ell_b = 5.27$ N and turn deposits $2n_r w_{\text{turn}} = 9.24$ N, summing to 14.51 N. This dead load is temporary and is absorbed by the root-interface capacity that matures over the $\beta\, t_{zz}$ deposition interval, verified by Check 1. **Bottom bar:** post-bond restoring force, with the gross transverse force $n_rR_\ell = 22.00$ N shown together with the branch self-weight $2n_r\mu_\ell\ell_b = 5.27$ N that is subtracted to give the net restoring action $n_r(R_\ell - 2\mu_\ell\ell_b) = 16.73$ N. The dotted reference line marks W_pN_c. The net restoring force exceeds the primary weight by a factor of roughly 1.95, leaving a per-cycle surplus of 8.13 N (Check 2 margin $\gamma_2 = +94.7\%$), which is what allows the build to repeat indefinitely.

Table 7: Physical interpretation of the selected zig-zag design $(n_r, d_r) = (7, 4)$ at $N_c = 1$. Check 1 evaluates the pre-bond penalty, when the zig-zag has not yet developed tensile action. Check 2 evaluates the post-bond force balance, when the diagonal branches mobilize tie-back action. Post-bond quantities are evaluated at $\eta = 0.50$, the lower illustrative engagement value used in fig. 13 and table 6.

State	Quantity	Value	Interpretation
Pre-bond	Primary moment $M^p_{I_0}$	0.060 N m	Moment from one unsupported primary filament
Pre-bond	Zig-zag branch dead-load moment	0.009 N m	Moment from unsupported branch material
Pre-bond	Turn-deposit moment	0.129 N m	Moment from excess material at reversals
Pre-bond	Total demand $M^p_{I_0} + \Delta M^{\text{load}}_\ell$	0.199 N m	Demand checked in Check 1
Pre-bond	Available capacity $M^{\text{cap,C1}}_{I_0}$	0.255 N m	Capacity at end of zig-zag pass
Post-bond	Primary weight per cycle W_pN_c	8.60 N	Downward force introduced by one primary filament
Post-bond	Effective nodal tie-back force R_ℓ at $\eta = 0.50$	3.14 N	Mobilized transverse force at one restraint node
Post-bond	Branch self-weight per node $2\mu_\ell\ell_b$	0.75 N	Dead load carried at each restraint node
Post-bond	Net restoring action $n_r(R_\ell - 2\mu_\ell\ell_b)$	16.7 N	Total post-bond upward/tie-back contribution
Post-bond	Net surplus	8.1 N	Reserve above primary weight; Check 2 satisfied

Check 1 governs ($\gamma_1 < \gamma_2$ at any feasible η). The zig-zag overhead is $L_{zz}/(N_cL_p) = 1.15$: each lamination cycle deposits approximately 15% more length than one primary filament. For the demonstrator, the conservative $N_c = 1$ schedule is adopted:

$$N_c = 1, \quad n_r^\star = 7, \quad d_r^\star = 4.$$

This deposits one primary filament ($t_p = 23$ s) followed by one zig-zag pass ($t_{zz} = 32$ s) per cycle, advancing I_0 by $p_p = 14$ mm.

3.4.1. *Sensitivity to calibrated and measured inputs*

The preceding η sweep examines the post-bond engagement factor, which enters Check 2 but is not identified directly from the Check 1 calibration experiments. We therefore also examined how the selected design responds to the calibrated and measured inputs that enter Check 1. These inputs are the interface capacity rate r_M, the capacity-impairment factor β, the turn-deposit weight w_{turn}, the primary and zig-zag TCP speeds v_p and v_{zz}, the bead areas A_p and A_ℓ, the primary spacing p_p, and the fresh density ρ.

A local one-at-a-time sensitivity screen was performed for the selected design $(n_r, d_r) = (7, 4)$. Each input was perturbed independently around its nominal value, while all other inputs were held fixed, and the resulting change in the Check 1 margin γ_1 was recorded. The perturbation bands were chosen to reflect the relative confidence in the different input classes rather than to represent a statistical uncertainty distribution. The empirically calibrated or process-identified quantities r_M, β, and w_{turn} were varied by ±20%, because they are obtained from a limited set of calibration or measurement tests. Prescribed process and bead-geometry inputs were varied by ±10% (v_p, v_{zz}, A_p, and A_ℓ), while the directly controlled primary spacing and measured fresh density were varied by ±5% (p_p and ρ). This analysis is therefore not a probabilistic uncertainty propagation; it is a local diagnostic used to identify which inputs most strongly control the pre-bond deposition-survivability margin.

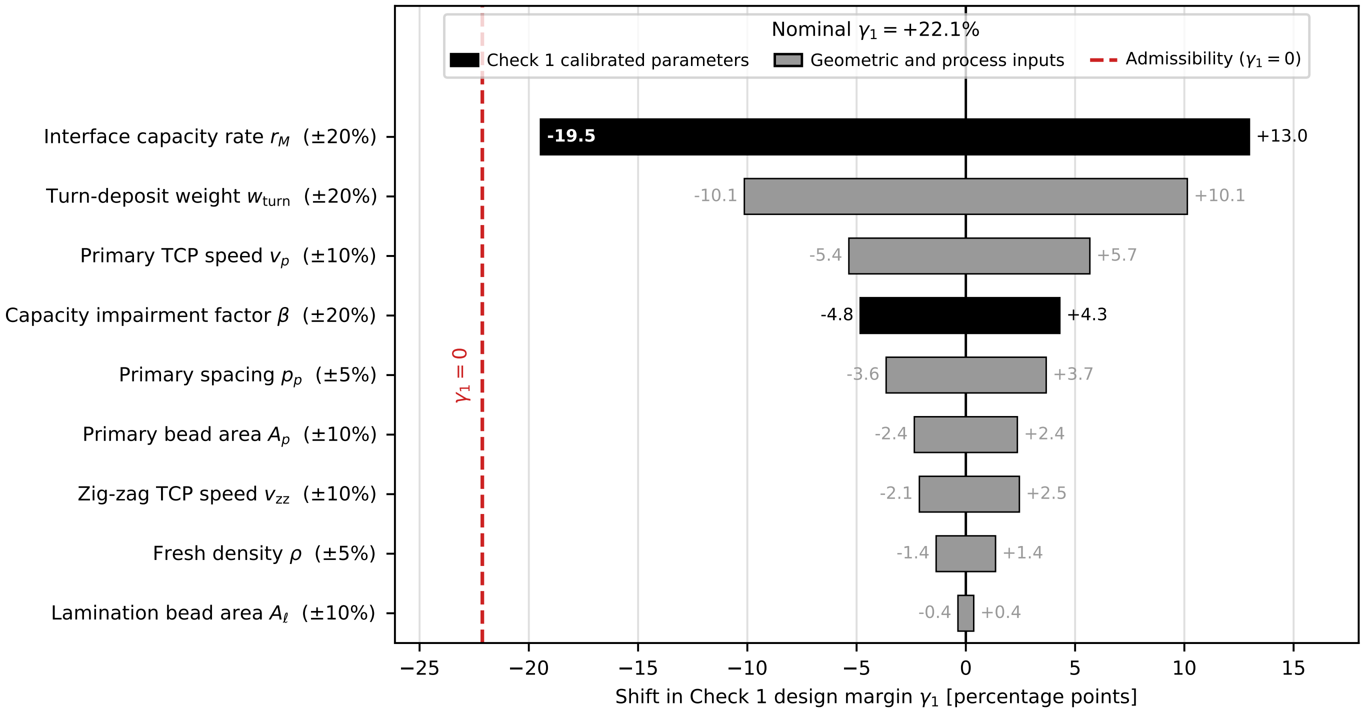

Figure 15: Local one-at-a-time sensitivity of the Check 1 margin γ_1 for the selected design $(n_r, d_r) = (7, 4)$. Each bar shows the shift in γ_1, in percentage points, caused by independently perturbing one input over the indicated range while holding all other inputs fixed. The nominal margin is $\gamma_1 = +22.1\%$; therefore, the dashed red line at -22.1 percentage points marks the admissibility boundary $\gamma_1 = 0$. Positive shifts increase the Check 1 reserve, while negative shifts reduce it. The analysis is a local sensitivity screen for the calibrated process state, not a statistical uncertainty propagation.

Figure 15 shows that the selected design remains on the admissible side of Check 1 for all one-at-a-time perturbations considered. Within the adopted perturbation ranges, the largest adverse shift occurs for the interface capacity rate r_M: a 20% reduction in r_M lowers the margin from $\gamma_1 = +22.1\%$ to $\gamma_1 = +2.7\%$. Thus, although the design remains admissible under this perturbation, it lies close to the Check 1 boundary with respect to the capacity-growth calibration. This is consistent with the role of r_M in Check 1, where the pre-bond demand is compared directly against the time-dependent root-interface capacity credited during primary and zig-zag deposition.

The next most influential inputs are w_{turn}, v_p, and β. This ranking is physically consistent with the structure of Check 1. The turn-deposit weight w_{turn} directly increases the pre-bond dead-load moment, while v_p, v_{zz}, β, and r_M control the amount of root-interface capacity credited during the continuous primary–zig-zag deposition sequence. In particular, faster deposition reduces the elapsed time available for capacity growth, whereas a smaller value of β reduces the fraction of zig-zag deposition time that is credited while the interface is already loaded.

The remaining inputs have a weaker effect on the selected design over the imposed perturbation ranges. Variations in p_p, A_p, v_{zz}, and ρ shift the Check 1 margin by only a few percentage points. The lamination bead area A_ℓ has the smallest influence on Check 1 because, for the selected design, the branch dead-load term is small relative to the total pre-bond demand. This does not imply that A_ℓ is unimportant overall: its primary role is in the post-bond force balance through the branch capacity and the Check 2 condition, which is accounted for separately through the η-dependent admissibility analysis.

Within the perturbation ranges adopted here, the largest reductions in γ_1 occur when the interface capacity rate r_M is decreased or when the turn-deposit weight w_{turn} is increased. This result should be interpreted as a local margin check for the calibrated process state, rather than as a global ranking of parameter importance, since the perturbation bands are intentionally different for calibrated, measured, and prescribed inputs. The result nevertheless identifies the inputs that most directly affect the Check 1 balance: r_M controls the root-interface capacity credited during deposition, while w_{turn} controls a concentrated component of the pre-bond dead-load penalty. These quantities should therefore be remeasured or recalibrated when changes in material formulation, nozzle behavior, deposition speed, or turn execution are expected to alter bond development or turn deposition.

3.4.2. Span scaling of the admissibility domain

The local sensitivity analysis in section 3.4.1 examined how the Check 1 margin responds to perturbations around the calibrated process state for a single geometry. We next examine how the admissibility domain itself depends on the span of the printed vault, with all material and process parameters held fixed at their calibrated values.

Span enters the framework through the primary filament length L_p, which for a barrel vault is the arc length of the cross-section. For a vault with fixed rise-to-span proportions, L_p scales linearly with the span S; for the demonstrator geometry (rise/span = 1/4), $L_p \approx 1.16\,S$. The primary demand $M^{p}_{I_0}$, the zig-zag dead-load moment $\Delta M^{\text{load}}_{\ell}$, the effective nodal force R_ℓ, and the branch length ℓ_b all depend explicitly on L_p. Sweeping L_p over a range of spans and recomputing the admissibility checks therefore shows how the admissible domain $\mathcal{D}_\ell(N_c)$ shifts with structural scale, with no change to the formulation or to the calibrated inputs.

Figure 16 shows the $\gamma_{\min} = 0$ envelope at thirteen spans from 0.6 m to 2.4 m, with the demonstrator span $S = 1.2$ m highlighted. The envelope encloses the admissible region in (n_r, d_r) space at each span, and shifts systematically as span grows: its vertex moves to larger n_r, and the binding check transitions from Check 1 at small spans to Check 2 at large spans. The demonstrator design $(n_r, d_r) = (7, 4)$ sits inside the envelope at $S = 1.2$ m, falls on the inadmissible side of the smaller-span envelopes where Check 1 binds because the primary deposition time is too short to credit sufficient interface capacity, and falls on the inadmissible side of the larger-span envelopes where Check 2 binds because the cantilever weight has grown faster than the available restoring force. A distinct selected design $(n_r^\star, d_r^\star)$ therefore exists at each span, and the framework identifies it directly from the same calibrated inputs.

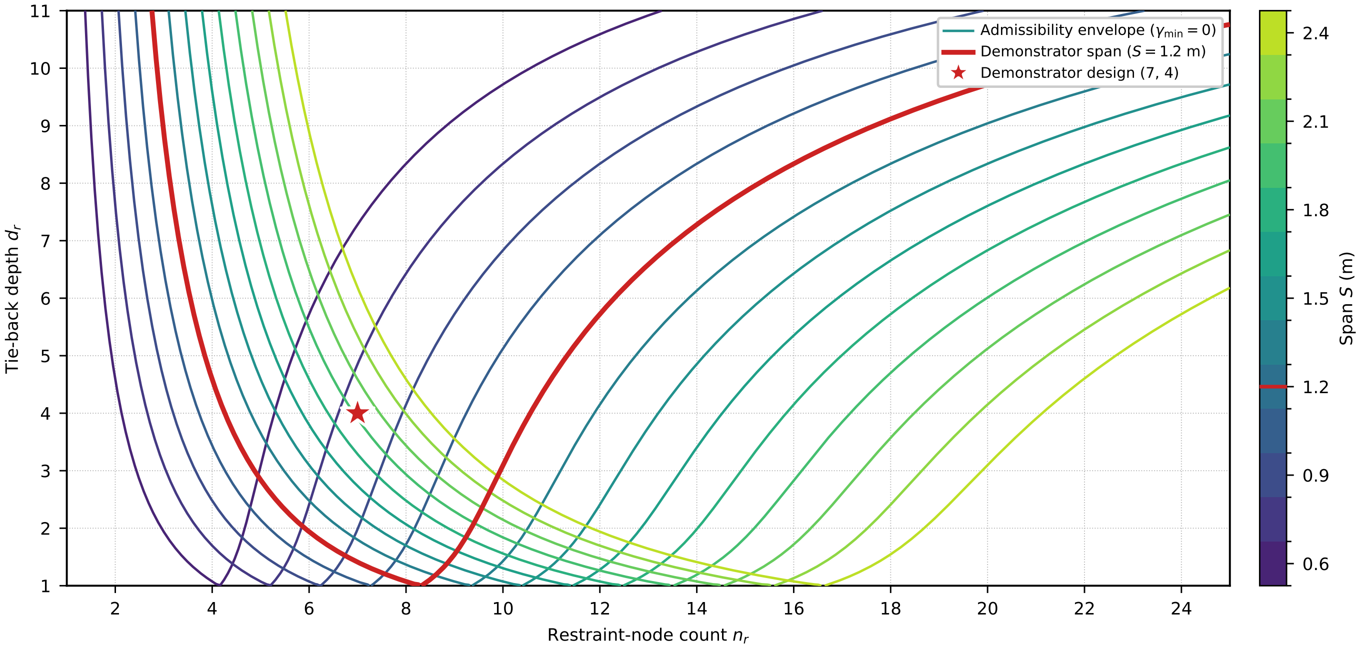

Figure 16: Span scaling of the admissibility envelope in the (n_r, d_r) design space, with all material and process parameters held fixed at their calibrated values. Each colored curve is the $\gamma_{\min} = 0$ envelope at one span, with span value indicated by the colorbar; the envelope encloses the admissible region at that span. The demonstrator span $S = 1.2$ m is overlaid in red and the demonstrator selected design $(7, 4)$ is marked.

3.5. *Validation (Experiment B)*

Experiment B implemented the selected $(7, 4)$ design at $N_c = 1$. Figure 17a shows the alternating primary and zig-zag deposition during printing; fig. 17b shows the completed vault.

Figure 17: Experiment B: full barrel-vault print with $(n_r, d_r) = (7, 4)$ at $N_c = 1$. **(a)** Photograph during deposition: the primary filament and zig-zag lamination pass are visible; I_0 advances by $p_p = 14$ mm after each completed cycle. **(b)** Top-view photograph of the zig-zag pattern on the vault surface.

The vault sustained the entire build sequence with no delamination or visible distress. This confirms the Check 1 prediction ($\gamma_1 = +22.1\%$) and verifies the necessary condition $\eta \geq 0.32$ for Check 2. Table 8 collects the complete experimental program.

Table 8: Complete experimental summary and model predictions at $\beta = 0.30$, $w_{\text{turn}} = 0.66$ N. Depths measured from $I_0^{(0)}$.

	Experiment	(n_r, d_r)	γ_1	**Role**
A	Primary only	—	—	r_M calibration
C1	Zig-zag	$(10, 3)$	-0.4%	β calibration (binding)
C2	Zig-zag	$(12, 4)$	-8.2%	cross-validation
C3	Zig-zag	$(35, 2)$	-151%	cross-validation
B	**Full vault**	$(\mathbf{7}, \mathbf{4})$	$\mathbf{+22.1\%}$	**validated;** $\eta \geq 0.32$ **confirmed**

4. Conclusion

This paper introduced and validated a zig-zag lamination strategy for internal-scaffold-free robotic 3D concrete printing of vaulted structures that undergo cantilevering construction states. The method addresses the construction-stage instability that occurs when freshly deposited primary filaments extend beyond the springing/startup zone and generate an increasing peel moment at the root interface. In the barrel-vault demonstrator, primary-only deposition failed after three unsupported filaments (42 mm beyond the initial root interface). With the selected zig-zag lamination design, the same geometry and material system were printed to completion without delamination or visible distress.

The validation demonstrates the construction-stage function of the zig-zag lamination. Prior to bond engagement, the secondary deposited path increases the demand at the root interface through its added dead load. After sufficient wet-on-wet bond develops, the inclined branches mobilize transverse force components that redirect part of the cantilever load back into the stabilized region. In the reduced-order model, this stabilizing action is represented as a discrete tie-back mechanism rather than as a smeared continuous layer. The result is therefore not only an increase in local unsupported length, but a repeatable construction cycle in which the root interface can advance while remaining admissible under both pre-bond and post-bond conditions.

To make this construction strategy designable, the paper developed a reduced-order analytical framework that parameterizes the zig-zag lamination by the restraint-node count n_r, the tie-back depth d_r, and the cycle depth N_c. The framework reduces the construction-stage admissibility problem to two checks. Check 1 evaluates deposition survivability by comparing the primary self-weight demand plus the unbonded zig-zag dead-load penalty against the time-dependent root-interface capacity. Check 2 evaluates chain sustainability by comparing the post-bond restoring action of the engaged branches against the primary weight introduced in each cycle. These checks provide an explicit screening criterion for candidate lamination designs without requiring numerical simulation.

The experimental program supports the internal consistency of the framework for the tested material, geometry, and process conditions. The interface capacity rate r_M was calibrated from the primary-only benchmark, the capacity-impairment factor β was bounded from zig-zag calibration prints, and the engagement fraction η was bounded from the validation print. With these sequentially identified inputs, the framework classified the primary-only failure, the three deliberately failing zig-zag calibration experiments, and the successful full-vault validation print. The selected design $(n_r, d_r) = (7, 4)$ was predicted to satisfy both admissibility checks and was subsequently validated experimentally.

The closed-form progressive analysis is validated here for a barrel vault with uniform filament spacing, constant primary-filament length, and a fixed material/process system. Under these conditions, a single zig-zag lamination design can be used throughout the build and the residual force balance reduces to a closed-form chain criterion. For geometries with non-uniform spacing, varying filament length, or spatially varying process conditions, the same local pre-bond and post-bond balances remain applicable, but the residual force must be accumulated sequentially on a cycle-by-cycle basis. Similarly, the calibrated parameters are not universal material constants; they must be re-identified when the mix design, accelerator dosage, bead geometry, deposition speed, nozzle configuration, or fabrication environment changes materially.

The present contribution is limited to construction-stage buildability. It does not constitute service-stage structural qualification of the printed vault, nor does it address long-term durability, cured-material strength, or full-scale load-bearing performance. The zig-zag lamination may influence the stiffness, thrust-line admissibility, composite action, and service-stage capacity of the hardened laminated vault, but those effects were not evaluated in this study and require separate structural testing or analysis after hardening. The present framework therefore makes no positive or negative claim about the service-stage structural contribution of the lamination; it evaluates its role as a construction-stage stabilization mechanism.

Within that scope, the framework provides a direct link between geometry, process parameters, early-age material response, and construction-stage stability. This link is essential for scaling the method: filament weight, cycle time, turn-deposit mass, tie-back depth, branch length, and branch capacity all enter explicitly into the admissibility checks. The framework can therefore be used to evaluate how changes in scale, material behavior, or fabrication parameters affect the admissible design domain, provided that the calibrated inputs are re-identified for the target system.

Future work should extend the validation in three directions. First, spatially resolved finite-element simulation can be coupled with the reduced-order framework to evaluate deformation, mixed-mode interfacial response, and non-uniform material state beyond the assumptions of the scalar root-interface model. Second, experimental validation should be extended to non-uniform and doubly curved vault geometries, where the cycle-by-cycle form of the framework becomes necessary. Third, the admissibility checks can be embedded into robotic planning and process control, allowing the lamination schedule to respond to measured material state, deposition quality, and geometric error during fabrication. These extensions would support the transition from design-stage screening to adaptive construction-stage stabilization for internal-scaffold-free robotic 3D concrete printing of vaulted structures.

CRediT Authorship Contribution Statement

Abdallah Kamhawi: Conceptualization; Data curation; Formal analysis; Investigation; Methodology; Project administration; Resources; Software; Validation; Visualization; Writing – original draft; Writing – review & editing.

Mania Aghaei Meibodi: Conceptualization; Data curation; Formal analysis; Investigation; Methodology; Project administration; Resources; Validation; Writing – review & editing.

Funding

This work was supported by the University of Michigan Research Catalyst and Innovation Program.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data Availability

The experimental data, input parameters, and computational scripts Supporting the findings of this study are openly available at https://doi.org/10.5281/zenodo.19203841.

References

[1] P. Block, Thrust network analysis: exploring three-dimensional equilibrium, Ph.D. thesis, Massachusetts Institute of Technology (2009).
URL https://hdl.handle.net/1721.1/49539

[2] P. Block, M. DeJong, L. Davis, J. Ochsendorf, Tile vaulted systems for low-cost construction in Africa, ATDF Journal 7 (1/2) (2010) 4–13.

[3] G. R. Collins, The transfer of thin masonry vaulting from Spain to America, Journal of the Society of Architectural Historians 27 (3) (1968) 176–201. doi:10.2307/988501.

[4] R. Guastavino, Essay on the Theory and History of Cohesive Construction, Applied Especially to the Timbrel Vault, Ticknor, Boston, Massachusetts, 1893.

[5] J. Ochsendorf, Guastavino Vaulting: The Art of Structural Tile, Princeton Architectural Press, New York, 2010, photographs by Michael Freeman.

[6] M. Nuh, R. Oval, J. Orr, Designing and prototyping of a reconfigurable segmented fan concrete shell as a flooring system, Architecture, Structures and Construction 5 (1) (2025) 8. doi:10.1007/s44150-024-00129-7.

[7] L. He, H. E. Fairclough, M. Gilbert, A. Liew, K. Bołbotowski, Minimum material vault designs generated via adaptive layout optimization, Engineering Structures 334 (2025) 120213. doi:10.1016/j.engstruct.2025.120213.

[8] J. Heyman, The stone skeleton, International Journal of Solids and Structures 2 (2) (1966) 249–279. doi:10.1016/0020-7683(66)90018-7.

[9] A. M. Neville, Properties of Concrete, 5th Edition, Pearson Education, Harlow, England, 2011.

[10] P. K. Mehta, P. J. M. Monteiro, Concrete: Microstructure, Properties, and Materials, 4th Edition, McGraw-Hill Education, New York, 2014.

[11] L. C. Lancaster, Innovative Vaulting in the Architecture of the Roman Empire: 1st to 4th Centuries CE, Cambridge University Press, 2015.

[12] C. Meyer, M. H. Sheer, Do concrete shells deserve another look?, Concrete International 27 (10) (2005) 43–50.

[13] G. Tang, An overview of historical and contemporary concrete shells, their construction and factors in their general disappearance, International Journal of Space Structures 30 (1) (2015) 1–12. doi:10.1260/0266-3511.30.1.1.

[14] D. López López, P. Roca, A. Liew, T. Van Mele, P. Block, Tile vaults as integrated formwork for reinforced concrete: Construction, experimental testing and a method for the design and analysis of two-dimensional structures, Engineering Structures 188 (2019) 233–248. doi:10.1016/j.engstruct.2019.03.034.

[15] B. Kromoser, J. Kollegger, Building concrete shells without formwork and falsework, in: Proceedings of fib Symposium 2016: Performance-based Approaches for Concrete Structures, Cape Town, South Africa, 2016.

[16] L. Scheder-Bieschin, T. Van Mele, P. Block, Bending-active formwork systems for concrete shells – a classification and state-of-the-art review, Structures 67 (2024) 106841. doi:10.1016/j.istruc.2024.106841.

[17] M. Rippmann, T. Van Mele, M. Popescu, E. Augustynowicz, T. M. Echenagucia, C. J. C. Barentin, U. Frick, P. Block, The Armadillo Vault: Computational design and digital fabrication of a freeform stone shell, in: Advances in Architectural Geometry 2016, vdf Hochschulverlag AG, 2016, pp. 344–363. doi:10.3218/3778-4_23.

[18] M. Vatandoost, A. Ekhlassi, M. Golabchi, M. Rahbar, P. von Buelow, Fabrication methods of shell structures, Automation in Construction 165 (2024) 105570. doi:10.1016/j.autcon.2024.105570.

[19] E. P. G. Bruun, R. Pastrana, V. Paris, A. Beghini, A. Pizzigoni, S. Parascho, S. Adriaenssens, Three cooperative robotic fabrication methods for the scaffold-free construction of a masonry arch, Automation in Construction 129 (2021) 103803. doi:10.1016/j.autcon.2021.103803.

[20] A. Kamhawi, M. A. Meibodi, Techniques and strategies in extrusion based 3D concrete printing of complex components to prevent premature failure, Automation in Construction 168 (2024) 105768. doi:10.1016/j.autcon.2024.105768.

[21] A. Kamhawi, J. Brown, A. Fahmy, M. Aghaei Meibodi, Waste-free production of ultra-thin concrete panels via robotic 3D printing and CNC dynamic bed device, in: RILEM International Conference on Concrete and Digital Fabrication, Springer, 2024, pp. 99–106. doi:10.1007/978-3-031-70031-6_11.

[22] A. Kamhawi, Y. Lin, C. Watson, K. Barton, M. A. Meibodi, A framework for process anomaly detection in 3D concrete printing, Additive Manufacturing (2025) 105056doi:10.1016/j.addma.2025.105056.

[23] M. A. Meibodi, Y. Lin, H. Chen, Hybrid approaches towards 3D concrete printing for lightweight reinforced concrete structures, in: Digital Concrete 2024 – Supplementary Proceedings, 2024.

[24] Y. Lin, A. Bayramvand, M. Aghaei Meibodi, Towards lightweight structure: Coupling topology optimization with non-planar 3D concrete printing, in: The International Conference on Computational Design and Robotic Fabrication, Springer, 2024, pp. 3–12. doi:10.1007/978-981-99-8405-3_1.

[25] Y. Lin, A. Bayramvand, M. A. Meibodi, Branch wall: Developing topology informed non-planar toolpath and variable deposition rate for 3D concrete printing of topology optimized load bearing wall, SSRN Electronic Journal (2023). doi:10.2139/ssrn.4663173.

[26] N. Roussel, Rheological requirements for printable concretes, Cement and Concrete Research 112 (2018) 76–85. doi:10.1016/j.cemconres.2018.01.019.

[27] S. Muthukrishnan, S. Ramakrishnan, J. Sanjayan, Technologies for improving buildability in 3D concrete printing, Cement and Concrete Composites 122 (2021) 104144. doi:10.1016/j.cemconcomp.2021.104144.

[28] V. Mechtcherine, F. P. Bos, A. Perrot, W. R. L. da Silva, V. N. Nerella, S. Fataei, R. J. M. Wolfs, M. Sonebi, N. Roussel, Extrusion-based additive manufacturing with cement-based materials – production steps, processes, and their underlying physics: A review, Cement and Concrete Research 132 (2020) 106037. doi:10.1016/j.cemconres.2020.106037.

[29] R. A. Buswell, W. R. L. De Silva, S. Z. Jones, J. Dirrenberger, 3D printing using concrete extrusion: A roadmap for research, Cement and Concrete Research 112 (2018) 37–49. doi:10.1016/j.cemconres.2018.05.006.

[30] B. N. Turner, R. Strong, S. A. Gold, A review of melt extrusion additive manufacturing processes: I. Process design and modeling, Rapid Prototyping Journal 20 (3) (2014) 192–204. doi:10.1108/RPJ-01-2013-0012.

[31] T. DebRoy, H. L. Wei, J. S. Zuback, T. Mukherjee, J. W. Elmer, J. O. Milewski, A. M. Beese, A. Wilson-Heid, A. De, W. Zhang, Additive manufacturing of metallic components – process, structure and properties, Progress in Materials Science 92 (2018) 112–224. doi:10.1016/j.pmatsci.2017.10.001.

[32] N. Roussel, A thixotropy model for fresh fluid concretes: Theory, validation and applications, Cement and Concrete Research 36 (10) (2006) 1797–1806. doi:10.1016/j.cemconres.2006.05.025.

[33] A. Perrot, D. Rangeard, A. Pierre, Structural built-up of cement-based materials used for 3D-printing extrusion techniques, Materials and Structures 49 (4) (2016) 1213–1220. doi:10.1617/s11527-015-0571-0.

[34] T. Wangler, E. Lloret, L. Reiter, N. Hack, F. Gramazio, M. Kohler, M. Bernhard, B. Dillenburger, J. Buchli, N. Roussel, R. Flatt, Digital concrete: Opportunities and challenges, RILEM Technical Letters 1 (2016) 67–75. doi:10.21809/rilemtechlett.2016.16.
URL https://letters.rilem.net/index.php/rilem/article/view/16

[35] A. S. J. Suiker, Mechanical performance of wall structures in 3D printing processes: Theory, design tools and experiments, International Journal of Mechanical Sciences 137 (2018) 145–170. doi:10.1016/j.ijmecsci.2018.01.010.

[36] R. J. M. Wolfs, F. P. Bos, T. A. M. Salet, Early age mechanical behaviour of 3D printed concrete: Numerical modelling and experimental testing, Cement and Concrete Research 106 (2018) 103–116. doi:10.1016/j.cemconres.2018.02.001.

[37] F. Bos, R. Wolfs, Z. Ahmed, T. Salet, Additive manufacturing of concrete in construction: Potentials and challenges of 3D concrete printing, Virtual and Physical Prototyping 11 (3) (2016) 209–225. doi:10.1080/17452759.2016.1209867.

[38] D. Jiao, R. De Schryver, C. Shi, G. De Schutter, Thixotropic structural build-up of cement-based materials: A state-of-the-art review, Cement and Concrete Composites 122 (2021) 104152. doi:10.1016/j.cemconcomp.2021.104152.

[39] C. Borg Costanzi, Z. Y. Ahmed, H. R. Schipper, F. P. Bos, U. Knaack, R. J. M. Wolfs, 3D printing concrete on temporary surfaces: The design and fabrication of a concrete shell structure, Automation in Construction 94 (2018) 395–404. doi:10.1016/j.autcon.2018.06.013.

[40] C. A. Battaglia, K. Verian, M. F. Miller, DE:Stress pavilion, in: M. Yablonina, A. Marcus, S. Doyle, M. del Campo, V. Ago, B. Slocum (Eds.), ACADIA 2020: Distributed Proximities, Volume II: Projects, ACADIA, Online and Global, 2020, pp. 202–207.

[41] S. Bhooshan, V. Bhooshan, J. Megens, T. Casucci, T. Van Mele, P. Block, Print-path design for inclined-plane robotic 3D printing of unreinforced concrete, in: Robotic Fabrication in Architecture, Art and Design 2022, Springer International Publishing, Cham, Switzerland, 2022, pp. 188–197. doi:10.1007/978-3-031-13249-0_16.

[42] S. Bodea, A. Dell'Endice, V. Aloutsanidis, T. Bürgin, T. Van Mele, P. Block, Phoenix 3D-concrete-printed bridge – falsework system and construction process development, Proceedings of IASS Annual Symposia 2024 (18) (2024) 1–11.
URL https://www.ingentaconnect.com/content/iass/piass/2024/00002024/00000018/art00036

[43] A. Brodmann, E. Damtsas, C. Schult, B. Spaeth, M. Herrmann, DesignBuild project of a 3D printed concrete shell with a segmented and stereotomic construction concept, in: O. Kontovourkis, M. C. Phocas, G. Wurzer (Eds.), Proceedings of the 42nd Conference on Education and Research in Computer Aided Architectural Design in Europe (eCAADe 2024), Volume 1: Data-Driven Intelligence, eCAADe, Nicosia, Cyprus, 2024, pp. 105–114. doi:10.52842/conf.ecaade.2024.1.105.

[44] A. Dell'Endice, S. Bouten, T. Van Mele, P. Block, Structural design and engineering of Striatus, an unreinforced 3D-concrete-printed masonry arch bridge, Engineering Structures 292 (2023) 116534. doi:10.1016/j.engstruct.2023.116534.

[45] A. Lin, A. Goel, D. H. A. Wong, C. Yeo, J. Chung, S. D. Pang, C.-H. Wang, H. Taylor, H. W. Kua, Compressive load-dominated concrete structures for customized 3D-printing fabrication, Automation in Construction 141 (2022) 104467. doi:10.1016/j.autcon.2022.104467.

[46] L. Prasittisopin, T. Sakdanaraseth, V. Horayangkura, Design and construction method of a 3D concrete printing self-supporting curvilinear pavilion, Journal of Architectural Engineering 27 (3) (2021). doi:10.1061/(ASCE)AE.1943-5568.0000485.

[47] P. Carneau, R. Mesnil, N. Roussel, O. Baverel, An exploration of 3D printing design space inspired by masonry, in: Proceedings of the IASS Annual Symposium 2019, 2019. doi:10.5281/zenodo.3563672.

[48] P. Carneau, R. Mesnil, N. Roussel, O. Baverel, Additive manufacturing of cantilever structures: From masonry to concrete 3D printing, Automation in Construction 116 (2020) 103184. doi:10.1016/j.autcon.2020.103184.

[49] A. Curth, T. Brodesser, L. Sass, C. Mueller, Multi-objective optimization of 3D printed shell toolpaths, Proceedings of the IASS Annual Symposia 2020 (3) (2020) 1–13.
URL https://www.ingentaconnect.com/content/iass/piass/2020/00002020/00000003/art00019

[50] F. Chovghi, C. Richter, P. D'Acunto, Fabrication-aware structural form-finding for additive manufacturing: An equilibrium-based approach, Automation in Construction 178 (2025) 106386. doi:10.1016/j.autcon.2025.106386.

[51] G. Duarte, N. C. Brown, J. Pinto Duarte, Workflow for generating, simulating, and optimizing form and toolpath in 3D concrete printing of vaults, Automation in Construction 166 (2024) 105590. doi:10.1016/j.autcon.2024.105590.

[52] A. Jipa, A. Anton, L. Gebhard, B. Dillenburger, The Nubian slab: 3D concrete printed stay-in-place formwork for vaulted slabs, in: Robotic Fabrication in Architecture, Art and Design 2024, Springer, Cham, Switzerland, 2024, pp. 389–396. doi:10.1007/978-3-031-70031-6_45.

[53] M. Motamedi, R. Mesnil, R. Oval, O. Baverel, Scaffold-free 3D printing of shells: Introduction to patching grammar, Automation in Construction 139 (2022) 104306. doi:10.1016/j.autcon.2022.104306.

[54] V. Ruitinga, E. C. Avramiea, Custom-made printheads empowering 3D printed concrete: Innovations in digital design and fabrication of complex prefabricated elements, in: Robotic Fabrication in Architecture, Art and Design 2024, Springer, Cham, Switzerland, 2024, pp. 239–261. doi:10.1007/978-3-031-47913-7_13.

[55] R. Guastavino, Essay on the Theory and History of Cohesive Construction, Applied Especially to the Timbrel Vault, Ticknor, Boston, Massachusetts, 1893.

[56] Cámara, Proceso de construcción de las bóvedas, Photograph, villanueva de la Cañada, Madrid, Spain (1941).

[57] G. R. Collins, The transfer of thin masonry vaulting from Spain to America, Journal of the Society of Architectural Historians 27 (3) (1968) 176–201. doi:10.2307/988501.

[58] KUKA robotics and automation solutions, accessed: 2025-01-12 (2025).
URL https://www.kuka.com/en-us

[59] MAI® MULTIMIX-3D: Precision and innovation in 3D concrete printing, accessed: 2025-01-12 (2025).
URL https://mai.at/en/product-range/3d-printing/mai-multimix-3d

[60] 3D concrete printing, accessed: 2025-01-12 (2025).
URL https://usa.sika.com/en/construction/concrete/3d-concrete-printing.html