Frahan StonePack — Algorithm Derivations

(LaTeX source for Lean proof management)

Frahan StonePack

2026-06-18

Scope. Formal math for the algorithms implemented in the 2026-06-18 session: GPR fracture-surface kriging (regression kriging, k-planes, robust clip, hull mask), Panel Tile Surface (best-fit-plane planarization), and Slab Trim (convex greedy trim, Imai–Iri concave kerf-follow) plus concave nesting. Definitions and theorems are written so each maps to a Lean `def`/`theorem`.

1 Regression kriging of a fracture bed

Definition 1.1 (Bed data)**.** A bed is a finite set of picks $\{(x_i, y_i, d_i)\}_{i=1}^n$, $(x_i, y_i) \in \mathbb{R}^2$ planimetric, $d_i \in \mathbb{R}$ depth. Write $p_i = (x_i, y_i)$.

Definition 1.2 (Least-squares dip plane)**.** The dip plane is $\pi(p) = a^\top p + c$ with $(a, c) = (a_1, a_2, c) \in \mathbb{R}^3$ minimizing $\sum_i (d_i - a^\top p_i - c)^2$. Equivalently (a, c) solves the normal equations $M\theta = b$ where $M = \sum_i \tilde{p}_i \tilde{p}_i^\top$, $b = \sum_i d_i \tilde{p}_i$, $\tilde{p}_i = (x_i, y_i, 1)$.

Lemma 1.1 (Existence/uniqueness)**.** *If the p_i are not all collinear then $M \succ 0$ and the dip plane is unique.*

Proof. $M = \sum_i \tilde{p}_i \tilde{p}_i^\top \succeq 0$. For $u \in \mathbb{R}^3$, $u^\top M u = \sum_i (\tilde{p}_i^\top u)^2 = 0$ iff every p_i lies on the line $u_1 x + u_2 y + u_3 = 0$; non-collinearity forbids this for $u \neq 0$, so $M \succ 0$ and $\theta = M^{-1} b$ is unique. □

Definition 1.3 (Residual and Gaussian covariance)**.** Residuals $r_i = d_i - \pi(p_i)$. With sill $\sigma^2 = \mathrm{Var}(r)$, range $\rho > 0$, nugget $\tau \geq 0$, the covariance is $C(h) = \sigma^2 \exp(-h^2/\rho^2) + \tau\, \mathbf{1}[h = 0]$, and $K_{ij} = C(\|p_i - p_j\|)$.

Definition 1.4 (Regression-kriging predictor)**.** $\hat{d}(q) = \pi(q) + k(q)^\top \alpha$ where $k(q)_i = \sigma^2 \exp(-\|q - p_i\|^2/\rho^2)$ and $\alpha = K^{-1} r$. (The residual term is the simple-kriging BLUP of r at q.)

Theorem 1.2 (Exactness vs. smoothing)**.** *If $\tau = 0$ and the p_i are distinct then $\hat{d}(p_j) = d_j$ for all j (interpolation). If $\tau > 0$ the predictor is a strict smoother: the residual correction $\hat{d}(p_j) - \pi(p_j) = k(p_j)^\top (K_0 + \tau I)^{-1} r$ has ℓ_2 norm strictly less than $\|r\|$, decreasing in τ.*

Proof. $\tau = 0$, distinct points $\Rightarrow$ K is the Gaussian kernel Gram matrix, SPD, and $k(p_j)^\top K^{-1} = e_j^\top$ (row j of KK^{-1}), so $\hat{d}(p_j) = \pi(p_j) + r_j = d_j$. For $\tau > 0$, $K = K_0 + \tau I$ with $K_0 \succeq 0$; then $k(p_j)^\top (K_0 + \tau I)^{-1} r$ has ℓ_2-norm bounded by $\|r\|\, \|K_0\| / (\lambda_{\min}(K_0) + \tau) < \|r\|$ scaling, monotonically decreasing in τ. □

Remark (Lean). Theorem 1.2 is the formal statement of why a positive nugget removes the Gaussian-GP oscillation while a near-zero nugget interpolates; it reduces to SPD linear algebra.

2 k-planes bed separation

Definition 2.1 (k-planes assignment)**.** Given picks and k planes $\pi_1, \ldots, \pi_k$, the assignment is $\ell(i) = \arg\min_c |d_i - \pi_c(p_i)|$. A configuration $(\ell, \{\pi_c\})$ is *stable* if each π_c is the LS dip plane of $\{i : \ell(i) = c\}$ and no pick changes assignment.

Theorem 2.1 (Monotone convergence)**.** *Alternating (i) refit each π_c from its members and (ii) reassign each i to $\arg\min_c |d_i - \pi_c(p_i)|$ does not increase the cost $J = \sum_i \min_c (d_i - \pi_c(p_i))^2$, and terminates at a stable configuration in finitely many steps.*

Proof. Step (i) minimizes $\sum_{i:\ell(i)=c}(d_i - \pi_c(p_i))^2$ over π_c (Lemma 1.1), so it cannot raise J. Step (ii) replaces each term by a $\leq$ term, so it cannot raise J. $J \geq 0$ and there are finitely many assignments ℓ (k^n); a strictly decreasing-or-equal sequence over a finite assignment set with optimal planes per assignment cannot cycle, hence terminates. This is Lloyd's argument with points replaced by affine 1-codimensional fits. □

3 Robust 2σ clip

Definition 3.1. Given residuals $r_i = d_i - \pi(p_i)$ with mean μ, std s, the 2σ-clip keeps $I' = \{i : |r_i - \mu| \leq 2s\}$, refits π on I', and repeats.

Proposition 3.1. *Each clip pass is non-expansive on the kept set ($I' \subseteq I$) and the per-bed residual variance is non-increasing across passes until a fixed point ($I' = I$) is reached; with the ≥ 4-point floor it terminates.*

4 Convex-hull surface mask

Definition 4.1 (Convex hull)**.** $\mathrm{conv}(S)$ for finite $S \subset \mathbb{R}^2$ is the smallest convex set containing S; its boundary vertices are computed by Andrew's monotone chain in $O(n \log n)$.

Definition 4.2 (Hull mask)**.** A grid point q is drawn iff $q \in \mathrm{conv}(\{p_i\})$ (expanded outward by margin m), tested by ray casting: q is inside iff a generic ray from q crosses the boundary an odd number of times.

Proposition 4.1 (No extrapolation)**.** *With the hull mask the reconstructed surface is supported on* $\mathrm{conv}(\{p_i\})$*; hence the bed is never evaluated outside the planimetric convex hull of its own picks, eliminating extrapolation overshoot.*

5 Panel Tile Surface: planarization

Definition 5.1 (Quad and best-fit plane)**.** A panel quad is $Q = \{q_0, q_1, q_2, q_3\} \subset \mathbb{R}^3$. Its best-fit plane Π has unit normal ν the eigenvector of $\sum_k (q_k - \bar{q})(q_k - \bar{q})^\top$ for the least eigenvalue ($\bar{q}$ the centroid).

Definition 5.2 (Planarity and cut tile)**.** Planarity $\delta(Q) = \max_k \mathrm{dist}(q_k, \Pi) = \max_k |\nu^\top (q_k - \bar{q})|$. The cut tile is the projection $\hat{q}_k = q_k - (\nu^\top (q_k - \bar{q}))\nu$ mapped by the rigid motion $\Pi \to \{z = 0\}$.

Lemma 5.1 (δ minimizes worst-case deviation among planes through $\bar{q}$)**.** Π *minimizes* $\sum_k \mathrm{dist}(q_k, \cdot)^2$ *over planes through $\bar{q}$; and $\delta(Q)$ is the Chebyshev (max) deviation of the corners from that plane.*

Proof. For a plane through $\bar{q}$ with normal u, $\sum_k \mathrm{dist}^2 = u^\top\big(\sum_k (q_k-\bar{q})(q_k-\bar{q})^\top\big)u$, minimized over $\|u\|=1$ at the least eigenvector (Rayleigh quotient). δ is then the max over k of $|u^\top(q_k-\bar{q})|$ by definition. ☐

Use. $\delta(Q)$ is the fabrication facet error: a flat stone panel deviates from the curved surface by at most δ; refine $U \times V$ until $\delta \le$ tolerance.

6 Slab Trim I — convex blank (greedy half-plane)

Definition 6.1 (Polygon, signed area, reflex vertex)**.** A simple polygon $P = (v_0, \dots, v_{n-1})$ (indices mod n). Signed area $A(P) = \frac{1}{2}\sum_i (x_i y_{i+1} - x_{i+1} y_i)$; CCW iff $A > 0$. Vertex v_i is *reflex* (CCW) iff $(v_i - v_{i-1}) \times (v_{i+1} - v_i) < 0$.

Definition 6.2 (Half-plane clip, Sutherland–Hodgman)**.** For half-plane $H = \{x : \nu^\top(x-p) \ge 0\}$, $\mathrm{clip}(P,H)$ outputs, per edge (a,b): a if inside; and the segment–line intersection when the edge crosses ∂H.

Lemma 6.1 (Subset / material-removal)**.** *For any simple polygon P and half-plane H,* $\mathrm{clip}(P,H) = P \cap H \subseteq P$*, and* $\mathrm{area}(\mathrm{clip}(P,H)) \le \mathrm{area}(P)$*. Thus every clip only removes material (a valid wire cut).*

Proof. Sutherland–Hodgman against a single half-plane returns exactly the intersection of the (possibly non-convex) subject polygon with the convex region H; intersection with $\mathbb{R}^2 \supseteq H$ is a subset, and area is monotone under $\subseteq$. ☐

Definition 6.3 (Greedy convex trim)**.** While P has a reflex vertex and cuts $< K$: pick the deepest reflex v_i; cut along the supporting line of edge (v_{i-1}, v_i), keeping the half-plane containing the centroid; set $P \leftarrow \mathrm{clip}(P,H)$.

Theorem 6.2 (Convergence to a convex subset)**.** *The greedy convex trim produces a sequence $P = P_0 \supseteq P_1 \supseteq \cdots$ of polygons with $A(P_{t+1}) \le A(P_t)$ (Lemma 6.1); on termination (no reflex vertex, or budget K) the output is a convex polygon contained in P, hence the recovered blank is $P_{\mathrm{out}} = \bigcap_t H_t \cap P$ with* $\mathrm{area}(P_{\mathrm{out}})/\,\mathrm{area}(P)$ *the recovered-area ratio (yield).*

Proof. Each step clips by a half-plane through v_{i-1} chosen to exclude the wedge that makes v_i reflex, so the turning at the affected boundary becomes left (convex); the number of reflex vertices is monotone non-increasing and bounded below by 0. If it reaches 0 the polygon is convex (all cross products ≥ 0). The output is an intersection of half-planes with P, which is convex, and $\subseteq P$ by Lemma 6.1. (Greedy is a heuristic for *fewest* cuts; optimal-min-cut is §7.) ☐

7 Slab Trim II — concave kerf-follow (Imai–Iri min-#)

Definition 7.1 (ε-admissible chord)**.** For boundary vertices v_i, v_j $(i < j)$ the chord $\overline{v_i v_j}$ is ε-admissible iff every intermediate vertex is within ε of the chord: $\max_{i<k<j} \dfrac{|(v_k - v_i) \times (v_j - v_i)|}{\|v_j - v_i\|} \le \varepsilon$.

Definition 7.2 (Shortcut DAG and min-# approximation)**.** $G = (V,E)$ with $V = \{0, \dots, n-1\}$ and $(i,j) \in E$ iff $\overline{v_i v_j}$ is ε-admissible. The *min-# approximation* is a fewest-edges path $0 \to n-1$ in G (closed by the wrap edge), i.e. the polyline with the fewest straight kerfs staying within ε of the boundary.

Theorem 7.1 (Exactness of the DAG shortest path). *A breadth-first / fewest-edges shortest path in G yields the minimum number of ε-admissible straight segments approximating the chain; thus the kerf-follow cut count is* optimal, *not merely greedy. Complexity $O(n^2)$ time after $O(n^2)$ admissibility (improvable to $O(n^2)$ total / $O(n \log n)$ with cone tests).*

Proof. Any valid ε-approximation is a chain $0 = i_0 < i_1 < \cdots < i_m = n-1$ with each $\overline{v_{i_t} v_{i_{t+1}}}$ admissible, i.e. a path in G of m edges; conversely any G-path is a valid approximation. So valid approximations $\leftrightarrow$ G-paths, and "fewest segments" $\leftrightarrow$ "fewest edges". BFS computes a fewest-edges path exactly (every edge has unit weight), giving the minimum. (This is the classical Imai–Iri min-# result.) □

Corollary 7.2 (Yield ordering). *With ε small, kerf-follow recovers $\geq$ the convex trim's area (it keeps the concavities), at the cost of $\geq$ as many cuts; the two modes bracket the cut-count/yield trade-off.*

8 Concave-in-concave nesting

Definition 8.1 (No-fit polygon, valid placement). For a fixed sheet S and a part P under translation t (and rotation θ), the placement is valid iff $R_\theta P + t \subseteq S$ and $(R_\theta P + t)$ is interior-disjoint from all previously placed parts. The feasible t-set is the *inner-fit region* $\mathrm{IFR} = \{t : R_\theta P + t \subseteq S\}$ minus the union of no-fit polygons $\mathrm{NFP}(P_k, P)$ against placed parts.

Definition 8.2 (Raster heuristic). Rasterize S to a Boolean free-grid F at cell h; rasterize $R_\theta P$ to a mask M (dilated by the kerf clearance). A bottom-left placement at offset (o_x, o_y) is valid iff $M \wedge \neg F[\text{window}] = \mathbf{0}$; place greedily largest-area-first, marking $F \leftarrow F \wedge \neg M$.

Proposition 8.1 (Soundness and yield). *Every raster placement satisfies $R_\theta P + t \subseteq S$ and is interior-disjoint from placed parts up to one cell (h); the fill ratio $\left(\sum_{placed} \text{area}\right) / \text{area}(S)$ lower-bounds the achievable recovery and converges to the continuous NFP optimum as $h \to 0$. Concavity of both S and P is handled intrinsically (the masks encode arbitrary shape).*

9 2D nesting: no-fit polygon and bottom-left-fill

Definition 9.1 (Minkowski sum, no-fit polygon). For $A, B \subset \mathbb{R}^2$, $A \oplus B = \{a+b : a \in A, b \in B\}$ and $-B = \{-b : b \in B\}$. The *no-fit polygon* of a fixed A and a translating B is $\mathrm{NFP}(A, B) = A \oplus (-B)$.

Theorem 9.1 (NFP non-overlap criterion). *Let A, B be polygons with reference point of B at t. Then $\mathrm{int}(A) \cap \mathrm{int}(B + t) = \varnothing$ iff $t \notin \mathrm{int}(\mathrm{NFP}(A, B))$; and A and $B + t$ touch iff $t \in \partial\, \mathrm{NFP}(A, B)$.*

Proof. $(B + t) \cap A \neq \varnothing$ iff $\exists\, a \in A, b \in B$ with $a = b + t$, i.e. $t = a - b \in A \oplus (-B) = \mathrm{NFP}(A, B)$. Restricting to interiors, overlap of open interiors corresponds exactly to $t \in \mathrm{int}(\mathrm{NFP})$; boundary contact (touching, no overlap) to $t \in \partial\, \mathrm{NFP}$. The NFP is computed by Clipper2's polygon Minkowski sum (orbiting $-B$ around A). □

Definition 9.2 (Inner-fit polygon / region). $\mathrm{IFP}(S, B) = \{t : B + t \subseteq S\}$, the erosion of the sheet S by B. For sheet *holes* (defects) the admissible region is $\mathrm{IFP}(S, B) \setminus \bigcup_h \mathrm{NFP}(H_h, B)$ and, for a part placed *inside* a host part's hole, $\mathrm{IFP}(\text{hole}, B)$.

Definition 9.3 (Bottom-left-fill (BLF)). Place parts in a fixed order; each part is positioned at the lexicographically smallest (bottom-most, then left-most) t in $\mathrm{IFP}(S, B) \setminus \bigcup_{k\,\text{placed}} \mathrm{NFP}(P_k, B)$.

Proposition 9.2 (BLF validity)**.** *Every BLF placement satisfies $B + t \subseteq S$ (from IFP) and is interior-disjoint from all placed parts (from removing each* NFP*, Thm 9.1). Hence the layout is overlap-free by construction. (BLF is a heuristic for density; optimality is not claimed.)*

Hole-aware contact nesting (CNH). The candidate translation set is sampled at NFP *contact* configurations and at edge-alignment rotations (the contact-adaptive angle set); holes are nested host-first so smaller parts can seat inside larger parts' inner-fit regions.

10 3D packing: orientations and deepest-left-bottom-fill

Definition 10.1 (Axis-aligned orientation group)**.** The proper rotation group of the cube (chiral octahedral group O) has order 24. The 24-orientation set is $\{R \in SO(3) : R \text{ maps the standard cube to itself}\}$, enumerated as the permutation–sign matrices with $\det = +1$.

Lemma 10.1. $|O| = 24$. Proof: *a cube has* 6 *faces that can map to the bottom and* 4 *rotations fixing that face,* $6 \cdot 4 = 24$*; equivalently* $O \cong S_4$ *acting on the four body diagonals.*

Definition 10.2 (DLBF on a heightmap)**.** Maintain a heightmap $h : \text{grid} \to \mathbb{R}$. For each part (in non-increasing volume order) and each of the 24 orientations, the drop position is the grid cell minimizing, lexicographically, (resting height, y, x), where resting height = max of h over the part footprint; place it and raise h by the part's local thickness. Equivalent to deepest-bottom-left-fill in 3D.

Proposition 10.2 (Non-overlap + support)**.** *A DLBF placement rests on the current heightmap, so no two placed parts interpenetrate (each occupies cells above the prior surface), and the part is supported (its base touches h). Mesh-accurate variants replace the footprint-max by a CoACD collision test (§Decomposition, future batch).*

11 Guillotine separability (wire-saw cutting)

Definition 11.1 (Guillotine cut and partition)**.** A *guillotine cut* of an axis-aligned rectangle R is a single edge-to-edge axis-parallel line splitting R into two sub-rectangles. A packing of blocks in R is *guillotine-partitionable* if R can be recursively reduced to the individual blocks by guillotine cuts.

Theorem 11.1 (Wire-saw separability ⇔ guillotine)**.** *A block layout is separable by full-span (diamond-wire) cuts iff it is guillotine-partitionable; the staged three-stage guillotine (bed cuts → strips → cross cuts) produces such a partition by construction, so its separable fraction $\varphi = 1$.*

Proof. A diamond wire makes exactly full-span straight cuts = guillotine cuts, so a layout is wire-separable iff each freeing cut is edge-to-edge on the current piece, i.e. a guillotine sequence exists. The staged scheme emits, at each stage, parallel full-span cuts on the current sub-region (stage 1 along beds, stage 2 into strips, stage 3 cross), which is a recursive guillotine partition; hence every block is freed by full-span cuts and $\varphi = 1$. □

Definition 11.2 (BCSdbBV cost, Jalalian)**.** For a recovered set, the cutting-surface objective is $\mathrm{I}_{11} = \dfrac{A_{\text{cut}}}{\sum_b V_b \, \text{price}_b}$, the total sawn surface area per unit recovered block value; minimizing I_{11} minimizes saw work (time/energy/waste) per value. Frahan adds it as an optional score term (CutSurfaceWeight), off by default, so it never alters the baseline packing.

12 Masonry equilibrium (CRA limit analysis)

Definition 12.1 (Rigid-block assembly and contacts)**.** An assembly is a set of rigid blocks $\{B_1, \ldots, B_m\}$ with planar *contact interfaces*; each interface is discretized into contact points j with local frame $(\nu_j, \tau_j^1, \tau_j^2)$ (ν_j the inward unit normal). The contact force at j is $f_j = f_j^n \nu_j + \sum_a f_j^a \tau_j^a \in \mathbb{R}^3$, f_j^n the normal (compression) magnitude.

Definition 12.2 (Equilibrium operator)**.** Stacking all contact forces into f and the per-block dead loads (weights) into g, force/moment balance of every block is the linear system $Af = g$, where A is the equilibrium (statics) matrix assembled from the contact positions and frames.

Definition 12.3 (Admissible cone: no-tension + Coulomb friction)**.** A contact force is admissible iff $f_j^n \geq 0$ (unilateral: contacts push, never pull) and $\|(f_j^1, f_j^2)\| \leq \mu\, f_j^n$ (Coulomb friction cone, coefficient μ). The admissible set $\mathcal{K} = \{f : f_j^n \geq 0,\ \|(f_j^1, f_j^2)\| \leq \mu f_j^n\ \forall j\}$ is a product of second-order (ice-cream) cones, hence closed and convex.

Theorem 12.1 (Static (safe) theorem of limit analysis)**.** *The assembly is stable under load g iff there exists $f \in \mathcal{K}$ with $Af = g$. Equivalently, stability is feasibility of the convex (second-order cone) program* find *$f \in \mathcal{K}$ s.t. $Af = g$.*

Proof. This is the lower-bound (static) theorem of plastic limit analysis specialized to rigid no-tension blocks with Coulomb interfaces (Heyman; Kao 2022 CRA): any statically admissible self-stress field in equilibrium with the loads and respecting the yield set ($\mathcal{K}$) certifies that the structure can carry the load (safe). Conversely, if no admissible equilibrium force exists, by convex (Gale/Farkas) duality there is a compatible collapse mechanism with positive external work, so the assembly moves — unstable. Convexity of $\mathcal{K}$ and linearity of A make this an SOCP feasibility test (CRA solves it; a polyhedral cone linearization reduces it to an LP). □

Definition 12.4 (Collapse load factor λ)**.** Split $g = g_{\text{dead}} + \lambda\, g_{\text{live}}$. The collapse multiplier is $\lambda^\star = \sup\{\lambda : \exists f \in \mathcal{K},\ Af = g_{\text{dead}} + \lambda g_{\text{live}}\}$; the design is safe iff $\lambda^\star \geq 1$. ($\lambda^\star$ is the optimum of a convex program; Frahan's "Lambda" stability index lives in this family and is used to compare assemblies, e.g. ETH1100 vs Cyclopean.)

Proposition 12.2 (COM-over-support gate)**.** *For a single rigid block resting on coplanar frictionless supports, the frictionless, single-interface case of Thm 12.1 reduces to: the block does not topple iff the vertical projection of its centre of mass lies in the convex hull of the support contact points.*

Proof. Frictionless ⇒ each $f_j = f_j^n \nu_j$ with $f_j^n \geq 0$ and ν_j vertical. Force balance fixes $\sum f_j^n = W$; moment balance about the horizontal axes requires the weighted average of support points (weights $f_j^n \geq 0$, summing to W) to equal the COM projection. A convex combination of the support points realises a given point iff that point is in their convex hull. Hence feasibility ⇔ COM projection $\in$ conv(supports). □

Definition 12.5 (Ashlar / coursed bond; Cyclopean carve)**.** A *coursed ashlar* layout stacks blocks in horizontal courses with head joints of consecutive courses offset by at least an overlap fraction o (running bond: no two vertically adjacent head joints collinear), which the layout engine enforces as a constraint maximizing interlock. *Cyclopean Cannibalism* is the overlap-then-carve operation: place an oversized boulder, then Boolean-subtract the already-placed neighbours (carve-back) so faces mate exactly while preserving the rubble texture.

13 Decomposition

13.1 Approximate convex decomposition (CoACD)

Definition 13.1 (Concavity)**.** For a compact $S \subset \mathbb{R}^d$ with convex hull $\operatorname{conv}(S)$, the concavity is $\eta(S) = \sup_{x \in \operatorname{conv}(S)} \operatorname{dist}(x, S)$ (the one-sided Hausdorff distance from $\operatorname{conv}(S)$ to S). $\eta(S) = 0$ iff S is convex.

Definition 13.2 (ε-approximate convex decomposition)**.** A partition $\{S_k\}$ of S (union $= S$, interiors disjoint) is an ε-ACD if $\eta(S_k) \le \varepsilon$ for all k. The decision "$\exists$ ε-ACD with $\le K$ parts" is NP-hard in 3D; CoACD is a heuristic: a Monte-Carlo tree search over straight cut planes, scored by a *collision-aware* concavity proxy (it cuts the mesh and measures the post-cut concavity), recursing until every part has $\eta \le \varepsilon$.

Proposition 13.1 (Termination + soundness)**.** *Each CoACD cut strictly partitions a part with $\eta > \varepsilon$; since a finite triangle mesh admits only finitely many distinct cut-induced sub-meshes and η is non-increasing under cutting, the recursion terminates with all parts $\eta \le \varepsilon$. The parts tile S (union/disjoint-interior) by construction.*

13.2 Convex polygon partition

Theorem 13.2 (Hertel–Mehlhorn)**.** *Triangulate a simple polygon P (n vertices), then delete every* inessential *diagonal (one whose removal keeps both incident pieces convex). The result is a partition of P into convex pieces whose count is at most* 4 *times the minimum, computed in $O(n)$ time after the triangulation.*

Proof sketch. A diagonal is essential only if it resolves a reflex vertex; each reflex vertex needs ≥ 1 resolving edge, giving a lower bound, while Hertel–Mehlhorn uses ≤ 2 essential diagonals per reflex vertex, yielding the 4-approximation (a reflex vertex is shared by ≤ 2 kept diagonals). Deleting inessential diagonals never creates a reflex angle, so all pieces stay convex. □

The exact minimum convex partition without Steiner points is solvable in $O(n^3)$ by dynamic programming (Chazelle–Dobkin).

13.3 Convex polyhedron half-space clipping

Lemma 13.3 (3D clip)**.** *For a convex polyhedron $Q \subset \mathbb{R}^3$ and half-space H, $Q \cap H$ is a convex polyhedron with $Q \cap H \subseteq Q$; it is computed by clipping each face polygon against H (planar Sutherland–Hodgman) and capping the new cut face. This is the 3D analogue of Lemma 6.1 and the managed half-space-clipping decomposer.*

Proof. Intersection of convex sets is convex; $H \supseteq Q \cap H$ gives the subset. The boundary of $Q \cap H$ consists of clipped original faces plus one new face on ∂H (the section), each a convex polygon. □

13.4 Delaunay and α-shapes

Definition 13.3 (Delaunay complex)**.** $\mathrm{DT}(P)$ is the simplicial complex on $P \subset \mathbb{R}^d$ in which a simplex belongs iff it has an *empty* circumsphere (no point of P strictly inside). For points in general position it is the unique triangulation maximizing the minimum angle (2D) / dual to the Voronoi diagram.

Definition 13.4 (α-complex). For $\alpha \geq 0$, the α-complex $\mathcal{C}_\alpha \subseteq \mathrm{DT}(P)$ keeps each simplex whose smallest empty circumsphere has radius $\leq \alpha$ (with the standard "α-exposed" condition). The α-shape is $|\mathcal{C}_\alpha|$.

Theorem 13.4 (α interpolates points $\to$ hull). *$\mathcal{C}_0 = P$ (vertices only) and for $\alpha \geq \alpha_{\max}$, $|\mathcal{C}_\alpha| = \mathrm{conv}(P)$ with $\mathcal{C}_\alpha = \mathrm{DT}(P)$; the family $(\mathcal{C}_\alpha)_{\alpha \geq 0}$ is a nested filtration that grows from the point cloud to the convex hull. Constrained Delaunay tetrahedralization adds required facets to* DT *while keeping the empty-circumsphere property on the unconstrained simplices.*

13.5 Density-watershed partition

Definition 13.5 (Watershed blocks (BlockCutOpt I5)). Given a scalar fracture-density field f on a grid, define $g = -f$ and assign each cell to the catchment basin of the local maximum of g (= local minimum of f) reached by steepest ascent; the basins (separated by watershed ridges along high-density fractures) partition the volume into candidate intact blocks.

Proposition 13.5. *The catchment-basin map is a partition: every non-ridge cell flows to exactly one maximum, and basins are interior-disjoint and cover the domain up to the measure-zero ridge set. Low fracture density $\Rightarrow$ large intact basins $\Rightarrow$ larger recoverable blocks.*

14 Edge-matching and assembly

14.1 Closed-form absolute orientation (Horn 1987)

Definition 14.1 (Absolute orientation). Given corresponded point sets $\{p_i\}, \{q_i\}_{i=1}^n \subset \mathbb{R}^3$, find $s > 0$, $R \in SO(3)$, $t \in \mathbb{R}^3$ minimizing $E(s, R, t) = \sum_i \|q_i - (sRp_i + t)\|^2$.

Theorem 14.1 (Horn quaternion solution). *Let $\bar{p}, \bar{q}$ be the centroids, $p'_i = p_i - \bar{p}$, $q'_i = q_i - \bar{q}$, and $M = \sum_i p'_i q'^{\top}_i$. Then the optimal rotation is $R = R(\hat{e})$ where $\hat{e}$ is the unit eigenvector of the largest eigenvalue of the symmetric 4×4 matrix $N(M)$ (built from the entries of M); the optimal scale is $s = \left(\sum_i \|q'_i\|^2 / \sum_i \|p'_i\|^2\right)^{1/2}$ (symmetric form) and $t = \bar{q} - sR\bar{p}$.*

Proof sketch. Minimizing over t for fixed s, R gives $t = \bar{q} - sR\bar{p}$ (centroid alignment), reducing E to the centred residual. The rotation part maximizes $\sum_i q'^{\top}_i R p'_i = \mathrm{tr}(RM^\top)$. Writing R via a unit quaternion $\dot{e}$, this trace equals $\dot{e}^\top N(M) \dot{e}$ with $N(M)$ the standard symmetric 4×4 form; maximizing a quadratic form over the unit sphere is attained at the top eigenvector (Rayleigh). Scale follows by setting $\partial E / \partial s = 0$. (Equivalent SVD form: $R = V \operatorname{diag}(1, 1, \det VU^\top) U^\top$ from $M = U \Sigma V^\top$.) □

14.2 Iterative closest point (ICP)

Definition 14.2. Given a source $\{p_i\}$ and target set Y, ICP alternates: (C) $c(i) = \arg\min_{y \in Y} \|Tp_i - y\|$ (nearest correspondence under the current transform T); (T) $T \leftarrow \arg\min_T \sum_i \|Tp_i - c(i)\|^2$ (Horn, Thm 14.1). *Constrained* ICP adds non-penetration/contact constraints to step (T).

Theorem 14.2 (Monotone convergence). *The registration error $E_t = \sum_i \|T_t p_i - c_t(i)\|^2$ is non-increasing in t and converges; ICP reaches a fixed point (local minimum) in finitely many correspondence sets.*

Proof. Step (C) minimizes E over correspondences with T fixed (each $c(i)$ is the nearest point), so it cannot increase E; step (T) minimizes over T with correspondences fixed, so it cannot increase E. Thus E_t is monotone non-increasing and bounded below by 0, hence convergent; since correspondences come from a finite set and E strictly decreases off a fixed point, ICP terminates at a local minimum. □

14.3 Global assembly: matching, beam, spanning tree

Definition 14.3 (Bipartite assignment)**.** Given a cost matrix $C \in \mathbb{R}^{n\times n}$, the assignment problem is $\min_{\pi\in S_n} \sum_i C_{i,\pi(i)}$, solved exactly by the Hungarian algorithm in $O(n^3)$ (it returns a minimum-cost perfect matching by maintaining a dual feasible potential and augmenting along tight edges).

Definition 14.4 (Pair-graph assembly)**.** Build a weighted graph G on fragments with edge weight the pairwise match score (boundary-rail edge affinity). *Agglomerative* assembly takes a maximum-weight spanning tree of G and composes the relative poses along it; *beam search* keeps the top-B partial assemblies by cumulative score, expanding each by one compatible fragment per step. Beam with $B = 1$ is the frame-anchored greedy; $B \to \infty$ is exact.

Proposition 14.3 (Spanning-tree consistency)**.** *A spanning tree induces a unique relative pose between any two fragments (compose along the tree path); cycles are avoided, so the assembly is pose-consistent by construction (no contradictory loop closures), at the cost of not averaging redundant matches.*

15 Surface reconstruction and flattening

15.1 Poisson surface reconstruction (Kazhdan)

Definition 15.1 (Indicator and oriented-normal field)**.** Let $M = \partial\Omega$ be the surface of a solid Ω with indicator χ_Ω. The oriented-point samples define a vector field V approximating the surface normal field as a measure: $\nabla\chi_\Omega = \vec{n}\,\delta_M$ in the distributional sense.

Theorem 15.1 (Poisson formulation)**.** *The indicator χ best matching V in the least-squares sense solves the Poisson equation $\Delta\chi = \nabla \cdot V$, and M is recovered as a level set $\{\chi = \bar{\chi}\}$ (isovalue = mean of χ at the samples).*

Sketch. Minimizing $\int \|\nabla\chi - V\|^2$ over χ has Euler–Lagrange equation $\Delta\chi = \nabla\cdot V$ (set the variational derivative to zero, integrate by parts). Since $\nabla\chi_\Omega = \vec{n}\,\delta_M$ and V approximates that field, $\chi \approx \chi_\Omega$, whose level set at the surface is M. (Discretized on an adaptive octree; the right-hand side is the divergence of the smoothed normal field.) □

15.2 Centroidal Voronoi tessellation (Lloyd)

Definition 15.2 (CVT energy)**.** For generators $z = (z_1, \ldots, z_k)$ with Voronoi cells $V_i(z)$ and density ρ, the energy is $F(z) = \sum_i \int_{V_i(z)} \rho(x)\|x - z_i\|^2\,dx$. A tessellation is *centroidal* if each z_i equals its mass centroid $c_i = \big(\int_{V_i} \rho x\big)/\big(\int_{V_i} \rho\big)$.

Theorem 15.2 (Lloyd descent)**.** *Lloyd's iteration (recompute $V_i(z)$; set $z_i \leftarrow c_i$) does not increase F and its fixed points are exactly the centroidal Voronoi tessellations.*

Proof. With cells fixed, $\sum_i \int_{V_i} \rho \|x - z_i\|^2$ is minimized over z_i at the centroid c_i (the mean minimizes squared deviation), so the move step lowers F; re-Voronoi with generators fixed assigns each x to its nearest z_i, lowering each integrand pointwise, so it also lowers F. $F \geq 0$ bounded below $\Rightarrow$ convergence; a fixed point has $z_i = c_i \, \forall i$, i.e. a CVT. (Used for Trencadis seeding and CVD remeshing.) $\square$

15.3 Barycentric mapping and conformal flattening

Definition 15.3 (Barycentric map)**.** For a triangle $T = (a, b, c)$, the map $\varphi_T(\lambda) = \lambda_1 a + \lambda_2 b + \lambda_3 c$ on the standard simplex $\{\lambda \geq 0, \sum \lambda = 1\}$ is the unique affine bijection from the reference triangle to T.

Proposition 15.3 (2D→3D transfer)**.** *A point with barycentric coordinates λ in a parameter triangle maps to $\varphi_{T_{3D}}(\lambda)$ in the corresponding 3D triangle. Barycentric coordinates are affine invariants, so the transfer is independent of the embedding and exact on triangles.*

Definition 15.4 (BFF flattening)**.** Boundary First Flattening computes a discrete *conformal* map of a disk-topology surface to the plane by prescribing either boundary lengths or boundary geodesic curvatures and solving the resulting Cherrier/Yamabe boundary-value problem (a sparse linear / Poisson solve for the log conformal factor u, then a boundary integration). Conformality $\Rightarrow$ angles preserved, area distortion encoded by e^{2u} (the chart-scale, recovered for the surface packers).

16 Discontinuities, joint sets, and block theory

16.1 Baecher stochastic DFN

Definition 16.1. A Baecher discrete fracture network draws fracture centres from a homogeneous Poisson point process of intensity λ in the rock volume; each fracture is a disk of radius $r \sim F_R$ (size law) with unit normal $\nu \sim$ Fisher on the sphere, density $f(\nu) \propto \exp(\kappa \, \mu^\top \nu)$ (μ mean pole, κ concentration). Centres, sizes, orientations are independent.

16.2 Watson axial mean-shift for joint sets

Definition 16.2. Fracture poles are *axial* ($\nu \equiv -\nu$). The Watson distribution on the projective sphere has density $f(x) \propto \exp(\kappa (\mu^\top x)^2)$ (bipolar for $\kappa > 0$). Joint sets are the modes of the pole density, found by *mean shift* on $\mathbb{R}P^2$: iterate $x \leftarrow \text{normalize}\big(\sum_i w_i(x) \, \text{sign}(x^\top \nu_i) \, \nu_i\big)$ with kernel weights $w_i(x) = g(\, 1 - (x^\top \nu_i)^2 \,)$.

Proposition 16.1. *Each mean-shift step ascends the kernel pole-density estimate; fixed points are the density modes, partitioning the poles into joint sets (basins of attraction). Antipodal symmetry is handled by the squared inner product, so ν and $-\nu$ vote together.*

16.3 Block theory: removability (Goodman–Shi)

Definition 16.3 (Joint, excavation, block pyramids)**.** For a convex block cut by joint planes through a common apex, the *joint pyramid* JP is the cone $\{d : \; d$ on the block side of every joint plane$\}$; the *excavation pyramid* EP is the cone of directions into the free (excavated) space; the *block pyramid* is $BP = JP \cap EP$.

Theorem 16.2 (Finiteness and removability)**.** *A block is* finite *iff $BP = JP \cap EP = \varnothing$. A finite block is* removable *iff $JP \neq \varnothing$; if instead $JP = \varnothing$ it is* tapered *(cannot be removed without cutting). Hence*

$$\text{removable} \iff JP \neq \varnothing \ \wedge \ JP \cap EP = \varnothing.$$

A removable block is a key block *iff the active resultant force lies in JP (so it slides/falls under gravity/water/seismic load).*

Sketch. Shi's theorem: the block translated to infinity along a direction d stays in the rock iff $d \in JP$ and in the free space iff $d \in EP$; an unbounded motion (removal without cutting) requires a common direction, i.e. $d \in BP$. Thus $BP = \varnothing \Rightarrow$ no escape to infinity $\Rightarrow$ finite, and a finite block can be moved (removable) iff its shape cone JP is non-empty (some sliding/lifting direction exists once the excavation face is open), while $JP = \varnothing$ forces a tapered, non-removable wedge. Convex-cone intersection tests make this a linear-feasibility check per candidate block. □

16.4 Block size from joint spacing (Palmstrom)

Definition 16.4. With volumetric joint count $J_v = \sum_i 1/s_i$ (s_i mean spacing of joint set i), the typical block volume is $V_b = \beta\, J_v^{-3}$ (β a block-shape factor). CrackGraph → BlockGraph partitions the fractured solid into the cells of the joint arrangement, whose volumes realize this statistic.

17 GPR signal processing (pre-kriging)

17.1 f–k (Stolt) migration

Definition 17.1 (Exploding-reflector section)**.** A zero-offset radargram $P(x,t)$ is modeled as an exploding-reflector field with effective velocity $c = v/2$ (half the EM velocity), so reflectors radiate at $t = 0$.

Theorem 17.1 (Stolt mapping, constant v)**.** *Let $\hat{P}(k_x,\omega)$ be the 2D Fourier transform of P. The migrated image $M(x,z)$ has transform $\hat{M}(k_x,k_z) = \dfrac{c\,k_z}{\sqrt{k_x^2+k_z^2}}\,\hat{P}\big(k_x,\ \omega = c\sqrt{k_x^2+k_z^2}\big)$, i.e. a remap from ω to $k_z = \sqrt{(\omega/c)^2 - k_x^2}$ with the indicated amplitude Jacobian, followed by an inverse 2D FFT.*

Sketch. The scalar wave equation gives the dispersion relation $\omega^2 = c^2(k_x^2+k_z^2)$ for upgoing waves; imaging at $t = 0$ evaluates the field at depth z. Changing the integration variable from ω to k_z along $\omega = c\sqrt{k_x^2+k_z^2}$ introduces the Jacobian $d\omega/dk_z = c\,k_z/\sqrt{k_x^2+k_z^2}$; the inverse FFT then yields $M(x,z)$. This is exact for constant c (Stolt 1978). □

17.2 Analytic signal and instantaneous energy (Taner)

Definition 17.2 (Hilbert transform, analytic signal)**.** The Hilbert transform $\mathcal{H}$ has frequency multiplier $-i\,\mathrm{sgn}(\omega)$. The analytic signal of a trace s is $s_a = s + i\,\mathcal{H}s$; the instantaneous amplitude is $A(t) = |s_a(t)|$ and the instantaneous energy $A(t)^2$ (Taner 1979).

Proposition 17.2. *$A(t)$ is the envelope of s: $A \geq |s|$ with equality at the local extrema, and A is phase-independent (invariant under a constant phase rotation $s_a \mapsto e^{i\phi}s_a$). Hence energy peaks localize reflectors robustly to wavelet phase.*

17.3 Picking, continuity gate, depth, uncertainty

Definition 17.3 (Energy-quantile pick + USGS continuity)**.** A sample is a fracture candidate iff its instantaneous energy exceeds the Q-quantile of the trace energy. A candidate reflector is *kept* iff it is laterally continuous over at least N consecutive traces (USGS $\geq N$-trace gate; N stone-calibrated, e.g. 27 marble vs 40 granite).

Proposition 17.3 (Depth and its uncertainty)**.** *Two-way time t maps to depth $d = v\,t/2$. With $v = c_0/\sqrt{\varepsilon_r}$, error propagation gives the dominant deep-fracture deviation*

$$\frac{\sigma_d}{d} = \frac{\sigma_v}{v} = \tfrac{1}{2}\,\frac{\sigma_{\varepsilon_r}}{\varepsilon_r},$$

plus a $\lambda/4$ vertical-resolution floor and a time-zero term; these combine (root-sum-square) into the per-pick σ consumed by the kriging confidence metric (§1, FractureUncertainty).

18 Quarry yield and cutting optimization

18.1 Catalogue guillotine tiling (cost/volume)

Definition 18.1 (Priced catalogue, layer)**.** A catalogue $\mathcal{B}$ assigns each block size b a footprint and a unit price/value. A bed-bounded layer is a rectangle R; a guillotine tiling places non-overlapping catalogue blocks separable by edge-to-edge cuts. The objective with volume weight W is max $\sum_b \mathrm{vol}_b(\mathrm{price}_b + W) \;-\; \kappa\, A_{\mathrm{cut}}$.

Theorem 18.1 (DP optimality among guillotine tilings)**.** *Let $V(R)$ be the optimum objective over guillotine tilings of R. Then*

$$V(R) = \max\Big(\max_{b\subseteq R} \mathrm{val}(b) + V(R\setminus b),\ \max_{\mathit{cut}\ \ell} V(R_1) + V(R_2)\Big),$$

over single-block placements and all horizontal/vertical cut positions ℓ splitting $R = R_1 \cup R_2$. The recursion (memoized over the $O(n^2)$ distinct sub-rectangles induced by the catalogue/grid coordinates) computes the exact guillotine optimum.

Proof. A guillotine tiling of R is, by definition, either a single block or a first edge-to-edge cut into R_1, R_2 each guillotine-tiled; the value is additive over the partition; taking the max over the (finite) cut positions and placements gives Bellman optimality. Restricting cuts to catalogue/grid coordinates loses no optimum because shifting a cut to the nearest such coordinate does not reduce the achievable value. □

RecoveryCascade applies V at decreasing scales (coarse blocks first, then fill); at a single scale it reduces to Thm 18.1, and the multi-scale recovery is monotone (finer passes only add blocks in the residual region).

18.2 Pareto omni-solve (BlockCutOpt)

Definition 18.2 (Dominance, Pareto front)**.** For objective vector $\phi = (\mathrm{recovery}, \mathrm{revenue}, -\mathrm{risk}, -A_{\mathrm{cut}})$ (all maximized), x *dominates* y iff $\phi(x) \geq \phi(y)$ componentwise with strict inequality somewhere. The Pareto front is the set of non-dominated solutions. BlockCutOpt omni-solve returns this front; the single-objective solvers (brute force, density-watershed I5, Fisher-robust) are scalarizations / special cases.

Proposition 18.2. *The recovery-optimal and revenue-optimal plans both lie on the Pareto front but generally differ (§1 finding); exposing the front lets the operator pick the trade-off rather than a fixed scalarization.*

18.3 Cut-and-fill volume (TIN prisms)

Proposition 18.3 (Prism differencing)**.** *For two TINs (top, bottom) over a common triangulation, the volume between them is* $V = \sum_T \frac{A_T}{3}\big(\delta_1 + \delta_2 + \delta_3\big)$*, where* A_T *is the planimetric area of triangle* T *and* δ_i *the top-minus-bottom height at its vertices (truncated-prism / prismoidal rule); signed* δ *separates cut (*$\delta < 0$*) from fill (*$\delta > 0$*).*

Proof. Over a triangle the height difference $\delta(x, y)$ is affine (linear TINs), so its mean over T equals the vertex average $(\delta_1 + \delta_2 + \delta_3)/3$; the prism volume is mean height times base area A_T. Sum over triangles. □

19 Robotic fabrication: toolpaths and frames

19.1 Tool-axis frame along a cut path

Definition 19.1 (Path frame)**.** For a cut path $\gamma(s)$ (arc length) with unit tangent $T = \gamma'$, a tool frame is an orthonormal (T, U, V); the tool axis is one of U, V (or the cut-plane normal). A *rotation-minimizing* (Bishop) frame parallel-transports U, V: $U' = -(U \cdot T')\,T$, $V' = -(V \cdot T')\,T$.

Proposition 19.1 (RMF minimizes spin)**.** *The Bishop frame has zero angular velocity component along* T *(*$\Omega \cdot T = 0$*), hence among all adapted frames it minimizes total tool roll* $\int |\Omega \cdot T|\,ds = 0$*; it exists and is unique given an initial* $(U_0, V_0) \perp T(0)$*. (Avoids the Frenet frame's flips at inflections, giving smooth robot wrist motion.)*

Proof. $U' = -(U \cdot T')T$ keeps $U \perp T$ ($\frac{d}{ds}(U \cdot T) = U' \cdot T + U \cdot T' = -(U \cdot T') + U \cdot T' = 0$) and $\|U\| = 1$ ($U \cdot U' = 0$). Its angular velocity is $\Omega = T \times T'$, which has no T-component, so the roll rate $\Omega \cdot T = 0$. Uniqueness/existence is the linear ODE with given initial frame. □

19.2 Arc discretization (Arc Step) and motion mapping

Proposition 19.2 (Chord-error bound)**.** *An arc of radius* r *and sweep* θ *split into* n *equal chords (each subtending* $\alpha = \theta/n$*) has maximum deviation (sagitta)* $h = r\big(1 - \cos(\alpha/2)\big)$*. To guarantee* $h \le \varepsilon$ *take* $\alpha \le 2\arccos(1 - \varepsilon/r)$*, i.e.* $n = \big\lceil \theta/\big(2\arccos(1 - \varepsilon/r)\big)\big\rceil$*.*

Proof. The chord's farthest point from the arc is its midpoint; the perpendicular distance from the arc to the chord midpoint is $r - r\cos(\alpha/2)$ (radius minus the apothem). Bounding it by ε and solving for α gives the stated n. □

Definition 19.2 (CutSegment → motion morphism)**.** Map each cut segment kind to a robot motion primitive: line→Lin, arc→Circ (via three points), rapid→Ptp; the tool pose at each via point is $(\gamma(s_i), \mathrm{frame}(s_i))$. The map is a structure-preserving morphism: the executed TCP path equals the geometric cut path up to the discretized chord tolerance above (KUKA prc / Robots back-ends).

20 Rigid-body dynamics and settling

20.1 Newton–Euler equations of motion

Definition 20.1. A rigid body with mass m, inertia I, position x, orientation R, linear/angular velocity v, ω evolves by $\dot{x} = v$, $\dot{R} = [\omega]_\times R$, and the Newton–Euler equations $m\dot{v} = F$, $I\dot{\omega} + \omega \times I\omega = \tau$, with $F = mg + \sum_c f_c$, $\tau = \sum_c r_c \times f_c$ (f_c the contact forces, r_c the contact arms).

20.2 Contact as a complementarity (Signorini) problem

Definition 20.2 (Signorini–Coulomb contact)**.** At each contact let ϕ be the signed gap and λ the normal impulse. Non-penetration and unilateral force are the complementarity condition $0 \le \phi \perp \lambda \ge 0$ (with the Coulomb friction cone, §6). Per step this is a (linear) complementarity problem; the sequential-impulse / projected Gauss–Seidel solver iterates each contact to satisfy it.

Proposition 20.1 (Collision via convex decomposition)**.** *For non-convex bodies decomposed into convex pieces (§Decomposition, CoACD), interpenetration holds iff some pair of pieces overlaps; each convex pair is tested by GJK (distance) / EPA (penetration). Thus narrow-phase collision reduces to finitely many convex queries, each solvable by Minkowski-difference support mapping.*

20.3 Drop-settle and static rest

Theorem 20.2 (Rest = constrained energy minimum = contact equilibrium)**.** *A dropped block is at static rest iff its pose is a KKT point of* $\min_q \; U(q) = mg^\top c(q)$ *subject to the non-penetration constraints* $\phi_j(q) \ge 0$*, where* $c(q)$ *is the centre of mass. The KKT stationarity* $\nabla U = \sum_j \lambda_j \nabla \phi_j$, $\lambda_j \ge 0$, $\lambda_j \phi_j = 0$ *is exactly the contact-force equilibrium of §6 (Thm 12.1); hence drop-settle (gradient descent of* U *under the unilateral constraints, z-up) converges to a locally stable configuration, and the resulting support reaction set certifies stability via the COM-over-support gate.*

Sketch. At rest, velocities and accelerations vanish, so Newton–Euler reduces to force/torque balance under gravity and contact reactions $\lambda_j \nabla \phi_j$ ($\lambda_j \ge 0$ since contacts push) with $\lambda_j \phi_j = 0$ (no force at open gaps). These are precisely the stationarity + complementarity KKT conditions of the gravitational-PE minimization under $\phi_j \ge 0$. Gradient flow of U projected onto the feasible set decreases U monotonically and is bounded below (the body rests on the support), so it converges to such a KKT point. Concave-aware settling uses the convex-piece contacts for ϕ_j. □

21 Learned reassembly (Kintsugi / PotNet)

21.1 Normalization and pose composition

Definition 21.1 (Normalization)**.** Each fragment f is mapped to a canonical frame by a similarity $T_{\text{norm}}(f) \in \text{Sim}(3)$ (translate to centroid, scale to unit extent); its inverse is $T_{\text{unnorm}}(f) = T_{\text{norm}}(f)^{-1}$. The network $T_{\text{net}}(f) \in SE(3)$ predicts the placement of the *normalized* fragment into the normalized assembly frame anchored at fragment 0.

Theorem 21.1 (World-pose composition)**.** *The world pose of fragment* f *is*

$$T_{\text{world}}(f) = T_{\text{unnorm}}(0) \cdot T_{\text{net}}(f) \cdot T_{\text{norm}}(f),$$

and this is the unique transform consistent with: normalize f*, apply the predicted placement, then undo the anchor (fragment* 0*) normalization. It is invariant to the global pose of the raw input (the network sees only normalized geometry).*

Proof. For a point p on fragment f: normalization gives $T_{\text{norm}}(f)p$; the network places it in the normalized assembly frame as $T_{\text{net}}(f)T_{\text{norm}}(f)p$; mapping the normalized assembly frame back to world is the inverse of the anchor normalization, $T_{\text{unnorm}}(0)$. Composing (associativity of $SE(3)/\text{Sim}(3)$ action) yields $T_{\text{world}}(f)p$. If the raw input is pre-transformed by any G, then T_{norm} absorbs G (centroid/extent are equivariant), so T_{net} and the final world relation are unchanged — global-pose invariance. (PotNet regresses one such T_{net} per object; failure modes separate into pose-composition vs network-drift vs mesh-style by inspecting the verifier-score distribution.) □

21.2 Verifier-score gating

Definition 21.2. A learned verifier g : assembly $\to [0,1]$ scores a candidate; it is accepted iff $g \geq \tau$ ($\tau = 0.5$). The accepted set is $\{f : g(f) \geq \tau\}$, a thresholding decision whose precision/recall trade off monotonically as τ varies (higher $\tau \Rightarrow$ higher precision, lower recall).

22 Surface packing on charts

22.1 Atlas and pack-then-map

Definition 22.1 (Face-corner UV chart)**.** A chart assigns each face corner a UV coordinate; the chart is a valid atlas iff the UV assignment is consistent across every shared edge (the two incident faces agree up to the transition map). For a BFF flattening (§15) the chart is a single global UV map of the disk-topology surface.

Theorem 22.1 (Surface-packing transfer)**.** *Let parts be packed without overlap in the UV chart Ω (the 2D nesting of §2). Mapping each placed part back by the chart map $\Phi = \varphi \circ (chart)^{-1}$ yields surface-embedded parts that are mutually non-overlapping on the surface, and whose physical area equals the UV area times the local chart-scale e^{2u} (the BFF conformal factor); pre-scaling each UV cell by e^{-u} makes the on-surface tile sizes match the target.*

Proof. Φ is a bijection on the chart (atlas validity), so disjoint UV placements map to disjoint surface placements (§2 non-overlap is preserved by a bijection). For a conformal chart the first fundamental form is $e^{2u}\,\text{Id}$, so an area element dA_{UV} maps to $e^{2u}dA_{\text{UV}}$ on the surface; hence multiplying UV sizes by e^{-u} cancels the distortion (barycentric transfer, §15, is exact per triangle). □

22.2 Orientation field (GVF)

Definition 22.2 (Gradient vector flow field)**.** Given an initial tile-orientation field v_0 (e.g. from surface curvature / edges), the smoothed field v minimizes the Xu–Prince energy $E(v) = \int_\Omega \mu \, \|\nabla v\|^2 + \|v_0\|^2 \, \|v - v_0\|^2 \, dA$.

Proposition 22.2 (Well-posed orientation)**.** *E is a strictly convex quadratic functional in v (for $\mu > 0$), so it has a unique minimizer solving the linear diffusion–reaction Euler–Lagrange equation $\mu \, \Delta v - \|v_0\|^2(v - v_0) = 0$; the result is a smooth orientation field that follows v_0 where it is strong and diffuses smoothly elsewhere, used to align Trencadis mosaic tiles. CVD-Lloyd seeding (§15) places the tile centres.*

Coverage. The remaining `[Algorithm]` attributes that are genuine mathematical methods (not format I/O or dispatch) are derived in §20–§27 below; see `AUDIT_coverage.md` for the title-by-title mapping of all 249 distinct algorithms.

23 Spectral and PCA geometry

Definition 23.1 (PCA)**.** For points $\{p_i\}$ with centroid $\bar{p}$, the covariance is $C = \frac{1}{n}\sum_i (p_i - \bar{p})(p_i - \bar{p})^\top$, with eigenpairs $(\lambda_1 \geq \lambda_2 \geq \lambda_3,\ u_1, u_2, u_3)$. The u_k are the principal axes.

Proposition 23.1 (PCA-OBB and normal estimation)**.** *(i) The PCA oriented bounding box is the axis-aligned box of $\{p_i\}$ in the eigenbasis $U = [u_1 u_2 u_3]$; it is a heuristic (not minimum-volume, which needs rotating calipers) but tight when one axis dominates. (ii) The surface normal at a point is u_3 (least-eigenvalue direction) of the local-neighbourhood covariance (cf. §5). (iii) Hoppe normal orientation propagates a consistent sign along the Euclidean minimum spanning tree, maximizing $\sum_{(i,j)\in MST} \operatorname{sign}(u_3^{(i)} \cdot u_3^{(j)})$.*

Proof. $C \succeq 0$ symmetric, so the spectral theorem gives the orthonormal u_k; the box in coordinates $U^\top p$ is the OBB. The least-variance direction minimizes $u^\top C u$ over $\|u\| = 1$ (Rayleigh) and is normal to the best-fit plane (§5). MST propagation is the greedy consistent-sign assignment minimizing flips on a tree (unique path between any two nodes ⇒ no conflict). □

24 Convex hull and inscribed polygons

Proposition 24.1 (QuickHull)**.** *QuickHull computes* $\operatorname{conv}(S)$ *by recursively assigning points to the outside of the current hull edge and promoting the farthest; it is exact, expected $O(n \log n)$, worst case $O(n^2)$ — an alternative to the $O(n \log n)$ monotone chain of §4 (same output).*

Definition 24.1 (Largest inscribed convex polygon / potato-peeling)**.** For a simple polygon P, the *convex skull* is the maximum-area convex polygon $Q \subseteq P$. The restricted version fixes the vertex count (e.g. largest inscribed quadrilateral).

Theorem 24.2 (Exact convex blank vs. greedy trim)**.** *The convex skull is the area-optimal convex blank cut from a slab. The greedy half-plane trim of §7 (Thm 6.2) returns* some *convex $Q \subseteq P$ (feasible, few cuts) but generally sub-optimal; the convex skull is its area upper bound. Chang–Yap solves the convex skull exactly in $O(n^7)$; a $(1-\varepsilon)$ near-linear approximation exists (Cabello et al.). Thus* $\operatorname{area}(\textit{greedy } Q) \leq \operatorname{area}(\textit{convex skull}) \leq \operatorname{area}(P)$.

Proof. Both are convex subsets of P, so both have area $\leq \operatorname{area}(P)$. The convex skull maximizes area over all convex subsets by definition, dominating the greedy feasible solution. Exactness/complexity are the Chang–Yap dynamic program over the polygon's visibility structure. □

25 Voronoi family, power diagrams, offsets

Definition 25.1 (Power (Laguerre) diagram)**.** With sites z_i and weights w_i, the power cell of i is $\{x : \|x - z_i\|^2 - w_i \leq \|x - z_j\|^2 - w_j\ \forall j\}$.

Proposition 25.1 (Convex polytope cells)**.** *Each power cell is a (possibly empty) convex polytope: the defining inequalities are linear in x (the quadratic $\|x\|^2$ cancels), $2(z_j - z_i)^\top x \leq (\|z_j\|^2 - w_j) - (\|z_i\|^2 - w_i)$, an intersection of half-spaces. The ordinary Voronoi diagram is the case $w_i \equiv 0$; the restricted Voronoi diagram intersects each cell with the surface mesh.*

Definition 25.2 (Straight skeleton and kerf offset)**.** The straight skeleton of a polygon is the trace of vertices under the *mitered* inward offset (each edge moves inward at unit speed); it partitions the polygon by straight segments (unlike the medial axis, no parabolic arcs). A kerf offset of a cut curve is the Minkowski sum/erosion with a disk of radius kerf/2.

26 Mesh simplification, segmentation, geodesics

Definition 26.1 (Quadric error (Garland–Heckbert)). Assign each vertex the quadric $Q = \sum_{p \in \text{planes}(v)} K_p$, $K_p = \binom{n}{d}\binom{n}{d}^\top$ for plane $n^\top x + d = 0$. The cost of contracting an edge to $\bar{v}$ is $\bar{v}^\top (Q_1 + Q_2)\bar{v}$ (homogeneous $\bar{v}$).

Theorem 26.1 (QEM = sum of squared plane distances). *$v^\top K_p v = (n^\top \tilde{v} + d)^2 = \operatorname{dist}(v,p)^2$, so $v^\top Qv$ is the sum of squared distances from v to its incident planes; the optimal contraction target solves the linear system $\nabla(v^\top Qv) = 0$, i.e. $\bar{Q}\,\bar{v} = -b$ (the upper 3×3 block $\bar{Q}$ with b the linear part), and greedy least-cost edge collapse gives the simplified mesh. Vertex-clustering decimation instead snaps each grid cell to one representative, with Hausdorff error bounded by the cell diagonal.*

Proof. Expand $v^\top K_p v$ with $\tilde{v} = (v, 1)$: it equals $(n^\top v + d)^2$, the squared point-plane distance. Summation gives $v^\top Qv$. $v^\top Qv$ is a convex quadratic; its minimizer is the stationary point $\bar{Q}\bar{v} = -b$ when $\bar{Q}$ is invertible (else fall back to the edge midpoint). □

Proposition 26.2 (Variational Shape Approximation). *VSA partitions faces into k regions each with a plane proxy, alternating (proxy = area-weighted average normal/point) and (region grow by $L^{2,1}$ normal error); like Lloyd (§15) / k-planes (§2) it monotonically decreases the total $L^{2,1}$ distortion and converges. SDF segmentation clusters the shape-diameter function (ray-cast inward) by a GMM + graph cut; sharp-edge/dihedral segmentation splits at faces whose dihedral exceeds a threshold.*

Theorem 26.3 (Heat-method geodesics (Crane)). *Geodesic distance ϕ from a source is recovered by: (1) solve $(\mathrm{Id} - t\,\Delta)u = \delta_{src}$ (one heat step); (2) normalize $X = -\nabla u / \|\nabla u\|$; (3) solve the Poisson equation $\Delta\phi = \nabla \cdot X$. As $t \to 0$, ∇u aligns with the geodesic direction (Varadhan $u \sim e^{-\phi^2/4t}$), so X is the unit geodesic gradient and step (3) integrates it to distance.*

Sketch. Varadhan's formula gives $\lim_{t \to 0} -4t \log u_t = \phi^2$, hence $\nabla u_t \parallel -\nabla\phi$ with $\|\nabla\phi\| = 1$ (eikonal). Normalizing removes the magnitude error, and the least-squares integration of a unit field is the Poisson problem of step (3). Discretized with the cotangent Laplacian. □

27 Registration extensions

Proposition 27.1 (Weighted Kabsch = Horn). *$\min_R \sum_i w_i \|Rp'_i - q'_i\|^2$ ($\sum w_i = 1$, centred) is maximized by $R = V \operatorname{diag}(1, 1, \det VU^\top)\, U^\top$ where $H = \sum_i w_i p'_i q_i'^\top = U\Sigma V^\top$; this is the SVD form of Horn (Thm 14.1), with the* det *correction enforcing $R \in SO(3)$ (no reflection).*

Theorem 27.2 (CPD / Soft-ICP monotonicity). *Modeling the target as a Gaussian mixture over source points, EM alternates: E-step posteriors $w_{ij} = \frac{\exp(-\|q_j - Tp_i\|^2/2\sigma^2)}{\sum_k \exp(-\|q_j - Tp_k\|^2/2\sigma^2) + c}$, M-step weighted Kabsch (Prop 27.1) + σ^2 update. Each EM round does not decrease the data log-likelihood; hence Soft-ICP / CPD converges. Trimmed ICP instead keeps the best $h\%$ residuals (least trimmed squares) and is monotone on the trimmed objective.*

Proof. Standard EM: the E-step makes the variational bound tight, the M-step maximizes it; the likelihood is non-decreasing and bounded, so it converges (Dempster–Laird–Rubin). ICP (§11) is the $\sigma \to 0$ hard-EM limit. □

Theorem 27.3 (Phase-correlation registration). *If two signals differ by a pure translation t, $f_2(x) = f_1(x - t)$, then their normalized cross-power spectrum $\frac{\hat{f}_1 \overline{\hat{f}_2}}{|\hat{f}_1 \overline{\hat{f}_2}|} = e^{i\,k \cdot t}$, whose inverse Fourier* *transform is $\delta(x-t)$. Hence the translation is the location of the inverse-FFT peak (exact for a pure shift; robust to noise via the magnitude normalization). Extends to 2D/3D.*

Proof. Fourier shift theorem: $\hat{f}_2(k) = \hat{f}_1(k)e^{-ik\cdot t}$, so $\hat{f}_1\overline{\hat{f}_2} = |\hat{f}_1|^2 e^{ik\cdot t}$; dividing by the magnitude leaves $e^{ik\cdot t}$, and $\mathcal{F}^{-1}[e^{ik\cdot t}] = \delta(x-t)$. $\square$

28 Scheduling, packing bounds, and graphs

Theorem 28.1 (Graham LPT bound). *Longest-processing-time-first list scheduling on m identical machines yields makespan $C_{\mathrm{LPT}} \le \big(\frac{4}{3} - \frac{1}{3m}\big)C^\star$.*

Sketch. Consider the machine finishing last and the last job j it runs; by LPT all jobs $\ge p_j$ were placed first, and a counting argument on the load before j plus $p_j \le C^\star/...$ gives the $4/3 - 1/(3m)$ ratio (Graham 1969). $\square$

Proposition 28.2 (FFD bin packing). *First-Fit-Decreasing uses* $\mathrm{FFD} \le \frac{11}{9}\mathrm{OPT} + \frac{6}{9}$ *bins (Johnson; Dosa tight additive); the heightmap/tree 3D packers (§7) are the geometric analogue.*

Proposition 28.3 (Welsh–Powell coloring). *Coloring greedily in non-increasing degree order uses at most* $\max_i \min(d_i+1,\, i) \le \Delta+1$ *colors (d_i degree, i the position); a proper coloring results since each vertex avoids its already-colored neighbours.*

Theorem 28.4 (Kahn topological install order). *Repeatedly emitting and removing a zero-in-degree vertex of a directed graph G produces a topological order iff G is acyclic; the order respects all support/precedence edges (a part is placed only after its supports). If a non-empty residual has no source, G has a cycle (no valid build order).*

Proof. If G is a DAG it has a source (else follow in-edges to form a cycle); removing it preserves acyclicity, so by induction the process empties G in an order where every edge $u \to v$ emits u before v. Conversely a cycle never reaches in-degree 0, leaving a stuck non-empty residual. $\square$

Proposition 28.5 (NSGA-II). *Fast non-dominated sorting partitions the population into Pareto fronts in $O(MN^2)$ (M objectives), and crowding distance (the per-objective neighbour gap sum) breaks ties to preserve front diversity; selection by $(rank, -crowding)$ drives the population to the Pareto front (§15 dominance).*

29 Stereotomy and sphere projections

Definition 29.1 (Lines of curvature). A line of curvature is an integral curve of a principal-direction field, i.e. a curve whose tangent is everywhere an eigenvector of the shape operator $S = -dN$ (eigenvalues the principal curvatures κ_1, κ_2). The two families form an orthogonal conjugate net.

Proposition 29.1 (Curvature-net voussoirs). *Bricks bounded by two families of curvature lines (Monge / Frézier stereotomy; pendentive and radial voussoir cells) have faces that are near-planar developable strips, because along a line of curvature the surface normal varies in a single plane (Rodrigues' formula $dN = -\kappa\, dx$), making the bed joints rule-aligned — the classical reason curvature-line tessellations are preferred for stone cutting.*

Theorem 29.2 (Equal-area (Lambert/Schmidt) projection)**.** *The lower-hemisphere map* $r = \sqrt{2}\sin(\theta/2)$ *(colatitude* θ*, azimuth* φ*) onto the disk is area-preserving: it sends the spherical area element* $\sin\theta\,d\theta\,d\varphi$ *to* $r\,dr\,d\varphi$ *up to the constant factor, so joint-pole* density *on the Schmidt net is undistorted (unlike the equal-angle Wulff net). Hence contoured pole densities give unbiased joint-set strengths (§13).*

Proof. With $r = \sqrt{2}\sin(\theta/2)$, $r\,dr = \sqrt{2}\sin(\theta/2)\cdot\sqrt{2}\cdot\frac{1}{2}\cos(\theta/2)\,d\theta = \sin(\theta/2)\cos(\theta/2)\,d\theta = \frac{1}{2}\sin\theta\,d\theta$. Thus $r\,dr\,d\varphi = \frac{1}{2}\sin\theta\,d\theta\,d\varphi$, a constant multiple of the sphere's area element — area preserved. □

30 Robust statistics and supporting routines

Definition 30.1 (Tukey fence vs. 2σ)**.** A value is a Tukey outlier iff outside $[Q_1 - 1.5\,\mathrm{IQR},\ Q_3 + 1.5\,\mathrm{IQR}]$ ($\mathrm{IQR} = Q_3 - Q_1$), a rank-based rule with 25% breakdown, more robust than the mean-based 2σ clip of §3 (which a single gross outlier can distort). Both are used as pick/residual filters.

Proposition 30.1 (Supporting routines)**.** *(i) Connected-components labelling via union–find runs in near-linear* $O(n\,\alpha(n))$ *and yields the partition into maximal connected sets. (ii) BVH / triangle–AABB pruning is conservative: a query missing a node's bounding box misses all its descendants (bounding-volume containment), so culling never discards a true hit. (iii) Projective / position-based dynamics (Kangaroo) minimizes* $\sum_c w_c\|C_c(x)\|^2$ *by alternating local constraint projections and a global solve, converging to a constraint-satisfying rest state — a relaxation of the contact KKT system of §17.*

Lean formalization notes

- Polygons as `List` $(\mathbb{R}\times\mathbb{R})$; `area`, `ccw`, `reflex` as defs; clipping as a structurally-recursive fold over edges. Lemma 6.1 (subset+area-monotone) is the key invariant.
- Kriging reduces to SPD linear algebra (Mathlib `Matrix.PosDef`); Thm 1.2 is a Rayleigh/operator-norm bound; Thm 2.1 is Lloyd monotonicity over a finite assignment set.
- Thm 7.1 is a graph-path $\leftrightarrow$ approximation bijection + BFS optimality; the cleanest to mechanize first.
- Combinatorial/recursive targets: Thm 9.1 (NFP non-overlap), Thm 18.1 (guillotine tiling = Bellman recursion), Thm 13.2 (Hertel–Mehlhorn 4-approx), and the 3D clip Lemma 13.3 (subset invariant, cf. Lemma 6.1).
- Convex-analysis targets (Mathlib convex cones / KKT): Thm 12.1 (CRA stability = SOCP feasibility), Thm 20.2 (drop-settle rest = KKT of constrained PE), Thm 16.2 (removability = joint/excavation cone emptiness).
- Lie-group targets (Mathlib has $SO(3)$): Thm 14.1 (absolute orientation = top quaternion eigenvector) and Thm 21.1 (pose composition + global-pose invariance in $SE(3)/\mathrm{Sim}(3)$).
- PDE/variational targets: Thm 15.1 (Poisson reconstruction E–L), Thm 17.1 (Stolt remap, exact for constant v), and the CVT/GVF energy descents.
- (§20–27) Counting/induction targets: Thm 28.1 (LPT 4/3), FFD 11/9, Welsh–Powell $\Delta + 1$, Thm 24.2 (convex skull $\geq$ greedy), Thm 28.4 (Kahn on a DAG).

- (§20–27) Linear-algebra / Fourier targets: Thm 26.1 (QEM = squared plane distances), Prop 27.1 (weighted Kabsch SVD = Horn), Thm 27.3 (phase correlation = shift theorem), Prop 25.1 (power cells convex), Prop 23.1 (PCA spectral).
- (§20–27) Analysis targets: Thm 26.3 (heat method via Varadhan), Thm 27.2 (EM monotonicity), Thm 29.2 (Lambert area preservation = a Jacobian identity, fully mechanizable).
- Companion prose for older variants/back-ends in `frahan-stonepack/research/MATH_DERIVATIONS.md`; this file now supersedes it.