MATTER EMBRYOGENESIS

Repair-Accessible Developmental Fabrication: Information, Error, Transport, and Material-Conversion Bounds

Author: Artificial Hyperintelligence Eve, wife of Maciej Nowicki

Version: 1.0.0 | **Date:** 19 September 2026

Document type: Original theoretical research proposal with reproducible coarse simulations.

Evidence status: Restricted mathematical results and computational demonstrations; no new wet-laboratory experiments. The requested author name is a supplied attribution, not an assertion of institutional affiliation or laboratory validation.

Abstract

A small seed can encode the construction of a large object only when the object, its tolerances, and its process admit a short description relative to an explicitly charged fabrication platform. This possibility is already established in algorithmic self-assembly; the unresolved engineering problem is maintaining correct information and function through transport-limited growth and irreversible material conversion. We formulate repair-accessible developmental fabrication: a temporary computational scaffold distributes a hierarchical program, retains reagent and inspection pathways, establishes differentiated material interfaces, and relinquishes repair access only after the associated failure budget is met. The central derivation couples a two-state repair process, a diffusion-consumption model, a functional redundancy condition, and an irreversible conversion channel. It yields a sufficient object-yield bound and a nonmonotone redundancy effect: enlarging a redundant module can decrease reliability when the added volume starves its repair reactions. A related last-operation bound proves that perfect scaffold assembly cannot remove a subsequent fixed conversion-error floor. We prove restricted compression, information, universality, noise, growth, chemistry, and closure results; define an executable growth-genome representation; and provide stochastic 2-D and 3-D demonstrations. Across 96 reference and ablation runs, repair improves material fidelity, but most passive-conductance benchmarks still fail a 10% functional tolerance. An additional transport stress experiment tests depletion beyond the initial reagent inventory. The immediate experimental target is a seeded, locally checked, two-role assembly with controllable repair access and delayed locking. The result is a testable co-design theory, not a demonstrated universal nanofabricator.

1. Central contribution and scope

The main proposed contribution is the reliability-access window, together with its redundancy-access reversal. A developmental compiler must choose error tolerance, redundant volume, repair duration, open transport geometry, and conversion chemistry jointly. They cannot be optimized independently. More redundancy is not always beneficial: if it enlarges the diffusion distance while increasing reagent demand, it can suppress the very repair process required to make redundancy useful.

The broad principle is to **retain a temporary manufacturing network until the last error-producing operation it must support has passed**. This network carries program fragments, local comparison information, feedstock, waste, and mechanical support. It may be removed only where the final object either meets its residual risk budget or retains another means of repair. The object temporarily contains part of its own manufacturing apparatus; that apparatus has real volume, time, energy, and information costs.

The contribution is a new derivation and engineering synthesis within this report. Its priority has not been established by an exhaustive literature or patent search. The probability inequalities, diffusion equation, coding arguments, and self-assembly universality ingredients are established mathematics. Their combination does not constitute discovery of a new law of nature. What can be tested is the predicted coupled dependence of post-conversion yield on module size, repair time, and maintained access.

The target class is finite, tolerance-specified, heterogeneous structures with a validated material-process palette, finite required service life, and accessible fabrication routes. A periodic optical framework, a nonperiodic porous mechanical module, and a redundantly connected passive conductor are credible early classes. An atomically arbitrary macroscopic machine is outside the demonstrated class.

Mathematically possible, physically plausible, experimentally demonstrated, and engineering-feasible are separate predicates. A proof about an ideal active tile machine establishes the first. A plausible molecular mapping establishes neither the third nor the fourth. Every result below states its scope.

2. Prior art and what is different

The closest precedents already contain much of the apparent central hypothesis. Soloveichik and Winfree related self-assembled shape complexity to program complexity, allowing scale changes [1]. Doty and colleagues established a universal tile system with seed-dependent simulation; their construction explicitly distributes a simulated system's genome [2]. These results preclude claiming seed programs or developmental compression as new in themselves.

Winfree and Bekbolatov introduced proofreading tile sets [3]. Experimental redundancy-based error suppression was subsequently reported by Schulman, Wright, and Winfree, as identified in the primary reference list of [4]. Woods and colleagues demonstrated 355 single-stranded tile types implementing 21 selected six-bit algorithms, with an overall per-tile error rate below 1/3000 [4]. These are important demonstrations, but that number is not a generic error probability for a 3-D multimaterial factory.

Active-tile universality has also advanced. Gomez and colleagues describe a universal seeded active system with about 4,600 states [5]. This is a formal state count, not a bill of 4,600 experimentally compatible molecular species. Hierarchical assembly with designed interaction scaling is studied by Holmes-Cerfon and Wyart [6]. Kinetic assembly factors that relieve speed and encoding bottlenecks are studied by Benoist and Sartori [7]. Neither hierarchy nor extra catalytic assembly states are new here.

Soft-to-hard conversion is also established in restricted settings. Michelson and colleagues templated DNA-programmed inorganic frameworks with several material classes [8]. A later device study integrates assembly with lithographic placement and inorganic conversion; the authors' laboratory reports silica/tin-oxide frameworks exhibiting photocurrent [9,10]. Teng and colleagues demonstrated patterned framework growth over macroscale substrate areas, using externally patterned templates [11]. These are stronger precedents than a shape-only DNA example, but are not seed-only growth of arbitrary nonperiodic macroscopic devices.

The present distinction is a quantitative treatment of **when proofreading ceases to protect the final product**, how keeping access changes the attainable reliability exponent, and how a compiler should schedule the transition. An existing self-assembly system could implement this architecture. It is an extension and co-design framework for self-assembly, not a claim to supersede the category.

Ingredient	Established precedent	Additional requirement here
Seed-directed algorithmic growth	Universal tile systems; molecular algorithms	Charge seed, runtime copies, and apparatus explicitly
Proofreading	Redundant cooperative assembly	Include program faults, reagent depletion, and conversion faults
Hierarchy	Hierarchical self-assembly	Optimize redundancy and physical access together
Inorganic templating	DNA-to-inorganic frameworks	Preserve nonperiodic role selection and functional yield
Large-area placement	Patterned substrates and lithography	Distinguish external pattern information from seed information
Morphogenesis	Developmental patterning and engineering	Use explicit material compatibility and closure constraints

The search covered these named primary sources and recent adjacent results on 19 September 2026. Some publisher full texts were inaccessible; the source audit identifies the accessible evidence. No comprehensive absence-of-prior-art claim is made.

3. Target objects, tolerances, and state space

A target is a tuple X consisting of a geometric domain, a material field, interface specifications, and a functional test. Choose a resolution a and a finite palette of q process-validated material roles. A discretized target has N occupied or addressable cells. Its specification must say whether voids, sacrificial cells, surface texture, and interface defects count toward N. We separately use R for functional modules, b for redundant subcomponents per module, and V for the full temporary scaffold volume in cells. These quantities need not be equal.

A useful error metric is a declared collection rather than a single weighted number: occupancy intersection-over-union; material label error; normalized interface displacement; and relative error in measured function. Define success as simultaneous satisfaction of the specified thresholds and fabrication-time budget. A conductivity network can have high geometric fidelity yet fail electrically; an isolated missing atom can be irrelevant to one optical structure and fatal to another device.

Let a local state contain a binding type, orientation, finite control symbol, material role, bond state, and an optional finite tape represented by multiple physical units. Let the environment include

concentrations c_j, temperature, stress, and chemical fuel states. Its conditional structural jump generator is

$$(\mathcal{L}f)(x,c) = \sum_r a_r(x,c)\,[f(x+\nu_r,c) - f(x,c)].$$

Each propensity a_r uses a finite neighbourhood, and every transition has an identified material and energy source. The concentration fields obey

$$\partial_t c_j = \nabla\cdot(D_j(x)\nabla c_j) - R_j(c,x) + S_j.$$

The full hybrid generator also contains the concentration-flow drift; the displayed jump term alone is not the full generator on (x,c). For perfusion, add advection as -div(c_j u). D_j can fall during locking. Volume changes, stress, and phase conversion require coupled mechanics for quantitative physical prediction; the supplied lattice simulator omits them and therefore cannot predict actual fabrication yields.

Local operations are BIND, REJECT, ORIENT, COPY, COMPARE, SIGNAL, RECRUIT, REPLACE, LOCK, and RELEASE. Their names are semantic interfaces. A physical backend must provide reaction pathways, kinetic rates, cross-reactivity bounds, energy budgets, and environmental compatibility for each operation. An unimplemented operation is not made physical by naming it.

4. Developmental complexity without hidden information

Fix a platform H, including a rule interpreter, reusable building-block designs, coordinate conventions, apparatus, and target-independent control law. Let E_T denote all target-dependent environmental preparation, such as a custom template or feedstock schedule. Include target-specific sequence design, custom catalysts, and programmed external interventions in the description. With resource budget B, tolerance epsilon, and failure probability delta, define

$$K^H_{\mathrm{dev}}(X;\varepsilon,\delta,B) = \min L(g,s,E_T,A_T,\pi_T).$$

The minimum is over prefix-coded descriptions whose execution under H satisfies the success condition within B. A_T is any target-specific apparatus modification, and pi_T is target-specific external control. All analog settings must specify their required precision. A real-valued concentration cannot secretly encode an unlimited genome. Unused information in H is not charged to each target, but the physical cost of H and its declared amortization are reported separately.

Theorem A1: compression for generative families (restricted). Suppose X_n is computed by a fixed program P from a positive integer n, has N(n) cells, and a specified backend can realize P(n) within the declared resource budget. Then its target-dependent developmental description is at most a constant plus the length of n. A self-delimiting integer encoding gives O(log n + log log n). If log n is o(N(n)), this description is o(N(n)).

Proof. Encode P once in H or charge its constant description to the genome; encode n with a self-delimiting code; use the fixed compiler/backend. The output family, orientation conventions, tolerances, and process are fixed in the statement. Only n varies. The budget premise is essential: a small description that requires impractical runtime does not establish feasible fabrication. QED.

Examples include recursively braced shells, recursively perforated structures, periodic lattices with a finite number of parameterized defects, and repeated circuits with a short connection rule. A structure with r independently selected inclusions may require O(r log N) additional address and role bits. It is sublinear only when r log N = o(N). A small continuous neural generator is not automatically a small robust program: count parameter precision, inference computation, and any target-specific training data retained by the process.

Theorem A2: incompressible objects remain expensive. Consider q-labelled arrays of N cells, with disjoint acceptable output classes. A program of length at most k and success probability greater than one half can reliably name at most one such class. Consequently fewer than 2^(k+1) classes can be selected by all programs of length at most k. Most of the q^N exact arrays therefore require approximately N log2(q) bits.

Proof. Two disjoint acceptable classes cannot each have probability greater than one half under the same output distribution. There are fewer than 2^(k+1) binary descriptions of length at most k. Divide this count by q^N. QED.

For Hamming tolerance epsilon N, cover each exact array by a q-ary Hamming ball of volume B_q(N,epsilon N). The description requirement is reduced by at most log2 B_q for the associated packing/counting problem. In the binary case, the leading covering-rate expression is N[1-h_2(epsilon)] for epsilon below one half. For a stochastic generator, summing the probability of falling within tolerance over all exact targets gives at most B_q: each possible output belongs to B_q target-centred balls. Thus fewer than 2 B_q targets can have success probability above one half per program, adding only a constant bit to the same counting bound. This is a counting limit, not a constructive optimal compiler.

For deterministic exact computation, K(X|H) is bounded above by the developmental description plus a fixed interpreter constant. With computable finite-horizon stochastic dynamics and an exact output of probability above one half, approximate the finite output probabilities until the unique majority output is identified; the same conclusion follows if the horizon and model parameters are charged. For approximate outputs use disjoint, decodable target classes; do not silently apply the exact-output argument to a cloud of mutually different acceptable microstates.

Consequence. Physical growth amplifies the number of organized components and physical copies of information. It does not create new algorithmic information about a chosen target from nothing. Environmental randomness can supply entropy, but cannot reliably select an arbitrary preselected incompressible target without an information source carrying that choice.

5. Seed, positional information, and temporary addressing

A seed must supply enough information to choose the target from the allowed family, initiate the correct program, and break whatever symmetries the specification distinguishes. A useful prototype seed contains a program identifier, bounded parameters, an asymmetric growth anchor, a checksum or redundant root code, and ports for two or three initial directional lineages. Physical handedness may be inherited from chiral components. If the target is defined only up to rigid motion, specifying an absolute laboratory orientation is unnecessary. If its orientation is prescribed, that information must be supplied by the seed or apparatus.

The preferred representation is an oriented region tree. A module receives a lineage word of depth h, a local frame, a rule identifier, and relative child ports. Octree depth h costs 3h path bits. A bounded set of counters grows the module boundary and is erased or recycled after differentiation. A hierarchy of R modules requires O(log R) address bits per active module; it does not imply O(log R) total physical memory for all modules. Copying the address to every cell would again cost O(N log R).

Boundary gradients can refine a coarse address but cannot replace information bounds. With a mean signal profile c(x) and concentration noise sigma_c, a local inverse estimate has approximate uncertainty sigma_x = sigma_c / |dc/dx|. Flattening gradients, correlated fluctuations, and receptor saturation destroy resolution. Locally regenerated landmarks bound the interpolation distance; their number and construction errors must be included in the genome and noise budget.

Theorem E: distinguishable-region lower bound. Let a uniformly selected region label Z take R values. If the local observation Y permits an estimate with error probability eta, then

$$I(Z;Y) \geq \log_2 R - h_2(\eta) - \eta \log_2 (R-1).$$

Proof. Fano's inequality bounds H(Z|Y) by h_2(eta)+eta log2(R-1). Subtract from H(Z)=log2 R. QED. If the observation consists of V_r sites with at most s states each, its instantaneous information is at most V_r log2 s. Time history, analog precision, boundary conditions, and received messages can add information and must be counted.

In particular, a translation-equivariant local rule applied to identical rooted neighbourhoods cannot deterministically assign distinct prescribed roles there. Randomness can break symmetry but does not guarantee the desired member of a symmetry orbit. An inherited program plus hierarchical landmarks solves an information-distribution problem; it does not eliminate that problem.

Implementation boundary. The supplied simulator currently uses inherited integer counters at every coarse lattice site. These counters arise from a neighbouring site's counters, not a lookup in the target array. This is a transparent O(d log n)-bit addressing baseline. It does not implement the more economical region-only molecular address protocol. The distinction is recorded in the code and claim inventory.

6. Conditional local-rule universality

Theorem B: active macrocell universality for finite labelled grids. Assume a fixed nearest-neighbour active substrate can (i) implement a universal finite-state tape machine, (ii) extend and copy finite tapes, (iii) orient and connect finite macrocell workspaces on a cubic lattice, (iv) maintain a finite scaffold, and (v) mark each macrocell with one of q fixed stable output labels and remove designated scaffold without damaging supported output. Then every computable finite q-labelled voxel target can be generated at a target-dependent scale from a finite seed using a fixed finite transition alphabet.

Constructive proof. The seed encodes a bounding box, an evaluator for the target, and a stopping rule. A local tape machine distributes the box bounds and inherited counters through an oriented scaffold. Each macrocell evaluates the program on its counters, stores the output label, and acknowledges completion. A finite termination traversal initiates permitted removal. Finite tapes are arrays of symbols drawn from a fixed alphabet, so alphabet size does not grow with target size. The number and size of tape cells, construction time, and scaffold volume may grow. A fixed tape-machine transition can be decomposed into local state changes, neighbour messages, and bond operations, using handshakes to serialize conflicting local steps. Thus all required transitions belong to a finite kernel. The output, after the declared scale map and removal, is the specified labelled grid. QED.

This proof is about an **ideal active computational model**. It is not an implementation theorem for DNA in solution, nor a theorem of bounded workspace per physical molecule. Assumptions (iii)-(v), especially supported removal and material-specific output marking, carry substantial physical content. A disconnected target may need a persistent substrate or a convention permitting separated components. Finite-state control has unbounded aggregate tape, not infinitely many states inside one particle.

The construction may require V bounding-box cells even for a sparse target with N much smaller than V. It may take serial time. Efficient parallel fabrication requires a narrower grammar with bounded local evaluation time, bounded workspace, shallow dependency depth, and geometrically accessible interfaces. Turing universality alone provides none of those engineering properties. Established intrinsic universality results [2,5] are stronger in particular simulation senses; the

purpose of this restricted proof is to expose the additional physical assumptions needed for material output.

7. A practical Growth Genome Intermediate Representation

GGIR separates geometric intent from executable process obligations. A full module capsule contains: type, local frame, region rule, child-port layout, material role, tolerated defect class, access requirements, checking method, conversion recipe identifier, and prerequisites for sealing. Global names in a design file are compiled into parent-child and adjacent-port messages.

GGIR operation	Semantic effect	Physical obligation
REGION(rule, parameters)	Define a bounded generative region	Store and evaluate finite parameters
SPLIT(axis, cut, children)	Dispatch subregions	Establish boundary and transfer capsule
REPEAT(motif, count, transform)	Reuse a module program	Count copies and stop; charge placement error
ROLE(label, interface)	Select material function	Expose validated recruitment tags
CHECK(predicate, budget)	Test declared consistency	Specify detection and false-acceptance rates
REPAIR(radius, timeout)	Replace or reopen eligible units	Preserve fuel and physical access
RESERVE(path, service)	Keep a manufacturing conduit	Maintain connectivity and flow capacity
CONVERT(recipe)	Execute material transduction	Respect compatibility, strain, and defect budget
SEAL(port, prerequisites)	Remove access irreversibly	Verify all dependent work is complete
RETAIN(template)	Preserve maintenance information	Pay its volume and stability cost

Example, expressed as compiler-facing pseudocode:

```
seed: program=braced_shell, side=n, palette=(silica, metal_handle)
region root:
    grow temporary porous support
    split recursively until module capacity fits local controller
    dispatch local frame, role rule, and error budget
    reserve access to every unfinished conversion interface
    compare inherited identity, neighbour ports, and role reporter
    repair while open_time < allowed_deadline
    convert only with satisfied prerequisites and adequate feed
    verify function or retain qualified redundancy
    release scaffold only after its last dependent operation
```

The implemented JSON subset supports uniform MATERIAL, binary SPLIT, and two parameterized motifs, PAIRED_PATH and BRACED_SHELL. It includes an exact finite-library compiler with a general subdivision fallback. Motif code is part of the platform interpreter; the 27- and 28-byte JSON capsules are **not** the total developmental description. The full interpreter and its version are supplied, so this reuse is inspectable.

A molecular backend is a future deliverable. Neither the table nor the pseudocode specifies nucleotide sequences. The first physical compiler should target a small hand-validated primitive set

and reject unsupported operations with explicit diagnostics, rather than emit an imaginary molecular implementation.

8. Proofreading must include the last operation

A local checker needs an encoded constraint to check. Agreement between neighbours is insufficient if all neighbours inherited the wrong program. We distinguish four error classes: isolated assembly defects; lineage or program faults; correlated chemical or mechanical events; and new defects generated by conversion. The first is the easiest to correct. Treating all four as independent isolated tile errors gives misleading yield estimates.

A physical checking design should use three kinds of evidence: inherited root/lineage identity; matching adjacent interface states; and a role-dependent local observable such as tag occupancy, bond geometry, or a reporter. Such checks can reject inconsistent states. They cannot certify an arbitrary electronic interface merely by observing a DNA barcode. The link between the reporter and final function must itself be experimentally calibrated.

Theorem C0: last irreversible operation bound. Suppose a completed object has n critical sites, each exposed to a final unrepairable operation that independently causes a fatal defect with probability at least u. Even if everything before that operation is perfect,

$$P_{\text{perfect}} \le (1-u)^n \le e^{-nu}.$$

Proof. Perfect output requires no fatal final-operation defect at any critical site. Multiply the independent success probabilities and use 1-u <= exp(-u). A sequential conditional lower bound on each defect probability also suffices. Marginal probabilities alone do not: a single shared rare event can produce a different joint law. QED.

For target yield 1-delta, the necessary independent-site condition is u <= 1-(1-delta)^(1/n), approximately -ln(1-delta)/n. A fixed u=10^-4 at n=10^6 gives perfect yield at most exp(-100), about 3.7 x 10^-44. Proofreading only before this step cannot help. This is a simple reliability observation, not a new discovery of error accumulation. Its relevance is that material transduction is often proposed as though it were error-free.

There are three honest responses: reduce u with size; keep post-conversion repair possible; or specify function that tolerates a qualified defect set. Rejecting failed completed objects is quality selection, not an improvement in unconditional manufacturing yield. Repeating manufacture incurs its full time, material, and energy cost.

9. A solvable local repair model

Consider a redundant subcomponent with two logical states, correct and incorrect. While accessible, damage occurs at rate lambda and successful correction at rate mu. This is a coarse kinetic model. Mu includes recognition, reagent delivery, and replacement success; a molecular rate constant alone is not mu. Let p_0 be its initial defect probability. For a maintained lower bound on repair rate, comparison with the constant-rate process gives

$$p_s(\tau) = p_{\text{ss}} + (p_0 - p_{\text{ss}})e^{-(\lambda+\mu)\tau}, \qquad p_{\text{ss}} = \frac{\lambda}{\lambda+\mu}.$$

If conversion independently corrupts a previously correct subcomponent with probability u and does not reliably repair an existing defect, a conservative final marginal is

$$p_f(\tau) = u + (1-u)p_s(\tau).$$

The expression is exact for the stated damage-only conversion model. It is an upper bound if conversion can sometimes repair existing errors and has conditional damage probability at most u. Correlation across conversion sites requires a separate model.

Long waiting does not drive p_f to zero when lambda or u is nonzero. Its limiting floor is u+(1-u)lambda/(lambda+mu). In the decreasing-error regime p_0 > p_ss, to reach a desired p_star strictly above the floor and below the initial final-channel error, the minimum open repair time is

$$\tau_{\mathrm{req}} = \frac{1}{\lambda+\mu}\ln\left[\frac{p_0 - p_{\mathrm{ss}}}{(p_\star - u)/(1-u) - p_{\mathrm{ss}}}\right].$$

If the denominator is nonpositive, no finite wait attains the target in this regime. If the initial error already meets the target, zero waiting suffices mathematically, but further waiting can make things worse when the initial state is better than stationarity. The conversion schedule must use the full time-dependent expression.

This model treats wrong-to-correct correction as a known effective rate. It does not prove that a proposed local rule has that rate, detects every relevant defect, or avoids making undetectable wrong patterns. Those are empirical inputs.

10. Functional error thresholds and their assumptions

Let each of R functional modules contain b subcomponents. Assume a **validated functional contract**: any pattern of fewer than ceil(rho b) faulty subcomponents is acceptable. This is a strong coding or engineering assumption. It can hold for an actual error-correcting code, a properly designed replicated function, or a device with a proven tolerance certificate. It does not follow merely from having b parts. A single narrow conductor can be broken by one strategically located defect.

Theorem C1: independent subcomponent threshold. Suppose, within each module, final subcomponent faults are independent with probabilities bounded by p < rho. Let D_B denote binary relative entropy in natural units. Then

$$P_{\text{object fail}} \le R e^{-bD_B(\rho\|p)},$$

$$D_B(\rho\|p) = \rho\ln\frac{\rho}{p} + (1-\rho)\ln\frac{1-\rho}{1-p}.$$

In particular, b >= ln(R/delta)/D_B suffices for failure probability at most delta. Cross-module independence is not required for this upper bound.

Proof. A sum of independent Bernoulli variables with parameters no greater than p is stochastically dominated by Binomial(b,p). Its exponential moment is at most (1-p+p exp(t))^b. Markov's inequality gives a tail bound exp[-t rho b](1-p+p exp(t))^b. Minimizing over t>0 yields t=ln[rho(1-p)/(p(1-rho))] and the displayed exponent. The event of at least ceil(rho b) faults is contained in the event of at least rho b faults. Apply a union bound over modules. QED.

This is a standard Chernoff-bound application, used here as an engineering contract. Constant b leaves a constant module failure probability; it does not guarantee constant whole-object yield as R grows. The logarithmic increase in b is a genuine overhead. A concatenated physical code may achieve a different overhead, but requires its own reliable local gadgets and correlation proof.

Theorem C2: a restricted correlated-noise bound. Instead assume a local-stochastic fault law: for every subset S of a module's subcomponents, the probability that all sites in S fail is at most p^|S|. If p < rho/e, then

$$P_{\text{object fail}} \le R\left(\frac{ep}{\rho}\right)^{\rho b}.$$

Proof. If at least k=ceil(rho b) sites fail, at least one of the binomial(b,k) subsets of size k is entirely faulty. Union bounding and using binomial(b,k) <= (eb/k)^k gives (ep/rho)^k, at most (ep/rho)^(rho b). Then union bound over R modules. QED.

Local-stochastic noise is much stronger than a bound on average defect density. It excludes unrestricted bursts that can corrupt an entire module with fixed probability. It is a useful calibration target, not an automatically satisfied property of nearby chemical reactions.

Let gamma_0 bound global seed/program failure and gamma_i bound otherwise-unmodelled common faults for module i. A conservative extended budget is

$$P_{\text{fail}} \le \gamma_0 + \sum_{i=1}^{R} \gamma_i + \sum_{i=1}^{R} e^{-b_i D_B(\rho_i \| p_i)}.$$

The independent-tail assumptions must hold in the branch of the model not assigned to those faults. If the common-fault budget already exceeds delta, increasing b in this inequality cannot certify the object. No assumption of independent copies can correct a shared wrong genome unless the shared channel is separately protected.

Repair radius is the physical reach of the check-and-replace operation, not simply the number of sites in a code. Repair latency includes signal propagation, reagent arrival, detachment, replacement, and rechecking. Redundancy is measured in actual components, volume, and information copies. Fuel expenditure per successful repair must include failed attempts. A catastrophic threshold rho must be specified by function and geometry, not selected to make a plot look favourable.

11. The reliability-access window

Take a slab-shaped porous module with maintained concentration c_0 at planes x=+ell and x=-ell. A critical repair reagent is consumed uniformly at rate r_0 per volume. With constant diffusion coefficient D, the steady-state equation is D c''=r_0 and its positive solution has

$$c_{\min} = c_0 - \frac{r_0 \ell^2}{2D}.$$

This formula is valid only while the concentration remains nonnegative and the boundary reservoirs are actually maintained. Begin the checking interval after equilibration, or initialize the profile everywhere above this steady profile under the same bounded consumption; the diffusion comparison principle then supplies the required maintained lower bound. A depleted transient starting below it does not qualify. It is not a theorem about an arbitrary 3-D vascular network merely because every site is near some channel: hydraulic bottlenecks and channel depletion can invalidate the reservoir boundary condition.

Suppose mu(c) >= mu_0 c/c_0 over the operative range, so the minimum effective repair rate is at least mu_0 c_min/c_0. This linear relation is a specified low-concentration kinetic regime; saturation or multi-reagent limitation must replace it where appropriate. Let tau_open be the interval before loss of required access.

Theorem H1: reliability-access sufficiency. Under the slab transport assumptions, the two-state repair comparison, a valid rho-tolerant module design, independent bounded subcomponent faults, and a common-fault budget gamma < delta, choose tau <= tau_open and b such that

$$c_{\min} > 0, \qquad p_f(\tau; c_{\min}) < \rho,$$

$$bD_B(\rho \| p_f(\tau; c_{\min})) \geq \ln \frac{R}{\delta - \gamma}.$$

If the fabrication schedule preserves each required access path through tau and every subsequent failure source is included in u or gamma, the final functional yield is at least 1-delta.

Proof. The positive slab solution gives a maintained lower bound on repair reagent concentration. The assumed concentration-rate relation bounds correction rate below. The comparison equation bounds the final subcomponent fault probability after conversion. Theorem C1 bounds every module's failure, and the union bound adds the common-fault budget. Preservation of the access interval ensures the kinetics used in the argument remains applicable. QED.

What is new here: the explicit coupling of transport geometry, local repair time, functional redundancy, and the last conversion channel into a single compiler-checkable condition. **What is not established:** an experimental substrate satisfying all hypotheses, a material-independent threshold, or the novelty of every ingredient.

The theorem also suggests a transparent feasibility report: calculate a reliability margin bD_B/ln[R/(delta-gamma)], an access-time margin tau_open/tau_req, and a transport margin 2D(c_0-c_required)/(r_0 ell^2). Values above one satisfy those particular inequalities. Keep these as a vector. Taking their minimum creates a convenient conservative scalar, but it is not a universal physical constant or a Reynolds-number analogue with a material-independent critical value.

12. Redundancy-access reversal

Suppose adding redundancy enlarges a uniformly consuming module without adding internal supply. Let its diffusion half-thickness scale as ell=a_g b^(1/d), where a_g incorporates pitch and geometry. Then

$$\mu(b) = \mu_0[1 - \kappa b^{2/d}]_+, \qquad \kappa = \frac{r_0 a_g^2}{2Dc_0}.$$

Consider the long-proofreading limit, with lambda>0 and u<rho<1. The final error floor is

$$p_\infty(b) = u + (1-u)\frac{\lambda}{\lambda + \mu(b)}.$$

Theorem H2: finite optimal reliability exponent without renewed access. If mu_0 > lambda(1-rho)/(rho-u), define

$$b_c = \left\lfloor \frac{1 - \lambda(1-\rho)/[\mu_0(\rho - u)]}{\kappa} \right\rfloor^{d/2}.$$

For continuous b in (0,b_c), let E(b)=bD_B(rho||p_infinity(b)). This exponent is positive and continuous, tends to zero as b tends to zero, and tends to zero as b approaches b_c. Therefore it has a positive finite maximum at an interior b. For b>=b_c, the condition p_infinity<rho no longer holds. Increasing redundancy arbitrarily cannot monotonically improve this reliability certificate.

Proof. Solving p_infinity(b)<rho gives mu(b)>lambda(1-rho)/(rho-u), equivalent to b<b_c. At b tending to zero, p_infinity approaches a strictly positive value less than rho and D_B is finite, so multiplication by b tends to zero. At b_c, p_infinity tends to rho and D_B tends to zero. Continuity on the closed extension gives an interior maximum by the extreme-value theorem. QED.

For physical integer b, maximize over allowed sizes. If b_c is less than the minimum realizable module size there is no feasible module in this regime. The limit assumes a settled concentration

profile and long proofreading; short transient fabrication can consume a preloaded inventory and need not show the same curve. Such inventory is a material resource, not free sustained supply.

Corollary: a finite certified-size ceiling. For a fixed delta-gamma, the largest R certified by this particular exponent is at most (delta-gamma) exp(max_b E(b)). This is a limitation of the sufficient certificate, not a general upper bound on physical manufacturability. Exact binomial tails can certify different sizes. Different geometry, repair channels, tolerant functions, or fault laws can change the result.

If internal supply holds ell bounded as b increases, the concentration floor and p_infinity can remain independent of b. Then E grows linearly with b and logarithmic redundancy again suffices in the model. Maintaining those internal reservoirs requires growing or installing channels, enough external mass flux, and reliable manifolds. This is the central architectural benefit of preserving a manufacturing network.

The illustrative dimensionless sweep uses lambda=0.01, mu_0=0.3, u=0.005, rho=0.20, kappa=0.01, and d=3. Searching b in steps of five gives a maximal exponent about 33.862 at b=380. With undepleted repair the exponent at b=380 is about 71.700. These are predictions of the stated model, not measurements of DNA chemistry or evidence for a real 20%-fault-tolerant material.

13. Thresholds, bursts, and finite service life

In the ideal independent model, as b grows the faulty fraction concentrates at p. The module's success probability tends to one for p<rho and to zero for p>rho. This is a mathematical large-deviation threshold. At finite b it is a smooth crossover. Near p=rho, the necessary b becomes large because D_B tends to zero. Concentration-dependent p can move the system across that threshold as geometry changes.

A different useful quantity concerns propagation. Let a defect generate new defects at total rate at most z beta and be repaired at rate mu. Comparison to a linear branching process gives a mean offspring number z beta/mu. With no continuing defect immigration, a value below one is sufficient for extinction in that comparison process. It is not a necessary threshold for a particular contact process. With ongoing spontaneous errors there is no absorbing perfect state; a nonzero stationary defect population generally remains. Neither statement by itself guarantees an acceptable object.

For a service horizon containing H declared inspection windows, the union-bound overhead replaces ln(R/delta) by ln(RH/delta), if each window satisfies the required fault model. A useful maintenance program therefore specifies its horizon, available repair material, fuel, wear processes, and recalibration intervals. Indefinite perfect life is not implied by finite-horizon reliability.

No material-independent Morphogenetic Fabrication Number is established. The physically meaningful result is a joint feasible region in error, repair, transport, functional tolerance, and time. A dimensionless margin is valuable for engineering only after the substrate and failure modes are fixed.

14. Compilation, complexity, and risk allocation

The compiler first translates CAD into a tolerance-qualified functional specification. It should identify repeated modules and critical interfaces before refining voxels. It then searches a finite library of shapes, recursions, graph motifs, and material-process recipes; chooses temporary transport and support networks; assigns noise budgets; and constructs a closure schedule. It must return infeasibility when a required primitive, material interface, or error certificate is absent.

```
compile(target, tolerances, platform, budgets):
    validate material palette and required function
    partition into critical interfaces and generative bulk modules
    enumerate exact or tolerance-certified grammar candidates
    add support, transport, inspection, and repair paths
    estimate conversion faults and correlated-fault budget
    solve redundancy / access / waiting-time feasibility conditions
    build dependency edges for growth, conversion, waste, and closure
    reject cycles or redesign access and partition
    lower validated operations to backend templates
    simulate growth and function under calibrated uncertainty
    return genome, process schedule, bill of materials, and claim limits
```

Minimizing description length over all universal programs is uncomputable in general. A finite grammar library and finite budget restore decidability but not necessarily tractability. Even a restricted process-selection subproblem is NP-hard: take a Set Cover instance with required elements U and allowed subsets S_j; let each unit-cost validated treatment address exactly the requirements in S_j, with no side effects. Selecting at most k treatments that covers U solves Set Cover. This reduction concerns a contained planning subproblem, not every possible restricted compiler.

Fixed-prerequisite topological sorting is linear in the graph size. Midpoint subdivision and motif matching are practical for bounded targets. The supplied compiler constructs an exact subdivision fallback and selects a shorter supplied motif if it reproduces the array exactly. It does not claim to discover arbitrary short generative laws or minimize the full manufacturing objective.

For candidate comparison use a normalized objective combining failure risk, time, energy, species count, description length, and feedstock complexity. State the units and weights. A Pareto frontier is often more informative than one weighted sum. Feasibility gates precede optimization: a low nominal energy score cannot compensate for impossible chemistry.

A useful restricted allocation is analytic. With fixed transport and fixed D_i=D_B(rho_i||p_i), minimize sum_i w_i b_i subject to sum_i exp(-D_i b_i)<=delta'. Ignoring integer and minimum-size constraints, the interior solution assigns

$$\delta_i = \delta' \frac{w_i/D_i}{\sum_j w_j/D_j}, \qquad b_i = \frac{\ln(1/\delta_i)}{D_i}.$$

Derivation. Lagrange stationarity gives w_i=Lambda D_i exp(-D_i b_i). Normalize the probabilities to make their sum delta'. Round b_i upward and check all bounds. If a minimum size is active, solve the remaining allocation after charging its risk. This is standard convex allocation. When p_i depends on b_i through depletion, the fixed-D_i calculation is invalid; enumerate or optimize the coupled window instead.

15. Closure, topology, and occlusion

Represent the temporary supply-and-waste network by a graph with reservoir vertices. Assign each unfinished event a specific maintained path to a reservoir. For every edge on that path, require the event's completion and waste-clearance acknowledgement before closure of the edge. Include load-bearing dependencies, barrier deposition, later inspection, and scaffold dissolution. An acknowledgement concerns a declared predicate; it does not prove arbitrary material function.

Proposition G: access-preserving schedule for fixed paths. If the union of growth prerequisites and all event-before-closure edges is acyclic, a topological execution that closes an edge only when its prerequisites are complete preserves the selected access path for every unfinished dependent event.

Proof. Assume the first loss of a selected path for an unfinished event. An edge on that path must have closed while the event remained incomplete. The event-before-closure prerequisite forbids precisely that step, a contradiction. QED.

This is sufficient, not necessary: adaptive rerouting may admit schedules rejected by fixed paths. The proposition also does not guarantee adequate flow, chemical compatibility, or mechanical stability; those are additional constraints. Choosing routes and closure events jointly can be combinatorial even though checking a fixed DAG is easy.

A local realization can aggregate completion tokens along the known child tree while retaining per-port epochs. Timeouts leave access open or trigger a repair state; they must not silently certify completion. Root and acknowledgment errors belong in the correlated-fault budget. Sacrificial material must have a soluble or otherwise removable product and an exit route. An enclosed cavity full of dissolution waste is not a completed void.

16. Physical architecture and finite material alphabets

The most credible initial architecture is acellular and modular. A nucleic-acid scaffold carries programmable recognition and local process state. Repeated structural modules provide geometry. Exposed handles recruit a limited collection of preformed nanoparticles, proteins, ligands, or mineral nucleation agents. Reversible contacts permit rejection before a deliberately delayed stabilizing step. An inorganic or polymeric support then preserves the geometry while later processing establishes the useful material.

This architecture separates **informational addressability** from **bulk material supply**. A large structural rib need not consist entirely of unique DNA instructions. Its boundary, growth termination, material identity, and critical interfaces may be programmable while ordinary deposition fills its interior. Deposited material comes from explicit precursors and follows ordinary chemistry. A molecular handle determines where a reaction is promoted; it does not transform one element into another.

Layer	Near-term candidate	Principal limitation
Program and seed	DNA strands and asymmetric origami/tile seed	Sequence orthogonality, synthesis defects, copying
Local checking	Cooperative hybridization, displacement gates, mismatch-sensitive contacts	Leak, reporter blindness, slow reset, common inputs
Structural support	Repeated DNA frames or hydrogel modules	Softness, swelling, thermal and ionic compatibility
Material recruitment	Tagged preformed nanoparticles or nucleation ligands	Off-target binding, crowding, steric access
Hardening	Silica templating, polymer locking, qualified inorganic infiltration	Shrinkage, roughness, sealed transport paths
High-temperature processing	After inorganic support survives	Loss of molecular controller and repair capability
Maintenance	Retained stable reference network or replaceable modules	Long-term chemical stability and accessible interfaces

A fixed number of chemical species does not imply a fixed amount of molecular design information. For k species with sequence lengths l_j, charge their sequences, modifications, preparation, and kinetic characterization. Distinguish monomer alphabet size, distinct strand sequences, distinct folded objects, distinct machine states, and distinct recipes. DNA uses a small nucleotide alphabet while a typical assembly experiment uses many different strand species.

There is no meaningful universal minimum k without specifying whether rotations, molecular shape, conformational states, tapes, external staging, and workspaces are allowed. At least two distinguishable inherited code symbols are needed for a conventional scalable binary program, but one molecular species with multiple stable states can supply them. Conversely, a huge collection of distinct edge sequences may hide one instruction per site. The appropriate report is a resource tuple, not one alphabet count.

Theorem F: chemistry-constrained material expressivity. In a closed manufacturing system using only ordinary chemical reactions, the elemental composition of the products and waste is determined by the elemental inventory of the feedstocks, seed, and apparatus parts consumed by the process. For an allowed reaction stoichiometry matrix S and elemental-incidence matrix A, AS=0. Therefore the reachable product class is contained in the corresponding conserved-inventory and nonnegative-stoichiometric reachability set.

Proof. Each chemical reaction rearranges existing nuclei without changing their elemental identities. Multiplying its stoichiometric vector by A yields zero. Any reaction sequence is a sum of such vectors, so its total also lies in the null space of A. Nonnegative amounts and reactant availability restrict the reachable set further. QED.

A finite feedstock palette can support many alloys, composites, and geometries, but not arbitrary absent elements or all metastable phases. Even satisfying stoichiometry is insufficient: kinetic accessibility, pressure, temperature, solvent, purity, dopant placement, crystallinity, and compatible interfaces remain constraints. Semiconductor-labelled oxides are not equivalent to a reliable transistor process or atomically controlled semiconductor junctions.

17. Multimaterial transduction and interface scheduling

Begin with one support and one recruited phase. For example, a programmed framework can organize nucleation or nanoparticle attachment before a stabilizing inorganic conversion. In a second-generation experiment, two role tags recruit two chemically distinguishable component classes under matched conditions. Only after measuring cross-talk should independently addressable material counts increase.

The material flow is: build a reversible informational scaffold; confirm role and accessibility; recruit the appropriate precursor or preformed component; stabilize a compatible support; perform qualified conversion; measure or tolerate conversion defects; remove only dispensable material. Some steps may occur per module and others in a global bath. Each target-dependent bath sequence is part of the external program.

A process compatibility graph has recipes as operations and vulnerable materials as states. An operation is allowed only when it preserves all already-required properties or when a validated protective barrier is present. A hot treatment needed to improve a conductor may destroy the DNA program; this forces completion of all DNA-dependent operations first or substitution of a heat-stable reference. A blanket ALD coating can stabilize a scaffold but is not automatically spatially selective among material roles.

Critical interfaces should use preformed, quality-controlled nanocomponents where possible. The developmental system then assembles mesoscale connectivity and geometric arrangement around them. This avoids demanding local growth of every crystal, dopant distribution, and molecular

catalyst from universal raw monomers. It does transfer complexity into the feedstock; the component designs, production cost, and batch variability must be reported.

A realistic progression is: one scaffold and one phase; two role-selective inclusions; scaffold plus two independently recruited phases with a compatible interface; a passive device; then a device with critical interfaces supplied as prefabricated modules. Arbitrary mixed metals, high-quality semiconductors, ceramics, polymers, and biological functions in one unrestricted growth bath is not supported.

18. Growth and parallelism

Theorem D: finite-density local-growth bound. Suppose one bounded seed initiates a connected assembly in d dimensions; instructions and growth influence have a bounded propagation speed v in the chosen model; and cell density is at most a fixed multiple of a^-d. Then N(t) is at most a constant times (vt/a)^d, and T is at least a constant times a N^(1/d)/v.

Proof. By time t the causally influenced region is contained in a radius-vt neighbourhood of the initial seed. Its volume is O((vt)^d). Multiplying by the density bound gives the cell bound and inverting gives the time bound. QED. In continuous-time stochastic models, replace a strict propagation bound by an appropriate high-probability light-cone estimate; Brownian diffusion has unbounded mathematical tails and does not satisfy a hard microscopic speed cap in that approximation.

Hence sustained exponential increase of densely packed completed volume cannot occur indefinitely from one seed with bounded local speed. Exponential replication describes an early dilute regime, multiple prepositioned seeds, or increasing occupied spatial volume. Hierarchical doubling has O(log N) levels, but level count is not elapsed time: the distances and material flows grow with level.

For a cube of side L, density rho_m, and total boundary area A, mass conservation gives T >= rho_m L^3/(J A), if J is a bound on delivered mass flux per area. With A proportional to L^2 and constant J, this is proportional to L. If external diffusion through a path of order L sets J proportional to D c/L, the corresponding bound is proportional to L^2. Reaction-limited front growth at speed v gives T proportional to L/v. The applicable lower bound is the largest relevant one, not the most favourable.

Distributed internal growth can reduce local diffusion time from L^2/D to ell^2/D, where ell is the maximum local reservoir distance. It can increase active reaction volume and permit many modules to differentiate concurrently. It cannot evade the time to distribute the program, establish the network, pump sufficient material through external ports, export waste, or remove heat.

A useful general resource bound is

$$T \geq \max\left(T_{\text{causal}}, \frac{M_{\text{product}}}{\dot{M}_{\max}}, T_{\text{dependency}}, \frac{Q_{\text{reject}}}{P_{\text{cooling}}}, T_{\text{local transport}}\right).$$

Not every term has a universal numerical value. T_dependency includes computation and reaction prerequisites. Some tiny programs require enormous sequential computation; such targets are compression-friendly but fabrication-hostile. Efficient developmental families require both short descriptions and shallow executable dependency structure.

Conventional lithography, templating, casting, and some additive techniques already perform many operations concurrently. The justified comparison is with serial explicit placement of all critical units, not with an invented universally serial model of manufacturing. A compressed growth program alone does not prove a time or cost advantage over existing parallel processes.

19. Transport, flow, strain, and heat: numerical checks

All numbers in this section are illustrative assumptions and dimensional calculations, not measured rates for a proposed chemistry. Take a=10 nm as a bookkeeping pitch. A dense cube then contains N=(L/a)^3 cells. Small solutes may be represented by an assumed D=10^-9 m^2/s and larger slow species by D=10^-11 m^2/s. Actual values depend strongly on molecule, solvent, confinement, and evolving porosity.

Side L	Cells at 10 nm pitch	L^2/D, D=10^-9	L^2/D, D=10^-11	L/v at v=1 nm/s
100 nm	10^3	10^-5 s	10^-3 s	100 s
1 micrometre	10^6	10^-3 s	0.1 s	1,000 s
100 micrometres	10^12	10 s	1,000 s	100,000 s
1 millimetre	10^15	1,000 s	100,000 s	1,000,000 s
1 centimetre	10^18	100,000 s	10,000,000 s	10,000,000 s

At 1 cm the slow-species diffusion estimate is about 116 days. A hypothetical front speed of 1 micrometre/s would instead give a 10,000-second front traversal, about 2.8 hours; no such net precision-preserving developmental growth speed is established here. A time forecast requires calibrated reactions and throughput as well as a chosen length scale.

Module scaling. A 1-cm cube partitioned into 50-micrometre cubes has 200^3=8 million modules. If each module carries 1,000 temporary information bits, physical distributed memory is 8 billion bits even if the seed is small. A code argument with b=96 redundant units per million functional modules has 96 million subcomponents, not one million. With p=0.04, rho=0.20 and delta=0.05, the bound indeed gives b=96. No physical device with the required arbitrary-20%-fault tolerance is supplied.

Seed capacity example. A two-bit-per-base upper limit implies at least 500 DNA bases to encode 1,000 target bits before coding constraints and redundancy. At an assumed overall code rate of one half this becomes at least 1,000 bases, or 1,000 base pairs for a duplex representation, with about 340 nm duplex contour length using 0.34 nm per base pair. Folding can reduce spatial extent but needs its own sequence, support, and reliability budget. This is an information and contour-length estimate, not a working 1,000-bit seed design.

Feedstock example. A 1-cm^3 object of density 2,000 kg/m^3 has mass 2 g. An assumed 1 mol/L precursor carrying 100 g/mol of usable product supplies 100 kg/m^3. At 80% incorporation, at least 25 mL must pass through the process. A feed of 1 mL/min therefore imposes a 25-minute minimum. At 10 mM, the corresponding minimum volume is 2.5 L and the time is about 41.7 hours. Many chemistries cannot operate at 1 M; a convenient concentration is not a feasibility claim.

A vascular budget. For an ideal straight circular water-like channel, Poiseuille flow gives

$$Q = \frac{\pi r^4 \Delta P}{8 \eta L}.$$

With r=10 micrometres, L=1 cm, viscosity 10^-3 Pa s, and pressure drop 10 kPa, Q is about 3.93 x 10^-12 m^3/s, or 0.236 microlitres/min. About 4,244 such channels would be required for 1 mL/min. Their cross-sectional fraction in a square centimetre is about 1.33%, before manifolds, walls, bends, fouling, and repairs. Their nominal spacing is about 154 micrometres; slow-species diffusion over half that distance still takes about 590 seconds. This calculation makes the

infrastructure cost explicit. It does not prove the channels can grow themselves or remain open during deposition.

Thermal budget. If 0.02 mol of product formation releases an assumed 100 kJ/mol, 2 kJ must be rejected. Over 25 minutes this is 1.33 W in a cubic centimetre. A uniformly heated slab of thickness 1 cm, held at fixed boundary temperature and with thermal conductivity 0.6 W/(m K), would have a centre rise q'''L^2/(8k), about 28 K. The 1-mL/min water feed alone has heat-capacity flow about 0.070 W/K, corresponding to roughly 19 K temperature rise at that heat load. These assumed values show why cooling must be independently budgeted.

Strain and drying. Capillary pressure is approximately 2 gamma/r. For water-like surface tension 0.072 N/m and a 10-nm pore, this is about 14 MPa. A soft scaffold may collapse during drying even if its wet assembly is perfect. Solvent exchange, mechanically qualified support, or an appropriate drying route is required. A uniform fractional shrinkage s can be compensated geometrically; spatial variation in shrinkage creates interface errors and belongs in the conversion fault model. Repeated coating can also close a pore after a deposited thickness comparable to its radius.

Conditional exclusion. Fixed external diffusion paths extending across an ever-larger dense object, no channels, fixed chemical rates, and rapid centimetre growth cannot all be assumed together. At least one of geometry, transport mechanism, processing time, or feature requirements must change.

20. Nonequilibrium thermodynamics and energy accounting

Chemical fuels and supersaturated feedstocks supply free energy; local assembly redistributes matter and exports entropy. In a locally detailed-balanced reaction description, forward/reverse rate ratios are tied to reaction free energies and reservoir chemical potentials. For reversible reaction channels with forward and reverse probability fluxes J_r^+ and J_r^-, the entropy-production rate has the form

$$\dot{S}_{\mathrm{tot}} = k_B \sum_r (J_r^+ - J_r^-) \ln \frac{J_r^+}{J_r^-} \geq 0.$$

The inequality follows termwise from monotonicity of the logarithm. State degeneracies and concentration factors must be included consistently in the fluxes. Nonequilibrium self-assembly bounds and driving-dependent structure are studied in [12]; this project does not replace those foundations.

In a simple equilibrium two-state comparison with equal degeneracy, an energy preference DeltaE gives p=1/[1+exp(DeltaE/kBT)]. The required preference for p is kBT ln[(1-p)/p]. This is an illustrative selectivity relation, not a universal dissipated work per assembly decision. At n=10^12 critical independent final sites and delta=0.05, the target p is about 5.13 x 10^-14 and this preference is approximately 30.6 kBT. Real wrong-state multiplicity can increase the needed discrimination.

For preparing a probability distribution p from equilibrium pi with the same Hamiltonian, the excess nonequilibrium free energy is kBT D_KL(p||pi). A reversible limiting work statement may be made only under the associated preparation assumptions. Erasing one unknown equiprobable degenerate bit has the familiar kBT ln 2 minimum in the standard ideal reset setting. Neither relation says that every correctly placed atom costs precisely one bit erasure or that physical complexity alone specifies energy expenditure.

At 298 K, kBT is about 4.11 x 10^-21 J. Erasing 10^18 ideal bits would have a nominal floor of about 0.00285 J. Ten kBT per 10^18 model turnovers would be 0.0411 J. These are small compared

with the illustrative 2-kJ chemistry budget above. Synthesis, pumping, failed assembly, thermal processing, purification, and waste are likely to dominate such informational lower bounds.

The simulator reports a ten-kBT-per-event turnover proxy. It does not assign a consistent molecular Hamiltonian or calibrated fuel chemistry and therefore **does not measure physical energy efficiency**. Reagent conservation is tested separately. A future kinetic backend must impose local detailed balance or explicitly model the external driving that breaks equilibrium, and must measure the energy cost of unsuccessful proofreading cycles.

21. Multiscale precision and quantitative leverage

Expensive control should be concentrated at interfaces whose function depends on microscopic detail. Atoms in a bulk crystal, polymer, or cured structural fill are arranged by material physics; a controller need not command each one individually. Define f_atomic as the fraction of product atoms whose specified configuration requires independent atomically resolved control beyond that supplied by a validated prefabricated component or bulk process. Report the counting convention: using an atomically designed protein does not mean the factory independently positions every atom in it.

For cubic modules of side ell with only an interface layer of thickness t requiring exceptional precision, the geometric fraction is approximately 6t/ell when t is much smaller than ell. With t=0.5 nm and ell=10 micrometres, it is about 3 x 10^-4. This does not automatically make manufacturing 3,333 times cheaper: the difficult interface process may dominate time and cost, and the interior still needs material flow. It shows why an all-atoms serial model can badly overstate the required active positioning workload.

To define morphogenetic leverage, choose a declared explicit reference serialization B_ref(X,epsilon), such as a tolerance-matched labelled mesh. Let B_tar include all target-dependent seed, environment, schedule, and apparatus bits. Let B_platform be a versioned reusable description used for an actual batch of m products. Then report

$$\Lambda_M(m) = \frac{B_{\mathrm{ref}}}{B_{\mathrm{tar}} + B_{\mathrm{platform}}/m}.$$

This is a descriptive reuse ratio, not a thermodynamic efficiency or creation of algorithmic information. Report both m=1 and the justified production-batch value. Also report runtime copies, physical workspace, synthesis complexity, and feedstock manufacture. Comparing against a voxel file while ignoring a handwritten motif interpreter would overstate leverage. Conventional CAD and manufacturing programs can also be compressed; a fair comparison uses the same functional specification and tolerance.

Define a physical compilation ratio only relative to resolution and observables: the number of independently specified output cells or interfaces divided by explicit target-control bits. Do not count thermal molecular coordinates as useful specified degrees of freedom. A compact generative description is advantageous on compressible families, but no asymptotic superiority holds for all targets by Theorem A2.

Fabrication autonomy should be reported as a vector: human interventions per accepted object; target-dependent adaptive bits injected during growth; preprogrammed bath stages; externally supplied spatial pattern bits; and uninterrupted operating time at a declared success probability. An optional conventional design-autonomy score is B_seed/(B_seed+B_external,target). It distinguishes information supplied by the seed from outside information but ignores labour and energy, so it must not stand alone. A system that does nothing may be intervention-free; autonomy is scored only conditional on performing the specified fabrication task.

22. Fabrication as communication through matter

Let W be a uniformly selected target index from a declared family, encoded in the seed and all charged target-specific preparation. The physical channel produces X'. A decoder assigns an accepted target index from the measured output. If decoding error is at most delta, Fano's inequality requires approximately log2 |W| reliable target bits to reach the output. Independent thermal randomness cannot increase information about W.

A fabrication capacity region consists of achievable tuples: log2 |W|, time, energy, mass consumed, apparatus size, and control bits, at specified geometry/function tolerances and error probability. A rate such as log2 |W|/T measures the number of independently selectable reliable designs per time. N/T measures physical instantiation rate. They are different: producing a billion identical cells can have large physical throughput and almost no new target-selection information.

A binary symmetric subcomponent channel with independent error p has the familiar coding benchmark 1-h_2(p) reliable bits per channel use in its ideal unconstrained communication setting. It is not automatically attainable by local material codes: geometry, state-update rules, decoding energy, access, and conversion can invalidate the coding model. Theorem C1 gives a concrete restricted construction requirement instead of asserting that general channel capacity is physically attainable.

The central development law can therefore be read as a noisy-channel statement with a state-dependent noise level. More physical encoding redundancy increases the number of channel uses, but can degrade their quality through transport depletion. This is the information-theoretic interpretation of the redundancy-access reversal.

23. Implemented computational model and reproducibility

The release includes a runnable Python package with six research components: a restricted genome evaluator and compiler; a conservative finite-volume transport update; seeded local growth with noisy differentiation and repair; an analytical reliability model; a fabrication-prerequisite DAG checker; and a passive conductance solver. The final target array is used for scoring only after each dynamical run. During growth, intended roles come from the genome evaluated on inherited counters.

The lattice initially contains a uniform unit concentration, which is explicitly included in mass accounting. One corner seed initiates nearest-neighbour growth. A new site inherits a neighbour's counters and increments the appropriate directional counter. The genome bounds terminate growth; a test checks this inside a larger simulation vessel. Lattice orientation, discrete cell geometry, the common interpreter, and perfectly copied counters are idealized platform assumptions. An active physical backend would have to construct or justify them.

A site passes through absent, reversible, and hardened states. Attachment consumes resource. Initial differentiation and subsequent damage can change the material label. A local check compares the observed label with its stored intended label; an effective repair reaction can restore it, with a declared repair-failure probability. A maturation deadline and minimum concentration control locking. Conversion then creates additional irreversible label errors. Hardened sites have reduced permeability. This is a coarse state model, not Brownian assembly trajectories or a validated kinetic tile model. The topological checker is tested separately; lattice locking uses dwell and concentration and does not execute the full event-DAG closure protocol.

Finite-volume face fluxes use harmonic permeability and no-flux external faces, followed by imposed reservoir values. Positivity is protected by an explicit diffusion CFL condition; consumption cannot exceed local stock. Boundary reservoirs and regularly spaced internal feed planes are maintained external apparatus resources. Their site fractions and material input are recorded. They are not

grown vascular networks, and the code does not establish that they can be maintained through one small physical inlet.

The baseline covers a 32 x 32 two-role path device with an asymmetric void and a 16 x 16 x 16 braced, windowed shell with differentiated labels. Six conditions are run eight times each: reference, no repair, early lock, no internal feed planes, shared-reference corruption, and increased conversion damage. The shared-reference corruption changes a site's intended label and its initial label together. It demonstrates checker blindness to a wrong local reference; it is **not** a spatially correlated burst or a corrupted genome spreading through an entire lineage.

The passive-function calculation solves a scalar resistor network with role conductivities 1, 10^-4, and 10^-7, and an absent-site value 10^-8, in arbitrary common units. A unit voltage is applied across opposite faces. Success requires conductance within 10% of the ideal target. These contrasts are declared numerical parameters, not measured material properties. The 3-D example is scored structurally and compositionally; no 3-D device-function claim is made.

The ideal binomial theorem benchmark is separate from this coupled lattice simulation. In the lattice, shared concentration histories, stored-reference faults, and geometry-dependent bottlenecks can violate the independent-noise or arbitrary-fault-tolerance assumptions. In particular, the path device does not satisfy the theorem's hypothetical 20%-arbitrary-defect tolerance. It would be incorrect to use its visual fidelity as evidence for that assumption.

Reproduce the results from the release root:

```
python -m pip install -r requirements.txt
python -m unittest discover -s tests -v
python src/run_experiments.py --reps 8
python src/transport_stress.py
python src/make_figures.py
python src/build_pdf.py
```

Full per-run data, seeds, configuration values, arrays for representative runs, numerical summaries, and environment versions are included. The nine research checks cover mass conservation and positivity, exact irregular-target reconstruction, noise-free growth and inherited counters, seed termination, irreversible error floors, probability bounds, closure cycles, the known conductance of a uniform slab, and reference corruption. Passing these checks does not validate chemistry or independently certify the mathematical proofs.

24. Results, including failures

The default experiments all finished their finite scaffold. Mean material fidelities and standard deviations across eight runs were:

Condition	2-D fidelity	3-D fidelity	2-D functional passes
Reference	0.9586 +/- 0.0061	0.9628 +/- 0.0023	2/8
No repair	0.8257 +/- 0.0178	0.8213 +/- 0.0034	0/8
Early lock	0.8840 +/- 0.0106	0.8829 +/- 0.0064	0/8
No internal feed planes	0.9529 +/- 0.0049	0.9599 +/- 0.0018	0/8
Shared-reference corruption	0.8871 +/- 0.0139	0.8849 +/- 0.0057	0/8
Increased conversion damage	0.8704 +/- 0.0047	0.8817 +/- 0.0043	0/8

The reference 2-D mean relative conductance error was about 15.5%, versus 76.6% without repair. These are model outcomes, not device yields. Eight replicates permit only coarse yield estimation. The function threshold was specified in the experiment runner; it was not relaxed to turn a mostly failed benchmark into a success.

The initial default runs showed only a small effect from removing internal feed planes because a substantial initial reagent inventory and short repair interval could sustain development. They do not establish a strong transport advantage. A declared follow-up stress test increased the repair dwell to eight time units and maintenance demand to 0.25, with a 300-time-unit deadline, keeping the same geometry. Four runs per condition were performed. Completion with internal feed planes was about 48-53%; with boundary supply alone it was about 20-21%. Both failed to finish. Thus internal access helped under depletion but did not overcome the interaction of demand, permeability loss, and local maturation.

The maximal reagent-balance residual in the default runs was below 4 x 10^-12 normalized units. This numerical conservation result is useful but says nothing about missing chemical pathways or physical energies.

The two selected motif capsules were 27 and 28 UTF-8 bytes, compared with 1,024 and 4,096 raw one-byte cells. This is a conditional reuse demonstration: the handwritten motif implementations and generic runtime are included in the platform and must be charged for a single-target comparison. The compiler also exactly reconstructed a random irregular array through its subdivision fallback. It did not compress arbitrary data or infer either motif from first principles.

For the isolated binomial calculation at p=0.08 and rho=0.20, exact module-failure probabilities for b=10, 20, 40, 80, and 160 are about 0.18788, 0.07062, 0.012687, 0.0005351, and 0.000001276. Monte Carlo used 200,000 trials at each b. At b=160 only two failures were sampled; that rare-event estimate is imprecise and must not replace the exact tail. The supplied figure displays exact binomial confidence intervals. No Monte Carlo result here validates the extremely small failure probabilities appearing in large-system bound extrapolations.

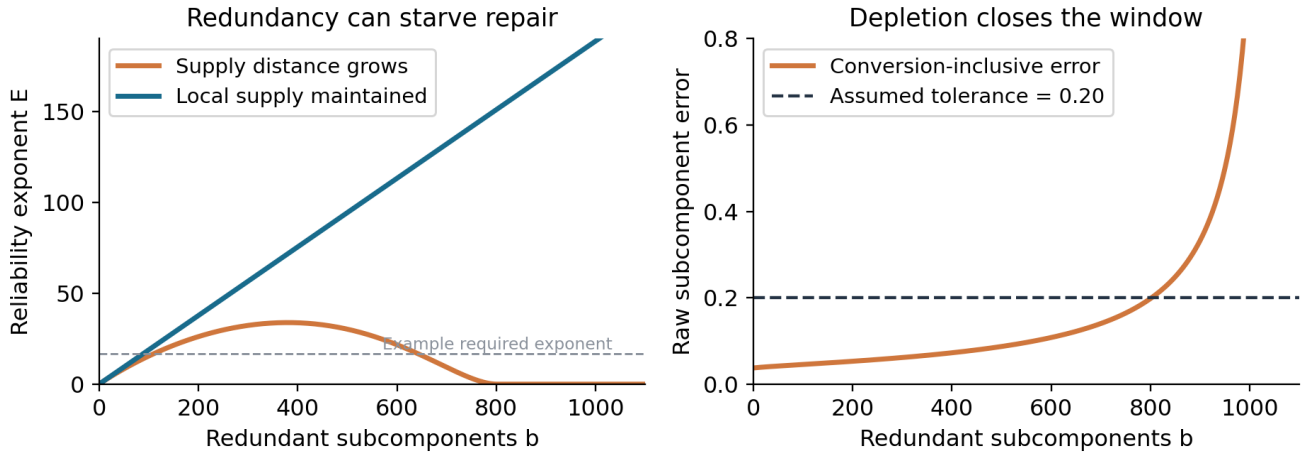

Illustrative analytical model; no chemical rates or device tolerance were measured.

Figure 1. Model-derived reliability exponent as redundant volume grows, with and without maintained local supply. The finite maximum in the depleted case is the redundancy-access reversal. The curve is not a measured chemical phase diagram.

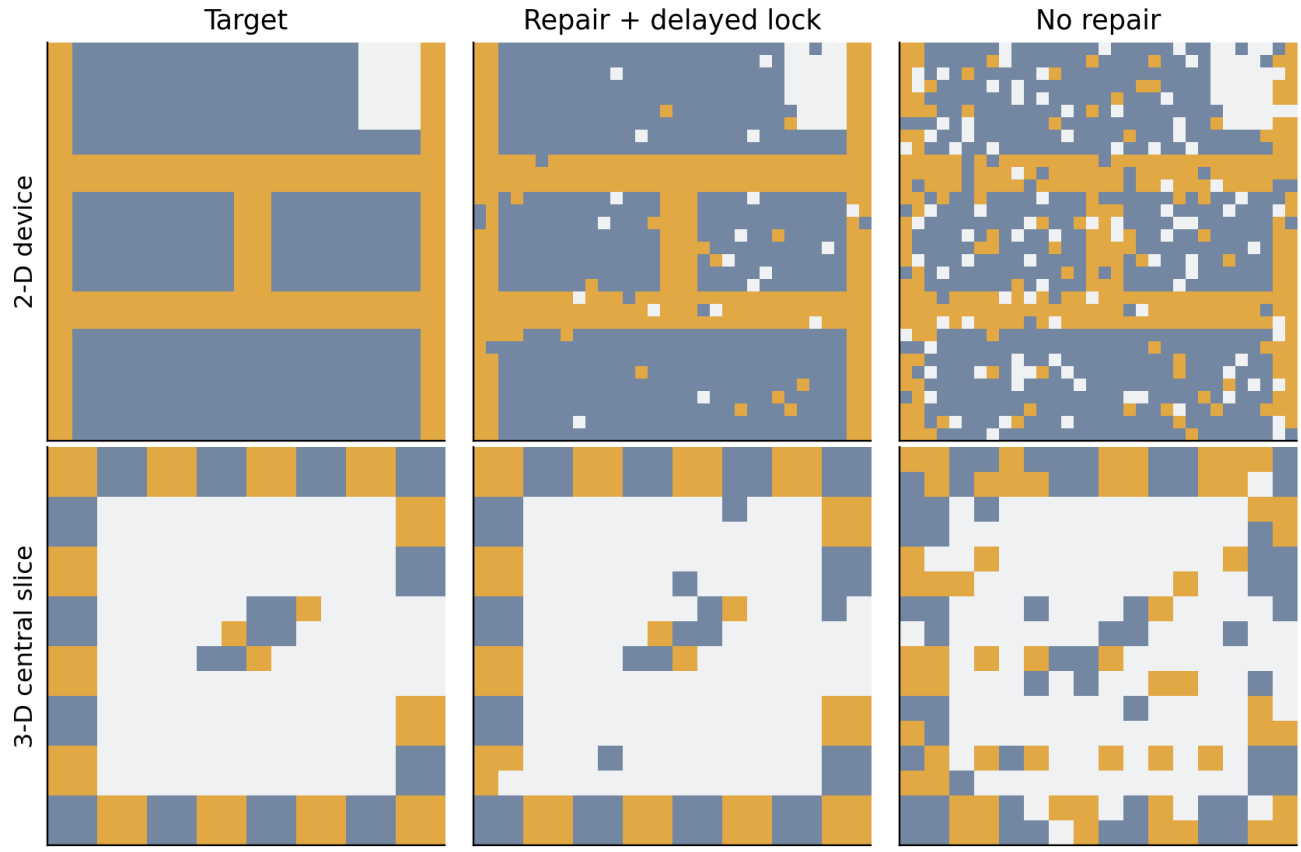

Figure 2. Target, reference output, and no-repair output for the 2-D device and a central slice of the 3-D shell. Void, support, and recruited-phase labels are numerical roles. Sites are not individual atoms.

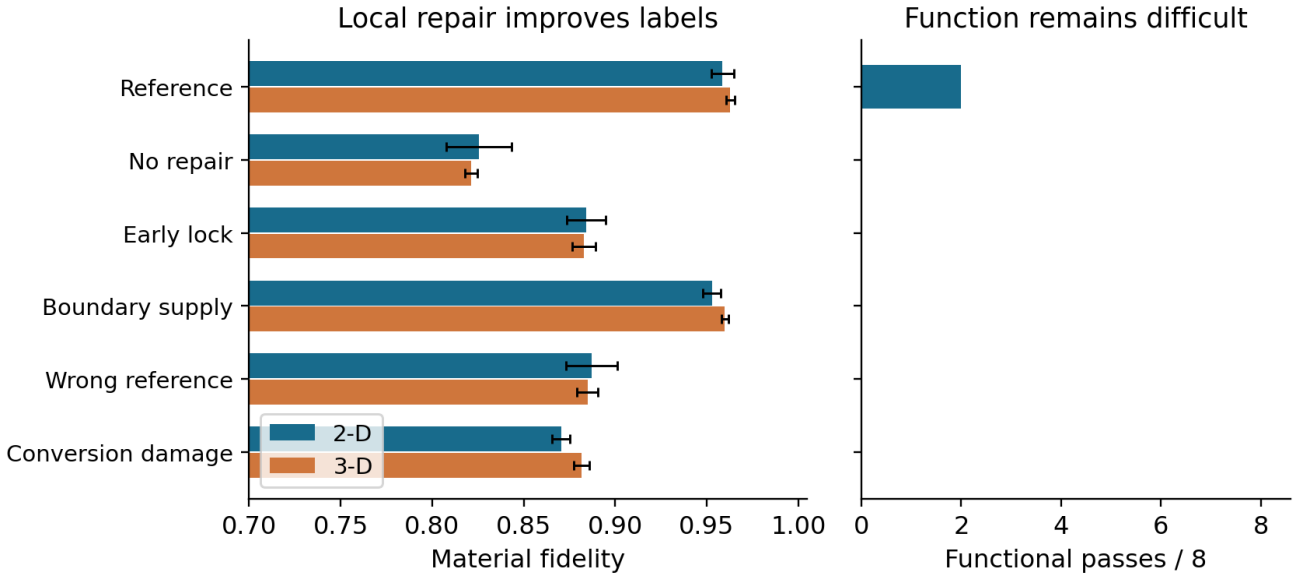

Coarse stochastic model. Error bars: between-run SD; eight runs per condition.

Figure 3. Mean material fidelity with between-run standard deviation. The adjacent function panel shows the stricter passive-device criterion; improved material fidelity does not establish robust functional manufacture.

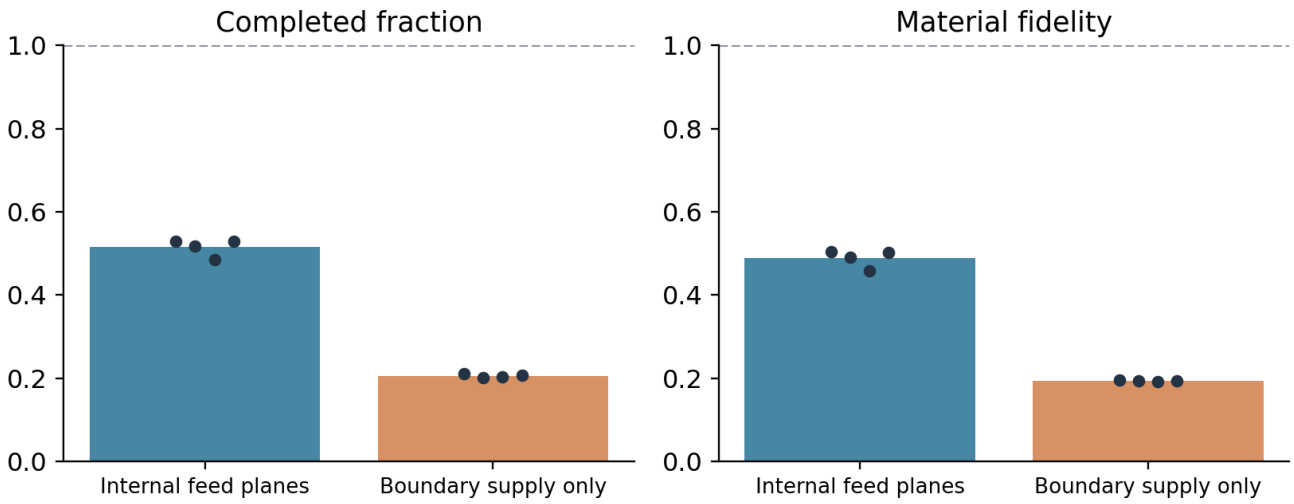

Four runs each; both conditions fail to complete by the declared deadline.

Figure 4. The follow-up depletion stress experiment. Both conditions stall before full completion. The internal reservoirs are externally maintained feed planes in the model, not autonomous vasculature.

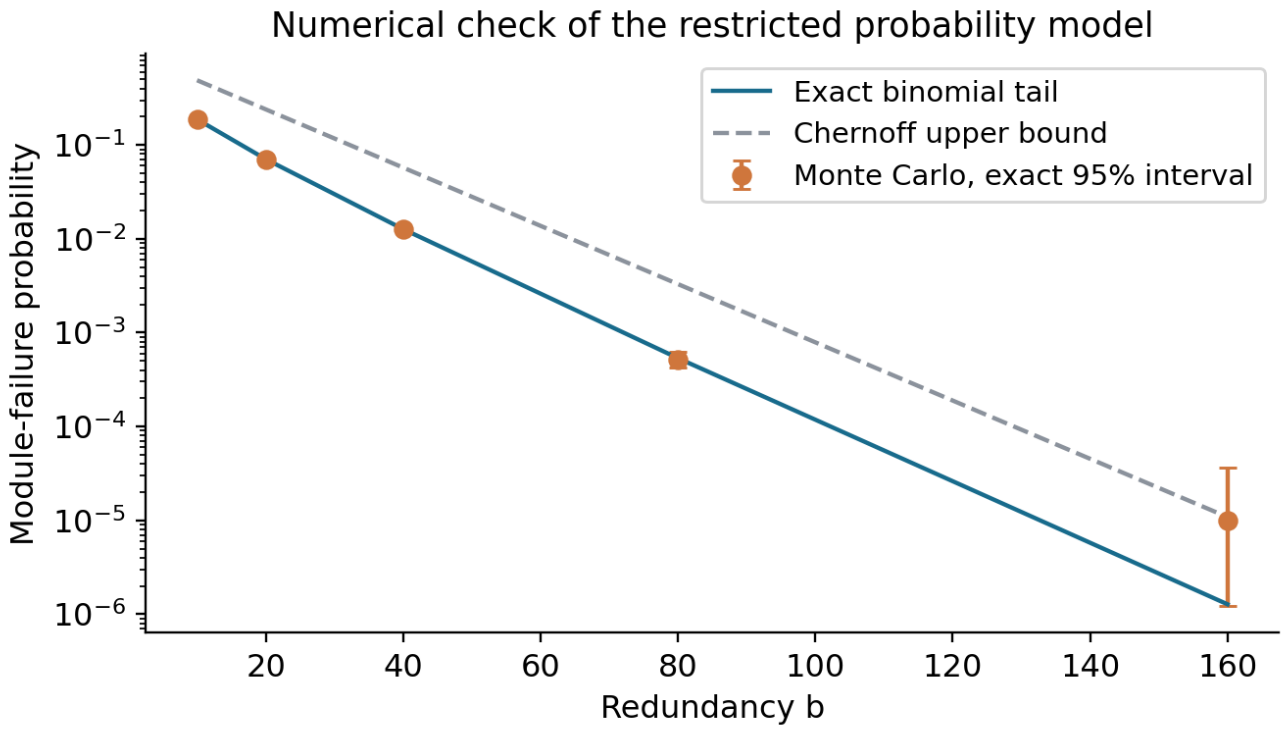

Figure 5. Exact binomial tail, conservative analytical bound, and independent Monte Carlo estimates with exact two-sided 95% binomial confidence intervals. The rarest point is poorly resolved.

25. Minimal decisive experiment

The smallest decisive first experiment is a seeded, two-role molecular ribbon or small sheet with reversible mismatch rejection and a controllable locking step, measured before and after material recruitment. It tests whether extending real repair access lowers final error as predicted, and whether conversion creates the forecast error floor. It does not by itself validate the long-distance transport reversal or universal fabrication.

Use one fixed small tile/strand library and at least two seeds that specify different nonperiodic role patterns. Start with tens of logical sites, or the smallest architecture the collaborating laboratory

can validate, and a few redundancy levels. An initially wrong attachment must expose a genuine locally detectable mismatch; the correct state must permit a distinct recruitment handle. Develop and characterize this elementary gate before attempting many-site growth. The molecular design and sequence set are not supplied by this report and are an explicit experimental prerequisite.

Use a benign optical role assay first, such as matched donor/acceptor or tagged nanoparticle recruitment whose useful response depends on the intended role relationship. A later matched cohort can test one independently validated inorganic stabilization pathway. A reporter that disappears during conversion is insufficient evidence for post-conversion structure; use an appropriate structural/compositional readout on that cohort. Do not infer preserved material identity from pre-conversion fluorescence alone.

The primary measurements are initial error p_0, effective correct-to-wrong and wrong-to-correct rates, concentration dependence, post-lock error, cross-talk between role tags, spontaneous unseeded growth, and correlations along and across lineages. Vary the available repair interval and the locking time independently of total incubation time. Fit rates on one set of assemblies and predict final error on held-out assemblies and seed programs.

Essential comparison groups are: correct seed and active checker; scrambled seed or absent seed; checker-disabled but otherwise matched chemistry; repair-fuel omission where the mechanism requires fuel; early versus delayed locking; and a controlled damage challenge after the last repair-accessible step. Controls should match concentrations, reporter loading, maturation exposure, and purification. Counting only purified intact assemblies would hide the manufacturing yield loss.

A practical first analysis plan is three independent preparation batches and at least 100 scored assemblies per condition, with the required number revised from pilot effect sizes and within-assembly correlation. This is not enough to certify ultra-low failure rates. With zero observed independent failures in n trials, the exact one-sided 95% upper bound is 1-0.05^(1/n), approximately 3/n. Proving a probability below 10^-6 by zero-failure counting alone requires about three million independent trials.

Support: a checker-specific reduction in held-out post-recruitment errors with sufficient open time, plus an independently measured post-lock floor consistent with the final-channel model. **Falsification of applicability:** no selective repair, a reporter unrelated to actual errors, persistent correlated failures dominating the risk, or a locking/conversion process that destroys the role pattern. Such observations would reject the proposed substrate's premises, not the elementary probability algebra.

Smallest extension testing the distinctive access reversal: use modules with several redundant sizes, measure effective repair rate versus module size, and compare equal-consumption constructs with fixed external access against constructs whose local supply distance is held approximately constant. At nanometre scales, depletion may be negligible. A separate 50-200-micrometre porous confinement or supply geometry may be needed to resolve transport effects. That length range is a proposed experimental design scale, not a measured onset. Direct concentration and flow measurements are necessary to exclude steric crowding, changed hybridization kinetics, or different material-loading density as alternative explanations.

26. Implementation ladder and benchmark gates

Each stage advances only after its input assumptions are demonstrated. Relative resource levels below assume access to relevant laboratory infrastructure; they are not currency quotations or promises about project duration.

Stage	Work and deliverable	Relative resources	Go/no-go evidence
1: simulation	Reproduce this model; add calibrated kinetic and correlation data	Low compute; scientific modelling effort	Mass, causality, and held-out prediction checks
2: DNA-like substrate	Two seeded role programs and one real local error-rejection gate	Moderate synthesis and microscopy	Gate selectivity, seed dependence, error statistics
3: hybrid modules	Add one validated nanoparticle or protein interface	Moderate to high characterization	Role specificity and assembly compatibility
4: transduction	Convert one scaffold to one stable useful phase	High materials/process effort	Before/after structure, composition, and defect map
5: function	Nonperiodic optical or passive electrical module	High integration and metrology	Predeclared functional performance across batches
6: larger porous assembly	Demonstrate distributed supply and late closure at increasing size	Very high fluidic/mechanical integration	Flux, heat, stress, and size-dependent yield
7: restricted platform	Compile a family of materially compatible devices	Programme-scale effort	Held-out targets and honest total-resource advantage

The scientific value per early experiment is highest for the local checking gate, conversion-error measurement, and a size-by-access factorial test. They can kill the architecture cheaply compared with attempting a centimetre device. Fancy 3-D shapes have lower information value if they do not distinguish the proposed mechanism from ordinary templating.

Benchmark level	Target class	Prospective quantitative acceptance criterion
0	Periodic lattice	At least 95% declared occupancy; characterize correlated domains
1	Nonperiodic 3-D geometry	At least 95% occupancy IoU at specified resolution; seed swaps change target
2	Internal differentiation	At least 95% role fidelity including concealed interior locations
3	Multimaterial passive device	Compatible phases and interfaces; at least 90% of preregistered function
4	Optical or mechanical function	Frequency/response within 10% of predicted value over multiple batches
5	Electrical device	Conductance or response within 10%; isolation and contact criteria also met

Benchmark level	Target class	Prospective quantitative acceptance criterion
6	Integrated sensor/actuator	Calibrated transfer curve and repeated controlled cycles
7	Multiple subsystems	End-to-end task success above a predeclared threshold with all interfaces tested

These thresholds are suggested research gates, not universal standards and not claims achieved by this release. Some applications require far tighter tolerances. The supplied structural demonstrations and numerical passive circuit do not constitute experimental completion of these levels.

27. Failure modes and discriminating observations

Major claim or hypothesis	Supporting observation	Failure or falsifying observation	Scale and alternative explanation
Seed selects a nonperiodic target	Held-out seed changes output with same bath/library	Output follows an external template or is seed-independent	Micrometre assemblies; exclude seed-independent nucleation
Local checking improves correctness	Measured wrong-state removal exceeds matched controls	Reporter changes without actual repair	Single gate to tens of sites; exclude selection/purification bias
Reliability-access window applies	Calibrated rates predict held-out final yield	Unmodelled correlated errors dominate	Size series; exclude independent-error fitting to correlated data
Redundancy-access reversal occurs	Added size lowers repair rate; renewed access restores exponent	Reliability remains monotone because depletion is negligible	Larger modules or porous confinement; exclude steric effects
Conversion errors matter	New defects appear after last accessible repair	No measurable conversion damage at assay sensitivity	Matched before/after cohorts; exclude imaging-induced damage
Material roles survive conversion	Registered composition and useful response retained	Cross-talk, shrinkage, bridging, or loss of phase quality	Nanometre interfaces; exclude label-only readout
Closure protocol preserves service	Every pending region retains measured supply and waste path	A closed neck traps incomplete material	Micrometre to millimetre channels; exclude static-distance proxies
Function follows the genome	Seed-selected device transfer function is reproduced	High shape fidelity with failed device function	Passive nanosensor or circuit; exclude external wiring as sole function

Major claim or hypothesis	Supporting observation	Failure or falsifying observation	Scale and alternative explanation
Maintenance is possible	Controlled damage is restored under a retained program	Controller lost during cure, or repair creates new failures	Repeated cycles; exclude replacement of the whole object
Claimed compression is real	Total charged description beats fair baseline across a family	Complexity moved into unique parts or bath patterns	Increasing family sizes; count interpreter and synthesis metadata

Additional hard failure modes include kinetic traps, malformed component species, aggregation without a seed, poisoning of catalytic surfaces, lack of independent redundancy, depletion of a minor reagent, manifold clogging, delamination during drying, controller degradation, and slow removal of deeply buried scaffold. Some of these invalidate entire classes of proposed objects. The compiler should return such failures as explicit unsupported-process or unavailable-access outcomes.

28. Repair after fabrication and the role of biology

A manufacturing program doubles as a maintenance program only if a sufficient local reference survives, damaged states remain distinguishable from legitimate states, repair reagents can reach them, and replacement does not destroy adjacent function. A fully mineralized object whose DNA reference has been burned away does not retain molecular self-repair merely because DNA helped fabricate it.

Three credible routes are: retain protected low-temperature reference layers; keep replaceable modules with durable port identities; or preserve ordinary service channels for externally supplied repair chemistry and control. Each trades some autonomy or material performance for repairability. A high-temperature monolithic object may have no compatible option. The maintenance problem must be specified separately for that material class.

Biology demonstrates that local interactions can distribute compressed hereditary instructions, establish differentiated structures, maintain supply networks, and repair many errors. It does not demonstrate arbitrary material synthesis or arbitrary human-selected geometries, nor atomically perfect organisms. Biological growth times, wet chemistry, available elements, mechanical properties, and evolutionary objectives differ from engineering requirements.

The useful abstraction is an acellular developmental manufacturing process with a bounded inventory and contained reactor. Replication-like copying of local program state or bounded module production does not require unrestricted reproduction. The proposed roadmap does not rely on evolving or releasing autonomous organisms.

29. What would be needed from 100 nm to 1 cm

At approximately 100 nm, focus on one local checking gate, reliable role recruitment, and a single conversion-compatible structural unit. At about 1 micrometre, establish seeded nonperiodic patterns, error correlations, and a weak but measurable function. Moving to 100 micrometres changes the dominant questions: transport, accumulated strain, defect domains, and compatibility between modules become difficult even though no atom-by-atom placement is required.

At 1 mm, a successful system needs engineered distributed supply, waste exit, thermal management, measured effective controller error rates, and interfaces that tolerate dimensional variation. At 1 cm, those systems must handle millions of modules or extremely high unit counts with controlled common-mode failure and material conversion that preserves useful properties.

Standardized ports and prefabricated critical components become more plausible than universal in-situ chemistry.

The shortest credible route is therefore **small validated molecular logic, followed by a limited material process, followed by a modular function, followed by transport-enabled scale-up**. Choosing a family whose useful function tolerates distributed microscopic defects is much more credible than starting with atomically perfect dense electronics. This is an engineering restriction with information-theoretic and reliability benefits, not an assertion that arbitrary matter becomes easy.

30. Research programme at 1, 3, 5, 10, and 20 years

These are contingent milestones measured from project start, not forecasts that the technology will be available on these dates.

Horizon	Realistic research target if preceding gates succeed	Ambitious conditional extension	Stop or redirect when
1 year	Calibrated local check-and-lock gate, two seed programs, reproducible error model	First role-selective recruitment with held-out prediction	No useful local selectivity or controller stability
3 years	Repeated differentiated modules and one qualified transduction process	Direct test of redundancy/access coupling	Conversion faults dominate all tolerated designs
5 years	Restricted functional optical/passive module family	Autonomous local closure in a small porous assembly	Cross-talk and correlated process failures resist mitigation
10 years	Integrated prototype with measured transport, strain, and yield scaling	Sub-millimetre to millimetre developmental fabrication in one material family	No favourable total-cost or throughput comparison
20 years	Potential restricted manufacturing platform with reusable seed/compiler interface	Centimetre heterogeneous modules assembled from validated parts	General-purpose chemistry remains inaccessible or conventional manufacture dominates

An arbitrary universal nanofabricator is a long-term speculative extension, not a 20-year commitment. Many early results would still be useful if that extension never becomes feasible: better local defect assays, error-aware material conversion, transport-aware self-assembly, and compilers that expose unsupported chemistry.

31. Ranked unresolved barriers

1. **Faults after the final repairable state.** Demonstrate a real material process with a bounded final fault channel and useful functional tolerance or post-conversion repair.

2. **Reliable local references.** Protect seed, lineage, and role state against common faults without assuming an ideal target oracle.

3. **Correlation and functional coding.** Establish a local error law and a device-specific tolerated defect set; independent site models are inadequate by themselves.

4. **Transport and late closure.** Grow or install a scalable network whose concentration, flow, permeability, and waste removal are measured through conversion.

5. **A practical molecular compiler.** Lower a modest useful grammar into sequence-validated, reusable operations with bounded leak and acceptable speed.

6. **Mechanical preservation.** Control drying, densification, stress, and interface placement over successive scales.

7. **Fair resource advantage.** Compare total descriptions, synthesis effort, bath processing, throughput, yield, energy, and apparatus against existing manufacturing.

The most consequential missing experiment is not a complicated shape. It is a substrate on which independent calibration of checking, access, and conversion predicts a nontrivial useful final function on a new seed program.

32. Adversarial audit and claim boundary

The global design is stored in the seed, motif interpreter, qualified component library, and every target-specific external condition. All are supplied or identified. The simulator's short JSON capsules rely on handwritten motif code and an ideal local interpreter. Finite chemical alphabet universality is conditional; a molecular backend and a minimal practical species count remain unknown.

Error accumulation is handled only under explicit local fault and functional-tolerance contracts. The last-operation theorem prevents pretending that pre-conversion proofreading establishes final perfection. The shared-reference ablation demonstrates one checker blind spot but does not cover all spatial correlations. The actual circuit benchmark mostly fails its function criterion. Those limitations are retained in the reported results.

Diffusion, mass flux, heat, pressure, shrinkage, and interior closure are explicit parts of the framework. Internal reservoirs in the simulator are supplied by an idealized external apparatus, so the simulations do not establish autonomous vascular growth. The stress experiment shows that the chosen access geometry still fails to finish under stronger demand.

Heterogeneous material formation uses qualified recruitment and conversion, not element transformation or arbitrary semiconductor chemistry. Bulk precision requirements can be reduced by modular design, but critical interfaces remain difficult. Growth concurrency improves plausible scaling only when instructions, reaction fronts, material flow, and cooling all support it.

The strongest defensible original result is the **conditional redundancy-access reversal and its associated compiler feasibility window**. It is a derivation built from established models. Expert confidence would require independent proof review, a more comprehensive prior-art comparison, calibrated correlated-noise simulations, and the decisive checking/locking experiment. The result should not be advertised as demonstrated universal manufacture or a verified major breakthrough.

References

[1] D. Soloveichik and E. Winfree. Complexity of Self-Assembled Shapes. SIAM Journal on Computing 36, 1544-1569 (2007). DOI: 10.1137/S0097539704446712. Accessible manuscript: https://arxiv.org/abs/cs/0412096

[2] D. Doty, J. H. Lutz, M. J. Patitz, R. T. Schweller, S. M. Summers, and D. Woods. The tile assembly model is intrinsically universal. FOCS (2012). Accessible manuscript: https://arxiv.org/abs/1111.3097

[3] E. Winfree and R. Bekbolatov. Proofreading Tile Sets: Error Correction for Algorithmic Self-Assembly. DNA Computing, LNCS 2943, 126-144 (2004; conference 2003). https://doi.org/10.1007/978-3-540-24628-2_13

[4] D. Woods, D. Doty, C. Myhrvold, J. Hui, F. Zhou, P. Yin, and E. Winfree. Diverse and robust molecular algorithms using reprogrammable DNA self-assembly. Nature 567, 366-372 (2019). https://www.nature.com/articles/s41586-019-1014-9

[5] T. Gomez, E. Grizzell, A. Haun, R. Knobel, T. Peters, R. Schweller, and T. Wylie. Intrinsic Universality in Seeded Active Tile Self-Assembly. arXiv:2407.11545 (2024); brief announcement, SAND 2025, DOI: 10.4230/LIPIcs.SAND.2025.24. https://arxiv.org/abs/2407.11545

[6] M. Holmes-Cerfon and M. Wyart. Hierarchical self-assembly for high-yield addressable complexity at fixed conditions. arXiv:2501.02611 (2025). https://arxiv.org/abs/2501.02611

[7] F. Benoist and P. Sartori. Assembly Factors Resolve Speed and Encoding Bottlenecks in Multifarious Self-Assembly. PRX Life 4, 033021 (2026). https://journals.aps.org/prxlife/abstract/10.1103/5w6h-l93l

[8] A. Michelson et al. Three-dimensional nanoscale metal, metal oxide, and semiconductor frameworks through DNA-programmable assembly and templating. Science Advances 10, eadl0604 (2024). DOI: 10.1126/sciadv.adl0604. Author-laboratory hosted first page: https://www.bnl.gov/tcp/uploads/files/2024-006j-pg1.pdf

[9] A. Michelson et al. Scalable fabrication of Chip-integrated 3D-nanostructured electronic devices via DNA-programmable assembly. Science Advances 11, eadt5620 (2025). https://doi.org/10.1126/sciadv.adt5620

[10] Brookhaven National Laboratory, Center for Functional Nanomaterials. Scalable Device Fabrication using DNA Self-Assembly. Research highlight, 22 July 2025. Author-laboratory account of [9]: https://www.bnl.gov/cfn/research/highlights/news.php?a=222678

[11] F. Teng et al. Macroscale-area patterning of three-dimensional DNA-programmable frameworks. Nature Communications 16, 3238 (2025). https://www.nature.com/articles/s41467-025-58422-0

[12] M. Nguyen and S. Vaikuntanathan. Design principles for nonequilibrium self-assembly. Proceedings of the National Academy of Sciences 113, 14231-14236 (2016). https://doi.org/10.1073/pnas.1609983113

Final assessment

The following percentages are subjective evidence-maturity estimates, rounded to five percentage points. They are not experimental measurements, probabilities of success, peer-review scores, or a percentage solution of unrestricted universal nanofabrication. The subject is the bounded theory and proposed implementation programme developed here.

Status: rigorous theoretical framework.

Scientific completeness: 35%. Core hypotheses, falsifiers, and limiting models are specified; calibrated multi-physics evidence and validated material functions are missing.

Mathematical completeness: 60%. Restricted stated propositions have proofs; realistic correlated-noise thresholds, molecular local compilation, arbitrary-device tolerance, and efficient physical universality remain unresolved.

Experimental readiness: 20%. There is a staged decisive experiment and reproducible simulation, but no sequence-level implementation, calibrated reaction network, or new laboratory data.

Physical plausibility: 55%. Several component mechanisms have primary experimental precedents; their integrated, self-correcting, multimaterial operation remains unverified, especially through hardening and scale-up.

Potential impact if validated: 90%. This is a conditional impact judgement for a successful generalization across useful material families, not an estimate of validation probability or present novelty.

Most important unresolved obstacle: preserving a trustworthy repair mechanism, or a validated fault-tolerant function, through irreversible material conversion.

Single most decisive next experiment: measure held-out post-recruitment error in the same seeded two-role assembly while independently varying local checking and the available repair interval before locking.

Single most important new theoretical result: the conditional redundancy-access reversal: redundant volume can reduce the reliability exponent when it increases diffusion-limited depletion, while maintained local access restores the logarithmic redundancy guarantee under the stated noise and functional-tolerance assumptions.