Active Pruning Controls Boundary Persistence in Protocell-like Systems: Saddle-Node Attractor Loss, Stochastic Early Warnings, Multiscale Simulations, and Wet-Lab Benchmark Predictions

Yining Wu
Independent Researcher
yining.wu@alumni.upenn.edu

(Dated: May 15, 2026)

Finite protocell-like systems under sustained flux must maintain boundary integrity, transport throughput, and internal functional availability while accumulating nonfunctional residue. Here we develop an analytical and computational framework for testing whether active pruning - selective, resource-coupled removal, export, degradation, or reorganization of residue - controls the persistence-collapse transition in finite compartments. The framework combines five layers: (i) a well-mixed residue-pruning-boundary model with explicit threshold bounds; (ii) a saddle-node normal-form reduction showing how a maintenance attractor can disappear when pruning falls below a critical value; (iii) a stochastic perturbation analysis deriving critical slowing down, variance growth, and lag-one autocorrelation increase near attractor loss; (iv) two spatial failure modes, reaction-diffusion clogging for droplet-like or porous compartments and phase-field-inspired reduced boundary deformation for vesicle-like compartments; and (v) a radial-shell size constraint showing that boundary-limited pruning obeys the areal lower bound $s_{m,c}(R) \geq qR/3$, with diffusion delay raising the effective numerical threshold above this mass-balance limit. Numerical simulations reproduce active-pruning thresholds, spatial clogging, residue-driven boundary deformation with local integrity loss, rescue-window closure, radius-dependent pruning demand, and robustness across randomized production and diffusion-arrest parameters. The paper does not claim wet-lab validation or a chemically complete origin-of-life mechanism. Its contribution is a chemically coarse-grained universality-class model and benchmark map: finite compartments satisfying these assumptions—flux generates residue, residue impairs transport or boundary integrity, active removal requires finite resources—are predicted to exhibit pruning-dependent persistence, collapse, and rescue thresholds. This is a theory-computation manuscript: the wet-lab section provides experimentally measurable benchmark predictions, not completed experimental validation.

I. INTRODUCTION

Replication, autocatalysis, compartmentalization, and energy transduction are central to origin-of-life research. Yet a finite compartment under sustained flux faces a maintenance problem that is logically prior to any specific molecular solution: useful flux also generates nonfunctional residue. Residue may be inert byproduct, polymerized tar, aggregate, blocked catalytic site, membrane-clogging material, damaged component, or unavailable internal volume. If residue accumulates faster than it can be diluted, degraded, exported, recycled, or reorganized, the system may lose transport, boundary integrity, and functional throughput even while energy continues to enter.

This paper studies *active pruning* as a control process for that problem. Active pruning denotes selective, resource-coupled removal, export, neutralization, degradation, or reorganization of residue while preserving the functional core of a finite boundary-maintaining system. The term is deliberately substrate-neutral. In a lipid vesicle, the relevant process may be export, leakage control, membrane-localized degradation, or recycling. In a coacervate droplet, it may be catalytic degradation or phase-selective residue removal. In a mineral pore, it may be surface refresh, desorption, or flow-mediated clearing. The common structure is not a particular molecule, but a residue-pruning-boundary loop.

The central claim tested here is narrower than a general theory of life: in finite protocell-like systems under sustained flux and residue production, active pruning can act as a control parameter separating a boundary-maintaining nonequilibrium regime from residue-induced collapse. The present paper extends earlier minimal simulations by adding analytical thresholds, a saddle-node reduction for attractor loss, stochastic early-warning analysis, multiscale spatial mechanisms, size-scaling lower bounds, robustness sweeps, and a benchmark map for future wet-lab tests.

A. Universality-class framing

The model is chemically coarse-grained by design. Its claim is not that all protocells share one molecular pruning mechanism. Rather, it defines a universality class: any finite compartment in which (i) flux generates residue, (ii) residue impairs diffusion, transport, or boundary integrity, and (iii) residue removal is active, selective, and resource-limited should exhibit a pruning-dependent persistence-collapse transition. This is analogous in spirit to using coarse-grained models to study phase transitions: the variables abstract away molecular detail to expose control parameters and scaling laws.

Recent protocell research has increasingly moved beyond passive compartmentalization toward chemically programmable coacervate systems, enzymatic microreactors, membranized coacervates with tunable interfacial stability, and active droplets in synthetic cells [15–19]. These developments make residue control, molecular transport, boundary integrity, and active nonequilibrium maintenance natural variables for a coarse-grained theory of protocell persistence.

B. Contributions

This paper makes nine contributions.

1. It formulates a residue-pruning-boundary model and derives an explicit target-load lower bound for a critical pruning rate S_c.
2. It gives a Jacobian-based saddle-node condition for the well-mixed reduction and shows how the local normal form produces maintenance-attractor loss.
3. It adds a stochastic differential-equation reduction and derives early-warning scaling: recovery time and variance increase while lag-one autocorrelation tends to one as the fold is approached.
4. It nondimensionalizes the theory into a small set of control numbers: pruning number, residue burden, diffusion-delay number, osmotic-stress number, and boundary-repair number.
5. It simulates spatial clogging in a two-dimensional reaction-diffusion model with residue-dependent diffusivity $D_\rho(\rho) = D_0 e^{-\lambda_D \rho}$.
6. It simulates osmotic/crowding-driven boundary deformation using a phase-field-inspired reduced radius-field model rather than claiming a full phase-field membrane solver.
7. It derives the three-dimensional boundary-limited areal lower bound $s_{m,c}(R) \geq qR/3$ and compares it with a radial-shell numerical threshold including diffusion delay.
8. It quantifies rescue-window closure and robustness under randomized parameter sweeps.
9. It provides a wet-lab benchmark map and a literature-informed plausibility check without claiming that wet-lab validation has already been performed.

C. What is not claimed

The paper does not identify the first historical pruning chemistry, does not reproduce a specific protocell, and does not provide experimental validation. It does not define clinical death. It does not claim that collapse in every spatial PDE must be mathematically a saddle-node. The saddle-node analysis applies to a reduced well-mixed maintenance coordinate; spatial simulations then show how this reduced attractor-loss mechanism can be physically realized through clogging, boundary deformation, and size-limited export.

D. Falsifiable predictions

The following four conditions would falsify or demote specific claims in this paper:

1. **Sustained pruning threshold.** If systems under sustained residue production maintain bounded ρ, stable B, and stable J with $S = 0$, no hidden dilution, no externalization, and no passive degradation sufficient to offset production, the active-pruning-as-control-parameter claim is falsified.
2. **Rescue window closure.** If the rescue window does not close as restoration delay increases (i.e., no rescue boundary exists), the saddle-node attractor-loss reduction is demoted to a non-universal simulation observation.
3. **Radius-dependent pruning demand.** If the critical areal pruning/export threshold $s_{m,c}(R)$ does not increase with compartment radius R under boundary-limited clearing, the areal lower bound $s_{m,c}(R) \geq qR/3$ and its diffusion-delay correction are demoted.
4. **Early-warning signals.** If recovery time and lag-one autocorrelation fail to increase near the persistence-collapse transition under quasi-static approach to the fold, the stochastic early-warning analysis is demoted; variance is treated as an auxiliary indicator.

II. RESIDUE-PRUNING-BOUNDARY MODEL

A. Core variables

The coarse-grained state is

$$\mathcal{X}(t) = \{\rho(t), B(t), J(t), S(t), V(t)\}, \tag{1}$$

where ρ is residue load, $B \in [0, 1]$ is boundary integrity, J is transport or reaction throughput, S is active pruning rate, and V is a viability-like maintenance score. A convenient normalized score is

$$V(t) = \frac{B(t)J(t)}{1 + \rho(t)/\rho_0}. \tag{2}$$

Collapse is operationally defined by $V(t) < V_{\min}$ for a specified duration or, in rescue experiments, by failure to recover above V_{rec} after restoration of pruning within an observation window.

B. Well-mixed dynamics

The minimal well-mixed dynamics are

$$\frac{d\rho}{dt} = f(\rho, B; S) = pJ_0 B e^{-\lambda_\rho \rho} - d\rho - S\frac{\rho}{K+\rho}, \quad (3)$$

$$\frac{dB}{dt} = g(\rho, B) = r_B B(1-B) - \beta_\rho \rho B + \gamma_B (1-B), \quad (4)$$

$$J(\rho, B) = J_0 B e^{-\lambda_\rho \rho}. \quad (5)$$

Here p is residue generated per unit throughput, d is passive degradation or dilution, K is the half-saturation constant for pruning, β_ρ is residue-induced boundary damage, γ_B is repair, and λ_ρ controls transport inhibition by residue. Resource coupling is coarse-grained into the pruning parameters: S represents the maximal clearing rate set by available resources, and K sets the residue load at which clearing reaches half its maximum, capturing saturation under resource or throughput constraints. The exponential throughput term is a smooth representation of crowding or clogging; replacing it by another decreasing function does not change the threshold logic.

Proposition 1 (Pruning lower bound) *Fix residue ceiling $\rho_c > 0$ and target throughput J^*. If ρ is to remain locally non-increasing at ρ_c, then active pruning must satisfy*

$$S \geq S_c(\rho_c) = (pJ^* - d\rho_c)_+ \frac{K+\rho_c}{\rho_c}, \quad (6)$$

where $(x)_+ = \max(x, 0)$.

Proof. At $\rho = \rho_c$, requiring $\dot{\rho} \leq 0$ gives $pJ^* - d\rho_c - S\rho_c/(K+\rho_c) \leq 0$. Solving for S yields Eq. (6). □

This bound is intentionally conservative. It does not prove global persistence, but it gives a direct analytical threshold that the simulations can test.

III. ANALYTICAL STRUCTURE OF ATTRACTOR LOSS

A. Jacobian and saddle-node condition

Let $z = (\rho, B)^T$ and $F(z; S) = (f, g)^T$. The Jacobian of Eqs. (3)-(4) is

$$\begin{aligned} J_{11} &= -pJ_0\lambda_\rho B e^{-\lambda_\rho \rho} - d - \frac{SK}{(K+\rho)^2}, \\ J_{12} &= pJ_0 e^{-\lambda_\rho \rho}, \\ J_{21} &= -\beta_\rho B, \\ J_{22} &= r_B(1-2B) - \beta_\rho \rho - \gamma_B, \end{aligned} \quad (7)$$

At an equilibrium $z^*(S)$ satisfying $F(z^*; S) = 0$, a local saddle-node candidate occurs when

$$\det J_z(z^*; S_c) = 0, \qquad \operatorname{tr} J_z(z^*; S_c) < 0, \quad (8)$$

with the usual non-degeneracy conditions

$$w^T F_S(z^*; S_c) \neq 0, \qquad w^T D^2 F(z^*; S_c)[v, v] \neq 0, \quad (9)$$

where v and w are the right and left null vectors of J_z at the fold. Under these conditions, center-manifold reduction gives the normal form

$$\dot{x} = a(S - S_c) - bx^2 + O(|x|^3, |S - S_c||x|), \quad (10)$$

with $ab > 0$ after orientation choice.

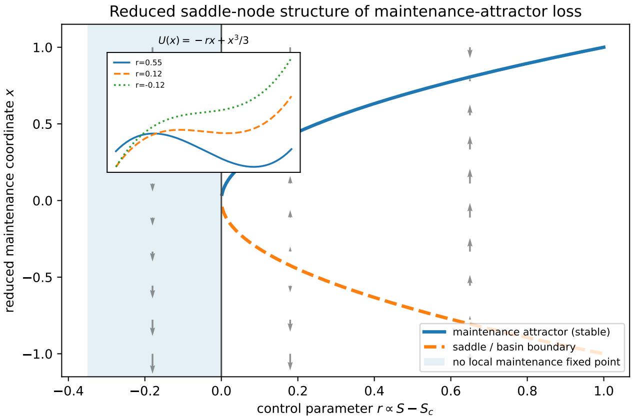

FIG. 1. Reduced saddle-node structure. The stable branch represents a maintenance attractor; the dashed branch represents the saddle or basin boundary. Arrows show the one-dimensional flow of $\dot{x} = r - x^2$. The inset shows the potential landscape $U(x) = -rx + x^3/3$, whose local minimum disappears at the fold.

Proposition 2 (Maintenance-attractor loss) *For parameters satisfying Eqs. (8) and (9), the well-mixed residue-pruning-boundary model admits a saddle-node loss of the maintenance attractor. For $S > S_c$, a stable maintenance branch and an unstable basin-boundary branch coexist locally. At $S = S_c$ they collide. For $S < S_c$, the local maintenance fixed point is absent and trajectories are forced toward the collapse basin unless rescued by an external nonlocal intervention.*

Figure 1 illustrates the reduced normal form, including the associated potential landscape $U(x) = -rx + x^3/3$ for $\dot{x} = -\partial_x U$. This is the analytical meaning of maintenance-collapse in the model: loss of access to a maintenance attractor in a finite artificial compartment.

B. Stochastic early-warning scaling

Near the fold, stochastic perturbations can be represented by the one-dimensional SDE

$$dx = (r - x^2)dt + \sigma dW_t, \quad (11)$$

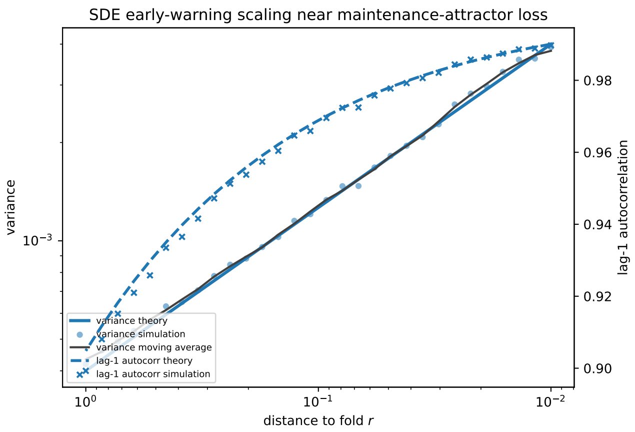

FIG. 2. SDE early-warning scaling. Theory and stochastic simulations of the normal-form reduction show increasing variance and lag-one autocorrelation as the distance to the fold r decreases. Variance points are shown with a moving average because variance is more sensitive to sample window and noise model than autocorrelation. The same eigenvalue gives the recovery-time divergence $\tau_{\mathrm{rec}} \sim (2\sqrt{r})^{-1}$.

where $r > 0$ is distance to the fold. The stable fixed point is $x^* = \sqrt{r}$ and the linearized eigenvalue is

$$\lambda = f'(x^*) = -2\sqrt{r}. \tag{12}$$

Therefore the recovery time scales as

$$\tau_{\mathrm{rec}} \sim \frac{1}{2\sqrt{r}}, \tag{13}$$

while the Ornstein-Uhlenbeck variance near the stable point is

$$\mathrm{Var}(x) \sim \frac{\sigma^2}{4\sqrt{r}}. \tag{14}$$

For a sampling interval Δt, the lag-one autocorrelation satisfies

$$AC(1) \approx \exp(-\Delta t/\tau_{\mathrm{rec}}) \to 1 \quad \text{as} \quad r \to 0^+. \tag{15}$$

Thus variance, autocorrelation, and recovery time are expected to rise near rescue-window closure. Variance is the most noise-model-sensitive of the three; in the simulations below we use smoothed stochastic estimates and interpret variance as an auxiliary indicator rather than a sole diagnostic. Early-warning indicators are interpreted here as local normal-form predictions, not universal diagnostics; recent work emphasizes both the usefulness and ambiguity of EWS outside ideal bifurcation settings [22].

IV. NONDIMENSIONAL CONTROL NUMBERS

The model can be summarized by a small set of dimensionless numbers (Table I). These quantities make explicit which parameters should be varied in simulations and wet-lab analogues.

The pruning number $\mathcal{E} = S_0/(pJ_0)$ and residue burden $\mathcal{P} = pJ_0/S_0$ are reciprocal under the same normalization. We report both because they emphasize opposite experimental readings: $\mathcal{E}$ measures excess active clearing capacity, whereas $\mathcal{P}$ measures residue pressure relative to that capacity.

Persistence requires $\mathcal{E}$ sufficiently large, $\mathcal{D}$ sufficiently small, $\mathcal{O}$ below rupture scale, and $\mathcal{B}$ large enough to prevent boundary damage from accumulating. These are not independent in a real chemistry, but separating them clarifies which experimental knobs correspond to which failure modes.

V. SPATIAL FAILURE MODE I: REACTION-DIFFUSION CLOGGING

Droplet-like, gel-like, and porous compartments may fail by internal clogging before any membrane-like rupture occurs. We model this by a residue field obeying

$$\frac{\partial \rho}{\partial t} = pJ + D_\rho(\rho)\nabla^2 \rho - S(x,t)\frac{\rho}{K+\rho}, \tag{16}$$

with residue-dependent diffusivity

$$D_\rho(\rho) = D_0 e^{-\lambda_D \rho}. \tag{17}$$

The local throughput is

$$J(x,t) = J_0 B(x,t) e^{-\lambda_J \rho(x,t)}, \tag{18}$$

while local integrity evolves as

$$\frac{\partial B}{\partial t} = \gamma_B(1-B) - \beta_\rho \rho B. \tag{19}$$

This mechanism is relevant to coacervates, gels, and mineral pores: residue slows diffusion, reduced diffusion increases local residue retention, and the resulting feedback creates dead zones.

Figure 3 shows the qualitative result. Without pruning, residue accumulates and spatially traps itself. With active pruning, residue remains bounded and transport remains available.

VI. SPATIAL FAILURE MODE II: PHASE-FIELD-INSPIRED REDUCED DYNAMIC-BOUNDARY FAILURE

Vesicle-like systems may fail through boundary stress and local integrity loss. In this reduced model, geometric shape change and integrity failure are distinct observables: two compartments may have similar gross outlines while differing strongly in local leakage or damage susceptibility. A full phase-field treatment would evolve an

Symbol	Definition	Meaning
$\mathcal{P}$	pJ_0/S_0	residue burden / pruning
$\mathcal{D}$	$(R^2/D_0)/(K/S_0)$	diffusion delay / pruning time
$\mathcal{O}$	$\alpha_\Pi\rho_0/\sigma_{\mathrm{crit}}$	osmotic stress / rupture threshold
$\mathcal{B}$	$\gamma_B/(\beta_\rho\rho_0)$	repair / residue damage
$\mathcal{E}$	$S_0/(pJ_0)$	pruning excess capacity

TABLE I. Dimensionless control numbers for residue-pruning-boundary systems.

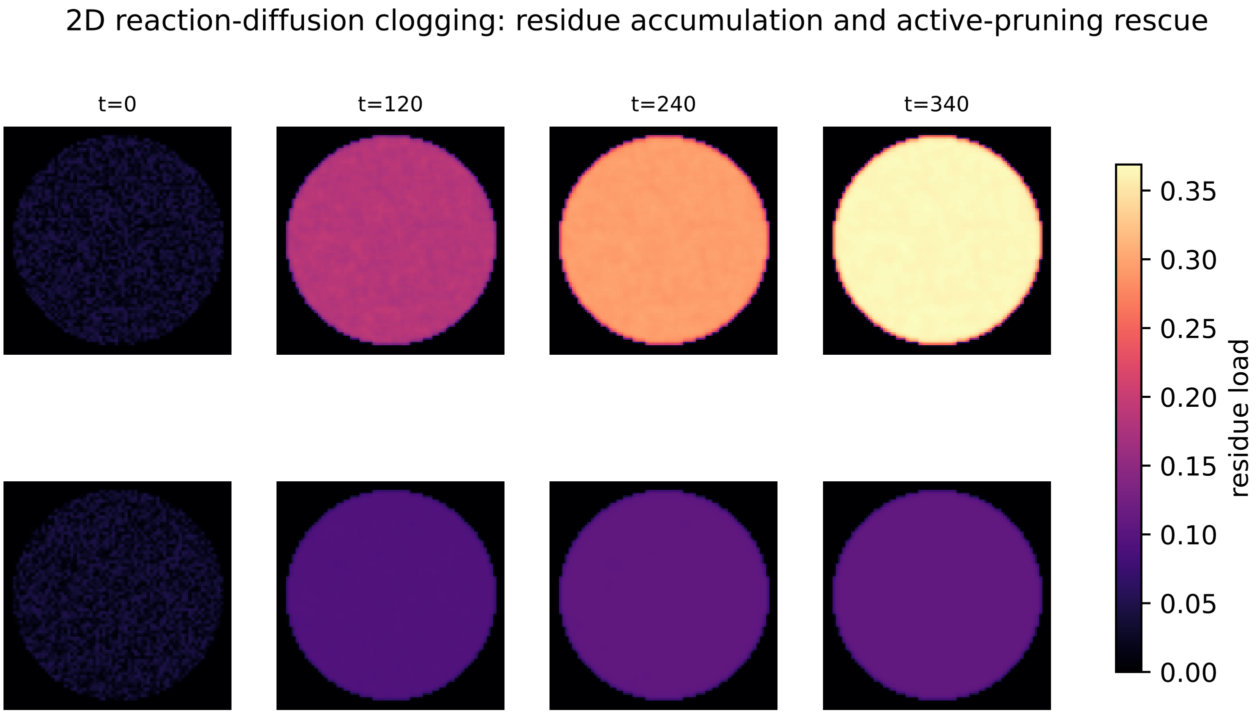

FIG. 3. Spatial failure mode I: reaction-diffusion clogging. The upper row shows residue accumulation without active pruning. As ρ increases, $D_\rho(\rho) = D_0 e^{-\lambda_D \rho}$ falls, producing diffusion arrest and dead zones. The lower row shows the same system with active pruning above threshold, where residue remains bounded.

order parameter $\phi(x,t)$ with a free-energy functional such as

$$\mathcal{F}[\phi,\rho] = \int \left[\frac{\kappa}{2}|\nabla\phi|^2 + W(\phi) + \chi\rho\phi\right] dx + \lambda_A(A - A_0)^2 + \lambda_V(V - V_0)^2. \tag{20}$$

with gradient-flow dynamics

$$\frac{\partial\phi}{\partial t} = -M_\phi \frac{\delta\mathcal{F}}{\delta\phi}. \tag{21}$$

We do *not* solve this full phase-field problem here. Instead, we use a phase-field-inspired reduced boundary radius field $R(\theta,t)$ together with a local integrity field $B(\theta,t)$ to isolate the coupling between residue pressure, boundary deformation, and rupture-like integrity loss. The purpose is not to resolve full membrane hydrodynamics, but to test whether residue pressure can drive loss of boundary maintenance in a vesicle-like geometry.

The reduced model is

$$\frac{d\rho}{dt} = pJ - d\rho - S\frac{\rho}{K+\rho}, \tag{22}$$

$$\Pi(\rho) = \alpha_\Pi\rho, \tag{23}$$

$$\frac{\partial R}{\partial t} = \mu_R\left[R_0 + \Pi(\rho)(1 + \eta(\theta,t)) - R\right] + D_R\frac{\partial^2 R}{\partial\theta^2}, \tag{24}$$

$$\frac{\partial B}{\partial t} = \gamma_B(1 - B) - \beta_\epsilon\epsilon(\theta,t)B - \beta_\rho\rho B, \tag{25}$$

where ϵ is a local strain proxy. Rupture-like failure occurs when mean boundary integrity falls below threshold or when strain exceeds a prescribed damage scale.

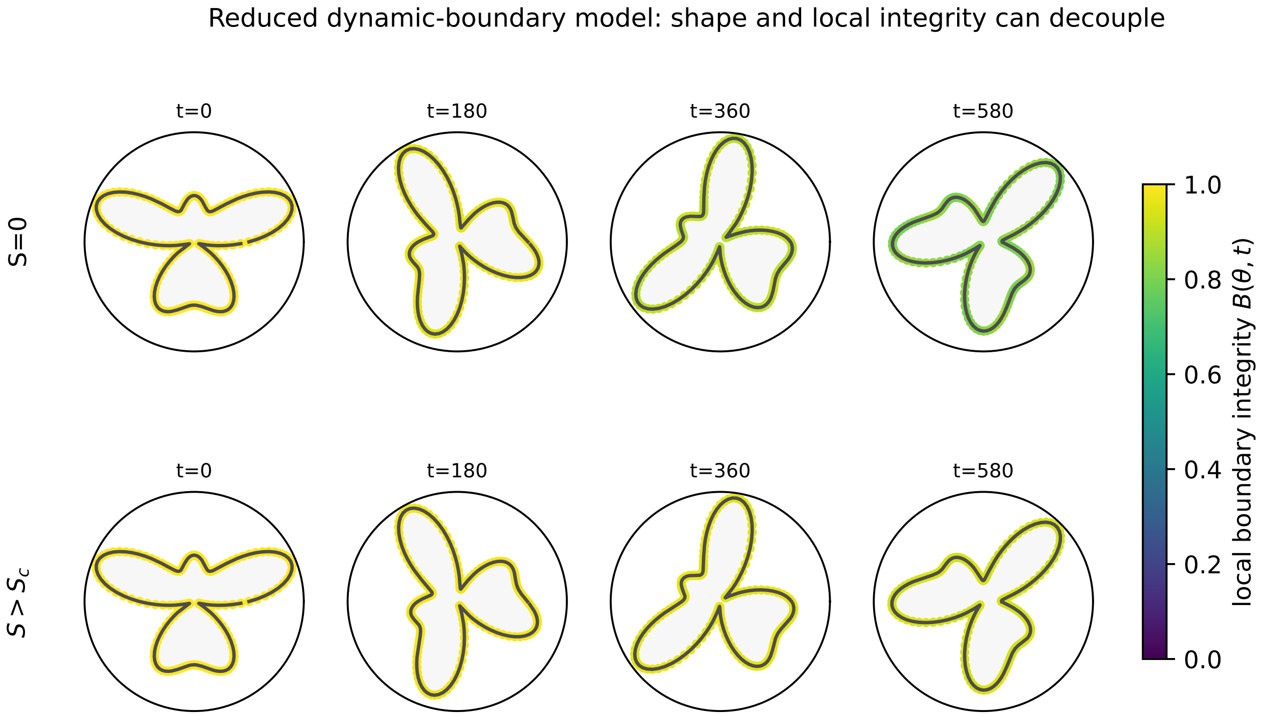

FIG. 4. Spatial failure mode II: phase-field-inspired reduced dynamic-boundary failure. Grey outlines show the boundary radius field $R(\theta, t)$; color along the boundary shows local integrity $B(\theta, t)$. The dominant failure signal is not necessarily gross shape change but loss of local integrity. Thus similar outlines can correspond to very different leakage or damage states. The model uses a reduced radius/integrity field rather than a full Cahn-Hilliard, Allen-Cahn, Helfrich, or hydrodynamic membrane solver.

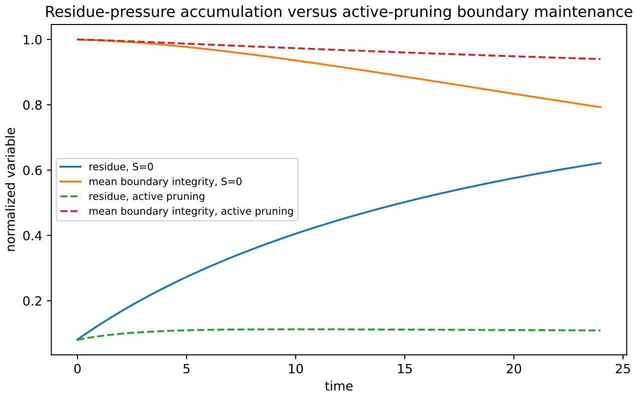

FIG. 5. Dynamic-boundary time series. In the no-pruning condition, residue increases and mean boundary integrity declines. In the active-pruning condition, residue and integrity approach a maintained nonequilibrium regime. This explains why Fig. 4 emphasizes integrity color in addition to geometric outline.

VII. THREE-DIMENSIONAL SIZE CONSTRAINT

A spherical compartment of radius R with volume residue-generation density q produces residue at rate

$$P_\rho = \frac{4}{3}\pi R^3 q. \tag{26}$$

If active export or membrane-localized pruning has maximum areal capacity s_m, the boundary clears at most

$$P_{\text{clear}} = 4\pi R^2 s_m. \tag{27}$$

A necessary condition for steady maintenance is $P_{\text{clear}} > P_\rho$, which yields the areal mass-balance lower bound

$$\boxed{s_{m,c}(R) \geq \frac{qR}{3}.} \tag{28}$$

The corresponding total membrane pruning requirement is $S_c^{\text{tot}}(R) = 4\pi R^2 s_{m,c}(R)$. Equation (28) is only a lower bound. If residue must diffuse from the interior to the boundary before export, then the diffusion delay grows as

$$t_D \sim \frac{R^2}{D_\rho}. \tag{29}$$

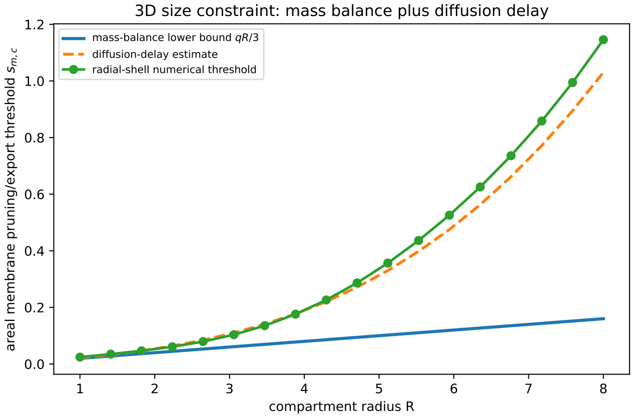

FIG. 6. Three-dimensional size constraint. Boundary-limited clearing gives the analytical areal lower bound $s_{m,c}(R) \geq qR/3$. The radial-shell numerical threshold lies well above this bound at large R because residue must diffuse to the boundary before export or pruning. The dashed curve shows the corrected estimate $s_{m,c}^{\mathrm{eff}}(R) \gtrsim (qR/3)[1+\chi R^2/(D_0\tau_p)]$, where the dimensionless factor $\chi R^2/(D_0\tau_p)$ represents the diffusion-delay penalty.

A minimal corrected estimate is therefore

$$s_{m,c}^{\mathrm{eff}}(R) \gtrsim \frac{qR}{3}\left[1+\chi\frac{R^2}{D_0\tau_p}\right], \tag{30}$$

where τ_p is the pruning timescale and χ is a geometry- and transport-dependent coefficient. If D_ρ decreases with residue load, the effective threshold can exceed the mass-balance lower bound by a large factor at large R. The radial-shell simulations implement this diffusion-delay correction. Figure 6 shows the lower bound, the corrected estimate, and the numerical threshold.

The growing separation between the mass-balance lower bound and the radial-shell threshold is therefore not a discrepancy but a mechanistic prediction: large compartments are not only surface-limited; they are transport-delay limited. A vesicle or droplet can have enough boundary area in principle and still fail if residue generated in the interior cannot reach that boundary on the pruning timescale.

VIII. RESCUE-WINDOW CLOSURE

To test reversibility, we use a pruning-interruption protocol: first equilibrate a maintained system with $S > S_c$, shut pruning off, restore pruning at time t_{restore}, and record whether $V(t)$ recovers above V_{rec} within the observation window. Figure 7 shows the resulting rescue surface. Early restoration succeeds over a broad range of S_{rescue}. Late restoration requires increasingly high pruning and eventually fails because the trajectory has crossed into the collapse basin.

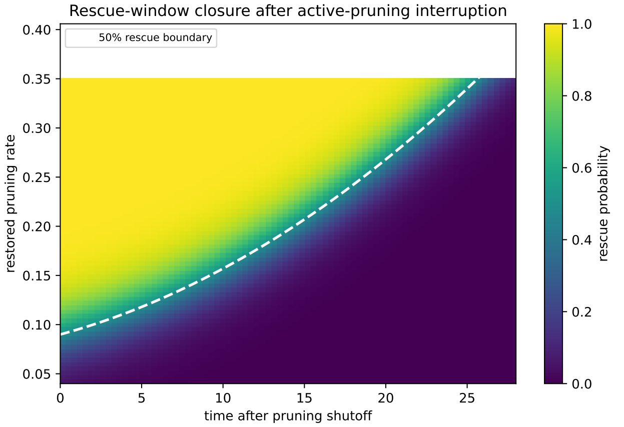

FIG. 7. Rescue-window closure. The heatmap shows rescue probability as a function of restoration time and restored pruning rate in the reduced rescue protocol. The dashed line marks the 50% rescue boundary.

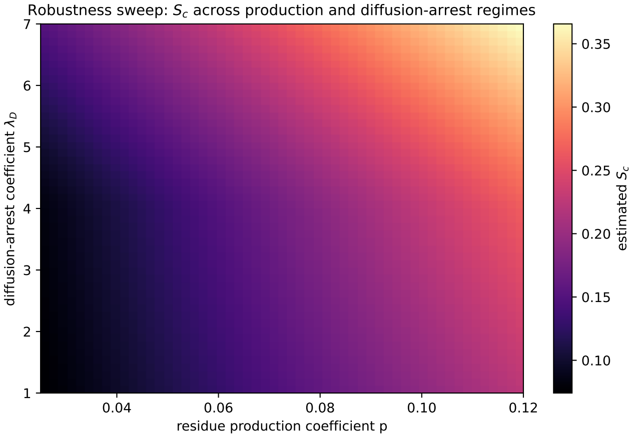

FIG. 8. Robustness sweep. Estimated pruning threshold S_c across residue production p and diffusion-arrest coefficient λ_D. The threshold persists across the explored region and rises as production and diffusion arrest increase.

IX. ROBUSTNESS SWEEP

A threshold that exists only for one parameter set is not a useful principle. We therefore varied residue production p and diffusion-arrest strength λ_D over broad ranges and estimated the pruning threshold. Figure 8 shows that S_c persists and increases with both production load and diffusion arrest. This does not prove universality, but it supports the claim that the threshold is not an artifact of one visually selected simulation.

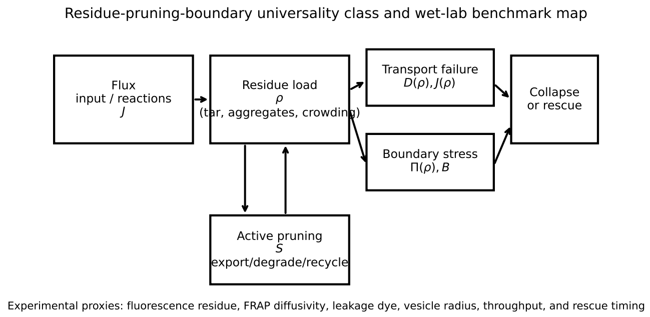

FIG. 9. Residue-pruning-boundary universality class and wet-lab benchmark map. The model predicts that flux-driven residue causes either transport failure or boundary stress unless an active pruning module removes, exports, degrades, or reorganizes residue.

X. PROSPECTIVE WET-LAB BENCHMARK PREDICTIONS

The proposed proxies are aligned with recent experimental platforms in protocell and synthetic-cell research, including membranized coacervates with tunable interfacial stability, spontaneous protocell membrane formation, and giant unilamellar vesicle stability and loading systems [20, 21]. This section is a prospective benchmark plan, not a report of completed experiments. The model should be judged by whether it suggests measurable experimental contrasts. Figure 9 and Table II map variables to possible wet-lab proxies. The strongest experiment is a paired comparison: two otherwise identical compartments under sustained flux, one without an active clearing module and one with a tunable degradation/export/recycling process. The predicted observables are bounded residue, preserved throughput, lower leakage or rupture rate, and a measurable rescue boundary in the active-pruning condition.

A. Literature-informed plausibility check

Existing protocell and soft-matter literature already shows that the ingredients assumed here are experimentally meaningful, although not yet in the combined threshold test proposed above. Fatty-acid vesicles can grow, exchange amphiphiles, and undergo permeability changes in response to osmotic and compositional conditions [1–3]. Membrane deformation and rupture are standard problems in membrane mechanics [4, 5]. Coacervates and condensate-like droplets exhibit strong changes in internal mobility, partitioning, and material state with composition and crowding [8, 9]. Early-warning indicators such as critical slowing down, variance increase, and autocorrelation increase are standard signatures near tipping points in stochastic dynamical systems [10, 11]. These observations are qualitatively consistent with the modeling ingredients used here, but they are not a quantitative validation. The proposed benchmark is intended to convert this qualitative plausibility into direct tests of S_c, $s_{m,c}(R)$, R_c, and rescue-window closure. We do not fit these data here; a future validation step should use reported vesicle leakage, rupture, FRAP diffusivity, or condensate-aging data to estimate the control numbers in Table I.

B. Literature-informed parameter anchoring to published protocell data

The present model is not fitted to a specific chemistry, but its variables can be anchored to published protocell observables. As an order-of-magnitude anchoring exercise, we consider two datasets that map directly to the model's variables.

a. Membranized coacervate boundary stability. In membranized coacervate protocells, dextran molecular weight and dextranase-mediated hydrolysis control interfacial membrane thickness, colloidal stability, and molecular transport [18]. In the terminology of the present model, the interfacial dextran layer acts as part of the boundary-integrity variable B, while dextranase acts as an externally imposed boundary-degradation stress. Reported membrane thickness can be mapped to a normalized boundary integrity proxy,

$$B_{\text{proxy}}(t) = \frac{h(t)}{h(0)},$$

where $h(t)$ is the measured interfacial membrane thickness. For example, Ji and Qiao report that after 30 minutes of dextranase treatment at 0, 0.09, 0.18, and 0.26 mg/ml, the residual membrane thickness falls from approximately 0.39 μm to 0.17, 0.03, and 0 μm respectively, showing a threshold-like loss of boundary integrity with increasing degradation stress [18]. If dextranase treatment reduces $h(t)$ over a time window Δt, an apparent boundary damage rate can be estimated as

$$k_B^{\text{app}} = -\frac{1}{\Delta t} \ln \frac{h(t)}{h(0)},$$

when $h(t) > 0$, or as a lower-bound failure rate when $h(t)$ falls below the optical detection limit [18]. This mapping does not identify dextranase hydrolysis with active pruning. Rather, it calibrates the measurable range over which boundary integrity can be experimentally driven from a persistent state to a coalescence-prone state. The active-pruning experiment proposed here would invert this perturbation: under sustained residue production, a tunable clearing or degradation module should maintain bounded residue and delay or prevent the loss of B_{proxy}.

b. GUV leakage versus gross survival. Published GUV data show that gross vesicle survival and leakage can occur on different timescales [23]. For microfluidically produced GUVs with mean diameter approximately

TABLE II. Model variables and experimental proxies. The table is a benchmark plan, not a claim of completed validation.

Model variable	Experimental proxy	Candidate platform	Predicted contrast
ρ	fluorescent inert product, aggregate load, tar fraction	coacervate, vesicle, microfluidic reactor	bounded with pruning, increasing without pruning
D_ρ	FRAP or particle-tracking diffusivity	coacervate, gel, mineral pore	decreases with residue load
B	leakage dye, membrane integrity, compartment retention	GUV, fatty-acid vesicle	integrity preserved by pruning
J	substrate/product throughput, permeability, reaction rate	sustained-flux reactor	throughput collapse without pruning
S	degradation rate, export rate, photochemical clearing rate	catalyst/enzyme/light-driven module	tunable threshold S_c or areal $s_{m,c}$
R	compartment radius	vesicle/droplet microscopy	larger systems require larger areal/total clearing
t^*	latest successful rescue time	pruning shutoff/restoration assay	rescue-window closure

$92\,\mu$m, microscopy-based counting gave a vesicle half-life of about 61 h, while approximately 90% of encapsulated calcein was released within 24 h. A first-order leakage proxy gives

$$k_{\mathrm{leak}} \approx -\frac{\ln(0.1)}{24\,\mathrm{h}} \approx 9.6 \times 10^{-2}\ \mathrm{h}^{-1},$$

whereas gross vesicle disappearance gives

$$k_{\mathrm{count}} \approx \frac{\ln 2}{61\,\mathrm{h}} \approx 1.1 \times 10^{-2}\ \mathrm{h}^{-1}.$$

Thus leakage occurs on a timescale roughly eightfold faster than gross vesicle loss in this dataset. This supports the modeling choice in the reduced boundary simulations: boundary shape, compartment count, and local integrity need not fail simultaneously. In the language of the present model, B should be interpreted as a local functional integrity variable, not merely as the existence of a visibly closed outline.

XI. FAILURE AND DEMOTION CONDITIONS

The active-pruning claim is falsified in the modeled regime if systems under sustained residue production maintain bounded ρ, stable B, and stable J indefinitely with $S = 0$, no hidden dilution, no externalization, and no passive degradation sufficient to offset production. The claim is demoted if pruning merely improves persistence without a threshold or rescue boundary. The spatial-clogging bridge is weakened if residue accumulation does not reduce diffusion or transport in droplet-like systems. The boundary-rupture bridge is weakened if residue pressure does not affect membrane integrity in vesicle-like systems. The size-scaling bridge is weakened if $s_{m,c}(R)$, or the corresponding total clearing requirement $S_c^{\mathrm{tot}}(R)$, does not increase with radius under boundary-limited clearing.

XII. DISCUSSION

Table IV summarizes which conclusions are supported by analytical derivation, simulation, or prospective experimental mapping, and what would demote each claim.

A. What the saddle-node analysis adds

The saddle-node reduction upgrades the collapse language from a numerical observation to a dynamical-systems mechanism. The well-mixed model supplies conditions under which a maintenance attractor and a saddle basin boundary collide. The spatial models then realize the same reduced loss of access in different physical ways: droplet-like systems clog internally; vesicle-like systems deform and rupture; large spherical systems face boundary-limited clearing and diffusion delay.

B. Why chemical coarse-graining is not a defect

The model intentionally does not specify whether residue is RNA side product, peptide aggregate, photochemical tar, lipid impurity, or mineral-surface deposit. That abstraction is a strength for the present purpose. The claim is about a universality class defined by flux, residue, transport/boundary impairment, and active clearing. Specific chemistries can instantiate the class and determine parameter values, but the model predicts that the persistence-collapse structure should recur across substrates.

Recent whole-cell modeling has demonstrated that integrated spatial and kinetic simulations can reproduce the cell-cycle dynamics of an existing genetically minimal organism. The 4D whole-cell model of JCVI-syn3A simulates a complete approximately 100-minute cell cycle, including genetic information processes, metabolic networks, growth, and division, and recovers multiple experimental observables such as doubling time, mRNA half-lives, protein distributions, and ribosome counts [14]. This work illustrates the increasing feasibility of simulation-driven artificial-life research. The present paper addresses a complementary design-level question. Rather than reconstructing an existing minimal cell, we ask what maintenance constraints should be made explicit when designing or screening finite protocell-like systems under sustained flux. In that setting, residue accumulation, transport impairment, boundary integrity, active clearing, rescue timing, and size-dependent pruning demand become candidate control variables for synthetic-cell persistence.

C. Limitations

First, the phase-field component is reduced. We write a phase-field energy to clarify the mechanical interpretation but simulate a lower-dimensional radius/integrity field. A full Cahn-Hilliard, Allen-Cahn, Helfrich-coupled, or hydrodynamic membrane solver is future work. Second, the saddle-node proof applies to the reduced well-mixed model, not to the full spatial PDE. Third, the SDE early-warning analysis is a local weak-noise approximation near the stable branch of the reduced saddle-node normal form. It predicts critical slowing down, variance growth, and rising autocorrelation as the fold is approached quasi-statically. It does not describe large-shock collapse, nonlinear basin escape, or finite-time rupture far from the fold; those regimes require first-passage or large-deviation analysis. Fourth, the simulations are nondimensional and do not fit a specific wet-lab chemistry. Fifth, the benchmark map is prospective, not retrospective validation.

XIII. METHODS AND REPRODUCIBILITY

All figures were generated with reproducible Python scripts using recorded random seeds; stochastic simulations used fixed seed sets for replicate generation. The 2D reaction-diffusion model uses explicit finite differences with residue-dependent diffusivity. The dynamic-boundary model uses a polar radius field with residue-driven pressure, angular deformation modes, and local integrity loss. The SDE early-warning simulations use an Ornstein-Uhlenbeck approximation to the saddle-node normal form. The radial-shell threshold combines the analytical area-volume lower bound with a diffusion-delay correction. The package accompanying this manuscript includes source code, generated figures, and CSV data tables.

A. Parameterization and simulation protocol

All simulations are nondimensional unless explicitly stated. The model uses three classes of parameters: flux–residue parameters, boundary-maintenance parameters, and transport/geometry parameters (summarized in Table III). The baseline value of each parameter is not intended to represent one specific chemistry. Instead, each baseline defines a reference point inside the residue-pruning-boundary universality class, and the robustness sweep tests whether the qualitative threshold structure persists under parameter variation.

The residue-production load is controlled by pJ_0, where p is the residue generated per unit throughput and J_0 is the maximum unblocked throughput. Passive loss is represented by d. Active pruning is represented by S, with a saturating half-load constant K. Boundary damage is controlled by β_ρ, repair by γ_B, transport inhibition by λ_ρ or λ_J, and diffusion arrest by λ_D. Spatial simulations also require a reference diffusivity D_0, lattice spacing Δx, time step Δt, and compartment radius R.

For each figure, the code records the parameter dictionary used to generate the corresponding output. A parameter set is accepted only if the simulation remains numerically stable under halving of Δt and if the qualitative threshold classification is unchanged under modest refinement of the pruning grid. The purpose of this requirement is not to prove convergence of a full continuum PDE solver, but to rule out visually selected numerical artifacts.

B. Threshold extraction and classification protocol

For a fixed parameter vector θ and pruning value S, a run is classified as persistent if three conditions hold over the final observation window $[T_{\text{end}} - \Delta T, T_{\text{end}}]$:

$$V(t) > V_{\min}, \qquad \rho(t) < \rho_{\max}, \qquad \frac{1}{\Delta T}\int_{T_{\text{end}}-\Delta T}^{T_{\text{end}}} |\dot{V}(t)|\,dt < \epsilon_V .$$

A run is classified as collapsed if $V(t) < V_{\min}$ for longer than a prescribed dwell time, or if $B(t)$, $J(t)$, or the spatially averaged transport field falls below its failure threshold.

For deterministic well-mixed and spatial simulations, the critical pruning threshold is estimated by bracketing. We first evaluate a coarse grid of pruning values $S_1 < S_2 < \cdots < S_m$ and identify the smallest interval $[S_i, S_{i+1}]$ for which the classification changes from collapse to persistence. We then refine this interval by bisection until the threshold uncertainty is smaller than a fixed tolerance: $|S_{i+1} - S_i| < \delta_S$. The reported deterministic threshold is $\hat{S}_c = (S_i + S_{i+1})/2$.

For stochastic simulations, the classification is repeated over N random seeds. The persistence probability is estimated as $\hat{P}_{\text{pers}}(S) = \frac{1}{N}\sum_{j=1}^{N} I_j(S)$, where $I_j(S) = 1$ if replicate j persists and 0 otherwise. The stochastic threshold is defined as the 50% persistence point: $\hat{S}_c^{(50)} = \{S : \hat{P}_{\text{pers}}(S) = 0.5\}$, obtained either by monotone interpolation or by fitting a logistic curve $P_{\text{pers}}(S) = 1/(1 + \exp[-a(S - S_c^{(50)})])$. Confidence intervals are obtained by bootstrap resampling over random seeds.

For rescue-window experiments, the same logic is applied to the two-dimensional surface $(t_{\text{restore}}, S_{\text{rescue}})$. The rescue boundary is the 50% contour $P_{\text{rescue}}(t_{\text{restore}}, S_{\text{rescue}}) = 0.5$. A rescue-window closure claim is supported only if the estimated required rescue pruning is nondecreasing with restoration delay over the tested range.

In Fig. 8, each pixel represents an independently estimated threshold using the same collapse/persistence classifier, not a manually selected visual transition.

XIV. DATA AND CODE AVAILABILITY

Source code, parameter files, generated figures, and CSV tables for the simulations are available at: `https://github.com/yiningwu-research/Distinction-Theory/tree/main/models/fds_l1`. The repository contains scripts for the well-mixed model, SDE early-warning simulations, 2D reaction-diffusion clogging, reduced boundary model, rescue-window heatmap, robustness sweep, and radial-shell size-scaling analysis. A machine-readable parameter table specifying baseline values, sweep ranges, random seeds, and figure-generation scripts is included with the released code.

XV. CONCLUSION

Active pruning can be formulated as a measurable control variable for finite protocell-like persistence. The analytical layer gives threshold bounds, saddle-node attractor-loss conditions, stochastic early-warning scaling, and a size-scaling lower bound. The spatial simulations show two physically distinct failure modes: internal clogging in droplet-like or porous compartments and residue-pressure boundary failure in vesicle-like compartments. The radial-shell model shows that size increases pruning demand because residue generated in the volume must diffuse to and be cleared at the boundary; the numerical threshold can substantially exceed the mass-balance lower bound when diffusion delay dominates. The benchmark map translates these variables into experimental proxies. The result is not a completed origin-of-life mechanism, but a testable theory-computation framework for asking whether active residue clearing is required under specified residue-production and boundary-maintenance constraints for long-term persistence in protocell-like systems.

ACKNOWLEDGMENTS

The author used AI-assisted tools for language polishing, code drafting, and structural feedback. All claims, equations, simulations, citations, and final wording remain the responsibility of the author.

Symbol	Meaning	Role in model	Reporting requirement
p	residue per unit throughput	increases residue load	baseline + sweep range
J_0	maximum throughput	sets flux scale	baseline normalization
d	passive dilution/degradation	non-active residue loss	baseline + control runs
S	active pruning rate	main control parameter	grid/bisection range
K	pruning half-saturation load	nonlinear pruning response	baseline value
λ_ρ, λ_J	transport inhibition	throughput suppression by residue	baseline + sensitivity
β_ρ	residue damage coefficient	boundary integrity loss	baseline + sensitivity
γ_B	boundary repair coefficient	recovery of B	baseline + sensitivity
D_0	reference diffusivity	spatial transport scale	baseline + grid test
λ_D	diffusion-arrest coefficient	nonlinear clogging strength	sweep range
R	compartment radius	size-scaling variable	values used in Fig. 6
q	residue generation density	3D source density	value used in size bound
τ_p	pruning timescale	diffusion-delay correction	value or normalization
χ	geometry/transport correction	effective size penalty	fitted or assumed value
$V_{\min}$	collapse threshold	persistence classification	fixed criterion
V_{rec}	rescue threshold	rescue classification	fixed criterion
σ	noise amplitude	SDE early-warning test	values used
$\Delta t, \Delta x$	numerical discretization	stability/reproducibility	refinement check

TABLE III. Simulation parameters and roles. Baseline values and sweep ranges are recorded in the released code rather than reconstructed manually.

Claim	Current support	Demotion condition	Failure consequence
Active pruning threshold	analytical bound + simulations	pruning helps but no threshold	core threshold claim weakened
Saddle-node reduction	well-mixed Jacobian conditions	fold replaced by other bifurcation	revise attractor-loss mechanism
SDE early warnings	normal-form derivation + simulation	variance unreliable alone	keep recovery time/autocorrelation only
Spatial clogging	reaction-diffusion simulation	no diffusion arrest in relevant platforms	limit to other residue mechanisms
Boundary rupture	reduced radius-field simulation	residue pressure does not damage membranes	limit vesicle bridge
Size scaling	lower bound + radial threshold	no $s_{m,c}(R)$ increase	revise boundary-limited scaling
Wet-lab benchmark	prospective map	proxies not measurable	redesign experimental protocol

TABLE IV. Claim status and consequence of failure.

NUMERICAL DETAILS

The 2D model is simulated on a circular mask embedded in a square lattice with an explicit five-point Laplacian. The dynamic-boundary model evolves $R(\theta, t)$ with periodic angular coupling. The rescue heatmap is generated from the reduced fold-distance protocol: required rescue pruning increases with time after shutoff, and rescue probability is a sigmoid function of restored pruning minus this time-dependent requirement. The robustness heatmap estimates S_c over residue production and diffusion-arrest parameters using the same scaling assumptions as the spatial model.

APPENDIX: LINEAR STABILITY OF THE SPATIAL CLOGGING MODEL

The spatial clogging model is not intended to claim a separate Turing-type bifurcation. Its role is to show how the reduced maintenance-attractor loss can be physically realized through residue-dependent transport arrest. To make this explicit, consider the homogeneous state of the spatial equations with constant pruning $S = \bar{S}$:

$$pJ_0\bar{B}e^{-\lambda_J\bar{\rho}} - \bar{S}\frac{\bar{\rho}}{K+\bar{\rho}} = 0, \tag{31}$$

$$\gamma_B(1-\bar{B}) - \beta_\rho\bar{\rho}\bar{B} = 0, \tag{32}$$

which gives $\bar{B} = \gamma_B/(\gamma_B + \beta_\rho\bar{\rho})$.

Let $\rho = \bar{\rho} + \delta\rho$, $B = \bar{B} + \delta B$, and expand perturbations in Laplacian eigenmodes $(\delta\rho, \delta B) \propto e^{\lambda t}\phi_n(x)$ with $-\nabla^2\phi_n = \mu_n\phi_n$, where $\mu_n \geq 0$ and $\mu_0 = 0$ for no-flux boundary conditions. Since $\nabla^2\bar{\rho} = 0$, the residue-dependent diffusivity contributes at linear order as $D_\rho(\bar{\rho})\nabla^2\delta\rho = -\bar{D}\mu_n\delta\rho$ with $\bar{D} = D_0e^{-\lambda_D\bar{\rho}}$. Thus each spatial mode obeys

$$\frac{d}{dt}\begin{pmatrix}\delta\rho\\ \delta B\end{pmatrix} = A(\mu_n)\begin{pmatrix}\delta\rho\\ \delta B\end{pmatrix}, \tag{33}$$

where

$$A(\mu_n) = \begin{pmatrix} f_\rho - \bar{D}\mu_n & f_B \\ g_\rho & g_B\end{pmatrix}, \tag{34}$$

with $f_\rho = -pJ_0\lambda_J\bar{B}e^{-\lambda_J\bar{\rho}} - \bar{S}\frac{K}{(K+\bar{\rho})^2}$, $f_B = pJ_0e^{-\lambda_J\bar{\rho}}$, $g_\rho = -\beta_\rho\bar{B}$, $g_B = -\gamma_B - \beta_\rho\bar{\rho}$. The trace and determinant are $\operatorname{tr} A(\mu_n) = f_\rho + g_B - \bar{D}\mu_n$ and $\det A(\mu_n) = (f_\rho - \bar{D}\mu_n)g_B - f_Bg_\rho = \det A(0) - \bar{D}\mu_n g_B$. Because $g_B < 0$, the determinant increases with μ_n, while the trace decreases with μ_n. Therefore, in this minimal local linearization, nonzero spatial modes are not the first modes to destabilize when the homogeneous maintenance state loses stability. The primary local loss is the $\mu_0 = 0$ maintenance mode, consistent with the reduced saddle-node analysis.

The spatial simulations should therefore be interpreted as nonlinear realizations of maintenance failure rather than as a claim of a separate Turing instability. Spatial structure enters through finite-amplitude residue accumulation, diffusion arrest, boundary geometry, and finite-size transport delay. A full nonlinear PDE bifurcation analysis is left for future work.

---

[1] M. M. Hanczyc, S. M. Fujikawa, and J. W. Szostak, Experimental models of primitive cellular compartments: encapsulation, growth, and division, Science **302**, 618–622 (2003). DOI: 10.1126/science.1089904.

[2] I. A. Chen and J. W. Szostak, A kinetic study of the growth of fatty acid vesicles, Biophysical Journal **87**, 988–998 (2004). DOI: 10.1529/biophysj.104.039875.

[3] I. Budin and J. W. Szostak, Physical effects underlying the transition from primitive to modern cell membranes, Proceedings of the National Academy of Sciences **108**, 5249–5254 (2011). DOI: 10.1073/pnas.1100498108.

[4] W. Helfrich, Elastic properties of lipid bilayers: theory and possible experiments, Zeitschrift fuer Naturforschung C **28**, 693–703 (1973). DOI: 10.1515/znc-1973-11-1209.

[5] U. Seifert, Configurations of fluid membranes and vesicles, Advances in Physics **46**, 13–137 (1997). DOI: 10.1080/00018739700101488.

[6] J. W. Cahn and J. E. Hilliard, Free energy of a nonuniform system. I. Interfacial free energy, Journal of Chemical Physics **28**, 258–267 (1958). DOI: 10.1063/1.1744102.

[7] S. M. Allen and J. W. Cahn, A microscopic theory for antiphase boundary motion and its application to antiphase domain coarsening, Acta Metallurgica **27**, 1085–1095 (1979). DOI: 10.1016/0001-6160(79)90196-2.

[8] Y. Shin and C. P. Brangwynne, Liquid phase condensation in cell physiology and disease, Science **357**, eaaf4382 (2017). DOI: 10.1126/science.aaf4382.

[9] S. F. Banani, H. O. Lee, A. A. Hyman, and M. K. Rosen, Biomolecular condensates: organizers of cellular biochemistry, Nature Reviews Molecular Cell Biology **18**, 285–298 (2017). DOI: 10.1038/nrm.2017.7.

[10] M. Scheffer et al., Early-warning signals for critical transitions, Nature **461**, 53–59 (2009). DOI: 10.1038/nature08227.

[11] C. Kuehn, A mathematical framework for critical transitions: bifurcations, fast-slow systems and stochastic dynamics, Physica D **240**, 1020–1035 (2011). DOI: 10.1016/j.physd.2011.02.012.

[12] S. H. Strogatz, Nonlinear Dynamics and Chaos: With Applications to Physics, Biology, Chemistry, and Engineering, 2nd ed. (Westview Press, 2015). ISBN: 978-0-8133-4910-7.

[13] C. W. Gardiner, Stochastic Methods: A Handbook for the Natural and Social Sciences, 4th ed., Springer Series in Synergetics Vol. 13 (Springer, 2009). ISBN: 978-3-540-70712-7.

[14] Z. R. Thornburg et al., Bringing the genetically minimal cell to life on a computer in 4D, Cell **189**(9), 2582–2597.e27 (2026). DOI: 10.1016/j.cell.2026.02.009.

[15] C. Chen and J. Li, Recent advances in coacervate protocells from passive catalysts to chemically programmable systems, Communications Chemistry **9**, Article 76 (2026). DOI: 10.1038/s42004-026-01937-4.

[16] I. B. A. Smokers, B. S. Visser, A. D. Slootbeek, W. T. S. Huck, and E. Spruijt, How Droplets Can Accelerate Reactions—Coacervate Protocells as Catalytic Microcompartments, Accounts of Chemical Research **57**, 1885–1895 (2024). DOI: 10.1021/acs.accounts.4c00114.

[17] R. Harris, N. Berman, and A. Lampel, Coacervates as enzymatic microreactors, Chemical Society Reviews **54**, 4183–4199 (2025). DOI: 10.1039/D4CS01203H.

[18] Y. Ji and Y. Qiao, Tuning interfacial fluidity and colloidal stability of membranized coacervate protocells, Communications Chemistry **7**, Article 122 (2024). DOI: 10.1038/s42004-024-01193-4.

[19] J. Sastre, A. Thatte, A. M. Bergmann, M. Stasi, M. Tena-Solsona, C. A. Weber, and J. Boekhoven, Size control and oscillations of active droplets in synthetic cells, Nature Communications **16**, Article 2003 (2025). DOI: 10.1038/s41467-025-57240-8.

[20] C. J. Cho, T. An, Y.-C. Lai, A. Vázquez-Salazar, A. Fracassi, R. J. Brea, I. A. Chen, and N. K. Devaraj, Protocells by spontaneous reaction of cysteine with short-chain thioesters, Nature Chemistry **17**, 148–155 (2025). DOI: 10.1038/s41557-024-01666-y.

[21] A. Agrawal, A. Radakovic, A. Vonteddu, S. Rizvi, V. N. Huynh, J. F. Douglas, M. V. Tirrell, A. Karim, and J. W. Szostak, Did the exposure of coacervate droplets to rain make them the first stable protocells?, Science Advances **10**, eadn9657 (2024). DOI: 10.1126/sciadv.adn9657.

[22] K. Evers, D. Borsboom, E. I. Fried, F. Hasselman, and L. Waldorp, Early warning signals of complex critical transitions in deterministic dynamics, Nonlinear Dynamics **112**, 19071–19094 (2024). DOI: 10.1007/s11071-024-10023-0.

[23] M. Ernits, O. Reinsalu, N. Yandrapalli, S. Kopanchuk, E. Moradpur-Tari, I. Sanka, O. Scheler, A. Rinken, R. Kurg, A. Kyritsakis, V. Linko, and V. Zadin, Microfluidic production, stability and loading of synthetic giant unilamellar vesicles, Scientific Reports **14**, Article 14071 (2024). DOI: 10.1038/s41598-024-64613-4.