# Cosmic Ecosystem Simulation

*Live Ecology in the Browser: Architecture, Dynamics, and ODE Analysis*

Amber Miller | Personal Project | 2024–2025

**Abstract**

This report documents the design, implementation, and mathematical analysis of a real-time browser ecosystem simulation in which six species across four trophic levels hunt, breed, and die inside a procedurally-generated cosmic canvas. Each creature runs its own priority-ordered drive hierarchy every animation frame. Multi-tab synchronisation is managed through a Firebase Realtime Database backend with automatic master-election logic. Population dynamics are modelled with an extended Lotka–Volterra ODE system whose parameters are derived directly from source-code constants and validated against stochastic simulation data. A key finding of the analysis was that the original spatial encounter function required sigma re-calibration (0.001 $\rightarrow$ 50000) to produce physically meaningful interaction rates; beta coefficients were then back-solved analytically to guarantee a stable fixed point at the target equilibrium H*=40, C*=15, A*=3.

| JavaScript | Firebase | Simulation | Lotka–Volterra | Dynamical Systems | SDE / Ecology |
|---|---|---|---|---|---|

## 1. Project Overview

The Cosmic Ecosystem is a browser-native, real-time simulation where autonomous creatures inhabit a scrolling canvas populated with stars, planets, suns, galaxies, comets, and nebulae. Each creature perceives its local neighbourhood through a spatial-hash grid, selects a behaviour via a priority-ordered drive hierarchy, and updates its position at approximately 60 fps. The simulation supports simultaneous viewing across multiple browser tabs through Firebase Realtime Database sync with master-election logic.

### Key Features

- Six species across four trophic levels: jellyfish, manta, seahorse (herbivores); shark, anglerfish (carnivores); leviathan (apex)
- Procedural cosmic background with day/night cycle, food bloom spawning, and 90 individually tracked stars
- Lévy-flight wander, spiral mate-search, and energy-gated reproduction with heritable traits
- Firebase Realtime Database multi-tab sync with automatic master-election and offline fallback
- God Mode panel for real-time parameter tuning: food multiplier, aggression, mutation rate, day speed

- Live population graph, trait evolution panel, per-creature inspection panel with Q-table readout
- "Chaos UI" easter egg: mouse-trail sparkles, cycling neon headers, Zerg Rush mini-game

## 2. System Architecture

The codebase is split into seven modules loaded sequentially in the HTML entry point:

| Module | File | Responsibility |
|---|---|---|
| Sync | eco-sync.js | Firebase init, master election, push/poll, snapshot serialisation |
| Canvas | eco-canvas.js | Canvas setup, background rendering, celestial objects, spatial hash |
| Creatures | eco-creatures.js | Species definitions, spawn logic, update loop, energy, reproduction, drawing |
| Behaviour | eco-behaviour.js | Drive hierarchy: flee > hunt > feed > mate > wander; Lévy flight; kill check |
| UI | eco-ui.js | Inspect panel, population graph, trait panel, God Mode, HUD, input handling |
| Main | eco-main.js | Init, main animation loop, celestial state injection |
| Chaos UI | chaos-ui.js | Sparkle cursor, marquees, animated headers, Zerg Rush mini-game |

### Spatial Hash Grid

Naïve $O(N^2)$ neighbour search is replaced by a fixed-cell spatial hash. Each frame all creatures are bucketed by (x, y) position; a nearby-query returns only creatures in the same and adjacent cells. This reduces lookup from O(N) to O(k) where k is local density, enabling smooth 60 fps with 80+ creatures. Spark effects use a DOM element pool (up to 30 reused nodes) and an 80 ms throttle to keep rendering cost low.

### Firebase Multi-Tab Sync

On load each tab queries the database. If the saved snapshot is recent (age < 12 s) the tab enters viewer mode and applies incoming snapshots via a real-time onValue listener. If no valid save exists, or no push has been seen within the timeout, the tab claims master status and begins pushing full snapshots every 5 s. Float32Array Q-tables are serialised to plain arrays for JSON compatibility, then re-cast on deserialisation.

## 3. Species Definitions & Food Web

Following the rebalancing pass, the food web forms a strict four-level chain. Apex predators hunt carnivores exclusively — a change from the initial design where leviathan could eat herbivores directly. This establishes proper top-down regulation: apex controls carnivore density, which in turn controls herbivore density, preventing apex-driven herbivore extinction.

| Species | Trophic Level | Diet | Target Pop. | Max Energy | Key Trait |
|---|---|---|---|---|---|
| Jellyfish | Herbivore | herb | ~14 | 150 | Low speed, schooling |
| Manta | Herbivore | herb | ~14 | 150 | Mid speed, wide sense |
| Seahorse | Herbivore | herb | ~14 | 150 | Slow, hides near planets |
| Shark | Carnivore | carn | ~8 | 300 | High speed, lone hunter |
| Anglerfish | Carnivore | carn | ~8 | 300 | Ambush, low light emit |
| Leviathan | Apex | apex | ~3 | 500 | Hunts carnivores only |

### Rebalancing Changes

- **Apex diet restriction:** Apex now eats carnivores only — creates proper four-level trophic control and prevents herbivore collapse
- **Energy capacity:** Carnivore max energy 150→300; apex max energy 150→500 to buffer starvation between kills
- **Food availability:** All herbivore energy gains roughly doubled: planets +82%, suns +80%, stars +150%, food blooms +67%
- **Kill efficiency:** Kill energy multiplier doubled 6×→12×; carnivores now sustain on fewer kills per minute
- **Population cap removed:** Hard cap of 60 removed; population limits emerge naturally from predation and starvation
- **Respawn removed:** Automatic per-species respawn removed; extinction is now possible, enabling true trophic cascades
- **Starting populations:** Herbivores 14×3=42, carnivores 8×2=16, apex 3 — all near target equilibrium values

## 4. Behaviour System

| | | |
|---|---|---|
| **1. Flee** | HIGHEST | Predator within sense/2 → steer away at 1.5× turn cap; fear state accumulates |
| **2. Hunt/Feed** | HIGH | Hungry or below mating threshold: herbivores seek food blooms; carnivores/apex chase prey and run kill-check on contact |
| **3. Seek Mate** | MEDIUM | Energy above seekMate threshold: spiral outward search, then approach; orbit slowly when within breedRange |

| | | |
|---|---|---|
| **4. Wander** | FALLBACK | Lévy-flight: alternates slow local meander and long power-law jumps; edge avoidance via centre-pull |

## Lévy-Flight Wander

When no drive fires, creatures use a two-state Lévy-flight wander. In the *explore* state the creature meanders slowly with small angular perturbations (±0.025 rad/frame). With probability proportional to time spent exploring (capped at 25%), it switches to *relocate*: a jump target is sampled from a power-law distance distribution (dist ~ $rand^{-0.5}$ × sense × 3) at a random bearing, then clamped to the canvas interior. On arrival, or after a 100-frame timeout, the creature returns to explore. Edge avoidance overrides both states by pulling the heading toward the canvas centre with strength proportional to proximity, preventing wall clustering.

# 5. Mathematical Formulation

The simulation can be approximated by a system of coupled ODEs analogous to extended Lotka–Volterra dynamics. Let F(t), H(t), C(t), A(t) denote food bloom density, herbivore, carnivore, and apex population counts respectively.

## 5.1 Core Dynamical System

```
dF/dt = r_F(t) · F · (1 - F/K_F) - alpha_HF · H · F · eps(H, F)
dH/dt = beta_HF · H · F · eps(H,F) · Theta(E_H - 150) - mu_H · H - alpha_CH · C · H · eps(C,H)
dC/dt = beta_CH · C · H · eps(C,H) - mu_C · C - alpha_AC · A · C · eps(A,C)
dA/dt = beta_AC · A · C · eps(A,C) - mu_A · A
```

The spatial encounter probability function eps(N1, N2) = 1 – exp(–sigma · N1·N2 / A_total) captures finite sensing range: sparse populations rarely meet. Circadian forcing enters as r_F(t) = r_day · [0.5 + 0.5 · sin(omega · t)], driving periodic food bloom growth with a 30 s period. Energy-gated reproduction is modelled by the Heaviside factor Theta(E_H – 150), where 150 is the seekMate energy threshold read directly from the source code.

## 5.2 Sigma Calibration — Fixing the Encounter Function

**Bug found during analysis:** The original script used sigma = 0.001 with A_total = 960000. At typical equilibrium populations (H=40, C=15), this gives eps(40, 15) = 1 – exp(–0.001 × 600 / 960000) ≈ 6.25 × $10^{-7}$ — effectively zero. All predator–prey interaction terms vanished, herbivores experienced no predation pressure, and the simulation collapsed monotonically rather than oscillating.

The correct sigma was derived by requiring eps ≈ 0.90 at the target apex–carnivore pair (A=3, C=15), which represents the sparsest realistic interaction. Solving 1 – exp(–sigma × 3 × 15 / 960000) = 0.90 gives sigma ≈ 48800. The script uses sigma = 50000 (rounded), which gives the following encounter probabilities at equilibrium:

| Pair | N1 | N2 | eps (sigma=0.001) | eps (sigma=50000) | Effect |
|---|---|---|---|---|---|
| Herb–Food | 40 | 100 | 4.2 × 10■■ | ≈ 1.000 | Herbivores always find food |
| Carn–Herb | 15 | 40 | 6.3 × 10■■ | ≈ 1.000 | Carnivores hunt effectively |
| Apex–Carn | 3 | 15 | 4.7 × 10■■ | ≈ 0.904 | Apex controls carnivores |

## 5.3 Equilibrium-Consistent Parameter Derivation

Rather than guessing conversion efficiencies (beta values), all four were back-solved analytically from the condition dX/dt = 0 at the target fixed point (F*=100, H*=40, C*=15, A*=3). This guarantees the system orbits around the desired populations instead of drifting away. The derivation proceeds as follows:

### From dA/dt = 0:

```
beta_AC = mu_A / (C* · eps(A*, C*)) = 0.00584 / (15 × 0.904) = 0.000431
```

*Apex conversion rate set to exactly sustain 3 individuals.*

### From dC/dt = 0:

```
beta_CH = (mu_C + alpha_AC · A* · eps(A*, C*)) / (H* · eps(C*, H*)) = 0.005215
```

*Carnivore conversion balances mortality plus apex predation.*

### From dH/dt = 0:

```
beta_HF = (mu_H + alpha_CH · C* · eps(C*, H*)) / (F* · eps(H*, F*)) = 0.003478
```

*Herbivore conversion balances natural mortality plus carnivore predation.*

### From dF/dt = 0 (avg. over circadian):

```
alpha_HF = r_F · 0.5 · (1 − F*/K_F) / (H* · eps(H*, F*)) = 0.000938
```

*Food consumption rate balanced against average logistic growth.*

## 5.4 Fitted Parameter Set

| Parameter | Value | Source | Description |
|---|---|---|---|
| alpha_CH | 0.02263 | ODE fit | Carnivore attack rate on herbivores |
| alpha_AC | 0.07424 | ODE fit | Apex attack rate on carnivores |
| beta_HF | 0.003478 | Back-solved | Herbivore energy gain per food unit |
| beta_CH | 0.005215 | Back-solved | Carnivore energy gain per herbivore killed |
| beta_AC | 0.000431 | Back-solved | Apex energy gain per carnivore killed |
| alpha_HF | 0.000938 | Back-solved | Food consumption rate by herbivores |
| mu_H | 0.00907 | From code | Herbivore mortality (metabolism + age) |
| mu_C | 0.00727 | From code | Carnivore mortality (metabolism + age) |
| mu_A | 0.00622 | From code | Apex mortality (metabolism + age) |
| sigma | 50000 | Calibrated | Spatial encounter scaling (eps ≈ 0.9–1.0 at equilibrium) |
| omega | 2π/30 | From code | Circadian angular frequency (30 s day/night period) |

*Mortality rates are computed directly from the per-frame energy drain formula in eco-creatures.js: mu_i = (0.05 + size × 0.005 + speed × 0.015) / avg_energy + 1/maxAge_seconds.*

# 6. Stochastic Simulation & Model Fitting

## 6.1 SDE Data Generation

Synthetic population data was generated using the Euler-Maruyama scheme applied to the ODE system with demographic noise. The SDE form is:

```
dX = f(X, t) dt + noise_scale · sqrt(X) · dW
```

where dW ~ N(0, sqrt(dt)) and noise_scale = 0.03. The sqrt(X) scaling is the standard van Kampen birth-death process approximation: variance grows with population size but the coefficient of variation decreases, matching real ecological stochasticity. A hard floor of 0.01 prevents negative populations. The simulation was run for 600 s at dt = 0.05 s (12001 time points) starting from y0 = [100, 42, 16, 3].

**Results from data generation:**

| Species | Mean | Min | Max | Interpretation |
|---|---|---|---|---|
| Herbivores (H) | 39.9 | 19.2 | 68.0 | Centred on H*=40; ±20 natural variation |
| Carnivores (C) | 14.7 | 8.8 | 24.6 | Near C*=15; tighter range due to apex buffering |
| Apex (A) | 3.1 | 2.4 | 3.8 | Very stable; slow dynamics, small amplitude |

## 6.2 ODE Model Fitting

The five free parameters (alpha_CH, alpha_AC, mu_H, mu_C, mu_A) were fitted by minimising the normalised sum of squared residuals between the deterministic ODE trajectory and the stochastic data. At each evaluation, all beta values and alpha_HF were re-derived analytically via the equilibrium back-solve, so the ODE always admits (F*, H*, C*, A*) as a fixed point regardless of the trial parameters. Fitting used a coarse $3^5$ = 243-point grid search (±20% around initial estimates) followed by Nelder-Mead local refinement.

| Parameter | Initial | Fitted | Change |
|---|---|---|---|
| alpha_CH | 0.02263 | 0.01753 | –22% |
| alpha_AC | 0.07424 | 0.09060 | +22% |
| mu_H | 0.00907 | 0.00699 | –23% |
| mu_C | 0.00727 | 0.00614 | –16% |
| mu_A | 0.00622 | 0.00766 | +23% |

The fitted alpha_CH is lower than the geometric estimate (carnivores hunt less aggressively in the stochastic data because noise sometimes scatters encounters), while alpha_AC is higher (apex is more effective than the initial estimate). The final residual error was 0.6276 on the normalised scale.

# 7. Stability Analysis

## 7.1 Linearised Stability at Equilibrium

The Jacobian of the ODE system was computed numerically at the fixed point (F*=100, H*=40, C*=15, A*=3) by finite differences with eps = $10^{-6}$. The four eigenvalues of the fitted system are:

| Eigenval | Real Part | Imaginary Part | Interpretation |
|---|---|---|---|
| λ■ | –0.0291 | +0.2793i | Complex pair → damped oscillation, period ≈ 22.5 s |
| λ■ | –0.0291 | –0.2793i | Complex conjugate of λ■ |
| λ■ | –0.0277 | 0 | Real → fast exponential decay (herbivore mode) |
| λ■ | –0.0109 | 0 | Real → slow decay (apex / food adjustment) |

**Stability conclusion:** All Re(λ_i) < 0, confirming the equilibrium is **asymptotically stable**. The complex pair λ■,■ produces damped oscillatory approach with intrinsic period T = 2π / |Im(λ)| = 2π / 0.2793 ≈ 22.5 s and half-life = ln(2) / 0.0291 ≈ 23.8 s. In the stochastic simulation, demographic noise continuously re-excites these oscillations, producing sustained fluctuations of amplitude ±20 (herbivores) and ±8 (carnivores) around equilibrium.

## 7.2 Intrinsic vs. Circadian Period

The Lotka–Volterra intrinsic oscillation period can be estimated analytically as:

```
T_LV = 2π / sqrt(alpha_CH · beta_CH · H* · C*) = 2π / sqrt(0.01753 × 0.00522 × 600) ≈ 24.4 s
```

This is close to but distinct from the circadian forcing period of 30 s. The power spectrum of herbivore counts peaks at 30 s — the circadian signal dominates at steady state because it is a persistent external forcing whereas the LV mode is rapidly damped (half-life ≈ 24 s). Under nominal parameters there is no resonance between the two periods, but if alpha_CH were increased by ~80% the intrinsic period would shift to ~30 s, potentially amplifying oscillations.

## 7.3 Comparison: Old vs. New Analysis

| Quantity | Original (broken) | Fixed analysis |
|---|---|---|
| sigma | 0.001 | 50000 |
| eps at equilibrium | ~10■■ (zero) | 0.90 – 1.00 |
| Beta derivation | Guessed | Back-solved from dX/dt = 0 |
| Starting conditions | Far from equilibrium | Near equilibrium (42, 16, 3) |
| Observed dynamics | Monotonic collapse | Stable bounded oscillations |
| H range (data) | Decays 40 → 0 | 19 – 68, mean = 40 |
| C range (data) | Decays 15 → 0 | 8.8 – 24.6, mean = 14.7 |
| Dominant period | 300.6 s (slow transient) | 30 s (circadian forcing) |
| Eigenvalues | Not computed | All Re(λ) < 0, stable |

| Quantity | Original (broken) | Fixed analysis |
| --- | --- | --- |
| Equilibrium guaranteed | No | Yes (by construction) |

## 8. Expected Population Dynamics

The corrected model predicts the following phases starting from the rebalanced initial conditions (H=42, C=16, A=3):

**Phase 1 — Transient (0–60 s)**
Herbivores fluctuate around 40–55 as carnivore hunting ramps up from the slightly above-equilibrium starting density. Apex remains near 3, conserving energy.

**Phase 2 — Damped Approach (60–200 s)**
The complex eigenvalue pair drives a damped spiral toward equilibrium. Herbivores and carnivores oscillate with decreasing amplitude; half-life ≈ 24 s means transient corrections are mostly resolved within 2–3 oscillation cycles (~70 s).

**Phase 3 — Noise-Sustained Oscillations (200+ s)**
Once the deterministic transient has decayed, demographic stochasticity continuously re-excites the oscillatory mode. Herbivores range from ~20 to ~68; carnivores from ~9 to ~25. Apex is buffered by the carnivore layer and varies only from ~2.4 to ~3.8.

**Long-term behaviour**
The system is ergodic around the fixed point. True extinction is possible but rare at nominal parameters — the demographic noise level (noise_scale=0.03) is calibrated to keep minimum populations well above zero.

### Emergent Phenomena

- Trophic cascades: removing apex → carnivore bloom → herbivore crash → food bloom (all eigenvalues shift)
- Near-resonance risk: if alpha_CH increases ~80%, T_LV ≈ 30 s circadian period → oscillation amplification
- Spatial pattern formation (Turing instability): predicted when D_herb/D_carn ≠ 1 in the PDE extension
- Extinction threshold: demographic noise pushes a species below ~5 individuals at which stochastic extinction becomes likely

## 9. Analysis Results

The figure saved as ecosystem_fit.png shows four panels from the corrected analysis (ecosystem_analysis.py). All panels now show a stable, bounded ecosystem rather than the monotonic collapse visible in the original plot.

### Population Dynamics (top-left)

Stochastic data (dots) and deterministic ODE trajectory (lines) for all three consumer trophic levels. Both hover around the target equilibrium values (H*=40, C*=15, A*=3, shown as dotted lines) throughout the 600 s run. The deterministic model tracks the mean of the noisy data closely.

### Phase Space H vs C (top-right)

Carnivore population plotted against herbivore population. The noisy data forms a cloud centred on (40, 15) while the ODE trajectory spirals inward toward the equilibrium star marker. This spiral structure is the signature of the complex eigenvalue pair (damped oscillation).

### Power Spectrum (bottom-left)

FFT of herbivore counts computed on the second half of the time series (300–600 s) to exclude the initial transient. Three period markers are shown: the dominant peak at 30 s (circadian forcing), the LV theory period at ~24 s (from fitted parameters), and the labelled fitted peak. The circadian signal dominates because the LV mode is damped (half-life ~24 s) while the forcing is persistent.

### Parameter Summary (bottom-right)

Table of all fitted parameters including derived beta values, the intrinsic LV period of the fitted system, and the target equilibrium values for reference.

## 10. Technical Challenges & Solutions

- **Performance at scale:** $O(N^2)$ neighbour queries caused frame-rate collapse above ~40 creatures. Solution: fixed-cell spatial hash grid reduces neighbour lookup to $O(k)$, $k \approx 5$–10. Additional: spark element pool (30 reused DOM nodes), 80 ms throttle on sparkle rate, background pre-rendered to an off-screen canvas and only re-drawn on dayT change.
- **Population instability in simulation:** Initial implementation: apex ate both herbivores and carnivores, causing rapid herbivore extinction. Solution: apex diet restriction to carnivores only, doubled food availability and kill efficiency, removed hard population cap to allow natural oscillation rather than hitting an artificial ceiling.
- **Spatial encounter function miscalibration:** The analysis script originally used sigma=0.001, giving eps ≈ 10■■ at all equilibrium populations — effectively disabling all interactions. This caused the ODE to produce monotonic population collapse rather than oscillations. Solution: sigma re-calibrated to 50000 so that eps ≈ 0.9–1.0 at typical population densities, restoring physically meaningful predator–prey coupling.
- **Beta coefficient instability:** Guessing beta values led to equilibria far from the target, with the ODE trajectory diverging from the simulation data regardless of alpha tuning. Solution: back-solve all four beta values analytically from $dX/dt = 0$ at $(F^*, H^*, C^*, A^*)$, guaranteeing the target is a fixed point by construction. This is enforced at every fitting evaluation.
- **Multi-tab state synchronisation:** Multiple tabs running independent simulations produced divergent state. Solution: master-election protocol with 12 s timeout. Master pushes every 5 s; viewers subscribe via real-time onValue listener (6 s poll fallback). Float32Array Q-tables serialised to plain arrays for JSON transport.
- **Lévy-flight boundary conditions:** Power-law jump sampling produced targets outside the canvas, causing wall clustering. Solution: jump target clamped to (edgePad, W–edgePad); edge-avoidance state overrides all wander modes with a centre-pull force proportional to proximity.

## 11. Comparison with Classical Lotka–Volterra

| Feature | Classical L–V | This Ecosystem |
|---|---|---|
| Encounter model | Mass action (N1·N2) | Spatial saturation; eps = 1 – exp(–sigma·N1·N2/A) |
| Food source | Constant / logistic | Periodic: r_F(t) = r_day·[0.5 + 0.5·sin(ωt)] |
| Reproduction | Instantaneous | Energy-gated threshold (E > 150); spiral mate-search |
| Mortality | Constant rate | Code-derived: 0.05 + size×0.005 + speed×0.015 per frame |
| Space | Well-mixed | Explicit 2-D; spatial hash; Lévy-flight wander |
| Stochasticity | Deterministic | Demographic noise: dX += noise·√X·dW |
| Trophic levels | 2 | 4 (food → herb → carn → apex) |
| Equilibrium guarantee | By parameter choice | Analytical back-solve enforces fixed point |

| Feature | Classical L–V | This Ecosystem |
|---|---|---|
| Oscillation period | Neutrally stable | T_LV ≈ 22.5 s (intrinsic), 30 s (circadian dominant) |

## 12. Tools & Technologies

| Category | Technology |
|---|---|
| Frontend rendering | HTML5 Canvas 2D API, vanilla JavaScript (ES2020) |
| Spatial indexing | Custom fixed-cell spatial hash (O(k) neighbour query) |
| State synchronisation | Firebase Realtime Database v10 (dynamic ESM CDN import) |
| Creature learning | Per-creature Q-table (tabular RL, Float32Array, 9 actions × N states) |
| ODE integration | Python / SciPy solve_ivp (RK45), Euler-Maruyama SDE |
| Parameter fitting | 3■ grid search + Nelder-Mead local refinement; equilibrium back-solve |
| Spectral analysis | SciPy FFT on second half of time series to exclude transient |
| Visualisation | Matplotlib; 4-panel figure (dynamics, phase portrait, FFT, params) |
| Deployment | Static file hosting; offline-capable with Firebase fallback |