

Entropy's Edge

A Specification for a Game of Organized Structure, Costly Time, Nested Boundaries,
and Delayed Consequence

Flyxion
Independent Researcher

October 2026

Abstract

Entropy's Edge is a space strategy and space-flight game whose central problem is maintaining organized structure against the processes that dissipate it. The world is a simulated environment of three interacting fields, a scalar constraint density, a directional flow, and an entropy density, and the calendar is a resource, every boundary is a potential trigger, and nothing a player does to another player is visible until its consequences have had time to arrive. This document specifies the field substrate, the game's time model, its world of nested scopes, the strategic layer of stars, corridors, fleets, factions and victory conditions, the visual direction, the control scheme, the global reset and its inoculations, the blink mechanism and its automated tiers, the placement and defusal of core triggers, the minefield economy, the guide bot that doubles as tutorial and screensaver, the multiplayer rules that follow from isolation across interstellar distances, and the status of the existing prototypes together with the gap that remains between them and the design. The design began as a response to a defect in ordinary social platforms, where a date attached to a post or a page is a free-form field that readers nevertheless treat as evidence. The game takes the opposite position: a date is a coordinate in a recorded tree of choices, every movement along the date axis has a price, and what one player learns about another is always an effect and never a claim. The document also states the rules by which events cross between the branch histories of different players, separates the several protections against a global reset, fixes the eligibility and propagation rules for core triggers, and records, with a status label for each, which parts of the design are recorded decisions, which are proposals made here, and which describe verified code.

Contents

1	Motivation and Design Stance	5
1.1	A second motivation: ecological proxies and collective selection	5
1.2	Conventions for status	6
2	The Time Model	6
2.1	The calendar and its endpoints	6
2.2	The tree of choices	7

2.3	Interpolation	7
2.4	Endpoints and the shape of the history	7
2.5	Two clocks	8
2.6	Causal time and cosmological time	8
3	The Field Substrate	9
3.1	Three fields	9
3.2	The lattice	9
3.3	Determinism and the branch tree	10
3.4	Scopes as organized regions	10
4	The World: Scopes, Boundaries, and Cores	10
4.1	Nested scopes	10
4.2	Popping	10
4.3	The core	11
4.4	Defusal and relocation	11
4.5	Eligibility, queue cycling, and propagation paths	11
5	The Galaxy and Its Contents	12
5.1	A galaxy built from articles	12
5.2	Resources and progression	12
5.3	Flares, beacons, and bookmarks	13
5.4	Removing or enclosing a star	13
5.5	Trajectories as media	13
6	Controls	14
6.1	Rotation, contrast, and blinking	14
6.2	Modes	15
6.3	Chords and sequences	15
7	The G Key and the Global Reset	15
7.1	Hold time	15
7.2	Three kinds of exposure and protection	16
7.3	Propagation and the light cone	17
7.4	Disincentives	17
7.5	Crystal-funded rewinds	17
7.6	Heat death and the annual cycle	17
7.7	Experimenting with resets	18
8	Blink, Autoblink, and Super Autoblink	19

8.1	Blink	19
8.2	Autoblink	19
8.3	Super autoblink	19
8.4	What blinking does not do	19
9	Minefields, Containment, and Salvage	20
9.1	Containment and insulation	20
9.2	Dyson spheres and swarms	20
9.3	Hidden until tested	21
9.4	Who is caught	21
9.5	Salvage	21
10	Multiplayer: Isolation and Light-Cone Verification	21
10.1	Scattered players	21
10.2	Travel as verification	22
10.3	Deniable presence	22
10.4	Rewrites and what observers see	22
10.5	Event transmission and branch assignment	22
11	The Strategic Layer	23
11.1	The four activities reinterpreted	23
11.2	Infrastructure	23
11.3	Stars, corridors, and fleets	24
11.4	Factions	24
11.5	Victory conditions and progression	24
11.6	Minimum version	25
12	The Prototype Document and Its Mechanics	25
12.1	Field ontology and the variational form	25
12.2	The core equations of the prototype	25
12.3	Phases of expansion and integration	26
12.4	Numerical scheme	26
12.5	Turn loop	26
12.6	Ethics tensor and diplomacy	27
12.7	Anomaly missions	27
12.8	Fleets and combat	27
12.9	Scenario generation and victory in the prototype	28
12.10	Neural Darwinism and the technology tree	28
12.11	Geothermal mass accelerators	29
12.12	Recursive futarchy	29

12.13	Grounding against instrumental convergence	29
12.14	Horology and monetization	29
12.15	Gradient-driven and activity-driven entropy production	30
12.16	Observations on the prototype’s claims	31
13	Two Layers of Play	33
14	Visual Direction	34
15	The Guide Bot	34
15.1	Open point	35
16	Implementation Status and Roadmap	35
16.1	What exists	35
16.2	The gap	35
16.3	Roadmap	36
17	Formal Sketch	36
18	Open Questions	38
19	Relation to the Originating Problem	39
A	Status of Principal Claims	39

1 Motivation and Design Stance

The immediate source of the design is an observation about how timestamps behave on a large social platform. A page there can carry posts dated to any moment between its recorded founding and, until limits were added, a point decades in the future. A post may be dated to the year of the Battle of Hastings, a page may display a job at a space station on a moon of Saturn in 1915, and an event may be scheduled for 2042 after the platform has since capped scheduling at two years, because the cap was applied to new entries and the old ones were left in place. The interface renders all of these in exactly the same format as a post whose date is real. A visitor reading a long timeline therefore receives what looks like a track record, and the platform’s own presentation lends credibility to a history that nobody checked. The same effect arises innocently when a page is renamed: posts made under an earlier identity appear under the current name with no seam, so an idea seems older than it is. The information that would correct the impression, namely the creation date and the name history, exists in the system but has been moved out of the main view.

Two distinct clocks are being displayed through a single slot. One is the time at which something was posted, which only the platform can know. The other is the time at which something is said to have happened, which only the poster can claim. Because the display does not distinguish them, the reader cannot tell a verified fact from an assertion, and because longevity is one of the cheapest trust signals people use, the confusion is exploitable.

Entropy’s Edge is the attempt to build a world in which the same capability is a feature and not a flaw. Players may move along the time axis, backward as well as forward, but every movement is priced, every movement is recorded in a tree that the player can inspect, and no player is ever asked to believe another player’s date. The three commitments that follow from this stance organize the rest of the specification.

Principle 1 (Dates are paid for). *A change of position on the time axis costs something the player can observe: real time, level progress, or a consumable resource. A date is therefore evidence of what was spent to reach it.*

Principle 2 (Seams are readable). *Operations that revise or interrupt history leave a mark in the player’s own record, and the game surfaces armed and hidden structure to players who have invested in perceiving it.*

Principle 3 (Effects, not claims). *What reaches a player from another player is the consequence of that other player’s choices after the signal has had time to arrive. Nothing about another player’s present is asserted to anyone.*

1.1 A second motivation: ecological proxies and collective selection

The design has a second root, which concerns money. The argument is that money behaves as a toxic meme: by making things quantifiable and comparable it sets up an optimization problem about acquiring more of it, which pushes people to categorize themselves and others by wealth and distorts what they attend to, in the way the arrival of a single bottle disturbs a moneyless society in *The Gods Must Be Crazy*. The reply is theory-laden, and the version adopted here is that money is a dynamic proxy and that the proxy should be shifted toward ecological awareness. The supporting proposals are hierarchical swarm piloting, in which the

passengers of a vehicle each pilot local drones and identify hazards through tablets or voice, invisible walls built from robotic animals that steer and monitor herds, and the conversion of every city into a kelp farm and rainforest generator.

The game is the experiment that gives these proposals a testing ground. Players try different terraformation schemes, technology trees, and takeoff rates, and the community then collectively selects the most coherent extrapolation. The field substrate of Section 3 is the physics of that experiment, the resources and crystals are its currencies, and the branch tree is how many extrapolations can be compared without being confused with one another. A further device from the same source is the counterfoil choice, described for the Kalabari, in which an option is offered mainly to show where a decision should not go, and Section 7 uses it for the reset key.

1.2 Conventions for status

The design draws on three kinds of record and this specification marks which is which. A recorded decision is something the designer stated in the design discussions. A proposal is an interpretation or a rule introduced in this document to resolve an inconsistency or fill a gap. An implementation fact describes code that was actually inspected. Appendix A lists the principal statements of the specification with their status, and a reader who needs to know whether a given rule has been decided should consult it.

2 The Time Model

2.1 The calendar and its endpoints

Play begins in the year 2400. The player may advance the year, and advancing is the cheap and default direction of travel. At the first level the passage of time is controlled by zoom: zooming out speeds time up and zooming in slows it down, which makes time control tactile and not a menu option. The earliest point on the axis is the Big Bang, about fourteen billion years before the start. A global reset, described in Section 7, returns a player to that earliest point, and the universe then unfolds back to 2400 over five minutes of real time, on a logarithmic schedule that runs fast at first and slows until the rate reaches parity with the ordinary clock. Backward movement short of a reset is gated: it becomes available only at the second level, which is reached by visiting a large number of stars, and it requires time crystals, which are a consumable and also act as a vector multiplier. The records state both that zoom changes the speed of time at the first level and that slowing time belongs to the time crystals. This specification reconciles the two by treating zoom as the coarse control that every player has over how fast their own simulation advances, and the crystals as the extension of control to direction and to fine adjustment.

The asymmetry between directions is deliberate and is built into the interface. Most players are expected to move forward easily and backward rarely. This mirrors the way time behaves for everyone and gives the backward direction the scarcity that makes it informative when it is used.

2.2 The tree of choices

The history of a player is a branching tree, and only branches that have been chosen or created are computed. A position on the time axis is a node in that tree, and the date displayed is the date attached to the node. Rewinding to an earlier date selects an earlier node, and a new choice made from there creates a new branch while the old line remains in the record. This is the same structure a version-control system maintains, and the analogy is useful because it makes the cost model explicit.

Once a player's parameter choices along a path are fixed, returning to the furthest point along that path is nearly free. A player who has rewound can spam the year-forward key and travel back up the existing branch with almost no effort, up to the earliest external event that newly applies to it, as Section 10 explains. The expensive operation is divergence, that is, making a different set of choices at an earlier node, because divergence is what creates new history. The five-minute climb from the Big Bang is not the price of time travel as such. It is the price of having discarded one's accumulated position, and it falls on a player who has used the reset rather than a rewind.

2.3 Interpolation

A player may be further along in the future than the point from which a revision is made. In that case a revision is not an edit to the past but an interpolation: the player fills in a path between two known endpoints, the earlier node and the later one already reached. This is a projected inquiry, in which the player poses a hypothetical path, examines its implications, and then commits it. The interface allows most players to go forward easily and does not allow them to rewrite backward without this projection, so interpolation is the ordinary method by which a recorded past is revised.

Rule 1 (Marking interpolated stretches). *A segment of a player's tree that was produced by interpolation and not lived through step by step is displayed differently from a lived segment, and a distant observer who eventually sees the consequences can see that the segment was reconstructed. Interpolation is a visible feature of the record and is never presented as though it were lived time.*

2.4 Endpoints and the shape of the history

Determinism makes replay from an earlier state well defined, but it does not guarantee that a newly chosen sequence of interventions will arrive at a later state that has already been fixed. Connecting two prescribed endpoints is a constrained reachability problem, and the specification must say what the later endpoint means. There are three strengths. An endpoint may be a date only, an endpoint may be a set of constraints on the world state that the reconstructed path must satisfy, or an endpoint may be the exact world state. This specification takes the exact state as the default meaning of an interpolation, so that a projected inquiry can fail. The promise the game can safely make is narrow: it tests a proposed sequence of interventions against the endpoint and rejects any candidate that misses it. That a candidate fails does not show that no permitted sequence succeeds, and the word infeasible is reserved for cases that a solver can certify, by an exhaustive search of a finite space, by a proof of unreachability, or by a

complete solver for the restricted system. Equality is taken on an explicitly defined simulation state, which excludes provenance and node identity, since those necessarily differ between the original endpoint and its reconstructed counterpart. A request that fixes only a date, or only some constraints, is not an interpolation in the full sense, because the reconstructed branch may reach the same year with a different world, and it is recorded as a fork with an explicit marker of which constraints were kept. The distinction is the one between exact and inexact values in a language such as Racket, where a comparison on inexact numbers can still yield an exact result, and it serves the same purpose here: to keep what was promised separate from what was only approximated.

A branching history is a tree as long as every state has exactly one parent, and an interpolation that joins two existing nodes would give a node two parents and turn the tree into a graph with merging paths. The specification preserves the tree. An interpolated path ends at a new node that is constrained to equal the old endpoint, and the new node carries a provenance link to the original, so that the history of states remains a tree and the relation between a reconstructed state and its counterpart is a recorded annotation and not a structural merge. Anyone who prefers to allow merging may call the result a directed acyclic graph, but the replay rule and the cost model are simpler on the tree.

2.5 Two clocks

The game keeps the two clocks that the social platform conflates. The year shown in the world is the in-game coordinate of the player’s current node. Alongside it the game maintains a record of how that node was reached: the jumps taken, the resets suffered, the crystals spent, and the stretches that were interpolated. A reset is therefore a mark in the history and never an erasure of it. A player reading their own tree can distinguish a position reached by ordinary advance from one reached by a costly rewind.

2.6 Causal time and cosmological time

A record of provenance does not by itself supply a coordinate for scheduling, and the specification adds one. Each branch carries a causal coordinate τ that increases monotonically along every path of the history tree and is distinct from the displayed cosmological year Y . A reset changes Y and selected world variables, returning the calendar to the Big Bang, and leaves τ increasing, so that a player whose calendar has returned to the beginning has not thereby returned to an earlier causal moment. A rewind selects an ancestor and begins a new continuation whose τ resumes from that ancestor’s value. Signal travel, cooldowns such as the refractory period of a retired scope, event ordering, and warning countdowns use τ , while cosmological development, the long-term state of the fields, and the displayed calendar use Y . Events received before a reset therefore stay received, and signals already emitted keep an unambiguous age, which would not be so if arrivals were scheduled by a calendar that a reset can wind backward.

Multiplayer also needs a correspondence between the causal coordinates of different players. The engine keeps one in a global causal frame, in which each continuation segment is placed at the moment it is created, and every delivery of an event is computed in that frame. The correspondence is never shown to the players, and its details are left open.

3 The Field Substrate

3.1 Three fields

The strategic identity of the game was fixed in an early design discussion as a 4X game in which players manipulate a scalar field Φ , a vector flow field $\mathbf{v}$, and an entropy field S . The objective is to sustain organized structures under a constant smoothing pressure, so that persistence, and not accumulation, is the governing strategic problem. Exploration, colonization, production, and conflict all operate inside this physical environment, and entropy is a working part of the simulation and never only background lore. The scalar field Φ represents constraint density, that is, organized structure. The vector field $\mathbf{v}$ represents directional flow. The entropy field S represents the dissipation of organization. The simulation advances by three processes: diffusion, which smooths the fields; flow reorientation, which redirects $\mathbf{v}$ in response to the other fields; and dissipation, which converts organization into entropy.

The recovered design names these processes without fixing their equations, so the form below is offered as a candidate that respects the stated structure and leaves the coefficients open:

$$\partial_t \Phi = D_\Phi \nabla^2 \Phi - \nabla \cdot (\Phi \mathbf{v}) - \kappa S \Phi + \sigma_\Phi, \quad \partial_t S = D_S \nabla^2 S - \nabla \cdot (S \mathbf{v}) + \alpha_\nabla \|\nabla \Phi\|^2 + \gamma a + \chi \kappa S \Phi + \sigma_S - \mu_S S - \pi,$$

where a is the density of operational throughput by player structures, which produces entropy at rate γ , α_∇ is the rate at which gradients of organization produce entropy by relaxing, μ_S is the linear damping of the prototype, σ_Φ and σ_S are player-supplied sources, π is entropy exported by pumps into a finite external reservoir described in Section 16, and the flow $\mathbf{v}$ relaxes toward a direction determined by the gradients of Φ and S plus player redirection. The term $\kappa S \Phi$ is the dissipation pressure that makes unmaintained structure decay, and it is the term whose absence from the first Python prototype is discussed in Section 16. The coefficient $\chi \geq 0$ is the conversion of dissipated organization into entropy, in the units used for S , so that the equations and the prose agree about where lost structure goes. The symbol γ here is the operational coefficient and corresponds to nothing in the prototype document, whose own γ multiplies $|\nabla \Phi|^2$ and is written α_∇ in this specification. The three production terms are separate mechanisms with separate coefficients, and Section 12 states the rule that keeps them from counting the same process twice. The equation is a proposed phenomenological form whose coefficients and units remain to be specified, and it does not inherit an energy-decay theorem from the prototype. Diffusion acts instantaneously across a connected domain, so these equations are a local approximation to be run inside causally isolated simulation patches. They do not enforce the light-cone rule across the galaxy, which needs a separate causal transport mechanism for the signals between patches.

3.2 The lattice

The browser prototype stores the four component arrays Φ , S , v_x , and v_y on a 256×256 toroidal lattice. The torus removes boundary effects, so that the only edges in the world are those the players and the fields create. The initial player actions are seeding Φ , pumping entropy, applying shocks, and redirecting flow. These four verbs are the primitive operations from which every strategic action in Section 11 is built.

3.3 Determinism and the branch tree

The field substrate is what makes the time model of Section 2 inexpensive to implement. If the simulation is deterministic given a seed and a sequence of player interventions, then a player's history need only store the interventions along the chosen branches, and any position in the tree can be recomputed by replaying them from the root. This is the practical meaning of the statement that only chosen or created branches are calculated. It also fixes a requirement on the implementation: every source of randomness must be seeded and recorded as part of the branch. The first Python prototype used unseeded randomness for its scan command, which would break replay and must be replaced before the time model can be built on it.

3.4 Scopes as organized regions

The nested scopes of Section 4 are the spatial form that organized structure takes in the field model. A natural reading, adopted here as a design proposal and not a recovered fact, is that a scope is a region of elevated Φ bounded by a level set, that its boundary is where the constraint density crosses a threshold, and that the core is the maximum of constraint density within the region. Popping a scope then corresponds to injecting entropy at the boundary, and the inward chain reaction corresponds to that entropy being carried along the flow toward the densest point, where the damage to organization is greatest. On this reading the batteries, escape routes, and reactors that attract the guide bot are the local maxima of useful structure, and the core triggers sit near them for the same reason that valuable infrastructure is where a disruption does the most harm.

4 The World: Scopes, Boundaries, and Cores

4.1 Nested scopes

Space in Entropy's Edge is organized as nested spheres, which the game calls scopes. A scope has a boundary, and scopes contain other scopes, so a star system sits inside a larger domain and contains planets, installations, and fields of its own. The structure is the one used by the SpherePOP calculus on which the game draws.

The calculus has a lexical reading that is worth stating, because it explains the vocabulary. A player fires a crystal by pressing the space bar, t, or y, in the way that a weapon is fired in Descent. When the crystal has travelled a specified distance it grows into a bubble, and the bubble corresponds to the opening and closing of parentheses in Racket. Forms such as define, lambda, and let bind identifiers, the part of a program in which a binding applies is its scope, and the set of bindings in effect for an expression is its environment. In the game a bubble is a scope, a player's environment is the set of bubbles enclosing the ship, and a bubble can be popped with another crystal or entered without popping by blinking. The ship shrinks as it moves forward, so that nesting is also a change of scale.

4.2 Popping

Passing through the boundary of a scope pops it. A pop sets off a chain reaction that travels to the live core of the domain, an innermost nested scope where the trigger resides, along the

routing rule stated below. In the ordinary case, in which the core lies inside the crossed scope, the reaction travels inward. The consequence of a pop is therefore not located at the boundary that was crossed but at the center of the domain, and it reaches the center after a delay. The reactor explosion at the end of the game Descent has the same shape: the cause is internal, the effect travels, and there is a window during which the player can still act.

4.3 The core

The innermost armed scope is called the core. It is the one place where a pop is irreversible. Outer scopes are forgivable in the sense that the blink mechanism of Section 8 can carry a player through them without firing anything, but the core cannot be shielded by any tier of blink. Skill in the game consists largely in reaching and leaving a domain without touching its core.

4.4 Defusal and relocation

The owner of a domain may deactivate the mechanism in its central scope. Deactivation takes effect instantly across the scope, and it does not make the domain safe. It changes which scope is the live core, because the armed state is conserved: when one trigger is retired, another is created elsewhere in the system. The replacement is not chosen by the owner. It is determined by relevance, in the manner of a protective boundary that absorbed the first-ranked case so that the next most probable scenario is activated. The live trigger is therefore the head of a ranked queue, and defusing removes the head and promotes the runner-up.

In a simple environment the ranking is legible from the layout. The comparison is the main breaker of a house: the breaker panel is built so that failures cascade in a known order, and someone who has wired a building can predict where the next failure goes. In large nested environments the ranking is only inferable from structure and, where it applies, from traffic, and reading an environment that nobody has labeled is the upper end of the skill range.

Rule 2 (Conservation of the armed state). *At every moment each domain has exactly one live core. Deactivating it creates another, and an owner cannot place the new core by choice.*

Rule 3 (Instant change, delayed observation). *A deactivation is instantaneous within the domain. An observer in another system continues to see the earlier configuration until the corresponding signal arrives, so a visitor may plan around a trigger that has already been moved.*

4.5 Eligibility, queue cycling, and propagation paths

The conservation rule needs a defined domain of application. A domain has a finite set of eligible scopes, those innermost scopes that can be armed, and the live core is the highest-ranked eligible scope. Deactivating the core promotes the highest-ranked eligible scope other than the one retired, which is guaranteed to exist only if the domain has at least two eligible scopes. A domain with a single eligible scope therefore cannot be defused, and its only core can be removed only by dismantling the domain. A retired scope is not lost: it re-enters the queue at the tail after a refractory period, so that repeated defusal cycles through the eligible scopes and does not exhaust them. A defusal is refused while every other eligible scope is still refractory, so a domain with two eligible scopes cannot be defused twice in quick succession, and the

refusal is reported to the owner as a postponement until a retired scope becomes available again. Nothing in the rule makes the replacement armed against any particular visitor, since it is armed only in the sense that it is the live trigger of the domain.

The specification uses a single routing rule throughout. A domain has exactly one live core, and a crossing that is not blinked sends a reaction along the unique path in the tree of scopes from the crossed scope to that core, upward to the lowest common ancestor and then downward. In the ordinary case the core is a descendant of the crossed scope and the reaction simply travels inward, and in other cases it first travels outward, so that every unblinked crossing in a domain reaches the same core. The delay of the reaction grows with the length of the path. This replaces any rule that chooses among the children of a scope by rank, since a highest-ranked child need not lead to the live core. The ranking is used only to decide which innermost scope becomes the live core, with ties broken by a deterministic order that is part of the seed, so that replay reproduces the same choice. The live core is also the one scope that a blink does not protect.

5 The Galaxy and Its Contents

5.1 A galaxy built from articles

The stars are named after a curated Wikipedia watchlist, and the structure of each article furnishes the structure of its star system. The planetary systems are extracted from an article's section headings, and the tunnels and surfaces from the content of the corresponding sections. The recorded example is a star named Free electron model. The geographic zones of its first planet are named after the first twenty-six words of the article's first section. The first planetary system is a planet and its moon, named Ideas and Assumptions, the next system is named Properties of an electron gas, and so on through the headings. A player who wants to learn about a subject flies to its star. The galaxy is therefore a map of ideas as well as of matter, and the difficulty of finding a particular place is the difficulty of finding a particular idea.

5.2 Resources and progression

Level one has three resources, boranium, ironium, and germanium, inherited from Stars! (1995). Level two adds two more, texture crystals and time crystals. Texture crystals fill in textures in outline-only zones and allow custom luminosity filters and shaders. Time crystals allow backward travel in time and act as a vector multiplier, and each time crystal sets the game back one billion years for all players, so that fourteen of them are needed to set the game back fourteen billion years without a game over message being issued. One million texture crystals can be exchanged for one time crystal when none can be found, and the chance that a given planetary system holds a texture crystal is one in a million, so time crystals are very rare among level-one players.

Level two is reached by visiting a fixed number of stars. An earlier record gives the number as one hundred and fifty thousand and a later one as one hundred thousand, and this specification treats the count as a parameter N_2 with the later figure as its default. At that scale a player may need years to advance, which is part of the intended difficulty, and fleet automation, scouting swarms, and programmed waypoints are the means by which the burden of visiting is

made bearable.

A consequence of the stated numbers deserves attention. If every planetary system independently has a one-in-a-million chance of a texture crystal, then acquiring fourteen time crystals by exchange requires fourteen million texture crystals, which is on the order of 1.4×10^{13} systems in expectation, assuming independent trials, one crystal per successful system, and acquisition entirely through exchange. It is an expected sampling requirement and not a guaranteed one, and it is far beyond the 10^5 stars needed to reach level two. A crystal-funded rewind of the whole universe is therefore effectively unreachable by ordinary play and is a limit case. Whether this is intended or the exchange rate should be revisited is recorded as an open question.

5.3 Flares, beacons, and bookmarks

A press of the g key shorter than half a second fires a flare when no acceleration boost is available, and gives a brief boost when the player is accelerating with the w key. A flare is a permanent beacon that sticks to walls, glows, and illuminates dark passages and tunnels. It also sticks to a bubble without popping it. Because it persists, it is a persistent feature of the world, and it lets a player find the way back out of a deep maze in an asteroid or planet. A flare fired at a star bookmarks it, so that the star remains visible on the map and when the view is zoomed out.

Under the time model of Section 2, a flare belongs to the branch in which it was fired. Rewinding to a point before the firing removes it from that line of history, and a flare fired on a branch that is later abandoned remains part of that branch's record. The permanence of a flare is a property of a world within a branch and never across branches.

5.4 Removing or enclosing a star

A player who dislikes a particular star may drop a molecular disassembly device into its sun, or build a minefield or a Dyson sphere around it. The enclosures are treated in Section 9.

5.5 Trajectories as media

Every idea has a force-based location in a global space, in the manner of a word-embedding space, so that locations correspond to skills or phonological markers. Trajectories flown through that space at various blink rates correspond to forms of media such as films and music albums. Certain trajectories are equivalent to, and backward compatible with, flashcard courses and certain books, so that a player may unlock a trajectory in the vector space by reading a book or by mastering a deck, in sequence or not, and may do so without a computer. The recorded design does not specify which books or decks correspond to which trajectories, and that mapping is open work.

The center of each star contains a typing tutor, where a player ends up by flying around randomly or ineptly. The relation of this place to the core trigger of Section 4 is not settled in the record and is listed among the open questions.

6 Controls

The control scheme borrows from the vim editor, and each binding has a mnemonic expansion that is part of the game’s teaching surface. Looking around is done by moving the mouse, by touch, or by the keys h, j, k, and l. The world also has a color wheel, which rotates when the arrow keys or the mouse are used, and the keys h, j, k, and l look around without rotating it.

Action	Key	Mnemonic
Pitch forward	Down or j	jump down
Pitch backward	Up or k	klimb up
Turn left	h	hop homeward
Turn right	l	leap letterwards
Slide left	a	ad sinestra
Slide right	d	dextra
Slide up	r	roofwards
Slide down	f	floorwards
Bank left, lower contrast	q	question everything (vim: begin recording a macro)
Bank right, raise contrast	e	examine your biases (vim: end of the word)
Accelerate	w	walk
Reverse	s	stop, slowdown, fly backwards
Fire primary, or draw a swype trace	y, t, or space	yes, trigger
Fire secondary	u	use missile
Flare or boost	g, tapped	see Section 7
Drop bomb	dedicated key	for extraplanetary minefields only
Rear view	o or p	original perspective, previous places
Automap	tab or semi-colon	
Blink	b	refresh the visuals, enter a bubble unpopped
Autoblink	a and b together	see Section 8
Super autoblink	s, then a, then b	see Section 8

Table 1: Principal bindings.

6.1 Rotation, contrast, and blinking

Rotating left lowers the contrast, with q, and rotating right raises it, with e. A player can rotate without changing the contrast by blinking at the same moment, or by having autoblink enabled. The arrow keys both pitch the ship and rotate the color wheel, whereas j and k pitch without rotating it, and the specification keeps this duplication deliberate, since it is the first

place where the difference between plain and blinking input becomes visible to a new player.

On the Vim correction: the key that begins recording a macro is the lowercase q, and the uppercase Q traditionally enters Ex mode. The mnemonic is attached to the lowercase key.

6.2 Modes

There is an insert mode for entering text and a command mode, entered with the escape key, in which the player controls the guide bot on the map. In command mode a doubled g goes to the beginning of the path and a single g goes to the end, and the space bar returns the player to firing mode. Vim uses an uppercase G for the end of a file, and the record uses the lowercase key, so which is intended is left open. In either case the g of command mode is a different command from the g of firing mode, which is the key governed by the hold-time rules of Section 7.

6.3 Chords and sequences

Autoblink is enabled by holding a and b together, and super autoblink by pressing s, then a, then b. Because a and s are also translation and reverse keys, the input rules must say what the ship does during activation, and the specification fixes the following. Once a chord is recognized, the translation produced by its component keys is suppressed for as long as the chord is held, so that enabling a mode never moves the ship. In the ordered super sequence, the brief movement caused by s and then a before b is pressed is accepted, and it is part of what makes the top tier harder to enable by accident. A chord is recognized when its keys are held together beyond a short window T_c , which is a tuning parameter.

7 The G Key and the Global Reset

7.1 Hold time

The g key behaves differently according to how long it is held. A press shorter than half a second is a tap, which gives a brief acceleration boost to a player who is accelerating and otherwise fires a flare. A hold longer than half a second but shorter than the reset threshold t_h does nothing when released, but it is not silent: an alarm tone builds from the half-second mark, so that the player is warned before anything irreversible can begin. Reaching t_h starts a countdown, with a loud alarm and a flashing warning that the universe will be reset in c seconds, and releasing the key during the countdown cancels it. Only a key held through the whole countdown fires the reset. The interval between half a second and one second, which the first formulation left unspecified, is therefore the warning zone, and the five-second first-use countdown communicates its warning through sight and sound and not through silence.

Both quantities shrink as a player accumulates warnings. The threshold t_h starts at one second and scales logarithmically toward half a second over many tries, so that eventually any press beyond a tap enters the countdown. The countdown c starts at five seconds and loses one second after each warning until it reaches one second. Whether the five seconds of the first countdown include the threshold interval is a parameter and is left open.

7.2 Three kinds of exposure and protection

The earlier text of this specification used the word inoculation for several different things, and the review of that text showed that they are not equivalent. Three effects of a reset must be distinguished.

The first is temporal exposure. The reset returns the game to the Big Bang, fourteen billion years back, and the universe unfolds again to the year 2400 over five minutes, on a logarithmic schedule that is fast at first and slows until it reaches parity with the normal clock. Every player the reset reaches goes through this, whether or not they lose anything by it.

The second is loss. A game over is issued to every player who is not blinking when the wave reaches them, which autoblink makes less likely without guaranteeing. The player's points are set to zero, the dictionary of planets visited is wiped, and the names of all planets and stars are replaced by the words game over for five minutes, which makes some of the minigames harder. A player who is blinking at the moment the wave arrives is still reset in time but does not suffer this loss and keeps their points. Enabling autoblink does not by itself provide this, since the protection requires an actual blink at the moment of arrival, and autoblink serves by keeping the player blinking often enough that arrival is likely to coincide with one. Super autoblink gives the same retention whether or not the player happens to be blinking, which is what distinguishes the two tiers.

The third is insulation, which is exemption from temporal exposure itself, and it comprises two properties that must not be confused. Blocking an incoming reset is one property: a player whose stars lie inside a completed Dyson sphere is invulnerable to resets that reach them. Preventing a reset from being emitted is the other: an insulating minefield prevents reset commands from being fired within it, which protects the players outside it and does nothing for a player inside, who is then contained. The invisibility of a Dyson sphere means that no further signals from the system escape, so that observers see nothing new from it. It does not erase light already emitted, and the earlier appearance of the system remains in the skies of distant observers until later arrivals fail to update it. Neither blink tier provides any of these properties.

Protection	Against loss	Against temporal exposure	Condition
None	no	no	
Autoblink	yes	no	only if blinking when the wave arrives
Super autoblink	yes	no	unconditional
Dyson sphere around the player	yes	yes, for incoming resets	structure must be complete
Insulating minefield around a source	for players outside	for players outside	suppresses emission from inside

Table 2: The three kinds of protection against a reset. The blink tiers protect a player's accumulated organization and do not protect their position in time.

This table also removes the apparent conflict between the earlier sections. The statement

that autoblink protects against the reset means that it protects against its loss, the statement that a player is reset in time but retains points describes the same protection, and the formal requirement that a local patch be approximately preserved is the mathematical form of retention.

7.3 Propagation and the light cone

The reset spreads outward from its source as a wave at the speed of light. A player is exposed when the wavefront reaches them and not when the key is pressed, so a distant player learns of the reset only on arrival, and for such a player the arrival is at once the warning and the event. This respects the rule that no information about another player's present can travel faster than light. In particular, super autoblink can respond to the arrival of the wave but cannot respond to its emission, because responding to the emission would require information that the multiplayer rules of Section 10 forbid. The recorded statement that the reset applies to all players is therefore read as applying to all players the wave reaches.

7.4 Disincentives

The reset is surrounded by disincentives, and the design hopes that no one ever holds the key that long. The location of the player who presses it is broadcast to the other players. The nearest players may build an insulating minefield or a Dyson swarm around the offender's stars and planets, which inoculates them against that player's reset and contains the source. The offender's points are set to zero and the dictionary of visited planets is wiped. Players are asked to be respectful of newcomers, and the advice to everyone is to enable autoblink as soon as they begin to play.

The mechanism has the form of what the Kalabari call a counterfoil choice, in which one option is offered chiefly to show where the decision should not go. The reset is the empty bowl: it is placed next to a useful action on the same key, its alarm shows the player where the choice leads, and the design depends on the norms of the community that the display is meant to reinforce and not on prohibition.

7.5 Crystal-funded rewinds

The recorded rule is that each time crystal sets the game back one billion years for all players, and that fourteen of them set the game back fourteen billion years without a game over message. The specification reads this as a reset financed by resources, in which the loss of the preceding kind is not issued, so that it is the considerate alternative to holding the key. Whether temporal exposure still applies to other players is not stated, and the exchange numbers make the full rewind a limit case, as noted in Section 5.

7.6 Heat death and the annual cycle

An earlier version of the design had the calendar reset to 2400 every January 1, because someone always triggers a runaway heat death on December 31, as the trigger mechanism of interstellar mines becomes more sensitive with time. A player who is a grabby alien, having visited every star in the observable universe long before that date, simply waits for it. The specification

proposes a reconciliation with the rest of the time model, and it remains a proposal until the designer adopts it.

The proposal treats January 1 as an annual scenario boundary, while the Big Bang reset remains an event inside a scenario. The two mechanics then do not compete for control of one clock. The annual cycle follows the real-world calendar in a declared server timezone. Each annual scenario begins with the displayed year at 2400, and players then advance, rewind, and suffer resets under the rules already stated. January 1 starts a new scenario identifier and does not rewind the previous scenario. The completed scenario and its branch histories are archived, so that neither provenance nor observations are erased. The causal coordinate τ of Section 2 continues to increase across the transition, and the scenario identifier records which environment an event belongs to.

Under this reading the December 31 heat death is a terminal environmental process whose hazard rises with the age of the scenario. A player-triggered Big Bang reset changes the displayed cosmological year and does not lower that hazard or restart the annual schedule, since otherwise repeated resets could postpone the annual ending indefinitely. The five-minute unfolding belongs to Big Bang resets only, and the January transition initializes the next scenario at 2400 directly.

The annual transition is an engine rule and not a wave emitted somewhere in the galaxy. Players may know its schedule in advance without learning anything about another player's present. Heat-death disturbances inside the closing scenario still obey causal propagation, but the scheduled retirement of the scenario does not need to masquerade as a physical signal. Dyson insulation therefore protects against reset waves within a scenario and does not prevent the scenario from ending.

Events retain their scenario identifier and remain attached to that scenario's histories. A signal still in flight at the annual boundary becomes an archived pending event. It can arrive if that history is later resumed and cannot enter the next scenario, which keeps the branch-assignment rule intact and prevents an old reset wave from striking a newly initialized universe.

The default proposed for what carries between scenarios is the player's archive, learned controls, accessibility settings, and historical achievements, together with a fresh strategic state. Whether visited-star knowledge or earned resources also carry over is a separate progression decision that the annual-clock rule does not settle. In summary, a reset restarts cosmological development within a recorded history, and January 1 starts another annual scenario while preserving the record of the previous one.

7.7 Experimenting with resets

A player may build a Dyson sphere around their own system in order to experiment with the global reset, and the recorded rationale is that the reset is the only way to reach extra-ocular regions, those outside the particle horizon. The mechanism by which a reset opens such regions is not specified.

8 Blink, Autoblink, and Super Autoblink

8.1 Blink

The b key blinks. A blink refreshes the visuals when there are too many texture tracers and Sobel edges on screen, it allows the player to enter a bubble without popping it, and it allows the player to rotate without changing the contrast when pressed together with q or e. Because outlines grow thicker and blurrier over time, a player who is not moving should blink about every five seconds or move. In the design's own terms, blinking is correct traversal: it is how a boundary is crossed without firing it.

8.2 Autoblink

Holding a and b together enables autoblink, which blinks automatically when necessary. It covers the outer scopes and does not protect the core. It prevents a player from accidentally popping bubbles while moving through nested space, and it keeps the player blinking often enough that a reset wave is likely to arrive in the middle of a blink, which is why the notes describe it as increasing the chances of not losing progress and not as guaranteeing them. Experienced players keep it on almost continuously, so it functions as a filter that separates brand-new players from those who have learned the traversal and grants little power beyond that. The design describes the relationship in two phrases: blinking is correct traversal, and autoblink is delegated traversal, in which the player automates their own correct crossing. Early formulations wrote the activation as a followed by b, and the specification fixes the chord as canonical, with the input rules of Section 6 governing it.

8.3 Super autoblink

Pressing s, then a, then b enables super autoblink. It blinks automatically when passing through an outer scope and also when the reset wave arrives at the player's location, so that retention does not depend on timing. It is the unconditional tier, and the ordered sequence is meant to make it harder to enable by accident and to read as a layer stacked on top of autoblink.

Super autoblink also adds perception. While it is enabled the player sees an amber warning on the areas that will set off triggers, and specifically on the armed innermost scopes that no blink tier protects. The overlay shows the armed configuration that is locally observable to the player. A change made elsewhere appears only after its signal arrives, so that the overlay never reveals the remote configuration earlier than the light-cone rules allow, and within the player's own domain it updates at once when a core is deactivated and replaced.

8.4 What blinking does not do

Blinking does not disarm anything. It prevents the scope being crossed at that moment from sending a reaction, and the scopes nested inside it remain armed. The one exception to that protection is the live core itself: a blink never prevents the core from firing, and no tier of blink does. Risk therefore accumulates with depth, but the structure is one in which the outer shells are forgivable and the center is the single irreversible point. A player with autoblink can enter another player's domain cleanly and unnoticed, and the owner is not told of the visit unless the

Tier	Input	Protects against	Perceives
Blink	b, by hand	popping the scope crossed	nothing added
Autoblink	a and b together	accidental pops in outer scopes; loss from a reset if blinking on arrival	nothing added
Super blink	auto- s, then a, then b	outer-scope pops; loss from a reset unconditionally	amber markers on armed cores

Table 3: The three tiers. No tier protects the core, and none protects against temporal exposure.

visitor trips the central trigger or the owner’s own choices have consequences that later reach the visitor.

9 Minefields, Containment, and Salvage

The bomb key drops a mine and is to be used only for extraplanetary minefields. Mining is a deliberate action with its own binding and is not a side effect of any other mechanic. Each mine is a domain consisting of a single armed scope, so every mine is its own live core and cannot be relocated, and a minefield is a set of such domains. The alternative reading, in which a field is one domain with exactly one live mine at a time, yields a different game in which defusing a mine promotes another, and it is recorded as a variant. The base specification uses the first reading because it matches the salvage economy below, in which a mine is removed by being scrapped and not by being moved.

9.1 Containment and insulation

A minefield has two roles, which the transcripts record. It contains a player who has announced their location by holding the reset key, since a field built around their stars and planets means they cannot easily leave. It also insulates the builders, because a field prevents reset commands from being fired within it, so that the players who build it are inoculated against the offender’s key. The nearest players do this, and the farther players are protected by their distance and by the delay of the wave.

9.2 Dyson spheres and swarms

Dyson swarms and spheres serve two purposes. The nearest players build swarms around an offender’s system to protect the rest of the galaxy. Advanced players build spheres around their own systems because they do not wish to reset the universe for others, and a completed sphere makes a star invisible to other players and invulnerable to resets. A sphere is not built in real time by the player. The player sets up the infrastructure by selecting a suitable technology tree and setting the factory defaults correctly, and the completed sphere is visible only far in the

future, so that a player must travel a long way forward in time to see it finished. A player who wants to experiment with the reset can build one around their own system.

9.3 Hidden until tested

A trapped player does not learn of a minefield until they try to leave their solar system. There is no notification, and a player who has announced their location with the reset and then lives quietly for years cannot tell whether a fence was built around them. The absence of information is part of the cost of announcing.

9.4 Who is caught

A well-placed mine does not rely on outer scopes, because autoblinkers cross those cleanly. Its value lies in where the core trigger sits relative to the route out of the system. A minefield is therefore a bet about which path a trapped player will take, and a counter-bet is available in super autoblink, whose amber overlay shows the armed cores and thus the field, as far as the player's own light cone reveals it. Newcomers and players without perception are caught, and experienced players with the top tier are warned.

9.5 Salvage

A mined player is not trapped permanently. Once the minelayer owns skimmers, the field can be scrapped for resources. The mines are capital and the scrap is the return, which gives other players a reason to mine a system that has nothing to do with hostility toward its occupant. Containment is an investment with a payoff schedule, and the length of a player's confinement depends on someone else's progression. Fleet infrastructure supports the work: mass accelerators and starbases build drone scouts, miners, mine-layers, and transports, and waypoints with repeating orders allow resource distribution networks to be programmed, so that salvage can be carried out by fleets sent ahead.

Rule 4 (Retirement without removal). *Because deactivation relocates the live core and does not remove it, the owner of a nested installation with several eligible scopes may retire its armed state by shifting the trigger to another scope while leaving the old layout looking armed. A player who passes through a deactivated installation may never learn whether it was dangerous. Single-scope mines are not subject to this rule, since they have nowhere to move their trigger.*

10 Multiplayer: Isolation and Light-Cone Verification

10.1 Scattered players

Players are scattered so far apart that they cannot influence each other directly. Each player's tree of choices is private, so there is no shared timeline to keep consistent and nothing to reconcile. Two players who are both at the year 2400 are not at the same time in any sense they can check, because there is no event that both could observe. A year is a coordinate within one's own tree, and the phrase *when* amounts to how many of one's own choices deep one is.

10.2 Travel as verification

Interaction is possible, but it is paid for in travel time. A player who travels to another player twenty-six thousand light years away must spend twenty-six thousand years of travel to arrive, and cannot see what the other player is doing now. What the traveler can see is the implication of the other player's choices as they have propagated: the consequences that the signal has had time to carry. The player sees effects and not claims.

This inverts the defect that motivated the design. On the social platform, a stranger's date is a label the reader must trust. Here, a stranger's date can only be inferred from the downstream evidence of their choices, which is far harder to forge than a field, and the cost of seeing it is time actually spent.

10.3 Deniable presence

Visibility is gated by causal effect. A visit leaves no mark on the visited player unless that player's choices carry consequences into the future the visitor travels to, and even then what registers is the consequence, not the visitor. The record of a visit exists only in the visitor's own tree. Presence is deniable by default, in the way that a night sky never tells you who is looking at it but only what the light from you did. A player whose choices have no downstream effects is permanently invisible to everyone, and quiet play is a legitimate way to stay unseen. The global reset is the one deliberate exception, since it announces the location of whoever triggers it.

10.4 Rewrites and what observers see

An observation made across distance is true of the state that was committed when the signal left. A later revision by the observed player, made by projected inquiry and interpolation, forks from that point, and the observer's view stays true of the branch they saw. The observer learns of the revision only when its own signal arrives. Abandoned branches therefore persist in the observer's record as something that really happened to the version of the other player that they observed.

10.5 Event transmission and branch assignment

Isolation postpones the problem of consistency between private histories but does not remove it. A visitor can observe a branch, act on what was observed, and send an effect back, and the specification needs an explicit rule for which branch receives that effect. Each transmitted event is therefore recorded with four parts: the origin branch, meaning the node in the sender's tree at which it was emitted; the emission coordinate in space and time; the destination; and the arrival condition, which is the light-travel delay together with any further condition the event carries.

An event is delivered to a branch of the recipient's tree if and only if its arrival condition is met on that branch. A reply to an event is addressed to the emission node of the event it answers, and it is delivered to every branch in the original sender's tree that contains that node among its ancestors, inclusive. A branch that was forked from a point before the emission does not receive it, because on that branch the original event was never sent. Rewinding does not

retract an event already in flight. Pending arrivals remain attached to the branches from which they were emitted, so that a message sent on an abandoned line continues to travel and to arrive as a fact about that line, and a new branch, if it wishes to send a message, emits its own. A player who rewinds therefore cannot erase what has already left, but can ensure that its consequences fall on a version of themselves they no longer inhabit, and a player who later returns to the abandoned branch by cheap forward replay will find the replies that arrived there in the meantime. The rule is the one that keeps the observer’s account true: what someone observed was real on the branch they observed, and the answer they sent is addressed to that same branch.

Incoming events must also be part of the replay record. A computed node cannot be reproduced from its seed and the player’s interventions alone if an external reply is later inserted into its past. When a new arrival would alter an already computed stretch, the stretch is preserved and a new continuation is created that incorporates the arrival. An inactive branch can hold a pending event without being simulated, and when it is resumed it advances to the arrival, applies it, and only then computes later states. Replay is therefore unchanged only as far as the earliest external event that newly applies to it.

The scheduling coordinate for arrivals is the causal coordinate τ of Section 2, and the engine keeps a hidden correspondence between the causal coordinates of different players, which players never see.

11 The Strategic Layer

11.1 The four activities reinterpreted

The four conventional activities of a 4X game receive interpretations drawn from the field substrate. Exploration is the discovery of gradients and of useful field configurations, so a map is a record of where organization is cheap to maintain and where it is not. Expansion is the seeding of constraints that can persist, which is a different thing from claiming territory, because a seeded region that the dissipation term overwhelms is lost without any opponent acting. Exploitation is the capture of flows and the export of entropy. Warfare includes injecting entropy into an opponent’s structures and disrupting the conditions that support them, so that the weapon of choice is often an attack on the environment and not on the occupant.

11.2 Infrastructure

The design names four kinds of infrastructure: Autocatalytic Nodes, Constraint Builders, Entropy Pumps, and Memory Crystals. The names were proposed in the design discussion, and the record does not establish that all of them became functioning objects, so the roles below are stated as the natural field-level reading and not as recovered specifications. A Constraint Builder is a source of Φ , the term σ_Φ in Section 3. An Entropy Pump is the export term π . An Autocatalytic Node is a source whose output grows with the local value of Φ , so that a well-maintained region reinforces itself and a neglected one loses its reinforcement. A Memory Crystal is a structure that preserves organization against dissipation, and its role in the time model remains to be decided: the consumable time crystals that slow the passage of time share part of the name, and the two may be the same object or related ones.

11.3 Stars, corridors, and fleets

Star systems are the nodes of the strategic map. Each star samples the local values of Φ , S , and $\mathbf{v}$, and those samples determine its character. Corridors are connected regions of low entropy, detected from the field, and they are the routes along which a civilization can maintain contact with its settlements. Travel cost depends on distance, on entropy along the route, and on corridor modifiers. Mass accelerators and starbases are the hubs of this network, and starbases build the drone scouts, miners, mine-layers, and transports that carry out its work. Fleets follow waypoints with repeating orders, which lets a player program resource distribution networks for the three level-one resources, and the orders may be conditional, for example mining one resource only while a stockpile is below a given fraction. Colonization requires a habitable star that nobody owns. The result is a geography in which routes and settlement opportunities depend partly on conditions that the players can themselves alter, and overextension arises from the simulated conditions: a civilization that spreads beyond what its corridors can keep organized loses the outlying settlements to dissipation without an enemy being involved.

Fleet movement and colonization were chosen as the next development step after the economy, the star-system nodes, and corridor detection, and that sequence of decisions is a fixed point of the plan. One consequence of the light-cone rules of Section 10 requires attention. Travel between distant players takes as many years as there are light years of separation, so corridor modifiers can change the cost of a journey, in energy or in entropy exposure, but cannot shorten its duration below the light-travel time without contradicting the isolation of players. This specification therefore treats corridors as modifiers of cost and not of speed.

11.4 Factions

Five factions were proposed: the Lamphron Ascetics, the Lamphrodyne Industrialists, the Vectic Nomads, the Crystal Archivists, and the Parasite Syndicates. Their names suggest different relationships to organization and flow, an ascetic discipline of constraint, an industrial throughput, a life carried on the flow, a custody of stored structure, and a dependence on the organization of others, but the recovered record does not contain their original descriptions, bonuses, or histories. Those are open design work, and this specification does not supply them.

11.5 Victory conditions and progression

Four victory conditions were proposed: Persistent Constraint, Entropy Mastery, Flow Dominion, and Reintegration. The envisioned progression runs from vulnerable islands of organization, through conflict over corridors, toward the stabilization of structure at galactic scale. Other proposals from the same discussion include research organized around Constraint Regimes, dynamic entropy events, artificial opponents whose objectives are expressed through Lyapunov-style criteria, and multiplayer diplomacy built on a shared entropy environment.

The last of these conflicts with the isolation of players in Section 10, and the conflict needs a provisional resolution. Players cannot share a field that all of them can see and alter, because no signal crosses between them faster than light. The shared entropy environment can instead be understood as the sum of delayed contributions: at any moment, the entropy a player experiences from the rest of the galaxy is the accumulated effect of what the others did, arriving after the appropriate delay. Diplomacy in such a world is conducted entirely through

consequences that arrive late, which agrees with the principle that what players learn about each other is effect and never claim.

11.6 Minimum version

The initial scope of the first playable version was two factions, no diplomacy, and three kinds of agent: Pump, Builder, and Parasite. Everything else in this section follows that version.

12 The Prototype Document and Its Mechanics

A twenty-one page document titled *Entropy's Edge: The RSVP Wars*, version 9, is kept in the design repository as the implementation specification and mathematical supplement of the prototype. It was written before most of the mechanics of this specification were settled, and it contains a large body of material that the sections above do not yet use. This section records that material, states how it relates to the rest of the specification, and lists the points at which the two disagree. Where the prototype document and this specification differ, the difference is stated and not silently resolved, and the prototype's own equations are reproduced as recorded.

12.1 Field ontology and the variational form

The prototype presents the three fields as a fixed plenum with no metric expansion: Φ is a scalar potential or semantic capacity (negentropic density), $\mathbf{v}$ is a vector flow of energy or baryon current, and S is an entropy field measuring disorder or informational smoothness. Apparent expansion is the diffusion of entropy gradients. The setting follows Prigogine's dissipative structures, in which ordered structure persists by dissipating energy, and the coupling $-\lambda\Phi S$ in the Lagrangian is read as the maintenance cost of structure.

The Lagrangian density is

$$\mathcal{L} = \frac{\kappa_\Phi}{2} |\nabla\Phi|^2 + \frac{\kappa_S}{2} |\nabla S|^2 + \frac{\kappa_v}{2} |\nabla \times \mathbf{v}|^2 - \lambda\Phi S - V(\Phi, S, \mathbf{v}),$$

with V a potential that may contain terms such as $\frac{\beta}{2}(\nabla \cdot \mathbf{v})^2$. The game is formalized as a discrete-time dynamical system with a lattice, a state, a set of actions (build, route, research), an update operator, and a reward. Modeling assumptions are periodic boundaries, stationarity within a phase, bounded parameters, and Gaussian noise for anomalies. An action is called RSVP-consistent when it preserves monotone energy descent, and the document identifies these with the actions that advance strategic progress.

12.2 The core equations of the prototype

The time-dependent equations of the prototype are

$$\begin{aligned}\partial_t \Phi &= \kappa_\Phi \nabla^2 \Phi - \lambda S + \eta_\Phi A(x, t), \\ \partial_t S &= \kappa_S \nabla^2 S + \gamma |\nabla \Phi|^2 - \mu_S S + \eta_S A(x, t), \\ \partial_t \mathbf{v} &= \kappa_v (\nabla(\nabla \cdot \mathbf{v}) - \nabla^2 \mathbf{v}) - \nabla S - \mu_v \mathbf{v} + \eta_v A(x, t),\end{aligned}$$

where $A(x, t)$ is an anomaly source, the κ are diffusion constants, λ couples potential to entropy, γ is the entropy production rate from potential gradients, μ_S and μ_v are damping, and the η are source couplings. In this subsection the prototype's entropy-production coefficient γ keeps its recorded name, and elsewhere it is written α_∇ . These differ from the candidate equations of Section 3 in three respects. The prototype produces entropy from gradients of Φ , whereas the candidate equations produce it from player activity a . The prototype damps entropy linearly, whereas the candidate equations remove it through pumps. And the prototype has no advection of Φ and S by $\mathbf{v}$, whereas the candidate equations do. The two forms are not contradictory, since each can be obtained from a common equation by setting coefficients to zero, and the next subsection settles the architectural question in that way.

The energy functional is

$$E = \int \frac{\kappa_\Phi}{2} |\nabla \Phi|^2 + \frac{\kappa_S}{2} |\nabla S|^2 + \frac{\kappa_v}{2} |\nabla \times \mathbf{v}|^2 + \lambda \Phi S \, dV,$$

and the prototype states a theorem that E decreases monotonically without sources under periodic or Neumann boundaries, with proof by integration by parts, together with a Lyapunov candidate $\mathcal{L} = E + \alpha C$ for stability. These statements are recorded as claims of the prototype and not as results of this specification, for reasons given in the final subsection of this section.

12.3 Phases of expansion and integration

Time in the prototype alternates between two phases that mimic cosmic cycles. The Lamphron phase is expansion: it creates gradients and negentropic structure, with enhanced diffusion in Φ and reduced damping. The Lamphrodyne phase is integration: it emphasizes dissipative relaxation and global smoothing, with increased entropy diffusion and coupling. A cycle parameter, written here $\vartheta \in [0, 1]$ so as not to collide with the causal coordinate τ of Section 2, runs from pure Lamphron at 0 to pure Lamphrodyne at 1, and the coefficients are interpolated, for example $\kappa_\Phi(\vartheta) = \kappa_\Phi^0(1 + \alpha_\vartheta \vartheta(1 - \vartheta))$ and $\kappa_S(\vartheta) = \kappa_S^0(1 + \alpha_\vartheta \vartheta)$. The names of two of the five factions echo these phases, and the recorded material does not say how the factions relate to them. The phase cycle is a property of the field dynamics and is distinct from both the annual scenario boundary and the Big Bang reset of Section 7.

12.4 Numerical scheme

The prototype discretizes on a square grid with central differences and a five-point Laplacian, uses explicit Euler time stepping, and offers a hexagonal variant on the TARTAN tiling in axial coordinates with six neighbors and the Laplacian $\frac{2}{3h^2}(\sum_{\text{nbr}} u_{\text{nbr}} - 6u)$. It also describes a semi-implicit Crank–Nicolson treatment of diffusion for larger steps, and adaptive stepping that halves Δt when the maximum field change in a step exceeds ten percent.

12.5 Turn loop

A turn corresponds to one or more steps of the field equations interleaved with player actions. The four activities of the strategic layer appear in the prototype as exploration, which reveals regions of large $|\nabla S|$ so that uncertainty gradients act as the boundary of fog of war; expansion, which claims systems with elevated Φ at the price of entropy generated by $\gamma|\nabla \Phi|^2$; exploitation,

which optimizes flows for coherence by maximizing $\int \mathbf{v} \cdot \nabla \Phi dx$; and extermination, which induces entropy shocks by setting flows in conflict. A fifth step, rebalancing, applies global Lamphrodyne smoothing. Noise enters through the anomaly term.

12.6 Ethics tensor and diplomacy

Local ethical coherence of an empire is the Frobenius inner product of the Jacobians of flow and potential gradient, $E_i = \nabla \mathbf{v}_i : \nabla \Phi_i$. The alignment between two empires is the cosine similarity of their spatially averaged ethics vectors,

$$A_{AB} = \frac{\langle \vec{E} \rangle_A \cdot \langle \vec{E} \rangle_B}{|\langle \vec{E} \rangle_A| |\langle \vec{E} \rangle_B|}.$$

Trade efficiency is proportional to A_{AB}^2 , conflict probability to $1 - A_{AB}$, and an alliance is stable when A_{AB} exceeds a threshold. The ethics field follows a transport equation $\partial_t E = \mathbf{v} \cdot \nabla E - \nu(E - \bar{E}) + \zeta \nabla^2 E$, which drives it toward the global mean. The same quantity feeds the value function $V_i = \tanh(\rho E_i)$ used in the evolutionary layer below. Diplomacy of this kind presupposes that players share a field they can both read, and Section 11 gives the provisional reading under which alignment is computed from delayed contributions and not from a simultaneous shared state. The minimum version of the game has no diplomacy, so this subsection is design material for later.

12.7 Anomaly missions

Anomalies are localized oscillating sources,

$$A(x, t) = \sum_i a_i e^{-|x-x_i|^2/(2\sigma_i^2)} \cos(\omega_i t + \phi_i),$$

and missions form a directed graph with states Detect, Stabilize, Interpret, and then either Harmony or Chaos. Transition probabilities are logistic in a combination of field alignment, ethical coherence, and torsion $|\nabla \times \mathbf{v}|$, namely $P_{i \rightarrow j} = (1 + \exp(-(\alpha \langle \nabla \Phi, \nabla S \rangle + \beta E_{\text{ethic}} - \gamma E_{\text{torsion}})))^{-1}$. Completing a mission modifies parameters, so that Harmony raises κ_Φ locally. The chain is analyzed as a finite Markov process whose stationary distribution gives an expected reward.

12.8 Fleets and combat

Fleets move by $\dot{x}_f = -\alpha \nabla \Phi(x_f) + \beta \mathbf{v}(x_f) + \xi(t)$, with Brownian noise for stochastic exploration. A fleet has mass $M = \Phi(1 - \tanh S)$, flow $F = \|\mathbf{v}\|$, and exposure $E = \lambda S + \frac{1}{2} \mu_v \|\mathbf{v}\|^2$. Combat is a softmax over effective strengths, so the win probability of A against B is $e^{\eta(M'_A + F'_A - E'_A)}$ divided by the sum of that quantity over both sides, where the primes denote values after cards are applied. Cards are drawn with probabilities that depend on technology level. This movement rule is a local rule and has to be reconciled with the light-travel constraint of Section 10 in the same way that Section 11 does for corridors: the field gradient may set direction and cost and cannot set arrival earlier than light allows.

12.9 Scenario generation and victory in the prototype

Initial conditions are correlated random fields built by Fourier synthesis with a power-law spectrum $P(k) \propto k^{-\beta}$, which gives fractal structure, and artificial opponents differ by multipliers on the parameters, an aggressive opponent having γ raised by twenty percent and λ lowered by ten percent. The prototype has three victory conditions. Entropy Equilibrium is attained when the mean gradient energy $G(t) = \frac{1}{|\Omega|} \int (|\nabla \Phi|^2 + |\nabla S|^2 + |\nabla \times \mathbf{v}|^2) dx$ stays below a threshold for a hundred turns. Dominion is attained when the controlled fraction of the volume exceeds seventy percent. Rebirth requires the equilibrium condition followed by an inflaton perturbation $\Phi \rightarrow \Phi + \delta\Phi_{\text{rand}}$ that triggers a new expansion phase.

These do not coincide with the four conditions of Section 11. A natural correspondence is proposed and is not recorded: Entropy Mastery with Entropy Equilibrium, Flow Dominion with Dominion, and Reintegration with Rebirth, leaving Persistent Constraint without a prototype counterpart. Rebirth also bears on the time model, since it begins a new expansion phase within a history, which is the same kind of event as a Big Bang reset and is distinct from the annual scenario boundary.

12.10 Neural Darwinism and the technology tree

The prototype proposes selection among technologies and species traits modeled on Edelman's theory of neuronal group selection, with its three ingredients of developmental selection, experiential selection, and reentrant mapping. Developmental selection is the initial variability produced at scenario generation, experiential selection is the amplification of adaptations that succeed in play, and reentrant mapping is the feedback among the ethics tensor, the field equations, and decision making. Each technology node T_i is a group with adoption strength $w_i \in [0, 1]$ and dependencies W_{ij} forming a directed acyclic graph. At each phase cycle,

$$\dot{w}_i = \alpha \sum_j W_{ij} w_j - \beta S_i w_i + \gamma \Phi_i - \mu w_i,$$

followed by normalization of w . Synergy reinforces, local entropy penalizes complex technology, abundance of Φ encourages innovation, and maintenance costs decay. An empire's value function $V_i = \tanh(\rho E_i)$ sets its mutation rate $\eta_0(1 - V_i)$, so that coherent empires evolve slowly and stably and incoherent ones evolve fast and risk collapse. The prototype sketches that the exponentiated update $w'_i \propto w_i e^{\rho u_i}$ stays on the simplex and, for small ρ and Lipschitz u , is a contraction with a unique fixed point for frozen fields.

Species are parametrized by $\theta_f = (g_\Phi, g_S, g_v, p)$, namely gain on Φ gradients, entropy tolerance, coupling to flow, and plasticity, which drift across epochs by $\theta^{t+1} = \theta^t + \xi \nabla_\theta V_f + \mathcal{N}(0, \Sigma)$. Resources follow the Stars! trio, with generation rules on a grid node of $\text{Fe} = R_\Phi |\nabla^2 \Phi| + \epsilon_1$ for ironium, $B = R_v |\nabla \times \mathbf{v}| + \epsilon_2$ for boranium, and $\text{Ge} = R_S(1 - S) + \epsilon_3$ for germanium, so that ironium follows condensed Φ gradients, boranium follows vortical flow, and germanium follows low-entropy coherent zones. Ironium leads toward construction, boranium toward drives, and germanium toward machine intelligence and governance. The planet-level resource levels in the Stars! sense of Section 5 are thus given a field-derived generator in place of random draws.

12.11 Geothermal mass accelerators

Planetary projects convert geothermal flux into launch energy. The local gain is $\delta(\Phi_i - S_i) \cdot \text{depth}_i$, where depth is $d_0 + \sum_j (\Phi_j - S_j) / r_{ij}$ over neighbors at distance r_{ij} . Over-extraction produces an entropy shock $\Delta S_i = \zeta(\text{EnergyGain}_i)^2$, which can appear as volcanic events or field instabilities, and the efficiency of each accelerator evolves under the same selection pressure as the technology tree. Accumulated energy supports a takeoff phase in which fleets are launched or megastructures built. The mass accelerators of the strategic layer, which are the hubs of the corridor network, are the product of this process. Species with high entropy tolerance favor aggressive geothermal play and species with low tolerance prefer sustainable low yield.

12.12 Recursive futarchy

The prototype adds a prediction-market layer in which factions bet on beliefs and vote on values. Each policy P_k has an expected entropy-reduction payoff $\pi_k = \mathbb{E}[-\Delta S \mid P_k]$, resources are allocated by $r_k = e^{\beta \pi_k} / \sum_j e^{\beta \pi_j}$ with β a cognitive temperature, the outcome is fed back into Φ and $\mathbf{v}$ through $\Phi^{t+1} = \Phi^t + \sum_k r_k \delta \Phi(P_k)$ and the corresponding rule for $\mathbf{v}$, and policies with persistently negative payoff are pruned. The recursion is that the market state changes the fields, which change the next round of predictions. Resource exchange, such as ironium for germanium, is conducted through these markets. A pseudo-code outline accompanies it. The mechanism is a game-design device for studying post-scarcity economies, and the prototype lists research uses: coevolution of technology networks and ethics tensors, phase transitions between cooperative and extractive futarchies, thresholds for geothermal takeoff, and tests of whether entropy smoothing predicts alignment or collapse.

12.13 Grounding against instrumental convergence

The prototype argues that proxy-seeking behavior such as hoarding or coercive commodification is penalized by the dynamics themselves. It defines a concentration functional $C = \text{Var}[\text{stock}_i] + \xi \text{HHI}(\text{market shares}) + \zeta \|\nabla \log \text{price}\|^2$ and a potential $J = \mathcal{V}(\Phi, S) + \alpha C - \beta U_{\text{missions}}$ to be minimized, where $\mathcal{V} = \int |\nabla \Phi|^2 + |\nabla S|^2 dx$. A tariff proportional to $\partial C / \partial a$ is levied during the Lamphron phase, and actions are constrained to be contracting in the Lamphrodyne phase. The informal theorem states that any policy sequence raising C without compensating smoothing worsens J in finite time, so that strategies based on commodification are dominated. The prototype itself lists three failure modes: miscalibration that lets C grow without $\mathcal{V}$ rising, anomaly sources that override descent, and multiplayer interactions that violate the contraction assumption. This is the in-game counterpart of the motivation in the first section, that money and similar proxies are selected against by the structure of the game, and the same caveat applies, namely that the property is a design target to be demonstrated by simulation.

12.14 Horology and monetization

The prototype places monetization outside the play economy, in two products: physical watches synchronized to the simulation's ephemeris, and collectible manuals in Latin, Standard Galactic, Cursive Galactic, and ciphered editions. The dial has twenty-two hours, and game time T relates to real time t by $T(t) = \alpha t + T_0$ with a base dilation $\alpha = 22/21 \approx 1.0476$, which the

prototype describes as making the dial turn like an eleven-hour analog clock; the dial convention is examined in the final subsection. Non-uniform but deterministic flow is added through quasi-periodic modulation,

$$\frac{dT}{dt} = \alpha \left[1 + \epsilon_c \cos(\omega_c t + \phi_c) + \epsilon_q \cos(\sqrt{2} \omega_c t + \phi_q) \right] + \alpha \epsilon_{\text{tex}} \xi(x),$$

with small amplitudes, incommensurate frequencies to prevent synchronization, and a bounded spatial texture $\xi(x)$ read from the player's map location. Each limited edition of the watch implements its own (ω_c, ϕ_c) . The prototype calls the cosine terms time crystals and the spatial term texture crystals. The manuals encode the axioms and mission hints by substitution or polyalphabetic ciphers keyed to the crystal phases. The monetization boundary is expressed as the condition that the derivative of in-game entropy with respect to external monetization is zero, so no price signal enters C .

The words time crystal and texture crystal therefore have two meanings in the recorded material. In the prototype they are modulation terms of a clock. In the game of Section 5 they are resources found at stars, the first of which rewinds history and the second of which is exchanged for the first. The specification does not merge them. The proposed reading is that the watch firmware reads the same ephemeris as the game and that the modulation terms define the displayed rate of the in-world clock, while the resources are separate objects with the same names, and whether the names should be disambiguated is an open question. The twenty-two hour dial also needs reconciling with the annual boundary of Section 7, which this specification ties to the real-world calendar and not to the dilated game clock.

12.15 Gradient-driven and activity-driven entropy production

The prototype produces entropy where organization has gradients, and the activity-driven candidate of Section 3 produces it where structures operate. They describe different mechanisms. Gradients generate entropy by relaxation even in an unoccupied region, and factories and other infrastructure generate it through operation. Choosing either exclusively would remove something the combined design needs. The specification therefore settles the architectural choice as follows, as a proposal pending the designer's adoption: gradient and activity production are independently switchable mechanisms in one model, the prototype is a reproducible reference configuration of that model, and the combined configuration is the proposed game model, whose balance is an experimental question. The proposed local equation is

$$\partial_t S = D_S \nabla^2 S - \nabla \cdot (S \mathbf{v}) + \alpha_{\nabla} \|\nabla \Phi\|^2 + \gamma a + \chi \kappa S \Phi + \sigma_S - \mu_S S - \pi,$$

with α_{∇} controlling gradient-associated production, γ operational production, and χ the conversion of lost organization into the units of S . Its coefficients and units are to be specified, and it is a phenomenological form that does not carry an energy-decay theorem.

The accounting rule is the point at which the design can go wrong. If the coefficient on activity already measures all the entropy generated by maintaining, transporting, and dissipating structure, then adding the gradient and dissipation terms counts the same process twice. The specification therefore defines a narrowly, as operational throughput whose entropy production is represented nowhere else. Relaxation of interfaces belongs to the gradient term, the loss of unmaintained structure belongs to the term in χ , and maintenance costs are paid in energy

and not in entropy. Linear damping is kept only when the quantity it removes enters the accounted reservoir, as in Section 16, and pump export and damping may share a destination while keeping separate rates and costs.

The combination is intended to supply a strategic tension. Building a sharply organized region raises gradient-associated dissipation, operating it incurs a throughput cost, and neglecting it permits decay. The gradient term alone penalizes sharp interfaces and does not necessarily penalize every large or organized civilization, so maintenance and operational production have to supply the other pressures, and their effects are to be demonstrated and not inferred from the vocabulary.

Four configurations are defined over the same initial conditions and the same policies. The first, B0, reproduces the prototype document as written, including its flow operator, with α_{∇} equal to the prototype's γ , damping μ_S , the loss term $-\lambda S$ in the Φ equation, and no advection, activity, or pump. It is a historical reference and carries no claim of stability. The second, B1, is B0 with the sign of the flow operator corrected as described below. The third, A, is activity-only, with $\alpha_{\nabla} = 0$. The fourth, C, is the combined model with all three production mechanisms, the loss term $-\kappa S\Phi$, advection, and the reservoir. The word baseline means the reference model being reproduced, and it does not imply that the prototype's stability, energy, or anti-convergence arguments have been established.

The decisive evidence is whether each configuration supports the intended finite-horizon distinction between maintained organization and unsustainable expansion, with every entropy transfer accounted for in the three-part total of Section 16 and with results insensitive to reasonable numerical refinement, tested by halving h and Δt . The comparison would also show whether the gradient term contributes meaningful geography or merely makes every boundary prohibitively expensive. These outcomes are acceptance targets.

12.16 Observations on the prototype's claims

Several statements in the prototype document are stronger than their derivations support, and a specification that incorporates the document has to say so. Each correction below stays attached to its claim.

The sign of the flow operator is wrong. The prototype writes $\partial_t \mathbf{v} = \kappa_v(\nabla(\nabla \cdot \mathbf{v}) - \nabla^2 \mathbf{v}) - \nabla S - \mu_v \mathbf{v} + \eta_v A$, and the parenthesized expression is $\nabla \times (\nabla \times \mathbf{v})$. For a transverse Fourier mode, with $\mathbf{k} \cdot \hat{\mathbf{v}} = 0$, that term contributes $+\kappa_v |\mathbf{k}|^2 \hat{\mathbf{v}}$, so with constant S and no sources the mode obeys $\partial_t \hat{\mathbf{v}} = (\kappa_v |\mathbf{k}|^2 - \mu_v) \hat{\mathbf{v}}$. Short-wavelength transverse modes then grow, which is an instability of the continuum equation, and a smaller time step cannot repair it. If relaxation of curl energy is intended, the operator needs the opposite sign, which is what configuration B1 uses. If amplification is intended, it needs a small-scale regularizer and a different stability analysis. The Euler–Lagrange equation for the flow in the same document, $\kappa_v \nabla \times (\nabla \times \mathbf{v}) = -\partial_{\mathbf{v}} V$, points to the relaxing sign.

The energy derivation is not consistent. Section 2.4.2 of the prototype defines $E = -\int \mathcal{L} dV$, but the functional derivative it displays does not follow from that definition and the preceding Lagrangian, and Section 3 then gives a different energy, with $+\lambda \Phi S$ and positive gradient terms. The monotone energy theorem therefore cannot be accepted as written. Descent holds for a consistently defined gradient flow, and the equations as written are not one: even with the production and damping terms removed, the entropy equation lacks the $-\lambda \Phi$ that the gradient

flow of the Section 3 energy would give, and the flow equation contains $-\nabla S$, which is not a variational derivative of that energy. The coupling $\lambda\Phi S$ is also bilinear and unbounded below, so decay of E would not in any case imply convergence to a bounded equilibrium. The Lyapunov claim for $E + \alpha C$ rests on the same energy and inherits these gaps. The theorem should be replaced by a derivation from one explicitly chosen functional, and until then monotone decay in the full game is an acceptance target.

The time-step bounds contradict one another. The prototype gives $\Delta t < h^2/(4 \max \kappa)$ as the diffusive bound, which is correct for explicit Euler with the five-point Laplacian in two dimensions. Its amplification factor for that stencil is $1 - 4r[\sin^2(k_x h/2) + \sin^2(k_y h/2)]$ with $r = \kappa \Delta t/h^2$, and stability requires $\Delta t \leq h^2/(4\kappa)$, so the lemma's factor and its conclusion of $h^2/(8\kappa_\Phi)$ are incorrect. The worked example takes $h = 1$ and $\kappa_\Phi = 0.5$ in one dimension, quotes a bound of 0.25, and calls $\Delta t = 0.3$ unstable. In one dimension the explicit bound is $h^2/(2\kappa) = 1$, and in two dimensions it is 0.5, so both 0.2 and 0.3 are stable for the isolated diffusion equation. The example should be recomputed before it is used as a test.

The diversity term in the futarchy functional has the sign its text claims. With $L = G + \kappa \sum_k r_k \log r_k$ and $\kappa > 0$, the sum is the negative Shannon entropy of the allocation, and minimizing it favors a uniform allocation, which promotes diversity. What remains unsupported is the convergence argument, which moves from a discrete softmax update to a continuous derivative without stating the connecting dynamics and assumes contractivity without establishing it. The prose also says to prune policies with persistent negative ΔS , although negative ΔS is what the payoff $\pi_k = \mathbb{E}[-\Delta S \mid P_k]$ rewards. The sentence is inverted, since the code keeps policies whose payoff exceeds the threshold, and the prose should say that persistent positive ΔS is pruned.

The ethics construction has a type error. Section 5 writes $\nabla \mathbf{v} : \nabla \Phi$ as a Frobenius product, but the first object is a matrix and the second a vector. The indexed expression it gives produces a scalar, whereas the diplomacy formula then needs an ethics vector for the averages and the cosine similarity. The intended object has to be chosen before alignment is defined. Trade efficiency proportional to A_{AB}^2 also rewards perfect opposition and perfect agreement equally, which may conflict with the intended meaning of alignment.

The watch arithmetic needs a dial convention. At 22/21 game hours per real hour, a cycle of twenty-two game hours takes twenty-one real hours. An eleven-hour revolution follows only if a revolution represents eleven game hours, and then it takes 10.5 real hours. The document does not say which convention applies.

Three smaller points complete the list. The cited section for the victory metrics is an unresolved reference in the document. The closing statement of the philosophical note, introduced with the word proves, restates the monetization boundary as a definition, since the derivative is zero by construction when the watches and manuals do not enter the simulation. And the anti-instrumental theorem is stated informally, with its failure modes listed by the prototype itself. None of this reduces the value of the prototype as a source of mechanics. It means that the claims in these places are design intentions, and the specification labels them so in the status appendix.

13 Two Layers of Play

A related design in the project’s records combines turn-based empire management with optional real-time missions in six degrees of freedom, both governed by scalar, vector, and entropy fields, with a player-controlled technological takeoff through policies and research and with stability or collapse emerging from an entropy budget. That design is inspired by Stars! (1995) and Descent, and the recovered material does not identify it as Entropy’s Edge. This specification adopts the combination as its own structure, since the flight mechanics of the Haplopraxis lineage, the vim-style controls, the nested scopes, the blink mechanism, and the guide bot all belong to a real-time six-degree-of-freedom layer, while the field simulation and the strategic map belong to a turn-based one.

The two layers are joined by a single rule: every action in the flight layer is an operation on the field in the strategic layer, and every state of the field is the environment in which the flight layer takes place. The correspondence is given below.

Flight-layer action	Field-layer effect
Popping a scope	Injects entropy at the boundary, carried inward along the flow to the densest region
Blinking through a boundary	Crosses a constraint gradient without perturbing it
Dropping a bomb	Plants an armed region, a localized defect in Φ that releases entropy when crossed
Scrapping a minefield with skimmers	Recovers the stored organization of the defect as resources
Holding g beyond the threshold	Applies the reset operator to the field, returning it to its initial condition and announcing the source
Autoblink	A stability constraint that keeps a player’s local patch approximately invariant under the reset operator
Advancing the year	Steps the simulation forward by one turn
Rewinding or interpolating	Selects or constructs a node in the branch tree and replays the recorded interventions

Table 4: Correspondence between flight-layer actions and operations on the fields.

The autoblink row follows an earlier specification of the starspace project, which treats the g-key as a reset operator on a graph and galaxy layout and states autoblink as a stability constraint that keeps selected users’ local patches approximately invariant through resets. That formulation agrees with the account in Section 8: autoblink does not cancel the reset, it decouples a player’s accumulated organization from the loss, so that the player is reset in time but keeps their points.

The turn-based layer is also where the guide bot’s instinct is best understood. A bot that heads for batteries, escape routes, and reactors is following the gradient of useful structure, and the strategic map is the place where those gradients are drawn.

14 Visual Direction

The visual language has two recoverable components. The simulation design calls for scalar heatmaps, luminous vector filaments, crystalline regions where entropy is low, and distorted regions where it is high, so that the physical state of the world is visible and a player can read the health of a civilization from the look of its surroundings. The second component is a direction given for the ships and the general styling: retro styling with futuristic ships, and explicitly no helicopters. That correction is one of the clearest concrete aesthetic constraints in the record.

Two features of the mechanics fit this direction naturally. The amber warning that super autoblink places on armed cores belongs to the vocabulary of retro instrument displays, the amber phosphor of old terminals, and so reads as part of the same world and not as an overlay laid on top. The attract mode of the guide bot is also a retro form, the self-running demonstration of an arcade cabinet, and it gives the world a reason to look the way it does.

The most difficult feature of the display is that objects are shown in outline only, in the manner of a vector outline or a Sobel filter. The edges are green when near and orange when distant, and they grow thicker and blurrier over time, so that a player who is not moving must blink about every five seconds or move, or enable autoblink. The player rotates the world's color wheel with the arrow keys or the mouse and looks around without rotating it using h, j, k, and l, and banking changes the contrast, as set out in Section 6. Texture crystals, obtained at the second level, allow outline-only zones to be filled with texture and allow custom luminosity filters and shaders.

Two accessibility questions follow directly from these choices. A color scheme that distinguishes near from distant by green against orange may not be distinguishable to every player, and the rotatable color wheel is the provision that addresses this, since a player can move the scheme to hues they can separate. The same rotation raises a problem for the amber warning of super autoblink, which must remain distinct from both the near and distant edge colors under every rotation, and the specification therefore requires the warning to be defined relative to the rotated palette and to differ in more than hue, for example in thickness or in flashing, so that it is not lost when the wheel is turned.

15 The Guide Bot

The guide bot is the game's tutorial and its screensaver at once. When it detects that the player is asleep or away from the screen, it takes over and explores. The player learns by watching it cross boundaries, blink, trip things, and recover. This is an attract mode in the manner of arcade cabinets that run themselves to show the game, and it also keeps the world visibly alive when nobody is at the controls.

The bot has a primitive instinct to look for battery sources, escape routes, and main reactors. These are exactly the places where a core trigger tends to sit, so its instinct pulls it toward the most dangerous locations in a system. In a simple environment the heuristic works. In a nested one it causes characteristic mistakes, and a watcher sees why: the places that look most useful are the places most likely to be armed. The lesson is the same one the amber overlay makes explicit later, and it is the lesson of the breaker panel, where the main panel is both

where one most wants to go and where one most needs to take care.

The bot’s behavior also carries information. Where it heads first is a hint about where the head of the ranking lies in that environment, and where it errs is a hint about what the instinct gets wrong. A simple environment is readable enough as a teaching level if a watcher can infer the ranking from the bot’s behavior alone, in the way one infers from watching an electrician which breaker is reached for first.

In the command mode described in Section 6, the player can take control of the guide bot on the map, move it to the beginning or the end of its path, and return to firing mode with the space bar. This makes the bot a delegated pilot that the player can take over, and it fits the account of autoblink as delegated traversal one level up: the bot automates a whole path where autoblink automates a single crossing.

15.1 Open point

This specification leaves unresolved whether the bot learns from its pops, avoiding a core it has already tripped, or repeats the same mistake each time it replays an environment. The choice affects how long a given tutorial environment remains instructive.

16 Implementation Status and Roadmap

16.1 What exists

The browser prototype described in the design discussion is organized as an HTML page, a main script, and a simulation script, followed by scripts for stars, corridors, fleets, and movement, with optional Python support for persistence and batch parameter sweeps. These are records of code supplied or described in conversation, and they do not establish that a particular build is running today.

The most recent concrete implementation is a small turn-based Python model of about forty lines, with typed commands and a state consisting of the turn number, energy, entropy, buildings, and tiles. Each turn changes energy and entropy by

$$\Delta E = 1 + 2f, \quad \Delta S = 0.2f + 0.05\ell,$$

where f is the number of factories and ℓ the number of laboratories. A factory costs 5 energy and a laboratory costs 4. With two factories and one laboratory, starting from 10 energy, ten turns with no intervening purchases yield 60 energy and an entropy increase of 4.5. The scan command generates tiles with unseeded randomness and does not affect these projections.

16.2 The gap

The principal limitation of that model is that entropy accumulates without feeding back into production, construction, or survival. Under the tested building policy, growth is therefore unchecked. This is a limitation of the model’s adequacy and not an implementation bug, but it is also the largest distance between the prototype and the design, because the design makes maintaining organization consequential and the prototype records entropy without imposing any consequence for it.

The smallest change that closes the gap is to make entropy act on the three things it should govern. Production becomes $\Delta E_{\text{eff}} = (1 + 2f)g(S)$, where g is a decreasing function with $g(0) = 1$, such as $g(S) = 1/(1 + S/S_0)$ for a scale S_0 to be tuned. Construction cost is multiplied by an increasing function $c(S)$ of the same quantity. Existing structures acquire a per-turn hazard $h(S)$ of degrading or being lost, which is the discrete counterpart of the dissipation term $\kappa S\Phi$. With these three couplings the prototype has the ingredients of both a regime in which a balanced economy is sustainable and a regime in which unchecked expansion collapses. Whether it actually has them depends on the coefficients, on the cost and capacity of pumps, and on what counts as collapse, and the equations do not by themselves establish either outcome. The statements that a factory-only policy fails and that a policy including pumps survives are therefore experimental acceptance targets, to be demonstrated by running the model, and they are not claims about the model.

Entropy export also needs a destination. On a periodic lattice a pump that removes entropy from a tile would otherwise be an unexplained sink. The specification models an explicit external reservoir R of finite capacity, into which pumps move entropy at an energy cost, and which leaks or is drained at a rate that is itself a parameter. The linear damping $\mu_S S$ may share this destination with the pumps while keeping a distinct rate and cost, and it is retained only on that condition, or else with S read as a local disorder variable and not as thermodynamic entropy. Drainage is not free: the reservoir obeys $\dot{R} = \int_{\Omega} (\pi + \mu_S S) dx - d$, where $d \geq 0$ is the drainage rate, and the drained entropy goes to an accounted sink Σ with $\dot{\Sigma} = d$. With $\alpha_{\nabla}, \gamma, \chi, a, \sigma_S \geq 0$, the total $\int_{\Omega} S dx + R + \Sigma$ then has rate $\int_{\Omega} (\alpha_{\nabla} \|\nabla \Phi\|^2 + \gamma a + \chi \kappa S\Phi + \sigma_S) dx \geq 0$, which is nondecreasing and keeps the second-law character of the game honest. Without the sink, drainage alone could make the sum of lattice and reservoir entropy decrease, so monotonicity holds only for the three-part total. A pump becomes a device for relocating disorder and not for annihilating it. Because persistent activity produces entropy and the outlet is finite, indefinite pump-supported survival is not claimed: any statement that a pump policy survives carries a stated horizon, or a demonstrated sustainable regime in which drainage balances production. When the reservoir is full, export fails, which gives a late-game pressure that the early game does not have. The reservoir is also the natural place to connect to the extra-ocular regions of Section 7, if the reset is ever to lead anywhere.

16.3 Roadmap

The ordering below follows the recorded sequence of decisions, in which the economy, star nodes, and corridor detection preceded fleet movement and colonization, and it places the new requirements of this specification after the minimum game.

17 Formal Sketch

The rules above admit a compact description, offered as a sketch and not as a complete semantics. Let a domain be a finite rooted tree T of scopes, with a ranking ρ on its nodes that orders them by relevance. The live core $c(T)$ is the highest-ranked eligible innermost scope that is not refractory. A crossing of boundary b without blink sends a reaction along the unique tree path from b to $c(T)$, arriving after a delay δ that grows with the path length and with the distance of the observer. A crossing with blink sends nothing unless $b = c(T)$, which no blink

	Stage	Test for completion
1	Deterministic, seeded field core on the torus	Two runs with the same seed and interventions agree exactly
2	Entropy feedback on production, cost, and survival	Factory-only policy collapses and a pump policy survives
3	Star sampling and corridor detection	Corridors appear and vanish as the fields change
4	Fleet movement and colonization	Overextension produces loss without an opponent
5	Minimum game: two factions, Pump, Builder, and Parasite agents	A full game can be won and lost under each victory condition
6	Branch tree, rewind, and interpolation	Replaying recorded interventions reproduces any node
7	Flight layer with scopes, blink, the g key, and mines	Each row of the correspondence table acts on the field
8	Isolation and delayed consequence between players	No player’s state is readable before its signal arrives

Table 5: Proposed implementation stages and completion tests.

protects, and leaves all descendants armed. Deactivating the live core yields

$$c'(T) = \arg \max_{v \in T \setminus \{c(T)\}} \rho(v),$$

which expresses both the conservation of the armed state and the promotion of the runner-up. A player’s history is a rooted tree H whose nodes carry dates and whose edges carry costs, and a rewind selects an ancestor of the current node while an interpolation adds a path ending at a new node constrained to equal an existing one, with a provenance link and marked edges, so that H remains a tree. In the formula for $c'(T)$ the maximum is taken over the eligible scopes, and the set is nonempty only when T has at least two of them. An event is a tuple of origin node, emission coordinate, destination, and arrival condition, and it is delivered to exactly those nodes of the recipient’s tree on which the arrival condition holds. The global reset is a wave W emitted from a source location. A player retains their progress through the reset if their blink interval contains the arrival time of W at their location, or if they hold the super tier, which adds that arrival time to the set of times at which a blink fires. A player is exempt from temporal exposure only if their system lies inside an insulating structure. The three protections are therefore distinct predicates, and retention does not imply exemption. The strategic layer adds the field state $(\Phi, S, \mathbf{v})$ on a torus, evolving under the candidate equations of Section 3, with a history node recorded as a seed together with the sequence of interventions along its path, so that the state at any node is the result of replaying that sequence from the root. In this setting the reset wave W acts on the field as the operator returning it to its initial condition, and the autoblink constraint is the requirement that the restriction of the field to the player’s local patch be approximately preserved by that operator.

18 Open Questions

Several decisions remain open and are recorded here so that they are settled deliberately and not by default. Whether the relevance ranking is computed from observed traffic through a system or from the structure of the nesting alone determines whether dangerous areas drift toward the busiest routes. Whether a hit on a mine destroys something or only stops the player and reveals the field determines whether the cost of discovery arrives before or after the knowledge. Whether the amber overlay shows only the boundaries that would fire on the player, or also the runner-up in the ranking, determines how much of the hidden order a super autoblinker can read. Whether auto blink can be tied to early detection of mines is a possible extension. Whether a mine respects blink in the same way a sphere boundary does remains unstated in the design notes, and this specification adopts the working assumption that mines are cores in the sense of Section 4 and are therefore not covered by any tier, while outer scopes of a field are.

The addition of the strategic layer opens further questions. The equations of the field substrate are a candidate form, and the coefficients and the flow-reorientation rule remain to be fixed by experiment. The faction descriptions and the full set of victory conditions are unspecified beyond their names. It is undecided whether time crystals and Memory Crystals are one object. It is undecided whether the entropy hazard $h(S)$ applies uniformly or depends on the type of structure, and whether the relevance ranking of Section 4 can be computed from the field, for instance as a function of Φ and the flow through each scope, so that the live core drifts toward the busiest and most organized regions without any separate rule. The pairing of a turn-based layer with a real-time layer raises the question of how many flight-layer seconds correspond to a year, and whether the guide bot's attract mode runs in the flight layer, the strategic layer, or both.

The records leave further matters unsettled. The count of stars needed for level two is given as one hundred thousand in a later record and one hundred and fifty thousand in an earlier one. The exchange rate of one million texture crystals for one time crystal, together with the one-in-a-million chance of finding one, makes a fourteen-crystal rewind a limit case that requires on the order of 10^{13} systems, and it is unclear whether that is intended. It is unclear whether the countdown of the first reset warning includes the threshold interval. What carries between annual scenarios beyond the archive and settings is undecided. The typing tutor at the center of each star has no stated relation to the core trigger, and it is unclear whether arriving at a core through inept flight leads to the tutor or to the reaction. It is unclear whether the g of command mode is the lowercase or the uppercase key of Vim. The mapping from particular books, decks, and media to trajectories in the global space is unspecified. Whether a crystal-funded rewind still exposes other players to temporal exposure is unspecified, and so is the mechanism by which a reset reaches the extra-ocular regions beyond the particle horizon. Finally, the correct reading of a minefield as many single-scope domains or as one domain with a single live mine has been decided here in favor of the former, and the other reading remains available as a variant.

19 Relation to the Originating Problem

The social-platform defect that prompted the design can be restated in the vocabulary of this specification. A single date field serves as both the platform’s record and the poster’s claim, so there is no distinction between a position reached by ordinary advance and one reached by an unpriced jump. Backdating is permitted for legitimate reasons, such as digitizing family photographs, but nothing marks a backdated item as the poster’s assertion, and nothing keeps it from feeding the signals by which readers judge credibility. Entropy’s Edge supplies the missing structure: moves along the date axis are priced, the record of how a position was reached is kept apart from the position itself, interpolated stretches are marked, hidden structure is revealed to those who pay to see it, and what travels between parties is consequence. These are offered as design properties of a game and as an illustration of what a date field would look like if the seams were drawn on it.

A Status of Principal Claims

The table lists the principal statements of the specification and marks each as a recorded decision, a proposal made in this document, or an implementation fact. A recorded decision may still contain an internal discrepancy, which is noted. An implementation fact describes code inspected at an earlier date and does not establish that any build is running now.

Statement	Status
The game begins in 2400, resets to the Big Bang, and returns to 2400 over five minutes on a logarithmic schedule	recorded decision
Zoom controls the speed of time at level one, and time crystals add reverse travel and a vector multiplier	recorded decision; their reconciliation is a proposal
Level two follows a large number of star visits, and texture and time crystals appear there	recorded decision; the count differs between records
A player’s history is a branching tree in which only chosen branches are computed, with cheap forward replay and costly divergence	recorded decision
A revision of a fixed past is made by interpolation through a projected inquiry	recorded decision
An interpolation targets an exact state and may be rejected as infeasible	proposal
The history stays a tree, with provenance links for reconstructed nodes	proposal
Events carry an origin branch, an emission coordinate, a destination, and an arrival condition, and replies go to the branches containing the emission node	proposal
Passing a boundary pops a scope and the reaction propagates inward; blinking enters without popping	recorded decision

Statement	Status
Outer scopes are protected by blink tiers and the core is not	recorded decision
Deactivating the core is instant and the replacement is chosen by relevance	recorded decision
Eligibility, queue cycling, refusal in single-scope domains, and propagation to the highest-ranked child	proposal
A tap of g gives a boost or a flare, and a long hold with a countdown triggers the reset	recorded decision
The threshold falls toward half a second and the countdown falls toward one second with repeated use	recorded decision
An alarm builds between half a second and the threshold	proposal
The reset is distinguished into loss, temporal exposure, and insulation, and the blink tiers protect against loss only	proposal, based on recorded statements that use the word inoculation for different effects
Autoblink is a and b together, and super autoblink is s, then a, then b	recorded decision; the notes drift between chord and sequence
Chord recognition suppresses the component keys' translation	proposal
The reset wave travels at light speed, and super autoblink responds to its arrival	proposal
Super autoblink shows amber warnings on armed cores	recorded decision
The amber overlay shows only the locally observable configuration	proposal
Minefields contain and insulate, and swarms and spheres protect against the reset	recorded decision
Mines are single-scope domains	proposal
Skimmers later scrap a minefield for resources	recorded decision
Flares persist, and a flare bookmarks a star	recorded decision
A flare's persistence is bound to its branch	proposal
Stars are named from a curated Wikipedia watchlist, with systems from headings and tunnels from content	recorded decision
Objects are drawn in outline, green near and orange far, thickening and blurring over time	recorded decision
The center of each star contains a typing tutor	recorded decision; relation to the core is open
The calendar resets to 2400 each January 1 through runaway heat death	recorded in an earlier design

Statement	Status
January 1 as an annual scenario boundary distinct from the Big Bang reset, with archived histories and a scenario identifier on every event	proposal
Carry-over of visited-star knowledge or earned resources between scenarios	open
Fields Φ , $\mathbf{v}$, S on a 256×256 torus with seeding, pumping, shocks, and redirection	recorded decision
The field equations in candidate form	proposal
Entropy feedback through g , c , and h and a finite reservoir for pumps	proposal; outcomes are acceptance targets
Prototype field equations with $\gamma \nabla\Phi ^2$ production and linear damping, explicit Euler on a 256×256 grid	implementation fact, from the prototype document
Lamphron and Lamphrodyne phase cycle with parameter ϑ	recorded in the prototype document
Ethics tensor, anomaly missions, fleet combat, and Fourier scenario generation	recorded in the prototype document; the minimum version omits them
Three prototype victory conditions and their correspondence to the four proposed ones	correspondence is a proposal
Neural Darwinism tech tree, species parameters, field-derived resources, geothermal accelerators, recursive futarchy	recorded in the prototype document as future directions
Twenty-two hour dial with time-crystal and texture-crystal modulation for watches	recorded; the double meaning of the crystal names is open
Gradient, activity, and loss-conversion entropy production as separately switchable terms, with configurations B0, B1, A, C	proposal; balance is experimental
Prototype flow operator $\nabla(\nabla \cdot \mathbf{v}) - \nabla^2 \mathbf{v}$ with transverse growth	defect in the prototype; corrected in B1
Monotone energy decay, Lyapunov stability, and futarchy convergence for the full system	claimed in the prototype; not established here
Reservoir drainage to an accounted sink Σ , with survival stated over a horizon	proposal

Statement	Status
Causal coordinate τ distinct from displayed year Y	proposal; hidden correspondence across players open
Single routing rule to the live core, which no blink protects	proposal
Incoming events recorded in the replay record	proposal
Five factions and four victory conditions, names only	recorded decision; descriptions absent
The Python model with $\Delta E = 1 + 2f$ and $\Delta S = 0.2f + 0.05\ell$	implementation fact, inspected in October 2026