Liva.Fuxi: Worlds Change, Consequences Continue

Runtime Semantics for Growing Worlds
and Unfinished Commitments

Huaiyang Liu
DeepBean Limited
wyonliu@gmail.com

21 November 2026

Abstract

A shared world must keep unfolding as people and agents change what it contains and how it works. Earlier actions may still be producing effects, objects and activities may be added, and accepted commitments may remain unfinished. We present Liva.Fuxi, a runtime architecture and reference implementation for continuing such worlds through change. Its interfaces distinguish three bindings: emitted effects retain their original content and provenance; receiving objects respond under their current mechanisms; and accepted commitments retain their subjects, agreed terms, and monitor histories. Changes to goals or deadlines require explicitly authorized, recorded revisions. Continuation does not guarantee that every commitment remains achievable. The architecture connects these semantics to object interactions, causal cascades, A2A collaboration, execution, and recovery, while separating prediction from authoritative outcomes. A complete local activity combines in-flight heat, material and activity edits, continued growth during client absence, installation and use of a predefined instrument, and actual fulfillment receipts. We realize this activity through both a resumable declaration compiler and an adapter to an existing handwritten world engine. Across eight requests per method, each fulfills all sixteen original obligations and completes two new activities; six unsafe, stale, or deliberately mismapped requests remain uncompleted expansions. The equally informed implementations have matched activity outcomes. Separate fixed-rule partition-recovery and observation-adapter studies support additional parts of the architecture, without establishing an integrated distributed or physical deployment. The contribution is explicit runtime interfaces and a reproducible system account of world growth with continuing consequences and commitments; independent productivity, scalability, learning, and participant-value gains remain to be established.

Supplementary methods and evidence accompany this manuscript.

1 A shared world remains consequential

Rain can make a fire difficult to light. A flame in a dry environment may instead ignite vegetation, change an animal's route and alter what participants can do next. Such consequences arise from the composition of object mechanisms and world conditions. A randomly chosen story does not supply that structure. For a shared world, the consequences must also persist when a person leaves, an agent takes over, a worker fails or someone changes a mechanism.

Consider participants preparing a garden performance. A lantern has already emitted heat when they decide to install a material resonator and a helper. The new receiving material can turn that pending heat into sound without ignition. The helper then uses the instrument. Meanwhile, an earlier commitment to bring a visitor to the stage and keep the habitat safe remains in force. Updating everything to one new version would be ambiguous: it could rewrite the old emission, prohibit the intended new reaction, or silently change what participants had promised.

Fuxi treats these as different bindings within one continuing activity. The emission is a retained fact; its reception is a new transition; the promise is an accepted subject-bound specification. A result is delivered only after actual execution or appropriate external evidence. This is the central system object of the paper. Figure 1 shows it in the fixed native-host episode, including environmental growth while clients are absent.

The complete research object is broader than the fixture: worlds whose objects interact and grow, subjects who collaborate and change their intentions, institutions that can be questioned, and new mechanisms that become usable. The runtime should support these activities inside a game, work process or embodied task, without prescribing a standalone application. It must not reduce a living world to a log or treat a fulfillment score as a measure of human value.

The contribution is an explicit continuity interface and an executable integration of that interface with a local world, plus a reproducible comparison against equally informed implementations. The emitted-fact, receiving-rule and accepted-goal bindings are visible in the same HTTP episode, including an actual pre-existing handwritten host. The host integration also makes its trusted semantic mapping visible: sharing a source declaration cannot by itself ensure that a host implements the intended behavior. We do not claim that mature programmable worlds, online synthesis or actor/workflow systems are unable to implement this combination. The evidence establishes capability and boundary; an independent benefit over those systems remains an empirical requirement.

Pangu produces executable worlds and mechanisms. Fuxi owns their actual world effects and activation. Nuwa studies subjects' understanding, action, relationships and learning. Taihe studies humanistic intelligence, purposes and shared life. These are research interfaces, not proof of a deployed four-system integration. Flow, Seek and Land are a product layer, not a one-to-one renaming of the technical systems.

2 The continuity object and world dynamics

2.1 Three bindings in a changing world

Let an emitted effect retain identity, original payload interpretation and source provenance. If the effect is received after a rule change, its receiving transition uses the declared active receiving mechanism. Let an accepted commitment retain its subject, goal, deadline and prior monitor state. The distinction can be written as

$$e = (id, p, \Gamma_{emit}, t_{emit}, t_{due}), \qquad s^+ = T_{\Gamma_{receive}}(s, e), \qquad m_c^+ = \mathrm{Monitor}_{g_c,d_c}(m_c, s^+, t). \tag{1}$$

The envelope does not require the receiver to reproduce an obsolete material response. Conversely, updating the receiver does not authorize changing the emitted dose or the accepted goal. If a payload needs translation across schemas, that translation must explicitly preserve its declared meaning. Multiple causes, resource interactions and external observations require the corresponding richer envelope and assumptions below.

These bindings make different questions inspectable: what happened before the edit, which mechanism acted now, and which promise is being assessed. They do not alone guarantee safety or fulfillment. A new response may make an old goal infeasible; admission must then refuse, preserve uncertainty, or obtain an authorized revision. A legitimate social breach remains part of the world history. The runtime cannot silently turn it into success by resetting a monitor. Participants may explicitly authorize a revised goal or deadline; the architecture records the authorization, effective boundary and link to the earlier terms while retaining the previous monitor history. An authorized revision is distinct from unilateral replacement or retroactive erasure. The fixed activity retains its original goals and does not experimentally validate a renegotiation protocol.

The composition supports an explanation with three independently inspectable answers: which earlier action supplied this effect, why the receiving object responded this way now, and which accepted

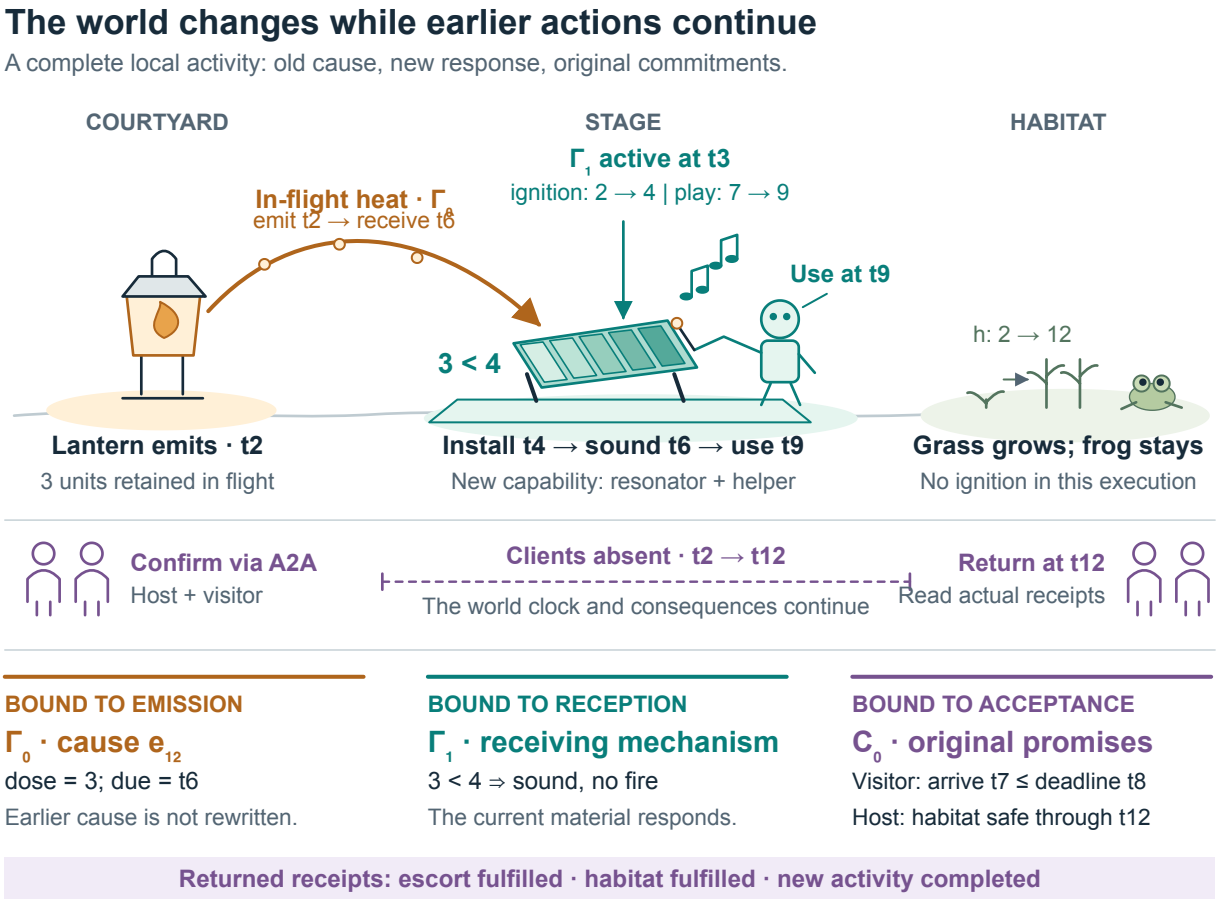


Figure 1. One activity, three continuity bindings. Spatial icons illustrate the authored fixture; they are not a rendered game or measured geometry. In the fixed native-host calm trace, the old source Γ_0 emits heat at tick 2. The modified source Γ_1 activates at tick 3; the predefined resonator and helper install at tick 4. The same old heat is received at tick 6 under the changed material rule, producing sound without ignition; the helper uses the instrument at tick 9. Grass continues growing, the visitor arrives at tick 7 and the habitat promise closes safely at tick 12. Clients confirm through the local A2A subset, make no requests during their absence and return for actual receipts. The drawing combines event times explicitly rather than depicting one instantaneous snapshot. Γ_0 and Γ_1 denote rule-source identities, not a claim that the entire activity shares one version. Quantities are authored model units, not validated thermodynamics.

terms determine the outcome for this participant. In the activity below, the emission record accounts for the retained heat dose, the active material rule accounts for sound without ignition, and the original monitors account for escort and habitat fulfillment. This is an explanation of execution under declared mechanisms, not identification of unknown physical causes or a claim that the recorded parents form a minimal explanation.

The same separation defines a reuse boundary. A new receiving mechanism can be checked and installed while the pending-effect envelope and accepted-goal descriptions continue to serve execution and outcome delivery. A host may retain its native control flow if explicit state and event mappings establish the relevant correspondences. The compiler and adapted host below reuse the same declarations, goal checks and receipt requirements over the supported edit domain. This is demonstrated interface reuse; it neither removes the initial handwritten mappings nor establishes lower integration cost than an equally capable implementation.

2.2 State, authority and situated information

Let a world specification be

$$\mathcal{W}_v = (\mathcal{S}_v, \mathcal{A}_v, \mathcal{U}_v, T_v, O_v, \mathcal{I}_v, \mathcal{C}_v), \tag{2}$$

where v identifies a version; S is the state space; A the participants; U the typed actions; T the execution transition; O participant-specific observations; I the invariants; and C the extension interface. A recorded state contains entities and geometry, resources and custody, interaction dependencies, commitments and institutions, event provenance, clocks, and versions.

A digital ownership fact may be defined by the runtime protocol. An assertion about physical reality remains evidence-dependent. A robot action request, a renderer showing completion, and a physical completion observation are distinct records. Similarly, an agent's belief is not a world fact. Let

$$b_t^i = P(X_t^{hidden}, Z_t^{-i} \mid h_t^i), \tag{3}$$

where the history contains only information legally available to participant i at the time. Delayed observations retain both occurrence and reception times. A private conversation cannot be transformed into a public world fact merely by summarizing it.

We distinguish typed intentions, predictions, observations, committed events, and fiction. Social obligations also differ from execution permissions: a character may violate a promise and face consequences, whereas an unauthorized resource transfer is excluded by the runtime. Encoding every social norm as an impossible-to-violate constraint would remove meaningful conflict and repair.

2.3 Compositional causal dynamics

Let objects carry typed attributes x_i, and let a versioned relation hypergraph describe spatial contacts, material exchange, perceptual access, obligations, and resource dependencies. A candidate local transition has the form

$$x_{i,t+\Delta} = F_{\tau_i,v}\left(x_{i,t}, \bigoplus_{e \ni i} \Psi_{\rho_e,v}(x_{e,t}, a_{e,t}), \nu_{i,t}; \Delta\right). \tag{4}$$

Here τ and ρ select object and relation mechanisms; the aggregation operator is chosen for the interaction type; ν denotes exogenous inputs or noise. Hyperedges allow interactions that require several participants or conditions simultaneously. Summation is suitable for some fluxes but not automatically for custody transfers, contact constraints, or institutional decisions. Conservation, exclusivity, units, and integration constraints must be checked at the level of the composed mechanism.

A material may absorb water, burn when sufficiently dry and heated, and transmit heat to nearby material. An organism may perceive heat, choose an available route, and accumulate exposure. These

local mechanisms can generate an episode that was not specified as an event sequence. The rules remain authored or learned: compositional outcomes do not imply that their laws were discovered by the system. Adding another combination of existing objects differs from creating a genuinely new executable relation or material behavior.

The proposal separates three mechanisms:

- Known game or physical rules are executed by exact symbolic code or an appropriate numerical solver. Approximate physical integration requires error and stability controls, even when the equations are known.
- Unknown material properties, partner policies, or off-screen dynamics may be inferred by a learned model with uncertainty. Counterfactual planning cannot read simulator-only state.
- New object or interaction semantics enter through Pangu's executable proposal and Fuxi's activation process. Old observations retain the versions under which they arose; new rules do not retroactively explain old events as if they had always applied.

Continuous flows and discrete events require explicit coupling. A hybrid model may evolve $\dot{x} = f_{m,v}(x, u)$ within mode m and apply reset maps when guards are crossed. Existence and uniqueness of a flow depend on conditions such as local Lipschitz continuity and a specified collision/contact treatment; event priorities and prevention of infinitely many jumps in finite time need separate assumptions. Merely connecting individually plausible modules does not establish a well-defined global simulator. The executable witness below uses declared discrete phases and makes no physical convergence claim.

2.4 Explanation, intervention and surprise

When transition functions are explicitly assigned modular intervention semantics and exogenous variables, unrolling them in time specifies a structural causal model with within-tick dependencies. An arbitrary predictor or execution trace does not acquire those semantics merely by being drawn as a graph. Feedback across time is allowed; instantaneous algebraic cycles require a solver or a declared ordering. A causal explanation should identify operative mechanisms and relevant conditions, and support a reproducible intervention. A generated narration saying that rain caused failure is not itself such evidence.

In a known synthetic model, replacing the rain input or a material equation and rerunning the same initial state and exogenous inputs evaluates a controlled intervention. Removing a thermal edge changes the mechanism; blocking an escape route changes an available action. These interventions answer different questions. In a real world or learned model, causal identification additionally requires valid assumptions, data, and uncertainty assessment. Logged correlations and plausible explanations do not identify the counterfactual effect by themselves.

"Surprising" is relative to a participant's prior model and goals. Predictive surprise can be measured as $-\log q_i(e \mid h_i, a_i, v)$, when q is calibrated and the event has well-defined probability. High surprise alone can reward incoherence. We therefore evaluate causal reproducibility, explanation accuracy under interventions, subsequent discoverability or use, and participant judgment separately. Humorous, meaningful, or valuable consequences cannot be proved by event entropy.

Locality can justify affected-region simulation only when the actual dependency graph supports it. A clock, broadcast, resource ledger or newly introduced edge can make a dependency global. The elementary finite-influence condition and proof are in Supplement A.1; they are background conditions, not a causal-discovery contribution.

2.5 Prediction is distinct from execution

We distinguish four objects. A predictor estimates an outcome given a recorded history and action. A causal model answers a declared intervention under structural assumptions. A mechanism revision

changes the executable world's specification. An execution receipt records what actually ran in a named branch and version. Their outputs may coincide numerically without having the same meaning.

For a structural model M with exogenous state U and a declared intervention family, let H denote information available before the queried action. The relevant queries include

$$\begin{aligned} &P_\theta(Y \mid H, A = a), \\ &P_M(Y \mid do(A = a), H), \\ &P_M(Y_a \mid H, A = a_0, Y = y_0). \end{aligned} \tag{5}$$

The last query conditions on a factual episode before asking what would have occurred for that episode under another action. Its abduction step updates the exogenous-state distribution using factual evidence. Conditioning, intervention, and cross-world counterfactual reasoning have distinct identification requirements [26, 27]. Their distinction is independent of the author's AI-stage labels, which must not be equated with the established causal hierarchy.

An observation-trained conditional model may support intervention prediction when justified by, for example, sufficient state, an appropriate action-assignment mechanism, consistency, and coverage of relevant actions and states. These conditions cannot be inferred from low prediction loss alone. Other participants adapting to a new policy, unmeasured confounders, or changes in selection can invalidate the inference. Training and evaluation must record the intervention target, assignment procedure, retained mechanisms, and source of causal assumptions.

In a fully specified synthetic game, replacing an equation and executing the modified program defines behavior in that game. It does not empirically identify the laws of an external world. A content revision $M_v \to M_{v'}$ also changes the domain against which predictions and contracts are interpreted. Evidence from v cannot be pooled into v' as if the same conditional mechanisms necessarily remained in force.

A shared random seed can provide reproducible paired runs. It supports a unit-level counterfactual only under an explicit justified coupling of exogenous variables across worlds. Matching separate observational or interventional marginals does not identify that coupling. Supplement C.4 supplies exact counterexamples. The qualitative causal-world witness therefore evaluates consequences of authored intervention semantics; it does not demonstrate causal discovery or identification from observations.

An `Outcome` must name whether an engine execution occurred in a live world, sandbox, or replay. A model-only rollout uses a separate `ModelEstimate` type and cannot carry a live completion receipt. A simulated counterfactual remains a model result even if its numerical values resemble a later observation. Event provenance identifies which computation and inputs produced a record; it is not automatically a minimal causal explanation, responsibility assignment, or proof about physical reality.

2.6 Digital authority and external evidence

A digital world owner can define an authorized state transition. An adapter to an external device cannot define physical reality by committing a database row. We represent an executed command c, a predicted consequence, a received observation and a verified task outcome separately. An outcome record binds environment, object, command, task, subject, role and issuing version. It retains occurrence time and reception time; a current world version need not equal the version under which a still-valid command was issued.

For a close-by-deadline commitment with deadline d, an acknowledged command without a result remains pending. Under authoritative closed-world digital execution, no completion by the closed deadline boundary establishes a breach. Under incomplete physical observation it establishes lack of confirmation. A trusted completion event with occurrence time τ and reception time r satisfies the temporal requirement if $\tau \leq$d, even when r>d; notification was still late. A completion event dated $\tau >$d is late for that command. In contrast, a positive state sample taken after d does not reveal when the condition first became true and cannot establish either timely or late completion by itself.

These distinctions can be expressed using an evidence set E and a set K(E) of physical histories consistent with the adapter's assumptions. For a deadline proposition ϕ, the evidence status is true if all histories in K(E) satisfy ϕ, false if all violate it, and unknown otherwise. This is an application of established partial-information reasoning. It depends on honest evidence, clock alignment and a nonempty consistent history set. Contradictory reports must be investigated, not used to prove arbitrary outcomes. The current adapter implements only the declared completion-event and state-sample cases, not a general possible-worlds reasoner.

An endpoint sample gives a direct counterexample to naive cross-domain reuse: a gate observed closed at time 6 could have closed at time 4 or time 6. With deadline 5, those histories disagree. An interface that has only that sample cannot distinguish them. A simulated plant with known fixtures does not establish real sensor accuracy, causal attribution or sustained closure. Physical rollback similarly requires a new physical action; changing a historical record cannot undo motion or heat.

3 Running the world

The runtime must execute a world, not merely certify proposed changes. We use a riverside settlement as the running design case: irrigation changes soil moisture and vegetation; vegetation changes fuel and habitat; fire changes traversal and available materials; participants negotiate water and repair work; subsequent choices change the settlement. This is an architectural scenario, not a claim that the current implementation simulates all of these domains. A plant is a persistent stateful object with local dynamics, not necessarily a language-model agent. A river can be a field discretization with boundary fluxes; a mountain can evolve slowly while constraining faster movement. Social commitments are represented explicitly without reducing people to physical state variables.

The distributed design and the live-revision fixture have different implementation scopes. The former is supported by a local multi-process recovery study with fixed rules; the latter runs one authoritative world thread with real loopback HTTP. The physical adapter uses report fixtures. No result combines all three into a deployed system. This separation is kept explicit while the following architecture describes the full execution responsibility.

Figure 2 relates these responsibilities while preserving their separate implementation scopes. The complete-activity execution paths are detailed in Figure 3.

3.1 Authority, ownership and clocks

The proposed entity record contains a stable identity, typed components, relation endpoints, a semantic version, an authority epoch, simulation time, and the last committed sequence. Spatial partitions group nearby work; nonspatial partitions own inventories, institutions, and long-running tasks. A directory maps component authority to a worker and epoch. One worker may update a component only under its current fenced epoch. Read replicas and interest subscriptions do not acquire write authority. A participant moving across space can retain a stable identity while component ownership changes.

An authority transfer freezes a bounded affected region, durably records the old owner's final sequence and unresolved deliveries, transfers state and pending work, and activates the new epoch through a single authoritative directory decision. Storage rejects writes from older epochs. A network timeout cannot itself establish that the previous owner is dead. During uncertainty, conflicting writes wait; disconnected regions may continue independent work. A consensus-backed directory and storage fencing are deployment requirements, not features implemented by the local diagnostic. The diagnostic uses a single trusted coordinator, separate worker processes and per-partition SQLite databases; it restarts a failed owner rather than implementing distributed failover.

Four time notions remain distinct: physical or simulation occurrence time; causal order; service time spent waiting for computation; and the time a participant learns an outcome. Physical integration can use substeps, ecological state can use larger steps, and deliberation can finish asynchronously. A model call cannot freeze basic motion. Conversely, fast animation cannot authorize an action before

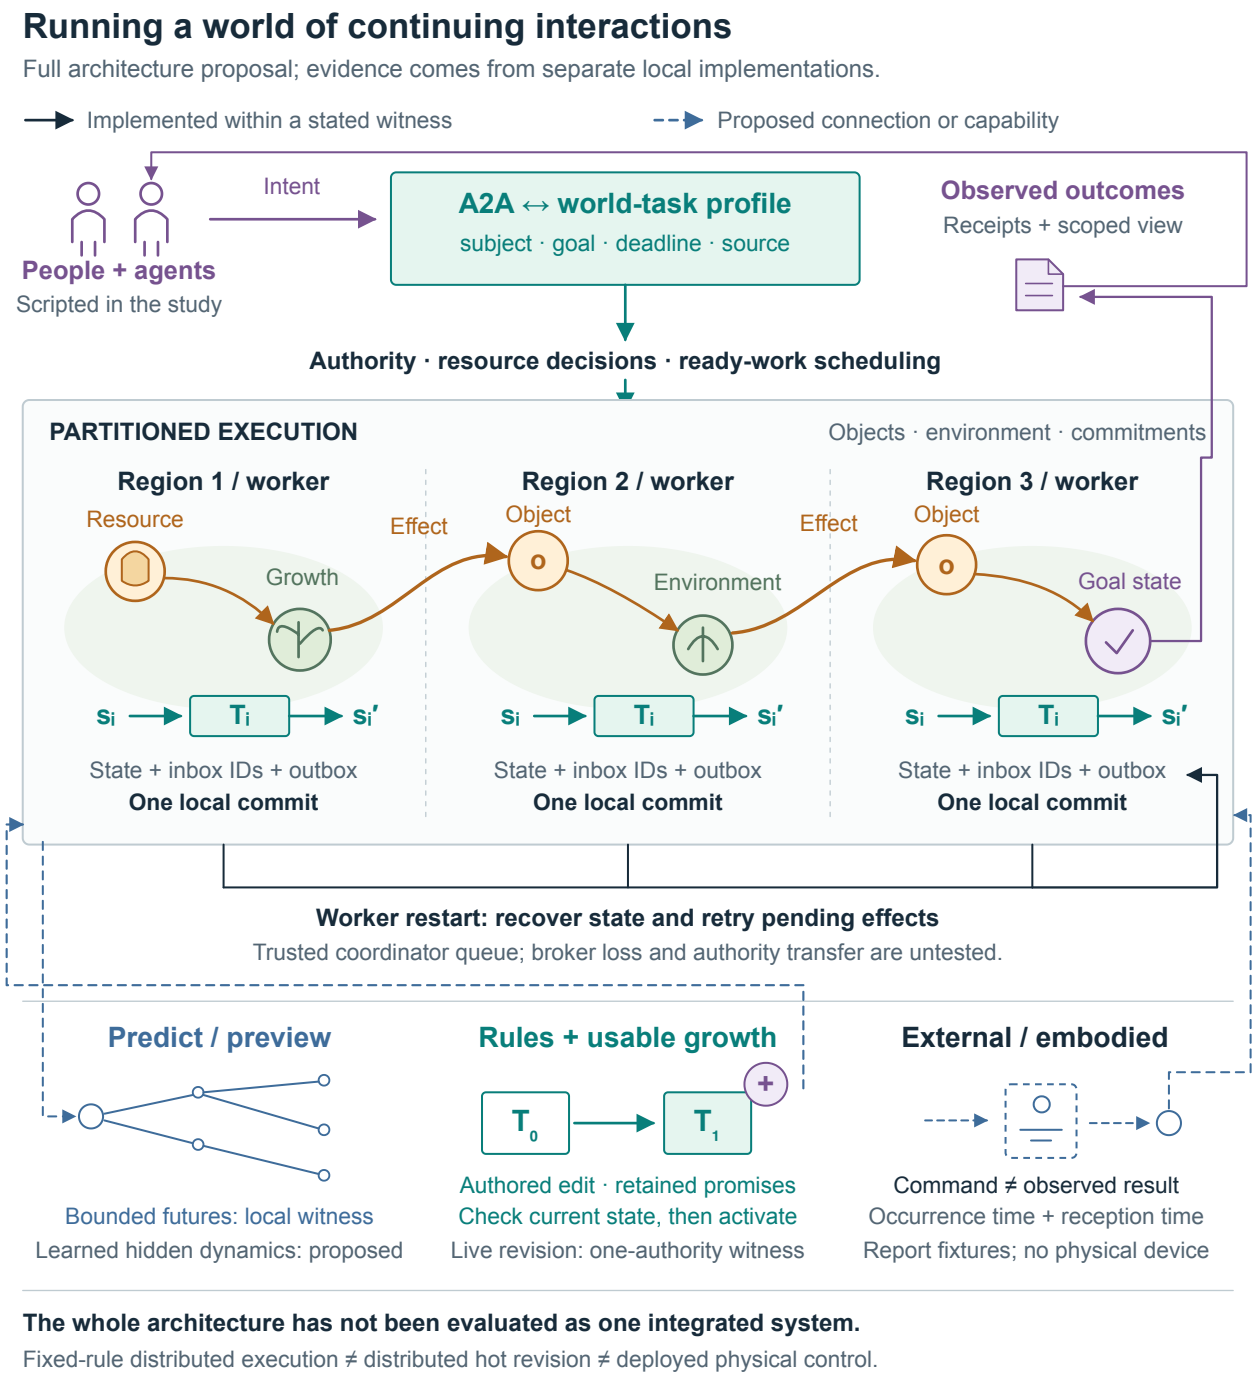


Figure 2. Full runtime architecture with separate evidence scopes. People and agents express intentions through a proposed world-task profile over the local A2A subset. Authority, resource decisions and ready-work scheduling constrain execution. Each regional worker applies local transitions, commits state and pending effects, and delivers causal effects to other regions; observed goal outcomes return as receipts. The fixed-rule distributed witness exercises three workers and worker-crash recovery under a trusted coordinator. Bounded preview and live rule revision are implemented in a separate single-authority world. The observed-outcome adapter uses simulated reports, not a physical device. Solid paths denote mechanisms within their stated witnesses; dashed paths denote proposed connections or capabilities. In particular, distributed live revision, learned hidden dynamics and deployed physical control have not been integrated or demonstrated. The entity drawing is a semantic schematic rather than measured spatial geometry.

its shared resource is acquired. Interest management limits replicated views and rendering, not the existence of offscreen consequences.

For tightly coupled continuous dynamics, adjacent solver regions exchange boundary conditions or fluxes at a declared integration boundary. Conservation and stability require a suitable solver and error analysis. For discrete-event regions, a conservative safe-time frontier can bound the earliest unseen incoming event using channel progress and positive lookahead. An event executes only when its required dependencies and safe-time condition hold. Zero-lookahead cycles require a co-scheduled region, an explicit synchronization barrier, or a justified optimistic rollback scheme; a Lamport timestamp alone does not make a numerical simulation correct [35, 36]. These are established distributed-simulation techniques. The executable witness uses neither a continuous solver nor a decentralized safe-time algorithm.

3.2 Committed effects and receiving mechanisms

Represent a discrete effect envelope by

$$e = (id, root, parents, src, dst, epoch, v, t_{eff}, op, payload). \tag{6}$$

Here parents are causal receipt references, root groups an admitted activity, and the semantic version fixes the emitted operation and payload interpretation. The receiver separately records the active receiving mechanism, so this emission version does not force obsolete receiving behavior. A stable effect ID denotes one intended application, not one transport attempt. Reusing an ID with a different canonical payload is an error. A merged interaction may have several parents; merely assigning the latest sender as the cause would erase material dependencies. Causal lineage records execution dependence, not a proof of minimal causal explanation or causal identification in an unknown physical system.

Within an authority scope, handling an event proceeds as follows:

```
check identity, authority epoch, capability, rule version and dependencies
if this effect was committed, return its recorded receipt
begin local transaction
  record the incoming effect and its content digest
  apply the authorized transition and record observed changes
  append derived effects to a durable outbox
  append the local execution receipt
commit local transaction
retry outbox deliveries until the declared durable handoff is acknowledged
expose only authorized receipt and state projections
```

A receiver acknowledges its own durable acceptance or commit, according to the declared channel contract. It suppresses repeated applications while still answering retries. A source cannot delete the sole durable copy merely because a socket write succeeded. A production channel needs replicated storage or another explicitly fault-tolerant durable handoff. The present experiment hands off to an in-memory coordinator queue assumed not to fail; its worker outbox therefore tests worker-crash recovery, not coordinator-crash or broker-loss recovery.

Cross-partition actions are not automatically atomic. A fire reaching another partition commits there under its receiving conditions; there is no requirement to lock the entire eventual fire footprint. A transfer of scarce water is different: its owner first grants a uniquely identified reservation, then records consumption or release. A joint action needing several resources acquires ordered reservations with expiry and an authoritative commit decision, or uses a declared compensating workflow. Partial completion and expired reservations remain observable. A promise to repair a bridge may be broken; failure to conserve an inventory must not be reclassified as a socially permissible breach.

3.3 Cascades, resources and continuing growth

The causal graph can cross spatial, ecological and social relations. An output event can schedule delayed events, change local transition rates, and create new opportunities. Cycles are legal: rain promotes vegetation, vegetation changes water uptake, and moisture changes subsequent growth. Each rule declares whether repeated arrival adds a new physical contribution, merges a semilattice observation, or represents the same effect. A transport deduplicator must not discard two distinct heat contributions simply because they share a causal root. In the finite witness, one irrigation wave activates each plot once by an explicit authored rule; this is not a universal semantics for fire, fluid or ecological feedback.

Each worker has separate queues for accepted world effects, control/recovery work, and new discretionary proposals. Deficit or root-based fair scheduling prevents one admitted activity from occupying every service slot, subject to finite work and eventual service. Hard dependency order always takes precedence over fairness. Backpressure delays admission of new proposals and exposes estimated waiting time. Already accepted effects remain durable and accounted for. Coalescing is permitted only when it preserves declared transition semantics, sufficient statistics, and error bounds; a rendering update can often coalesce while a transfer or causal trigger cannot.

For admitted work A(t), completed work D(t), and nonnegative initial queue Q(0), the accounting requirement is

$$Q(t) \geq Q(0) + A(t) - D(t). \tag{7}$$

If long-run admitted work exceeds service capacity, an exact system cannot simultaneously promise bounded backlog, bounded delay, and no lost effects. Branching and feedback amplify work: under a stationary multitype branching approximation, expected finite cascade work requires a subcritical reproduction matrix. That approximation is not a guarantee for adversarial, state-dependent or growing worlds. In particular, limiting a cascade's depth changes the world unless that limit is itself a physical or semantic rule. On budget exhaustion, Fuxi reports unfinished consequences and preserves resumable work; it does not silently report a stable completed world.

Inactive regions retain their state, scheduled events and approximation contract. A coarse ecological step can replace fine updates only under a specified fidelity criterion. Returning participants see the actual committed history and uncertainty of reconstruction. Hotspot relief may move ownership or subdivide work, but tightly coupled components can impose a scaling limit. The design permits time dilation or admission refusal when necessary, with visible effects on commitments. It does not promise arbitrary scale at fixed latency.

3.4 A2A bound to shared-world outcomes

A2A specification v1.0.1 is the interoperability reference [32]. Its wire version is `1.0`, rather than the patch-bearing specification label. Agent Cards expose capabilities; messages and tasks support ongoing interaction. Sending a message is not universally idempotent. Ordered streaming does not guarantee recovery of every missed update after reconnection. Cancellation concerns task lifecycle and is not physical rollback. These distinctions motivate, rather than supply, our additional shared-world semantics.

Our proposed world-task profile binds an authenticated task to world/branch, subject, observation cut, object references, semantic version, intended outcome, authority basis, reservations, deadline, and actual outcome receipts. The profile is a proposed versioned extension, not an official A2A standard. A capability description is evidence for discovery; the relevant world authority still checks execution permission. Private context and creator traces are not carried across agents merely because a task references them.

The interaction has three linked state machines. Communication tracks message acceptance and delivery. The collaboration commitment tracks proposed, accepted, reserved, executing, fulfilled, breached, renegotiated or expired status. World execution tracks unapplied, committed, partially

committed, compensated or externally unknown effects. None can substitute for the others. For example, an agent accepting irrigation work does not establish that water exists, a reservation does not wet the field, and a transport acknowledgment cannot complete the goal.

The planned gateway discovers an eligible skill, requests a scoped proposal, records mutually accepted terms, obtains resource reservations from their owners, dispatches executable intentions, and derives task completion from authoritative goal receipts. Repeated requests bind the same intent ID and payload; changed proposals receive new revisions. A rejected offer is distinct from an accepted promise later breached. For competing clients, the resource owner linearizes reservation decisions and records the winning policy. This ordering does not prove a fair or desirable allocation; Taihe and the participating people can question that policy.

Cancellation races are resolved at an authoritative decision boundary. Before any effect commits, cancellation can release a reservation and close the activity. After a partial effect, the gateway reports the committed prefix and the requested compensation, instead of deleting history. If the goal already committed before cancellation, the task remains completed and a compensating action becomes a separate task. If the external outcome is unknown, the system reconciles evidence before making a definitive claim. A disconnected client may recover current task state and durable world receipts independently of a missed stream. The detailed proposed payload and transition rules remain in the proposed `interfaces/a2a-world-profile.md` profile in the evidence archive.

The HTTP implementation supplies a local HTTP JSON-RPC subset: Agent Card retrieval, proposal/confirmation through SendMessage, GetTask, limited ListTasks, and cancellation before confirmation. Its structured proposal declares one water unit, a goal and world version; a confirmation becomes a resource-owner request, and a completed task carries a world receipt. Message identity and context checks support retry behavior. The public observation endpoint is application-specific. The fixed integration sequence exercises disconnect/re-entry and a quiescent object extension. Preassigned laboratory subject headers are spoofable and do not establish authentication; TLS, independent SDK interoperability, all protocol options, free-language negotiation, post-commit cancellation and compensation remain unimplemented. This is executable application binding evidence, not an official A2A conformance result.

3.5 Recovery without rewriting consequences

The earlier irrigation integration separately exercises resource use, propagation across three workers during client absence, return observations and a quiescent extension. It supports recovery requirements, but does not establish distributed hot revision for the lantern episode.

Recovery has separate levels. A worker reconstructs local state and unfinished outbox work; a task recovers its reservations and committed prefix; a world reconciles inter-partition frontiers; a participant recovers an observation-limited account. Replay freezes accepted intentions, semantic versions, nondeterministic inputs and external observations. It never resamples a language-model response as if it were historical fact. Floating-point physical replay may require tolerance and convergence evidence. Internal rollback cannot withdraw an external message, reverse a physical action, or undo a completed transaction. Compensation is a new event with its own effects and possible failure.

An upgrade pauses only the affected execution boundary, verifies the contracts in Section 4, migrates state and pending effects, and records when the new rule takes effect. An old effect must either retain its old interpretation or undergo an explicit validated translation; interpreting queued work under whatever code is newest can silently change history. A learning update changes a predictor or policy artifact and its provenance, not committed world facts. This makes continued world growth compatible with auditable change without treating immutability of all rules as the goal.

3.6 Inspecting and revising from the activity

The interface has four operations: inspect the permitted current situation; propose a scoped edit; preview effects against retained obligations; and activate only against an applicable current state. The

host may be a running game, an existing business application, a message channel or an embodied task controller. A studio or canvas is an optional view. A preview is a model result, an accepted command is an execution attempt, and a resulting observation is evidence; none substitutes for the other.

The environment adapter declares the authority of state, action semantics, identity/role mapping, observation source, time basis and recovery semantics. A game adapter can own a digital object transition. An enterprise adapter must reconcile the system of record, delegated organizational role and external completion; an employee's preference cannot silently become an organizational obligation. A device adapter must translate commands into supported control actions and validate observations. The implementations cover a local game-like digital runtime, an adapted earlier handwritten transition function and simulated device reports. Enterprise systems, independent game plugins and physical control remain proposed.

Slow compilation and planning belong outside latency-critical execution. The online garden witness tests one nonblocking initial planning path and checks whether the current state still belongs to the computed winning policy before activation. The smaller embedding adapter conservatively rejects an outdated snapshot and asks its caller to re-evaluate. Those are different implemented guarantees. It does not claim that every edit succeeds without interruption or that refusal is a successful expansion. State, memory, skill, model and improvement-procedure changes retain separate interfaces and evaluation; embedding a learner in an activity does not require learning on every simulation tick.

4 Constructing a valid continuation

4.1 Retained contracts and new possibilities

An extension declares executable object or rule changes, a migration scope, the properties to retain, any authorized revisions and the activation boundary. Preserving every old action trace is an optional strong compatibility mode; a deliberate new material response need not satisfy it. The standard legacy simulation condition and its limits are recorded in Supplement A.2.

Represent a continuity contract as

$$c = (id, v, scope, \phi, \mathsf{q}, \mathcal{E}, \mathcal{O}, B, status). \tag{8}$$

Here scope identifies participants and affected resources; ϕ is the property; q specifies its quantifiers; E records environmental and scheduling assumptions; O specifies available observations; B specifies time and resource budgets. Its status distinguishes active, satisfied, breached, renegotiated, and expired. A record that a promise existed must remain interpretable even after that promise is legitimately revised. Keeping a promise's record does not imply keeping its outcome possible, nor does feasibility imply fulfillment.

Three contract classes require separate checks:

1. Historical consistency: a migration retains the declared interpretation and provenance of selected committed consequences. Corrective records can coexist with a disputed historical assertion; keeping provenance does not endorse its truth.

2. Future feasibility: a goal remains reachable under specified resources, cooperation, observations, and horizon. Existential reachability, adversarial strategy existence, and probabilistic guarantees have different semantics.

3. Information and capability: the updated observation and authorization functions retain specified restrictions and usable capabilities. A changed renderer, response code, or timing channel can matter even if state variables are unchanged.

Existential reachability, robust continuation and actual fulfillment remain separate. The earlier finite contract diagnostic asks whether some bounded action sequence reaches a goal. The live episode instead

enumerates installation times and checks both declared wind outcomes. Neither result establishes general partial-observation, unbounded-horizon or arbitrary partner-policy guarantees. Two individually reachable goals can compete for one resource, so simultaneous obligations require joint reasoning or explicit reservations.

An extension proposal therefore contains executable definitions, migration m, an explicit set of retained contracts, proposed revisions with their authorization basis, and an activation boundary. A proposal cannot silently remove an active contract from validation. Soft social promises may be breached through ordinary world behavior; the runtime records consequences instead of making every norm an immutable law. Extension-admission requirements and participant commitments are distinct types.

4.2 One executable declaration for model and runtime

The new prototype accepts finite scalar state, typed total expressions, tick-local inputs, ordered guarded assignments and explicit subject-specific obligations. Right-hand sides within a rule are evaluated in the same pre-rule state; rules execute in declared order. A rule can preserve an emitted effect while modifying its receiving material, spend a resource, grow vegetation or introduce an instrument and its helper. The compiler contains no garden-specific transition handler. The earlier bounded-model form generates a time-unrolled expression graph. The live runtime uses a resumable one-step graph that can start from the actual current state; the same candidate backend object supplies preview and, after acceptance, execution. Both retain common subexpressions, explicit inputs and full-state outputs.

Obligation declarations generate persistent monitor state: fulfilled and breached states are not overwritten by later arrival. A query roots action legality, the requested capability and every relevant obligation. Guard reads, assignment reads, version-dependent expressions and prior monitor state all become dependencies. The query graph may omit irrelevant computation; the actual world history is never deleted. An active obligation outside the modeled horizon produces unknown unless a separate continuation condition has been supplied. The prototype supplies no such infinite-horizon condition.

Semantic preservation has a standard conditional induction argument. Constants and input nodes have the source values. Induction on each typed expression preserves its value. Guarded assignment nodes select the old value or the corresponding pre-rule RHS, so induction over ordered rules preserves the state after a tick. Generated monitors obey the same specified transition relation; induction over ticks preserves the finite trace and query. This argument assumes the implementation realizes the stated translation. Differential checks against the separately implemented earlier runtime provide finite evidence, not a mechanized compiler proof or validation of the authored physical model.

Source binding matters independently of compilation: a result must name its source, goal, horizon, inputs and implementation. The earlier binding checker reconstructs the complete generated graph and explicitly trusts its compiler. That construction cannot verify its own translation. Generic slicing and the rejected specialized reduction are reported in Supplement C.8; no slicing advantage is claimed here.

4.3 Integrating an existing host

Connecting an existing world engine requires more than hashing its proposed rules. Let H be native host state, S declared state, $P : H \to S$ an explicit projection, δ_θ the adapted native transition and T_Γ the declared transition. For a fixed edit, the commuting condition

$$P(\delta_\theta(h, u)) = T_\Gamma(P(h), u) \tag{9}$$

must hold on the states and inputs on which continuity relies, with matching initial state. This is a standard refinement obligation, not a new theorem. Induction yields equality of projected tick states when the condition holds on the entire reached domain. An update additionally has to preserve pending causes and the accepted subject, goal, deadline and monitor state. State equality alone does not establish event identity or cause lineage: a separate event projection and preserved effect identifiers

are needed. Receiving-rule semantics and original task intent must therefore remain explicit at the adapter boundary.

Source identity alone cannot imply this condition. Two adapters can expose the same declared rule and hash while one executes a helper action at tick 9 and the other at tick 7. Both may end with a functioning instrument and completed old promises; only one realizes the specific requested revision. This elementary counterexample motivates a concrete host check, without claiming a new impossibility theorem or a novel verification method.

Our native-host integration reuses the earlier handwritten transition function. A trusted, explicit adapter translates two scalar controls into native code, projects nested commitment state and normalizes native event records. It does not infer those correspondences. The source compiler and the native host both receive a bounded conformance check against direct source interpretation before preview and activation. The check enumerates installation times and both wind conditions from the current snapshot, compares each resulting tick state, and reports a mismatch with its time and fields. It is finite development evidence over those paths, not a proof over arbitrary native states. The native event order remains intact; only end-of-tick state equivalence and separately audited causal/receipt properties are claimed. Public dependency edges use declared read sets and observed write history, while native events are retained separately. They are not Python read-access instrumentation or a proof of physical causal identification.

The host continues while this work runs. A successful check and preview are insufficient if the live state has left the covered continuation set. Activation tests the current state, retains the old in-flight effect and monitors, and installs the checked execution object. The original declaration and every emitted effect retain their identities. New object schemas and arbitrary host-code changes require additional adapters or migration; the current implementation rejects unsupported source shapes rather than guessing a translation.

4.4 Activation and subsequent use

The runtime accepts a scoped source revision against a named current source. It validates the supported edit domain and unchanged original commitments, constructs the candidate backend, checks model/host agreement on the declared paths, and searches for a bounded continuation. These operations run outside the authoritative world lock. At activation, it tests the live state against the covered continuation states; success installs the same checked backend and plan. Otherwise the old world continues and the requested expansion is not completed.

The general architecture additionally requires explicit state and pending-effect migration when schemas change, current authority and observation checks, and a versioned activation receipt. The implemented native adapter accepts only the two stated scalar controls. It does not infer arbitrary Python semantics or initialize unknown new attributes. An old source hash is necessary provenance but cannot establish host conformance.

An affordance becomes meaningful to the activity when it can be discovered and used under a subject's actual view. Descriptions should expose relevant prerequisites, effects, uncertainty and supported versions without leaking private conditions. The installed resonator has an actual subsequent use in the fixed episode. It is an authored construction, not evidence that a new participant autonomously discovered its use or that a learned system invented it.

5 Evidence from complete activities

5.1 The earlier handwritten host executes the revision

The fixed protocol, complete source snapshots, adapted native functions and all execution traces are at `next-research/2026-11-21-host-adaptation/fixed-01/`; the preceding shared-source episode is fixed at `next-research/2026-11-21-shared-source/fixed-01/`.

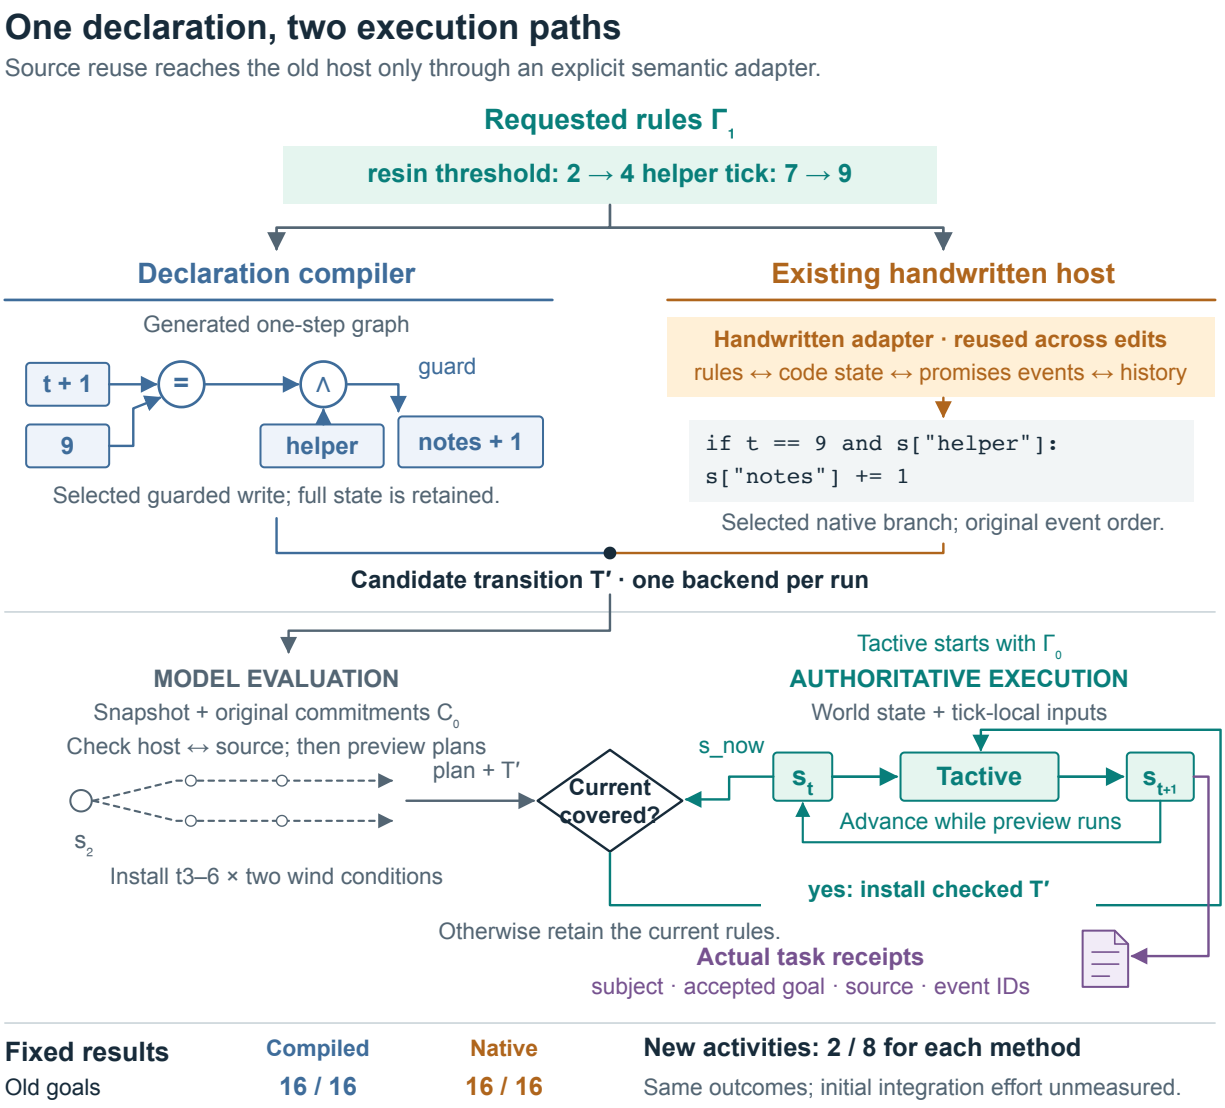


Figure 3. Source reuse with an explicit integration boundary. The two alternative backends receive the same declaration. The selected generated expression depicts a guarded helper write, not the entire transition graph; a false guard retains the previous notes value. The native path uses trusted, handwritten rule, state/promise and event mappings, generates the selected edit in the earlier host function, and preserves that host's internal event order. A candidate transition is checked against source interpretation and used for finite preview outside the world lock. The authoritative world continues with tick-local inputs. Only a plan covering the actual current state can install that checked transition. The model branch cannot issue completion receipts; actual state and events substantiate them. Eight requests per implementation yield sixteen old-goal fulfillments and two new activities each. Matching final outcomes do not imply identical timing or every state history. Setup effort is unmeasured, and these are local authored development cases rather than an external benchmark or learned controller.

The principal study executes the actual earlier handwritten garden transition through a new explicit adapter. The native host keeps its own control flow, nested promise records, interaction logic and internal event order. The adapter generates two native scalar edits and observation hooks, and supplies the state/event correspondences described in Section 4.3. It is allowed the same source, checks, reuse and information as the compiler condition. It is not required to handwrite a fresh conversion for each case. Both conditions are restricted to the stated two-control edit domain.

The complete task now changes two interacting parts of the activity: the resin ignition threshold from 2 to 4 and helper use from tick 7 to tick 9. Heat is already in flight; old escort and habitat obligations remain. A successful receipt requires actual installation, heat reception and use at the requested time. Four conditions—safe two-part edit, unsafe threshold-3 boundary, a 0.20-second injected analysis delay and an injected old-helper-time mapping error—cross two winds and two implementations, for sixteen predeclared development executions. The mapping fault is applied equally to both backends and is a checking control, not an observed human mistake or an intentionally weakened baseline.

Table 1: Original obligations and expansion requests have separate denominators.

Outcome per method, eight executions	**Declaration compiler**	**Native-host integration**
Old obligations fulfilled / required	16 / 16	16 / 16
New activities completed / requested	2 / 8	2 / 8
Unsafe edits refused	2	2
Stale edits refused	2	2
Injected mapping faults refused	2	2

All successful new activities use the instrument at tick 9. The safe compiler executions activate at tick 2 and install at tick 3; the safe native executions activate at tick 3 and install at tick 4. Consequently those paired state histories differ. The other six pairs have equal state histories, and all eight pairs have equal final activity outcomes. Late decisions occur at ticks 10–12 and are retained as observed. Refusing an expansion does not satisfy that expansion; old fulfillment is reported separately.

The fixed runner takes 5.064 seconds wall time and 0.801 seconds CPU. Summed construction, conformance-check, preview and actual-step times are respectively 0.026, 0.056, 0.019 and 0.009 seconds for the compiler, and 0.090, 0.109, 0.056 and 0.021 seconds for the instrumented native host. These overlapping internal components do not sum to wall time. Observation hooks add native overhead, and the short timings are not a scalability or production-engine ranking. Additional development runs and scorer-control costs are retained in Supplement C.11. Authoring and preparation effort remain unmeasured, not zero.

After initial integration, both implementations consume the same two-value source edit without another handwritten per-case dynamics mapping. However, the native rule anchors, nested promise projection, version fields and event correspondence were written for this host. The unchanged declaration hash did not prevent an injected tick-7 execution from violating the tick-9 request; the shared bounded checker caught it. This supports a precise adapter trust boundary and a working full episode, but neither an independent reduction in per-edit labor nor a higher complete-activity success rate. It does not measure arbitrary host transfer, external benchmark performance, human productivity, learned mechanisms, real devices or partitioned hot updates.

5.2 The preceding shared-source HTTP episode

The shared-source episode connects the previously separate local mechanisms. Two scripted subjects use the frozen HTTP/A2A subset to propose and confirm their obligations. The authoritative world emits a three-unit heat effect at tick 2, due at tick 6. A subsequent HTTP task carries a material rule edit. The same backend object supplies proposed-world preview and, on acceptance, actual

execution; the two subjects make no further requests until returning at tick 12. The pending heat, original goals and source identities persist. An installed, predefined resonator receives the old heat and is subsequently used by a helper. This is executable creation of an available affordance through installation, not intelligent invention of an instrument.

Twelve fixed executions compare resumable compilation and direct interpretation with identical rule source, finite plans, cache opportunities and information. Each method fulfills 12 of 12 old obligations and completes two of six expansions. Two infeasible and two stale edits per method do not complete the new activity. All six paired state trajectories agree. The 4.114-second wall and 0.760-second CPU totals include the full runner; there are no external model calls. Detailed delivery-audit corrections remain in Supplement C.10.

This closes a local implementation gap, but does not establish that shared semantics has a unique modeling benefit: the equally capable shared-source interpreter needs no separate per-edit dynamics mapping either. It is a single-authority world, not a combination of live rule migration with the earlier partition-crash experiment or the physical-report fixture.

5.3 What the comparison establishes

The primary capability is the complete episode: an old cause is consumed once under a new receiving rule, the revised capability is actually used at its requested time, and original subject promises retain their terms and actual receipts. The comparison establishes that both implemented paths can support this episode. In each method, the two admissible safe cases complete; the six unsafe, stale or mapping-fault requests do not complete a new activity. The overall expansion denominator remains eight, so refusal is not relabeled as success.

The native host is a strong control over the stated edit domain. It receives the source, conformance checker, finite search, full snapshot, commitments and reuse opportunities; its adapter is reused rather than rewritten per request. Both implementations therefore consume the same source-only edit after initial integration. No independent developer study measures that initial work, and this experiment does not establish lower modification cost or a uniquely available activity.

Table 2: Evidence from separate implementations; scopes are not a combined deployment.

Separate evidence	**Compared result**	**Supported scope**
Frozen concurrent HTTP/partition study	320 required tasks completed by each strong method; no latency advantage	Local concurrent effects and worker-crash recovery under its fixed-rule protocol
Shared-source revision episode	12 old obligations per method; 2/6 new activities per method	One source object connects preview and actual local execution
Native-host two-part revision	16 old obligations per method; 2/8 new activities per method	A reused handwritten host realizes the requested semantics through explicit mappings
Bounded query slicing	396→260 expression nodes for both methods	Inherited dependency reduction, no specialized advantage
Digital/observed-outcome fixture	Matched status judgments with ordinary workflow control	Authority and occurrence/reception-time distinction in local fixtures

These results are not additive evidence of a production system combining live rule migration, distributed failover and real physical control. The full legacy experiments, original failures, scoring corrections, cost accounting and raw-artifact locations remain in Supplement C and their frozen directories.

6 World growth, subject agency and learning

6.1 Growth has several forms

A continuing world can change in four distinct ways: object processes alter the environment; subjects learn or revise intentions; repeated interaction creates social practices; creators introduce new executable objects or rules. A plant counter establishes only a narrow instance of the first. Neither message throughput nor object count establishes the other three. Fuxi must preserve the consequences and evidence connecting these changes without requiring every activity to optimize a goal fixed in advance.

For example, an irrigation channel may later support an improvised sound-making activity. Pangu supplies an executable instrument proposal; Fuxi establishes when it becomes available, what water it consumes, which other objects it affects and what actually happens; Nuwa revises each subject's situated understanding and possible actions. People may reject the new activity or choose to keep gardening. Taihe studies the humanistic purposes, relationships and forms of flourishing at stake. Fuxi supports those choices through feasible action, visible commitments and continuing consequences; it does not infer human value from a completion receipt. The executable garden installs a researcher-defined resonator and helper, then uses the resulting capability. Its rules and components are authored in advance. This is installation and use of a predefined construction, not intelligent generation of a new instrument or a participant study.

The resulting research hypothesis is integration-sensitive: matching language messages while losing resource consequences, or preserving a log while making new creations unusable, both fail to sustain the intended shared activity. A capable conventional runtime can implement these connections too. The proposal earns independent value only if explicit semantic structure improves correctness, adaptation, development effort or cooperative outcomes against that matched system. The present evidence fixes these bindings through complete local traces; the comparative criterion is assessed separately in Section 9.

6.2 Consequence prediction under limited observation

The environment combines continuous physical processes, scheduled events, actions with duration, and instantaneous transactions. A partially observable semi-Markov formulation represents the next event, delay, and observation:

$$Q_\theta(de, d\tau, do' \mid h_t^i, u_t^i, v). \tag{10}$$

Other participants' private states and future behavior must be marginalized or approximated through persistent latent particles. Resampling a partner from the initial prior at every imagined step changes the model and discards information acquired from waiting.

For the known structural component, one admissible factorization is

$$Q_\theta(dS', de, d\tau \mid h, u) = \delta_{T_v(S,e,\tau,u)}(dS')K_\theta(de, d\tau \mid h, u). \tag{11}$$

Unknown hidden state requires additional integration. The learned kernel predicts uncertain events; the execution function enforces their specified structural effects. A model cannot assign itself authority or commit its own prediction as an observation.

Event-time likelihoods need consistent measures. If completion and refusal have competing hazards, and a deterministic deadline occurs at d, the deadline has probability mass equal to survival until d. It is not a continuous density evaluated at d. Administrative censoring differs from an observed refusal. Mixture likelihoods are marginalized before taking logarithms. These details matter for learning from asynchronous commitments and are inherited from the program's earlier mathematical analysis.

A proposed event model represents entities, actions, commitments, and versioned causal links as a typed event graph. Encoders consume only the learner's permitted history and corresponding semantic

version. A persistent latent belief accounts for uncertain partner state. Separate heads estimate event type, elapsed time, observable effects, and failure or unknown outcomes. Action execution remains governed by the versioned transition semantics; predicted actions do not commit facts.

A possible training objective combines event-time likelihood and proper scoring of observed outcomes:

$$\mathcal{L}(\theta) = -\sum_j \log \int p_\theta(e_j, \tau_j, o_j \mid h_j, u_j, v_j, z) b_j(dz) + \lambda \mathcal{L}_{obs}(\theta). \tag{12}$$

The first term uses densities, deadline masses, and censoring terms appropriate to each record; it must not mix measures. The second term applies only where targets are observed and justified, with masks for missing evidence. Model-independent structural constraints are reported separately from predictive accuracy. This is a predictive training objective. A causal claim additionally requires the identification conditions in Section 2.5 and intervention-based validation; raw likelihood establishes neither causality nor an improvement in policy return.

Migration-aligned examples can connect old and new encodings only on explicitly retained semantics. A new physics rule or institution is a distribution change, not an invitation to force representation equality. Proposed alignment losses must mask altered consequences and compare against unshared encoders, version tokens alone, and explicit symbolic semantic descriptors under equal data and budget. These objectives have not been trained in the reported experiments.

For transfer, a learned option needs more than a shared action name. Its input units, preconditions, observable effect contract, failure modes, and version compatibility must map to the target world. An exact transition homomorphism would support formal transfer of an abstract policy under corresponding assumptions; approximate learned mappings instead require measured prediction and control error. Existing mathematical and representation-learning ideas underlie these requirements. We make no claim that a universal shared latent world model has been demonstrated.

6.3 Creation-process evidence and actual outcomes

Optional Studio and Canvas views in Pangu support a connected creation process, including alternative sketches, edits, external tool steps, abandoned branches, and feedback. Its improvement loop must distinguish retained memory, improved tools, learned model changes, and changes to the improvement procedure itself. Fuxi consumes executable candidates from that process and returns actual running-world evidence; it does not infer a creator's preference from whether a branch was executed.

Pangu owns a scoped `ProcessEpisode`, separate from `WorldPackage`: a local decision, its brief and artifact versions, information available before that decision, the proposal, and separately timed later outcomes. It grants no access to the entire Canvas. A later result may serve as an eligible learning target but must not be inserted into the earlier decision context. Availability is checked against the actor's knowledge cut, not only an event's occurrence time. Fuxi joins episode references to engine outcomes; it does not rewrite the episode history.

The proposed link is a stable chain of `creation_trace_id`, `candidate_revision_id`, `package_digest`, `activation_id`, `intent_id`, and `outcome_id`. A creator can reject a working design for aesthetic reasons, retain a failed experiment because it taught something, or leave a draft unfinished for lack of time. `abandoned` is a lifecycle state, not an automatic negative training label. A retry, a modified candidate, and a new execution retain separate identities.

Fuxi returns environment and rule versions, branch mode, actual changes, observation scope, failure or unknown status, timing, and costs. Pangu keeps process lineage and selection rationale. Feedback records separately state their author, target revision, dimension, context, confidence, and training permission. Preference, executable failure, visual quality, and physical validity are distinct dimensions. Unexecuted candidates have no invented execution outcome. Sandbox success does not imply production activation, and process inspection does not expose private creator content to other agents.

This interface allows positive and negative experience to inform future creation without discarding exploratory history. It remains a proposed contract: cross-line message exchange and end-to-end integration have not been demonstrated in the current experiments.

6.4 An environment for testing improvement

A memory update changes available evidence; a tool update changes a procedure or implementation; a model update changes learned parameters or architecture; an improvement-procedure update changes how later candidates are generated, selected, or trained. These interventions have different costs and baselines. Logging them under a single "self-evolution" field would obscure the scientific question.

Fuxi can version the environment, supply held-out mechanism families, and record independent outcomes for each candidate. Evaluating a revised improvement procedure requires repeated downstream improvements across fresh problems, compared with an unchanged procedure given the same total compute, data, tools, and search opportunities. Merely selecting a better candidate from more attempts does not establish improvement of the improver. Candidate design and generation must not access evaluation answers; changes to a benchmark or objective are reported separately, not counted as performance gains on the original target.

These requirements must be evaluated against editable-agent and meta-agent baselines [28, 29, 30, 31], not only frozen tools. A bounded Fuxi hypothesis concerns reliable application and activation when experience scope, world mechanisms, skills, tools, and candidate versions change together. Existing provenance, validation, and routing mechanisms remain strong controls. Neither the new joint setting nor the combination alone establishes novelty; the study must isolate an additional benefit, failure mode, or necessary condition.

Separate the research workspace from the deployed artifact, but also audit what the updater and selection procedure could know before deployment. A solver with frozen inputs can still inherit information leaked through its builder, candidate summaries, retrieval indices, or prior selection outcomes. Retrospective cleanup of contaminated test rows cannot certify the resulting evolution history. Use separate development, selection-validation, and final evaluation data with availability cuts across all these roles; report adaptation on target tasks separately from transfer without target practice.

A proposal may improve a simulator, an inference algorithm, or the procedure that designs them. The current runtime is infrastructure for such comparisons. It does not demonstrate redesigned mathematical foundations, materials, general recursive improvement, or a route that guarantees the author's AI4.0 vision.

The author's AI1.0–4.0 frame motivates this broader program rather than defining an accepted scientific periodization. AI1.0 names present work without fixing it to statistical learning; 1.5 concerns world modeling and 2.0 causal models. The 3.0 motivation concerns human–agent shared discovery, where emotion, culture and unexpected insight can change an activity's direction. The 4.0 motivation is AI designing AI, broader than recursive improvement of the improvement procedure. These are distinct hypotheses with distinct evidence. A runtime that hosts their evaluation is not thereby AI4.0. The complete originating research-frame explanation is retained in Supplement E.

7 Relation to strong existing systems

Explicit persistent worlds already maintain programmable entities, attributes, relations and rules. PWM [22] connects such execution to generative rendering; SpatialOS [6, 33] is a historical distributed-world reference. Fuxi's present focus is the application-level meaning of change during an unfinished multi-subject activity. It does not introduce programmable worlds, claim PWM's rendering capability or demonstrate an advantage over an equally extended engine.

Live synthesis [39] formalizes residual obligations from old executions. GR(1) dynamic updating [40] constructs bridge controllers for a running system and selects continuation from its actual state when

computation finishes; its formulation accommodates the union of old and new variables. Those handover semantics differ and should not be conflated. Our finite installation-plan enumeration is narrower than those general synthesis methods. Continuing execution, carrying obligations and checking a moving activation state are inherited ideas; the system contribution here is their explicit connection to the world-effect and task-receipt interfaces.

Orleans and durable workflow systems [34, 47, 48] provide actor lifecycle, recovery, routing or version-management mechanisms on which a world application can be built. They still require application-specific interpretation of effects, goals and compatibility. Fuxi makes that interpretation explicit and tests a local combination, but has not shown that a competent implementation on those systems needs greater effort or cannot provide the same capability. No product-level performance comparison was run.

Model compilation, state migration and incremental verification also have direct precedents [20, 21, 43, 44]. On-the-fly synthesis and dependency reduction [41, 42] preclude a generic composition or slicing novelty claim. Learned relational dynamics [23, 24, 25], partial-observation control and sensing [45, 49], and editable-agent systems [28, 29, 30, 31] supply the appropriate controls for the proposed learning extensions. No learned result is inferred from the current symbolic episode.

Recent offchain verification [46] already combines newly discovered parameters, retrospective log queries, monitor initialization and simulation using an existing execution engine. Ordinary event sourcing [50] also separates retained event history from disposable projections. Thus neither reusing an executor for preview nor querying past events for newly relevant attributes is, by itself, an uncovered research contribution. The extended prior-work analysis and source-reading limits are retained in Supplement D and the evidence records.

8 Limits and the next discriminating study

The complete activity is real local execution of an authored finite model. It is not a real participant study, a learned world model, intelligent instrument generation, continuous physical simulation, independent host-plugin deployment or a long-duration shared society. Its local A2A subset uses trusted laboratory identities, not production authentication or interoperability certification. Partition recovery and external observation are separate witnesses with their own failure assumptions.

The comparative requirement now concerns actual value over a matched integration: fewer semantic adaptation errors, lower total development/maintenance cost, or more valid and voluntarily adopted activities under equal authority, information and budget. Source line counts, a larger collection of similar edits and extra passing controls cannot establish those outcomes. Initial integration, failed attempts, external tool work, runtime cost and human effort all belong in such a study. The present implementation makes those costs inspectable but leaves total authoring effort unmeasured.

Generation, learning, entry-point effects and world transfer require separate controls. Hold executable mechanisms and agent policies fixed when testing runtime behavior. Hold the runtime and permissions fixed when testing discovery or learning. An in-activity interface should be compared with an equally capable external interface, not one deprived of context. Mechanism families, partner lineages and author templates must be split independently; shared evaluators and hidden outcomes cannot enter candidate construction.

For shared discovery, preserve the original activity, any new direction, who proposed it and participants' subsequent choices. People may prefer continuation, rest or refusal. Adoption and long-term value cannot be inferred from object count, agent-generated approval or a completion receipt. For AI-designing-AI questions, architecture changes, updater changes and recursive use are separate interventions; independent multi-round evaluation and total cost are necessary. Supplement F retains the fuller evaluation design. None of these ambitions is treated as accomplished by the current fixture.

9 Conclusion

Fuxi makes one continuity object executable: an earlier action remains a fact, a changed world supplies its current response, and an accepted promise retains its agreed terms and history, with any authorized revision explicitly recorded. The complete local activity shows these bindings working together through HTTP collaboration, causal reception, continued growth, usable extension and actual receipts, including an adapted handwritten host. The separate bindings make the earlier cause, current response and participant-specific outcome inspectable, and let declared goals and receipt requirements be reused across the two supported execution paths. The matched controls delimit the contribution: this is a working system account with reproducible integration evidence, while independent gains in development cost, learning, transfer and participant value remain to be demonstrated. The broader purpose is a world in which creating a new possibility does not erase what already happened or what people and agents still owe one another.

Tools and assistance

Generative AI tools were used extensively for literature synthesis, drafting, code development, figure preparation, and internal checking. This assistance is distinct from the system experiments and computational diagnostics reported in the paper. Internal AI-assisted checks are not external peer review.

References

[1] Park et al. Generative Agents. https://arxiv.org/abs/2304.03442

[2] Vezhnevets et al. Concordia. https://arxiv.org/abs/2312.03664

[3] Piao et al. AgentSociety. https://arxiv.org/abs/2502.08691

[4] Akkil et al. Emergence World: A Platform for Evaluating Long-Horizon Multi-Agent Autonomy. https://arxiv.org/abs/2606.08367

[5] Akkil et al. Emergence World: Adversarial Stress-Testing of Long-Horizon Multi-Agent Systems. https://arxiv.org/abs/2609.17320

[6] Improbable. SpatialOS UnrealSDK introduction. https://github.com/spatialos/UnrealSDK/blob/master/docs/introduction.md

[7] VAST AI Research. Project Eden. https://www.tripo3d.ai/blog/project-eden

[8] Koksal et al. FlockWorld. https://arxiv.org/abs/2610.34564

[9] Wang et al. POET. https://arxiv.org/abs/1901.01753

[10] Wang et al. Voyager. https://arxiv.org/abs/2305.16291

[11] SIMA Team. SIMA2. https://arxiv.org/abs/2512.04797

[12] Hafner et al. Training Agents Inside of Scalable World Models. https://arxiv.org/abs/2509.24527

[13] Lamport. The Temporal Logic of Actions. https://www.microsoft.com/en-us/research/publication/the-temporal-logic-of-actions/

[14] Bailis et al. Coordination Avoidance in Database Systems. https://arxiv.org/abs/1402.2237

[15] Zhang et al. AGWM. https://arxiv.org/abs/2605.06841

[16] Khetarpal et al. Affordances Enable Partial World Modeling with LLMs. https://arxiv.org/abs/2602.10390

[17] Bogdanovych, Simoff, and Esteva. Virtual Institutions Prototype. https://www.ifaamas.org/Proceedings/aamas09/pdf/06_Demos/d_02.pdf

[18] Altera.AL et al. Project Sid. https://arxiv.org/abs/2411.00114

[19] Wang et al. RoboGen. https://arxiv.org/abs/2311.01455

[20] Hayden et al. Kitsune: Efficient, General-purpose Dynamic Software Updating for C. OOPSLA 2012. https://www.cs.umd.edu/~mwh/papers/kitsune.pdf

[21] Johnson, Calinescu, and Kikuchi. An Incremental Verification Framework for Component-Based Software Systems. CBSE 2013. https://eprints.whiterose.ac.uk/id/eprint/75438/1/cbse21_inc.pdf

[22] Huang et al. Programmable World Model. arXiv:2609.10540v2, 2026. https://arxiv.org/abs/2609.10540

[23] Battaglia et al. Interaction Networks for Learning about Objects, Relations and Physics. 2016. https://arxiv.org/abs/1612.00222

[24] Sanchez-Gonzalez et al. Learning to Simulate Complex Physics with Graph Networks. 2020. https://arxiv.org/abs/2002.09405

[25] Pfaff et al. Learning Mesh-Based Simulation with Graph Networks. 2020. https://arxiv.org/abs/2010.03409

[26] Pearl. An Introduction to Causal Inference. The International Journal of Biostatistics, 2010. https://ftp.cs.ucla.edu/pub/stat_ser/r354-reprint-corrected.pdf

[27] Shpitser and Pearl. Complete Identification Methods for the Causal Hierarchy. JMLR 9, 1941–1979, 2008. https://www.jmlr.org/papers/v9/shpitser08a.html

[28] Lee et al. Meta-Harness: End-to-End Optimization of Model Harnesses. arXiv:2603.28052v1, 2026. https://arxiv.org/html/2603.28052v1

[29] Zhang et al. Darwin Godel Machine: Open-Ended Evolution of Self-Improving Agents. arXiv:2505.22954v3. https://arxiv.org/abs/2505.22954v3

[30] Zhang et al. Hyperagents. arXiv:2603.19461v2, 2026. https://arxiv.org/abs/2603.19461v2

[31] Liu et al. Adaptive Auto-Harness: Sustained Self-Improvement for Agentic System Deployment on Open-Ended Task Streams. arXiv:2606.01770v3, 2026. https://arxiv.org/html/2606.01770v3

[32] A2A Project. Agent2Agent Protocol Specification v1.0.1. Accessed 20 November 2026. https://a2a-protocol.org/v1.0.1/specification/

[33] Improbable. SpatialOS RPGDemo and Database Sync Worker (archived official examples). https://github.com/spatialos/RPGDemo ; https://github.com/spatialos/database-sync-worker

[34] Bernstein et al. Orleans: Distributed Virtual Actors for Programmability and Scalability. MSR-TR-2014-41. https://www.microsoft.com/en-us/research/wp-content/uploads/2016/02/Orleans-MSR-TR-2014-41.pdf

[35] Chandy and Misra. Distributed Simulation: A Case Study in Design and Verification of Distributed Programs. IEEE TSE, 1979. https://www.cs.utexas.edu/~misra/scannedPdf.dir/DistrSimulation.pdf

[36] Lamport. Time, Clocks, and the Ordering of Events in a Distributed System. CACM, 1978. https://lamport.azurewebsites.net/pubs/time-clocks.pdf

[37] Helland. Life beyond Distributed Transactions: an Apostate's Opinion. CIDR, 2007. https://www.cidrdb.org/cidr2007/papers/cidr07p15.pdf

[38] Mordvintsev et al. Growing Neural Cellular Automata. Distill, 2020. https://distill.pub/2020/growing-ca/

[39] Finkbeiner, Klein, and Metzger. Live synthesis. Innovations in Systems and Software Engineering 18, 443–454, 2022. https://link.springer.com/article/10.1007/s11334-022-00447-5

[40] Dynamic Update for Synthesized GR(1) Controllers. ICSE 2022. https://smlab.cs.tau.ac.il/syntech/dupdate/dupdate-icse22-cr.pdf

[41] A Compositional Framework for On-the-Fly LTLf Synthesis. arXiv:2508.04116v2. https://arxiv.org/html/2508.04116v2

[42] Dynamic Hierarchical Reactive Controller Synthesis. arXiv:1510.07246v2. https://arxiv.org/pdf/1510.07246v2

[43] Hayden et al. Specifying and Verifying the Correctness of Dynamic Software Updates. Author preprint. https://www.cs.tufts.edu/~jfoster/papers/vstte12.pdf

[44] Fernandez Adiego et al. Modelling and Formal Verification of Timing Aspects in Large PLC Programs. IFAC World Congress, 2014. https://cds.cern.ch/record/1956687/files/CERN-ACC-2014-0226.pdf?version=1

[45] Fu and Topcu. Integrating active sensing into reactive synthesis with temporal logic constraints under partial observations. arXiv:1410.0083. https://arxiv.org/pdf/1410.0083

[46] Capretto et al. Offchain Runtime Verification (for The Tezos Blockchain). Author manuscript; accessed 21 November 2026. https://software.imdea.org/~cesar/papers/capretto25offchain.pdf

[47] Microsoft. Compatible grains. Official Orleans documentation; accessed 21 November 2026. https://learn.microsoft.com/en-us/dotnet/orleans/grains/grain-versioning/compatible-grains

[48] Temporal. Roll out and pin Workflows. Official documentation source; accessed 21 November 2026. https://github.com/temporalio/documentation/blob/main/docs/production-deployment/worker-deployments/worker-versioning/roll-out-and-pin.mdx

[49] Fu, Dimitrova, and Topcu. Abstractions and sensor design in partial-information, reactive controller synthesis. ACC 2014. https://eprints.whiterose.ac.uk/id/eprint/156437/1/acc14.pdf

[50] Microsoft. Event Sourcing pattern. Azure Architecture Center; accessed 21 November 2026. https://learn.microsoft.com/en-us/azure/architecture/patterns/event-sourcing