Shared Clocks and Historical Targets

B-Theory research programme

Explanatory appendix, version 0.2
9 November 2026

Prepared with AI assistance. Internal AI review is distinct from external peer review.

Purpose and placement

This appendix connects the static MIND construction to retrospective trajectory checks and to a small, fully charged archive-access problem. It consolidates distributed-temporal research from 4 November 2026. The disposition is an explanatory appendix linked from MIND M4 and Acquirable Predictive Models; it is not a new standalone research paper.

MIND §5.1–5.3, Proposition 3 (pp.6–7), already establishes the static parity obstruction and requires matched solver comparisons. Acquirable Predictive Models §7, Proposition 10 and Examples 11–12, already distinguishes acquiring a current predictive class from recovering an origin. Diagnosing While Repairing v0.2.5 §1.4, Proposition 1.2, already supplies a general erasure certificate; §1.5, Proposition 1.3, warns that fixed-order normalization can increase cost. Reflective Self-Models R4 §4 already establishes equality of its matched tracker mappings. We use those results as antecedents, rather than present them again as discoveries [1–4].

The companion's §3, Theorem 3.2, already gives an adaptive-versus-fixed repair-order separation. The archive example below instead compares finite costs of nonmutating query selections; it does not restate that destructive-repair theorem.

The additional content is a shared-clock model and an explicit historical-target example with exact communication accounting. Temporal constraints, signed-graph consistency, variable elimination and adaptive test selection are established methods [5–9]. Historical priority of this particular specialization is not established.

A. Compatibility of records is a trajectory question

Let the body have finite state space $\mathcal{X}$, horizon $0, \dots, H$, initial set I, and specified transition relations R_t. A report asserts a relation about the state at its acquisition time. Receipt time determines which evidence is available, not which state the report describes. Omitted acquisition time is already a recognized source of sensor inconsistency [10]. Write O_t for the states satisfying all reports acquired at t. No reports means $O_t = \mathcal{X}$. Define

$$F_0 = I \cap O_0, \qquad F_{t+1} = O_{t+1} \cap \{y : \exists x \in F_t,\ (x, y) \in R_t\}.$$

Induction on the prefix shows that $F_H \neq \varnothing$ exactly when some complete trajectory satisfies all records and transitions. This is ordinary finite-state reachability, not a new inference principle.

For example, the parity reports $A_0 + B_0 = 0$, $B_0 + C_0 = 0$, $C_1 + A_1 = 1$ fit $000 \to 100$ when an A-flip is permitted. Erasing the time labels creates a false static odd-parity obstruction. Moving the final report to time zero produces a genuine same-time contradiction. Keeping the original times but requiring persistence produces a transition inconsistency despite nonempty individual slices. A report can remain true about the past while being unsuitable as a statement about the present.

Existence of a compatible trajectory does not establish which history occurred, or show that clocks caused an observed disagreement. Conversely, learning a target requires that every surviving history have the same target label. Neither statement is a future dynamic-controllability guarantee.

A.1 One offset per module

For integer uncertainty bound κ, module j has one constant offset $d_j \in [-\kappa, \kappa]$. Every record from that module obeys

$$h_\rho = \tau_\rho + d_j.$$

With authoritative receipt time r_ρ and end-to-end delay bound D, require $0 \le \tau_\rho \le H$ and $0 \le r_\rho - \tau_\rho \le D$. Its exact module-wide domain is

$$K_j = [-\kappa, \kappa] \cap \bigcap_{\rho \in j} [h_\rho - H, h_\rho] \cap \bigcap_{\rho \in j} [h_\rho - r_\rho, h_\rho - r_\rho + D],$$

intersected with the integers. An empty domain is a timing-contract conflict.

For one module reporting complete n-bit states y_i at sorted local times h_i, assume a free initial state, no further state constraints, and at most one changed body bit per tick; no-change ticks are allowed. A shared-offset trajectory exists exactly when $K_j \neq \varnothing$ and

$$\mathrm{Ham}(y_i, y_{i+1}) \le h_{i+1} - h_i \quad \text{for every consecutive pair.} \tag{A1}$$

Necessity follows because the common offset cancels from each time difference. For sufficiency, choose any offset in K_j, flip each differing coordinate once between successive reports, pad with no-change steps, and hold the endpoint states outside those segments. This works for every allowed offset. Equal stamps therefore require equal complete reports. Partial observations, another module, or a restricted initial set can invalidate this sufficiency.

In the two-bit example $H = 3, \kappa = 1, D = 2$, one module reports 00 at $h = 1, r = 2$ and 11 at $h = 2, r = 3$. Its domain is $\{-1, 0, 1\}$, but the two-bit distance exceeds the one-tick gap. Independent per-record retiming accepts, for example, offsets $+1, -1$, acquisition times $0, 3$, and path $00, 01, 11, 11$. That is a relaxation, not a valid common-clock explanation. The preserved exact enumeration finds zero shared-offset solutions and 18 relaxed offset/trajectory pairs.

A.2 What the elimination-width bound actually says

Reject an empty K_j as the timing-contract conflict described above. Modules with no records need no clock variable in the elimination. For the remaining modules, represent each body slice X_t as one 2^n-valued variable. A record factor involves d_j and its candidate body slices $X_{h_\rho - d}$, $d \in K_j$. Process that factor at its latest candidate time u_ρ. Put $\ell = \max(1, 2\kappa)$. Retain body slices from $\max(0, t - \ell)$ through t, and retain clock d_j from its first through its last factor-completion layer, including long intervening gaps.

If at most a clocks are alive in any bag, bag and separator domains are bounded by

$$V = 2^{n(\ell+1)}(2\kappa+1)^a, \qquad W = 2^{n\ell}(2\kappa+1)^a. \tag{A2}$$

Every report span is at most 2κ, every transition fits because $\ell \geq 1$, and each variable occurs in consecutive bags. These facts prove a valid path decomposition. At a layer, join its factors with the incoming separator table and project onto the next separator. Induction identifies the outgoing table with exactly the prefix-extendible assignments. The final empty-variable relation is nonempty exactly for a satisfiable instance. Backtracking retained choices yields one trajectory and one offset per module. A negative trace must verify equality of every projected table, including alleged zero entries; merely checking listed positive tuples is insufficient.

This is a concrete application of classical variable elimination [6,7]. It is not an optimal-width theorem. A dense decision frontier uses $O(W)$ bits, in addition to input/tuple storage, while enumeration uses bag domains bounded by V. Predicate evaluation and arithmetic have their ordinary costs. Polynomial dependence on an explicit horizon for fixed n, κ, a does not mean polynomial dependence on a binary-encoded large horizon. Small clock uncertainty alone is insufficient: a can equal the total number of modules, and n remains in an exponent. Input routing, table transmission, computation and proof delivery must be charged under a stated network contract. No generic scalability or architectural advantage follows.

B. A historical target with a fully specified access contract

Consider eight declared histories $w = (a, u, v) \in \{0,1\}^3$, and fix the audit target before selecting queries:

$$f(a,u,v) = \begin{cases} u, & a = 0, \\ v, & a = 1. \end{cases} \tag{B1}$$

The body consists of two calibration bits and flags U, V. For times $0, \ldots, 4$, the calibration pair follows $00, 00, 01, 11, 11$ when $a = 0$, and $11, 11, 01, 00, 00$ when $a = 1$; $U = u, V = v$ remain fixed. Records $P : 00, h = 2, r = 3$ and $Q : 11, h = 2, r = 3$ have offsets $(+1, -1)$ or $(-1, +1)$, respectively, with $\kappa = 1, D = 2$. Thus a specifies their physical order.

At times 5 and 6 reset U and V, and at times 7 and 8 reset the first and second calibration coordinates. Every step changes at most one bit. From time 8 onward the body is 0000. This eight-history template is a new declared restriction of the model, not a deduction that the earlier 36-explanation example has only eight possibilities.

Two passive archives already store U_0 and V_0. Their acquisition and storage are common sunk costs for all comparators. Four different leaves have only their respective ports: A returns a, using a preexisting trusted clock reference; B returns archived u; C returns archived v; D reports whether the current body is nonzero, always zero during querying. Establishing the reference is a precondition, not free synchronization.

Queries start after time 8. A shared half-duplex bus permits one outstanding serialized transaction. Each request has a two-bit query/version header and two-bit port address; each response has a two-bit header and one answer bit. Capacity is one bit per round, with one additional leaf read/compute round and no other propagation delay. Each query therefore costs seven transmitted bits and eight rounds. Root selection/reporting is a zero-additional-round combinational operation equally allowed to both controllers. There are no batches, overlapping transactions, unsolicited pushes, interleaf exchanges, hidden timing/silence codes, failures or retransmissions. A leaf may

encode its own answer, but cannot read the other private inputs. A joint multiplexer port would change the contract.

B.1 Exact optimum within this contract

Ask A, then ask B when $a = 0$, otherwise C, and report the archive answer. Every history is identified with two transactions: exactly 14 bits and 16 rounds. No single port identifies f, so every guaranteed policy has a branch requiring at least two transactions. This proves the adaptive optimum.

A fixed identifying selection must contain all of A, B, C. If B is absent, histories $(0, 0, 0)$ and $(0, 1, 0)$ agree on all remaining ports but have opposite targets. If C is absent, use $(1, 0, 0)$ and $(1, 0, 1)$. If A is absent, use $(0, 0, 1)$ and $(1, 0, 1)$. Conversely A, B, C suffice. Hence the fixed optimum is exactly 21 bits and 24 rounds. Even a fixed query order with answer-dependent early stopping has worst-case depth three: each choice of two distinct informative ports has an ambiguous answer cell. The saving comes from selecting which archive to read after learning a.

These are incremental wire-bit and deadline optima under the stated interface, not universal network bounds, Shannon-information costs, or total system construction costs. Both controllers leave the irrelevant archive bit unknown. Recovering the complete history is unnecessary.

For general finite histories and stable nonmutating deterministic ports, a fixed selection identifies a target exactly when it separates every differently labelled pair. An adaptive decision tree succeeds exactly when some affordable first query has successful continuations for every nonempty answer cell, subtracting its charged bits and duration. These are standard test-selection/decision-tree principles [8,9], not new general theorems. The available-port set must be included in the recursion; changing permissions, consuming archives or time-varying answers requires additional state.

B.2 Clocks cannot recreate missing records

If only A and fresh current-body readings are available, histories $(0, 0, 0)$ and $(0, 1, 0)$ have the same offsets and current body, but opposite targets. Every available answer agrees. Inductively, an adaptive policy receives the same transcript and cannot be correct on both histories. Zero-error randomization cannot repair this; condition on the same private tape. No larger communication budget helps.

The distinction from the companion's erasure theorem is essential. If historical flags survive in archives, the complete physical state includes those archives. They must not be silently omitted to assert irreversible loss of all information. When archives exist but are inaccessible, the obstruction is actual-port insufficiency. When they do not exist and no other retained variable carries the flag, the example is an instance of the general erasure principle already proved in Diagnosing While Repairing §1.4, Proposition 1.2. Current certainty after reset is likewise already distinguished from origin recovery in Acquirable Predictive Models §7. The archive construction illustrates these principles; it does not replace them.

B.3 The architecture comparison is matched

A central coordinator and a distributed implementation have identical input placement, ports, local encoding rights, bus, deadlines and output root. Either can ask A and conditionally select B/C; neither can beat the two-query lower bound. Thus 14 versus 21 compares adaptive with fixed selection. It does not compare distributed with centralized architecture.

R4 §4 already demonstrates equality under identical evidence-processing mappings; its external

observer communicates freely. Here communication is explicitly priced, and matched controllers still attain the same optimum. Moving archives into one leaf or adding a joint target port, batching or concurrency changes the allowed interface and requires new accounting.

C. Verification, interpretation and release scope

The portable package retains copies of the original standard-library checkers, run in separate reproduction directories so the source records remain unchanged. Iteration 1 checks finite trajectory/parity constructions; iteration 2 checks shared-offset witnesses and complete negative traces; iteration 3 checks all eight templates and all sixteen port selections, 4,096 version-space/selection comparisons, and 1,020 version-space/depth comparisons against a separate explicit tree enumerator. A separately implemented iteration-3 checker confirms the same optima. Exact run outcomes and hashes are supplied in the release manifest. Enumeration corroborates the finite examples; the arguments above carry the stated quantified claims.

B-Theory here names the Being/Becoming research programme. No premise about the philosophical B-theory of time is used. The appendix formalizes states, changes and available historical evidence. It supplies no subjective-consciousness criterion, inherent distributed-architecture superiority, biological validation or unrestricted scalability result. No new experiment or broader temporal research is needed for this inclusion decision.

References

[1] B-Theory research programme. The MIND: Five Finite Constructions. Research edition with P1 and P2, 5 October 2026, §5, Proposition 3, pp.6–7. https://www.b-theory.com/manuscripts/mind-five-finite-constructions-r1-2026-10-05.pdf

[2] Brian Theory. From Observable States to Acquirable Predictive Models. Revised report, 7 October 2026, §§3,6–8. https://www.b-theory.com/manuscripts-html/acquirable-predictive-models/

[3] B-Theory research programme. Diagnosing While Repairing. Companion v0.2.5, 8 November 2026, §§1.4–1.5,3. https://www.b-theory.com/manuscripts/diagnosing-while-repairing-companion-v0.2.5.pdf

[4] B-Theory research programme. When Does an Explicit Self-Model Help? R4, 14 October 2026, §§1,3–4. https://www.b-theory.com/manuscripts-html/reflective-self-models/

[5] F. Harary. On the number of bi-colored graphs. Pacific Journal of Mathematics 8 (1958), §6, p.754. https://msp.org/pjm/1958/8-4/pjm-v8-n4-p11-p.pdf

[6] R. Dechter, I. Meiri, J. Pearl. Temporal constraint networks. Artificial Intelligence 49 (1991), 61–95. https://doi.org/10.1016/0004-3702(91)90006-6

[7] A. Petcu, B. Faltings. A Scalable Method for Multiagent Constraint Optimization. IJCAI 2005, 266–271. https://www.eecis.udel.edu/~decker/courses/886f06/pubs/DPOP.pdf

[8] D. Golovin, A. Krause, D. Ray. Near-Optimal Bayesian Active Learning with Noisy Observations. NeurIPS 2010. Its Bayesian expected-cost setting differs from the exact worst-case contract here. https://arxiv.org/abs/1010.3091

[9] S. Javdani et al. Near Optimal Bayesian Active Learning for Decision Making. AISTATS 2014, PMLR 33:430–438. https://proceedings.mlr.press/v33/javdani14.html

[10] M. Robinson. Sheaves are the canonical data structure for sensor integration. 2016, §5.2 discusses omitted acquisition time as an inconsistency source. https://arxiv.org/abs/1603.01446