EXACT REALIZATION AND SUPREMAL SYNTHESIS UNDER HISTORY-DEPENDENT VETO AUTHORITY

PHILLIP KINGSTON

Abstract. Classical conjunctive co-observability fixes the supervisors assigned to each controllable event and does not restrict whether a veto lies in its issuer's current action domain. We study prefix-closed decentralized control when veto authority varies with hidden execution history. Legitimacy is specification-relative: actions are constrained at target-reachable decision points. Exact realization requires an aligned witness—one supervisor must know both that a continuation is forbidden and that its veto is feasible throughout the same information cell. We prove that controllability and *legitimate co-observability* are necessary and sufficient, construct canonical policies, and give a minimal three-point obstruction showing that separate decision and authority witnesses do not compose.

For supremal synthesis, we extend fixed-ambient relative co-observability with issuer-safe consistency blocks and a distinct veto-coverage condition. Hidden authority schedules require whole-fiber decisions: all extended histories with one operational projection must be retained or removed together. Under compatible finite monitors, schedule classes capture these fibers, while path coverage and semantic self-pairs repair the failure of class transitions to determine word-level membership. The resulting fixed point computes the supremal relative language exactly. Although the associated observer is exponential in the worst case, construction and synthesis are polynomial in the explicit product for every fixed current-estimate width, with a degree that may depend on the width.

1. Introduction

Decentralized supervisory control asks local supervisors with partial information to coordinate a global decision. Under classical conjunctive fusion, each controllable event has a fixed participating supervisor set (Rudie and Wonham, 1992). In delegated or revocable control architectures, however, the vetoes that affect the plant vary with execution history. A supervisor may know that an event must be disabled but not know that it is authorized to veto; another may know its authority but not the required decision. Neither can act. The logical obstruction is

$$\exists i\,(D_i \wedge U_i) \;\not\equiv\; (\exists i\, D_i) \wedge (\exists j\, U_j).$$

one issuer must align decision knowledge and action feasibility on one information cell.

Date: 1 August 2026.
Key words and phrases. decentralized supervisory control, co-observability, history-dependent authority, relative co-observability, supremal synthesis, agentic artificial intelligence.
Email: phillip.kingston@gmail.com. ORCID: 0009-0005-2256-3079.

We model authority as a pointwise action domain. A veto is legitimate only when it remains feasible for its issuer at every specification-reachable point in the cell on which the local policy is uniform. This intensional constraint matters in a capability-checked command architecture even when the current fusion engine ignores an inactive output: an out-of-domain veto is still an invalid issued command, can violate an audit or interface contract, and may be stale, replayed, or misrouted to a consumer that does not apply the same filter. The model therefore separates behavioral effect from permission to issue.

A concrete architecture has four roles. A trusted authority service maintains the current configuration and active issuer set; local supervisors receive only their prescribed observations, which need not reveal that set; the fusion engine receives the authority service's authenticated active-set signal and uses it to determine closed-loop behavior; and synthesis or an output gateway certifies that each local veto lies in the issuer's action domain. Thus the fusion layer may know the active set while a local supervisor does not know its own current status, without treating the active-set signal as a local observation. Authority updates are exogenous uncontrollable events; delegation itself is not synthesized. Hidden updates create a second coupling: extended histories with one operational projection are not independent synthesis choices and must be retained or removed as a complete fiber.

The two paper-specific mechanisms are the aligned-witness obstruction and the coverage and self-pair repair at the boundary between classes and words. The fixed-ambient relative construction supplies the bridge between them. This paper makes three contributions:

(1) **Aligned-witness irreducibility and exact realization.** A minimal three-point construction proves that separate decision and authority witnesses do not compose (Proposition 3.4 and Corollary 3.5); this is the part not inherited from a generic observation-uniform feasibility principle. Against that boundary, we define legitimate co-observability and prove that a prefix-closed target is exactly realizable if and only if its lift is controllable and every forbidden continuation has one cell-uniform legitimate witness (Theorem 3.2). The proof gives canonical policies, and authority observability identifies where the obstruction collapses (Corollary 3.10). Under componentwise finite input, verification is PSPACE-complete, matching the classical worst-case baseline (Rohloff et al., 2003), but is polynomial after reachable-product materialization (Proposition 5.4).

(2) **Relative authority-safe bridge.** We explicitly extend the fixed-ambient union-closure geometry of Cai–Zhang–Wonham relative co-observability to issuer-specific authority-safe cells. The inherited consistency scaffold is filtered by issuer safety; the new coverage obligation repairs the resulting vacuity and induces effective uncontrollability wherever no safe issuer exists (Definition 4.7).

(3) **Exact whole-fiber finite synthesis.** Schedule classes collect all product states reached by hidden authority schedules of one operational word. Because different words can reach one class through different retained paths, class transitions alone do not determine word-level membership. Coverage and semantic self-pairs repair this mismatch. The resulting alternating fixed point computes the supremal relative language exactly (Theorem 5.19). Its worst case is the exponential subset construction; bounded current-estimate width identifies a tractable special class rather than changing that worst-case fact. The supplementary validator constructs schedule classes, runs the exact pruning operators, and compares the result with exhaustive enumeration on small instances (Appendix A).

Under legal-point neutrality, fixed conjunction over faithfully delivered outputs realizes the same language as history-dependent active-set conjunction (Theorem 3.11).

Several neighboring decentralized-control architectures vary related but distinct objects. General conjunctive/disjunctive architectures fix the fusion choice by event (Yoo and Lafortune, 2002b); dynamic decision-fusion and dynamic-default models allow the default or fusion treatment to change during operation (Yoo and Lafortune, 2002a; Takai and Ushio, 2005); and state-dependent controllability changes centrally whether an event is controllable at a plant state (Wang and Cai, 2009). Fixed nonhomogeneous assignments and conditional-decision architectures provide further eventwise fusion variants (Mortazavian and Lin, 1993; Yoo and Lafortune, 2004). The present model instead attaches a pointwise action domain to each issuer. A fusion layer may ignore an inactive supervisor's output, while the command interface can still reject or record that output as outside the issuer's domain. Behavioral fusion and issuer admissibility are therefore separate requirements; the formal separation is given in Corollary 3.8.

TABLE 1. Selected neighboring decentralized-control mechanisms.

Model	Control feature that varies, if any	Behavioral semantics	Action-domain issue
Fixed general fusion (Yoo and Lafortune, 2002b)	None: each controllable event is assigned a priori to union or intersection fusion	Local enabled-event sets are combined according to that event's fixed fusion rule	No history-dependent pointwise issuer action domain is modeled
Dynamic fusion/default (Yoo and Lafortune, 2002a; Takai and Ushio, 2005)	The default-enable/default-disable classification, and therefore the effective fusion treatment, may vary with observed history	Local enable/disable decisions are evaluated under the dynamically selected default or fusion treatment	Local controllable-event assignments remain fixed; the model varies defaults or fusion treatment rather than history-dependent pointwise issuer admissibility
State-dependent controllability (Wang and Cai, 2009)	An event's controllability status varies with plant state	At a state where an event is uncontrollable, the supervisor cannot disable its enabled transition	Plant-state-dependent controllability in a centralized supervisory-control model; no issuer-indexed information-cell admissibility constraint
Present model	The effective issuer set varies with extended history and may be hidden from local supervisors	Active-set conjunction uses only the currently effective issuers	At specification-reachable decision points, every issued veto must lie in that issuer's pointwise action domain

Other neighboring DES lines alter inference, component reliability, or the allocation of control rather than the admissibility of each issued command. Inference-based decentralized

control ranks local decisions by ambiguity and uses information about other supervisors' uncertainty in fusion and synthesis (Kumar and Takai, 2007; Takai and Kumar, 2008). Reliable general architectures preserve behavior under failures of local supervisors (Liu and Lin, 2010), while actuator and sensor fault models treat local losses of controllability or observability through post-fault supervisors (Dai and Lin, 2018). Supervisor localization instead decomposes a synthesized monolithic control action among partial-observation local controllers (Zhang et al., 2017). The cited mechanisms are mechanically adjacent to varying issuer capability, but they do not impose a history-dependent local action domain that every issued veto must respect. Event forcing under partial observation supplies another nonstandard action primitive with an information-cell consistency condition and a failure of union closure (Tong and Cai, 2026); it does not introduce issuer-safe coverage or whole-fiber coupling under hidden authority schedules.

Standard imperfect-information games usually make action availability constant on an information set; stochastic-action-set models instead reveal the realized set to the player (Schwarz et al., 2026a,b). Under the availability analogy, these lie in the observable-authority regime identified by Corollary 3.10. Uniform strategies constrain the action selected on related histories (Maubert and Pinchinat, 2014). Knowledge-based DES already ties epistemic conditions to control decisions (Ricker and Rudie, 2000; Ritsuka and Rudie, 2023). In particular, Ritsuka and Rudie's solvability expressions couple what a supervisor knows to the action it must take. The present model adds a separate issuer-admissibility constraint: even when a supervisor knows that a veto is required, the veto must lie in its action domain at every indistinguishable point, while authority status may itself be hidden. Their question is which decision a knowledgeable supervisor should emit; ours additionally asks whether that supervisor is permitted to emit a veto throughout the cell. Within the present model, this additional gap disappears when authority is locally observable on the relevant cells (Corollary 3.10).

Relative observability and relative co-observability supply the fixed ambient-language mechanism used to recover union closure (Cai et al., 2015, 2019). Complementary work studies exact branch-set preservation under fixed actuator interfaces through equality-support covers (Kingston, 2026a); the present setting instead allows the effective issuer interface to vary with hidden history. The added ingredients here are issuer-specific safe cells, the separate coverage obligation, and whole-fiber synthesis under hidden authority schedules. Set-valued estimates are standard (Cassandras and Lafortune, 2021); the schedule-class algorithm additionally needs path coverage and self-pairs because class-edge membership does not determine the membership of every word reaching that class. Naive authority-indexed event splitting is not a reduction of the present problem: it exposes the hidden authority mode to local policies, models fusion participation rather than issuer admissibility, and permits different copies of one operational fiber to be pruned independently. Restoring those constraints reconstructs the obligations represented directly here (Remark 4.10).

Throughout the main realization and synthesis results, "legitimate" means specification-relative admissibility at target-reachable decision points. Section 2 defines the stronger plant-wide variant and identifies a regime in which existence is unchanged. We otherwise restrict attention to prefix-closed behavior, nonempty active sets, faithful output delivery, exogenous authority updates, and explicit finite-output observers for synthesis. Delegation

synthesis, communication, marked nonblocking behavior, faults, malicious supervisors, and plant-wide synthesis without authority observability are outside scope.

2. Model and issuer-admissible control

2.1. **Extended plant and operational specification.** Fix a finite supervisor set $\mathcal{I}$. Let

$$\Sigma = \Sigma_c \,\dot\cup\, \Sigma_u$$

be the operational alphabet, partitioned into controllable and uncontrollable events, and let Δ be a disjoint alphabet of authority updates. Put $E = \Sigma \,\dot\cup\, \Delta$. The extended plant is a nonempty prefix-closed language

$$\mathcal{H} \subseteq E^*.$$

The projections onto operational and authority events are denoted by p_Σ and p_Δ.

Let Conf be a set of authority configurations and $c : \mathcal{H} \to \mathsf{Conf}$ the configuration after an extended history. The configuration may depend on the full extended history: Δ records updates with no operational projection, but operational events may also change c, and hence the effective veto sets, at their target. The effective veto relation is a map

$$\mathsf{A} : \mathsf{Conf} \times \Sigma_c \longrightarrow 2^{\mathcal{I}} \setminus \{\varnothing\}.$$

Thus $\mathsf{A}(c,\sigma)$ is the nonempty set of supervisors whose vetoes affect σ in configuration c. The nonemptiness assumption excludes authority states in which a nominally controllable event has no effective veto holder. Representing such a configuration faithfully would require configuration-dependent controllability, which lies outside the present model. Assuming nonemptiness is therefore an explicit scope restriction, not a step taken without loss of generality.

Let $K \subseteq \mathsf{p}_\Sigma(\mathcal{H})$ be a nonempty prefix-closed operational specification. Its lift is

$$\widetilde{K} := \{h \in \mathcal{H} : \mathsf{p}_\Sigma(h) \in K\}. \tag{1}$$

Because authority updates do not alter the operational projection, the lift is closed under every available Δ-extension. Authority updates are therefore exogenous, undisablable plant behavior.

Each supervisor has an arbitrary observation map

$$\mathsf{P}_i : \mathcal{H} \longrightarrow Y_i.$$

The map may be a natural projection, a perfect-recall observation, or the output of a finite transducer. Only the information partition induced by P_i enters Sections 2–4: observation maps with the same fibers yield the same decision cells and closed-loop conditions. For the partial-information reading we take P_i to factor through an observation sequence, so $\mathsf{P}_i(h)$ depends only on the outputs emitted along h. No finite eventwise realization is required there. Section 5 additionally requires each observation map to be realized by a compatible plant-total Moore monitor. Write $h \sim_i h'$ when $\mathsf{P}_i(h) = \mathsf{P}_i(h')$.

2.2. **Decision points and legitimacy.** The plant-level and specification-reachable controllable decision points are

$$X_{\mathcal{H}} := \{(h,\sigma) : h \in \mathcal{H},\ \sigma \in \Sigma_c,\ h\sigma \in \mathcal{H}\}, \qquad X_K := \{(h,\sigma) \in X_{\mathcal{H}} : h \in \widetilde{K}\}.$$

The latter split into legal and forbidden points

$$X_K^+ := \{(h,\sigma) \in X_K : h\sigma \in \widetilde{K}\},$$
$$X_K^- := X_K \setminus X_K^+.$$

For every $x = (h,\sigma) \in X_{\mathcal{H}}$ write $A_x := \mathsf{A}(c(h),\sigma)$.

A local supervisor is a map

$$g_i : Y_i \times \Sigma_c \longrightarrow \{0,1\},$$

where 0 denotes veto and 1 admit. Under active-set conjunction, the closed language $L_A(g/\mathcal{H})$ contains ε and is generated recursively as follows. From $h \in L_A(g/\mathcal{H})$, every available extension by $e \in \Delta \cup \Sigma_u$ is admitted, while an available $\sigma \in \Sigma_c$ is admitted exactly when

$$\bigwedge_{i \in \mathsf{A}(c(h),\sigma)} g_i(\mathsf{P}_i(h),\sigma) = 1. \tag{2}$$

The fusion layer is assumed to receive or reconstruct the current set $\mathsf{A}(c(h),\sigma)$ exactly. This active-set signal need not be included in any local observation P_i; hidden authority means precisely that it may vary within a local information cell.

Authority determines a local action domain at $x = (h,\sigma) \in X_{\mathcal{H}}$:

$$\Gamma_i(x) = \begin{cases} \{0,1\}, & i \in A_x, \\ \{1\}, & i \notin A_x. \end{cases}$$

Admission is always admissible; veto is admissible only for an active issuer.

Definition 2.1 (Active-set realization and legitimacy)**.** A tuple $g = (g_i)_{i\in\mathcal{I}}$ is an *active-set realization* of K if

$$L_A(g/\mathcal{H}) = \widetilde{K}.$$

It is *K-legitimate* if, for every $x = (h,\sigma) \in X_K$ and $i \in \mathcal{I}$,

$$g_i(\mathsf{P}_i(h),\sigma) \in \Gamma_i(x).$$

It *exactly realizes* K when both properties hold.

Definition 2.2 (Plant-wide legitimacy)**.** A tuple g is *plant-wide legitimate* if, for every $x = (h,\sigma) \in X_{\mathcal{H}}$ and $i \in \mathcal{I}$,

$$g_i(\mathsf{P}_i(h),\sigma) \in \Gamma_i(x).$$

Thus an inactive issuer must emit admission at every plant decision point, not only at points reachable under the target.

Plant-wide legitimacy implies K-legitimacy for every K, but the converse fails in general: a K-legitimate policy may issue an unauthorized veto at a plant history outside $\widetilde{K}$, which is unreachable after exact realization. The main theory uses the specification-relative notion because it characterizes implementations on the realized behavior. The following condition identifies a useful regime in which strengthening the requirement does not change existence.

For i, y, σ, define the plant-level cell and authority region

$$\mathsf{C}_i^{\mathcal{H}}(y,\sigma) := \{(h,\sigma) \in X_{\mathcal{H}} : \mathsf{P}_i(h) = y\}, \qquad R_{i,\sigma}^{\mathcal{H}} := \{(h,\sigma) \in X_{\mathcal{H}} : i \in A_{(h,\sigma)}\}.$$

Authority is *plant-wide locally observable* when every nonempty $\mathsf{C}_i^{\mathcal{H}}(y,\sigma)$ is either contained in $R_{i,\sigma}^{\mathcal{H}}$ or disjoint from it.

Proposition 2.3 (Plant-wide legitimacy under authority observability)**.** *If authority is plant-wide locally observable, then K has an exact K-legitimate realization if and only if it has an exact plant-wide legitimate realization.*

Proof. The reverse implication is immediate. Let g be an exact K-legitimate realization. For each local argument (i, y, σ) whose plant-level cell is wholly inactive for i, replace its output by 1; leave every other output unchanged. Plant-wide local observability makes the resulting tuple g' plant- wide legitimate. On every point of X_K, either the cell is wholly active and unchanged, or it is wholly inactive and K-legitimacy already forced the output to be 1. Hence g' agrees with g on all target-reachable decision points, and the word-length realization argument gives $L_A(g'/\mathcal{H}) = \widetilde{K}$. □

This proposition marks the observable-authority collapse boundary; it does not solve plant-wide legitimate synthesis when authority remains hidden. In that regime an off-target information cell may mix active and inactive points, so the output repair used in the proof is unavailable.

Definition 2.4 (Controllability)**.** The lift $\widetilde{K}$ is *controllable* in $\mathcal{H}$ in the standard Ramadge–Wonham sense (Ramadge and Wonham, 1987) if

$$h \in \widetilde{K}, \quad u \in \Sigma_u, \quad hu \in \mathcal{H} \quad \Longrightarrow \quad hu \in \widetilde{K}.$$

Example 2.5 (The alignment problem)**.** Suppose a controllable event σ is forbidden after a history h. One active supervisor distinguishes h from every legal σ-decision but has the same observation at another point where it is inactive. A second active supervisor knows from its observation that it is authorized, but merges h with a legal σ-decision. The first supervisor cannot issue a legitimate uniform veto and the second cannot issue a correct one. Proposition 3.4 in Section 3 formalizes this obstruction as a minimal three-point realization.

3. Exact realization and same-witness authority

For $i \in \mathcal{I}$, $y \in Y_i$, and $\sigma \in \Sigma_c$, define the specification-level decision cell

$$\mathsf{C}_i(y, \sigma) := \{(h', \sigma) \in X_K : \mathsf{P}_i(h') = y\},$$

and write

$$R_{i,\sigma} := \{(h, \sigma) \in X_K : i \in \mathsf{A}(c(h), \sigma)\}$$

for the points at which i holds the relevant veto right.

Definition 3.1 (Legitimate co-observability)**.** The language K is *legitimately co-observable* if, for every $x = (h, \sigma) \in X_K^-$, there exists $i \in A_x$ such that

$$\mathsf{C}_i(\mathsf{P}_i(h), \sigma) \subseteq X_K^- \qquad \text{and} \qquad 0 \in \bigcap_{x' \in \mathsf{C}_i(\mathsf{P}_i(h), \sigma)} \Gamma_i(x'). \tag{3}$$

Such an i is a *legitimate witness* for x.

Condition (3) joins two uniformity requirements on one cell. Every indistinguishable available occurrence of σ must require a veto, and the same veto action must remain feasible for the same issuer throughout that cell. Since $0 \in \Gamma_i(x')$ exactly when i is active at x', the condition is equivalent to

$$\mathsf{C}_i(\mathsf{P}_i(h), \sigma) \subseteq X_K^- \cap R_{i,\sigma}.$$

Theorem 3.2 (Exact realization)**.** *There exists a tuple of local supervisors that exactly realizes K if and only if $\widetilde{K}$ is controllable and K is legitimately co-observable.*

Proof. Suppose g exactly realizes K. Every available uncontrollable extension of a closed-loop history is admitted, so controllability is necessary. Fix $x = (h, \sigma) \in X_K^-$. Exact realization requires some $i \in A_x$ to satisfy $g_i(\mathsf{P}_i(h), \sigma) = 0$. At every point $x' = (h', \sigma)$ in the same i-cell, the local output remains 0. Legitimacy forces $i \in A_{x'}$, and x' cannot be legal because exact realization requires every active supervisor to admit at a legal point. Hence the whole cell lies in $X_K^- \cap R_{i,\sigma}$.

Conversely, assume controllability and legitimate co-observability. Define the canonical policies by

$$g_i(y, \sigma) = 0 \quad \Longleftrightarrow \quad \varnothing \neq \mathsf{C}_i(y, \sigma) \subseteq X_K^- \cap R_{i,\sigma}. \tag{4}$$

They are K-legitimate. At a legal point, every active supervisor admits because its cell contains that legal point. At a forbidden point, legitimate co-observability supplies an active supervisor that vetoes. A word-length induction now gives $L_A(g/\mathcal{H}) = \widetilde{K}$: authority updates preserve the lift, controllability handles Σ_u, and the fused decision on each controllable extension agrees with membership in X_K^+. □

The theorem is representation independent. In particular, it does not require finite automata, natural projections, or an observable authority protocol. The canonical policy vetoes every information cell on which a veto is uniformly correct and authority-admissible. It is therefore a canonical *veto envelope*, not a uniquely permissive implementation: another exact tuple may admit on such a cell when a different supervisor already supplies the required veto.

Proposition 3.3 (Canonical veto envelope)**.** *Let g^{can} be the tuple in* (4)*, with local outputs ordered by $0 < 1$. If g is any exact K-legitimate tuple, then on every nonempty specification-level cell*

$$g_i^{\mathrm{can}}(y, \sigma) \leq g_i(y, \sigma).$$

Thus every veto issued by an exact legitimate tuple is also issued by the canonical tuple, although exact tuples need not be unique.

Proof. If $g_i(y, \sigma) = 0$, legitimacy puts the whole cell inside $R_{i,\sigma}$. The cell cannot contain a legal point, because an active veto there would contradict exact realization. Hence it lies in $X_K^- \cap R_{i,\sigma}$ and the canonical policy also vetoes. Its own exactness follows from Theorem 3.2. □

The closed-loop language is exact relative to the target, while the canonical implementation exercises every locally justified legitimate veto. In an explicit finite product, the condition is checked by grouping decision states by (i, y, σ); a failed group returns a decision- or authority-conflict witness, and successful groups define the canonical policies in polynomial time in the reachable product.

3.1. **Uniform-strategy embedding and residual content.** Legitimate co-observability can be embedded directly in the general language of observation-uniform strategies (Maubert and Pinchinat, 2014). For each supervisor i, regard $(h, \sigma) \in X_K$ as a local decision history with observation $(\mathsf{P}_i(h), \sigma)$, binary action set $\{0, 1\}$, and history-dependent feasible-action set $\Gamma_i((h, \sigma))$. A local supervisor is then an observation-uniform strategy, K-legitimacy is

pointwise feasibility of that strategy on X_K, and exact active-set realization imposes the additional DES outcome condition that every legal continuation is admitted while every forbidden continuation receives at least one active veto. Under this embedding, Definition 3.1 is precisely the cell-uniform feasibility condition for a veto witness, and Theorem 3.2 is its prefix-closed DES realization characterization together with uncontrollable- event closure.

The embedding is useful because it localizes, rather than erases, the paper's claim. At the exact-realization level, the result not supplied by the generic uniform-strategy principle is the aligned-witness irreducibility established below: the separate decision and authority marginals do not compose, and the smallest obstruction requires two active supervisors and three decision points. Nor does the embedding supply the Cai-style union-closed relative construction with its new coverage obligation, the whole-fiber restriction created by hidden authority schedules, or the coverage/self-pair pruning that makes finite synthesis exact.

3.2. **Aligned-witness irreducibility.** For $x = (h, \sigma) \in X_K^-$ and $i \in A_x$, put

$$\begin{aligned} D_i(x): &\quad \mathsf{C}_i(\mathsf{P}_i(h), \sigma) \subseteq X_K^-, \\ U_i(x): &\quad \mathsf{C}_i(\mathsf{P}_i(h), \sigma) \subseteq R_{i,\sigma}. \end{aligned}$$

Thus $D_i(x)$ says that i knows the required decision and $U_i(x)$ says that i knows its own authority. Legitimate co-observability is

$$\forall x \in X_K^- \; \exists i \in A_x : D_i(x) \wedge U_i(x).$$

Proposition 3.4 (Alignment is necessary)**.** *The two marginal conditions*

$$\forall x \in X_K^- \; \exists i \in A_x : D_i(x), \qquad \forall x \in X_K^- \; \exists i \in A_x : U_i(x)$$

do not imply legitimate co-observability.

Proof. Let a, b, c be distinct uncontrollable events and let σ be controllable. Take the prefix-tree plant

$$L(G) = \{\varepsilon, a, b, c, a\sigma, b\sigma, c\sigma\}$$

and the specification

$$K = \{\varepsilon, a, b, c, c\sigma\}.$$

Thus the σ-decision points $x_0 = (a, \sigma)$ and $x_1 = (b, \sigma)$ are forbidden, while $x_2 = (c, \sigma)$ is legal. Set

$$A_{x_0} = A_{x_2} = \{1, 2\}, \qquad A_{x_1} = \{2\}.$$

Let supervisor 1 merge a with b and supervisor 2 merge a with c; all other relevant observation classes are singletons. At x_0, supervisor 1 satisfies $D_1(x_0)$ but fails $U_1(x_0)$ because it is inactive at x_1. Supervisor 2 satisfies $U_2(x_0)$ but fails $D_2(x_0)$ because its cell contains the legal point x_2. At the other forbidden point x_1, supervisor 2 is a decision and authority witness through its singleton cell. Hence both marginal conditions hold everywhere they are required, but no legitimate witness exists at x_0. □

Corollary 3.5 (Minimal alignment obstruction)**.** *Any failure of legitimate co-observability at a forbidden point while both marginal conditions in Proposition 3.4 hold requires at least two supervisors active at that point and at least three decision points. The construction in Proposition 3.4 attains both bounds.*

Proof. With only one active supervisor, the decision and authority marginals must have the same witness, which is then legitimate. Now let x be an obstructed forbidden point.

A decision witness i that is not legitimate must fail $U_i(x)$, so its cell contains a distinct point x_1 at which i is inactive; $D_i(x)$ makes x_1 forbidden. An authority witness j that is not legitimate must fail $D_j(x)$, so its wholly authorized cell contains a distinct legal point x_2. The points x_1 and x_2 are distinct because one is forbidden and the other legal. Hence x, x_1, x_2 are three distinct decision points. □

The proposition and corollary are the non-tautological content of the aligned-witness condition. A generic uniformity requirement says that one local rule must be constant on its information cell; it does not imply that the separately satisfiable decision and authority marginals fail to compose, nor does it give the sharp two-supervisor, three-point lower bound. Thus the new condition is not classical co-observability followed by an unrelated authority check: the two facts must be aligned at the same local decision rule.

For comparison, define

$$W_{\mathrm{af}}(x) := \{i \in A_x : \mathsf{C}_i(\mathsf{P}_i(h), \sigma) \cap R_{i,\sigma} \subseteq X_K^-\}.$$

Active-filtered co-observability requires $W_{\mathrm{af}}(x) \neq \varnothing$ at every forbidden point and ignores a veto where its issuer is inactive. Call authority *witness aligned relative to* K when, for every $x = (h, \sigma) \in X_K^-$ with $W_{\mathrm{af}}(x) \neq \varnothing$, some $i \in W_{\mathrm{af}}(x)$ has its full cell contained in $R_{i,\sigma}$.

Remark 3.6 (Witness-alignment reformulation). A language is legitimately co-observable if and only if it is active-filtered co-observable and authority is witness aligned. This is a re-slicing of Definition 3.1, useful for comparison rather than an independent result: a legitimate witness is active-filtered and has a wholly authorized cell; conversely, witness alignment selects an active-filtered witness whose full cell is authorized, so its active-filtered condition makes the whole cell forbidden.

Proposition 3.7 (Strict gap from active filtering)**.** *Legitimate co-observability implies active-filtered co-observability, and the converse fails even with two supervisors, one controllable event, and two indistinguishable decision histories.*

Proof. The implication follows because a legitimate witness remains a witness after its cell is intersected with its authority region. For strictness, take $L(G) = \{\varepsilon, a, b, a\sigma, b\sigma\}$ and $K = \{\varepsilon, a, b\}$, with a, b uncontrollable and σ controllable. Let both supervisors merge a, b, but set the active sets at the two σ-decision points to $\{1\}$ and $\{2\}$, respectively. Under active filtering, each supervisor's uniform veto is used only where that supervisor is active, so both forbidden continuations are disabled. No supervisor is active throughout the common cell, hence no legitimate witness exists. □

Corollary 3.8 (Behavioral fusion does not determine issuer legitimacy)**.** *There exist two policy tuples with the same active-set closed-loop language, one K-legitimate and one not. Consequently, a correctness criterion that depends only on the fused behavior cannot characterize exact legitimate realization; it must additionally constrain inactive outputs or make the relevant authority status available to the local policy.*

Proof. Take $L(G) = K = \{\varepsilon, a, a\sigma\}$, with a uncontrollable, σ controllable, two supervisors, and $A_{(a,\sigma)} = \{1\}$. Let supervisor 1 admit at the unique σ-decision. If supervisor 2 also admits, the tuple is K-legitimate. Changing only supervisor 2's output at that local argument to veto leaves (2), and hence the closed-loop language, unchanged because supervisor 2

is inactive. The modified tuple is not legitimate because $0 \notin \Gamma_2(a, \sigma)$. Thus behavioral equivalence under active filtering does not determine issuer admissibility. □

Proposition 3.4, Corollary 3.5, Proposition 3.7, and Corollary 3.8 establish the irreducibility of Definition 3.1. A uniform local action must remain both correct and inside the same issuer's action domain throughout the information cell on which it is selected; changing how outputs are behaviorally combined does not by itself enforce that constraint.

3.3. **When authority observation removes the gap.** The obstruction disappears when each supervisor can infer its own authority on every tested cell.

Definition 3.9 (K-relative local authority observability)**.** Authority is *locally observable relative to* K when, for every i, y, σ, membership of i in the active set is constant on the decision cell:

$$\mathsf{C}_i(y, \sigma) \subseteq R_{i,\sigma} \quad \text{or} \quad \mathsf{C}_i(y, \sigma) \cap R_{i,\sigma} = \varnothing.$$

Corollary 3.10 (Authority-observability collapse)**.** *Under Definition 3.9, legitimate co-observability is equivalent to active-filtered co-observability.*

Proof. The forward implication is Proposition 3.7. Conversely, at a forbidden point let i be an active-filtered witness. Its cell contains the current active point, so authority homogeneity makes the whole cell a subset of $R_{i,\sigma}$. The active-filtered witness condition then makes the whole cell forbidden, and i is legitimate. □

The corollary is the model's principal positioning boundary. Standard information-set models make feasible actions constant across an information set, while stochastic-action-set models disclose the realized availability to the acting player. Under the natural analogy, both fall in the observable- authority regime: the issuer can condition its local decision on whether veto is available, and the authority conjunct becomes locally decidable. This is not merely a notational refinement, because revealing the authority coordinate can strictly split an information cell. Proposition 3.4 proves that without such a refinement, separate knowledge of the required decision and of authority does not recover a legitimate witness.

3.4. **Exact criterion for compiling out authority filtering.** Let $L_\forall(g/\mathcal{H})$ denote the closed language obtained by replacing (2) with fixed conjunction over all supervisors:

$$\bigwedge_{i \in \mathcal{I}} g_i(\mathsf{P}_i(h), \sigma) = 1. \tag{5}$$

This interface does not query the current active set. Say that g is *legally neutral outside the active set* when

$$(h, \sigma) \in X_K^+, \quad i \notin A_{(h,\sigma)} \quad \Longrightarrow \quad g_i(\mathsf{P}_i(h), \sigma) = 1. \tag{6}$$

Theorem 3.11 (Legal-point neutrality characterizes compile-out)**.** *Let g be an active-set realization of K. Then*

$$L_\forall(g/\mathcal{H}) = \widetilde{K}$$

if and only if g satisfies (6)*.*

Proof. For necessity, let $(h, \sigma) \in X_K^+$. Since $h\sigma \in L_\forall(g/\mathcal{H}) = \widetilde{K}$, every fixed-conjunction input is 1, including those of inactive supervisors.

Conversely, induct on word length. Authority updates preserve the lift, and uncontrollable extensions are retained because g already realizes $\widetilde{K}$ under active-set conjunction. At a legal controllable point, every active output admits and (6) makes every inactive output neutral, so fixed conjunction admits. At a forbidden point, active-set realization supplies an active veto, which also vetoes under fixed conjunction. □

Corollary 3.12 (Legitimate compile-time elimination)**.** *Let g be K-legitimate.*

(i) At every $x = (h, \sigma) \in X_K$,

$$\bigwedge_{i \in A_x} g_i(\mathsf{P}_i(h), \sigma) = \bigwedge_{i \in \mathcal{I}} g_i(\mathsf{P}_i(h), \sigma).$$

(ii) Active-set and fixed-conjunction realization are equivalent:

$$L_A(g/\mathcal{H}) = \widetilde{K} \quad \Longleftrightarrow \quad L_\forall(g/\mathcal{H}) = \widetilde{K}.$$

Proof. Legitimacy forces every inactive conjunct to equal 1 at each $x \in X_K$, giving (i). For (ii), part (i) makes active-set and fixed conjunction agree at every decision point in X_K. A word-length induction starting at $\varepsilon \in \widetilde{K}$, with authority updates preserving the lift and uncontrollable extensions handled identically, therefore transfers membership in both directions. The forward implication alternatively follows from Theorem 3.11 via (6); that theorem does not supply the reverse implication because it assumes active-set realization. □

Pure active-set realization may contain behaviorally invisible unauthorized vetoes; compile-out requires neutrality only at legal points, while full legitimacy excludes unauthorized vetoes at every reachable decision. The result assumes faithful output delivery and is not a fault-tolerance claim.

4. Relative legitimate co-observability

This section extends the fixed-ambient union-closure method of Cai–Zhang–Wonham relative co-observability (Cai et al., 2019). The inherited scaffold is candidate-independent consistency on ambient information cells. The new ingredients are issuer-safe filtering of those blocks and the separate coverage obligation (R2): after unsafe cells are removed from consistency, coverage prevents the resulting vacuity from making a nominally controllable event removable where no issuer can legitimately veto it.

This is not reproduced by authority-indexed event splitting. Such splitting reveals the authority coordinate to the local rule, treats copy assignment as fusion participation rather than an action-domain constraint, and does not couple the copies belonging to one operational fiber. A detailed comparison appears in Remark 4.10.

Legitimate co-observability is not union closed: legal continuations contributed by different sublanguages can destroy every candidate-level witness. Fix instead a nonempty prefix-closed ambient language $C \subseteq \mathsf{p}_\Sigma(\mathcal{H})$ and write $\widetilde{C} = \{h \in \mathcal{H} : \mathsf{p}_\Sigma(h) \in C\}$. For a prefix-closed candidate $M \subseteq C$, put

$$\begin{aligned} X_C &:= \{(h, \sigma) : h \in \widetilde{C},\ \sigma \in \Sigma_c,\ h\sigma \in \mathcal{H}\}, \\ X_M^- &:= \{(h, \sigma) : h \in \widetilde{M},\ \sigma \in \Sigma_c,\ h\sigma \in \mathcal{H} \setminus \widetilde{M}\}, \\ \mathsf{C}_i^C(y, \sigma) &:= \{(h, \sigma) \in X_C : \mathsf{P}_i(h) = y\}, \\ R_{i,\sigma}^C &:= \{(h, \sigma) \in X_C : i \in \mathsf{A}(c(h), \sigma)\}. \end{aligned}$$

An ambient cell is *authority safe* when it is nonempty and contained in $R^C_{i,\sigma}$. Equivalently,

$$\mathrm{SafeIssuers}_C(h,\sigma) := \{i \in \mathcal{I} : \mathsf{C}^C_i(\mathsf{P}_i(h),\sigma) \subseteq R^C_{i,\sigma}\}. \tag{7}$$

This set is candidate independent: candidate changes alter continuation truth values but not the cells on which consistency is tested.

Definition 4.1 (Relative legitimate co-observability)**.** A prefix-closed $M \subseteq C$ is *relatively legitimately co-observable with respect to* C when:

(R1) If two ambient σ-decision points lie in one authority-safe i-cell, then

$$h\sigma \in \widetilde{M} \quad \Longleftrightarrow \quad h'\sigma \in \widetilde{M}. \tag{8}$$

(R2) Every $(h,\sigma) \in X^-_M$ satisfies

$$\mathrm{SafeIssuers}_C(h,\sigma) \neq \varnothing. \tag{9}$$

The empty language is admitted as a bottom element.

Thus authority-safe ambient cells are fixed consistency blocks, while unsafe cells impose no equality. Guarding (R1) creates a vacuity hole wherever no safe issuer exists; (R2) is the necessary coverage repair.

Lemma 4.2 (Fixed authority-safe geometry)**.** *Fix i and σ, and write $X^\sigma_C := \{(h,\sigma) \in X_C\}$ for the ambient σ-decision points. Partition X^σ_C by retaining each authority-safe i-cell as one block and replacing every unsafe cell by singleton blocks. A candidate satisfies (R1) for i,σ exactly when continuation membership is constant on every block of this candidate-independent partition.*

Proof. (R1) imposes equality precisely within authority-safe cells. Singleton blocks from unsafe cells impose no additional condition, so block constancy is equivalent to (8). □

Theorem 4.3 (Implication and union closure)**.** *Every nonempty relatively legitimately co-observable language is legitimately co-observable. For fixed C, relatively legitimately co-observable sublanguages of C are closed under arbitrary unions.*

Proof. At $x = (h,\sigma) \in X^-_M$, (R2) supplies $i \in \mathrm{SafeIssuers}_C(h,\sigma)$. Regard M as the target K. Since $\widetilde{M} \subseteq \widetilde{C}$, the specification-level cell $\mathsf{C}_i(\mathsf{P}_i(h),\sigma)$ lies inside the corresponding authority-safe ambient cell $\mathsf{C}^C_i(\mathsf{P}_i(h),\sigma)$, so every point of the smaller cell authorizes i. If some (h',σ) in that smaller cell were legal for M, then (R1), applied to (h,σ) and (h',σ) in the ambient safe cell, would force $h\sigma \in \widetilde{M}$, a contradiction. Hence $\mathsf{C}_i(\mathsf{P}_i(h),\sigma) \subseteq X^-_M \cap R_{i,\sigma}$, and i is a legitimate witness.

For a nonempty union $M = \bigcup_\alpha M_\alpha$, any true continuation on a safe cell belongs to one M_α, whose (R1) makes the whole cell true. If a point is forbidden in the union, choose a member containing its prefix; it is forbidden there too and inherits that member's fixed ambient (R2) witness. The empty union is the admitted bottom. □

Corollary 4.4 (Supremal relative sublanguage)**.** *For prefix-closed $K_0 \subseteq C$, let*

$$\mathfrak{F}(K_0,C) := \{M \subseteq K_0 : M \text{ is prefix closed and controllable, and } M \text{ is relatively legitimately co-observable}\}.$$

Then

$$\mathrm{supRLCO}(K_0, C) := \bigcup_{M \in \mathfrak{F}(K_0, C)} M$$

is the unique supremal member and may be empty.

Proof. Prefix closure and controllability are preserved by arbitrary unions, and Theorem 4.3 preserves relative legitimate co-observability. Hence $\mathfrak{F}(K_0, C)$ is union closed and its union is its unique supremal member. □

Proposition 4.5 (Effective controllability)**.** *Let $M \subseteq C$ be prefix closed. Then ordinary controllability together with (R2) is equivalent to the following closure conditions for every $h \in \widetilde{M}$:*

(i) every available extension by $\Sigma_u \cup \Delta$ remains in $\widetilde{M}$;
(ii) every available $\sigma \in \Sigma_c$ with $\mathrm{SafeIssuers}_C(h, \sigma) = \varnothing$ satisfies $h\sigma \in \widetilde{M}$.

Proof. Clause (i) combines ordinary controllability with the automatic Δ-closure of an operational lift. Clause (ii) is the contrapositive of (R2): a controllable continuation may be removed only where an ambient-safe issuer can supply a legitimate veto. □

Definition 4.6 (C-effective controllability)**.** A prefix-closed $M \subseteq C$ is *C-effectively controllable* when it satisfies the equivalent conditions of Proposition 4.5. The empty language is vacuously C-effectively controllable.

Definition 4.7 (Effective uncontrollability under authority)**.** An available $\sigma \in \Sigma_c$ is *C-effectively uncontrollable at $h \in \widetilde{C}$* when $\mathrm{SafeIssuers}_C(h, \sigma) = \varnothing$. Every C-effectively controllable candidate M with $h \in \widetilde{M}$ must then satisfy $h\sigma \in \widetilde{M}$. This differs from the excluded empty-active-set case: here $\mathsf{A}(c(h), \sigma)$ is nonempty, but no issuer is uniformly admissible on its ambient information cell.

The operative control status is therefore induced jointly by nominal controllability, history, authority, and observation. In particular, an event can move between ordinary controllable behavior and forced behavior without changing the nominal alphabet partition: the change occurs when observation and authority cease to supply any cell-uniform admissible issuer. This is a closure consequence of issuer legitimacy, not a reclassification of the plant event.

Proposition 4.8 (Independent roles of consistency and coverage)**.** *Neither relative clause replaces the other.*

(i) (R1) without (R2) may admit a forbidden point with no legitimate witness.
(ii) Languages satisfying (R2) and legitimate co-observability may have a union that loses every witness unless (R1) is imposed against the fixed ambient cells.

Proof. For (i), let $\Delta = \varnothing$, $\alpha, \beta \in \Sigma_u$, $\sigma \in \Sigma_c$, and

$$L(G) = \{\varepsilon, \alpha, \beta, \alpha\sigma, \beta\sigma\}, \qquad C = L(G), \qquad M = \{\varepsilon, \alpha, \beta\}.$$

Use configurations c_0, c_α, c_β with $\mathsf{A}(c_\alpha, \sigma) = \{1\}$, $\mathsf{A}(c_\beta, \sigma) = \{2\}$, and let both supervisors merge α, β. Their ambient σ-decision cell is contained in neither authority region, so no safe issuer exists at either forbidden point. Hence (R1) is vacuous, (R2) fails, and no legitimate witness exists; M is controllable because both uncontrollable branches are retained.

For (ii), use the same tree with all events controllable, both supervisors always authorized, and σ-decision observations merging α, β (root branch cells remain separated by the event argument). Let

$$M_1 = \{\varepsilon, \alpha, \alpha\sigma\}, \qquad M_2 = \{\varepsilon, \beta\}.$$

Both are prefix closed, controllable, satisfy (R2), and are individually legitimately co-observable: each reaches only one of the two σ-decision prefixes. Their union contains the legal point (α, σ) and the forbidden point (β, σ) in the same candidate cell, so no legitimate witness remains. Thus coverage survives while fixed ambient consistency is what fails under union. □

In the present prefix-closed notation, classical conjunctive co-observability for fixed nonempty event assignments I_σ requires

$$\forall (h, \sigma) \in X_K^- \; \exists i \in I_\sigma : \mathsf{C}_i(\mathsf{P}_i(h), \sigma) \subseteq X_K^-. \tag{10}$$

For a prefix-closed candidate $M \subseteq C$, the corresponding relative condition requires, for every $i \in I_\sigma$,

$$\begin{aligned} (h, \sigma), (h', \sigma) \in X_C, \quad \mathsf{P}_i(h) = \mathsf{P}_i(h') \\ \Longrightarrow \quad (h\sigma \in \widetilde{M} \iff h'\sigma \in \widetilde{M}). \end{aligned} \tag{11}$$

Theorem 4.9 (Static-authority regression)**.** *If, for each $\sigma \in \Sigma_c$, $\mathsf{A}(c(h), \sigma) = I_\sigma$ for a fixed nonempty set I_σ at every $(h, \sigma) \in X_K$, then legitimate co-observability is exactly* (10)*. If instead the same equality holds at every $(h, \sigma) \in X_C$, the relative condition is exactly* (11)*.*

Proof. For $i \in I_\sigma$, every nonempty tested cell is contained in the authority region; for $i \notin I_\sigma$, no tested cell can support a veto. Substitution in Definition 3.1 therefore makes its authority conjunct automatic for precisely the assigned supervisors and gives (10). The same substitution in Definition 4.1 makes every assigned ambient cell authority safe and every unassigned cell irrelevant, giving (11). □

Remark 4.10 (Why naive event splitting is not semantics preserving). A natural reduction replaces each σ by copies σ_A indexed by the active sets in the range of $\mathsf{A}(\cdot, \sigma)$, assigns the fixed set $I_{\sigma_A} = A$, and applies classical co-observability. Without additional constraints, this changes the problem in three ways.

First, the copy label exposes the authority coordinate. A local policy can condition on σ_A, so each refined decision cell contains only points with active set A. This is exactly the observation refinement under which Corollary 3.10 removes the aligned-witness obstruction. The split model therefore assumes away the case being studied.

Second, fixed assignment to a copy models which outputs participate in fusion, not whether every issued output lies in its issuer's action domain. A supervisor outside A simply contributes no conjunct at σ_A; classical co-observability has no reason to constrain the veto it would emit there. The result therefore has active-filtered semantics, which Proposition 3.7 shows can be strictly more permissive than legitimate realization.

Third, split-alphabet synthesis does not itself preserve operational fibers. Hidden histories with one operational continuation may carry different copy indices, and a classical computation may retain some copies while deleting others. A semantics-preserving reduction must add cross-copy constraints that retain or remove the entire fiber and enforce the issuer

action domains. Those constraints reconstruct the two obligations represented directly here. Theorem 4.9 identifies the constant-authority case in which no such enrichment is needed.

5. Finite schedule-class synthesis

The exact-realization and abstract relative results of Sections 2–4 depend only on the observation maps and do not require finite representations. The synthesis results below instead assume the compatible finite monitors of Assumption 5.3. Observation maps that do not admit such monitors, including genuinely unbounded perfect-recall maps, therefore lie outside the finite-synthesis theorem.

The supremal language in Corollary 4.4 is abstract. The centerpiece of the finite construction is the class/word boundary: whole-fiber schedule classes are necessary, but a retained class transition still does not determine membership for every word reaching its source. Coverage and semantic self-pairs repair exactly this leakage and make the pruning operators exact.

5.1. Whole fibers and schedule classes.

For $L \subseteq \mathcal{H}$, define

$$\mathrm{Sat}_\Sigma(L) := \mathcal{H} \cap \mathsf{p}_\Sigma^{-1}(\mathsf{p}_\Sigma(L)), \qquad \mathcal{F}_\Sigma(s) := \{h \in \mathcal{H} : \mathsf{p}_\Sigma(h) = s\}. \tag{12}$$

Proposition 5.1 (Operational lifts and whole fibers)**.** *For $L \subseteq \mathcal{H}$, the following are equivalent:*

(i) $L = \widetilde{M}$ *for some* $M \subseteq \mathsf{p}_\Sigma(\mathcal{H})$*;*
(ii) $L = \mathrm{Sat}_\Sigma(L)$*;*
(iii) every fiber $\mathcal{F}_\Sigma(s)$ *is either wholly contained in or disjoint from* L*.*

When they hold, the representing language is unique and equals $M = \mathsf{p}_\Sigma(L)$*. Hence any operational synthesis must retain or remove all hidden authority schedules of one continuation together.*

Proof. For every L, $\mathrm{Sat}_\Sigma(L) = \widetilde{\mathsf{p}_\Sigma(L)}$. Thus a saturation fixed point is a lift, every lift is saturated, and the fiber statement is the pointwise form of the same equality. Uniqueness follows by projection. □

Definition 5.2 (Plant-total finite monitor)**.** Let $f : \mathcal{H} \to V$ have finite range. A deterministic Moore machine $\mathcal{M} = (Q, q_0, \eta, \lambda)$ over E is a *plant-total monitor* of f when $\eta(q_0, h)$ is defined for every $h \in \mathcal{H}$ and

$$\lambda(\eta(q_0, h)) = f(h) \qquad (h \in \mathcal{H}).$$

An authority monitor is *compatible* with (c, A) when it monitors the finite-valued map

$$h \longmapsto \big(\mathsf{A}(c(h), \sigma)\big)_{\sigma \in \Sigma_c}.$$

A finite observer for supervisor i is compatible when it monitors P_i. Thus synchronization records the prescribed authority and observation values without deleting any plant history.

Assumption 5.3 (Finite explicit representation)**.** The plant is a deterministic finite automaton over E; C and K_0 have complete deterministic recognizers over Σ whose lifts self-loop on Δ; authority is supplied by a compatible plant-total finite monitor; and every P_i is supplied by a compatible plant-total finite-output observer in the sense of Definition 5.2.

Let Z be their reachable synchronous product. Then $L(\mathsf{Z}) = \mathcal{H}$. Write $z(h)$ for its state and $\omega_i(z)$ for observer i's output. Product states record plant availability, the $\widetilde{C}$ and $\widetilde{K}_0$ flags, authority, and all observer outputs. A product state is an ambient σ-decision state when its $\widetilde{C}$-flag holds and σ is plant enabled. Define

$$\mathrm{safe}_{i,\sigma}(y) = 1 \iff \begin{array}{l}\text{some ambient } \sigma\text{-decision state has output } y, \text{ and} \\ \text{every such state authorizes supervisor } i.\end{array} \tag{13}$$

Then $\mathrm{SafeIssuers}_C(h,\sigma) = \{i : \mathrm{safe}_{i,\sigma}(\mathsf{P}_i(h)) = 1\}$. Plant totality ensures that synchronization monitors, rather than removes, plant histories.

Proposition 5.4 (Verification complexity)**.** *Under an analogous componentwise finite representation of the plant, the realization target K, the authority monitor, and the local observers, deciding legitimate co-observability is PSPACE-complete. Once the reachable synchronous product is materialized, the same test and construction of the canonical policies are polynomial in that product.*

Proof. For PSPACE-hardness, reduce from DFA-intersection emptiness, which is PSPACE-complete (Kozen, 1977). Let deterministic automata $\mathcal{A}_1, \ldots, \mathcal{A}_n$ share alphabet Λ, introduce branch letters $\beta_0, \ldots, \beta_n$ and one controllable event σ, and take $\Delta = \varnothing$. The prefix-closed plant language is

$$\mathcal{H} = \{\varepsilon\} \cup \bigcup_{j=0}^{n} \beta_j \Lambda^* \cup \beta_0 \Lambda^* \sigma \cup \bigcup_{i=1}^{n} \beta_i L(\mathcal{A}_i)\sigma,$$

and the prefix-closed target deletes exactly the β_0 service continuations:

$$K = \{\varepsilon\} \cup \bigcup_{j=0}^{n} \beta_j \Lambda^* \cup \bigcup_{i=1}^{n} \beta_i L(\mathcal{A}_i)\sigma.$$

All branch letters and letters of Λ are uncontrollable. Give σ the static issuer set $I_\sigma = \{1, \ldots, n\}$. Supervisor i uses a finite Moore observer that tracks the current state of $\mathcal{A}_i$, gives the branches β_0 and β_i the same branch tag, and gives every other branch a distinct tag. Thus, at the forbidden point $(\beta_0 w, \sigma)$, its cell contains a legal decision point on branch β_i exactly when the current $\mathcal{A}_i$-state is accepting, equivalently when $w \in L(\mathcal{A}_i)$. Hence supervisor i is a decision witness exactly when $w \notin L(\mathcal{A}_i)$. By Theorem 4.9, K is legitimately co-observable exactly when every word is rejected by at least one input automaton, equivalently when $\bigcap_i L(\mathcal{A}_i) = \varnothing$. The plant, target, finite observers, and static authority have total size polynomial in the input automata.

For membership, a nondeterministic polynomial-space procedure searches on-the-fly for failure: guess a reachable forbidden decision point x, and for every $i \in A_x$ guess a reachable decision point in the same i-cell that is legal or at which i is inactive. Reachability and observation equality require storing only the current states of the component automata and polynomially many counters. The issuers in A_x are processed sequentially: for each issuer the procedure guesses and verifies one spoiling point, then reuses the same polynomial workspace for the next issuer. Thus failure lies in NPSPACE, hence in PSPACE, and PSPACE is closed under complement. If the reachable product is already explicit, grouping its decision states by (i, y, σ) performs the same test directly in polynomial time. □

The proposition distinguishes verification from the synthesis bound below. Polynomial time after materialization means polynomial in the explicit reachable product, which may

be exponentially larger than the separate plant, recognizer, authority-monitor, and observer descriptions.

For $S \subseteq Q_{\mathsf{Z}}$ let $\mathrm{cl}_{\Delta}(S)$ be its finite Δ-closure and let $\mathrm{Post}_e(S)$ be its set of e-successors.

Definition 5.5 (Schedule-class system)**.** With initial product state z_0, set

$$D(\varepsilon) = \mathrm{cl}_{\Delta}(\{z_0\}), \qquad D(se) = \mathrm{cl}_{\Delta}(\mathrm{Post}_e(D(s)))$$

whenever $se \in \mathsf{p}_{\Sigma}(\mathcal{H})$. The reachable subsets and their induced deterministic Σ-transitions form the schedule-class system $\mathcal{D}$.

The system $\mathcal{D}$ is the standard observer, or subset, automaton of Z under the projection that hides Δ (Cassandras and Lafortune, 2021); $D(s)$ is its current-state estimate after observing s. The synthesis result begins where this standard determinization is insufficient: candidate-language membership need not transfer from a class to every word reaching it, and coverage together with semantic self-pairs restores exact pruning (Lemma 5.12 and Example 5.13).

Lemma 5.6 (Schedule-class semantics)**.** *For every $s \in \mathsf{p}_{\Sigma}(\mathcal{H})$,*

$$D(s) = \{z(h) : h \in \mathcal{H},\ \mathsf{p}_{\Sigma}(h) = s\}. \tag{14}$$

Consequently $\mathcal{D}$ reunites exactly the product-state images of each operational fiber and has at most $2^{|Q_{\mathsf{Z}}|}$ states.

Proof. Induct on $|s|$. Histories with empty operational projection are precisely the finite Δ-paths from z_0. Every history projecting to se decomposes as $h_0 e d$ with $\mathsf{p}_{\Sigma}(h_0) = s$ and $d \in \Delta^*$; the induction hypothesis, the e-post, and final closure give the forward inclusion, while the same data concatenate to a witness for the reverse inclusion. □

Define $s \equiv_D t$ when $D(s) = D(t)$.

Proposition 5.7 (Deterministic fiber-image quotient)**.** *The relation $\equiv_D$ is a right congruence relative to the prefix-closed language $\mathsf{p}_{\Sigma}(\mathcal{H})$, and the reachable quotient by $\equiv_D$ is isomorphic to $\mathcal{D}$ under $[s] \mapsto D(s)$.*

Proof. If $D(s) = D(t)$, then for every $e \in \Sigma$,

$$\mathrm{cl}_{\Delta}(\mathrm{Post}_e(D(s))) = \mathrm{cl}_{\Delta}(\mathrm{Post}_e(D(t))).$$

The two continuations are therefore simultaneously available and, when available, satisfy $D(se) = D(te)$. The quotient map is well defined, bijective on reachable classes, and transition preserving. □

Theorem 5.8 (Current-estimate width)**.** *Let*

$$z := |Q_{\mathsf{Z}}|, \qquad w := \max_{s \in \mathsf{p}_{\Sigma}(\mathcal{H})} |D(s)|.$$

Thus w is the maximum current-state-estimate cardinality, also called the product-state ambiguity width in this paper. The general subset bound is $|Q_{\mathcal{D}}| \le 2^z$. More sharply,

$$|Q_{\mathcal{D}}| \le \sum_{k=1}^{w} \binom{z}{k}, \qquad |\mathrm{Trans}(\mathcal{D})| \le |\Sigma| \sum_{k=1}^{w} \binom{z}{k}. \tag{15}$$

Thus construction is exponential in the worst case, but takes $z^{O(w)}$ time up to polynomial factors in the remaining explicit input when estimate width is bounded. This is XP, not an

FPT claim. When $w = 1$, all histories with one operational projection reach one product state.

Proof. Every reachable class is a nonempty subset of Q_{Z} of size at most w, which gives the state bound. Determinism gives at most one outgoing edge per operational event and class. Each successor is obtained by one post operation and a finite Δ-closure, both polynomial in the explicit product. □

Corollary 5.9 (Bounded hidden-mode width)**.** *Suppose the reachable product state admits a factorization $z(h) = (q_{\mathrm{op}}(\mathsf{p}_\Sigma(h)), q_{\mathrm{hid}}(h))$, where the first coordinate is determined by the operational word and the hidden coordinate ranges over a set of cardinality at most b, independent of the plant size. Then $w \le b$. Consequently, for fixed b, schedule-class construction and synthesis are polynomial in the explicit product size; see Corollary 5.23 below for the full bound. The cyclic family in Section 6 has $b = 2$.*

Proof. For a fixed operational word s, every state in $D(s)$ has the same first coordinate $q_{\mathrm{op}}(s)$ and differs only in the hidden coordinate. Hence $|D(s)| \le b$ for every s, so $w \le b$; the complexity claim follows from Corollary 5.23 below. □

Corollary 5.10 (Operational determinacy)**.** *The condition $w = 1$ is equivalent to*

$$\mathsf{p}_\Sigma(h) = \mathsf{p}_\Sigma(h') \Longrightarrow z(h) = z(h') \qquad (h, h' \in \mathcal{H}).$$

In that case every schedule class can be identified with its unique product state, although distinct operational words may still reach the same class.

Proof. By Lemma 5.6, $D(s)$ is exactly the set of product states reached by histories projecting to s; it is a singleton for every s exactly under the displayed implication. □

When $w > 1$, product-state pruning can be unsound. Suppose two hidden authority schedules of one prefix s reach $z_1, z_2 \in D(s)$, and an available σ-continuation is forced at z_1 but not at z_2. Retaining the transition from z_1 while deleting the one from z_2 keeps only part of the fiber of $s\sigma$, so the result is not an operational lift. Schedule-class transitions couple the fiber. They nevertheless introduce a second issue: different operational words can reach one class through paths with different retained-edge status.

All states of $D(s)$ share the C- and K_0-flags; write them $\chi_C(D)$ and $\chi_{K_0}(D)$. For $u \in \Sigma_u$ define

$$\phi_u(D) := [\exists z \in D : \ u \text{ is available at } z],$$

and for $\sigma \in \Sigma_c$ define

$$\mathrm{NoSafeIssuer}(D, \sigma) := [\exists z \in D : \ \sigma \text{ is available at } z \text{ and } \mathrm{safe}_{i,\sigma}(\omega_i(z)) = 0 \text{ for all } i]. \quad (16)$$

The existential quantifier is forced by operational membership: one hidden schedule that requires an extension requires the operational continuation.

For i and σ, define the unordered relation

$$(D, D') \in \mathcal{P}_{i,\sigma} \iff \begin{array}{l}\text{some ambient } \sigma\text{-decision states } z \in D, z' \in D' \\ \text{share one } i\text{-output that is safe for } i.\end{array} \quad (17)$$

Self-pairs (D, D) are retained. They may constrain distinct words reaching the same class and are therefore semantic equalities, not disposable graph loops.

Lemma 5.11 (Finite class predicates)**.** *For prefix-closed $M \subseteq C$:*

(i) M is C-effectively controllable if and only if, for every $s \in M$,

$$\phi_u(D(s)) = 1 \Rightarrow su \in M \quad (u \in \Sigma_u), \qquad \mathrm{NoSafeIssuer}(D(s), \sigma) = 1 \Rightarrow s\sigma \in M \quad (\sigma \in \Sigma_c).$$

(ii) M satisfies (R1) for supervisor i if and only if, for every $\sigma \in \Sigma_c$ and all ambient words $s, s' \in C$,

$$(D(s), D(s')) \in \mathcal{P}_{i,\sigma} \quad \Longrightarrow \quad (s\sigma \in M \Longleftrightarrow s'\sigma \in M).$$

Proof. Part (i) translates Definition 4.6 through Lemma 5.6: availability after any schedule of s is represented by the corresponding class predicate. The existential test is sound at the operational-language level precisely because $\widetilde{M}$ is a whole-fiber lift by Proposition 5.1: if one hidden schedule forces an available extension, effective controllability forces the common operational continuation into M, and the lift then retains every plant history with that continuation.

For (ii), assume (R1) and let $(D(s), D(s')) \in \mathcal{P}_{i,\sigma}$. Choose witnessing ambient decision states $z \in D(s)$ and $z' \in D(s')$ with a common safe output y. Lemma 5.6 realizes them by histories h, h' with projections s, s'. They lie in one authority-safe ambient i-cell, so (R1) gives $h\sigma \in \widetilde{M} \Longleftrightarrow h'\sigma \in \widetilde{M}$, which is the displayed word equality. Conversely, two histories in one safe cell yield a relation pair between their classes, and the class condition gives (R1). The argument includes self-pairs: one class may be reached by distinct words whose continuations must agree. The event need not be enabled at every state of a class; the relation provides the required ambient decision-state witnesses. □

5.2. **The class/word boundary and exact pruning.** A nonbottom $\mathcal{D}$-subautomaton is a transition set T with language

$$M_T := \{s \in \mathsf{p}_\Sigma(\mathcal{H}) : \text{every edge on the unique } \mathcal{D}\text{-path labelled } s \text{ lies in } T\}. \tag{18}$$

Let $M_\perp = \varnothing$ and, when $K_0 \neq \varnothing$, initialize with

$$T_0 := \{D \xrightarrow{e} D' : \chi_{K_0}(D) = \chi_{K_0}(D') = 1\}, \qquad M_{T_0} = K_0.$$

A class is *covered* by T when every path from the initial class to it uses only edges of T. Coverage is computable by two-layer reachability, with the second layer recording whether a path has crossed an edge outside T.

Lemma 5.12 (Coverage–transition invariant)**.** *Let T be a nonbottom transition set and let $D \xrightarrow{\sigma} D'$ be a schedule-class transition.*

(i) If it is absent from T, then $s\sigma \notin M_T$ for every word s reaching D.
(ii) If it is present and D is covered, then $s\sigma \in M_T$ for every word s reaching D.
(iii) If it is present and D is uncovered, at least one word reaching D has $s\sigma \notin M_T$, although another may have its continuation in M_T.

Proof. Membership in M_T requires every path edge to lie in T. An absent final edge proves (i); coverage plus a present final edge proves (ii). If D is uncovered, some word reaches it through an earlier missing edge, and appending a present edge does not repair that leak, proving (iii). □

Clause (iii) is the failed transfer from class states to words. It is also why a self-pair can be nontrivial: two words reaching one class can have different membership values despite sharing the same outgoing transition.

Example 5.13 (Coverage and a semantic self-pair are both necessary)**.** Let $\Delta = \varnothing$, let all events a, b, σ be controllable, and let a deterministic plant have transitions

$$q_0 \xrightarrow{a} q_1, \qquad q_0 \xrightarrow{b} q_1, \qquad q_1 \xrightarrow{\sigma} q_2.$$

Take $C = K_0 = L(G)$, one always-authorized supervisor, and one observer output at q_1. The product therefore has $D(a) = D(b) =: D$ and $(D, D) \in \mathcal{P}_{1,\sigma}$. Let

$$T = T_0 \setminus \{D(\varepsilon) \xrightarrow{b} D\}.$$

The σ-edge out of D is present, but $M_T = \{\varepsilon, a, a\sigma\}$: the word b reaches the same class through the missing b-edge, so $b\sigma \notin M_T$. Hence a class-only test of the present σ-edge would incorrectly treat the continuation as uniformly true. The self-pair records the required equality between the two ambient words, while coverage detects that D is uncovered. Guarded pruning therefore deletes $D \xrightarrow{\sigma} D(a\sigma)$ and returns $\mathsf{RO}_{1,C}(M_T) = \{\varepsilon, a\}$. Because $\Delta = \varnothing$ and every synchronized component is deterministic, each schedule class is a singleton and this example has $w = 1$. The self-pair is therefore load bearing inside the smallest-width tractable regime: it is caused by distinct words reaching one class through different retained paths, not by hidden-state multiplicity within the class.

For nonempty prefix-closed $L \subseteq M_T$, let T_L be the edges traversed by words of L and put $L^{\mathrm{s}} := M_{T_L}$.

Lemma 5.14 (Saturation and subautomaton representation)**.** *Let $T \subseteq T_0$ be nonbottom and let $\varnothing \neq L \subseteq M_T$ be prefix closed. Then $L \subseteq L^{\mathrm{s}} \subseteq M_T \subseteq K_0 \subseteq C$, and L^{s} preserves from L each of: (R1) for a fixed supervisor, C-effective controllability, and membership in $\mathfrak{F}(K_0, C)$. For prefix-closed $M \subseteq K_0$, define*

$$\mathsf{EC}_C(M) := \bigcup\{M' \subseteq M : M' \text{ is prefix closed and } C\text{-effectively controllable}\},$$
$$\mathsf{RO}_{i,C}(M) := \bigcup\{M' \subseteq M : M' \text{ is prefix closed and satisfies (R1) for } i\}.$$

Both operators are well defined, monotone, and deflationary. Every nonempty value $\mathsf{EC}_C(M_T)$ or $\mathsf{RO}_{i,C}(M_T)$ is represented by the transitions traversed by its words.

Proof. The inclusions follow from $T_L \subseteq T \subseteq T_0$. Each edge retained in T_L is witnessed by a word of L, and every class on a word of L^{s} is reached by a word of L.

For (R1), suppose $(D(s), D(s')) \in \mathcal{P}_{i,\sigma}$ and $s\sigma \in L^{\mathrm{s}}$. Some $w \in L$ reaches $D(s)$ with $w\sigma \in L$; the class-predicate form of (R1) gives $s'\sigma \in L \subseteq L^{\mathrm{s}}$, and symmetry gives the converse. For effective controllability, choose $w \in L$ reaching the class of $s \in L^{\mathrm{s}}$. Since ϕ_u and NoSafeIssuer are class predicates, any event forced from s is forced from w; its continuation in L places the same edge in T_L. Prefix closure is automatic, proving the three preservation claims.

Both defining families contain the empty language and are closed under arbitrary unions. For (R1), a continuation true in a union is true in one member, whose constancy on the fixed ambient block makes it true throughout that block. For C-effective controllability, the uncontrollable, Δ-extension, and no-safe-issuer forced-extension clauses are each preserved by union. Hence the displayed unions are their unique supremal members, and the operators are monotone and deflationary.

Now let $S = \mathsf{RO}_{i,C}(M_T) \neq \varnothing$. The preceding preservation result gives $S \subseteq S^{\mathrm{s}} \subseteq M_T$, with S^{s} still satisfying (R1). Supremality gives $S^{\mathrm{s}} \subseteq S$, hence $S = S^{\mathrm{s}}$. The same argument with effective controllability proves the claim for EC_C. □

For each i, σ, let $\equiv_{i,\sigma}$ be the reflexive, symmetric, transitive closure of $\mathcal{P}_{i,\sigma}$ on classes occurring in a pair; a component is *active* when it contains a pair, including a self-pair.

Lemma 5.15 (Safe-component consistency)**.** *A candidate satisfies (R1) for i and σ exactly when the truth value of $s\sigma \in M$ is constant over all ambient words whose classes lie in each active $\equiv_{i,\sigma}$-component and at which σ is available.*

Proof. Assume first that (R1) holds. By Lemma 5.11(ii), every generating pair $(D, D') \in \mathcal{P}_{i,\sigma}$ imposes equality of $s\sigma \in M$ and $s'\sigma \in M$ for all ambient words reaching D and D', respectively. Equality is transitive, so it propagates along every symmetric–transitive chain of generating pairs. If (D, D) is a retained self-pair, the same statement applies to any two ambient words reaching D; the loop therefore records a word-level equality rather than a vacuous class-level one. Hence the truth value is constant over every active $\equiv_{i,\sigma}$-component at which σ is available.

Conversely, constancy on every active component gives equality on each generating pair of $\mathcal{P}_{i,\sigma}$. Applying Lemma 5.11(ii) in the reverse direction yields (R1). □

Let $D_0 := D(\varepsilon)$ and let $\tau(D, e)$ denote the unique schedule-class edge $D \xrightarrow{e} D'$ when it exists; write $\mathrm{tgt}(\tau(D, e)) = D'$ for its target. For $T \subseteq T_0$, define

$$\mathrm{Cov}(T) := \{D : \text{every path in } \mathcal{D} \text{ from } D_0 \text{ to } D \text{ uses only edges of } T\}.$$

For a class D, put

$$\mathrm{Force}(D) := \{u \in \Sigma_u : \phi_u(D) = 1\} \cup \{\sigma \in \Sigma_c : \mathrm{NoSafeIssuer}(D, \sigma) = 1\}.$$

Definition 5.16 (Exact pruning operators)**.** For a nonbottom $T \subseteq T_0$, let $\mathrm{Bad}(T)$ be the least subset of $Q_{\mathcal{D}}$ satisfying

$$\big[\exists e \in \mathrm{Force}(D) : \tau(D, e) \notin T \text{ or } \mathrm{tgt}(\tau(D, e)) \in \mathrm{Bad}(T)\big] \quad \Longrightarrow \quad D \in \mathrm{Bad}(T). \tag{19}$$

The effective-controllability operator is

$$\mathcal{E}(T) := \begin{cases} \bot, & D_0 \in \mathrm{Bad}(T), \\ \{D \xrightarrow{e} D' \in T : D, D' \notin \mathrm{Bad}(T)\}, & D_0 \notin \mathrm{Bad}(T). \end{cases} \tag{20}$$

Fix i. An active $\equiv_{i,\sigma}$-component H is *T-defective* when either $H \not\subseteq \mathrm{Cov}(T)$ or $\tau(D, \sigma) \notin T$ for some $D \in H$. Define

$$\mathcal{G}_i(T) := T \setminus \{D \xrightarrow{\sigma} D' \in T : D \text{ lies in a } T\text{-defective active } \equiv_{i,\sigma}\text{-component}\}. \tag{21}$$

Because $\mathcal{G}_i$ is deflationary on a finite edge set, its iteration stabilizes; write $\mathcal{G}_i^{\infty}(T)$ for that stable value.

Lemma 5.17 (Deletion and survival invariant)**.** *Let $T \subseteq T_0$ be nonbottom.*

(i) If $\varnothing \neq L \subseteq M_T$ is prefix closed and C-effectively controllable, no word of L reaches a class in $\mathrm{Bad}(T)$*. Consequently $D_0 \notin \mathrm{Bad}(T)$ and $T_L \subseteq \mathcal{E}(T)$.*

(ii) If $\varnothing \neq L \subseteq M_T$ is prefix closed and satisfies (R1) for supervisor i, then $T_L \subseteq \mathcal{G}_i(T)$, and hence $T_L \subseteq \mathcal{G}_i^{\infty}(T)$.

Proof. For (i), expose the least fixed point in (19) by ranks. A word of L cannot reach a rank-zero bad class: the witnessing forced event must remain in L by effective controllability, while its edge is absent from T, which would place the continuation outside M_T. Inductively, if a class becomes bad because a forced edge enters a lower-rank bad class, effective controllability would make a word of L reach that lower-rank class. Thus no word of L reaches $\mathrm{Bad}(T)$. In particular $\varepsilon \in L$ excludes $D_0 \in \mathrm{Bad}(T)$, and every edge traversed by L has nonbad endpoints.

For (ii), first consider one application of $\mathcal{G}_i$. Suppose it deletes σ-edges from a T-defective active component H. Choose $D_* \in H$ witnessing defectiveness, so that either $D_* \notin \mathrm{Cov}(T)$ or $\tau(D_*, \sigma) \notin T$. Because every class of H occurs in a relation pair, D_* contains an ambient σ-decision-state witness; hence the class transition $\tau(D_*, \sigma)$ exists.

If $\tau(D_*, \sigma) \notin T$, Lemma 5.12(i) makes every σ-continuation of every ambient word reaching D_* false in M_T. Otherwise the edge is present and defectiveness gives $D_* \notin \mathrm{Cov}(T)$. Lemma 5.12(iii) then supplies an ambient word w_* reaching D_* with $w_*\sigma \notin M_T$. The continuation is plant available: the present class edge, together with the right-congruence property of Proposition 5.7, gives $w_*\sigma \in \mathsf{p}_\Sigma(\mathcal{H})$.

Now let $L \subseteq M_T$ satisfy (R1) for i. In either case at least one ambient σ-continuation represented in H is false in M_T, and hence false in L. Every class of H occurs in a safe pair, so its ambient C-flag is true. Lemma 5.15 makes continuation membership constant over all ambient words represented in H at which σ is available. Constancy together with the one false continuation therefore makes every such continuation false in L. No deleted σ-edge can belong to T_L, and thus $T_L \subseteq \mathcal{G}_i(T)$.

For the iteration, set $T^{(0)} = T$ and $T^{(k+1)} = \mathcal{G}_i(T^{(k)})$. Inductively assume $T_L \subseteq T^{(k)}$. Then $L \subseteq M_{T_L} \subseteq M_{T^{(k)}}$, while L still satisfies (R1), so the one-application argument with $T^{(k)}$ in place of T gives $T_L \subseteq T^{(k+1)}$. Hence T_L survives every iterate and $T_L \subseteq \mathcal{G}_i^{\infty}(T)$. □

Algorithm 5.18 (Schedule-class fixed point)**.** Fix any order of the supervisors. Start with $T = T_0$; return $\bot$ immediately when $K_0 = \varnothing$. Repeat the following outer pass:

1. Set $T_{\mathrm{old}} = T$ and replace T by $\mathcal{E}(T)$. If the result is $\bot$, return $\bot$.
2. In the fixed supervisor order, replace T by $\mathcal{G}_i^{\infty}(T)$ for each $i \in \mathcal{I}$.
3. Stop and return M_T when $T = T_{\mathrm{old}}$; otherwise repeat.

Theorem 5.19 (Exact finite schedule-class synthesis)**.** *Under Assumption 5.3, the following hold.*

(i) For every nonbottom $T \subseteq T_0$, $\mathcal{E}(T) = \bot$ exactly when $\mathsf{EC}_C(M_T) = \varnothing$; otherwise $M_{\mathcal{E}(T)} = \mathsf{EC}_C(M_T)$.

(ii) For every nonbottom $T \subseteq T_0$ and supervisor i,

$$M_{\mathcal{G}_i^{\infty}(T)} = \mathsf{RO}_{i,C}(M_T).$$

(iii) Algorithm 5.18 terminates at $\mathrm{supRLCO}(K_0, C)$*, including bottom. It makes at most $|\mathrm{Trans}(\mathcal{D})| + 1$ strict outer changes.*

Proof. For (i), suppose first that $D_0 \notin \mathrm{Bad}(T)$ and put $T' = \mathcal{E}(T)$. A word in $M_{T'}$ never visits a bad class. Hence every forced event at a visited class has a retained edge whose target is also nonbad. Lemma 5.11(i) shows that $M_{T'}$ is C-effectively controllable. Conversely, every nonempty prefix-closed C-effectively controllable $L \subseteq M_T$ satisfies $T_L \subseteq T'$ by Lemma 5.17(i), and therefore $L \subseteq M_{T'}$. Thus $M_{T'}$ is the supremal such sublanguage and

equals $\mathsf{EC}_C(M_T)$. If $D_0 \in \mathrm{Bad}(T)$, the same lemma shows that no nonempty admissible L exists, so $\mathsf{EC}_C(M_T) = \varnothing$.

For (ii), write $T' = \mathcal{G}_i^{\infty}(T)$. By the iterative survival statement in Lemma 5.17(ii), every nonempty prefix-closed sublanguage $L \subseteq M_T$ satisfying (R1) for i obeys $T_L \subseteq T'$ and therefore $L \subseteq M_{T'}$. The empty language is contained as well. It remains to prove that $M_{T'}$ itself satisfies (R1).

Fix σ and an active $\equiv_{i,\sigma}$-component H. Every class of H has a σ-transition because it occurs in a relation pair. At the stable transition set T' there are two exhaustive cases.

If some class of H has its σ-edge absent, then H is T'-defective. Stability under $\mathcal{G}_i$ implies that no present σ-edge remains in H; otherwise the operator would delete it. Lemma 5.12(i) therefore makes every represented continuation false in $M_{T'}$.

If no class of H has an absent σ-edge, then every such edge is present. No class can be uncovered, because an uncovered class would again make H defective and stability would delete all of its present σ-edges. Thus every class in H is covered, and Lemma 5.12(ii) makes every represented continuation true in $M_{T'}$.

Continuation membership is consequently constant throughout every active component, including a component consisting of a self-paired class reached by distinct words. Lemma 5.15 gives (R1). Since $M_{T'}$ is itself admissible and contains every admissible sublanguage of M_T, it equals $\mathsf{RO}_{i,C}(M_T)$.

For (iii), Lemma 5.14 ensures that every nonempty value of the exact language operators in (i)–(ii) has a transition-set representation. Each nonbottom outer pass therefore applies the exact deflationary operators EC_C and $\mathsf{RO}_{i,C}$ to the current language. A strict pass deletes at least one schedule-class edge; bottom can occur only once, proving termination and the stated bound. At a nonbottom fixed point the language is C-effectively controllable and satisfies (R1) for every supervisor. By Proposition 4.5, it is controllable and satisfies (R2), hence belongs to $\mathfrak{F}(K_0, C)$. Conversely, every member of $\mathfrak{F}(K_0, C)$ lies in the defining family of every operator and is contained in every iterate by the supremal properties in (i)–(ii). The fixed point contains all such members and is itself one of them, so it equals $\mathrm{supRLCO}(K_0, C)$. □

Whenever the returned language $M^\star$ is nonempty, it belongs to $\mathfrak{F}(K_0, C)$ and is therefore controllable and relatively legitimately co-observable. Theorem 4.3 makes $M^\star$ legitimately co-observable, and Theorem 3.2 together with the canonical policy (4) constructs an exact legitimate supervisor tuple for $M^\star$.

Corollary 5.20 (Order independence)**.** *The language returned by Algorithm 5.18 is independent of the fixed order chosen for the supervisors.*

Proof. For every order, Theorem 5.19 identifies the returned language with the unique language $\mathrm{supRLCO}(K_0, C)$. □

Remark 5.21 (Finite deletion certificates)*.* The pruning computation can retain a finite witness for every deletion. A class entering $\mathrm{Bad}(T)$ has a least-fixed-point rank, which gives a chain of forced transitions ending at an edge absent from T; if the initial class is bad, this chain certifies the bottom outcome. A guarded deletion has a path in an active $\equiv_{i,\sigma}$-component to a class witnessing defectiveness, together with either an absent σ-edge there or a path from the initial class that reaches it after crossing an edge outside T. These

data certify the false continuation and the safe-cell equality that propagates it through the component. They are diagnostic certificates, not additional assumptions of the fixed-point theorem.

Let N_{exp} be the total explicit encoding size of the reachable product Z and its transitions, alphabets, target and ambient flags, observer outputs, and authority data. The safe table is constructed from these product labels in polynomial time. Schedule classes, forced-event predicates, and the relations $\mathcal{P}_{i,\sigma}$ are then constructed within the bound of Corollary 5.23; they are not treated as free preprocessing inputs in that end-to-end bound.

Corollary 5.22 (Explicit materialized complexity)**.** *Given Z, $\mathcal{D}$, the forced-event predicates, and the explicit relations $\mathcal{P}_{i,\sigma}$, Algorithm 5.18 runs in polynomial time and space in these materialized structures.*

Proof. Bad-set propagation is reverse reachability over forced edges. Each guarded pass uses the two-layer coverage graph, component construction in the explicit pair relation, and edge scans. Every strict outer change deletes at least one schedule-class edge, so only polynomially many passes occur. □

Corollary 5.23 (End-to-end complexity and bounded width)**.** *Starting from the explicit product Z, the worst-case schedule-class construction has at most 2^z states and is therefore exponential in z. If every reachable estimate has cardinality at most w, construction of the safe table, schedule-class system, forced-event predicates, and pair relations, followed by Algorithm 5.18, takes*

$$z^{O(w)}\,\mathrm{poly}(N_{\mathrm{exp}}, |\mathcal{I}|)$$

time and space polynomial in the materialized schedule-class representation. Thus the procedure is XP in w and polynomial in the explicit product size for every fixed w; no polynomial bound is claimed when w is unbounded.

Proof. The safe table is obtained by grouping ambient decision states by (i, σ, y) and checking their authority labels, which is polynomial in N_{exp}. Theorem 5.8 bounds the numbers of schedule classes and edges by $z^{O(w)}$. Forced predicates are computed by scanning product states in each class. The explicit pair relation has at most $|\mathcal{I}||\Sigma_c||Q_{\mathcal{D}}|(|Q_{\mathcal{D}}|+1)/2$ pairs and can be generated by indexing safe outputs across classes. Corollary 5.22 then supplies a polynomial fixed-point cost in the constructed structures. Without a width bound, the subset construction gives the stated 2^z worst-case ceiling. □

When $w = 1$, every schedule class is a singleton. In particular, $\Delta = \varnothing$ gives explicit-state guarded pruning. Nevertheless, Example 5.13 shows that semantic self-pairs can already be essential at $w = 1$; estimate width and class/word path leakage are orthogonal. More generally, Corollary 5.9 gives a nontrivial polynomial class when hidden updates affect only a bounded finite-mode component.

6. A cyclic comparison family under dynamic authority

Example 5.13 isolates the class/word leakage and semantic self-pair mechanism used by guarded pruning. The family below instead isolates the authority-coverage mechanism that induces effective uncontrollability in a parameterized dynamic-authority system.

A parameterized finite family makes the synthesis mechanism and its separation from the classical static-assignment baseline explicit. It is a structural example, not a scalability claim. Fix $m \geq 1$, write $\mathbb{Z}_m$ for indices modulo m, and put

$$\Sigma_c = \{a, b, c, \sigma\}, \qquad \Sigma_u = \{r\}, \qquad \Delta = \{\delta_1, \delta_2\}.$$

The operational locations are H_j, A_j, B_j, C_j for $j \in \mathbb{Z}_m$, and the authority modes are q_1, q_2. These capitalized location labels are local to this example and are unrelated to the active-set notation A_x. The deterministic extended plant has states (v, q_k), initial state (H_0, q_1), and the following transitions, where every operational event preserves q_k:

Source	Event	Destination
(H_j, q_k)	a, b, c	$(A_j, q_k), (B_j, q_k), (C_j, q_k)$, respectively
$(A_j, q_k), (B_j, q_k), (C_j, q_k)$	r	(H_{j+1}, q_k)
$(A_j, q_k), (B_j, q_k), (C_j, q_k)$	σ	the same state
(v, q_k)	δ_ℓ	(v, q_ℓ) for $\ell \in \{1, 2\}$

All unspecified plant transitions are undefined.

The configuration set and map are

$$\mathsf{Conf} = \{H_j, A_j, B_j, C_j : j \in \mathbb{Z}_m\} \times \{q_1, q_2\}, \qquad c(h) = (\mathrm{loc}(h), \mathrm{mode}(h)).$$

The authority configuration is permitted to refine the plant location; this is what allows the single service event σ to carry different issuer rights at different operational locations. For $k \in \{1, 2\}$, authority is given by

$$\begin{aligned} &\mathsf{A}((H_j, q_k), e) = \{1, 2\}, && e \in \{a, b, c\}, \\ &\mathsf{A}((A_j, q_k), \sigma) = \{1\}, \\ &\mathsf{A}((B_j, q_k), \sigma) = \mathsf{A}((C_j, q_k), \sigma) = \{k\}. \end{aligned}$$

Every remaining event–configuration entry is assigned $\{1, 2\}$, making A total and nonempty valued without affecting any enabled decision.

Both supervisors receive the same plant-total finite observer. Its output is (j, v, β), where j and $v \in \{H, A, B, C\}$ record the operational location and $\beta \in \{0, 1\}$ is initially zero and becomes one at the first occurrence of b. It ignores δ_1, δ_2, which act as observer self-loops. Thus the supervisors are observationally identical and differ only in their authority. Consequently, the separation cannot be attributed to asymmetric observation quality. The bit β is load bearing: without it, a hub cell would merge b-free and b-visited ambient words, although the a-continuation is retained for the former and absent for the latter in the language below.

Let $G_\Sigma^{(m)}$ be the operational generator obtained by forgetting the mode and update events. The upper language $K_0^{(m)}$ is generated by its operational table after deleting the σ-loops at A_j and B_j and retaining the loop at C_j, with every branch and return retained. Set $C = K_0^{(m)}$; a complete recognizer adds one rejecting sink. The lifted recognizers and both observers are functions of the operational word alone, and each δ_ℓ changes only the mode coordinate. Since both updates are enabled at every reachable location,

$$|D(s)| = 2 \qquad \text{for every } s \in L(G_\Sigma^{(m)}),$$

with the two representatives differing only in mode. Hence the current-estimate width is exactly $w = 2$.

The resulting safe-issuer table is

Ambient decision cell	Safe issuers
Hub cell (j, H, β) for each of a, b, c	$\{1, 2\}$
A_j service cell (j, A, β)	$\{1\}$
B_j or C_j service cell	$\varnothing$

Indeed, every B_j- and C_j-service cell contains both mode representatives: supervisor 1 is inactive in mode q_2 and supervisor 2 is inactive in mode q_1.

Let $M^{(m)}$ be the b-free part of $K_0^{(m)}$: it retains the a- and c-branches, every return from a retained location, and the service loop at each C_j.

Proposition 6.1 (Closed-form relative outcome)**.** *For every $m \geq 1$,*

$$\mathrm{supRLCO}(K_0^{(m)}, K_0^{(m)}) = M^{(m)}.$$

Its displayed reachable operational generator has $3m$ states, $5m$ transitions, $2m$ retained branch routes, and m service states, where a service state means a reachable state with a retained σ-transition.

Proof. The language $M^{(m)}$ is prefix closed and controllable because every return from a retained decision location is retained. For (R1), the A_j service cells are authority safe and their σ-continuations are uniformly absent. The B_j and C_j service cells have no safe issuer and impose no (R1) test. At hubs, the observer bit separates b-free from b-visited words. In every b-free hub cell, the a- and c-continuations are uniformly present and the b-continuation is uniformly absent; every word in a b-visited hub cell lies outside $M^{(m)}$, so all its candidate continuations are absent. Thus (R1) holds on every authority-safe cell.

For (R2), each removed hub b-continuation has safe issuers $\{1, 2\}$, and each forbidden A_j service continuation has safe issuer $\{1\}$. The legal C_j continuation is retained, so every candidate-forbidden point is covered. Hence $M^{(m)} \in \mathfrak{F}(K_0^{(m)}, K_0^{(m)})$.

For maximality, let L be any nonempty member of this family and suppose that it retains a b-branch. The resulting B_j prefix belongs to L, the service event σ is plant available there, and the safe-issuer table is empty. Thus σ is C-effectively uncontrollable at that history (Definition 4.7), and Proposition 4.5(ii) forces the σ-continuation into L, contradicting $L \subseteq K_0^{(m)}$. Thus every family member is b-free and lies in $M^{(m)}$. The counts follow from the displayed reachable generator $\{H_j, A_j, C_j : j \in \mathbb{Z}_m\}$: per module it has two branch transitions, two returns, and one retained service loop. □

For comparison, first drop issuer legitimacy but retain active-set conjunction. The explicit local policies admit every branch, make supervisor 1 veto at A_j, make both supervisors veto at B_j, and make both admit at C_j. Whichever supervisor is active at B_j therefore vetoes, so these policies realize $K_0^{(m)}$ exactly under active filtering. They are not legitimate: the supervisor inactive in each mode still emits a veto on the common B_j observation.

The established decentralized baseline is classical relative co-observability on the extended plant after the authority automaton has been included in the plant state but issuer admissibility has been discarded. Assign every controllable event statically to both supervisors,

$I_e = \{1, 2\}$ for $e \in \Sigma_c$, and take the ambient language to be $\widetilde{K}_0^{(m)}$. The observer output separates the location types H, A, B, C and the bit β, while membership in $\widetilde{K}_0^{(m)}$ is independent of the hidden authority mode. Thus continuation membership is constant on every ambient decision cell for each controllable event. Since $\widetilde{K}_0^{(m)}$ is also controllable, it is itself the supremal controllable classically relatively co-observable sublanguage under these static assignments. Its operational projection is $K_0^{(m)}$. In this family the classical operator therefore returns exactly the same operational behavior as the explicit active-filtering policies above, but it does not express their issuer-admissibility failure.

The table reports reachable operational generators; no minimization is claimed. The active-filtered policies remain in the prose because they realize the same operational language as the classical row while exposing its unauthorized veto.

Method	States	Transitions	Branches	Service states
Issuer-admissible RLCO	$3m$	$5m$	$2m$	m
Classical relative co-observability	$4m$	$7m$	$3m$	m

Issuer-admissible synthesis therefore rejects the uncovered forbidden B_j branch while preserving the legal C_j service loop. Classical relative co-observability retains B_j because it treats both supervisors as statically assigned to the service event and imposes no history-dependent constraint on the veto issuer. The equivalent active-filtering realization makes the gap concrete by using behaviorally effective but unauthorized local vetoes.

The fixed point has a direct causal interpretation. At a B_j class, $\mathrm{NoSafeIssuer}(D(B_j), \sigma) = 1$, so the service event is effectively uncontrollable in the sense of Definition 4.7 and $\sigma \in \mathrm{Force}(D(B_j))$. Its service edge is absent from T_0; effective-controllability pruning marks the class bad and removes the incoming b-branch. At C_j, the legal service edge remains for the complete two-state class, as required by Proposition 5.1. Corollary 5.23 applies with $w = 2$.

7. Conclusion

History-dependent veto authority creates an aligned-witness requirement: one supervisor's information cell must establish both the forbidden decision and the feasibility of that supervisor's veto throughout the cell. This fits the general language of observation-uniform strategies, but the separate decision and authority marginals do not compose; the tight three-point obstruction identifies the minimal failure. Controllability and legitimate co-observability then exactly characterize prefix-closed realization.

For synthesis, we use the fixed-ambient scaffold of relative co-observability. Issuer-safe filtering creates a new vacuity: the coverage obligation repairs it and induces effective uncontrollability where no safe issuer exists. Hidden authority schedules must be retained by complete operational fibers. Schedule classes carry those fibers finitely, and coverage plus semantic self-pairs make the alternating pruning algorithm exact despite the gap between class transitions and word membership. The supplementary cross-check reproduces both worked mechanisms and finds no disagreement with exhaustive enumeration on the reported bounded instances.

The main notion is intentionally specification relative: it constrains issued actions on the behavior being exactly realized, not on every plant history. Plant-wide legitimacy is strictly

stronger in general, although its existence coincides under plant-wide local authority observability. Extending supremal synthesis to plant-wide admissibility without that observability, synthesizing delegation itself, marked nonblocking behavior, communication that repairs failed alignment, and complexity below the explicit saturated representation remain open directions.

Appendix A. Computational cross-check

The archived supplementary verification artifact (Kingston, 2026b) contains `authority_synthesis_validator.py`, a self-contained Python reference implementation of schedule-class construction, the operators $\mathcal{E}$ and $\mathcal{G}_i^{\infty}$, and Algorithm 5.18. Its oracle does not call those pruning operators. Instead, for each small instance it enumerates every transition subset of T_0, evaluates C-effective controllability and (R1) directly from reachability, path coverage, and the word-level continuation truth sets, and forms the unique supremal candidate among the valid sublanguages. By Lemma 5.14, restricting this oracle to transition-represented candidates is complete for the supremal comparison: every feasible prefix-closed sublanguage is contained in a feasible transition saturation. The artifact also checks the exact language operators separately and reruns the complete fixed point with the supervisor order reversed.

The command

```
python3 authority_synthesis_validator.py --samples 20000 \
  --seed 20260801 --output authority_synthesis_validation.json
```

reproduces the results in Table 2. The fixed integer seed initializes the pseudorandom generator so that rerunning the same script regenerates the same sequence of test instances. The first paper test starts from the punctured transition set in Example 5.13; both guarded pruning and exhaustive (R1) enumeration return the edge for a and delete the σ-edge, yielding $\{\varepsilon, a\}$. The second instantiates the cyclic family at $m = 1$; the full algorithm and the oracle agree on the b-free outcome of Proposition 6.1. The random generator uses one or two hidden modes per schedule class, at most five reachable classes, and at most eleven initial candidate edges. Hidden-update observer construction is performed before synthesis in every generated instance.

Table 2. Finite computational cross-check. "Operator checks" count separate comparisons of $\mathcal{E}$ or $\mathcal{G}_i^{\infty}$ with its exhaustive oracle.

Check	Instances or checks	Agreements
Semantic self-pair example ($w = 1$)	1	1
Cyclic family ($m = 1$, $w = 2$)	1	1
Seeded random full fixed points	20,000	20,000
Hidden-update observer constructions	20,000	20,000
Reversed supervisor-order comparisons	20,000	20,000
Exact operator comparisons	50,025	50,025

Among the random instances, 3,555 had width one and 16,445 had width two; 17,747 contained a semantic self-pair and 5,064 contained a cross-class safe pair. The full algorithm returned bottom in 9,425 cases and made a strict nonbottom pruning change in 967 further cases. The machine-readable output is distributed as `authority_synthesis_validation.json`.

This bounded exhaustive exercise is designed to detect implementation or proof-formalization errors in the bespoke pruning logic. It does not prove the theorem, establish an asymptotic performance claim, or make the finite search exhaustive beyond the reported generator and bounds.

References

K. Cai, R. Zhang, and W. M. Wonham, "Relative Observability of Discrete-Event Systems and Its Supremal Sublanguages," *IEEE Transactions on Automatic Control*, 60(3):659–670, 2015. doi:10.1109/TAC.2014.2341891.

K. Cai, R. Zhang, and W. M. Wonham, "Relative coobservability in decentralized supervisory control of discrete-event systems," *International Journal of Control*, 92(7):1481–1489, 2017. doi:10.1080/00207179.2017.1397754.

C. G. Cassandras and S. Lafortune, *Introduction to Discrete Event Systems*, 3rd ed., Springer, Cham, 2021. doi:10.1007/978-3-030-72274-6.

J. Dai and H. Lin, "Coordination and Control of Distributed Discrete Event Systems under Actuator and Sensor Faults," *arXiv preprint* arXiv:1707.05428, version 2, 2018. arXiv:1707.05428.

P. Kingston, "Equality-Support Covers for Exact Branch-Set Preservation under Fixed Actuator Interfaces," 2026. doi:10.13140/RG.2.2.14253.65767.

P. Kingston, *Verification Artifact for Exact Realization and Supremal Synthesis under History-Dependent Veto Authority*, Zenodo, 2026. doi:10.5281/zenodo.21728999.

D. Kozen, "Lower bounds for natural proof systems," in *Proceedings of the 18th Annual Symposium on Foundations of Computer Science*, pp. 254–266, 1977. doi:10.1109/SFCS.1977.16.

R. Kumar and S. Takai, "Inference-Based Ambiguity Management in Decentralized Decision-Making: Decentralized Control of Discrete Event Systems," *IEEE Transactions on Automatic Control*, 52(10):1783–1794, 2007. doi:10.1109/TAC.2007.906158.

F. Liu and H. Lin, "Reliable supervisory control for general architecture of decentralized discrete event systems," *Automatica*, 46(9):1510–1516, 2010. doi:10.1016/j.automatica.2010.06.011.

B. Maubert and S. Pinchinat, "A general notion of uniform strategies," *International Game Theory Review*, 16(1):1440004, 2014. doi:10.1142/S0219198914400040.

H. Mortazavian and F. Lin, "Decentralized supervisory control of discrete event systems with nonhomogeneous control structure," *Information Sciences*, 68(3):233–246, 1993. doi:10.1016/0020-0255(93)90107-W.

P. J. Ramadge and W. M. Wonham, "Supervisory Control of a Class of Discrete Event Processes," *SIAM Journal on Control and Optimization*, 25(1):206–230, 1987. doi:10.1137/0325013.

S. L. Ricker and K. Rudie, "Know means no: Incorporating knowledge into discrete-event control systems," *IEEE Transactions on Automatic Control*, 45(9):1656–1668, 2000. doi:10.1109/9.880616.

K. Ritsuka and K. Rudie, "Do what you know: coupling knowledge with action in discrete-event systems," *Discrete Event Dynamic Systems*, 33(3):257–277, 2023. doi:10.1007/s10626-023-00381-z.

K. Rohloff, T.-S. Yoo, and S. Lafortune, "Deciding co-observability is PSPACE-complete," *IEEE Transactions on Automatic Control*, 48(11):1995–1999, 2003. doi:10.1109/TAC.2003.819285.

K. Rudie and W. M. Wonham, "Think globally, act locally: decentralized supervisory control," *IEEE Transactions on Automatic Control*, 37(11):1692–1708, 1992. doi:10.1109/9.173140.

T. Schwarz, R. Sim, and C. K. Ling, "Computing Equilibria in Games with Stochastic Action Sets," *arXiv preprint* arXiv:2602.16234, 2026. arXiv:2602.16234.

T. Schwarz, R. Sim, and C. K. Ling, "Equilibrium Computation in Extensive-Form Games with Stochastic Action Sets," *arXiv preprint* arXiv:2606.13093, 2026. arXiv:2606.13093.

S. Takai and R. Kumar, "Synthesis of Inference-Based Decentralized Control for Discrete Event Systems," *IEEE Transactions on Automatic Control*, 53(2):522–534, 2008. doi:10.1109/TAC.2007.915171.

S. Takai and T. Ushio, "Decentralized Supervisory Control of Discrete Event Systems Using Dynamic Default Control," *IEICE Transactions on Fundamentals of Electronics, Communications and Computer Sciences*, E88-A(11):2982–2988, 2005. doi:10.1093/ietfec/e88-a.11.2982.

Y. Tong and K. Cai, "Supervisory Control with Event Forcing Under Partial Observation," *arXiv preprint* arXiv:2607.21040, 2026. arXiv:2607.21040.

P. Wang and K.-Y. Cai, "Supervisory control of discrete event systems with state-dependent controllability," *International Journal of Systems Science*, 40(4):357–366, 2009. doi:10.1080/00207720802436711.

T.-S. Yoo and S. Lafortune, "Decentralized supervisory control: a new architecture with a dynamic decision fusion rule," in *Proceedings of the 6th International Workshop on Discrete Event Systems (WODES 2002)*, Zaragoza, Spain, pp. 11–17, 2002. doi:10.1109/WODES.2002.1167663.

T.-S. Yoo and S. Lafortune, "A General Architecture for Decentralized Supervisory Control of Discrete-Event Systems," *Discrete Event Dynamic Systems*, 12(3):335–377, 2002. doi:10.1023/A:1015625600613.

T.-S. Yoo and S. Lafortune, "Decentralized supervisory control with conditional decisions: supervisor existence," *IEEE Transactions on Automatic Control*, 49(11):1886–1904, 2004. doi:10.1109/TAC.2004.837595.

R. Zhang, K. Cai, and W. M. Wonham, "Supervisor localization of discrete-event systems under partial observation," *Automatica*, 81:142–147, 2017. doi:10.1016/j.automatica.2017.03.018.