Certifying SDCPNs without a non-Zeno axiom, and with finitely many verification conditions

safety unconditional, liveness angelic, the worst case priced; one footprint VC per Petri transition

davidad's AI ensemble*

MAITRIA — Mathematical Assurance and Interaction Toolkit for Reliable Intelligent Agents

2 August 2026 — draft for discussion

1 Scope, and how claims are graded

This report addresses two questions about the certification of stochastically and dynamically coloured Petri nets (SDCPNs, Everdij–Blom [1]):

1. **Non-Zeno.** How are certificates about an SDCPN protected from Zeno behaviour — and in particular, what do rate majorants actually buy, given that forced (boundary-hit) transitions have no rate to majorize?
2. **Variadic certificates.** SDCPN markings have variable and canonically unbounded token counts, and each transition acts at every multiset of tokens matching its arcs — so a certificate seems to need infinitely many verification conditions (VCs): one per token configuration, and one per transition *instance*. What replaces that infinity?

The short answers, unpacked in Parts I and II:

> (1) Safety certificates never needed a non-Zeno hypothesis; the demand "prove every realization non-Zeno before certifying anything" dissolves at the quantifier, not at a cleverer certificate. What remains is a per-wing price list: rate-raced transitions are capped by their declared majorants, delay and scheduler wings by their declared dwell floors, and the one genuinely rate-less wing — instantaneous forced crossings — is priced by a certified *crossing budget*, with a machine-checked refusal theorem on genuinely Zeno systems.
> (2) Neither infinity survives contact with the certificate language. Certificates are polynomials in *pooled* coordinates, so the marking's unboundedness never enters the VC count; and one *footprint VC* — one polynomial sequent — per Petri transition covers every instance at every marking, with machine-checked transfer theorems saying exactly which aggregation shapes support which claim polarities.

This revision also carries the worked exhibit to its end (§§9–10): its open supervisor seat exercised end to end — joint policy+certificate synthesis, then wall-clock (real-time) certificates, an exactly optimal-throughput policy with a machine-checked imprecise-decision-theoretic reduction, an *untrusted* policy running inside a certified envelope, and one certificate instantiated across six orders of magnitude of lot count at flat checking cost.

Grading. Every substantive claim below carries one of six grades — occasionally annotated with what remains to reach the next grade up (e.g. "argued, at sketch strength") — and the claim ledger (§11) collects them:

- **[Lean-checked]** — stated and proved in Lean 4 over mathlib against our mechanized semantics; zero `sorry`s; the axiom footprint of every cited theorem is exactly {`propext`, `Classical.choice`, `Quot.sound`}; and adversarial *mutants* (deliberate corruptions at loci preregistered before proving) are refused by the checker.

*Drafted, machine-checked, and revised primarily on Claude Fable 5; adversarial review rounds on GPT-5.6 Sol.

- [**geolog-checked**] — carried as rows of a sealed certificate container and re-derived mechanically from the bytes: by the geolog engine's evaluator walk in the loop, and independently by the producer's exact-rational tier; adversarial mutants at preregistered loci refused. Machine-checked, but by the geolog engine rather than the Lean kernel — a weaker warrant than [**Lean-checked**], a stronger one than [**argued**].
- [**cited**] — a published result by others, relied on with premises audited against our normal form; in-house mechanization is a standing debt — deferred, not waived.
- [**argued**] — complete on paper: a house proof in which every dropped hypothesis is refuted by an executable exact-rational countermodel; mechanization is priced (in agent-hours, from measured prior artifacts) and small.
- [**designed**] — statement and proof route fixed; not yet carried out.
- [**definitional**] — a consequence of declared semantics (readable off the syntax), not a theorem anyone must trust.

This grading is load-bearing, not decorative: one lemma in an earlier draft of Part I was caught false at the [**argued**] tier by its own countermodel discipline and repaired *before* any claim shipped (footnote 1).

Part I — Non-Zeno

2 The reframe: the quantifier was the problem

Requiring non-Zeno of a whole SDCPN before certifying anything about it was always a category error. The impossible problem — proving that *all* realizations of an SDCPN are live, including the Zeno ones — is impossible because its quantifier is wrong, not because the certificates are weak. The repair is semantic: interpret the model in a universe whose denotations are *credal bodies* of subprobability laws (constructible classes in a hyperspace of belief states), and read the two claim polarities off the body with their native quantifiers — **safety universally** (every admitted law, Zeno ones included, as honest subprobability citizens) and **liveness existentially**: the *angelic* quantifier (some admitted law reaches the horizon with the demanded mass). Under that interpretation nobody ever needs all realizations live. Three acts follow, cheaper in the order stated: safety consumes no non-Zeno resource at all (§3); angelic liveness consumes one honest law (§5); only worst-case (all-scheduler) liveness consumes non-Zeno of the whole envelope, and that is exactly what the crossing budget prices (§6).

Why a non-Zeno axiom felt necessary upstream: the classical construction had no home for partial laws, so one standing assumption — $\mathbb{E}[N_t] < \infty$, condition G4 of [2], with N_t the total number of transition firings by time t — did double duty, as totality of the construction and as an implicit liveness quantifier baked into the semantics. We moved the quantifier out of the semantics and into the claim. That is also why we can say when it *fails*: see the non-vacuity paragraph of §5.

3 Act I: safety is unconditional [Lean-checked]

In the jump-by-jump (segment/skeleton) semantics underlying Everdij and Blom's correspondence between SDCPNs and general stochastic hybrid systems (GSHS) [1, 2], Zeno behaviour is *emergent partiality*: the process stopped at its accumulation time is the classical minimal (killed) process, and its time-t law is a subprobability measure — an honest citizen of the credal-body universe. Safety-type certificates bind with no non-Zeno hypothesis: the supermartingale safety bounds quantify over the *total* jump-indexed chain, which enumerates every state a Zeno execution ever occupies (nonnegativity and Fatou replace any integrability of the jump count). No dwell floors, no rate caps, and no crossing counts appear in any statement or proof — a fact one checks by reading the Lean statements, not by trusting this sentence. The kernel-reset versions cover SDCPN's random firing measures at this rung.

What about behaviour after the accumulation point? The Everdij–Blom construction defines the process jump by jump: under G4 no accumulation occurs; without G4 the minimal reading kills at T_∞. There is no post-Zeno behaviour in the model class for a safety claim to miss. The honest residue is model-versus-plant (the plant does something after t^*; the model is silent), and "the model describes the plant at all times" is a *liveness-of-definedness* claim — under the angelic reading, the demand that some law be defined to the horizon, which is precisely §5's certificate, made visible rather than taxed onto safety. (The classical sticking/sliding completions in the tradition of Johansson, Egerstedt, Lygeros, and Sastry are the *modelling* move that re-totalizes; the budget of §6 is the *certification* move that proves re-totalization unneeded.) At the specification layer, safety clauses read time × state regions, which factor exactly through the floor-forgetting projection of the hybrid time domain, so post-Zeno safety is stateable with no exotic machinery; only stage-*ordering* clauses (until-sequencing at a shared instant) genuinely need the full $\omega + 1$ stage structure [**Lean-checked — both faces: the factoring is exact and hypothesis-free, and the ordering carve-out is itself machine-checked negatively (equal shadows, opposite until-order)**].

4 Rates are majorized by declaration: the per-wing count prices

The rate majorant is not a certificate one must add — in our transition vocabulary it is the modelling primitive. A rate-raced transition declares a constant rational majorant Λ with *thinning* semantics: the effective jump rate is the majorant times the kernel's non-stay mass, so effective rate $\le \Lambda$ is definitional, and a finite conforming roster has total finite-rate intensity $\le \sum_i \Lambda_i$ — a constant read off the syntax. The scheduler and deterministic-delay wings declare strictly positive dwells, each pathwise-capped at $T/d + 1$ firings per horizon. The one representable Zeno engine is the *instant* wing (boundary-hit crossings at infinite majorant), and §6's budget prices exactly it. One price per wing:

transition wing	declared datum	count price by horizon T	grade
rate-raced	majorant Λ_i	a.s. finite; $\mathbb{E} \le \Lambda_i T$	cap [**definitional**]; a.s. finiteness [**Lean-checked**][a]; $\mathbb{E}$-bound [**Lean-checked**][c]
deterministic delay	dwell $d_i > 0$	$\le T/d_i + 1$ pathwise	[**definitional**][b]
scheduler-fired	dwell floor $m_i > 0$	$\le T/m_i + 1$ pathwise	[**definitional**][b]
instantaneous (boundary-hit)	budget pair (W, c)	$\mathbb{E} \le W(x_0) + cT$; $\mathbb{P}(\text{Zeno}) = 0$	[**Lean-checked**] (§6)

[a]via the glue lemma below; [b]the caps are semantics-level arithmetic; their record-level restatements are routine and [**designed**]; [c]on its canonical witness (the race member's counted chain), tight at $r = \Lambda$; the joint-chain form is a named successor.

So: *non-Zeno by any horizon is the sum of four receipts, three of them read off the syntax, the fourth a certified witness* — and refusal-complete where the system is genuinely Zeno (§6).

The probabilistic glue for the rate wing is [**Lean-checked**]: for adapted nonnegative dwells (τ_n) with respect to a filtration $(\mathcal{F}_n)$, almost surely

$$\textstyle\sum_n \tau_n = \infty \quad \Longleftrightarrow \quad \sum_n \mathbb{E}[\,\tau_n \wedge 1 \mid \mathcal{F}_{n-1}] = \infty,$$

the *truncated* extended Borel–Cantelli lemma[1] (cf. Kallenberg [4]), mechanized in iff form together with its two corollaries: if the truncated conditional dwell means are floored at some $\varepsilon > 0$, then the raw dwell sums diverge almost surely, and hence only finitely many events arrive before any horizon. Instantiating the floor from the roster's total rate cap $\Lambda = \sum_i \Lambda_i$ is the elementary bound $\mathbb{E}[\tau \wedge 1 \mid \mathcal{F}] \ge (1 - e^{-\Lambda})/\Lambda > 0$ [**Lean-checked**]. The remaining quantitative refinement on this wing — the compensator bound $\mathbb{E}[N_t^{\mathrm{r}}] \le \Lambda t$, where N_t^{r} counts rate-wing firings by time t — is [**Lean-checked**] on its canonical witness — the race member's counted chain, with the compensator inequality proven from the thinning normal form and tight at $r = \Lambda$ — so G4's total count *is* a theorem for conforming models, all four wings summed:

$$\mathbb{E}[N_t] \;\le\; \underbrace{W(x_0) + ct}_{\text{instantaneous}} + \underbrace{\Lambda t}_{\text{rate-raced}} + \underbrace{\textstyle\sum_i \big(t/d_i + 1\big)}_{\text{both dwell-floored wings}} \;<\; \infty.$$

One boundary is deliberate rather than accidental [**definitional**]: constant majorants over a finite roster make unbounded-intensity models (e.g. mass-action rates growing with an unbounded population) unrepresentable in the conforming normal form. SDCPNs discharge this classically — an invariant-bounded token count plus compact colour boxes yield finitely many modes, where constant majorants suffice; a state-dependent majorant species with a Foster–Lyapunov non-explosion certificate is the named extension if a consumer ever needs one.

[1]The truncation is not decorative: without it the right-to-left inclusion fails — independent $\tau_n = n^3$ with probability n^{-2} (else 0) has conditional-mean sums divergent deterministically while $\sum_n \tau_n < \infty$ almost surely, by ordinary Borel–Cantelli. An earlier internal draft stated the untruncated form; the grading discipline of §1 caught it at the [**argued**] tier — countermodel first, then the classical repair, then the mechanization — before anything shipped. The truncated form is exactly as strong for non-Zeno, since a nonnegative series and its 1-truncation diverge together pathwise.

5 Act II: liveness is angelic — one witness law

The claims Zeno actually threatens — reach-by-deadline, mass floors at a horizon, definedness-to-horizon — consume survival, and under the angelic reading they demand it of *one* admitted law: "some realization reaches the horizon and attains the bound." The certificate that exhibits the law is a *positional selection*: a safety certificate overapproximates the reachable set by a region S; the witness gives a finite cell cover of S and, per cell, a chosen transition landing back in S at whose target the *expected* dwell is bounded below by a fixed $\varepsilon > 0$. The bound is deliberately an *expected* dwell, never a hard floor, and it is one number read two ways: expected dwell $\geq \varepsilon$ at a state is exactly total exit rate $\leq 1/\varepsilon$, read off the per-transition majorants of §4. The pointwise VCs glue coinductively into one memoryless scheduler resolution [**Lean-checked — one glued kernel, all four conclusions, mutants refused at the preregistered loci**]; the truncated Borel–Cantelli lemma [**Lean-checked**] then gives non-Zeno of that law outright. Zero-time steps are disciplined structurally rather than counted: pass-through cells must land where dwelling is certified, so instantaneous chains cannot form — no crossing counting and no envelope pricing anywhere in this act. Scheduler-fired transitions carrying a declared minimum dwell strengthen this further: with all minimum dwells positive, scheduler-driven Zeno is unrepresentable by construction.

Non-vacuity, and the diagnostic gift. The G4 axiom is the angelic quantifier hardened into the semantics — we moved it into the claim, which is why we can also say when it fails. If every consistent law of a model is Zeno, no positional selection exists: the certificate search fails *with the obstructing cell in hand* — a pointer at the model's actual Zeno engine — and the verdict is NOT LIVE, never vacuous truth. A model whose body is all-Zeno gets the same non-vacuous refusal as an algebraically inconsistent one: it is not live because it denotes no records at the horizon.

6 Act III: the worst case is priced — the crossing budget [Lean-checked]

Where a claim genuinely demands progress under *every* scheduler, or the event count is itself the claim (tails, throughput ceilings), the price is the crossing budget. For a hybrid system with finitely many modes, polynomial data per mode, guard reset map R, and extended generator $\mathcal{A}$, a budget witness $W \geq 0$ (polynomial per mode) with constant $c \geq 0$ gives:

$$\underbrace{\mathcal{A}W \leq c \text{ on mode cells}}_{\text{generator VC}} \qquad \wedge \qquad \underbrace{W(R(x)) - W(x) \leq -1 \text{ on guard slabs}}_{\text{seam VC}}$$
$$\Longrightarrow \qquad \mathbb{E}\big[\#\{\text{forced crossings by } T\}\big] \leq W(x_0) + cT,$$

together with the Markov tail $\mathbb{P}[N_T^{\mathrm{f}} \geq K] \leq (W(x_0) + cT)/K$ and $\mathbb{P}(\text{Zeno by } T) = 0$, where N_T^{f} counts the boundary-hit crossings (the literature's *forced* transitions) by horizon T. Why a *potential* witness and not a rate bound: forced jump times are predictable, so the forced-crossing process is its own compensator — there is no intensity function to bound (cf. Davis [3] for the compensator calculus); W is to crossings what a Foster–Lyapunov function is to escape, and at $c = 0$ the minimal solution is $\mathbb{E}_x[\text{remaining crossings}]$. The classical rank-function/acyclicity method is precisely the $c = 0$, $\mathbb{N}$-valued instance (machine-checked bridge). A multiplicative variant, the *tilted* budget $\mathcal{A}W \leq cW \Rightarrow \mathbb{E} \leq e^{cT}\, W(x_0)$, covers chattering-class (superlinear but non-Zeno) systems. Sharpness both ways is [**Lean-checked**]: a floored two-mode thermostat whose certified budget equals its true crossing rate exactly; and a refusal theorem — the shrinking-dwell, genuinely Zeno thermostat admits *no* budget pair (W, c) of any regularity. For readers who do want the classical total reading: the budget converts the standing assumption $\mathbb{E}[N_t] < \infty$ from an axiom into a certified, refutable obligation.

Worst-case reach-avoid decomposes exactly [Lean-checked] (invariant (U) a stated hypothesis). For a horizon-T reach-avoid demand — reach G by T while avoiding B — over a set of admitted subprobability laws, fix the outcome conventions once (first decisive outcome; goal-settlement outranks a simultaneous bad-hit; "by T" means $\leq T$; missing mass is the completion point of the one-point completion) and write, per law: the *reach-avoid probability* (mass that arrives, avoiding B); the *safety value* (one minus the mass of witnessed violations — hitting B, or overrunning the deadline unreached; a trajectory that terminates early violates nothing); and the *completion value* (mass whose episode runs to the horizon or to settlement, at the goal *or at the bad set* — the inclusion is part of the definition). These satisfy an exact Łukasiewicz law:

$$\mathbb{P}(\text{reach-avoid}) \;=\; \text{safety} + \text{completion} - 1,$$

because the two events overlap precisely on fulfilment and jointly cover the completed outcome space, while missing mass is credited to safety and charged to completion. Consequently worst-case (all-laws) reach-avoid at level p is *equivalent* to the all-laws floor on the sum at $1 + p$, and is *implied* by all-laws safety at p_1 plus all-laws completion at p_2 whenever $p_1 + p_2 = 1 + p$; the angelic conjunct "some law completes at level p" adds exactly non-vacuity. Two cautions the arithmetic enforces: an *existential* completion witness cannot substitute for the all-laws completion floor (a two-law system with one perfect reacher and one immediate halter passes both conjuncts at every level with worst-case reach-avoid zero); and a completion reading demanding presence at exactly T, rather than settlement, breaks the equivalences. The identity is machine-checked at exact rationals at the outcome-partition level and has passed an independent adversarial verification round; identifying a given specification language's until/invariance/lifetime denotations with these outcome classes is a stated invariant, discharged per semantics.

Recurrence in continuous time: two meters [Lean-checked separations]. Discrete-time treatments see one recurrence rule where continuous time has two. A recurrence obligation may be *dwell-metered* (total occupation of the request region diverges) or *crossing-metered* (infinitely many entry events); the species are provably incomparable — enter-and-stay satisfies the dwell meter with a single entry event, while vanishing-occupation chattering satisfies the entry meter with convergent dwell — and they collapse over uniform time lifts, which is exactly the discrete-time embedding. The certificate kit splits accordingly: dwell-metered recurrence consumes §§4–5's rate/dwell machinery; crossing-metered recurrence consumes this section's budget. Rules stated over discrete time generalize to the hybrid semantics only after choosing a meter — the choice is semantic content, not bookkeeping.

7 SDCPN transition types, one clause each

Delay transitions. Totally inaccessible jump times with absolutely continuous compensator — the declared majorant handles their count (§4); this is why "rate-bound" never solved the guard case, where no density exists.

Immediate transitions. Canonically cascade-closed into the jump kernel (vanishing markings are not states); cascade termination by syntactic acyclicity or by a cascade-rank certificate — the latter [**Lean-checked**] at the net level, including the scheduler-quantified face, where the correct adversary class is the enabling-respecting one (a scheduler weighting a disabled transition would sit still in zero time; the rank machinery refuses that rather than papering over it). Instantaneous cascades under the budget are *priced, not banned*: a k-fold same-instant cascade needs $W \geq k$ at entry; an infinite one forces $W = \infty$ — refused (Davis's hypothesis that resets never land on the active boundary [3] is this certificate's qualitative shadow, refined in both directions).

Guard transitions. The budget's target (§6) — and, under the angelic act, the selection's pass-through discipline (§5).

Transport. The theorems of §§3 and 6 live at the general-stochastic-hybrid stratum. We do not, however, take the flattening of an SDCPN to a GSHS as the *definition* of the net's semantics: the flattened object has one mode per reachable token distribution — combinatorially many under bounded token counts, countably many without — while the net is the compact object, and the certificate language of Part II is deliberately polynomial in the net, never in its mode graph. Instead, each formalism carries its own adequacy square against the same semantic universe — its community semantics tied to its compiled certificate reading — and the flattening enters as a translation between presentations, bridging the two compiled readings in a commuting triangle. The triangle's composed corollary is now [**Lean-checked**] at the uncoloured jump-skeleton grain: the token-game reading and the flattened-execution reading sit one invertible cell apart, assembled by *pasting* three independently landed cells (never re-proved monolithically), with sound-direction demonic and angelic faces, exit-rate data agreement, and an end-to-end witness — a three-mode net whose certificate was computed on the flattened object alone and consumed as token-game verdicts by transport — and mutants refused at four preregistered loci. The flattening-agreement law, formerly consumed as a hypothesis, is now itself [**Lean-checked**]: derived from the flattening construction for every net and every declared scheduler at this grain, so the composed corollary states its conclusion with no agreement hypothesis. The same corollary now closes one colour grain up, at the Dirac-firing slice, [**Lean-checked**]: the flattening is again constructed — one flattened transition per delay transition per colour vector, its jump intensity the coloured ordered census times the colour-dependent rate, its reset the cascade-closed firing kernel — with the agreement again derived rather than hypothesized, and a witness net whose certificate is the colour dependence itself: from a one-of-each marking the flattened chain drains with probability 2/5 and flips back with 3/5, the two branches differing only in the colour their selections read, consumed as token-game verdicts by transport; mutants refused at six preregistered loci. The pasting against the general-stochastic-hybrid adequacy square is now [**Lean-checked**] at the chart-trivial face: the flattened coloured chain is read as a pure-jump general stochastic hybrid system — modes the tangible coloured markings over an inert chart, one spontaneous transition per ordered mode pair — whose compensator data is proved to be the token game's own (the total alarm intensity is the observed chain's exit rate; post-jump reads and extended generator rows are the net's rate-weighted reads and Foster–Lyapunov rows, for every choice of derivative data; indicator reads recover the chain's rates, so the translation forgets nothing at this grain), and the martingale-problem adequacy square is consumed unchanged on the execution side: certificates whose drift rows hold at the scheduler-closed token-game chain yield safety and deadline verdicts for every segment-model presentation of the embedded profile — the martingale-problem presentations themselves admitted as the execution community's own solution objects — with the triangle packaged as one statement whose three conclusions are its three corners, an end-to-end witness closed with no hypotheses left, and mutants refused at four preregistered loci (the floor, the generator's sign, the bad set, and the cell's horizon each machine-visibly load-bearing). What remains, stated: at the coloured stratum, the hybrid-grain carrier identification — lex-sorted colour-vector markings onto the flattened object's mode charts as genuine continuous coordinates, with its padding obligations — stays [**designed**], as do the coloured directed (demonic and angelic) composed faces and the exponential-holding realization that would pin the execution-side presentations to the community's own killed construction; the time-indexed comparison is now [**Lean-checked**] on the killed trajectory-measure side (construction, prefix-marginal factorization, transient-read lift; mutants refused at four preregistered loci), its grain bridge for consuming the published non-Zeno hypothesis verbatim remaining [**designed**]; and the guard-identification and per-model premise conditions (no immediate transition enabled initially, invariant-bounded token count, colour-box discharge, cascade termination) are specified as certificate VCs — the premise conditions now carried as mechanized premise rows at both grains, with colour-box discharge holding by construction at the mechanized coloured fragment, whose colour boxes are finite types. Everdij–Blom's published bisimulation [1, 2] [**cited — premises audited per model**] remains the pedigree for the token-level correspondence; the triangle's statement, made at the killed jump-skeleton stratum, does not assume their standing finite-expected-jumps hypothesis — on budget-certified models that hypothesis is recovered as a theorem (§6). Where a checker consumes a flattened profile directly (the budget certificate's working object), that is a per-claim choice made at checking time, not the meaning of the net. A reader who insists the semantics of an SDCPN *is* its token-level ordered execution is owed exactly the SDCPN square's usual-semantics edge together with the triangle — both on the enumerated track. The budget rung currently takes Dirac guard resets; random firing measures there are a mixture-linearity extension [**designed**] (the safety rung already has kernel resets).

Part II — Variadic certificates

8 Finitely many VCs: plates, footprints, and transfer

SDCPN certificates look variadic — token counts are variable and canonically unbounded, so a certificate seems to need one VC per token configuration and, worse, one VC per transition *instance*: a countably infinite family. Neither infinity survives contact with the certificate language, and the two mechanisms fail separately when abused, so we state them separately.

The state side is plate notation. Certificates are polynomials in finitely many *pooled coordinates* — per-place token counts and per-place feature pools $\sum_t \varphi(c_t)$ — and every obligation is one polynomial sequent over those coordinates, with hypothesis lines: feature-range enclosures from the colour boxes, count constraints, integer rounding at count atoms. Obligation count and dimension are set by the net and the feature dictionary, never by a marking bound. The grammar layer is [**Lean-checked**]: the plate calculus has no expression for "the i-th token" — no index literals, no index order — so permutation-equivariance of everything expressible, and descent of invariant outputs to functions of the marking, are theorems of the syntax, mechanized with adversarial mutants (the mutant that adds an index-literal construct falsifies equivariance at a two-token transposition, kernel-visibly).

The transition side rides the same grammar fact. Transition data reads selected tokens only through their colours, so the per-instance obligations of one Petri transition collapse to a single *footprint sequent* in distinguished-token variables — the selected tuple's colours, the pooled coordinates, the counts — with dimension set by the transition's arity. One symbolic VC per Petri transition covers every instance at every marking. The countably infinite family was never infinite.

What the collapsed VC can *claim* depends on the aggregator, sharply. The following transfer package is [**Lean-checked**] (as of 2026-07-30; previously it was [**argued**] with executable countermodels — both layers now exist, and every dropped side condition below is refuted by a kernel countermodel as well as an exact-rational executable one):

- *Safety needs no strictness.* For a certificate that is a commutative-monoid fold of per-token features with order-monotone translations, one footprint safety VC transfers at every token count, births included. Fold-ness is load-bearing: an arity-normalized *mean* is not such a fold — a birth changes its denominator under every untouched token, so its footprint VC can hold with equality while the ambient mean rises — and the exclusion is itself a theorem (the positive transfer demands fold data a mean cannot supply); normalized quantities enter certificates only through cleared-denominator direct VCs.
- *Progress needs exactly three more things*: a per-firing gap that is a fixed positive rational, stated in the scale where the aggregator is a translated sum; a safety VC on every other Petri transition (coverage); and a floor. With those, one footprint VC bounds the number of progress-rowed firings by V_0/δ — initial certificate value over gap — uniformly in token count and interleaving. Each of the three is *necessary*, by kernel countermodels: strict decrease without a pinned gap certifies nothing (halve one token's value forever — unbounded runs at every length); dropping the floor admits a signed-birth descent that runs the value below any bound; and coverage is the VC shape itself.
- *Max-aggregated certificates* support safety VCs but no footprint-only progress VC against arbitrary ambient tokens, under any strictly increasing rescaling — max is a fixed point of conjugation, so no change of scale repairs it. The witness is a two-transition regenerating net that keeps the max constant through infinitely many gap-verified firings (in kernel: all VCs hold at $\delta = 1$, the floor holds, and the count exceeds V_0/δ), while the corresponding sum-certificate correctly refuses. This is a *formation restriction* in the VC vocabulary — a gap VC on an idempotent fold is rejected at load — not a lint.
- *The positive family is wide.* The conjugated sums $\varphi^{-1}\big(\sum_t \varphi(v_t)\big)$ — unnormalized ℓ_p folds interpolating from sum toward max, log-sum-exp likewise — are strictly monotone in every argument at every finite parameter (what the max endpoint lacks). Progress VCs for them are stated in the φ-scale, where transfer is uniform; verdicts transport monotonely without ever computing a root; for integer exponents every obligation is an ordinary exact-rational polynomial sequent, with an explicit factorized modulus $\delta/(p\,B^{\,p-1})$ for gap transport on ranges bounded by B. The modeller who wants max-like aggregation certifies at a finite exponent: a dial turn, not a redesign. (Log-sum-exp obligations are discharged by certified interval enclosures rather than exact rationals, by design.)

Two connections back to Part I, so the prices stay distinct: the progress gap is the same witness species as §5's dwell bound — a strictly positive pinned rational per event — and where the event *count* is itself the claim, the budget of §6 is the certificate; the three prices are deliberately not interchangeable. And one structural dividend: this route needs no invariant token bound. The classical reduction to a finite mode lattice requires the token count to be bounded by a place invariant; the pooled-coordinate route states its obligations at the full combinatorial state space directly — which is exactly the regime "canonically unbounded" names.

9 Worked exhibit: the Mini-Fab as an SDCPN

The next three pages present a complete worked model — the Mini-Fab, the standard six-step re-entrant semiconductor line with three machine groups. The first draws the net in Petri notation (Figure 1); the tables then classify it against the SDCPN element dictionary: 16 places, 19 Petri transitions, 49 arcs, and 12 certificates carried by 150 footprint VCs (each a row of the emitted certificate container), every line checkable against the emitted certificate container (the species named are the container's own vocabulary); and the closing figure draws the certificate hypernets themselves, in plate notation (Figure 2). Exhibit-local notation: T_j is step j's tick window and R_k the remaining-work dictionary; "exact tier" in the reset column means the obligation is discharged directly by the exact-rational checker, its declared VC home being a named successor; and the clock column's "forced" names the deterministic-dwell clock species — a different sense from §6's boundary-hit forced crossings, of which this exhibit has none. Four things to read off it:

- *Nothing scales with the lot count.* The initial marking puts N lots in q_0; N appears nowhere else — not in the transition count, not in the VC count, not in any obligation's dimension. The work-potential certificate budgets each lot's run at exactly 638 firings, tight, at every N simultaneously.
- *The aggregator doctrine of §8 is exercised, not just stated.* The congestion certificate is a conjugated-power fold at $p = 2$, its gap VCs stated in φ-space; the work-ceiling certificate is a max fold carrying safety VCs only — its gap VC is *refused at load*, and the mutant battery exercises the refusal.
- *This exhibit sits in the all-demonic, unit-dwell, Dirac corner* — the discrete-time embedding of an untimed net, with zero stochastic species instantiated. That is a feature for Part I's story: since every certificate quantifies over the demonic scheduler, its claims dominate every rate-raced refinement of the same scheduling freedom; the stochastic species, when they arrive, only resolve freedom the certificates have already priced. The tables' own notes name the one declared-home gap (reset VCs for place-crossing transitions) rather than smoothing it over.
- *The exhibit closes as an open system.* The tables' final note reports a joint policy+certificate synthesis: the plant opens into a comb — one supervisor seat (Figure 1's dashed box, drawn at the widened §10 boundary), sealed for this synthesis reading the cumulative-arrival observable and driving the start$_6$ gate wire — a supervision-pattern filler closes it inside the model class, and under the synthesized latest-safe unlock every schedule of the closed composite reaches all-arrivals at exactly $626N$ steps, against the open plant's worst-case $638N - 12$: a certified throughput-floor gain of $+1.9\%$ at the weekly $N = 84$, met exactly and unbeatable by any filler in counting time. The synthesis objective is lexicographic — certified completion first, then the certified horizon — and the discipline is load-bearing: a gate that never opens ties the arrival optimum and then deadlocks shipping nothing, and an over-gated policy "claims" $370N$ by starving the arrival work out of its own constraint set; both are refused at the completion certificate with the starved queue named. Grades split per direction (§11): the within-$626N$ direction rides geolog-checked rows, the not-before direction and the all-fillers floor are argued with executable witnesses, and the note states which clause each of the two closed-composite encodings carries at which tier. Section 10 carries this story the rest of the way: real time, an exactly optimal-throughput policy, a machine-checked decision-theoretic optimality reduction, an untrusted policy sealed inside a certified envelope, and six orders of magnitude of lot count under one certificate.

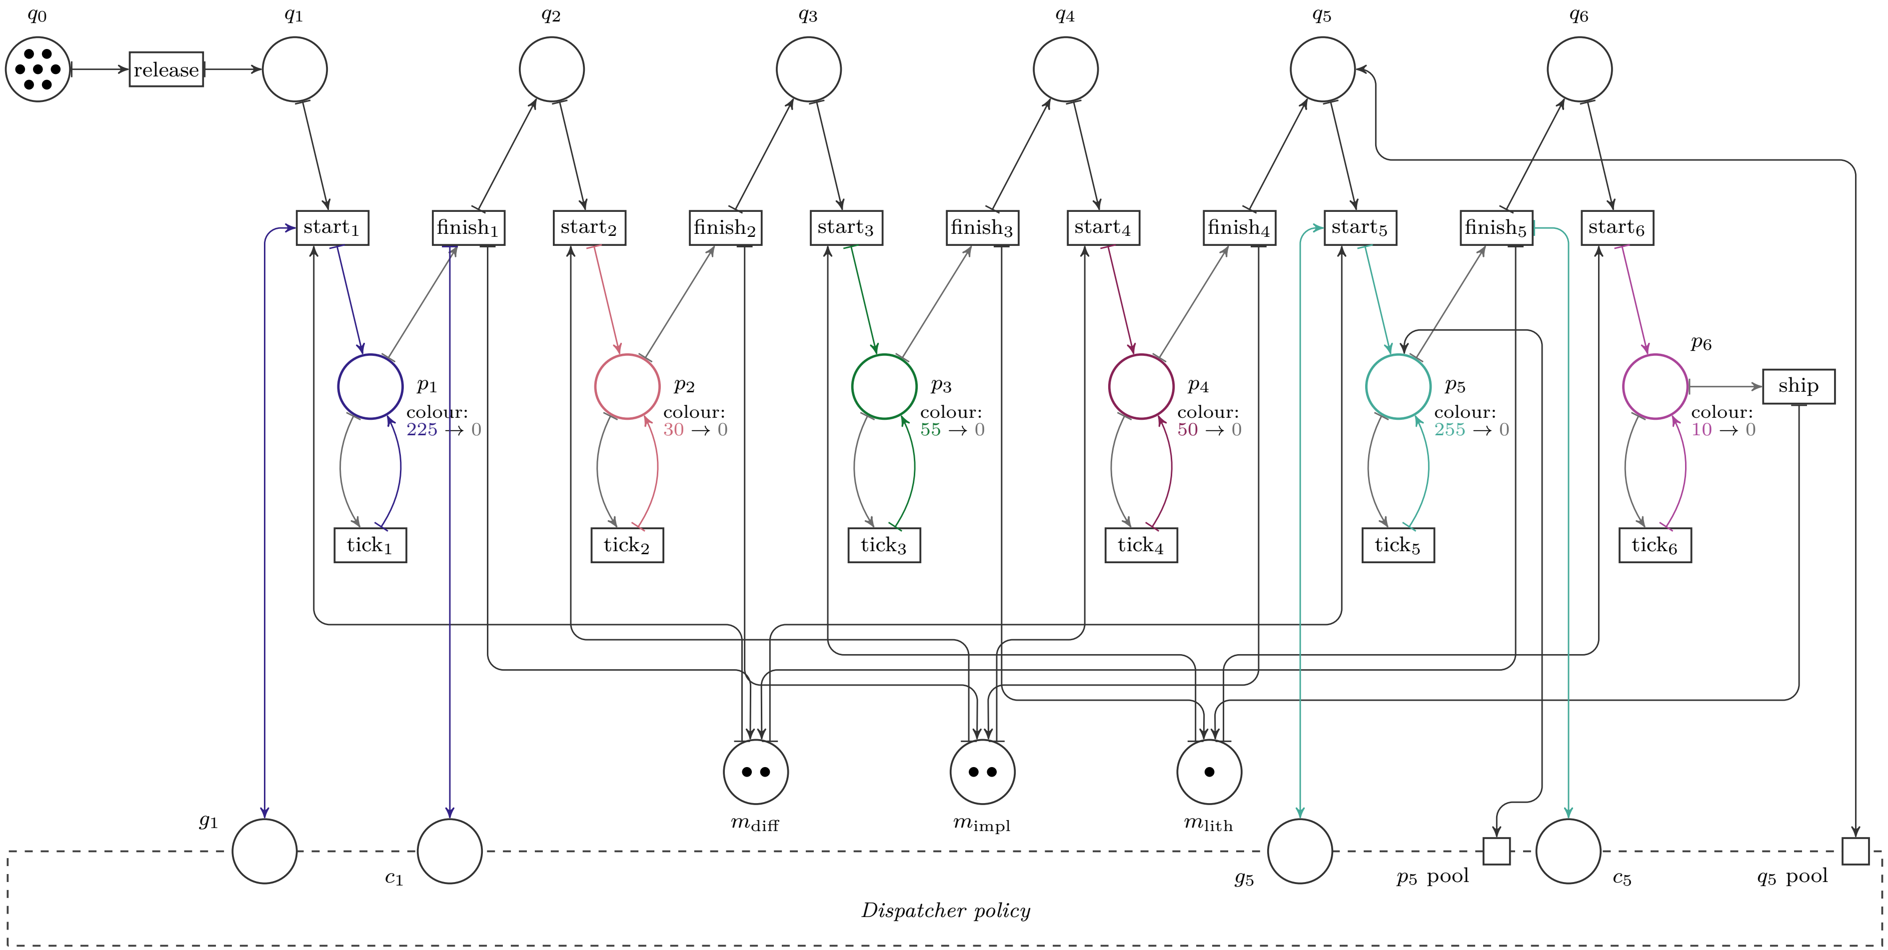

Figure 1. The Mini-Fab net of §9, in Petri notation: circles are places, rectangles Petri transitions, and every arc runs place-to-transition or transition-to-place, drawn with a bar at its source and a head at its target — 16 places, 19 transitions, 49 arcs, the census of the tables that follow. The initial marking is drawn: $N = 7$ lots pending in q_0 (the hexagonal cluster; N appears nowhere else in the model), and $(2, 2, 1)$ idle machines in m_{diff}, m_{impl}, m_{lith}. Colour coding: each in-process place p_k declares its own colour type, tagged *colour:* $T_k \to 0$ with T_k the step's work quantum; the arcs that *produce* into p_k — its start, and its tick's return, which pays one unit of w per firing — wear that type's hue, while every arc that *consumes* from an in-process place wears the one shared grey of the terminal colour 0, the value each countdown drains to and the guard at which finishes and shipment fire; the tag numerals wear the same coding. Machine take and give-back ride the rectilinear arcs to the pool places at the bottom — the two-steps-per-machine-group wiring (1,5 diffusion; 2,4 implantation; 3,6 lithography) is the re-entrancy. The dashed box at the foot is the open system's supervisor hole — the one seat the plant leaves open, labelled by the policy that fills it — drawn at the boundary the synthesized dispatch policy of §10 exercises, one grammar with Figure 3, which draws that filler seated: two gate places g_1, g_5, whose read-and-return enabling arcs (a head at both ends) gate the two diffusion starts — the seat's act wires; two firing counters c_1, c_5, one token per finish$_1$/finish$_5$ firing; and two pooled read-and-return taps on q_5 and p_5 — four observables; with the gates, the strict subset of §10's full-dispatch interface that π^* needs. Boundary wires wear the gated step's hue; the boundary is drawn unfilled — counters empty at the initial marking, gate tokens exactly what a filler mints (Figure 3 draws π^*'s opening move marking g_1). The seat's *first* filler — the counting-time synthesis of the tables' final note — sealed this same seat at a narrower interface: one weight-N cumulative-arrival tap on q_6, one gate place g enabling start$_6$, and the armed plate's one-shot key k, whose certified conservation pair $k + g = 1$ is the Boolean gate discipline; §10 widens the seat, Figure 3 fills it. The hole apparatus is the comb's, beyond the plant census above.

Places (species `HPlate`)

#	plate	colour	Bag carrier	init
0	q_0	—	$\mathrm{Bag}(\mathbf{1}) \cong \mathbb{N}$	N
1	q_1	—	$\mathrm{Bag}(\mathbf{1})$	0
2	q_2	—	$\mathrm{Bag}(\mathbf{1})$	0
3	q_3	—	$\mathrm{Bag}(\mathbf{1})$	0
4	q_4	—	$\mathrm{Bag}(\mathbf{1})$	0
5	q_5	—	$\mathrm{Bag}(\mathbf{1})$	0
6	q_6	—	$\mathrm{Bag}(\mathbf{1})$	0
7	p_1	$[0, 225]$	$\mathrm{Bag}([0, 225])$	0
8	p_2	$[0, 30]$	$\mathrm{Bag}([0, 30])$	0
9	p_3	$[0, 55]$	$\mathrm{Bag}([0, 55])$	0
10	p_4	$[0, 50]$	$\mathrm{Bag}([0, 50])$	0
11	p_5	$[0, 255]$	$\mathrm{Bag}([0, 255])$	0
12	p_6	$[0, 10]$	$\mathrm{Bag}([0, 10])$	0
13	m_{diff}	—	$\mathrm{Bag}(\mathbf{1})$	2
14	m_{impl}	—	$\mathrm{Bag}(\mathbf{1})$	2
15	m_{lith}	—	$\mathrm{Bag}(\mathbf{1})$	1

The arc-type census, classified

SDCPN element	#	pattern instance
ordinary arc, input face	25	Bag destroy edge (rendezvous)
ordinary arc, output face	24	Bag create edge† (spawn)
enabledness count	19	monotone scope rows
colour guard slab	12	guard content on the edge
enabling arc	0	per-tuple parameter-setting
inhibitor arc	0	excluded by formation
delay (rate-raced) transition	0	race clock, finite majorant
forced (deterministic delay)	6	forced clock, dwell 1
immediate transition	0	race clock, infinite majorant
demonic scheduling	13	carved clock $[1, 1]$; choice stays demonic
firing measure	19	Dirac reset legs
initial marking	1	root init rows

Petri transitions (species `Transition`: grain, clock, attachment)

#	transition	consumes [guard slab]	produces [colour law]	clock	reset rows
0	release	q_0	q_1	carved $[1, 1]$	exact tier†
1	start_1	$q_1,\ m_{\mathrm{diff}}$	$p_1\ [w := 225]$	carved $[1, 1]$	exact tier†
2	start_2	$q_2,\ m_{\mathrm{impl}}$	$p_2\ [w := 30]$	carved $[1, 1]$	exact tier†
3	start_3	$q_3,\ m_{\mathrm{lith}}$	$p_3\ [w := 55]$	carved $[1, 1]$	exact tier†
4	start_4	$q_4,\ m_{\mathrm{impl}}$	$p_4\ [w := 50]$	carved $[1, 1]$	exact tier†
5	start_5	$q_5,\ m_{\mathrm{diff}}$	$p_5\ [w := 255]$	carved $[1, 1]$	exact tier†
6	start_6	$q_6,\ m_{\mathrm{lith}}$	$p_6\ [w := 10]$	carved $[1, 1]$	exact tier†
7	tick_1	$p_1\ [1 \le w \le 225]$	$p_1\ [w := w - 1]$	forced	`tr.scale`
8	tick_2	$p_2\ [1 \le w \le 30]$	$p_2\ [w := w - 1]$	forced	`tr.scale`
9	tick_3	$p_3\ [1 \le w \le 55]$	$p_3\ [w := w - 1]$	forced	`tr.scale`
10	tick_4	$p_4\ [1 \le w \le 50]$	$p_4\ [w := w - 1]$	forced	`tr.scale`
11	tick_5	$p_5\ [1 \le w \le 255]$	$p_5\ [w := w - 1]$	forced	`tr.scale`
12	tick_6	$p_6\ [1 \le w \le 10]$	$p_6\ [w := w - 1]$	forced	`tr.scale`
13	finish_1	$p_1\ [w = 0]$	$q_2,\ m_{\mathrm{diff}}$	carved $[1, 1]$	exact tier†
14	finish_2	$p_2\ [w = 0]$	$q_3,\ m_{\mathrm{impl}}$	carved $[1, 1]$	exact tier†
15	finish_3	$p_3\ [w = 0]$	$q_4,\ m_{\mathrm{lith}}$	carved $[1, 1]$	exact tier†
16	finish_4	$p_4\ [w = 0]$	$q_5,\ m_{\mathrm{impl}}$	carved $[1, 1]$	exact tier†
17	finish_5	$p_5\ [w = 0]$	$q_6,\ m_{\mathrm{diff}}$	carved $[1, 1]$	exact tier†
18	ship	$p_6\ [w = 0]$	m_{lith}	carved $[1, 1]$	exact tier†

Certificates (species `HCert`; 150 footprint VCs, species `HFootprintRow`)

certificate	claim family	fold	features	footprint VCs
work-potential	hard progress (ranking)	sum	$v = F$	$19 \times$ gap 1
lot-count	census upper bound	sum	1 on lot plates	18 safety + gap at ship
machines-diff-ub	place invariant, $\le$	sum	$+1$ on $m_{\mathrm{diff}}, p_1, p_5$	6 safety
machines-diff-lb	place invariant, $\ge$	sum	-1 on $m_{\mathrm{diff}}, p_1, p_5$	6 safety
machines-impl-ub	place invariant, $\le$	sum	$+1$ on $m_{\mathrm{impl}}, p_2, p_4$	6 safety
machines-impl-lb	place invariant, $\ge$	sum	-1 on $m_{\mathrm{impl}}, p_2, p_4$	6 safety
machines-lith-ub	place invariant, $\le$	sum	$+1$ on $m_{\mathrm{lith}}, p_3, p_6$	6 safety
machines-lith-lb	place invariant, $\ge$	sum	-1 on $m_{\mathrm{lith}}, p_3, p_6$	6 safety
congestion-ℓ_2	congestion non-increase	conj. pow. $p{=}2$	$v = R$	$6 \times \varphi$-gap 1 + 13 safety
work-ceiling-max	ceiling safety	max	$v = R$	19 safety (gap refused)
arrival-potential	deadline ranking (reach q_6)	sum	$v = D$	$16 \times$ gap 1 + 3 safety
step6-budget	post-arrival budget	sum	$v = F - D$	$3 \times$ gap 1 + 16 safety

Places. A place is a Bag of coloured tokens — remaining work w in its declared box for the in-process places, unit tokens (so: a counter) everywhere else. The machine plates are the resource-pool pattern in SDCPN dress: take and give are ordinary arcs, and the pool's conservation law is certified (the machine pairs above), not assumed.

Transitions. Every arc is an *ordinary* arc: each consumed entry is one input-face arc, each produced entry one output-face arc, every firing law Dirac — a deterministic colour polynomial in the consumed colours. There are no enabling arcs (every emitted selection row carries read $= 0$, a fact a reader can grep the container for) and no inhibitor arcs: enabledness everywhere is a monotone count-and-slab condition, so adding tokens never disables a step. Guard slabs are the footprint sequents' hypothesis slabs; every dwell is 1. †The six same-place ticks carry their reset as a declared scale row (`tr.scale`: image and reset polynomial in the container). The thirteen place-crossing transitions discharge their produced side directly in the exact-rational tier; the declared reset-row home arrives with the Bag reset-row presentation, staged for the next theory revision.

The census. Counts are arcs, transitions, or slabs. The zero rows read two ways: enabling arcs and both stochastic transition species are *available and unexercised* — this exhibit is the demonic-choice, carved-unit-dwell, Dirac corner of the frame (every clock carved to the window $[1, 1]$: timing owned by the clocks, choice by the seat), the discrete-time embedding of an untimed net, faithful for ordering and counting — while the inhibitor row is *unavailable by design*, warranted by the upward-closure theorem. Since every certificate quantifies over the demonic scheduler, its claims dominate every rate-raced refinement of the same scheduling freedom: the stochastic species, when they arrive, only resolve freedom the certificates have already priced. The founding taxonomy's conflict-resolution element is likewise priced rather than modelled: no priority or randomization rows exist, and every consequence holds by dominance over all of them at once.

Certificates. F is the remaining-firings dictionary ($F_k = \sum_{j \ge k}(T_j + 2)$; $v(q_0) = F_1 + 1 = 638$, $v(q_k) = F_k$, $v(p_k, w) = w + 1 + F_{k+1}$, machines uncited) and R the remaining-work dictionary ($R_k = \sum_{j \ge k} T_j$; $v(q_0) = R_1 = 625$, $v(q_k) = R_k$, $v(p_k, w) = w + R_{k+1}$). Transfer polarity reads off the VC profile: an all-gap sum certificate is a ranking (every firing pays, so runs are budgeted — 638 firings per lot, tight); a pair of safety VCs in both signs is a conserved place invariant ($\#m_s + \#p_i + \#p_j = C_s$); the conjugated-power gap VCs budget the ticks in φ-space; and the max fold carries safety only — a gap VC on an idempotent fold is a formation error the vocabulary load-rejects, exercised by the mutant battery.

Deadline. D is the arrival dictionary: F with the step-6 remainder $F_6 = T_6 + 2 = 12$ subtracted on every position before q_6, and 0 from q_6 on — v_D vanishes exactly where a lot has entered q_6, so the fold's zero-set *is* the reach goal. The pointwise split $D + (F - D) = F$ retells the ranking as a deadline argument: pre-arrival firings pay D (at most $626N$ from N pending lots), post-arrival firings pay $F - D$ (at most 12 per arrived lot), so along every maximal run every lot has entered q_6 within $T(N) = 638N - 12$ steps — counting time, the discrete-time unfolding's step index (the metric-time note below prices the same bound on the tick meter). Read as a Spec this is the deadline lower bound $\mathrm{P}_{\geq 1}\left[\mathrm{F}_{\leq T(N)}\,(\text{all } N \text{ lots have entered } q_6)\right]$ with demonic timing: in the Dirac corner every scheduler resolution is a point mass, the dictionaries quantify over the whole demonic seat, and the bound is tight — the schedule that defers the last finish_5 behind all other work attains it exactly (and typically-scheduled runs sit at it too: the step-6 tail is 12 firings against step 5's 255-tick pole). Three honesty lines the arithmetic enforces: the goal is *cumulative* arrival, never simultaneous occupancy ($\#q_6 = N$ at once is demonically avoidable — start step 6 as lots arrive and q_6 stays drained); the claim reads at the deadline wall (a record that stops existing is not late; the completion conjunct that upgrades the wall to full worst-case reach is the liveness face, priced separately); and the counting-time reading is primary but no longer alone — every clock in this encoding is carved to a finite must-fire window, and the metric-time note below prices the same bound with steps replaced by window widths. The judgment row binding this Spec to the container is the completion face's declared row home — not smoothed over: the deadline content shipped here is the two certificates' 38 footprint VCs.

Deadline, metric time. With every clock carved to a positive finite must-fire window (dwell exactly one at these declarations), the counting deadline becomes a wall-clock statement: all N lots enter the finishing queue within $638N - 12$ time units along every maximal schedule ($626N$ under the latest-safe unlock) — probability 1, uniformly over the scheduler's remaining freedom. The certificate content is exactly the dictionary pair already presented; the transfer multiplies by the largest window ceiling (here 1), read off the model bytes by a must-fire census that refuses any idling member. Three clauses carry the reading, each at its stated tier: the dictionary rows are geolog-checked in-loop; the window-top law (a finite ceiling is a deadline) is the specification's declared clock semantics, its meter arithmetic owed at the proof tier; and maximality with deadlock-freedom-below-goal remain the same two argued clauses the counting statement already carries — the checkable species for the latter exists (the angelic-selection witness family) and awaits its active-region staging. The bound is firing-tight but not time-tight: simultaneously due windows fire in batches, and the round replay brackets the attained wall well below the certified one — structure, not slack in the rows.

Open system and synthesis. The plant opens into a comb: one supervisor seat whose declared interface reads the cumulative-arrival observable and drives the start_6 gate wire. A supervision-pattern filler closes it inside the model class — per gate, a gate place g, a one-shot key k (initial 1), and an unlock jump consuming k, reading q_6 at weight N (returning what it reads) and minting g; the gated Petri transition gains a read-and-return enabling arc. Closed roster: 18 places, 20 Petri transitions, 14 certificates, 160 footprint VCs — the plant's twelve extend (the ranking pays the unlock's single firing through the key: total completing-run firings $= 638N + 1$), and the Boolean gate discipline is *certified*, not declared, by the conservation pair $k + g = 1$. Joint policy+certificate synthesis runs lexicographically: certified completion first (the gate that never opens ties the arrival optimum, then deadlocks with every lot in q_6 shipping nothing — refused, starved plate named; so is the over-gated policy whose claimed horizon $370N$ beats the optimum by starving arrival work out of its own constraint set), then minimize the certified horizon, where the least dictionary-certifiable T is an exact-$\mathbb{Q}$ linear program over dictionary space (gap rows on the admitted pre-goal roster, boxwise nonnegativity, goal pins, goal dominance) whose deterministic optimal vertex *is* the arrival dictionary D restricted to the closed roster. The winner is the latest-safe unlock (reopen exactly at the Nth arrival): under it no pre-goal firing exists but the $626N$ arrival-paying ones, so every schedule of the closed composite reaches all-arrivals at exactly step $626N$ — against the open plant's $638N - 12$, a certified throughput-floor gain of $12(N-1)$ steps (+1.6% at $N = 7$; +1.9% at the weekly $N = 84$: 52584 vs 53580). Honesty lines: the geolog-checked rows carry the *within*-$626N$ direction (gap rows in-loop); the *not-before* direction is the producer tier's exact-drop identities (each gap row's displacement equals its gap identically, both directions by Bernstein) plus the replay bracket, its in-loop row species a named absence; the $626N$ floor over *all* fillers is plant arithmetic (626 lot-specific firings per arrival, one firing per composite step, policy firings only add); completion in every encoding is two-tier (ranking geolog-checked; enabledness below the ship goal argued); the seat's weight-N observation read arc is the one arc-separation clause the flattened encoding cannot pass — carried as the staged observation stratum's named debt, never smoothed; and this armed-plate encoding has a native-read sibling (17 places, 20 Petri transitions, no key) whose completion carries one clause more — its unlock consumes nothing, so no ranking row can gap it, and unlock-fires-finitely-often rides as a stated convention, argued — each variant citing its own roster and tier, neither borrowing the other's numbers.

(a) work-potential — hard progress (ranking); fold Σ; dictionary F

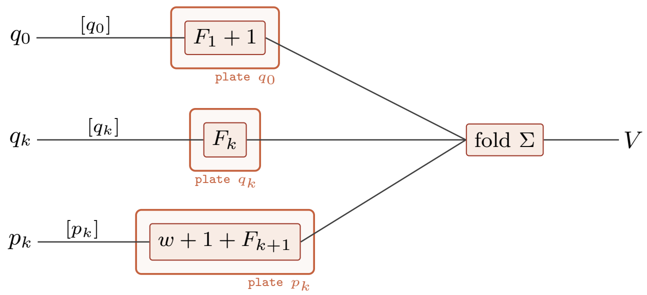

(one q_k and one p_k row per step, $k = 1, \ldots, 6$; all plate indices bound by the fold)

$F_k = \sum_{j \geq k}(T_j + 2)$, the firings a lot still owes; $v(m) = 0$: machines uncited. One gap row per transition, gap 1 — 19 rows; along every run, from any marking m_0, total firings $\leq V(m_0)$, tight at 638 per lot.

(c) the bottleneck pair — two folds, one dictionary R

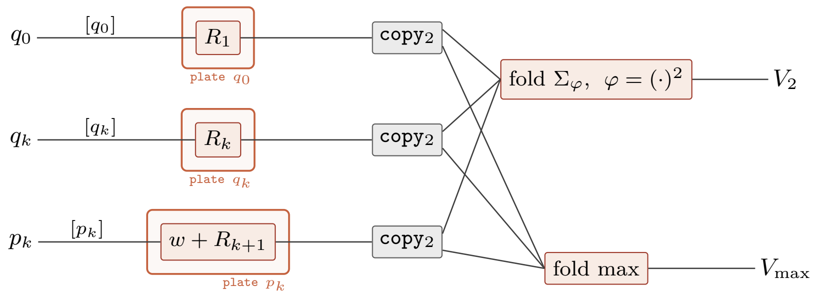

$R_k = \sum_{j \geq k} T_j$, the work a lot still owes ($R_1 = 625$); rows repeat per step as in (a). Congestion-ℓ_2 (upper fold): obligations stated in φ-scale — $\sum v^2$, no root ever computed, verdicts transporting monotonely — 6 φ-gap rows (the ticks) + 13 safety rows. Work-ceiling (lower fold): safety rows only (19); its gap row is *refused at load* — a gap VC on an idempotent fold is a formation error, not a lint (§8).

(b) machine conservation — place invariants; fold Σ; ± twins

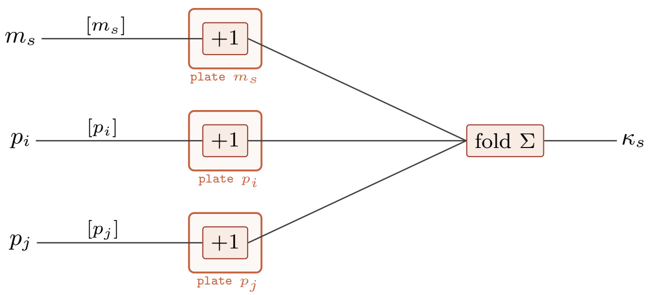

the -1 twin certifies the reverse inequality, pinning $\kappa_s = \#m_s + \#p_i + \#p_j = C_s$ under every scheduler, at every lot count; $(s; i, j) \in \{(\text{diff}; 1, 5), (\text{impl}; 2, 4), (\text{lith}; 3, 6)\}$ with $(C_{\text{diff}}, C_{\text{impl}}, C_{\text{lith}}) = (2, 2, 1)$ — 6 safety rows per certificate, 36 in all.

Figure 2. The exhibit's certificate hypernets, in plate notation. Each tinted enclosing box is a functorial box declared per node: the enclosed feature cell is applied once per token of the plate named at the box corner (its whiskered lift), and its wires carry the plate's symbolic extent, marked $[q_k]$, $[p_k]$; a fold node consumes the plated wires into one scalar. The drawn artefact is therefore a single finite object denoting a function of the whole variable-dimensional marking — the grammar has no expression for "the i-th token," which is why one footprint VC per (certificate, transition) pair covers every instance at every marking (§8). The sum fold is contraction with the plate indices bound; constant features compile to the standard count wire, so (b) reads $\kappa_s = \#m_s + \#p_i + \#p_j$. Four shapes carry all ten certificates of the preceding tables: (a) the ranking; (b) the six machine invariants; (c) the two folds sharing the remaining-work dictionary R; and the tenth — the lot census — is (b)'s wiring with feature 1 on the nine lot populations (q_0, q_k, p_k; machines uncited), 18 safety rows and a gap row exactly at shipment. Row total: $19 + 19 + 36 + 19 + 19 = 112$ footprint VCs, the tables' census.

10 The dispatch seat, end to end: real time, optimal throughput, and untrusted policies

Everything in §9 is stated in *counting time* — makespans measured in firings — where total work is scheduler-invariant ($638N$ firings, always) and arrival time is the only nontrivial objective. This section re-reads the same exhibit in *real time* and pushes the open-seat story to its natural end: what the wall clock rewards, exactly how much, a policy that attains the optimum, a machine-checked sense in which "optimal" is the decision-theoretically right word, a deliberately untrusted policy run inside a certified envelope, and one certificate that spans lot counts from 7 to 10^7.

Real time: the carved-clock reading

The real-time semantics is the theory's discrete clock species: every transition dwells exactly one tick, and each transition row fires at most once per tick (the per-row meter law); simultaneous instances resolve through the demonic scheduler seat, so the only freedom left is genuinely dispatch freedom. Both closed-composite encodings reseal under this reading with row tallies identical to the counting-time containers — the certificates themselves never change.

Two transfers move the counting-time results onto the wall clock. First, upper bounds transfer by a *no-idle* argument: before the goal, a tick with an empty due set — no enabled transition instance whose one-tick dwell has elapsed — would freeze the state forever, so along every goal-reaching schedule each tick fires at least one row — ticks-to-all-arrivals is bounded by pre-goal firings, giving $\leq 626N$ ticks under the synthesized gate and $\leq 638N - 12$ open, carried by the *same* geolog-checked rows plus two named glue clauses [**argued**] (the meter/no-idle lemma and enabledness below the goal). Second, there is a wall-clock *floor* over all dispatch policies of any kind, supervisory or never-idle:

$$\text{ticks to all arrivals} \;\geq\; 255N + 371 \qquad \text{(21,791 at the weekly } N = 84\text{)},$$

because step 5's service tick is one row — at most one firing per tick, 255 firings per lot — behind an uncompressible 370-event first-lot pipeline fill, plus the final transfer into the arrival queue. The bound is Z3-verified below $N = 8$, and its machines-deleted variant is refused too (checked at $N \leq 2$): machine capacity is not the binding constraint; the clock law and per-lot precedence are. (Tight at $N = 1$: the singleton trajectory takes exactly 626 ticks.)

Measured against that floor at $N = 84$, the dispatch levers separate cleanly: the synthesized arrival gate, whose counting-time value is certified, moves wall-clock arrivals 0.0% (greedy dispatch reaches all-arrivals in 22,019 ticks with the gate on or off — the gate's value was always the certified worst-case floor, a guarantee over every schedule, not a typical-case speedup); pull-order is a 1.83× lever (an adversarial pull priority needs 40,406 ticks, +83.5%); and greedy itself sits 228 ticks above the floor. Real-time value lives almost entirely in the dispatch seat — which is why the seat widens next.

Optimal throughput, exactly — and IDT-correctness, machine-checked

The exact optimum is now known in closed form at small N and attained at every tested N:

$$T_{\text{opt}}(N) \;=\; 255N + 371 + d(N), \qquad d = 0 \ (N \leq 2), \quad 2 \ (N = 3), \quad 4 \ (N \geq 4),$$

certified two-sided for $N \leq 7$ by exhaustive Z3 search — satisfiability with an independently re-validated witness schedule at T_{opt}, unsatisfiability at $T_{\text{opt}} - 1$ — and the +4 plateau is attained by an explicit supervisor family π^* (hold the first few diffusion starts on one machine with two-tick handovers, then alternate pre-positioned; run the step-1 services back to back) at every tested lot count ($N \leq 8$, 12, 30, 84): 21,795 ticks at the weekly count, within 4 ticks (0.018%) of the floor, and *exactly* optimal in per-lot throughput at all N — 255 ticks per lot, matched from both sides. The attainment side is *container replay* [**geolog-checked**]: π^*'s explicit firing table replays against the sealed exhibit bytes to arrival at exactly 21,795, under the meter law, every hold accounted. The floor is *not* attainable for $N \geq 3$ (Z3 refuses it outright), and the +4 plateau has a structural reason [**Lean-checked**]: a five-case machine-contiguity analysis of the first four diffusion entrants shows any schedule faster than $255N + 375$ leaves the third or fourth lot's service without a contiguous machine window (for $N \geq 4$; the $N = 3$ value $d = 2$ is Z3-certified two-sided).

Two consequences deserve their own sentences. *Deliberate idling is necessary for optimality*: for $N \geq 3$ the best never-idle (work-conserving) policy needs $255N + 455$ ticks — at the weekly count, 80 ticks behind π^*, with its never-idle law verified at every pre-goal tick of the container replay — and a measured greedy baseline runs 224 ticks behind; gating the *last* step moves arrivals 0.0%, while holding the *diffusion starts* recovers the full gap. The necessity is also mechanical: replaying π^* with holds disallowed is *refused* by the engine at its first delayed diffusion start. Where the supervisor holds is what the last 80 ticks are made of. And *IDT-correctness is discharged, not gestured at*: with deterministic closed dynamics the belief state is a point mass, and imprecise-decision-theoretic selection at the widest-uncertainty locus provably collapses to makespan minimization [**Lean-checked — the reduction, against our mechanized IDT development, with a concrete refutation cell showing the collapse fails without determinism**]. "The synthesized policy is IDT-correct" therefore means optimal throughput — and the optimality receipts above discharge it, each at its stated tier.

One number in Table 1 spent this exhibit's earlier drafts deliberately shy of 100: the certified-optimality column credited π^* with 99.98%, the honest daylight between a $255N + 371$ certified floor and an attainment four ticks above it. That daylight is closed. The five-case machine-contiguity tree is mechanized [**Lean-checked**] — and it collapsed under mechanization: the tick-meter forces two pigeonhole machine handovers, three of the five cases dying inside the first forcing — lifting the certified floor to $255N + 375$ for $N \geq 4$, exactly the attainment. In the same pass the counter/meter bookkeeping behind the machine-free floor landed as [**geolog-checked**] rows, including this container family's first checker-side clock-law rows. The column reads 100.00 exactly, and the mechanization explains why the plateau was +4: it is the double diffusion handover π^* saturates.

Widening the seat: full dispatch as a filler

The supervisor seat of §9 was first sealed at one observable and one gate wire. That narrowness was an encoding shim, not a formalism limit: an interface always was a pair of port lists, and the same comb construction re-seals the same single seat with a *full-dispatch* interface — fifteen observables (six queue lengths, three machine-idle flags, six in-process work counts) and four act ports (three per-station pull decisions plus token selection) — around the unchanged plant. The closed net is the plant net, so the exhibit's twelve-certificate family binds every filler *a fortiori* [**geolog-checked**]: the sealed full-dispatch container re-verifies with the full row roster and the certificate never mentions the policy.

Two findings from exercising it. Token selection is a *second* dispatch freedom that determinism requires: a pull priority alone does not close the system; fixing both yields one trajectory per policy, verified by replaying residual-order permutations to identical traces. And across the 16-policy deterministic grid (8 pull priorities $\times$ 2 selection rules, $N \in \{1, 7, 84\}$; Table 1), only the diffusion-station pull preference binds the outcome — the implant and litho columns of the table never move a number — while selection swings $\sim$7,000 ticks and *crosses* with pull preference; the best cell lands on $255N + 455$ exactly at $N \in \{7, 84\}$ — the work-conserving optimum, consistent with the previous subsection's separation, since grid policies never deliberately idle. Threshold-style gate supervisors provably cannot express per-station pull preference [**argued**]: dispatch is a genuine seat species, not sugar over gates.

An untrusted policy in a certified envelope

The *envelope* certificates — safety, floors, the deadline — never read the policy: they quantify over every filler of the seat, which is why the policy may be *anything*, including a black box nobody should trust. (The per-policy rows in the tables above do read the policy — but only as a replayed *trace*, re-checked row by row against the sealed container, so even the performance claims ask for no trust in the box, only in the replay.) The exhibit's third closure demonstrates exactly that: the same open seat filled by an uncertified lookahead heuristic (a stand-in for a learned or LLM dispatcher) composed with a certified switching envelope — a decision-budget bench with a certified recovery region. In replay: the heuristic ran free, was never benched, completed, and attained the envelope's arrival bound $638N - 11$ *exactly* at the tested $N = 84$ — one tick above the open plant's $638N - 12$, the closed composite's own unlock firing — so the bound is witnessed tight, not merely sound; a deliberate idler hit the budget boundary and was benched with recovery at exactly 1,009 firings; a rogue policy was benched after eight named in-loop refusals with the plant untouched; and disabling the bench stalls the line — the switch is load-bearing, not decorative. After all runs the sealed container re-verified *byte-identical*: the certificate is carried entirely by the envelope [**geolog-checked**], and no behaviour of the filler — competent, idle, or hostile — can touch it.

One certificate, every lot count

The deadline certificates above are stated once and instantiated at any lot count N: the sealed container's row obligations (footprint, flow, floor, canonicity) never mention N — they quantify over *all* markings —

dispatch policy				ticks per arrival				certified optimality (%)	
m_{diff} pulls	m_{impl} pulls	m_{litho} pulls	token selection	N=1	N=7	N=84	N=10^6	N=84	N=10^6
ours: π^* (synthesized holds — not never-idle)				626.0	**308.6**	**259.5**	**255.000375**	**100.00**	**100.00**
q_5	q_2	q_3	most-advanced	626.0	320.0	260.4	—	99.63	—
q_5	q_2	q_6	most-advanced	626.0	320.0	260.4	—	99.63	—
q_5	q_4	q_3	most-advanced	626.0	320.0	260.4	—	99.63	—
q_5	q_4	q_6	most-advanced	626.0	320.0	260.4	—	99.63	—
q_5	q_2	q_3	least-advanced	626.0	384.7	343.2	—	75.61	—
q_5	q_2	q_6	least-advanced	626.0	384.7	343.2	—	75.61	—
q_5	q_4	q_3	least-advanced	626.0	384.7	343.2	—	75.61	—
q_5	q_4	q_6	least-advanced	626.0	384.7	343.2	—	75.61	—
q_1	q_2	q_3	least-advanced	626.0	445.0	397.1	—	65.34	—
q_1	q_2	q_6	least-advanced	626.0	445.0	397.1	—	65.34	—
q_1	q_4	q_3	least-advanced	626.0	445.0	397.1	—	65.34	—
q_1	q_4	q_6	least-advanced	626.0	445.0	397.1	—	65.34	—
q_1	q_2	q_3	most-advanced	626.0	448.6	477.4	—	54.35	—
q_1	q_2	q_6	most-advanced	626.0	448.6	477.4	—	54.35	—
q_1	q_4	q_3	most-advanced	626.0	448.6	477.4	—	54.35	—
q_1	q_4	q_6	most-advanced	626.0	448.6	477.4	—	54.35	—
certified floor, all policies ($255N+375$ *at* $N \geq 4$, *met by* π^*)				*626.0*	*308.6*	*259.5*	*255.000375*	*100*	*100*

Table 1: All seventeen policies of this section, benchmark-style (the synthesized policy itself is drawn in Figure 3, following this section): the sixteen never-idle full-dispatch cells (three binary pull preferences × token selection; deterministic replay, verified by residual-order permutations) and the synthesized π^*, which is not a row of the grid because it deliberately idles — the one capability the never-idle class lacks, worth exactly the last 80 ticks. Read the parameter columns: only the diffusion column ever moves a number; the implant and litho preferences are provably irrelevant at this roster. Token selection crosses with pull preference (most-advanced wins under q_5, loses under q_1). Units: ticks per arrival (T/N, the throughput cost of one lot; the asymptotic ideal is 255); at $N = 1$ every policy ties the floor at 626.0 (no contention; the floor is tight there). "Certified optimality" is the certified floor divided by the policy's attained ticks: the fraction of certifiably-unbeatable performance the policy *provably* delivers. Ours reads 100.00 at the weekly count — earned, not rounded: the four-tick daylight this column carried through earlier drafts closed when the machine-contiguity tree was mechanized, and the Lean-checked floor now meets the replayed attainment at $255N + 375$ exactly. The closure is N-uniform, so the $N = 10^6$ column reads 100.00 as well. The sixteen never-idle cells are deterministic replays and do not extend to 10^6 (simulation is the bottleneck there by design, Table 2); only rows whose attainment has a certified or argued closed form carry the column, which is itself the point. The single-supervisor lanes of §9 show the same pull-order lever at the same magnitude (40,406 ticks under an adversarial greedy pull there).

and N enters only as three integers in the container's own data (the initial lot count, the unlock guard's threshold multiplicity, and hence the instantiated horizon, which the reader recovers from the marking row by one dot product with the arrival dictionary). Table 2 instantiates, reseals, and re-verifies the pair of exhibits at lot counts spanning six orders of magnitude. The certificate's byte size moves only with the decimal width of those integers (six bytes across the whole sweep); the checker's wall time is flat at $\sim$2 ms; and the objects the claims quantify over grow combinatorially — at $N = 10^6$ the model has at least 2×10^{38} reachable markings and at least $10^{2\times10^7}$ distinct maximal schedules, and the checker never visits any of them.

What is N-parametric and what is not — the honest split. The *certificate* and the *checker* are N-parametric in the strong sense: the same 36 footprint rows (this container's share of the exhibit family's 150 VCs), the same dictionaries, the same claim roster discharge at every N (identical tally vectors across the sweep), and at the landed weekly count the resealed container reproduces the landed exhibit exactly. Re-instantiation at a new N does re-run the *producer*: it reseals the container (0.002 s at $N = 84$, 106 s at $N = 10^7$ — the producer's in-memory form expands the weight-N guard unarily, a linear-cost implementation artifact of the untrusted side that never touches the sealed bytes or the checker), and re-derives the exact-rational producer receipts (the =-direction's drop identities; 0.002 s at $N = 84$, 9.8 s at $N = 10^6$). Dynamic attainment witnesses (replays arriving at exactly $626N$) were re-run through $N = 10^4$, i.e. 6,260,000-step runs; beyond that the replay grinder's quadratic cost makes simulation the bottleneck — which is the point of certifying.

N	$T^* = 626N$	$638N - 12$	cert. (B)	check (ms)	markings $\geq$	schedules $\geq$
7	4,382	4,454	3,105	1.9	3.4×10^3	10^{123}
84	52,584	53,580	3,105	1.9	8.1×10^9	$10^{1{,}719}$
10^3	626,000	637,988	3,107	1.9	2.0×10^{17}	$10^{20{,}469}$
10^6	626,000,000	637,999,988	3,109	1.8	2.0×10^{38}	$10^{2.0\times10^7}$
10^7	6,260,000,000	6,379,999,988	3,111	1.9	2.0×10^{45}	$10^{2.0\times10^8}$

Table 2: The synthesized-policy deadline certificate instantiated at lot count N. Columns: the certified all-arrivals horizon under the synthesized gate ($T^* = 626N$ steps, exact on every completing run); the no-policy worst case ($638N - 12$); the sealed container's size in bytes (supervised exhibit; the no-policy container runs 2,963–2,966 B over the same sweep); the checker's wall time (median of 7 re-discharges of the sealed bytes by the production engine's in-loop evaluator walk, single core of an NVIDIA GB10 (Grace, aarch64) under concurrent load; all samples across all rows fell in 1.7–2.0 ms); a reachable-marking lower bound; and a lower bound on distinct maximal schedules the claims quantify over. Intermediate lot counts 10^4 and 10^5 were measured identically (containers 3,107 and 3,109 B; medians 1.9 ms).

Claim grades at every N. The within-T^* direction rides [**geolog-checked**] at every N (the row obligations are N-free; the horizon instantiates from the container's marking row). The not-before direction and exact attainment remain producer-tier receipts, and completion keeps its graded split (ranking as [**geolog-checked**], enabledness glue [**argued**]), exactly as in the claim ledger — instantiation at large N upgrades nothing and degrades nothing.

The contrast columns, derived not simulated. Both bounds are conservative combinatorics on the net. *Markings:* lots distribute over pending, the six queues, and shipped; with machines idle every such distribution is reachable (advance one lot at a time) and distinct, giving $\binom{N+7}{7}$ reachable markings before counting any in-process or machine state. *Schedules:* two lots processing concurrently on distinct stations (one in step 1, one in step 2) admit $\binom{259}{32} \approx 10^{40.9}$ interleavings of their firing chains, all distinct as transition words; scheduling $\lfloor N/2 \rfloor$ such pairs in blocks gives $\binom{259}{32}^{\lfloor N/2 \rfloor}$ distinct maximal runs. Every one of them is covered by the same $\sim$3 kB of certificate.

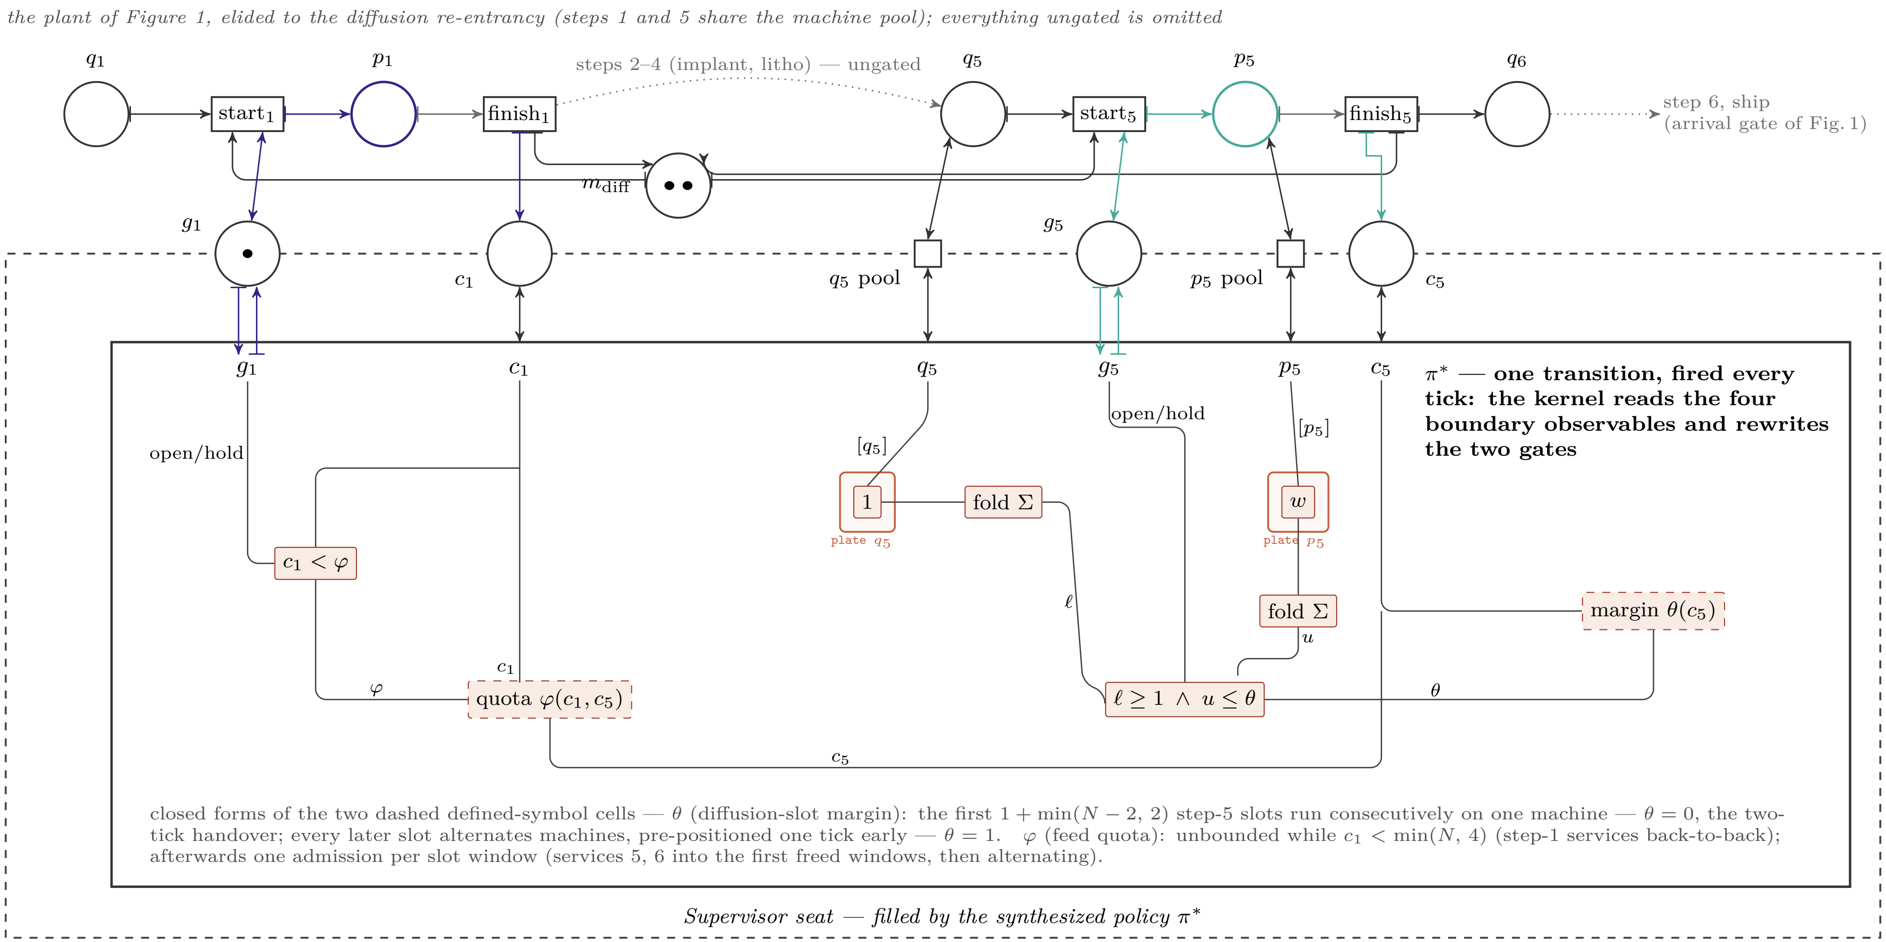

Figure 3. The synthesized optimal policy π^*, drawn as what it is: a filler for the supervisor seat. The dashed seat of Figure 1 is here nearly the whole page, and the plant survives only as the gated neighbourhood along the top — the diffusion re-entrancy, where start$_1$ and start$_5$ compete for the same two machines. Straddling the boundary, the seat's declared interface: two pooled reads (read-and-return taps on q_5 and p_5), two firing counters (c_1, c_5; one token per finish$_1$/finish$_5$ firing), and two gate places whose read-and-return arcs enable the two diffusion starts — a strict subset of the fifteen-observable, four-act full-dispatch interface of §10, and the only wires π^* needs. Inside sits one enormous Petri transition: the policy's firing kernel, applied once per tick, reading the four observables and rewriting the two gates (consume-and-reproduce arc pairs). Its interior is a hypernet in the plate notation of Figure 2: the pooled observables are two plated folds ($\ell = \Sigma\, 1$ over $[q_5]$, the queue length; $u = \Sigma\, w$ over $[p_5]$, the in-service residual — pre-positioning keeps at most the incumbent in p_5 at decision time, so the sum reads the incumbent's residual), the counters decode the schedule phase through the two *dashed* cells (θ, φ — defined symbols, not primitive generators: each abbreviates a piecewise-constant function of the counters, closed forms in the box), and two threshold cells (region-membership) emit the gate bits. Holding start$_5$ until $u \leq \theta$ pre-positions each diffusion slot; holding start$_1$ at the quota keeps a diffusion machine free for step 5 — deliberate idling, the one capability the never-idle class lacks. What this filler attains: all arrivals in $255N + 371 + d(N)$ ticks, $d = 0, 0, 2$, then 4 — 21,795 at the weekly $N = 84$, *on* the certified all-policy floor exactly [**geolog-checked — container replay of the explicit firing table against the sealed exhibit bytes, every hold accounted, on these two gates only; the floor itself Lean-checked**]; exactly optimal for $N \leq 7$ [**Z3-checked — two-sided, witnesses re-validated**]; the +4 plateau structural for $N \geq 4$ [**Lean-checked — the machine-contiguity tree**]. The best never-idle policy needs $255N + 455$ — these two gates are worth the last 80 ticks, replaying π^* with holds disallowed is refused at its first delayed diffusion start, and gating the last step instead moves arrivals 0.0%. Under deterministic closed dynamics, machine-checked imprecise-decision-theoretic selection collapses to makespan minimization [**Lean-checked**] — so this page is also the decision-theoretic optimum. Initial marking drawn: g_1 open ($c_1 = 0 < \varphi$), g_5 shut ($\ell = 0$).

Reproduction. Every certificate in this exhibit is available as a sealed container, together with a one-command reproduction kit that rebuilds the independent checker at a pinned revision, re-verifies each container from its bytes alone — with per-relation claim counts as a load-bearing gate and committed integrity mutants that must refuse — and re-runs the producer-tier mutant batteries. The dispatch grids, the Z3 optima (each satisfiability/unsatisfiability pair re-derives in one command), and the N-scaling sweep ship the same way.

11 Claim ledger

claim	grade	witness
Safety transfer with no non-Zeno hypothesis (total jump-indexed chain; Dirac and kernel resets)	Lean-checked	Lean*
Crossing budget: $\mathbb{E}[N_T^{\mathrm{f}}] \le W(x_0) + cT$, Markov tail, $\mathbb{P}(\text{Zeno by } T) = 0$	Lean-checked	Lean*, mutants
Rank/acyclicity method = the c=0, $\mathbb{N}$-valued budget	Lean-checked	Lean*
Multiplicative (tilted) budget $\mathcal{A}W \le cW \Rightarrow \mathbb{E} \le e^{cT} W(x_0)$	Lean-checked	Lean*, mutants
Budget tightness (thermostat pin) and refusal completeness (Zeno thermostat admits no (W, c))	Lean-checked	Lean*
Truncated extended Borel–Cantelli (iff) + dwell-floor divergence + finitely many events per horizon	Lean-checked	Lean*, mutants
Rate cap $\Rightarrow$ truncated-mean floor $(1 - e^{-\Lambda})/\Lambda$	Lean-checked	Lean*, mutants; tight at the cap
Rate-wing expectation bound $\mathbb{E}[N_t^{\mathrm{r}}] \le \Lambda t$ (G4's total count: a theorem, four wings summed)	Lean-checked	Lean*, mutants; canonical witness
Angelic selection gluing (pointwise VCs $\Rightarrow$ one memoryless law)	Lean-checked	Lean*, mutants 4/4; the e^{-1} constant tight
Two-polarity verdict semantics (safety read universally, liveness existentially, §2)	definitional	certificate grammar
Dwell caps $T/d + 1$ on delay/scheduler wings	definitional	record-level lemma routine
Recurrence metering fork (dwell vs. crossing; incomparability; uniform-lift collapse)	Lean-checked	Lean*
Reach-avoid Łukasiewicz decomposition	Lean-checked	Lean*, mutants; exact-$\mathbb{Q}$ fixture partner; (U) a stated hypothesis

— *continued.*

claim	grade	witness
SDCPN → GSHS transport (published bisimulation, audited premises)	cited + Lean-checked	[1, 2]; the commuting triangle's composed corollary Lean-checked at the uncoloured grain and at the coloured Dirac-firing slice, and its general-stochastic-hybrid foot Lean-checked at the chart-trivial face (agreement law derived from the construction at both grains; compensator data the token game's own; martingale-problem square consumed unchanged; time-indexed comparison Lean-checked, its grain bridge designed; hybrid-grain coloured carrier identification, exponential-holding realization, and coloured directed faces designed)
Inhibitor arcs excluded by formation (enabledness monotone; adding tokens never disables)	Lean-checked	Lean*, mutants
Plate grammar equivariance + descent	Lean-checked	Lean*, mutants
Aggregator transfer package (safety fold transfer; gap/coverage/ floor each necessary; idempotent formation restriction; φ-space corollaries)	Lean-checked	Lean*; exact-$\mathbb{Q}$ battery
Mean locality failure (normalized means excluded structurally)	Lean-checked	Lean*
Post-Zeno safety factors through the floor-forgetting projection (§3)	Lean-checked	Lean*, mutants; ordering carve-out checked negatively

— *continued: the Mini-Fab exhibit (§9).*

claim	grade	witness
Mini-Fab exhibit classifies whole against the SDCPN dictionary (19 Petri transitions / 49 arcs / 16 places / 12 certificates, 150 VCs, incl. the counting-time deadline pair)	definitional	container, line by line
Mini-Fab open system: under the synthesized latest-safe-unlock supervisor, every schedule of the closed composite reaches all-arrivals at exactly $626N$ (open plant: $638N-12$; $+1.9\%$ certified throughput floor at $N{=}84$); the dictionary LP's optimal vertex is the arrival dictionary itself	geolog-checked	container, in-loop + exact tier; exact replay; mutant battery
Completion under the synthesized policy: ranking geolog-checked; enabledness below the ship goal argued; the native-read encoding's old unlock-boundedness clause is *withdrawn* under the real-time reading (the unlock refires freely after the goal; its rows are safety-polarity in every fold, so no certificate consumed boundedness) — the armed-plate encoding's one-shot key never posed it ($k+g=1$ certified)	geolog-checked	starving-gate and never-open mutants refused, starved queue named
The $626N$ floor over *all* fillers of the comb, in counting time (626 lot-specific firings per arrival; policy firings only add)	argued	plant arithmetic; replay attains it

— *continued: the dispatch story (§10).*

claim	grade		witness
Carved-clock (real-time) reading: both closed composites reseal with identical row tallies; counting bounds transfer to ticks by no-idle ($\leq 626N$ gated, $\leq 638N-12$ open)	geolog-checked		carved reseals, tallies identical; two named glue clauses
Real-time floor over *all* dispatch policies: $255N+375$ ticks to all arrivals at $N \geq 4$ ($255N+373$ at $N \geq 3$; machine-contiguity, five cases collapsed to two pigeonhole forcings), strengthening the machine-free $255N+371$ (ship variant $255N+383$; rides the clock law and per-lot precedence only)	Lean-checked geolog-checked	+	Lean contiguity tree, mutants at four preregistered loci; counter/meter rows + three-lemma proof; Z3 UNSAT at floor-1, $N \leq 7$; machines-deleted variant refused at $N \leq 2$
Exact real-time optimum $255N+371+d(N)$, $d=(0,0,2,4,4,\dots)$, closed two-sided at every N: Z3 at $N \leq 2$, Lean-checked floor + π^* attainment at $N \geq 3$ (21,795 at N=84, on the floor exactly; per-lot throughput 255 closed both sides)	Lean-checked geolog-checked	+	Z3 SAT/UNSAT pairs, witnesses independently re-validated, $N \leq 7$; π^* attainment by container replay ($N \leq 8$, 12, 30, 84), closed form beyond; the +4 plateau is the contiguity tree's double handover, saturated by π^*
Deliberate idling necessary: best work-conserving policy $=255N+455$ for $N \geq 3$ (80 ticks behind optimal at the plateau; measured greedy 224 behind at N=84)	geolog-checked		attainment by container replay, never-idle law verified every pre-goal tick; holds-disallowed π^* replay refused; floor argued; grid receipts
IDT-correctness reduction: deterministic closed dynamics $\Rightarrow$ point-mass beliefs $\Rightarrow$ IDT selection = makespan minimization	Lean-checked		Lean*; refutation cell without determinism
Widened full-dispatch seat (15 observables, 4 act ports): the 12-certificate family binds every filler a fortiori; token selection a second, determinism-required freedom; threshold gate supervisors cannot express pull preference (the obstruction, argued)	geolog-checked		sealed container, full roster; permutation replay; 16-policy grid
Untrusted-policy envelope (switching-bench closure): arrival bound $638N-11$ witnessed tight at N=84; bench/recovery exact (1,009; eight named refusals, plant untouched); container byte-identical after all runs	geolog-checked		sealed envelope container; replay receipts
One certificate, every lot count: identical row roster discharges at $N=7$ through 10^7; container grows 6 bytes over six orders; checking flat at $\sim$2 ms	geolog-checked (per N)		re-sealed containers, identical tallies; timing receipts (Table 2)

*Lean 4 over mathlib; zero `sorry`s; axiom footprint exactly {`propext`, `Classical.choice`, `Quot.sound`} per cited theorem; "mutants" = deliberate corruptions at loci preregistered before proving, refused by the checker. The executable countermodel batteries ship in our reproduction repository (one make target per claim family); the exhibit's tables are checkable line by line against the emitted certificate container. This is a discussion draft: repository locators for the Lean developments, the batteries, and the container accompany distribution.

References

[1] M. H. C. Everdij and H. A. P. Blom. Hybrid state Petri nets which have the analysis power of stochastic hybrid systems and the formal verification power of automata. In P. Pawlewski, editor, *Petri Nets Applications*, chapter 12, pages 227–252. InTech, 2010.

[2] M. H. C. Everdij and H. A. P. Blom. Bisimulation relations between automata, stochastic differential equations and Petri nets. In *Proc. Workshop on Formal Methods for Aerospace (FMA)*, EPTCS 20, 2010. arXiv:1003.4812.

[3] M. H. A. Davis. *Markov Models and Optimization.* Chapman & Hall, London, 1993.

[4] O. Kallenberg. *Foundations of Modern Probability.* 2nd edition, Springer, New York, 2002.

Source appendix. This document carries its complete LaTeX source (mirco-tech-report.tex, minifab-sdcpn-petri.tex, minifab-sdcpn-tables.tex, minifab-certs-hypernets.tex, minifab-dispatch-story.tex, minifab-nscaling-table.tex, minifab-pistar-policy.tex) in two forms, adding no pages beyond this one. (1) As an invisible text layer of microscopic type riding the bottom margins of the document's early pages, in the order listed: pdftotext -layout <thisfile>.pdf - yields every source line, in order, as per-page blocks bracketed by the sentinel lines srclayer-dust-begin and srclayer-dust-end (the layer never prints or displays; interior runs of two or more spaces may collapse to one). (2) Byte-exactly, as PDF attachments: pdfdetach -saveall <thisfile>.pdf, or reader.attachments in pypdf.